

COMPLEXITY ASPECTS OF FUNDAMENTAL
QUESTIONS IN POLYNOMIAL OPTIMIZATION

JEFFREY ZHANG

A DISSERTATION

PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF
OPERATIONS RESEARCH AND FINANCIAL ENGINEERING

ADVISER: AMIR ALI AHMADI

SEPTEMBER 2020

© Copyright by Jeffrey Zhang, 2020.

All rights reserved.

Abstract

In this thesis, we settle the computational complexity of some fundamental questions in polynomial optimization. These include the questions of (i) finding a local minimum, (ii) testing local minimality of a candidate point, and (iii) deciding attainment of the optimal value. Our results characterize the complexity of these three questions for all degrees of the defining polynomials left open by prior literature.

Regarding questions (i) and (ii), we show that unless $P=NP$, there cannot be a polynomial-time algorithm that finds a point within Euclidean distance c^n (for any constant c) of a local minimum of an n -variate quadratic function over a polytope. This result answers a question of Pardalos and Vavasis that appeared in 1992 on a list of seven open problems in complexity theory for numerical optimization. By contrast, through leveraging techniques from algebraic geometry, we show that a local minimum of a cubic polynomial can be found efficiently by semidefinite programming. Interestingly, we prove that second-order points of cubic polynomials admit an efficient semidefinite representation, even though their critical points are NP-hard to find. We also give an efficiently-checkable necessary and sufficient condition for local minimality of a point for a cubic polynomial.

Regarding question (iii), we prove that testing whether a quadratically constrained quadratic program with a finite optimal value has an optimal solution is NP-hard. We also show that testing coercivity of the objective function, compactness of the feasible set, and the Archimedean property associated with the description of the feasible set are all NP-hard. The latter property is the assumption on which convergence of the Lasserre hierarchy relies. We also give a new characterization of coercive polynomials that lends itself to a hierarchy of semidefinite programs.

In our final chapter, we present a semidefinite programming relaxation for the problem of finding approximate Nash equilibria in bimatrix games. We show that for a symmetric game, a $1/3$ -Nash equilibrium can be efficiently recovered from any rank-2 solution to this relaxation. We also propose semidefinite programming relaxations for NP-hard problems related to Nash equilibria, such as that of finding the highest achievable welfare under any Nash equilibrium.

Acknowledgements

I would like to first and foremost thank my adviser, Amir Ali Ahmadi, for making all of this possible. Thank you for being an amazing adviser, and for teaching me everything you have. Thanks to you, I feel ready to be a professor and a researcher. Thank you for all the long hours, for not letting me take the easy road, and for all the feedback on every little thing. Thank you for giving me every opportunity to grow as a student, as a future teacher, and as a person. Thank you for broadening my horizons, both academically and literally. I know that so much of where I am now is due to you, and I know that all the lessons will stick with me for a long future to come.

I also want to thank the other professors who helped me along my way; Bob Vanderbei, for being there for me from the very beginning in so many ways, and I hope my life can be as interesting as yours one day. Nicolas Boumal and Anirudha Majumdar for also serving on my thesis committee. Jianqing Fan, for introducing me to ORFE which has become a home to me. Matt Weinberg, for a most interesting course and new research ideas. I would also like to thank all the members of ORFE who have made my experience what it was. Thank you Kim, Tara, Michael, Tiffany, and Melissa for putting up with me and answering questions when I had them.

I would also like to thank all the friends I made along the way; Bachir and Cemil, who made working with Amirali all the more interesting, Galen and Dan, for all the hours spent in the STWG, Thomas, Elahe, Yiqiao, Sinem, Yair, Kaizheng, Suqi, for making the time spent more interesting. A special thanks to Georgina, who was a guiding light in the early years when I needed one. Thank you also to David, Chris, Alan, Dan, Andre, Josh, Connor, Demi, and Jenny for keeping me grounded and always different perspectives on life.

Last but not least I want to thank my family. My brother Leon for all the motivation, ideas, and discussion, and being with me my whole life. My father Heping, for setting me in the path to being a professor, keeping me motivated, and all the advice over the years. And finally my mother Julan, without whose love and passion I may not be here today.

To my mother, Julian, for all her loving care.

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	x
List of Figures	xii
1 Introduction	1
1.1 Preliminaries	7
1.2 Outline of this thesis	10
1.3 Related Publications	11
2 On Local Minima and Related Notions in Unconstrained Polynomial Optimization	12
2.1 Introduction	12
2.1.1 Organization and Main Contributions of the Chapter	16
2.1.2 Preliminaries and Notation	17
2.2 NP-hardness Results	18
2.3 Checking Local Minimality of a Point for a Cubic Polynomial	21
2.3.1 Preliminaries on Cubic Polynomials	21
2.3.2 Local Minimality of a Point for a Cubic Polynomial	22
2.3.3 Examples	29
2.4 On the Geometry of Local Minima of Cubic Polynomials	31

2.4.1	Convexity of the Set of Local Minima	32
2.4.2	Local Minima and Solutions to a “Convex” Problem	37
2.4.3	Distinction Between Local Minima and Second-Order Points	40
2.4.4	Spectrahedra and Convexity Regions of Cubic Polynomials	44
2.5	Complexity Justifications for an Exact SDP Oracle	46
2.6	Finding Local Minima of Cubic Polynomials	50
2.6.1	Preliminaries from Semidefinite and Sum of Squares Optimization	50
2.6.2	A Sum of Squares Approach for Finding Second-Order Points	53
2.6.3	A Simplified Semidefinite Representation of Second-Order Points and an Algorithm for Finding Local Minima	57
2.7	Conclusions and Future Directions	70
2.7.1	Approximate Local Minima	70
2.7.2	Unregularized Third-Order Newton Methods	72
3	On the Complexity of Finding a Local Minimizer of a Quadratic Function over a Polytope	77
3.1	Introduction	77
3.1.1	Notation and Basic Definitions	80
3.2	The Main Result	81
3.2.1	Complexity of Deciding Existence of Local minima	81
3.2.2	Complexity of Finding a Local minimum of a Quadratic Function Over a Polytope	87
4	On Attainment of the Optimal Value in Polynomial Optimization	91
4.1	Introduction	91
4.1.1	Organization and Contributions of the Chapter	93
4.2	NP-hardness of Testing Attainment of the Optimal Value	94
4.3	NP-hardness of Testing Sufficient Conditions for Attainment	98

4.3.1	Coercivity of the Objective Function	98
4.3.2	Closedness and Boundedness of the Feasible Set	100
4.4	Algorithms for Testing Attainment of the Optimal Value	108
4.4.1	Compactness of the Feasible Set	108
4.4.2	Coercivity of the Objective Function	110
4.5	Summary and Conclusions	116
5	Semidefinite Programming Relaxations for Nash Equilibria in Bimatrix Games	118
	Games	118
5.1	Introduction	118
5.1.1	Organization and Contributions of the chapter	121
5.1.2	Notation	122
5.2	The Formulation of our SDP Relaxation	122
5.2.1	Nash Equilibria as Solutions to Quadratic Programs	123
5.2.2	SDP Relaxation	125
5.2.3	Valid Inequalities	127
5.2.4	Simplifying our SDP	131
5.3	Exactness for Strictly Competitive Games	133
5.4	Algorithms for Lowering Rank	136
5.4.1	Linearization Algorithms	137
5.4.2	Numerical Experiments	142
5.5	Bounds on ϵ for General Games	144
5.5.1	Lemmas Towards Bounds on ϵ	144
5.5.2	Bounds on ϵ	148
5.5.3	Bounds on ϵ in the Rank-2 Case	152
5.6	Bounding Payoffs and Strategy Exclusion in Symmetric Games	156
5.6.1	Bounding Payoffs	156
5.6.2	Strategy Exclusion	158

5.7	Connection to the Sum of Squares/Lasserre Hierarchy	160
5.7.1	Sum of Squares/Lasserre Hierarchy	161
5.7.2	The Lasserre Hierarchy and SDP1	161
5.8	Future Work	166
A	Appendices for Nash Equilibria	168
A.1	Statistics on ϵ from Algorithms in Section 5.4	168
A.2	Comparison with an SDP Approach from [60]	170
A.3	Lemmas for Extreme Nash Equilibria	171
A.4	A Note on Reductions from General Games to Symmetric Games	173
A.4.1	The Reduction of [45]	173
A.4.2	The Reduction of [54]	174
	Bibliography	176

List of Tables

1.1	Complexity of deciding whether a point is a local minimum of a POP, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table. ¹	3
1.2	Complexity of deciding whether a POP has a local minimum, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table.	4
1.3	Complexity of deciding whether a POP with a finite optimal value has an optimal solution, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table. Note that whenever the degree of the objective is at most three, and the feasible set is a polyhedron, there is no algorithm required; the answer is simply ‘yes’.	6
2.1	Complexity of deciding whether a given point is of a certain type, based on the degree of the polynomial. Entries without a reference are classical.	15
2.2	Complexity of deciding whether a polynomial has a point of a certain type, based on the degree of the polynomial. Entries without a reference are classical.	15
2.3	Iterations of the third-order Newton method (left) and the classical Newton method (right) on the function f in (2.34) starting at $x_0 = 1.5$	74

5.1	Statistics on ϵ for 20×20 games after 20 iterations.	143
5.2	Performance of LP2 and SDP4 on 5×5 games	160
5.3	Performance of LP2 and SDP4 on 10×10 games	160
A.1	Statistics on ϵ for 5×5 games after 20 iterations.	168
A.2	Statistics on ϵ for 10×5 games after 20 iterations.	168
A.3	Statistics on ϵ for 10×10 games after 20 iterations.	169
A.4	Statistics on ϵ for 15×10 games after 20 iterations.	169
A.5	Statistics on ϵ for 15×15 games after 20 iterations.	169
A.6	Statistics on ϵ for 20×15 games after 20 iterations.	169
A.7	Statistics on ϵ for 20×20 games after 20 iterations.	169
A.8	Statistics on ϵ for first level of the hierarchy in [60].	170
A.9	Statistics on ϵ for SDP2 with $\text{Tr}(M)$ as the objective function.	170

List of Figures

2.1	Contour plots of $x_1^2x_2$ (left) and $x_2^2 - x_1^2x_2$ (right) from Examples 2.3.1 and 2.3.2. The polynomials are zero on the black lines, positive on the gray regions, and negative on the white regions. The dashed line in the right-side figure denotes a descent parabola at the origin.	29
2.2	The critical points of the polynomial (2.9). One can verify that the set of critical points is $\{(x_1, x_2) \mid (x_1 + x_2)^2 = 1\}$, and that the set of local minima is $\{(x_1, x_2) \mid x_1 + x_2 = 1\}$. The points on the dashed line are local maxima.	34
2.3	The projection of the set of local minima of the polynomial in (2.10) onto the x_1 and x_2 variables. This example shows that $\overline{LM_p}$ is not always a polyhedral set.	35
2.4	The set of local minima (left) and second-order points (right) of the cubic polynomial $p(x_1, x_2) = x_1^2x_2$. Note that SO_p is the closure of LM_p (Theorem 2.4.7) and LM_p is the relative interior of SO_p (Theorem 2.4.10).	43
2.5	The plots of the function f in (2.34) and its second and third-order Taylor expansions around $x_0 = 1.5$. One can see that one iteration of the third-order Newton method in (2.33) brings us closer the global minimum of f compared to one iteration of the Newton method in (2.32).	75
2.6	Sensitivity of the limits of the iterates (2.32), (2.33), and (2.35) respectively to initial conditions. Regions with the same color denote initial conditions which converge to the same critical point.	76

5.1	Distribution of ϵ over numbers of iterations for the square root algorithm (left) and the diagonal gap algorithm (right).	143
5.2	The quality of the upper bound on the maximum welfare obtained by LP1 and SDP3 on 100 5×5 games (left) and 100 10×10 games (right).	158

Chapter 1

Introduction

In this thesis, we concern ourselves with *polynomial optimization problems (POPs)*, i.e., problems of the type

$$\begin{aligned} \inf_{x \in \mathbb{R}^n} \quad & p(x) \\ \text{subject to} \quad & q_i(x) \geq 0, \forall i \in \{1, \dots, m\}, \end{aligned} \tag{1.1}$$

where p, q_1, \dots, q_m are polynomial functions. In Chapters 2-4, we address the complexity of the following questions for a problem in the form of (1.1):

Q1: Is a given a point x a local minimum of (1.1)?

Q2: Does (1.1) have a local minimum (and if so, can one be found efficiently)?

Q3: If (1.1) has a finite optimal value, does it have an optimal solution?

Precise definitions and in-depth study of related problems can be found in their respective chapters.

POPs have wide modeling capabilities and arise ubiquitously in applications, either as exact models of objective functions or as approximations thereof. Perhaps the most well-known special case of POPs is linear programming, but POPs have much richer expressive power. For example, in full generality any decision problem in NP, a class of yes-no decision problems with the property that any yes answer can be certified efficiently, can be posed as

a POP.¹ An example of a POP we will see in this thesis is the search for Nash equilibria in bimatrix games, but other examples arise from optimal power flow [51], the quadratic assignment problem [65], robotics and control [6], and statistics and machine learning [106, 104]. Even when the goal is not to minimize a polynomial function, optimization algorithms that involve minimizing Taylor expansions of functions solve POPs as a subroutine.

With such modeling power comes the price of computational intractability, and unfortunately POPs become intractable to solve even when the degrees of the defining polynomials are low. The study of local minima in Q1 and Q2 is in large part motivated by this intractability of finding global minima in polynomial optimization. It is common for optimization algorithms to instead search for local minima, with the hope that local minima are easier to find. This notion is not new; for example [72] provides an explicit example of a class of POP where global minima are hard to find but local minima are not. There has also been renewed interest in finding local minima due to the growth of machine learning applications, where local minima of highly nonconvex functions are sought for in practice with simple first-order methods like gradient descent. Our goal in Chapters 2 and 3 of this thesis is to more formally understand the complexity of finding local minima (as well as some related questions).

We point out that a priori there are no complexity implications between questions Q1 and Q2 stated above. For example, there is no reason to expect that an efficient algorithm for verifying that a given point is a local minimum would provide any guidance on how one decide if a problem has a local minimum (this is in essence the dilemma of the famous question “ $P = NP?$ ”). Conversely, even if local minimality of a given point cannot always be efficiently certified, that does not rule out the existence of algorithms that can efficiently find particular local minima that are easy to certify; see e.g. Question 3 of [87]. Thus the complexities of these two questions need to be studied separately.

¹This follows straightforwardly from the fact that NP-complete problems can be formulated as POPs; see for example Section 2.2, Section 3.2, or Section 4.2.

One of the first hardness results on local minima in the literature is due to Murty and Kabadi [75], who show that the problems of deciding whether a given point is a local minimum of a quadratic program or a local minimum of a quartic polynomial are NP-hard (see Table 1.1). These were two of the first results indicating that local minima are not necessarily “easier” than global minima. In regards to Q1, along with prior classical results, they only leave open the complexity of deciding whether a given point is a local minimum of a cubic polynomial. We show that this last case is polynomial-time solvable in Section 2.3 of this thesis.

$\begin{matrix} \text{max } \text{deg}(q_i) \\ \text{deg}(p) \end{matrix}$	\emptyset	1	≥ 2
1	P	P	NP-hard
2	P	NP-hard [75],[85]	NP-hard
3	P (Theorem 2.3.3)	NP-hard	NP-hard
≥ 4	NP-hard [75]	NP-hard	NP-hard

Table 1.1: Complexity of deciding whether a point is a local minimum of a POP, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table.³

We next comment on Q2, the question of deciding whether a POP has a local minimum. The complexity of Q2 based on the degrees of the defining polynomials is presented in Table 1.2. This problem has not been as extensively studied in prior literature, though it

³The NP-hardness of the case of linear objective and quadratic constraints is implied by the NP-hardness of the quadratic programming case. Indeed, minimizing a quadratic function $p(x)$ over a polyhedron $Ax = b$ can be reduced to minimizing a variable γ over the set $\{(x, \gamma) \mid Ax = b, p(x) = \gamma\}$.

is more closely related to the problem of searching for a local minimum as compared to Q1. In fact, for the cases labeled “P” in Table 1.2, some natural algorithms that find local minima implicitly check that they exist; see Section 2.1 for details in the unconstrained case. In Chapter 2 of this thesis, we show that this is also the case for cubic polynomials. In particular, the problem of deciding if a cubic polynomial has a local minimum (and then finding one) can be done by solving a polynomial number of polynomially-sized semidefinite programs (SDPs), hence the label in Table 1.2. By contrast, we show that Q2 is intractable in the same cases that Q1 is.

$\begin{matrix} \text{max } \text{deg}(q_i) \\ \text{deg}(p) \end{matrix}$	\emptyset	1	≥ 2
1	P	P	NP-hard
2	P	NP-hard (Theorem 3.2.5)	NP-hard
3	SDP (Algorithm 2)	NP-hard	NP-hard
≥ 4	NP-hard (Theorem 3.2.1)	NP-hard	NP-hard

Table 1.2: Complexity of deciding whether a POP has a local minimum, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table.

In many settings the existence of a local minimum in a POP is guaranteed; the focus is then on finding a local minimum without needing to consider whether one exists. One way this can be the case is when the feasible set is bounded, as is commonly the case in both applications and POPs encoding classical combinatorial problems. For the specific case of

quadratic programs with bounded feasible sets, the question of the complexity of finding a local minimum has appeared explicitly in the literature in [87]:

“What is the complexity of finding even a local minimizer for nonconvex quadratic programming, assuming the feasible set is compact? Murty and Kabadi (1987) and Pardalos and Schnitger (1988) have shown that it is NP-hard to test whether a given point for such a problem is a local minimizer, but that does not rule out the possibility that another point can be found that is easily verified as a local minimizer.”

We settle this question in Chapter 3, where we show that unless $P=NP$, no polynomial-time algorithm can even find a point within a Euclidean distance of c^n (for any constant $c \geq 0$) of a local minimum.

The final complexity question we study for POPs is testing the existence of an optimal solution when the optimal value of the problem is finite. This problem is in part motivated by a question of Nie, Dummel, and Sturmfels [84], who provide an algorithm for solving an unconstrained POP under the assumption that the optimal value is attained. The authors remark

“This assumption is non-trivial, and we do not address (the important and difficult) question of how to verify that a given polynomial $f(x)$ has this property.”

Most prior work on this question has focused on identifying cases where the existence of optimal solutions is guaranteed. Perhaps the most classical example is the Bolzano-Weirstrauss theorem for continuous functions over compact sets. Such theorems in the setting of POPs are commonly referred to as Frank-Wolfe type theorems, due to the eponymous authors’ result that quadratic programs attain their optimal value when that value is finite [38]. For POPs, this result was extended to cubic programs in [10], where it is shown that finite optimal values are always attained if the objective is at most cubic, and the constraints are affine. However, we will show that these are the only cases where this is true, and that

deciding whether a POP attains its optimal value is NP-hard in the remaining cases (see Table 1.3).

$\begin{matrix} \text{max } \text{deg}(q_i) \\ \text{deg}(p) \end{matrix}$	\emptyset	1	≥ 2
1	YES	YES	NP-hard (Theorem 4.2.2)
2	YES	YES [38]	NP-hard
3	YES	YES [10]	NP-hard
≥ 4	NP-hard (Theorem 4.2.1)	NP-hard	NP-hard

Table 1.3: Complexity of deciding whether a POP with a finite optimal value has an optimal solution, based on the degree of the objective p and the maximum degree of any constraint function q_i . Entries without a reference are either classical or implied by a stronger hardness result in the table. Note that whenever the degree of the objective is at most three, and the feasible set is a polyhedron, there is no algorithm required; the answer is simply ‘yes’.

In this thesis, we also study an application of POPs to the problem of finding Nash equilibria in bimatrix games. Nash equilibria are a fundamental concept in economics, but they also arise frequently in disciplines such as biology and finance. Finding a Nash equilibrium however, is computationally intractable [31]. In the final chapter of this thesis, we formulate the problem of finding Nash equilibria in bimatrix games as a POP, and explore semidefinite programming relaxations for finding approximate Nash equilibria. We also apply these techniques to certain decision problems related to Nash equilibria.

1.1 Preliminaries

In this thesis, we will study the complexity of Q1-Q3 in the Turing model of computation. Since polynomial functions of a given degree are finitely parameterized, they allow for a convenient study of complexity questions in this setting. The size of a given instance is determined by the number of bits required to write down the coefficients of the polynomial (and, in the case of Q1, the entries of the point x), which are all taken to be rational numbers. For the purposes of analyzing the complexity of these three questions for POPs, we consider the relevant setting in applications where the degrees of any polynomials are fixed and the number of variables in the POP increases. We are interested in the existence or non-existence of efficient algorithms for solving Q1-Q3 in this setting, as established theory (e.g. quantifier elimination theory [105, 100]) already yields exponential-time algorithms for them. We also point out that our intractability results are in the strong sense, meaning that the problems remain NP-hard even if the bitsize of all numerical data is $O(\log(n))$, where n is the number of variables in the problem. Unless $P = NP$, not even a pseudo-polynomial time algorithm (an algorithm whose running time is polynomial in the magnitude of the numerical data, but not their bitsize) can exist that solves a strongly NP-hard problem on all instances. This is in contrast to problems such as knapsack [39], which can be solved tractably, e.g. by dynamic programming, when the size of the numerical data is “small”. See [39] or [5, Section 2] for more details on the distinction between weakly and strongly NP-hard problems.

A prevalent tool in this thesis will be sum of squares programming. More details will be provided in each chapter as they are used, but we provide an introduction here. We say that a polynomial p is a *sum of squares* (sos) if there exist polynomials q_1, \dots, q_r such that $p = \sum_{i=1}^r q_i^2$. This is an algebraic sufficient condition for global nonnegativity of a polynomial which is in general not necessary [49]. While deciding nonnegativity of a polynomial is in general NP-hard (e.g., as a consequence of [75]), deciding whether a polynomial is sos can be done via semidefinite programming. This is because a polynomial p of degree $2d$ in n variables is a sum of squares if and only if there exists an $\binom{n+d}{d} \times \binom{n+d}{d}$ positive semidefinite

matrix Q satisfying the identity

$$p(x) = z(x)^T Q z(x), \tag{1.2}$$

where $z(x)$ denotes the vector of all monomials in x of degree less than or equal to d . Note that because of this equivalence, one can also impose the constraint that a polynomial p with unknown coefficients is sos by semidefinite programming (see, e.g., [88]). Given a rank- r psd matrix Q that satisfies (1.2), one can write Q as $\sum_{i=1}^r v_i v_i^T$ (e.g. via a Cholesky factorization), and obtain an sos decomposition of p as $p = \sum_{i=1}^r (v_i^T z(x))^2$.

Sum of squares polynomials have gained interest in the field of polynomial optimization because the problem of finding the infimum of some polynomial p can be straightforwardly reformulated into the problem

$$\begin{aligned} & \sup_{\gamma \in \mathbb{R}} \quad \gamma \\ & \text{subject to } p(x) - \gamma \text{ is a nonnegative polynomial.} \end{aligned} \tag{1.3}$$

This formulation has the interpretation of finding the largest lower bound on a polynomial. Unfortunately, this problem cannot be efficiently solved, since testing whether a polynomial of degree at least 4 is nonnegative is NP-hard and thus the constraint “ $p(x) - \gamma$ is a nonnegative polynomial” cannot be imposed in a tractable fashion. Therefore, to obtain what is known as a “sum of squares relaxation”, this constraint is replaced by a sum of squares constraint, which can be imposed tractably:

$$\begin{aligned} & \sup_{\gamma \in \mathbb{R}} \quad \gamma \\ & \text{subject to } p(x) - \gamma \text{ is sos.} \end{aligned} \tag{1.4}$$

Since the constraint “ $p(x)$ is sos” is a semidefinite constraint, (1.4) is an SDP. As any sos polynomial is nonnegative, this gives a lower bound on the infimum of $p(x)$.

There are many extensions for constrained problems, with one of the more well-known being based on Putinar's Positivstellensatz [92]. Putinar's Positivstellensatz states that if the so-called Archimedean property is satisfied, then if a polynomial p is positive for all x in the set $\{x \in \mathbb{R}^n \mid q_i(x) \geq 0, \forall i = 1, \dots, m\}$, there exist sos polynomials $\sigma_0, \dots, \sigma_m$ such that

$$p(x) = \sigma_0(x) + \sum_{i=1}^m \sigma_i(x)q_i(x). \quad (1.5)$$

The Archimedean property requires existence of a scalar R such that the polynomial $R - \sum_{i=1}^n x_i^2$ belongs to the *quadratic module* of q_1, \dots, q_m , i.e., the set of polynomials that can be written as

$$\tau_0(x) + \sum_{i=1}^m \tau_i(x)q_i(x),$$

where τ_0, \dots, τ_m are sum of squares polynomials. This is an algebraic notion of the compactness of the set $\{x \in \mathbb{R}^n \mid q_i(x) \geq 0, \forall i = 1, \dots, m\}$, which is stronger than the geometric notion. Similar to the construction of (1.4), for any positive integer d the following problem gives a lower bound on the optimal value of (1.1):

$$\begin{aligned} \gamma^d := \sup_{\gamma \in \mathbb{R}, \sigma_i} \quad & \gamma \\ \text{subject to} \quad & p(x) - \gamma = \sigma_0(x) + \sum_{i=1}^n \sigma_i(x)q_i(x), \end{aligned} \quad (1.6)$$

σ_i is an sos polynomial of degree at most $2d, \forall i = 0, \dots, m$.

Taking $d = 1, 2, \dots$ defines a sequence of problems referred to as the *Lasserre hierarchy* [61]. There are two primary properties of the Lasserre hierarchy which are of interest. The first is that for any fixed d , the d -th problem in this sequence is an SDP of size polynomial in n . The second is that under the Archimedean property, $\lim_{d \rightarrow \infty} \gamma^d = p^*$, where p^* is the optimal value of the POP in (1.1). While this is a powerful property of the hierarchy, and in practice the hierarchy is often exact at low levels, in general the level of the hierarchy needed can be arbitrarily high, and the semidefinite programs involved become expensive very quickly.

Additionally, the Archimedean property is NP-hard to check, as we show in Section 4.3 of this thesis.

1.2 Outline of this thesis

Complexity Results in Polynomial Optimization Chapters 2-4 this thesis focus on fundamental problems in polynomial optimization from an algorithmic perspective. Chapter 2 concerns itself with the complexity of local minima and related notions in unconstrained polynomial optimization. In particular, it establishes that local minima and second-order points of cubic polynomials can be found by solving polynomially many semidefinite programs of polynomial size. In the negative direction, it establishes that the problems of deciding if quartic polynomials have second-order points and whether cubic polynomials have critical points are NP-hard. Notably, our approach for finding local minima of cubic polynomials relies on circumventing the search for critical points. Chapter 3 establishes our intractability results related to local minima in quadratic programming. In particular, we show that unless $P=NP$, no polynomial-time algorithm can find points within distance c^n , for any constant $c \geq 0$, of a local minimum in an n -variate quadratic program with a bounded feasible set. Chapter 4 focuses on the problem of testing attainment of optimal values, and settles the complexity of deciding whether a POP with a finite optimal value has an optimal solution.

Semidefinite Relaxations for Bimatrix Games In the second part of this thesis, we explore an application of semidefinite programming to game theory, in particular semidefinite programming relaxations for finding Nash equilibria in bimatrix games. We show that for a symmetric game, a symmetric 1/3-Nash equilibrium can be efficiently recovered from any rank-2 solution to this relaxation. We also present semidefinite programming relaxations for NP-hard decision problems related to Nash equilibria, such as that of finding the highest achievable welfare under any Nash equilibrium.

1.3 Related Publications

The material in this thesis is based on the following work:

Chapter 2: Ahmadi, A. A., and Zhang, J. Complexity Aspects of Local Minima and Related Notions. Available at arXiv:2008.06148.

Chapter 3: Ahmadi, A. A., and Zhang, J. On the Complexity of Finding a Local Minimizer of a Quadratic Function over a Polytope. Available at arXiv:2008.05558.

Chapter 4: Ahmadi, A. A., and Zhang, J. On the Complexity of Testing Attainment of the Optimal Value in Nonlinear Optimization. *Mathematical Programming*, 2019. Available at arXiv:1803.07683.

Chapter 5: Ahmadi, A. A., and Zhang, J. Semidefinite Programming and Nash Equilibria in Bimatrix Games. To appear in *INFORMS Journal on Computing*. Available at arXiv:1706.08550.

Chapter 2

On Local Minima and Related Notions in Unconstrained Polynomial Optimization

2.1 Introduction

In this chapter of the thesis, we address the complexity of questions Q1 and Q2 from Chapter 1, but more generally for the following four types of points for a given polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$:

- (i) a *critical point*, i.e., a point x where the gradient $\nabla p(x)$ is zero,
- (ii) a *second-order point*, i.e., a point x where $\nabla p(x) = 0$ and the Hessian $\nabla^2 p(x)$ is positive semidefinite (psd), i.e. has nonnegative eigenvalues,
- (iii) a *local minimum*, i.e., a point x for which there exists a scalar $\epsilon > 0$ such that $p(x) \leq p(y)$ for all y with $\|y - x\| \leq \epsilon$,
- (iv) a *strict local minimum*, i.e., a point x for which there exists a scalar $\epsilon > 0$ such that $p(x) < p(y)$ for all $y \neq x$ with $\|y - x\| \leq \epsilon$.

We note the following straightforward implications between (i)-(iv):

strict local minimum \Rightarrow local minimum \Rightarrow second-order point \Rightarrow critical point.

Notions (i)-(iv) appear ubiquitously in nonconvex continuous optimization as surrogates for global minima, since it is well understood that finding a global minimum of f is in general an intractable problem. With regard to each of these four notions, we restate Q1 and Q2:

Q1*: Given a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ and a point $x \in \mathbb{R}^n$, is x of a given type (i)-(iv)?

Q2*: Given a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, does p have a point of a given type (i)-(iv) (and if so, can one be found efficiently)?

As discussed in Chapter 1, a priori there are no complexity implications between these two questions.

Let us first comment on the complexity of Q1* and Q2* for some simple and classical cases. For Q1*, checking whether a given point is a critical point of a polynomial function (of any degree) can trivially be done in polynomial time simply by evaluating the gradient at that point. To check that a given point is a second-order point, one can additionally compute the Hessian matrix at that point and check that it is positive semidefinite. This can be done in polynomial time, e.g., by performing Gaussian pivot steps along the main diagonal of the matrix [76, Section 1.3.1] or by computing its characteristic polynomial and checking that the signs of its coefficients alternate [50, p. 403]. Since for affine or quadratic polynomials, any second-order point is a local minimum, the only remaining case of Q1* for them is that of strict local minima. Affine polynomials never have strict local minima, making the question uninteresting. A point is a strict local minimum of a quadratic polynomial if and only if it is a critical point and the associated Hessian matrix is positive definite (pd), i.e., has positive eigenvalues. The latter property can be checked in polynomial time, for example by computing the leading principal minors of the Hessian and checking that they are all positive. As for Q2*, the affine case is again uninteresting since there is a critical point

(which will also be a second-order point and a local minimum) if and only if the coefficients of all degree-one monomials are zero. For quadratic polynomials, since the entries of the gradient are affine, searching for critical points can be done in polynomial time by solving a linear system. A candidate critical point will be a second-order point (and a local minimum) if and only if the Hessian is psd, and a strict local minimum if and only if the Hessian is pd.

Other than the aforementioned cases, the only prior result in the literature that we are aware of is due to Murty and Kabadi [75], which settles the complexity of $Q1^*$ for degree-4 polynomials. Our contribution in this chapter is to settle the complexity of the remaining cases for both $Q1^*$ and $Q2^*$. A summary of the results is presented in Tables 2.1 and 2.2. Entries denoted by “P” indicate that the problem can be solved in polynomial time. The notation “SDP” indicates that the problem of interest can be reduced to solving either one or polynomially-many semidefinite programs (SDP) whose sizes are polynomial in the size of the input. (In fact, the reduction also goes in the other direction for second-order points and local minima; see Theorems 2.5.3 and 2.5.4.) Finally, recall that a strong NP-hardness result implies that the problem of interest remains NP-hard even if the size (i.e. bit length) of the coefficients of the polynomial is $O(\log(n))$, where n is the number of variables. Therefore, unless $P=NP$, even a pseudo-polynomial time algorithm (i.e., an algorithm whose running time is polynomial in the magnitude of the coefficients, but not necessarily their bit length) cannot exist for the indicated problems in these tables.

Q1*: property vs. degree	1	2	3	≥ 4
Critical point	P	P	P	P
Second-order point	P	P	P	P
Local minimum	P	P	P (Theorem 2.3.3)	strongly NP-hard [75] ¹
Strict local minimum	P	P	P (Corollary 2.3.5)	strongly NP-hard [75] ¹

Table 2.1: Complexity of deciding whether a given point is of a certain type, based on the degree of the polynomial. Entries without a reference are classical.

Q2*: property vs. degree	1	2	3	≥ 4
Critical point	P	P	strongly NP-hard (Theorem 2.2.1)	strongly NP-hard (Theorem 2.2.1)
Second-order point	P	P	SDP (Corollary 2.6.5)	strongly NP-hard (Theorem 2.2.2)
Local minimum	P	P	SDP (Algorithm 2)	strongly NP-hard (Theorem 3.2.1)
Strict local minimum	P	P	SDP (Algorithm 2, Remark 2.6.1)	strongly NP-hard (Corollary 3.2.4)

Table 2.2: Complexity of deciding whether a polynomial has a point of a certain type, based on the degree of the polynomial. Entries without a reference are classical.

¹The proof in [75] is based on a reduction from the “matrix copositivity” problem. However, [75] only shows that this problem (and thus deciding if a quartic polynomial has a local minimum) is weakly NP-hard, since the reduction to matrix copositivity there is from the weakly NP-hard problem of Subset Sum. Nonetheless, their result can be strengthened by observing that testing matrix copositivity is in fact strongly NP-hard. This claim is implicit, e.g., in [35, Corollary 2.4]. The NP-hardness of testing whether a point is a strict local minimum of a quartic polynomial is not explicitly stated in [75], though it follows in the weak sense from the weak NP-hardness of Problem 8 of [75]. Again, with some work, this can be strengthened to a strong NP-hardness result.

The majority of the technical work in this chapter is spent on the case of cubic polynomials. It is somewhat surprising that many of the problems of interest to us are tractable for cubics, especially the search for local minima. This is in contrast to the intractability of other interesting problems related to cubic polynomials, for example, minimizing them over the unit sphere [79], or checking their convexity over a box [5]. It is also interesting to note that second-order points of cubic polynomials are easier to find than their critical points, despite being a more restrictive type of point. This shows that the right approach to finding second-order points involves bypassing the search for critical points as an initial step.

2.1.1 Organization and Main Contributions of the Chapter

Section 2.2 covers the NP-hardness results from Table 2.2. The remainder of the chapter is devoted to our results on cubic polynomials, which fills in the remaining entries of Tables 2.1 and 2.2. In Section 2.3, we give a characterization of local minima of cubic polynomials (Theorem 2.3.1) and show that it can be checked in polynomial time (Theorem 2.3.3). In Section 2.4, we give some geometric facts about local minima of cubic polynomials. For example, we show that the set of local minima of a cubic polynomial p is convex (Theorem 2.4.3), and we relate this set to the second-order points of p and to the set of minima of p over points where $\nabla^2 p$ is positive semidefinite (Theorem 2.4.7 and Theorem 2.4.10). In Section 2.4.4, we show that the interior of any spectrahedron is the projection of the local minima of some cubic polynomial (Theorem 2.4.12). In Section 2.5, we use this result to show that deciding if a cubic polynomial has a local minimum or a second-order point is at least as hard as some semidefinite feasibility problems.

In Section 2.6, we start from a “sum of squares” approach to finding second-order points of a cubic polynomial (Theorem 2.6.2 and Theorem 2.6.3), and build upon it (Section 2.6.3) to arrive at an efficient semidefinite representation of these points (Corollary 2.6.5). This also leads to an algorithm for finding local minima of cubic polynomials by solving polynomially-many SDPs of polynomial size (Algorithm 2). In Section 2.7, we take preliminary steps

towards some interesting future research directions, such as the design of an unregularized third-order Newton method that would use as a subroutine our algorithm for finding local minima of cubic polynomials (Section 2.7.2).

2.1.2 Preliminaries and Notation

We review some standard facts about local minima; more preliminaries specific to cubic polynomials appear in Section 2.3.1. Three well-known optimality conditions in unconstrained optimization are the *first-order necessary condition* (FONC), the *second-order necessary condition* (SONC), and the *second-order sufficient condition* (SOSC). Respectively, they are that the gradient at any local minimum is zero, the Hessian at any local minimum is psd, and that any critical point at which the Hessian is positive definite is a strict local minimum. A vector $d \in \mathbb{R}^n$ is said to be a *descent direction* for a function $p : \mathbb{R}^n \rightarrow \mathbb{R}$ at a point $\bar{x} \in \mathbb{R}^n$ if there exists a scalar $\epsilon > 0$ such that $p(\bar{x} + \alpha d) < p(\bar{x})$ for all $\alpha \in (0, \epsilon)$. Existence of a descent direction at a point clearly implies that the point is not a local minimum. However, in general, the lack of a descent direction at a point does not imply that the point is a local minimum (see, e.g., Example 2.3.2).

Next, we establish some basic notation which will be used throughout the chapter. We denote the set of $n \times n$ real symmetric matrices by $\mathbb{S}^{n \times n}$. For a matrix $M \in \mathbb{S}^{n \times n}$, the notation $M \succeq 0$ denotes that M is positive semidefinite, $M \succ 0$ denotes that it is positive definite, and $\text{Tr}(M)$ denotes its trace, i.e. the sum of its diagonal entries. For a matrix M , the notation $\mathcal{N}(M)$ denotes its null space, and $\mathcal{C}(M)$ denotes its column space. All vectors are taken to be column vectors. For two vectors x and y , the notation (x, y) denotes the vector $\begin{pmatrix} x \\ y \end{pmatrix}$. The notation 0_n denotes the vector of length n containing only zeros. The notation e_i denotes the i -th coordinate vector, i.e., the vector with a one in its i -th entry and zeros everywhere else.

2.2 NP-hardness Results

In this section, we present reductions that show our NP-hardness results from Tables 2.1 and 2.2. For concreteness, we construct these reductions from the (simple) MAXCUT problem, though our proof can work with any NP-hard problem that can be encoded by quadratic equations with “small enough” coefficients. Recall that in the (simple) MAXCUT problem, we are given as input an undirected and unweighted graph G on n vertices and an integer $k \leq n$. We are then asked whether there is a cut in G of size k , i.e. a partition of the vertices into two sets S_1 and S_2 such that the number of edges with one endpoint in S_1 and one endpoint in S_2 is equal to k . It is well known that the (simple) MAXCUT problem is strongly NP-hard [39].

If we denote the adjacency matrix of G by $E \in \mathbb{S}^{n \times n}$, it is straightforward to see that G has a cut of size k if and only if the following system of quadratic equations is feasible:

$$\begin{aligned} q_0(x) &:= \frac{1}{4} \sum_{i=1}^n \sum_{j=1}^n E_{ij} (1 - x_i x_j) - k = 0, \\ q_i(x) &:= x_i^2 - 1 = 0, i = 1, \dots, n. \end{aligned} \tag{2.1}$$

Indeed, the second set of constraints enforces each variable x_i to be -1 or 1 , and any $x \in \{-1, 1\}^n$ encodes a cut in G by assigning vertices with $x_i = 1$ to one side of the partition, and those with $x_i = -1$ to the other. Observe that with this encoding, $x_i x_j$ equals 1 whenever the two vertices i and j are on the same side and -1 otherwise. The size of the cut is therefore given by $\frac{1}{4} \sum_{i=1}^n \sum_{j=1}^n E_{ij} (1 - x_i x_j)$, noting that every edge is counted twice.

Theorem 2.2.1. *It is strongly NP-hard to decide whether a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ of degree greater than or equal to three has a critical point.*

Proof. Let $d \geq 3$ be fixed. Given an instance of the (simple) MAXCUT problem with a graph on n vertices, let the quadratic polynomials q_0, \dots, q_n be as in (2.1), and consider the

following degree- d polynomial in $2n + 2$ variables $(x_1, \dots, x_n, y_0, y_1, \dots, y_n, z)$:

$$p(x, y, z) = \sum_{i=0}^n y_i q_i(x) + z^d.$$

Note that all coefficients of this polynomial take $O(\log(n))$ bits to write down. We show that $p(x, y, z)$ has a critical point if and only if the quadratic system $q_0(x) = 0, \dots, q_n(x) = 0$ is feasible. Observe that the gradient of p is given by

$$\begin{pmatrix} \frac{\partial p}{\partial x} \\ \frac{\partial p}{\partial y} \\ \frac{\partial p}{\partial z} \end{pmatrix} = \begin{pmatrix} \sum_{i=0}^n y_i \frac{\partial q_i}{\partial x_1}(x) \\ \vdots \\ \sum_{i=0}^n y_i \frac{\partial q_i}{\partial x_n}(x) \\ q_0(x) \\ \vdots \\ q_n(x) \\ dz^{d-1} \end{pmatrix}.$$

If $\bar{x} \in \mathbb{R}^n$ is a solution to (2.1), then the point $(\bar{x}, 0_{n+1}, 0)$ is a critical point of p . Conversely, if $(\bar{x}, \bar{y}, \bar{z})$ is a critical point of p , then, since $\frac{\partial p}{\partial y}(\bar{x}, \bar{y}, \bar{z}) = 0$, \bar{x} must be a solution to (2.1). \square

Theorem 2.2.2. *It is strongly NP-hard to decide whether a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ of degree greater than or equal to four has a second-order point.*

Proof. Let $d \geq 4$ be fixed. Given an instance of the (simple) MAXCUT problem with a graph on n vertices, let the quadratic polynomials q_0, \dots, q_n be as in (2.1), and consider the following degree- d polynomial in $3n + 3$ variables $(x_1, \dots, x_n, y_0, y_1, \dots, y_n, z_0, z_1, \dots, z_n, w)$:

$$p(x, y, z, w) = \sum_{i=0}^n (y_i^2 q_i(x) - z_i^2 q_i(x)) + w^d.$$

Note that all coefficients of this polynomial take $O(\log(n))$ bits to write down. We show that $p(x, y, z, w)$ has a second-order point if and only if the quadratic system $q_0(x) = 0, \dots, q_n(x) = 0$ is feasible.

Observe that $\frac{\partial^2 p}{\partial y^2}$ is an $(n+1) \times (n+1)$ diagonal matrix with $2q_0(x), \dots, 2q_n(x)$ on its diagonal. Similarly, $\frac{\partial^2 p}{\partial z^2}$ is an $(n+1) \times (n+1)$ diagonal matrix with $-2q_0(x), \dots, -2q_n(x)$ on its diagonal. Suppose first that $(\bar{x}, \bar{y}, \bar{z}, \bar{w})$ is a second-order point of p . Since $\nabla^2 p(\bar{x}, \bar{y}, \bar{z}, \bar{w}) \succeq 0$, and since $\frac{\partial^2 p}{\partial y^2}$ and $\frac{\partial^2 p}{\partial z^2}$ are both principal submatrices of $\nabla^2 p$, it must be that $q_0(\bar{x}) = 0, \dots, q_n(\bar{x}) = 0$.

Now suppose that $\bar{x} \in \mathbb{R}^n$ is a solution to (2.1). We show that $(\bar{x}, 0_{n+1}, 0_{n+1}, 0)$ is a second-order point of p . Note that $\frac{\partial p}{\partial x}$ is quadratic in y and z , $\frac{\partial p}{\partial y}$ is linear in y , $\frac{\partial p}{\partial z}$ is linear in z , and $\frac{\partial p}{\partial w} = dw^{d-1}$. Thus $(\bar{x}, 0_{n+1}, 0_{n+1}, 0)$ is a critical point of p . Now observe that the entries of $\frac{\partial^2 p}{\partial x^2}$ are quadratic in y and z or are zero, the entries of $\frac{\partial^2 p}{\partial x \partial y}$ are linear in y or are zero, the entries of $\frac{\partial^2 p}{\partial x \partial z}$ are linear in z or are zero, $\frac{\partial^2 p}{\partial w^2} = d(d-1)w^{d-2}$, $\frac{\partial^2 p}{\partial y^2}(\bar{x}, 0_{n+1}, 0_{n+1}, 0)$ and $\frac{\partial^2 p}{\partial z^2}(\bar{x}, 0_{n+1}, 0_{n+1}, 0)$ are both zero, and all other entries of $\nabla^2 p$ are zero. Thus $\nabla^2 p(\bar{x}, 0_{n+1}, 0_{n+1}, 0) = 0$, and we conclude that $(\bar{x}, 0_{n+1}, 0_{n+1}, 0)$ is a second-order point of p . \square

The remaining two NP-hardness results from Table 2.2 are stated next, but proven in Chapter 3, since a corollary of them is the main result of that chapter.

Theorem 2.2.3. *It is strongly NP-hard to decide whether a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ of degree greater than or equal to four has a local minimum. The same statement holds for testing existence of a strict local minimum.*

2.3 Checking Local Minimality of a Point for a Cubic Polynomial

As the reader can observe from Tables 2.1 and 2.2 from Section 2.1, the remaining entries all have to do with the case of cubic polynomials. To answer these questions about cubics, we start in this section by showing that the problem of deciding if a given point is a local minimum (or a strict local minimum) of a cubic polynomial is polynomial-time solvable. This answers the remaining cases in Table 2.1. We first make certain observations about cubic polynomials that will be used throughout the chapter.

2.3.1 Preliminaries on Cubic Polynomials

It is easy to observe that a univariate cubic polynomial has either no local minima, exactly one local minimum (which is strict), or infinitely many non-strict local minima (in the case that the polynomial is constant). Further observe that if a point $\bar{x} \in \mathbb{R}^n$ is a (strict) local minimum of a function $p : \mathbb{R}^n \rightarrow \mathbb{R}$, then for any fixed point $\bar{y} \in \mathbb{R}^n$, the restriction of p to the line going through \bar{x} and \bar{y} —i.e. the univariate function $q(\alpha) := p(\bar{x} + \alpha(\bar{y} - \bar{x}))$ —has a (strict) local minimum at $\alpha = 0$. Since the restriction of a multivariate cubic polynomial to any line is a univariate polynomial of degree at most three, the previous two facts imply that (i) if a cubic polynomial has a strict local minimum, then it must be the only local minimum (strict or non-strict), and that (ii) if a cubic polynomial has multiple local minima, then the polynomial must be constant on the line connecting any two of these (necessarily non-strict) local minima.

Observe that for any cubic polynomial p , the error term of the second-order Taylor expansion is given by the cubic homogeneous component of p . More formally, for any point $\bar{x} \in \mathbb{R}^n$ and direction $v \in \mathbb{R}^n$,

$$p(\bar{x} + \lambda v) = p_3(v)\lambda^3 + \frac{1}{2}v^T \nabla^2 p(\bar{x})v\lambda^2 + \nabla p(\bar{x})^T v\lambda + p(\bar{x}), \quad (2.2)$$

where p_3 is the collection of terms of p of degree exactly 3.

Note that the Hessian of any cubic n -variate polynomial is an affine matrix of the form $\sum_{i=1}^n x_i H_i + Q$, where H_i and Q are all $n \times n$ symmetric matrices and the H_i satisfy

$$(H_i)_{jk} = (H_j)_{ik} = (H_k)_{ij} \quad (2.3)$$

for any $i, j, k \in \{1, \dots, n\}$. This is because an $n \times n$ symmetric matrix $A(x) := A(x_1, \dots, x_n)$ is a valid Hessian matrix if and only if $\frac{\partial}{\partial x_i} A_{jk}(x) = \frac{\partial}{\partial x_j} A_{ik}(x) = \frac{\partial}{\partial x_k} A_{ij}(x)$ for all $i, j, k \in \{1, \dots, n\}$. If $\sum_{i=1}^n x_i H_i + Q$ is a valid Hessian matrix, then the cubic polynomial which gives rise to it is of the form

$$\frac{1}{6} \sum_{i=1}^n x^T x_i H_i x + \frac{1}{2} x^T Q x + b^T x + c. \quad (2.4)$$

In this chapter, it is sometimes convenient for us to parametrize a cubic polynomial in the above form. As the scalar term in (2.4) is irrelevant for deciding local minimality or finding local minima, in the remainder of this chapter, we take $c = 0$ without loss of generality. Observe that the gradient of the polynomial in (2.4) is $\frac{1}{2} \sum_{i=1}^n x_i H_i x + Qx + b$, or equivalently a vector whose i -th entry is $\frac{1}{2} x^T H_i x + e_i^T Qx + b_i$.

2.3.2 Local Minimality of a Point for a Cubic Polynomial

In this section, we give a characterization of local minima of cubic polynomials and show that this characterization can be checked in polynomial time. Recall that we use the notation p_3 to denote the cubic homogeneous component of a cubic polynomial p , and $\mathcal{N}(M)$ (resp. $\mathcal{C}(M)$) to denote the null space (resp. column space) of a matrix M .

Theorem 2.3.1. *A point $\bar{x} \in \mathbb{R}^n$ is a local minimum of a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ if and only if the following three conditions hold:*

- $\nabla p(\bar{x}) = 0$,

- $\nabla^2 p(\bar{x}) \succeq 0$,
- $\nabla p_3(d) = 0, \forall d \in \mathcal{N}(\nabla^2 p(\bar{x}))$.

Note that the first two conditions are the well-known FONC and SONC. Throughout the chapter, we refer to the third condition as the *third-order condition (TOC)* for optimality. This condition is requiring the gradient of the cubic homogeneous component of p to vanish on the null space of the Hessian of p at \bar{x} . We remark that the FONC, SONC, and TOC together are in general neither sufficient nor necessary for a point to be a local minimum of a polynomial of degree higher than three. The first claim is trivial (consider, e.g., $p(x) = x^5$ at $x = 0$); for the second claim see Example 2.3.3.

Remark 2.3.1. It is straightforward to see that any local minimum \bar{x} of a cubic polynomial p satisfies a condition similar to the TOC, that $p_3(d) = 0, \forall d \in \mathcal{N}(\nabla^2 p(\bar{x}))$. Indeed, if \bar{x} is a second-order point and $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$, then Equation (2.2) gives $p(\bar{x} + \lambda d) = p_3(d)\lambda^3 + p(\bar{x})$. Hence, if $p_3(d)$ is nonzero, then either d or $-d$ is a descent direction for p at \bar{x} , and so \bar{x} cannot be a local minimum. This observation was made in [9] for three-times differentiable functions, and is referred to as the “third-order necessary condition” (TONC) for optimality. Note that because p_3 is homogeneous of degree three, from Euler’s theorem for homogeneous functions we have $3p_3(x) = x^T \nabla p_3(x)$. We can then see that $\nabla p_3(d) = 0 \Rightarrow p_3(d) = 0$, and therefore the TOC is a stronger condition than the TONC. Indeed, the FONC, SONC, and TONC together are not sufficient for local optimality of a point for a cubic polynomial; see Example 2.3.2. Intuitively, this is because the FONC, SONC, and TONC together avoid existence of a descent direction for cubic polynomials, but as the proof of Theorem 2.3.1 will show, existence of a “descent parabola” must also be avoided.

We will need the following fact from linear algebra for the proof of Theorem 2.3.1.

Lemma 2.3.2. *Let $M \in \mathbb{S}^{n \times n}$ be a symmetric positive semidefinite matrix and denote its smallest positive eigenvalue by λ_+ . Then if $z \in \mathcal{C}(M)$ and $\|z\| = 1, z^T M z \geq \lambda_+$.*

Proof. Suppose M has eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > \lambda_{k+1} = \dots = \lambda_n = 0$ (so $\lambda_+ = \lambda_k$). Let v_1, \dots, v_n be a set of corresponding mutually orthogonal unit-norm eigenvectors of M . Observe that any $z \in \mathcal{C}(M)$ can be written as $z = \sum_{i=1}^n \alpha_i v_i$, for some scalars α_i with $\alpha_i = 0$ for $i = k+1, \dots, n$. This is because the column space is orthogonal to the null space, and the eigenvectors corresponding to zero eigenvalues span the null space.

Since v_1, \dots, v_k are mutually orthogonal unit vectors, we have

$$z^T M z = \left(\sum_{i=1}^k \alpha_i v_i \right)^T \left(\sum_{i=1}^k \lambda_i v_i v_i^T \right) \left(\sum_{i=1}^k \alpha_i v_i \right) = \sum_{i=1}^k \alpha_i^2 \lambda_i v_i^T v_i v_i^T v_i = \sum_{i=1}^k \alpha_i^2 \lambda_i,$$

and

$$1 = \|z\|^2 = \left(\sum_{i=1}^k \alpha_i v_i \right)^T \left(\sum_{i=1}^k \alpha_i v_i \right) = \sum_{i=1}^k \alpha_i^2 v_i^T v_i = \sum_{i=1}^k \alpha_i^2.$$

These two equations combined imply that $z^T M z \geq \lambda_k = \lambda_+$. \square

Proof (of Theorem 2.3.1). As any local minimum must satisfy the FONC and SONC, it suffices to show that a second-order point is a local minimum for a cubic polynomial if and only if it also satisfies the TOC.

We first observe that for any second-order point \bar{x} , scalars α and β , and vectors $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$ and $z \in \mathbb{R}^n$, the following identity holds:

$$\begin{aligned} p(\bar{x} + \alpha d + \beta z) &= p_3(\alpha d + \beta z) + \frac{1}{2}(\alpha d + \beta z)^T \nabla^2 p(\bar{x})(\alpha d + \beta z) + p(\bar{x}) \\ &= \beta^3 p_3(z) + \frac{\beta^2}{2} z^T \nabla^2 p_3(\alpha d) z + \beta \nabla p_3(\alpha d)^T z + p_3(\alpha d) + \frac{\beta^2}{2} z^T \nabla^2 p(\bar{x}) z + p(\bar{x}) \\ &= \beta^3 p_3(z) + \frac{\alpha \beta^2}{2} z^T \nabla^2 p_3(d) z + \alpha^2 \beta \nabla p_3(d)^T z + \alpha^3 p_3(d) + \frac{\beta^2}{2} z^T \nabla^2 p(\bar{x}) z + p(\bar{x}). \end{aligned} \tag{2.5}$$

The first equality follows from (2.2) and the FONC. The second equality follows from the Taylor expansion of $p_3(\alpha d + \beta z)$ around αd and using the fact that $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$. The last equality follows from homogeneity of p_3 .

(second-order point) + TOC \Rightarrow local minimum:

Let \bar{x} be any second-order point at which the TOC holds. Note that any vector $v \in \mathbb{R}^n$ can be written as $\alpha d + \beta z$ for some (unique) scalars α and β , and unit vectors $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$ and $z \in \mathcal{C}(\nabla^2 p(\bar{x}))$. Since from the TOC we have $\nabla p_3(d) = 0$ (which also implies that $p_3(d) = 0$, as seen e.g. by Euler's theorem for homogeneous functions mentioned above), the identity in (2.5) reduces to

$$p(\bar{x} + v) - p(\bar{x}) = \beta^2 \left(\beta p_3(z) + \frac{\alpha}{2} z^T \nabla^2 p_3(d) z + \frac{1}{2} z^T \nabla^2 p(\bar{x}) z \right). \quad (2.6)$$

Let $\lambda > 0$ be the smallest nonzero eigenvalue of $\nabla^2 p(\bar{x})$. From Lemma 2.3.2 we have that $z^T \nabla^2 p(\bar{x}) z \geq \lambda$. Thus, if α and β satisfy

$$|\alpha| + |\beta| \leq \lambda \cdot \left(\max_{\|z\|=1, \|d\|=1} \max\{z^T \nabla^2 p_3(d) z, 2p_3(z)\} \right)^{-1}, \quad (2.7)$$

the expression on the right-hand side of (2.6) is nonnegative. As the set $\{\|z\| = 1\} \cap \{\|d\| = 1\}$ is compact and p_3 is continuous and odd, the quantity

$$\gamma := \max_{\|z\|=1, \|d\|=1} \max\{z^T \nabla^2 p_3(d) z, 2p_3(z)\}$$

is finite and nonnegative, and thus λ/γ is positive (or potentially $+\infty$). Finally, note that for any $v \in \mathbb{R}^n$ such that $\|v\| \leq \lambda/\gamma$, the corresponding α and β satisfy (2.7), and thus $p(\bar{x} + v) - p(\bar{x}) \geq 0$ as desired.

Local minimum \Rightarrow TOC:

Note that if \bar{x} is a local minimum, then we must have $p_3(d) = 0$ whenever $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$ (see Remark 2.3.1). We also assume that p_3 is not the zero polynomial, as then the TOC would be automatically satisfied.

Now suppose for the sake of contradiction that there exists a vector $\hat{d} \in \mathcal{N}(\nabla^2 p(\bar{x}))$ such that $\nabla p_3(\hat{d}) \neq 0$. Consider the sequence of points given by

$$\hat{x}_i := \bar{x} + \alpha_i \hat{d} + \beta_i z, \quad (2.8)$$

where

$$z = -\frac{\nabla p_3(\hat{d})}{\|\nabla p_3(\hat{d})\|}, \alpha_i = \frac{1}{i} \sqrt{\frac{z^T \nabla^2 p(\bar{x}) z}{|\nabla p_3(\hat{d})^T z|}}, \beta_i = \frac{1}{i^2}.$$

Observe that $\hat{x}_i \rightarrow \bar{x}$ as $i \rightarrow \infty$. From (2.5), we have

$$p(\bar{x} + \alpha_i \hat{d} + \beta_i z) - p(\bar{x}) = p_3(z) \beta_i^3 + \frac{1}{2} z^T \nabla^2 p_3(\hat{d}) z \alpha_i \beta_i^2 + \nabla p_3(\hat{d})^T z \alpha_i^2 \beta_i + \frac{1}{2} z^T \nabla^2 p(\bar{x}) z \beta_i^2.$$

Note that because $\alpha_i \propto \sqrt{\beta_i}$, the third and fourth terms of the right-hand side of the above expression will be the dominant terms as $i \rightarrow \infty$. For our choices of α_i and β_i , the sum of these two dominant terms simplifies to $-\frac{1}{2i^4} z^T \nabla^2 p(\bar{x}) z$. Observe that for any $w \in \mathcal{N}(\nabla^2 p(\bar{x}))$ and any $\alpha \in \mathbb{R}$, $p_3(\hat{d} + \alpha w) = 0$. Since the gradient of p_3 is orthogonal to its level sets, we must then have $\nabla p_3(\hat{d})^T w = 0$ for any $w \in \mathcal{N}(\nabla^2 p(\bar{x}))$. Thus, $\nabla p_3(\hat{d})$ is in the orthogonal complement of $\mathcal{N}(\nabla^2 p(\bar{x}))$, i.e. in $\mathcal{C}(\nabla^2 p(\bar{x}))$, and hence $z^T \nabla^2 p(\bar{x}) z > 0$. Thus, for any sufficiently large i , $p(\hat{x}_i) < p(\bar{x})$, and so \bar{x} is not a local minimum. □

Remark 2.3.2. Note that the points \hat{x}_i constructed in (2.8) trace a parabola as i ranges from $-\infty$ to $+\infty$. Thus as a corollary of the proof of Theorem 2.3.1, we see that if a point $\bar{x} \in \mathbb{R}^n$ is not a local minimum of a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, then there must exist a “descent parabola” that certifies that; i.e. a parabola $q(t) : \mathbb{R} \rightarrow \mathbb{R}^n$ and a scalar $\bar{\alpha}$ satisfying $q(0) = \bar{x}$ and $p(q(\alpha)) < p(\bar{x})$ for all $\alpha \in (0, \bar{\alpha})$.

Theorem 2.3.1 gives rise to the following algorithmic result.

Theorem 2.3.3. *Local minimality of a point $\bar{x} \in \mathbb{R}^n$ for a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ can be checked in polynomial time.*

Proof. In view of Theorem 2.3.1, we show that the FONC, SONC, and TOC can be checked in polynomial time (in the Turing model of computation). Checking that the gradient of p vanishes at \bar{x} and that the Hessian at \bar{x} is positive semidefinite can be done in polynomial time as explained in Section 2.1. We give the following polynomial-time algorithm for checking the TOC:

Algorithm 1 Algorithm for checking the TOC.

- 1: **Input:** Coefficients of a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, a point $\bar{x} \in \mathbb{R}^n$
 - 2: Compute $\nabla^2 p(\bar{x})$
 - 3: Compute a rational basis $\{v_1, \dots, v_k\}$ for the null space of $\nabla^2 p(\bar{x})$
 - 4: Check if coefficients of $g(\lambda) := \nabla p_3(\sum_{i=1}^k \lambda_i v_i)$ are all zero
 - 5: **if** YES
 - 6: \bar{x} is a local minimum of p
 - 7: **if** NO
 - 8: \bar{x} is a not local minimum of p
-

Note that the entries of the function $g : \mathbb{R}^k \rightarrow \mathbb{R}^n$ that appears in this algorithm are homogeneous quadratic polynomials in $\lambda := (\lambda_1, \dots, \lambda_k)$, where k is the dimension of $\mathcal{N}(\nabla^2 p(\bar{x}))$. For the TOC to hold, g must be zero for all $\lambda \in \mathbb{R}^k$, which happens if and only if all coefficients of every entry of g are zero.

A rational basis for the null space of a symmetric matrix can be computed in polynomial time, for example through the Bareiss algorithm [15]. For completeness, we give a less efficient but also polynomial-time algorithm which solves a series of linear systems. The first linear system finds a nonzero vector $v_1 \in \mathbb{R}^n$ such that $\nabla^2 p(\bar{x})^T v_1 = 0$. The successive linear systems solve for nonzero vectors $v_i \in \mathbb{R}^n$ such that $\nabla^2 p(\bar{x})^T v_i = 0, v_j^T v_i = 0, \forall j = 1, \dots, i-1$. To ensure nonzero solutions, some entry of the vector is fixed to 1, and if the system is infeasible, the next entry is fixed to 1 and the system is re-solved. Once the only feasible vector is the zero vector, the basis is complete.

The next step is to compute the coefficients of g . To do this, one can first compute the coefficients of ∇p_3 . There are $n \times \binom{n+1}{2}$ coefficients to compute, and each is a coefficient of p_3 , multiplied by 1, 2, or 3. If the m -th entry of ∇p_3 is given by $\sum_{i=1}^n \sum_{j \geq i}^n c_{ij} x_i x_j$, then the m -th entry of g is equal to $g_m(\lambda) = \sum_{a=1}^n \sum_{b=1}^n (\sum_{i=1}^n \sum_{j \geq i}^n c_{ij} (v_a)_i (v_b)_j) \lambda_a \lambda_b$, where the vectors $\{v_i\}$ are our rational basis for $\mathcal{N}(\nabla^2 p(\bar{x}))$. Observe that g_m is a polynomial in λ whose coefficients can be computed with a polynomial number of additions and multiplications over polynomially-sized scalars, and thus checking if all these coefficients are zero for every m can be done in polynomial time. □

Let us end this subsection by also giving an efficient characterization of strict local minima of cubic polynomials.

Corollary 2.3.4. *A point $\bar{x} \in \mathbb{R}^n$ is a strict local minimum of a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ if and only if*

- $\nabla p(\bar{x}) = 0$,
- $\nabla^2 p(\bar{x}) \succ 0$.

Proof. The fact that these two conditions are sufficient for local minimality is immediate from the SOSC. To show the converse, in view of the FONC, we only need to show that positive definiteness of the Hessian is necessary. Suppose for the sake of contradiction that for some nonzero vector $d \in \mathbb{R}^n$, we have $d^T \nabla^2 p(\bar{x}) d = 0$ (note that in view of the SONC, we cannot have $d^T \nabla^2 p(\bar{x}) d < 0$). From (2.2), we have $p(\bar{x} + \alpha d) = p(\bar{x}) + p_3(d) \alpha^3$. Hence, $\alpha = 0$ is not a strict local minimum of the univariate polynomial $p(\bar{x} + \alpha d)$, and so \bar{x} is not a strict local minimum of p . □

Corollary 2.3.5. *Strict local optimality of a point $\bar{x} \in \mathbb{R}^n$ for a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ can be checked in polynomial time.*

Proof. This follows from the characterization in Corollary 2.3.4. Checking the FONC is straightforward as before. As explained in Section 2.1, to check that $\nabla^2 p(\bar{x})$ is positive definite, one can equivalently check that all n leading principal minors of $\nabla^2 p(\bar{x})$ are positive. This procedure takes polynomial time since determinants can be computed in polynomial time. \square

2.3.3 Examples

We give a few illustrative examples regarding the application and context of Theorem 2.3.1.




Figure 2.1: Contour plots of $x_1^2 x_2$ (left) and $x_2^2 - x_1^2 x_2$ (right) from Examples 2.3.1 and 2.3.2. The polynomials are zero on the black lines, positive on the gray regions, and negative on the white regions. The dashed line in the right-side figure denotes a descent parabola at the origin.

Example 2.3.1. A cubic polynomial with local minima

Consider the polynomial $p(x_1, x_2) = x_1^2 x_2$. By inspection (see Figure 2.1), one can see that points of the type $\{(x_1, x_2) \mid x_1 = 0, x_2 > 0\}$ are local minima of p , as p is nonnegative when $x_2 > 0$, zero whenever $x_1 = 0$, and positive whenever $x_2 > 0$ and $x_1 \neq 0$. As a sanity check, we use Theorem 2.3.1 to verify that the point $(0, 1)$ is a local minimum of p (the same reasoning applies to all other local minima).

Through straightforward computation, we find

$$\nabla p(x) = \begin{pmatrix} 2x_1x_2 \\ x_1^2 \end{pmatrix}, \nabla p_3(x) = \begin{pmatrix} 2x_1x_2 \\ x_1^2 \end{pmatrix}, \nabla^2 p(x) = \begin{bmatrix} 2x_2 & 2x_1 \\ 2x_1 & 0 \end{bmatrix}.$$

We can see that the FONC and SONC are satisfied at $(0, 1)$. The null space of $\nabla^2 p(0, 1)$ is spanned by $(0, 1)$. We have

$$\nabla p_3 \left(\alpha \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) = \begin{pmatrix} 2(0)(\alpha) \\ (0)^2 \end{pmatrix} = 0,$$

which shows that the TOC is satisfied, verifying that $(0, 1)$ is a local minimum of p .

One can also verify that $\{(x_1, x_2) \mid x_1 = 0, x_2 > 0\}$ are the only local minima. Indeed, the critical points of p are those where $x_1 = 0$, and the second-order points are those where $x_1 = 0$ and $x_2 \geq 0$. To see that $(0, 0)$ is not a local minimum, observe that $(1, 1) \in \mathcal{N}(\nabla^2 p(0, 0))$, but $\nabla p_3(1, 1) = (2, 1) \neq 0$, and thus the TOC is violated.

Example 2.3.2. A cubic polynomial with no local minima

We use Theorem 2.3.1 to show that the polynomial $p(x_1, x_2) = x_2^2 - x_1^2x_2$ has no local minima. We have

$$\nabla p(x) = \begin{pmatrix} -2x_1x_2 \\ 2x_2 - x_1^2 \end{pmatrix}, \nabla p_3(x) = \begin{pmatrix} -2x_1x_2 \\ -x_1^2 \end{pmatrix}, \nabla^2 p(x) = \begin{bmatrix} -2x_2 & -2x_1 \\ -2x_1 & 2 \end{bmatrix}.$$

Observe that $(0, 0)$ is the only second-order point of p . The null space of $\nabla^2 p(0, 0)$ is spanned by $(1, 0)$. We have

$$\nabla p_3 \left(\alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = \begin{pmatrix} -2(\alpha)(0) \\ -(\alpha)^2 \end{pmatrix} = \begin{pmatrix} 0 \\ -\alpha^2 \end{pmatrix} \neq 0,$$

which shows that the TOC is violated, and hence $(0, 0)$ is not a local minimum. Note that the TONC is in fact satisfied at $(0, 0)$, since $p_3(\alpha, 0) = 0$ for any scalar α .

It is also interesting to observe that there are no descent directions for p at $(0, 0)$ (this is implied, e.g., by satisfaction of the TONC, along with the FONC and SONC). However, we can use the proof of Theorem 2.3.1 to compute a descent parabola, thereby more explicitly demonstrating that $(0, 0)$ is not a local minimum. The column space of $\nabla^2 p(0, 0)$ is spanned by $(0, 1)$. Then, following the proof of Theorem 2.3.1 with $z = (0, 1)$ and $\hat{d} = (1, 0)$, we have $z^T \nabla^2 p(0, 0) z = 2$ and $|\nabla p_3(\hat{d})^T z| = 1$. The parabola prescribed is then the set $\{(x_1, x_2) \mid x_2 = \frac{1}{2}x_1^2\}$. Indeed, one can now verify that except at $(0, 0)$, p is negative on the entire parabola; see the dashed line in Figure 2.1.

Example 2.3.3. A quartic polynomial with a local minimum that does not satisfy the TOC

We show in this example that for polynomials of degree higher than three, the TOC is not a necessary condition for local minimality. Consider the quartic polynomial given by $p(x_1, x_2) = 2x_1^4 + 2x_1^2x_2 + x_2^2$. The point $(0, 0)$ is a local minimum, as $p(0, 0) = 0$ and $p(x_1, x_2) = x_1^4 + (x_1^2 + x_2)^2$ is nonnegative. However, the Hessian of p at $(0, 0)$ is

$$\nabla^2 p(0, 0) = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix},$$

which has a null space spanned by $(1, 0)$. We observe that $\nabla p_3(x_1, x_2) = \begin{bmatrix} 4x_1x_2 \\ 2x_1^2 \end{bmatrix}$ does not vanish on this null space, as it evaluates, for example, to $(0, 2)$ at $(1, 0)$.

2.4 On the Geometry of Local Minima of Cubic Polynomials

We have shown that deciding local minimality of a given point for a cubic polynomial is a polynomial-time solvable problem. We now turn our attention to the remaining unresolved

entries in Table 2.2 from Section 2.1, which are on the problems of deciding whether a cubic polynomial has a second-order point, a local minimum, or a strict local minimum. In Sections 2.5 and 2.6, we will show that these problem can all be reduced to semidefinite programs of tractable size. In the current section, we present a number of geometric results about local minima and second-order points of cubic polynomials which are used in those sections, but are possibly of independent interest. For the remainder of this chapter, we use the notation SO_p to denote the set of second-order points of a polynomial p , LM_p to denote the set of its local minima, and \bar{S} to denote the closure of a set S .

2.4.1 Convexity of the Set of Local Minima

We begin by showing that for any cubic polynomial p , the set LM_p is convex. We go through two lemmas; the first is a simple algebraic observation, and the second contains information about some critical points. Recall that the Hessian of a cubic polynomial p written in the form of (2.4) is given by $\sum_{i=1}^n x_i H_i + Q$. Furthermore, its gradient is given by $\frac{1}{2} \sum_{i=1}^n x_i H_i x + Qx + b$, or equivalently a vector whose i -th entry is $x^T H_i x + e_i^T Qx + b_i$.

Lemma 2.4.1. *Let $H_1, \dots, H_n \subseteq \mathbb{S}^{n \times n}$ satisfy (2.3). Then for any two vectors $y, z \in \mathbb{R}^n$,*

$$\left(\sum_{i=1}^n y_i H_i \right) z = \left(\sum_{i=1}^n z_i H_i \right) y.$$

Proof. Observe that for any index $k \in \{1, \dots, n\}$, we have

$$\begin{aligned} \left(\left(\sum_{i=1}^n y_i H_i \right) z \right)_k &= \sum_{i=1}^n \sum_{j=1}^n (H_i)_{kj} y_i z_j \\ &= \sum_{i=1}^n \sum_{j=1}^n (H_j)_{ki} y_i z_j \\ &= \left(\left(\sum_{j=1}^n z_j H_j \right) y \right)_k, \end{aligned}$$

where the second equality follows from (2.3). □

Lemma 2.4.2. *Let $\bar{x} \in \mathbb{R}^n$ be a local minimum of a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, and let $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$. Then for any scalar α , $\bar{x} + \alpha d$ is a critical point of p .*

Proof. Let p be given in our canonical form as $\frac{1}{6} \sum_{i=1}^n x^T x_i H_i x + \frac{1}{2} x^T Q x + b^T x$. We have

$$\begin{aligned}
\nabla p(\bar{x} + \alpha d) &= \left(\frac{1}{2} \sum_{i=1}^n \bar{x}_i H_i + \alpha d_i H_i \right) (\bar{x} + \alpha d) + Q(\bar{x} + \alpha d) + b \\
&= \left(\frac{1}{2} \sum_{i=1}^n \bar{x}_i H_i \right) \bar{x} + Q\bar{x} + b \\
&\quad + \frac{1}{2} \sum_{i=1}^n \alpha d_i H_i \bar{x} + \frac{1}{2} \sum_{i=1}^n \alpha \bar{x}_i H_i d + \alpha Q d \\
&\quad + \frac{\alpha^2}{2} \sum_{i=1}^n d_i H_i d \\
&= \nabla p(\bar{x}) + \alpha \nabla^2 p(\bar{x}) d + \alpha^2 \nabla p_3(d) \\
&= 0 + 0 + 0 = 0,
\end{aligned}$$

where the third equality follows from Lemma 2.4.1, and the last follows from the FONC and TOC. \square

Theorem 2.4.3. *The set of local minima of any cubic polynomial is convex.*

Proof. If for some cubic polynomial p , the set LM_p of its local minima is empty or a singleton, the claim is trivially established. Otherwise, let $\bar{x}, \bar{y} \in LM_p$ with $\bar{x} \neq \bar{y}$. Consider any convex combination $z := \bar{x} + \alpha(\bar{y} - \bar{x})$, where $\alpha \in (0, 1)$. We show that z satisfies the FONC, SONC, and TOC, and therefore by Theorem 2.3.1, $z \in LM_p$.

Note from (2.2) that the restriction of p to the line passing through \bar{x} and \bar{y} is

$$p(\bar{x} + \alpha(\bar{y} - \bar{x})) = p_3(\bar{y} - \bar{x})\alpha^3 + \frac{1}{2}(\bar{y} - \bar{x})^T \nabla^2 p(\bar{x})(\bar{y} - \bar{x})\alpha^2 + \nabla p(\bar{x})^T (\bar{y} - \bar{x})\alpha + p(\bar{x}).$$

Since this univariate cubic polynomial has two local minima at $\alpha = 0$ and $\alpha = 1$, it must be constant. In particular, the coefficient of α^2 must be zero, and because $\nabla^2 p(\bar{x})$ is psd, that implies $\bar{y} - \bar{x} \in \mathcal{N}(\nabla^2 p(\bar{x}))$. Hence, by Lemma 2.4.2, the FONC holds at z . To

show the SONC and TOC at z , note that because $\nabla^2 p(x)$ is affine in x , $\nabla^2 p(z)$ can be written as a convex combination of $\nabla^2 p(\bar{x})$ and $\nabla^2 p(\bar{y})$, both of which are psd. The SONC is then immediate. To see why the TOC holds, recall that the null space of the sum of two psd matrices is the intersection of the null spaces of the summand matrices. Thus $\mathcal{N}(\nabla^2 p(z)) \subseteq \mathcal{N}(\nabla^2 p(\bar{x}))$, and the TOC is satisfied. \square

As a demonstration of Theorem 2.4.3, Figure 2.2 shows the critical points and the local minima of the cubic polynomial

$$x_1^3 + 3x_1^2x_2 + 3x_1x_2^2 + x_2^3 - 3x_1 - 3x_2. \quad (2.9)$$

Note that the critical points form a nonconvex set, while the local minima constitute a convex subset of the critical points.




Figure 2.2: The critical points of the polynomial (2.9). One can verify that the set of critical points is $\{(x_1, x_2) \mid (x_1 + x_2)^2 = 1\}$, and that the set of local minima is $\{(x_1, x_2) \mid x_1 + x_2 = 1\}$. The points on the dashed line are local maxima.

Unlike the above example, LM_p (or even $\overline{LM_p}$ as LM_p is in general not closed) may not be a polyhedral² set for cubic polynomials. For instance, the polynomial

$$p(x_1, x_2, x_3, x_4) = -x_1x_3^2 + x_1x_4^2 + 2x_2x_3x_4 + x_3^2 + x_4^2, \quad (2.10)$$

has $LM_p = \{x \in \mathbb{R}^4 \mid x_1^2 + x_2^2 < 1, x_3 = x_4 = 0\}$ (see Figure 2.3). This is in contrast to quadratic polynomials, whose local minima always form a polyhedral set. We show in Theorem 2.4.5, however, that $\overline{LM_p}$ is always a spectrahedron³. We first need the following lemma.




Figure 2.3: The projection of the set of local minima of the polynomial in (2.10) onto the x_1 and x_2 variables. This example shows that $\overline{LM_p}$ is not always a polyhedral set.

Lemma 2.4.4. *For any cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, suppose $\bar{x} \in \mathbb{R}^n$ and $\bar{y} \in \mathbb{R}^n$ satisfy*

- $\bar{x} \in SO_p$,
- $\nabla^2 p(\bar{y}) \succeq 0$,

²Recall that a *polyhedron* is a set defined by finitely many affine inequalities.

³Recall that a *spectrahedron* is a set of the type $S = \{x \in \mathbb{R}^n \mid A_0 + \sum_{i=1}^n x_i A_i \succeq 0\}$, where A_0, \dots, A_n are symmetric matrices of some size $m \times m$ [110].

- $p(\bar{x}) = p(\bar{y})$.

Then $p(\bar{x} + \alpha(\bar{y} - \bar{x})) = p(\bar{x})$ for any scalar α , and $\bar{y} - \bar{x} \in \mathcal{N}(\nabla^2 p(\bar{x}))$.

Note in particular that this lemma applies if \bar{y} is simply a second-order point, since p must take the same value at any two second-order points. This is because any non-constant univariate cubic polynomial can have at most one second-order point.

Proof. Consider the Taylor expansion of p around \bar{x} in the direction $\bar{y} - \bar{x}$ (see (2.2)):

$$q(\alpha) := p(\bar{x} + \alpha(\bar{y} - \bar{x})) = p_3(\bar{y} - \bar{x})\alpha^3 + \frac{1}{2}(\bar{y} - \bar{x})^T \nabla^2 p(\bar{x})(\bar{y} - \bar{x})\alpha^2 + \nabla p(\bar{x})^T (\bar{y} - \bar{x})\alpha + p(\bar{x}).$$

Note that q is a univariate cubic polynomial which has a second-order point at $\alpha = 0$. It is straightforward to see that if a univariate cubic polynomial is not constant and has a second-order point, then any other point which takes the same function value as the second-order point must have a negative second derivative. As this is not the case for q (in view of $\alpha = 0$ and $\alpha = 1$), q must be constant, i.e., $p(\bar{x} + \alpha(\bar{y} - \bar{x})) = p(\bar{x})$ for any α . Now observe that for $p(\bar{x} + \alpha(\bar{y} - \bar{x}))$ to be constant, we must have $(\bar{y} - \bar{x})^T \nabla^2 p(\bar{x})(\bar{y} - \bar{x}) = 0$. As $\nabla^2 p(\bar{x}) \succeq 0$, we have $\bar{y} - \bar{x} \in \mathcal{N}(\nabla^2 p(\bar{x}))$. \square

Theorem 2.4.5. For a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, \overline{LM}_p is a spectrahedron.

Proof. If LM_p is empty, the claim is trivial. Otherwise, let $\bar{x} \in LM_p$. We show that \overline{LM}_p is given by the spectrahedron

$$M := \{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0, \nabla^2 p(\bar{x})(x - \bar{x}) = 0\}. \quad (2.11)$$

First consider any $\bar{y} \in LM_p$. From the SONC we know that $\nabla^2 p(\bar{y}) \succeq 0$ and from Lemma 2.4.4, we know that $\bar{y} - \bar{x} \in \mathcal{N}(\nabla^2 p(\bar{x}))$. Thus $\bar{y} \in M$. Since M is closed, we get that $\overline{LM}_p \subseteq M$.

Now consider any $\bar{y} \in M$. By the definition of M , \bar{y} satisfies the SONC, and by Lemma 2.4.2, it also satisfies the FONC. Since for any scalar $\alpha \in (0, 1)$, $\nabla^2 p(\bar{x} + \alpha(\bar{y} - \bar{x}))$ is a convex

combination of the two psd matrices $\nabla^2 p(\bar{x})$ and $\nabla^2 p(\bar{y})$, $\mathcal{N}(\nabla^2 p(\bar{x} + \alpha(\bar{y} - \bar{x}))) \subseteq \nabla^2 p(\bar{x})$ and thus $\bar{x} + \alpha(\bar{y} - \bar{x})$ satisfies the TOC (since \bar{x} does). Thus \bar{y} can be written as the limit of local minima of p (e.g. $\{\bar{x} + \alpha(\bar{y} - \bar{x})\}$ as $\alpha \rightarrow 1$). \square

Remark 2.4.1. We will soon show that for a cubic polynomial p , if LM_p is nonempty, then $\overline{LM_p} = SO_p$ (see Theorem 2.4.7). In Section 2.6, we will give other representations of SO_p , which in contrast to the representation in (2.11), do not rely on access to or even existence of a local minimum.

2.4.2 Local Minima and Solutions to a “Convex” Problem

In Section 2.6, we present an SDP-based approach for finding local minima of cubic polynomials. (We note again that the SDP representation in (2.11) is useless for this purpose as it already assumes access to a local minimum.) Many common approaches for computing local minima of twice-differentiable functions involve first finding critical points of the function, and then checking whether they satisfy second-order conditions. However, as discussed in the introduction and in Section 2.2, such approaches are unlikely to be effective for cubic polynomials as critical points of these functions are in fact NP-hard to find (see Theorem 2.2.1). Interestingly, however, we show in Section 2.6 that by bypassing the search for critical points, one can directly find second-order points and local minima of cubic polynomials by solving semidefinite programs of tractable size. The key to our approach is to relate the problem of finding a local minimum of a cubic polynomial p to the following optimization problem:

$$\begin{aligned} \inf_{x \in \mathbb{R}^n} \quad & p(x) \\ \text{subject to} \quad & \nabla^2 p(x) \succeq 0. \end{aligned} \tag{2.12}$$

The connection between solutions of (2.12) and local minima of p is established by Theorem 2.4.7 below. The feasible set of (2.12) has interesting geometric properties (see, e.g., Corollary 2.4.12) and will be referred to with the following terminology in the remainder of the chapter.

Definition 2.4.6. *The convexity region of a polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is the set*

$$CR_p := \{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0\}.$$

Observe that for any cubic polynomial, its convexity region is a spectrahedron, and thus a convex set. As p is a convex function when restricted to its convexity region, one can consider (2.12) to be a convex problem in spirit.

Theorem 2.4.7. *Let p be a cubic polynomial with a second-order point. Then the following sets are equivalent:*

(i) SO_p

(ii) Minima of (2.12).

Furthermore, if p has a local minimum, then these two sets are equivalent to:

(iii) \overline{LM}_p .

Proof. (i) \subseteq (ii).

Let $\bar{y} \in SO_p$ and \bar{x} be any feasible point to (2.12). If we consider the univariate cubic polynomial $q(\alpha) := p(\bar{x} + \alpha(\bar{y} - \bar{x}))$, i.e., the restriction of p to the line passing through \bar{x} and \bar{y} , we can see that $\alpha = 1$ is a second-order point of q . Note that if any univariate cubic polynomial has a second-order point, then that second-order point is a minimum of it over its convexity region. In particular, because \bar{x} is feasible to (2.12) and thus $\alpha = 0$ is in the convexity region of q , we have $p(\bar{y}) = q(1) \leq q(0) = p(\bar{x})$. As \bar{y} is feasible to (2.12) and has objective value no higher than any other feasible point, it must be optimal to (2.12).

(ii) \subseteq (i)

Let \bar{y} be a minimum of (2.12) (we know that such a point exists because we have shown SO_p is a subset of the minima of (2.12), and SO_p is nonempty by assumption). Let $\bar{x} \in SO_p$ and $d := \bar{y} - \bar{x}$. Observe that $p(\bar{y}) = p(\bar{x})$, and so by Lemma 2.4.4, we must have $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$.

It follows that $\nabla p(\bar{y}) = \nabla p_3(d)$ (cf. the proof of Lemma 2.4.2). Now suppose for the sake of contradiction that \bar{y} is not a second-order point. Since \bar{y} is feasible to (2.12), we must have $\nabla p(\bar{y}) = \nabla p_3(d) \neq 0$. As $p(\bar{x}) = p(\bar{x} + \alpha d)$ for any scalar α due to Lemma 2.4.4, we must have $p_3(d) = \frac{1}{6}d^T \nabla^2 p_3(d)d = 0$ (see (2.2)). Thus we can write

$$\begin{aligned}
(d - \alpha \nabla p_3(d))^T \nabla^2 p(\bar{y})(d - \alpha \nabla p_3(d)) &= (d - \alpha \nabla p_3(d))^T (\nabla^2 p(\bar{x}) + \nabla^2 p_3(d))(d - \alpha \nabla p_3(d)) \\
&= \alpha^2 \nabla p_3(d)^T \nabla^2 p(\bar{x}) \nabla p_3(d) - 2\alpha \nabla p_3(d)^T \nabla^2 p_3(d)^T d \\
&\quad + \alpha^2 \nabla p_3(d)^T \nabla^2 p_3(d) \nabla p_3(d) \\
&= \alpha^2 (\nabla p_3(d)^T \nabla^2 p(\bar{x}) \nabla p_3(d) + \nabla p_3(d)^T \nabla^2 p_3(d)^T \nabla p_3(d)) \\
&\quad - 4\alpha \nabla p_3(d)^T \nabla p_3(d),
\end{aligned}$$

where the last equality follows from that $\nabla p_3(d) = \frac{1}{2} \nabla^2 p_3(d)^T d$ due to Euler's theorem for homogeneous functions. Note that the right-hand side of the above expression is negative for sufficiently small $\alpha > 0$, and so $\nabla^2 p(\bar{y})$ is not psd, which contradicts feasibility of \bar{y} to (2.12).

For the second claim of the theorem, suppose that p has a local minimum. The following arguments will show $(i) = (ii) = (iii)$.

$$(iii) \subseteq (i)$$

Clearly any local minimum of p is a second-order point. Since the gradient and the Hessian of p are continuous in x and as the cone of psd matrices is closed, the limit of any convergent sequence of second-order points is a second-order point.

$$(ii) \subseteq (iii).$$

Let \bar{y} be any minimum of (2.12). Consider any local minimum \bar{x} of p and let $z_\alpha := \bar{x} + \alpha(\bar{y} - \bar{x})$. As both $\nabla^2 p(\bar{y})$ and $\nabla^2 p(\bar{x})$ are psd, any point z_α with $\alpha \in [0, 1)$ satisfies the SONC and TOC, by the same arguments as in the proof of Theorem 2.4.3.

Now note that since \bar{x} is a second-order point, it is also a minimum of (2.12) (as (i) \subseteq (ii)) and thus $p(\bar{y}) = p(\bar{x})$. From Lemma 2.4.4, we then have $\bar{y} - \bar{x} \in \mathcal{N}(\nabla^2 p(\bar{x}))$, and so from Lemma 2.4.2, z_α satisfies the FONC for any α . Thus, in view of Theorem 2.3.1, for any $\alpha \in [0, 1)$, z_α is a local minimum of p . Therefore \bar{y} can be written as the limit of a sequence of local minima (i.e., $\{z_\alpha\}$ as $\alpha \rightarrow 1$), and hence $\bar{y} \in \overline{LM_p}$. \square

Remark 2.4.2. Note that as a consequence of Theorems 2.4.5 and 2.4.7, if a cubic polynomial p has a local minimum, then SO_p is a spectrahedron. In fact, SO_p is a spectrahedron for *any* cubic polynomial p ; see Theorem 2.6.3. In that theorem, we will give a more useful spectrahedral representation of SO_p which does not rely on knowledge of a local minimum.

Corollary 2.4.8. *Let p be a cubic polynomial with a second-order point. Then the optimal value of (2.12) is the value that p takes at any of its second-order points (and in particular, at any of its local minima if they exist).*

Proof. This is immediate from the equivalence of (i) and (ii) in Theorem 2.4.7. \square

2.4.3 Distinction Between Local Minima and Second-Order Points

We have shown that the optimization problem in (2.12) gives an approach for finding second-order points of a cubic polynomial p without computing its critical points. However, not all second-order points are local minima, and so in this subsection, we characterize the difference between the two notions more precisely. We first recall the concept of the relative interior of a (convex) set (see, e.g., [97, Chap. 6]).

Definition 2.4.9. *The relative interior of a nonempty convex set $S \subseteq \mathbb{R}^n$ is the set*

$$ri(S) := \{x \in S \mid \forall y \in S, \exists \lambda > 1 \text{ s.t. } \lambda x + (1 - \lambda)y \in S\}.$$

This definition generalizes the notion of interior to sets which do not have full dimension. One can show that for a convex set S , $ri(S)$ is convex, $ri(\bar{S}) = ri(S)$, and $\overline{ri(S)} = \bar{S}$ [97].

In general, for a nonempty convex set S , we have $ri(\bar{S}) = ri(S) \subseteq S$, but we may not have $ri(\bar{S}) = S$. (For example, let S be a line segment with one of its endpoints removed.) It turns out, however, that for a cubic polynomial p with a local minimum, $ri(\overline{LM}_p) = LM_p$.

Theorem 2.4.10. *Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a cubic polynomial with a local minimum. Then the following three sets are equivalent:*

(i) LM_p

(ii) $ri(SO_p)$

(iii) Intersection of critical points of p with $ri(CR_p)$.

Proof. (ii) \subseteq (i)

Recall from Theorem 2.4.3 that LM_p is convex, and from Theorem 2.4.7 that $SO_p = \overline{LM}_p$. Then we have $ri(SO_p) = ri(\overline{LM}_p) = ri(LM_p) \subseteq LM_p$.

(i) \subseteq (ii)

We prove the contrapositive. Let \bar{x} be a point which is not in $ri(SO_p)$. If \bar{x} is not a second-order point, then it clearly cannot be a local minimum. Suppose now that $\bar{x} \in SO_p \setminus ri(SO_p)$. Then there is another second-order point \bar{y} such that $\bar{y} + \lambda(\bar{x} - \bar{y})$ is not a second-order point for any $\lambda > 1$. Note from Lemma 2.4.4 and the statement after it that $p(\bar{y} + \lambda(\bar{x} - \bar{y}))$ is a constant univariate function of λ . Now for any $\epsilon > 0$, define the point $\bar{z}_\epsilon := \bar{x} + \frac{\epsilon}{2\|\bar{x} - \bar{y}\|}(\bar{x} - \bar{y})$. Since \bar{z}_ϵ is not a second-order point and thus not a local minimum, there is a point z_ϵ satisfying $\|\bar{z}_\epsilon - z_\epsilon\| < \frac{\epsilon}{2}$ and

$$p(z_\epsilon) < p(\bar{z}_\epsilon) = p\left(\frac{\epsilon}{2\|\bar{x} - \bar{y}\|}(\bar{x} - \bar{y})\right) = p(\bar{x}).$$

Furthermore, by the triangle inequality, z_ϵ also satisfies $\|z_\epsilon - \bar{x}\| < \epsilon$. Thus, by considering $\{z_\epsilon\}$ as $\epsilon \rightarrow 0$, we can conclude that \bar{x} is not a local minimum.

(i) \subseteq (iii)

Consider any local minimum \bar{x} of p , which clearly must also be a critical point of p , and a member of CR_p . Suppose for the sake of contradiction that $\bar{x} \notin ri(CR_p)$. Then there exists

$y \in CR_p$ such that for any scalar $\alpha > 0$, $\nabla^2 p(\bar{x} + \alpha(\bar{x} - y))$ is not psd. In particular, for any $\alpha > 0$ there exists a unit vector $z_\alpha \in \mathbb{R}^n$ such that $z_\alpha^T \nabla^2 p(\bar{x} + \alpha(\bar{x} - y)) z_\alpha < 0$.

We now show that for any α , z_α can be taken to be in $\mathcal{C}(\nabla^2 p(\bar{x}))$. This is because, as we will show, if $z_\alpha = d + v$, where $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$ and $v \in \mathcal{C}(\nabla^2 p(\bar{x}))$,

$$(d + v)^T \nabla^2 p(\bar{x} + \alpha(\bar{x} - y))(d + v) = v^T \nabla^2 p(\bar{x} + \alpha(\bar{x} - y))v. \quad (2.13)$$

Observe that if p is written in the form (2.4), for any $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$, we have

$$\begin{aligned} d^T \nabla^2 p(\bar{x} + \alpha(\bar{x} - y))d &= d^T \left(\sum_{i=1}^n (\bar{x}_i + \alpha(\bar{x}_i - y_i)) H_i + Q \right) d \\ &= d^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) d + \alpha \sum_{i=1}^n (d^T H_i d) (\bar{x}_i - y_i) = 0, \end{aligned}$$

where the last equality follows from that $d \in \mathcal{N}(\nabla^2 p(\bar{x}))$, and the TOC, recalling that the i -th entry of $\nabla p_3(d)$ is $\frac{1}{2} d^T H_i d$. Note in particular that the expression above also holds for $\alpha = -1$, and so $d \in \mathcal{N}(\nabla^2 p(y))$. Now observe that because we can write

$$\nabla^2 p(\bar{x} + \alpha(\bar{x} - y)) = (1 + \alpha) \nabla^2 p(\bar{x}) - \alpha \nabla^2 p(y),$$

we have $\nabla^2 p(\bar{x} + \alpha(\bar{x} - y))d = 0$. Thus, we have shown (2.13), and we can take $z_\alpha \in \mathcal{C}(\nabla^2 p(\bar{x}))$.

Note that if $z_\alpha \in \mathcal{C}(\nabla^2 p(\bar{x}))$, then by Lemma 2.3.2 we have $z_\alpha^T \nabla^2 p(\bar{x}) z_\alpha \geq \lambda$, where λ is the smallest nonzero eigenvalue of $\nabla^2 p(\bar{x})$. Thus, for small enough α , the quantity $z_\alpha^T \nabla^2 p(\bar{x} + \alpha(\bar{x} - y)) z_\alpha$ is positive and so we arrive at a contradiction.

(iii) \subseteq (i)

Let \bar{x} be a critical point which is in $ri(CR_p)$. Clearly $\bar{x} \in SO_p$. Consider any local minimum

\bar{y} of p , and observe that for any $\alpha \neq 0$, we can write

$$\bar{x} = \frac{1}{\alpha}(\alpha\bar{x} + (1 - \alpha)\bar{y}) + \frac{\alpha - 1}{\alpha}\bar{y}. \quad (2.14)$$

As $\bar{x} \in \text{ri}(CR_p)$ and $\bar{y} \in CR_p$, $\alpha\bar{x} + (1 - \alpha)\bar{y} \in CR_p$ for some $\alpha > 1$. In particular, for that α , $\nabla^2 p(\alpha\bar{x} + (1 - \alpha)\bar{y}) \succeq 0$ and thus in view of (2.14), we can see that $\mathcal{N}(\nabla^2 p(\bar{x})) \subseteq \mathcal{N}(\nabla^2 p(\bar{y}))$. Hence, because the TOC holds at \bar{y} , it must also hold at \bar{x} . Thus \bar{x} is a local minimum. \square

Figure 2.4 demonstrates the relation between LM_p and SO_p for the polynomial $p(x_1, x_2) = x_1^2 x_2$. For this example, $SO_p = \{(x_1, x_2) \mid x_1 = 0, x_2 \geq 0\}$, and $LM_p = \{(x_1, x_2) \mid x_1 = 0, x_2 > 0\}$ (see Example 2.3.1).




Figure 2.4: The set of local minima (left) and second-order points (right) of the cubic polynomial $p(x_1, x_2) = x_1^2 x_2$. Note that SO_p is the closure of LM_p (Theorem 2.4.7) and LM_p is the relative interior of SO_p (Theorem 2.4.10).

Theorem 2.4.10 gives rise to the following interesting geometric fact about local minima of cubic polynomials.

Corollary 2.4.11. *Let \bar{x} and \bar{y} be two local minima of a cubic polynomial. Then*

$$\mathcal{N}(\nabla^2 p(\bar{x})) = \mathcal{N}(\nabla^2 p(\bar{y})).$$

Proof. It is known ([93, Corollary 1]) that for a spectrahedron $\{x \in \mathbb{R}^n \mid A_0 + \sum_{i=1}^n x_i A_i \succeq 0\}$ and any two points x and y in its relative interior, $\mathcal{N}(A_0 + \sum_{i=1}^n x_i A_i) = \mathcal{N}(A_0 + \sum_{i=1}^n y_i A_i)$. In view of the facts that for any cubic polynomial p , CR_p is a spectrahedron and $LM_p \subseteq ri(CR_p)$ (from Theorem 2.4.10), the result is immediate. \square

2.4.4 Spectrahedra and Convexity Regions of Cubic Polynomials

We end this section with a result relating general spectrahedra and convexity regions of cubic polynomials. Recall from the end of Section 2.3.1 that if $S := \{x \in \mathbb{R}^n \mid A_0 + \sum_{i=1}^n x_i A_i \succeq 0\}$ is a special spectrahedron, where A_0, \dots, A_n are $n \times n$ symmetric matrices satisfying

$$(A_i)_{jk} = (A_j)_{ik} = (A_k)_{ij}$$

for any $i, j, k \in \{1, \dots, n\}$, then S is the convexity region of the cubic polynomial

$$p(x) = \frac{1}{6} \sum_{i=1}^n x^T x_i A_i x + \frac{1}{2} x^T A_0 x.$$

The following theorem shows that if the number of variables is allowed to increase, then *any* spectrahedron can be represented by the convexity region of a cubic polynomial.

Theorem 2.4.12. *Let a spectrahedron $S \subseteq \mathbb{R}^n$ be given by $S := \{x \in \mathbb{R}^n \mid A_0 + \sum_{i=1}^n x_i A_i \succeq 0\}$, where $A_0, \dots, A_n \in \mathbb{S}^{m \times m}$. There exists a cubic polynomial p in at most $m + n$ variables such that S is a projection of its convexity region; i.e.,*

$$S = \{x \in \mathbb{R}^n \mid \exists y \in \mathbb{R}^m \text{ such that } (x, y) \in CR_p\}.$$

Furthermore, the interior of S is a projection of the set of local minima of p .

Proof. Let $A(x) := A_0 + \sum_{i=1}^n x_i A_i$. We first present a characterization of the interior of S following the developments in Section 2.4 of [93]. Let $\mathcal{N}_A := \mathcal{N}(A_0) \cap \dots \cap \mathcal{N}(A_n)$, and V

be a full-rank matrix whose columns span the orthogonal complement of \mathcal{N}_A . Suppose that \mathcal{N}_A is $(m - k)$ -dimensional. Then there exist matrices $B_0, \dots, B_n \in \mathbb{S}^{k \times k}$ with $\mathcal{N}(B_0) \cap \dots \cap \mathcal{N}(B_n) = \{0_k\}$ such that

$$B(x) := B_0 + \sum_{i=1}^n x_i B_i = V^T A(x) V.$$

In [93, Corollary 5], it is shown that $B(x)$

$$\{x \in \mathbb{R}^n \mid A(x) \succeq 0\} = \{x \in \mathbb{R}^n \mid B(x) \succeq 0\} \quad (2.15)$$

and that the set $\{x \in \mathbb{R}^n \mid B(x) \succ 0\}$ gives the interior of S . Now consider the following cubic polynomial in $n + k$ variables:

$$p(x, y) := y^T B(x) y. \quad (2.16)$$

Observe that the partial derivative of p with respect to y is $2B(x)y$, the partial derivative of p with respect to x_i is $y^T B_i y$, and the Hessian of p is

$$\nabla^2 p(x, y) = 2 \begin{bmatrix} 0 & C(y)^T \\ C(y) & B(x) \end{bmatrix},$$

where $C(y)$ is an $k \times n$ matrix whose i -th column equals $B_i y$. One can then immediately see that if $(\bar{x}, \bar{y}) \in CR_p$, then we must have $B(\bar{x}) \succeq 0$. Conversely, if $B(\bar{x}) \succeq 0$, then $(\bar{x}, 0_k) \in CR_p$. Hence, in view of (2.15), we have shown that the spectrahedron S is the projection of CR_p onto the x variables.

We now show that $LM_p = \{x \in \mathbb{R}^n \mid B(x) \succ 0\} \times \{0_k\}$. This would prove the second claim of the theorem. First let \bar{x} be such that $B(\bar{x}) \succ 0$. Note that $p(\bar{x}, 0_k) = 0$ and that for

any two vectors $\chi \in \mathbb{R}^n$ and $\psi \in \mathbb{R}^k$,

$$p(\bar{x} + \chi, \psi) = \psi^T \left(B(\bar{x}) + \sum_{i=1}^n B_i \chi_i \right) \psi.$$

Since $B(\bar{x}) \succ 0$, then for any χ of sufficiently small norm, $B(\bar{x}) + \sum_{i=1}^n B_i \chi_i$ is still positive definite, and hence for any ψ , $p(\bar{x} + \chi, \psi) \geq 0 = p(\bar{x}, 0_k)$. Thus $(\bar{x}, 0_k)$ is a local minimum of p .

Now let (\bar{x}, \bar{y}) be a local minimum of p . From the SONC, we must have $B(\bar{x}) \succeq 0$ and $C(\bar{y}) = 0$, which implies that $B_i \bar{y} = 0_k, \forall i \in \{1, \dots, n\}$. Since

$$\frac{\partial p}{\partial y}(\bar{x}, \bar{y}) = 2B(\bar{x})\bar{y} = 2 \left(B_0 + \sum_{i=1}^n \bar{x}_i B_i \right) \bar{y} = 2B_0 \bar{y} + 2 \sum_{i=1}^n \bar{x}_i (B_i \bar{y}),$$

it further follows from the FONC that $B_0 \bar{y} = 0$. As $\mathcal{N}(B_0) \cap \dots \cap \mathcal{N}(B_n) = \{0_k\}$ by construction, it follows that we must have $\bar{y} = 0_k$. Next, observe that $\mathcal{N}(\nabla^2 p(\bar{x}, 0_k)) = \mathbb{R}^n \times \mathcal{N}(B(\bar{x}))$. Let $d \in \mathcal{N}(B(\bar{x}))$, and note that for any $i \in \{1, \dots, n\}$, $(e_i, d) \in \mathcal{N}(\nabla^2 p(\bar{x}, 0_k))$ and $\frac{\partial p_3}{\partial y}(e_i, d) = B_i d$. Then from the TOC, we must have $B_i d = 0_k, \forall i \in \{1, \dots, n\}$. Furthermore, since $d \in \mathcal{N}(B(\bar{x}))$, it follows that $B_0 d = 0_k$ as well. Again, as $\mathcal{N}(B_0) \cap \dots \cap \mathcal{N}(B_n) = \{0_k\}$ by construction, it follows that we must have $d = 0_k$ and thus $B(\bar{x}) \succ 0$. \square

2.5 Complexity Justifications for an Exact SDP Oracle

In the next section, we show that second-order points and local minima of cubic polynomials can be found by solving polynomially-many semidefinite programs with a polynomial number of variables and constraints. One caveat however is that the inputs and outputs of these semidefinite programs can sometimes be algebraic but not necessarily rational numbers. As a result, we cannot claim that second-order points and local minima of cubic polynomials can be found in polynomial time in the Turing model of computation. In this subsection, we

give evidence as to why establishing the complexity of these problems in the Turing model is at the moment likely out of reach.

Definition 2.5.1. *The SDP Feasibility Problem (SDPF) is the following decision question: Given $m \times m$ symmetric matrices A_0, \dots, A_n with rational entries, decide whether there exists a vector $x \in \mathbb{R}^n$ such that $A_0 + \sum_{i=1}^n x_i A_i \succeq 0$.*

Definition 2.5.2. *The SDP Strict Feasibility Problem (SDPSF) is the following decision question: Given $m \times m$ symmetric matrices A_0, \dots, A_n with rational entries, decide whether there exists a vector $x \in \mathbb{R}^n$ such that $A_0 + \sum_{i=1}^n x_i A_i \succ 0$.*

Even though semidefinite programs can be solved to arbitrary accuracy in polynomial time [108], the complexities of the decision problems above remain as two of the outstanding open problems in semidefinite programming. At the moment, it is not known if these two decision problems even belong to the class NP [95, 90, 34]. We show next that the complexities of these problems are a lower bound on the complexities of testing existence of second-order points and local minima of cubic polynomials. (In Section 2.6, we accomplish the more involved task of giving the reduction in the opposite direction.)

Theorem 2.5.3. *If the problem of deciding whether a cubic polynomial has any second-order points is in P (resp. NP), then SDPF is in P (resp. NP).*

Proof. Given matrices $A_0, \dots, A_n \in \mathbb{S}^{m \times m}$, let $A(x) := A_0 + \sum_{i=1}^n x_i A_i$. By noting that the cubic polynomial $p(x, y) = y^T A(x) y$ has as its Hessian

$$\nabla^2 p(x, y) = 2 \begin{bmatrix} 0 & B(y)^T \\ B(y) & A(x) \end{bmatrix},$$

where $B(y)$ is an $m \times n$ matrix whose i -th column equals $A_i y$, we can see that if $A(\bar{x}) \succeq 0$ for some $\bar{x} \in \mathbb{R}^n$, then $\nabla^2 p(\bar{x}, 0_k) \succeq 0$. Since p is quadratic in the variables y , $\nabla p(\bar{x}, 0_k) = 0_{m+n}$, and hence $(\bar{x}, 0_k)$ is a second-order point of p . Conversely, if $A(x) \not\succeq 0$ for any $x \in \mathbb{R}^n$, then

clearly $\nabla^2 p(x, y) \not\geq 0$ for any $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$, and thus p cannot have any second-order points.

The above reduction shows that any polynomial-time algorithm (or polynomial-time verifiable certificate) for existence of second-order points of cubic polynomials translates into one for SDPF. \square

Theorem 2.5.4. *If the problem of deciding whether a cubic polynomial has any local minima is in P (resp. NP), then SDPSF is in P (resp. NP).*

Proof. Given matrices $A_0, \dots, A_n \in \mathbb{S}^{m \times m}$, let $A(x) := A_0 + \sum_{i=1}^n x_i A_i$ and consider the set $S := \{x \in \mathbb{R}^n \mid A(x) \succeq 0\}$. It is not difficult to see that there exists $\bar{x} \in \mathbb{R}^n$ such that $A(\bar{x}) \succ 0$ if and only if S has a nonempty interior and $\mathcal{N}_A := \mathcal{N}(A_0) \cap \mathcal{N}(A_1) \cap \mathcal{N}(A_2) \dots \cap \mathcal{N}(A_n) = \{0_m\}$.⁴ The latter condition can be checked in polynomial time by solving linear systems. The former can be reduced—due to the second claim of Theorem 2.4.12—to deciding if the cubic polynomial constructed in (2.16) has a local minimum. Note that the polynomial in (2.16) has coefficients polynomially sized in the entries of the matrices A_i , since the matrix V in the proof of Theorem 2.4.12 can be taken to be the identity matrix when $\mathcal{N}_A = \{0_m\}$. \square

In addition to the difficulties alluded to in the above two theorems, the following three examples point to concrete representation issues that one encounters in the Turing model when dealing with local minima of cubic polynomials. The same complications are known to arise for SDP feasibility problems [34].

Example 2.5.1. A cubic polynomial with only irrational local minima. Consider the univariate cubic polynomial $p(x) = x^3 - 6x$. One can easily verify that its unique local minimum is at $x = \sqrt{2}$, which is irrational even though the coefficients of p are rational.

⁴The “only if” direction is straightforward and the “if” direction follows from [93, Corollary 5].

Example 2.5.2. A cubic polynomial with an irrational convexity region. Consider the quintary cubic polynomial $p(x, y) = y^T A(x)y$, where

$$A(x) = \begin{bmatrix} 2 & x & 0 & 0 \\ x & 1 & 0 & 0 \\ 0 & 0 & 2x & 2 \\ 0 & 0 & 2 & x \end{bmatrix}.$$

One can easily verify that $x = \sqrt{2}$ is the only scalar satisfying $A(x) \succeq 0$. Since the matrix $2A(x)$ is a principal submatrix of $\nabla^2 p(x, y)$, any point in the convexity region of p must satisfy $x = \sqrt{2}$ (even though the coefficients of p are rational).

Example 2.5.3. A family of cubic polynomials whose local minima have exponential bitsize. Consider the family of cubic polynomials $p_n(x, y) = y^T A_n(x)y$ in $3n$ variables, where

$$A_n(x) = \begin{bmatrix} x_1 & 2 & 0 & 0 & \cdots & 0 & 0 \\ 2 & 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & x_2 & x_1 & \cdots & 0 & 0 \\ 0 & 0 & x_1 & 1 & \cdots & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \ddots & \cdots & \cdots \\ 0 & 0 & 0 & 0 & \cdots & x_n & x_{n-1} \\ 0 & 0 & 0 & 0 & \cdots & x_{n-1} & 1 \end{bmatrix}.$$

We show that even though these polynomials have some rational local minima, it takes exponential time to write them down. From the proof of Theorem 2.4.12, one can infer that the set of local minima of p_n is the set $\{x \in \mathbb{R}^n \mid A_n(x) \succ 0\} \times \{0_{2n}\}$. However, observe that to have $A_n(x) \succ 0$ (or even $A_n(x) \succeq 0$), we must have

$$x_1 \geq 4, x_2 \geq 16, \dots, x_n \geq 2^{2^n}.$$

Hence, any local minimum of p_n has bit length at least $O(2^n)$ even though the bit length of the coefficients of p_n is $O(n)$.

2.6 Finding Local Minima of Cubic Polynomials

In this section, we derive an SDP-based approach for finding second-order points and local minima of cubic polynomials. This, along with the results established in Section 2.2, will complete the entries of Table 2.2 from Section 2.1. We begin with some preliminaries that are needed to present the theorems of this section.

2.6.1 Preliminaries from Semidefinite and Sum of Squares Optimization

The Oracle E-SDP

Recall that a *spectrahedron* is a set of the type

$$\left\{ x \in \mathbb{R}^n \mid A_0 + \sum_{i=1}^n x_i A_i \succeq 0 \right\},$$

where A_0, \dots, A_n are symmetric matrices of some size $m \times m$. A *semidefinite representable set* (also known as a *spectrahedral shadow*) is a set of the type

$$\left\{ x \in \mathbb{R}^n \mid \exists y \in \mathbb{R}^k \text{ such that } A_0 + \sum_{i=1}^n x_i A_i + \sum_{i=1}^k y_i B_i \succeq 0 \right\}, \quad (2.17)$$

for some integer $k \geq 0$ and symmetric $m \times m$ matrices $A_0, A_1, \dots, A_n, B_1, \dots, B_k$. These are exactly sets which semidefinite programming can optimize over.

We show in Theorem 2.6.3 and Corollary 2.6.5 that the set of second-order points of any cubic polynomial is a spectrahedron and describe how a description of this spectrahedron

can be obtained from the coefficients of p only.⁵ Since relative interiors of semidefinite representable sets (and in particular spectrahedra) are semidefinite representable [81, Theorem 3.8], it follows from our Theorem 2.4.10 that the set of local minima of any cubic polynomial is semidefinite representable.

Due to the complexity results and representation issues presented in Section 2.5, we assume in this section that we can do arithmetic over real numbers and have access to an oracle which solves SDPs exactly. This oracle—which we call *E-SDP*—takes as input an SDP with real data and outputs the optimal value as a real number if it is finite, or reports that the SDP is infeasible, or that it is unbounded.⁶ The following lemma shows that E-SDP can find a point in the relative interior of a semidefinite representable set. This will be relevant for us later in this section when we search for local minima of cubic polynomials.

Lemma 2.6.1. *Let S be a nonempty semidefinite representable set in \mathbb{R}^n . Then a point in $ri(S)$ can be recovered in $2n$ calls to E-SDP.*

Proof. Consider the following procedure. Let $S_1 = S$, and for $i \in \{1, \dots, n\}$ let

$$S_{i+1} = S_i \cap \{x \in \mathbb{R}^n \mid x_i = x_i^*\},$$

where the scalar x_i^* is chosen to be any “intermediate” value of x_i on S_i . More precisely, let \bar{x}_i (resp. \underline{x}_i) be the supremum (resp. infimum) of x_i over S_i (these two values may or may not be finite). If $\bar{x}_i = \underline{x}_i$, then set $x_i^* = \bar{x}_i$. Otherwise, set x_i^* to be any scalar satisfying $\underline{x}_i < x_i^* < \bar{x}_i$. Note that for each i , x_i^* can be computed using 2 calls to E-SDP. Hence, after $2n$ calls to E-SDP, we arrive at a set S_{n+1} which is a singleton by construction.

We next show, by induction, that the point in S_{n+1} belongs to $ri(S)$. First note that as S is nonempty, $ri(S)$ is nonempty [97, Theorem 6.2], which implies that $S_1 \cap ri(S) = ri(S)$

⁵Recall that the results of Section 2.4 by contrast established spectrahedrality of the set of second-order points under the assumption of existence of a local minimum (see Remark 2.4.2). Furthermore, the spectrahedral representation that we gave there (see Theorem 2.4.5) required knowledge of a local minimum.

⁶Though this will not be needed for our purposes, it is straightforward to show that for an SDP with n scalar variables, the oracle E-SDP can be called twice to test attainment of the optimal value, and a total of $n + 1$ times to recover an optimal solution.

is nonempty. Now suppose that $S_i \cap ri(S)$ is nonempty for $i \in \{1, \dots, k\}$. We show that $S_{k+1} \cap ri(S)$ is nonempty.

First suppose that k is such that $\bar{x}_k = x_k^* = \underline{x}_k$. In this case, because $\forall x \in S_k, x_k = x_k^*$,

$$S_{k+1} \cap ri(S) = S_k \cap \{x \in \mathbb{R}^n \mid x_k = x_k^*\} \cap ri(S) = S_k \cap ri(S) \neq \emptyset.$$

Now suppose that $\underline{x}_k < x_k^* < \bar{x}_k$. By the definition of \bar{x}_k , there exists a sequence of points $\{y_j\} \subseteq S_k$ such that $(y_j)_k \rightarrow \bar{x}_k$. We recall that for any $z \in ri(S), y \in \bar{S}$, and $\lambda \in (0, 1]$, $\lambda z + (1 - \lambda)y \in ri(S)$ [97, Theorem 6.1]. Now let $z \in S_k \cap ri(S)$. Since S_k is convex, for any $y \in S_k \cap \bar{S}$ and $\lambda \in (0, 1]$, $\lambda z + (1 - \lambda)y \in S_k \cap ri(S)$. In particular, since $S_k \cap \bar{S} = S_k$, the sequence $\{z_j\} := \{\frac{1}{j}z + \frac{j-1}{j}y_j\}$ satisfies $\{z_j\} \subseteq S_k \cap ri(S)$ and $(z_j)_k \rightarrow \bar{x}_k$. Similarly, there exists a sequence of points $\{w_j\} \subseteq S_k \cap ri(S)$ such that $(w_j)_k \rightarrow \underline{x}_k$. As $S_k \cap ri(S)$ is convex, there must then be a point $x \in S_k \cap ri(S)$ satisfying $x_k = x_k^*$, and so

$$S_{k+1} \cap ri(S) = S_k \cap \{x \in \mathbb{R}^n \mid x_k = x_k^*\} \cap ri(S)$$

is not empty. □

Overview of Sum of Squares Polynomials

In order to describe our SDP-based approach for finding local minima of cubic polynomials, we also need to briefly review the connection between sum of squares polynomials and matrices to semidefinite programming. Approaches to finding local minima based on sum of squares have been studied before, such as in [83]. The SDP approach in this chapter, however, is based partially on finding critical points and does not formally study the case of cubic polynomials.

Recall that a (multivariate) polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is *nonnegative* if $p(x) \geq 0, \forall x \in \mathbb{R}^n$, and that a polynomial p is said to be a *sum of squares* (sos) if $p = \sum_{i=1}^r q_i^2$ for some polynomials q_1, \dots, q_r . The notion of sum of squares also extends to polynomial matrices (i.e.,

matrices whose entries are multivariate polynomials). We say that symmetric polynomial matrix $M(x) : \mathbb{R}^n \rightarrow \mathbb{R}^m \times \mathbb{R}^m$ is an *sos-matrix* if it has a factorization as $M(x) = R(x)^T R(x)$ for some $r \times m$ polynomial matrix $R(x)$ [47]. Observe that if M is an sos-matrix, then $M(x) \succeq 0$ for any $x \in \mathbb{R}^n$. One can check that $M(x)$ is an sos-matrix if and only if the scalar-valued polynomial $y^T M(x)y$ in variables $(x_1, \dots, x_n, y_1, \dots, y_m)$ is sos. Indeed, the “only if” direction is clear, the “if” direction is because when $y^T M(x)y = \sum_{i=1}^r q_i^2(x, y)$ for some polynomials q_1, \dots, q_r , each q_i must be linear in y and thus writable as $q_i(x) = \sum_{j=1}^m y_j q_{ij}(x)$ for some polynomials q_{ij} . Then if $R(x)$ is the $r \times m$ matrix where $R_{ij}(x) = q_{ij}(x)$, we will have $M(x) = R^T(x)R(x)$.

2.6.2 A Sum of Squares Approach for Finding Second-Order Points

We have shown in Theorem 2.4.7 that if a cubic polynomial p has a second-order point, the solutions of the optimization problem in (2.12) exactly form the set SO_p of its second-order points. The same theorem further showed that if p has a local minimum, then the solutions of (2.12) also coincide with $\overline{LM_p}$, i.e. the closure of the set of its local minima. Our goal in this section is to develop a semidefinite representation of SO_p which can be obtained directly from the coefficients of p (Corollary 2.6.5). To arrive to this representation, we first present an sos relaxation of problem (2.12), which we prove to be tight when SO_p is nonempty (Theorem 2.6.2). We then provide a more efficient representation of the SDP underlying this sos relaxation in Section 2.6.3. This will lead to an algorithm (Algorithm 2) for finding local minima of cubic polynomials which is presented in Section 2.6.3.

Theorem 2.6.2. *If a cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ has a second-order point, the optimal value of the following semidefinite program⁷ is attained and is equal to the value of p at all second-order points:*

$$\begin{aligned}
& \sup_{\gamma \in \mathbb{R}, \sigma(x), S(x)} \quad \gamma \\
\text{subject to} \quad & p(x) - \gamma = \sigma(x) + \text{Tr}(S(x)\nabla^2 p(x)), \\
& \sigma(x) \text{ is a degree-2 sos polynomial,} \\
& S(x) \text{ is an } n \times n \text{ sos-matrix with degree-2 entries.}
\end{aligned} \tag{2.18}$$

Proof. Let \bar{x} be a second-order point of p and γ^* be the optimal value of (2.18). Consider any feasible solution (γ, σ, S) to (2.18) (nonemptiness of the feasible set is established in the next paragraph). Since $\nabla^2 p(\bar{x}) \succeq 0$ and $S(\bar{x}) \succeq 0$, we have $\text{Tr}(\nabla^2 p(\bar{x})S(\bar{x})) \geq 0$. Since $\sigma(\bar{x}) \geq 0$ as well, it follows that $p(\bar{x}) \geq \gamma$. Hence, $p(\bar{x}) \geq \gamma^*$.

To show that $p(\bar{x}) \leq \gamma^*$ and that the value $\gamma^* = p(\bar{x})$ is attained, we establish that

$$(\gamma, \sigma, S) = \left(p(\bar{x}), \frac{1}{3}(x - \bar{x})^T \nabla^2 p(\bar{x})(x - \bar{x}), \frac{1}{6}(x - \bar{x})(x - \bar{x})^T \right)$$

is feasible to (2.18). Note that $\frac{1}{3}(x - \bar{x})^T \nabla^2 p(\bar{x})(x - \bar{x})$ is an sos polynomial (as $\nabla^2 p(\bar{x})$ can be factored into $V^T V$), and that $\frac{1}{6}(x - \bar{x})(x - \bar{x})^T$ is an sos-matrix by construction. To show that the first constraint in (2.18) is satisfied, consider the Taylor expansion of p around \bar{x} in the direction $x - \bar{x}$ (see (2.2), noting that $\nabla p(\bar{x}) = 0$):

$$p(\bar{x} + (x - \bar{x})) = p(\bar{x}) + \frac{1}{2}(x - \bar{x})^T \nabla^2 p(\bar{x})(x - \bar{x}) + p_3(x - \bar{x}). \tag{2.19}$$

⁷To clarify, x is not a decision variable in this problem. The decision variables are γ , the coefficients of σ , and the coefficients of the entries of S . The identity in the first constraint must hold for all x , and this can be enforced by matching the coefficient of each monomial on the left with the corresponding coefficient on the right.

Observe that if p is written in the form (2.4), then we have

$$\begin{aligned}
p_3(x - \bar{x}) &= \frac{1}{6}(x - \bar{x})^T \left(\sum_{i=1}^n (x_i - \bar{x}_i) H_i \right) (x - \bar{x}) \\
&= \frac{1}{6}(x - \bar{x})^T \left(\sum_{i=1}^n (x_i - \bar{x}_i) H_i + Q - Q \right) (x - \bar{x}) \\
&= \frac{1}{6}(x - \bar{x})^T \left(\sum_{i=1}^n x_i H_i + Q - \sum_{i=1}^n \bar{x}_i H_i - Q \right) (x - \bar{x}) \\
&= \frac{1}{6}(x - \bar{x})^T \nabla^2 p(x) (x - \bar{x}) - \frac{1}{6}(x - \bar{x})^T \nabla^2 p(\bar{x}) (x - \bar{x}).
\end{aligned}$$

Note further that due to the cyclic property of the trace, we have

$$\frac{1}{6}(x - \bar{x})^T \nabla^2 p(x) (x - \bar{x}) = \text{Tr} \left(\left(\frac{1}{6}(x - \bar{x})(x - \bar{x})^T \right) \nabla^2 p(x) \right).$$

Hence, (2.19) reduces to the following identity

$$p(x) - p(\bar{x}) = \frac{1}{3}(x - \bar{x})^T \nabla^2 p(\bar{x}) (x - \bar{x}) + \text{Tr} \left(\left(\frac{1}{6}(x - \bar{x})(x - \bar{x})^T \right) \nabla^2 p(x) \right), \quad (2.20)$$

and thus the claim is established. \square

Since (2.18) is a tight sos relaxation of (2.12) when SO_p is nonempty, it is interesting to see how an optimal solution to (2.12) can be recovered from an optimal solution to (2.18). This is shown in the next theorem, keeping in mind that optimal solutions to (2.12) are second-order points of p (see Theorem 2.4.7).

Theorem 2.6.3. *Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a cubic polynomial with a second-order point, and let $(\gamma^*, \sigma^*, S^*)$ be an optimal solution of (2.18) applied to p . Then, the set*

$$\Gamma := \{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0, \sigma^*(x) = 0, \text{Tr}(S^*(x) \nabla^2 p(x)) = 0\} \quad (2.21)$$

is a spectrahedron, and $\Gamma = SO_p$.

Proof. We first show that $\Gamma = SO_p$. Let \bar{x} be a second-order point of p . From Theorem 2.6.2 and the first constraint of (2.18) we have

$$0 = p(\bar{x}) - p(\bar{x}) = p(\bar{x}) - \gamma^* = \sigma^*(\bar{x}) + \text{Tr}(S^*(\bar{x})\nabla^2 p(\bar{x})).$$

As $\sigma^*(\bar{x})$ and $\text{Tr}(S^*(\bar{x})\nabla^2 p(\bar{x}))$ are both nonnegative, the above equation implies they must both be zero, and hence $SO_p \subseteq \Gamma$. To see why $\Gamma \subseteq SO_p$, let \bar{y} be a point in Γ and \hat{x} be an arbitrary second-order point (which by the assumption of the theorem exists). Observe from Theorem 2.6.2 and the first constraint of (2.18) that

$$p(\bar{y}) - p(\hat{x}) = p(\bar{y}) - \gamma^* = \sigma^*(\bar{y}) + \text{Tr}(S^*(\bar{y})\nabla^2 p(\bar{y})) = 0.$$

Additionally, because $\nabla^2 p(\bar{y}) \succeq 0$, it follows from Corollary 2.4.8 that \bar{y} is optimal to (2.12), and thus is a second-order point by Theorem 2.4.7.

Now we show that Γ is a spectrahedron by “linearizing” the quadratic and cubic equations that appear in (2.21). Since σ^* is a quadratic sos polynomial, it can be written equivalently as $\sigma^*(x) = \sum_{i=1}^m q_i^2(x)$ for some affine polynomials q_1, \dots, q_m . Similarly, since S^* is an sos-matrix with quadratic entries, it can be written as $S^*(x) = R(x)^T R(x)$ for some $k \times n$ matrix R with affine entries. First note that as $\nabla^2 p(x)$ is affine in x and σ^* is a sum of squares of affine polynomials, the set

$$\{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0, \sigma^*(x) = 0\} = \{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0, q_1(x) = 0, \dots, q_m(x) = 0\}$$

is clearly a spectrahedron.

Now let y be any point in $ri(CR_p)$. Such a point exists because CR_p is nonempty by assumption, and relative interiors of nonempty convex sets are nonempty [97, Theorem 6.2]. Now let r_i be the i -th column of the matrix R^T . We claim that Γ is equivalent to the

following set:

$$\{x \in \mathbb{R}^n \mid \nabla^2 p(x) \succeq 0, q_1(x) = 0, \dots, q_m(x) = 0, \nabla^2 p(y)r_1(x) = 0, \dots, \nabla^2 p(y)r_k(x) = 0\}. \quad (2.22)$$

Note that this set is a spectrahedron, and that the final k equality constraints are enforcing that each column of R^T be in the null space of $\nabla^2 p(y)$.

To prove the claim, first let x be in (2.22). Note that $\mathcal{N}(\nabla^2 p(y)) \subseteq \mathcal{N}(\nabla^2 p(x))$, as $y \in \text{ri}(CR_p)$ and so $\nabla^2 p(y) = \lambda \nabla^2 p(x) + (1 - \lambda) \nabla^2 p(z)$ for some $z \in CR_p$ and $\lambda \in (0, 1)$.

Then,

$$\text{Tr}(S^*(x) \nabla^2 p(x)) = \sum_{i=1}^k r_i^T(x) \nabla^2 p(x) r_i(x) = 0.$$

Hence (2.22) \subseteq (2.21).

To show the reverse inclusion, let x be a point in (2.21). It is easy to check that $\text{Tr}(AB) = 0$ for two psd matrices $A = C^T C$ and B if and only if the columns of C^T belong to the null space of B . Hence, we must have $r_i(x) \in \mathcal{N}(\nabla^2 p(x))$. Assume first that $x \in \text{ri}(CR_p)$. Then we must have $r_i(x) \in \mathcal{N}(\nabla^2 p(x)) = \mathcal{N}(\nabla^2 p(y))$ as CR_p is a spectrahedron and any two matrices in the relative interior of a spectrahedron have the same null space [93, Corollary 1]. To see why we must also have $r_i(x) \in \mathcal{N}(\nabla^2 p(y))$ for any $x \in CR_p \setminus \text{ri}(CR_p)$, observe that $\mathcal{N}(\nabla^2 p(y))$ is closed, the vector-valued functions r_i are continuous in x , and the preimage of a closed set under a continuous function is closed. \square

2.6.3 A Simplified Semidefinite Representation of Second-Order Points and an Algorithm for Finding Local Minima

In this subsection, we derive a semidefinite representation of the set SO_p , which will be given in (2.31). In contrast to the semidefinite representation in (2.22), which requires first solving (2.18) and then performing some matrix factorizations, the representation in (2.31) can be immediately obtained from the coefficients of p . To find a second-order point of an n -variate

cubic polynomial via the representation in (2.31), one needs to solve an SDP with $\frac{(n+2)(n+1)}{2}$ scalar variables and two semidefinite constraints of size $(n+1) \times (n+1)$. This is in contrast to finding a second-order point via the representation in (2.22), which requires solving two SDPs: (2.18) which has $\left(\frac{n(n+1)}{2} + 1\right) \left(\frac{(n+2)(n+1)}{2}\right) + 1$ scalar variables and two semidefinite constraints of sizes $(n+1) \times (n+1)$ and $n(n+1) \times n(n+1)$ (coming from the two sos constraints), and then the SDP associated with (2.22), which has n scalar variables and a semidefinite constraint of size $n \times n$. Another purpose of this subsection is to present our final result, which is an algorithm for testing for existence of a local minimum (Algorithm 2 in Section 2.6.3).

A Simplified Sos Relaxation

Recall from the proof of Theorem 2.6.2 that if p has a second-order point \bar{x} , then there is an optimal solution to (2.18) of the form

$$(\gamma, \sigma, S) = \left(p(\bar{x}), \frac{1}{3}(x - \bar{x})^T \nabla^2 p(\bar{x})(x - \bar{x}), \frac{1}{6}(x - \bar{x})(x - \bar{x})^T \right). \quad (2.23)$$

In particular, for this solution, the coefficients of σ and S can both be written entirely in terms of the entries of \bar{x} and the coefficients of p . In what follows, we attempt to optimize over solutions to (2.18) which are of the form in (2.23). However, imposing this particular structure on the solution requires nonlinear equality constraints (in fact, it turns out quadratic constraints suffice). Instead, we will impose an SDP relaxation of these nonlinear constraints and show that the relaxation is exact. We follow a standard technique in deriving SDP relaxations for quadratic programs, where the outer product xx^T of some variable x is replaced by a new matrix variable X satisfying $X - xx^T \succeq 0$. The latter matrix inequality that can be imposed as a semidefinite constraint via the Schur complement [22]. The variable \bar{x} will be represented by a variable $y \in \mathbb{R}^n$, and the symmetric matrix variable $Y \in \mathbb{S}^{n \times n}$ will represent yy^T . In addition, we will need another scalar variable z .

Assume p is given in the form (2.4), and let us expand σ in (2.23) (disregarding the factor $\frac{1}{3}$) as follows:

$$\begin{aligned} (x - \bar{x})^T \nabla^2 p(\bar{x})(x - \bar{x}) &= x^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) x - 2\bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) x + \bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) \bar{x} \\ &= x^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) x - 2 \sum_{i=1}^n \text{Tr}(H_i \bar{x} \bar{x}^T) x_i - 2\bar{x}^T Q x + \bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) \bar{x}, \end{aligned}$$

where in the last equality we used Lemma 2.4.1. If we replace any occurrence of \bar{x} with y , any occurrence of $\bar{x} \bar{x}^T$ with Y and any occurrence of $\bar{x}^T (\sum_{i=1}^n \bar{x}_i H_i + Q) \bar{x}$ with z , we can rewrite the above expression as

$$\sigma_{Y,y,z}(x) := \sum_{j=1}^n \sum_{k=1}^n \left(\sum_{i=1}^n (H_i)_{jk} y_i + Q_{jk} \right) x_j x_k - 2 \sum_{i=1}^n (\text{Tr}(H_i Y) + e_i^T Q y) x_i + z. \quad (2.24)$$

Similarly, the matrix S in (2.23) can be written as $xx^T - xy^T - yx^T + Y$ (disregarding the factor $\frac{1}{6}$). Note that if $Y - yy^T \succeq 0$, then the matrix $xx^T - xy^T - yx^T + Y$ is an sos-matrix (as a polynomial matrix in x). By making these replacements, we arrive at an SDP which attempts to look for a solution to the sos program in (2.18) which is of the structure in (2.23). This is the following SDP⁸:

$$\begin{aligned} &\sup_{\gamma \in \mathbb{R}, Y \in \mathbb{S}^{n \times n}, y \in \mathbb{R}^n, z \in \mathbb{R}} \gamma \\ \text{subject to} & \quad p(x) - \gamma = \frac{1}{3} \sigma_{Y,y,z}(x) + \frac{1}{6} \text{Tr}(\nabla^2 p(x)(xx^T - xy^T - yx^T + Y)), \\ & \quad \sigma_{Y,y,z} \text{ is sos,} \\ & \quad \begin{bmatrix} Y & y \\ y^T & 1 \end{bmatrix} \succeq 0. \end{aligned} \quad (2.25)$$

⁸Note that x is not a decision variable in this SDP as the first constraint needs to hold for all x .

Through straightforward algebra and matching coefficients, the first constraint (keeping in mind that p is as in (2.4)) can be more explicitly written as:

$$\begin{aligned} b_i &= -e_i^T Q y - \frac{1}{2} \text{Tr}(H_i Y), i = 1, \dots, n, \\ -\gamma &= \frac{1}{6} \text{Tr}(Q Y) + \frac{z}{3}. \end{aligned}$$

These constraints reflect that the coefficients of the linear terms and the scalar coefficient match on both sides; the cubic and quadratic coefficients are automatically the same. We can rewrite (2.24) as

$$\sigma_{Y,y,z}(x) = \begin{pmatrix} x \\ 1 \end{pmatrix}^T T(Y, y, z) \begin{pmatrix} x \\ 1 \end{pmatrix},$$

where

$$T(Y, y, z) := \begin{bmatrix} \sum_{i=1}^n y_i H_i + Q & \sum_{i=1}^n \text{Tr}(H_i Y) e_i + Q y \\ (\sum_{i=1}^n \text{Tr}(H_i Y) e_i + Q y)^T & z \end{bmatrix}.$$

The constraint in (2.25) that σ be sos is the same as the matrix T being psd. Putting everything together, the problem in (2.25) can be rewritten as the following SDP:⁹

$$\begin{aligned} & \inf_{Y \in \mathbb{S}^{n \times n}, y \in \mathbb{R}^n, z \in \mathbb{R}} \frac{1}{6} \text{Tr}(Q Y) + \frac{z}{3} \\ & \text{subject to} \quad \frac{1}{2} \text{Tr}(H_i Y) + e_i^T Q y + b_i = 0, \forall i = 1, \dots, n, \\ & \quad T(Y, y, z) \succeq 0, \\ & \quad \begin{bmatrix} Y & y \\ y^T & 1 \end{bmatrix} \succeq 0. \end{aligned} \tag{2.26}$$

It is interesting to observe that the first constraint is a relaxation of the quadratic constraint which would impose $\nabla p(y) = 0$, and that the constraint $T(Y, y, z) \succeq 0$ in particular

⁹Recall that the data to this SDP is obtained from the representation of p in the form of (2.4).

implies $\nabla^2 p(y) \succeq 0$. One can think of (2.26) as another SDP relaxation of (2.12) which is tight when p has a second-order point.

Combining the SDP in (2.26) with its Dual

In this subsection, we write down an SDP (given in (2.28)) whose optimal value can be related to the existence of second-order points of a cubic polynomial. To arrive at this SDP, we first take the dual of (2.26). It will turn out that the constraints in the dual follow a very similar structure to those in the primal, and that any feasible solution of the primal yields a feasible solution of the dual. We then combine the primal-dual pair of SDPs to arrive at a single SDP, which is the one in (2.28). To this end, let us write down the dual of (2.26):

$$\begin{aligned}
& \sup_{R, S, r, s, \lambda, \sigma, \rho, \gamma} \gamma \\
\text{subject to} \quad & \frac{1}{6} \text{Tr}(QY) + \frac{z}{3} - \gamma = \sum_{i=1}^n \lambda_i \left(\frac{1}{2} \text{Tr}(H_i Y) + e_i^T Q y + b_i \right) \\
& + \text{Tr} \left(\begin{bmatrix} Y & y \\ y^T & 1 \end{bmatrix} \begin{bmatrix} R & r \\ r^T & \rho \end{bmatrix} \right) + \text{Tr} \left(T(Y, y, z) \begin{bmatrix} S & s \\ s^T & \sigma \end{bmatrix} \right), \forall (Y, y, z) \\
& \begin{bmatrix} R & r \\ r^T & \rho \end{bmatrix} \succeq 0, \\
& \begin{bmatrix} S & s \\ s^T & \sigma \end{bmatrix} \succeq 0,
\end{aligned}$$

where $R, S \in \mathbb{S}^{n \times n}$, $r, s, \lambda \in \mathbb{R}^n$, and $\sigma, \rho, \gamma \in \mathbb{R}$. The right-hand side of the first constraint simplifies to

$$b^T \lambda + \rho + \text{Tr}(QS) + \text{Tr} \left(\left(\sum_{i=1}^n \left(\frac{1}{2} \lambda_i + 2s_i \right) H_i + R \right) Y \right) + \left(Q(\lambda + 2s) + \sum_{i=1}^n \text{Tr}(H_i S) e_i + 2r \right)^T y + \sigma z.$$

After matching coefficients, the dual problem can be rewritten as

$$\begin{aligned}
& \sup_{R,S,r,s,\lambda,\rho} && -b^T\lambda - \rho - \text{Tr}(QS) \\
& \text{subject to} && \sum_{i=1}^n \left(\frac{1}{2}\lambda_i + 2s_i\right)H_i + R = \frac{1}{6}Q, \\
& && Q(\lambda + 2s) + \sum_{i=1}^n \text{Tr}(H_i S)e_i + 2r = 0, \\
& && \begin{bmatrix} R & r \\ r^T & \rho \end{bmatrix} \succeq 0, \\
& && \begin{bmatrix} S & s \\ s^T & \frac{1}{3} \end{bmatrix} \succeq 0,
\end{aligned}$$

Substituting R and r using the first two constraints into the first psd constraint and then multiplying by 6, we arrive at the problem

$$\begin{aligned}
& \sup_{S,s,\lambda,\rho} && -b^T\lambda - \rho - \text{Tr}(QS) \\
& \text{subject to} && \begin{bmatrix} \sum_{i=1}^n (-3\lambda_i - 12s_i)H_i + Q & Q(-3\lambda - 6s) - 3\sum_{i=1}^n \text{Tr}(H_i S)e_i \\ (Q(-3\lambda - 6s) - 3\sum_{i=1}^n \text{Tr}(H_i S)e_i)^T & 6\rho \end{bmatrix} \succeq 0, \\
& && \begin{bmatrix} S & s \\ s^T & \frac{1}{3} \end{bmatrix} \succeq 0.
\end{aligned}$$

Replacing S with $\frac{1}{3}S$, s with $-\frac{1}{3}s$, and ρ with $\frac{1}{6}\rho$, we can reparameterize this problem and arrive at our final form for the dual of (2.26):

$$\begin{aligned}
& \sup_{S,s,\lambda,\rho} && -b^T\lambda - \frac{1}{6}\rho - \frac{1}{3}\text{Tr}(QS) \\
\text{subject to} & && \begin{bmatrix} \sum_{i=1}^n (4s_i - 3\lambda_i)H_i + Q & Q(2s - 3\lambda) - \sum_{i=1}^n \text{Tr}(H_i S)e_i \\ (Q(2s - 3\lambda) - \sum_{i=1}^n \text{Tr}(H_i S)e_i)^T & \rho \end{bmatrix} \succeq 0, \\
& && \begin{bmatrix} S & s \\ s^T & 1 \end{bmatrix} \succeq 0.
\end{aligned} \tag{2.27}$$

One can easily verify that if (Y, y, z) is feasible to (2.26), then (Y, y, y, z) is feasible to (2.27). Replacing (S, s, λ, γ) with (Y, y, y, z) in (2.27) gives an SDP whose constraints are the two psd constraints in (2.26) and whose objective function is $-b^T y - \frac{1}{6}z - \frac{1}{3}\text{Tr}(QY)$. We now create a new SDP, which has the same decision variables and constraints as (2.26), but whose objective function is the difference between the objective function of (2.26) and $-b^T y - \frac{1}{6}z - \frac{1}{3}\text{Tr}(QY)$. The optimal value of this new SDP is an upper bound on the duality gap of the primal-dual SDP pair (2.26) and (2.27). If our cubic polynomial p is written in the form (2.4) and

$$T(Y, y, z) = \begin{bmatrix} \sum_{i=1}^n y_i H_i + Q & \sum_{i=1}^n \text{Tr}(H_i Y)e_i + Qy \\ (\sum_{i=1}^n \text{Tr}(H_i Y)e_i + Qy)^T & z \end{bmatrix}$$

as before, the new SDP we just described can be written as

$$\begin{aligned}
& \inf_{Y \in \mathbb{S}^{n \times n}, y \in \mathbb{R}^n, z \in \mathbb{R}} \frac{1}{2} \text{Tr}(QY) + b^T y + \frac{z}{2} \\
& \text{subject to} \quad \frac{1}{2} \text{Tr}(H_i Y) + e_i^T Q y + b_i = 0, \forall i = 1, \dots, n, \\
& \quad T(Y, y, z) \succeq 0, \\
& \quad \begin{bmatrix} Y & y \\ y^T & 1 \end{bmatrix} \succeq 0.
\end{aligned} \tag{2.28}$$

The following theorem relates the optimal value of this SDP to the existence of second-order points of p .

Theorem 2.6.4. *For a cubic polynomial p given in the form (2.4), consider the SDP in (2.28). For any feasible solution (Y, y, z) to (2.28), the objective value of (2.28) is nonnegative. Furthermore, the optimal value of (2.28) is zero and is attained if and only if p has a second-order point.*

Proof. Suppose (Y, y, z) is a feasible solution to (2.28). Note that (Y, y, z) is feasible to (2.26) and (Y, y, y, z) is feasible to (2.27), and so

$$\frac{1}{2} \text{Tr}(QY) + b^T y + \frac{z}{2} = \frac{1}{6} \text{Tr}(QY) + \frac{z}{3} - \left(-b^T y - \frac{1}{6} z - \frac{1}{3} \text{Tr}(QY) \right) \geq 0$$

by weak duality applied to (2.26) and (2.27). Hence, the objective of (28) is nonnegative at any feasible solution.

Now suppose that p has a second-order point \bar{x} . We claim that the triplet

$$\left(\bar{x} \bar{x}^T, \bar{x}, \bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) \bar{x} \right)$$

is feasible to (2.28) and achieves an objective value of zero. Indeed, the first constraint of (2.28) is satisfied because its left-hand side reduces to $\nabla p(\bar{x})$, which is zero. The third constraint is satisfied since the matrix $(\bar{x}, 1)(\bar{x}, 1)^T$ is clearly psd. The second constraint is

satisfied since $T(\bar{x}\bar{x}^T, \bar{x}, \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)\bar{x})$ can be written as

$$\begin{bmatrix} \sum_{i=1}^n \bar{x}_i H_i + Q & (\sum_{i=1}^n \bar{x}_i H_i + Q)\bar{x} \\ \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q) & \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)\bar{x} \end{bmatrix} = \begin{bmatrix} (\sum_{i=1}^n \bar{x}_i H_i + Q)^{\frac{1}{2}} \\ \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} (\sum_{i=1}^n \bar{x}_i H_i + Q)^{\frac{1}{2}} \\ \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)^{\frac{1}{2}} \end{bmatrix}^T.$$

The objective value at $(\bar{x}\bar{x}^T, \bar{x}, \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)\bar{x})$ is

$$\begin{aligned} & \frac{1}{2}\text{Tr}(Q\bar{x}\bar{x}^T) + b^T \bar{x} + \frac{1}{2}\bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) \bar{x} \\ &= \frac{1}{2}\bar{x}^T Q \bar{x} - \left(\frac{1}{2} \sum_{i=1}^n \bar{x}_i H_i \bar{x} + Q \bar{x} \right)^T \bar{x} + \frac{1}{2}\bar{x}^T \left(\sum_{i=1}^n \bar{x}_i H_i + Q \right) \bar{x} \\ &= 0. \end{aligned}$$

Since we have already shown that the objective function of (2.28) is nonnegative over its feasible set, it follows that when p has a second-order point, the optimal value of (2.28) is zero and is attained.

To prove the converse, suppose the optimal value of (2.28) is zero and is attained. Let (Y^*, y^*, z^*) be an optimal solution to (2.28). We will show that y^* is a second-order point for p . Clearly $\nabla^2 p(y^*)$ is psd, since $T(Y^*, y^*, z^*) \succeq 0$. To show that $\nabla p(y^*) = 0$, let us start by letting $D := Y^* - y^* y^{*T}$, and $d := \sum_{i=1}^n \text{Tr}(H_i D) e_i$. Note that

$$\frac{1}{2}\text{Tr}(H_i y^* y^{*T}) + \frac{1}{2}\text{Tr}(H_i D) + e_i^T Q y^* + b_i = 0 \stackrel{(2.4)}{\Rightarrow} -2(\nabla p(y^*))_i = \text{Tr}(H_i D),$$

or equivalently $d = -2\nabla p(y^*)$. In the remainder of the proof, we show that $d = 0$.

Since $\sum_{i=1}^n y_i^* H_i y^*$ is the vector whose i -th entry is $y^{*T} H_i y^*$, we have that

$$\sum_{i=1}^n \text{Tr}(H_i Y^*) e_i + Q y^* = \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + d. \quad (2.29)$$

Then from the generalized¹⁰ Schur complement condition applied to $T(Y^*, y^*, z^*)$, we have

$$\begin{aligned} z^* &\geq \left(\left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + d \right)^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ \left(\left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + d \right) \\ &= y^{*T} \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + 2d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ d. \end{aligned}$$

It is not difficult to verify that since $T(Y^*, y^*, z^*) \succeq 0$, we have

$$\sum_{i=1}^n \text{Tr}(H_i Y^*) e_i + Q y^* \in \mathcal{C} \left(\sum_{i=1}^n y_i^* H_i + Q \right),$$

and thus (2.29) implies $d \in \mathcal{C}(\sum_{i=1}^n y_i^* H_i + Q)$. Therefore, there exists a vector $v \in \mathbb{R}^n$ such that $d = (\sum_{i=1}^n y_i^* H_i + Q)v$. We then have

$$\begin{aligned} &d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* \\ &= v^T \left(\sum_{i=1}^n y_i^* H_i + Q \right) \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* \\ &= v^T \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* \\ &= d^T y^*. \end{aligned}$$

Now let

$$\delta := z^* - y^{*T} \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* - 2d^T y^* - d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ d$$

and observe that $\delta \geq 0$. We can then write the objective value of (2.28) at (Y^*, y^*, z^*) in terms of D, d , and δ :

¹⁰Here, A^+ refers to any pseudo-inverse of A , i.e. a matrix satisfying $AA^+A = A$.

$$\begin{aligned}
& \frac{1}{2}\text{Tr}(QY^*) + b^T y^* + \frac{1}{2}z^* \\
&= \frac{1}{2} (y^{*T} Q y^* + \text{Tr}(QD)) + \sum_{i=1}^n \left(-e_i^T Q y^* - \frac{1}{2}\text{Tr}(H_i y^* y^{*T}) - \frac{1}{2}\text{Tr}(H_i D) \right) y_i^* \\
& \quad + \frac{1}{2} \left(y^{*T} \left(\sum_{i=1}^n y_i^* H_i + Q \right) y^* + 2d^T y^* + d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ d + \delta \right) \\
&= \left(\frac{1}{2} - 1 + \frac{1}{2} \right) y^{*T} Q y^* + \left(-\frac{1}{2} + \frac{1}{2} \right) \sum_{i=1}^n y_i^{*T} H_i y_i^* \tag{2.30} \\
& \quad + \frac{1}{2}\text{Tr}(QD) + \left(-\frac{1}{2} + 1 \right) \sum_{i=1}^n \text{Tr}(H_i D) y_i^* + \frac{1}{2} d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ d + \frac{\delta}{2} \\
&= \frac{1}{2}\text{Tr} \left(\left(\sum_{i=1}^n y_i^* H_i + Q \right) D \right) + \frac{1}{2} d^T \left(\sum_{i=1}^n y_i^* H_i + Q \right)^+ d + \frac{\delta}{2} \\
&\geq 0,
\end{aligned}$$

where in the last inequality we used the facts that $D \succeq 0$ and that the pseudo-inverse of a psd matrix is psd.

Since the left-hand side of the above equation is zero by assumption, and since all three terms on the right-hand side are nonnegative, it follows that $(\sum_{i=1}^n y_i^* H_i + Q)^+ d = 0$. As the null space of $(\sum_{i=1}^n y_i^* H_i + Q)^+$ is the same as the null space of $(\sum_{i=1}^n y_i^* H_i + Q)$, we have $(\sum_{i=1}^n y_i^* H_i + Q)d = 0$. However, because $d \in \mathcal{C}(\sum_{i=1}^n y_i^* H_i + Q)$, it must be that $d = 0$.

□

An Algorithm for Finding Local Minima

Theorem 2.6.4 leads to the following characterization of second-order points of a cubic polynomial.

Corollary 2.6.5. *Let $p : \mathbb{R}^n \rightarrow \mathbb{R}$ be a cubic polynomial written in the form (2.4). Then the set of its second-order points is equal to*

$$\{y \in \mathbb{R}^n \mid \exists Y \in \mathbb{S}^{n \times n}, z \in \mathbb{R} \text{ such that}$$

$$\frac{1}{2}\text{Tr}(QY) + b^T y + \frac{z}{2} = 0, \frac{1}{2}\text{Tr}(H_i Y) + e_i^T Qy + b_i = 0, \forall i = 1, \dots, n,$$

$$T(Y, y, z) \succeq 0, \begin{bmatrix} Y & y \\ y^T & 1 \end{bmatrix} \succeq 0\}.$$
(2.31)

Proof. Recall from the proof of Theorem 2.6.4 that if \bar{x} is a second-order point of p , then the triplet $(\bar{x}\bar{x}^T, \bar{x}, \bar{x}^T(\sum_{i=1}^n \bar{x}_i H_i + Q)\bar{x})$ is feasible solution to (2.28) with objective value zero. Hence any second-order point belongs to (2.31). Conversely, recall that if (Y, y, z) is a feasible solution to (2.28) with objective value zero, then y is a second-order point of p . Therefore any point in (2.31) is a second-order point of p . \square

In view of Theorem 2.4.7, we observe that if p has a local minimum, the set in (2.31) is a semidefinite representation of \overline{LM}_p . This observation gives rise to the following algorithm which tests if a cubic polynomial has a local minimum.

Algorithm 2 Algorithm for finding a local minimum of a cubic polynomial using a polynomial number of calls to E-SDP.

1: **Input:** A cubic polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ in the form (2.4)

2: **TEST1** test using E-SDP if (2.31) is empty

3: **if** YES

4: **return** NO LOCAL MINIMUM

5: **if** NO

6: Find (via Lemma 2.6.1) a point x^* in the relative interior of (2.31)

7: **TEST2** test (via Theorem 2.3.3) if x^* is a local minimum

8: **if** YES

9: **return** x^*

10: **if** NO

11: **return** NO LOCAL MINIMUM

Complexity and correctness of Algorithm 2. By design, if p has no local minimum, Algorithm 2 will return NO LOCAL MINIMUM since TEST2 answers NO for every point. If p has a local minimum, then SO_p is nonempty. Since SO_p is given by (2.31) due to Corollary 2.6.5, TEST1 answers YES. Then, by Theorem 2.4.10, any point in the relative interior of (2.31) is a local minimum. Hence x^* will pass TEST2. Note that this algorithm makes $2n + 1$ calls to E-SDP, and then runs Algorithm 1.¹¹

Remark 2.6.1. Finding strict local minima. If we are specifically interested in searching for a strict local minimum of a cubic polynomial, we can simply check if the point x^* returned by Algorithm 2 satisfies $\nabla^2 p(x^*) \succ 0$. If the answer is yes, we return x^* ; if the answer is no, we declare that p has no strict local minimum. Clearly, if a local minimum x^* satisfies $\nabla^2 p(x^*) \succ 0$, it must be a strict local minimum due to the SOSC. Furthermore, recall from

¹¹In fact, the number of calls to E-SDP can be reduced to $2n$ if the very first call to E-SDP uses x_1 as the objective function.

Section 2.3.1 that if p has a strict local minimum, then it has a unique local minimum, and thus that must be the output of Algorithm 2.

2.7 Conclusions and Future Directions

In this chapter, we considered the notions of (i) critical points, (ii) second-order points, (iii) local minima, and (iv) strict local minima for multivariate polynomials. For each type of point, and as a function of the degree of the polynomial, we studied the complexity of deciding (1) if a given point is of that type, and (2) if a polynomial has a point of that type. See Tables 2.1 and 2.2 in Section 2.1 for a summary of how our results complement prior literature. The majority of our work was dedicated to the case of cubic polynomials, where some new tractable cases were revealed based in part on connections with semidefinite programming. In this final section, we outline two future research directions which also have to do with cubic polynomials.

2.7.1 Approximate Local Minima

In Sections 2.5 and 2.6, we established polynomial-time equivalence of finding local minima and second-order points of cubic polynomials and some SDP feasibility problems (see Corollary 2.6.5, Algorithm 2, Theorem 2.5.3, Theorem 2.5.4). Unless some well-known open problems around the complexity of SDP feasibility are resolved (see Section 2.5), one cannot expect to make claims about finding local minima of cubic polynomials in polynomial time in the Turing model of computation. Nonetheless, it is known that under some assumptions, one can solve semidefinite programs to arbitrary accuracy in polynomial time (see, e.g. [94, 8, 108, 90, 80, 46]). It is therefore reasonable to ask if one can find local minima of cubic polynomials to arbitrary accuracy in polynomial time. This is a question we would like to study more rigorously in future work. We present a partial result in this direction in Theorem 2.7.1 below.

Recall from Section 2.6.2 that our ability to find local minima of a cubic polynomial p depended on our ability to minimize p over its convexity region CR_p . We show next that we can find an ϵ -minimizer of p over CR_p by approximately solving a semidefinite program.

Theorem 2.7.1. *For a cubic polynomial p given in the form (2.4), consider the SDP in (2.28). If the objective value at a feasible point (Y, y, z) is $\epsilon \geq 0$, then $p(y) \leq p(x) + \frac{2}{3}\epsilon$, $\forall x \in CR_p$.*

Proof. Consider a feasible solution (Y, y, z) to (2.28). Let γ^* be the infimum of p over CR_p . Observe that

$$-\frac{1}{6}\text{Tr}(QY) - \frac{z}{3} \leq \gamma^*.$$

This is because the SDPs in (2.28) and (2.26) have the same constraints, and the optimal value of (2.26) is the negative of the optimal value of (2.25), which by construction is a lower bound on γ^* . Similarly as in the proof of Theorem 2.6.4, let $D := Y - yy^T$, $d := \sum_{i=1}^n \text{Tr}(H_i D)e_i$, and

$$\delta := z - y^T \left(\sum_{i=1}^n y_i H_i + Q \right) y - 2d^T y - d^T \left(\sum_{i=1}^n y_i H_i + Q \right)^+ d.$$

We can then write:

$$\begin{aligned}
\frac{1}{6}\text{Tr}(QY) + \frac{z}{3} &= \frac{1}{6}\text{Tr}(QY) + \frac{z}{3} - \sum_{i=1}^n \left(\frac{1}{2}\text{Tr}(H_i Y) + e_i^T Q y + b_i \right) y_i \\
&= \frac{1}{6} (\text{Tr}(Qyy^T) + \text{Tr}(QD)) \\
&+ \frac{1}{3} \left(y^T \left(\sum_{i=1}^n y_i H_i + Q \right) y + 2d^T y + d^T \left(\sum_{i=1}^n y_i H_i + Q \right)^+ d + \delta \right) \\
&- \frac{1}{2} \left(\text{Tr} \left(\sum_{i=1}^n y_i H_i y y^T \right) + \text{Tr} \left(\sum_{i=1}^n y_i H_i D \right) \right) - y^T Q y - b^T y \\
&= -\frac{1}{6} \sum_{i=1}^n y^T y_i H_i y - \frac{1}{2} y^T Q y - b^T y \\
&+ \frac{1}{6} \text{Tr} \left(\left(\sum_{i=1}^n y_i H_i + Q \right) D \right) + \frac{1}{3} \left(d^T \left(\sum_{i=1}^n y_i H_i + Q \right)^+ d \right) + \frac{\delta}{3} \\
&= -p(y) + \frac{1}{6} \text{Tr} \left(\left(\sum_{i=1}^n y_i H_i + Q \right) D \right) + \frac{1}{3} \left(d^T \left(\sum_{i=1}^n y_i H_i + Q \right)^+ d \right) + \frac{\delta}{3} \\
&\leq -p(y) + \frac{2}{3}\epsilon,
\end{aligned}$$

where the first equality is due to the first constraint in (2.28), and the last inequality follows from the last equation of (2.30) with (Y^*, y^*, z^*) replaced by (Y, y, z) and the fact that $\sum_{i=1}^n y_i H_i + Q$ and D are both psd matrices. We therefore conclude that

$$p(y) - \frac{2}{3}\epsilon \leq -\frac{1}{6}\text{Tr}(QY) - \frac{z}{3} \leq \gamma^*.$$

We then have that $p(y) \leq p(x) + \frac{2}{3}\epsilon, \forall x \in CR_p$ as desired.

□

2.7.2 Unregularized Third-Order Newton Methods

We end our chapter with an interesting application of the problem of finding a local minimum of a cubic polynomial. Recall that Newton's method for minimizing a twice-differentiable

function proceeds by approximating the function with its second-order Taylor expansion at the current iterate, and then moving to a critical point¹² of this quadratic approximation. It is natural to ask whether one can lower the iteration complexity of Newton’s method for three-times-differentiable functions by using third-order information. An immediate difficulty, however, is that the third-order Taylor expansion of a function around any point will not be bounded below (unless the coefficients of all its cubic terms are zero). In previous work (see, e.g. [78]), authors have gotten around this issue by adding a regularization term to the third-order Taylor expansion. In future work, we aim to study an unregularized third-order Newton method which in each iteration moves to a local minimum of the third-order Taylor approximation by applying Algorithm 2. We would like to explore the convergence properties of this algorithm and conditions under which the algorithm is well defined at every iteration.

As a first step, let us consider the univariate case. For a function $f : \mathbb{R} \rightarrow \mathbb{R}$, the iterations of (classical) Newton’s method read

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)}. \tag{2.32}$$

The update rule of a third-order Newton method, which in each iteration moves to the local minimum of the third-order Taylor approximation, is given by

$$x_{k+1} = x_k - \frac{f''(x_k) - \sqrt{f''(x_k)^2 - 2f'(x_k)f'''(x_k)}}{f'''(x_k)}. \tag{2.33}$$

We have already observed that in some settings, these iterations can outperform the classical Newton iterations. For example, consider the univariate function

$$f(x) = 20x \arctan(x) - 10 \log(1 + x^2) + x^2, \tag{2.34}$$

¹²If the function to be minimized is convex, this critical point will be a global minimum of the quadratic approximation.

which is strongly convex and has a (unique) global minimum at $x = 0$, where $f(x) = 0$; see Figure 2.5. The first three derivatives of this function are

$$f'(x) = 20 \arctan(x) + 2x,$$

$$f''(x) = 2 + \frac{20}{1+x^2},$$

$$f'''(x) = \frac{-40x}{(1+x^2)^2}.$$

One can show that the basin of attraction of the global minimum of f under the classical Newton iterations in (2.32) is approximately $[-1.7121, 1.7121]$. Starting Newton's method with $|x_0| \geq 1.7122$ results in the iterates eventually oscillating between ± 13.4942 . In contrast, the iterates of our proposed third-order Newton method in (2.33) are globally convergent to the global minimum of f . The iterations of both methods starting at $x_0 = 1.5$ are compared in Table 2.3 and Figure 2.5, showing faster convergence to the global minimum for the third-order approach.

k	x_k	$f(x_k)$	k	x_k	$f(x_k)$
0	1.5	19.9473	0	1.5	19.9473
1	-.2327	.5910	1	-1.2786	15.1411
2	-.0030	1.0014e-4	2	.8795	7.7329
3	-8.3227e-9	1.4546e-15	3	-.3396	1.2477
4	2.3490e-9	1.1587e-16	4	.0230	.0058

Table 2.3: Iterations of the third-order Newton method (left) and the classical Newton method (right) on the function f in (2.34) starting at $x_0 = 1.5$.

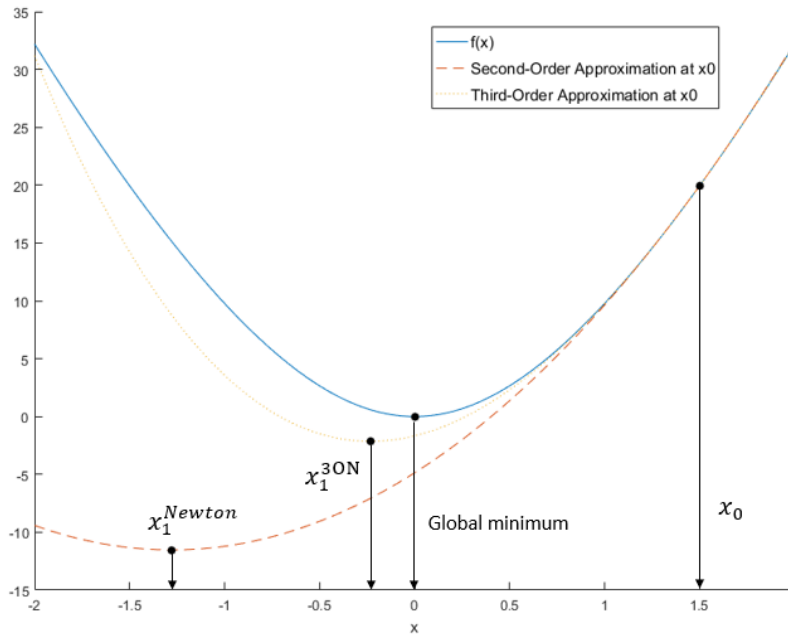


Figure 2.5: The plots of the function f in (2.34) and its second and third-order Taylor expansions around $x_0 = 1.5$. One can see that one iteration of the third-order Newton method in (2.33) brings us closer the global minimum of f compared to one iteration of the Newton method in (2.32).

In addition to potential benefits regarding convergence, we have also observed that the behavior of the algorithm can be less sensitive to the initial condition when compared to Newton’s method. As an example, we used Newton’s method to find the critical points $\{1, -1, i, -i\}$ of $f(x) = x^5 - 5x$ on the complex plane, using the iterates (2.32), (2.33), and iterates given by

$$x_{k+1} = x_k - \frac{f''(x_k) + \sqrt{f''(x_k)^2 - 2f'(x_k)f'''(x_k)}}{f'''(x_k)}, \quad (2.35)$$

which can be interpreted as the iterates for moving to the local maximum of a third-order approximation of f . For each of the three iterations, the plots below demonstrate which initial conditions converge to the same critical point. As can be seen, sensitivity of Newton’s

method to the initial condition demonstrates fractal behavior, while the third-order iterates do not.

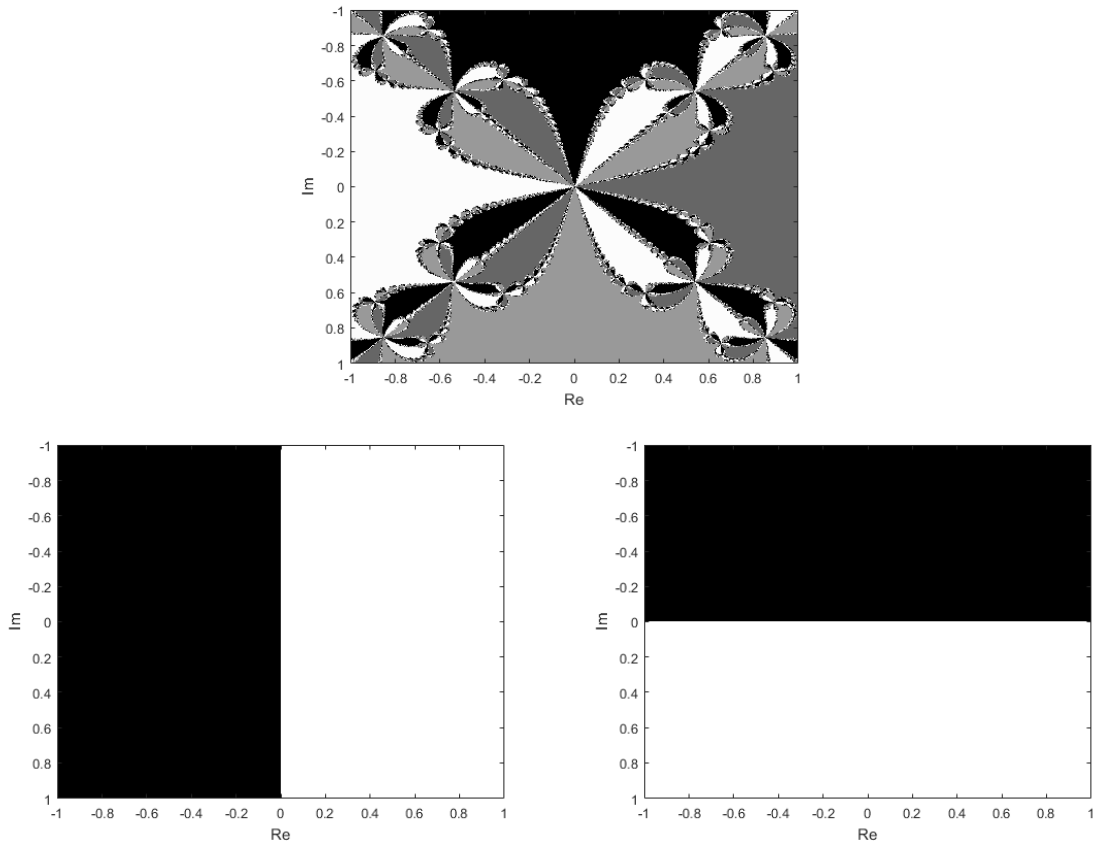


Figure 2.6: Sensitivity of the limits of the iterates (2.32), (2.33), and (2.35) respectively to initial conditions. Regions with the same color denote initial conditions which converge to the same critical point.

Chapter 3

On the Complexity of Finding a Local Minimizer of a Quadratic Function over a Polytope

3.1 Introduction

In this chapter of the thesis, we consider *quadratic programs*, which are polynomial optimization problems of the form (1.1) where the objective function p is quadratic and all constraint functions q_i are affine. Recall that a *local minimum* of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ over a set $\Omega \subseteq \mathbb{R}^n$ is a point $\bar{x} \in \Omega$ for which there exists a scalar $\epsilon > 0$ such that $p(\bar{x}) \leq p(x)$ for all $x \in \Omega$ with $\|x - \bar{x}\| \leq \epsilon$. In the case where p and all the constraint functions q_i are affine (i.e., linear programming), it is well known that a local minimum (which also has to be a global minimum) can be found in polynomial time in the Turing model of computation [56, 57]. Perhaps the next simplest constrained optimization problems to consider are

quadratic programs, which can be written as

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & x^T Q x + c^T x \\ \text{subject to} \quad & a_i^T x \leq b_i, \forall i \in \{1, \dots, m\}, \end{aligned} \tag{3.1}$$

where $Q \in \mathbb{R}^{n \times n}$, $c, a_1, \dots, a_m \in \mathbb{R}^n$, and $b_1, \dots, b_m \in \mathbb{R}$. The matrix Q is taken without loss of generality to be symmetric. When complexity questions about quadratic programs are studied in the Turing model of computation, all these data are rational and the input size is the total number of bits required to write them down. It is well known that finding a global minimum of a quadratic program is NP-hard, even when the matrix Q has a single negative eigenvalue [86]. It is therefore natural to ask whether one can instead find a local minimum of a quadratic program efficiently. In fact, this precise question appeared in 1992 on a list of seven open problems in complexity theory for numerical optimization [87]:

“What is the complexity of finding even a local minimizer for nonconvex quadratic programming, assuming the feasible set is compact? Murty and Kabadi (1987, [75]) and Pardalos and Schnitger (1988, [85]) have shown that it is NP-hard to test whether a given point for such a problem is a local minimizer, but that does not rule out the possibility that another point can be found that is easily verified as a local minimizer.”

A few remarks on the phrasing of this problem are in order. First, note that in this question, the feasible set of the quadratic program is assumed to be compact (i.e., a polytope). Therefore, there is no need to focus on the related and often prerequisite problem of deciding the *existence* of a local minimum (since any global minimum e.g. is a local minimum). The latter question makes sense in the case where the feasible set of the quadratic program is unbounded; the complexity of this question is also addressed in this chapter (Theorem 3.2.5). Second, as the quote points out, the question of finding a local minimum is also separate from a complexity viewpoint from that of testing if a given point is a local minimum. This related

question has been studied more extensively and its complexity has already been settled for optimization problems whose objective and constraints are given by polynomial functions of any degree; see [75, 85, 7].

To point out some of the subtle differences between these variations of the problem more specifically, we briefly review the reduction of Murty and Kabadi [75], which shows the NP-hardness of deciding if a given point is a local minimum of a quadratic program. In [75], the authors show that the problem of deciding if a symmetric matrix Q is *copositive*—i.e. whether $x^T Q x \geq 0$ for all vectors x in the nonnegative orthant—is NP-hard. From this, it is straightforward to observe that the problem of testing whether a given point is a local minimum of a quadratic function over a polyhedron is NP-hard: Indeed, the origin is a local minimum of $x^T Q x$ over the nonnegative orthant if and only if the matrix Q is copositive. However, it is not true that $x^T Q x$ has a local minimum over the nonnegative orthant if and only if Q is copositive. Although the “if” direction holds, the “only if” direction does not. For example, consider the matrix

$$Q = \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix},$$

which is clearly not copositive, even though the point $(1, 0)^T$ is a local minimum of $x^T Q x$ over the nonnegative orthant.

Our main results in this chapter are as follows. We show that unless $P=NP$, no polynomial-time algorithm can find a point within Euclidean distance c^n (for any constant $c \geq 0$) of a local minimum of an n -variate quadratic program with a bounded feasible set (Theorem 3.2.6). See also Corollaries 3.2.7 and 3.2.8. To prove this, we show as an intermediate step that deciding whether a quartic polynomial or a quadratic program has a local minimum is strongly NP-hard¹ (Theorems 3.2.1 and 3.2.5). Finally, we show that

¹This implies that these problems remain NP-hard even if the bitsize of all numerical data are $O(\log n)$, where n is the number of variables. For a strongly NP-hard problem, even a pseudo-polynomial time algorithm—i.e., an algorithm whose running time is polynomial in the magnitude of the numerical data of the problem but not necessarily in their bitsize—cannot exist unless $P=NP$. See [39] or [5, Section 2] for more details.

unless $P=NP$, there cannot be a polynomial-time algorithm that decides if a quadratic program with a bounded feasible set has a unique local minimum and if so returns this minimum (Theorem 3.2.9).

Overall, our results suggest that without additional problem structure, questions related to finding local minima of quadratic programs are not easier (at least from a complexity viewpoint) than those related to finding global minima. It also suggests that any efficient heuristic that aims to find a local minimum of a quadratic program must necessarily fail on a “significant portion” of instances; see e.g. Corollary 2.2 of [48] for a more formal complexity theoretic statement.

3.1.1 Notation and Basic Definitions

For a vector $x \in \mathbb{R}^n$, the notation x^2 denotes the vector in \mathbb{R}^n whose i -th entry is x_i^2 , and $diag(x)$ denotes the diagonal $n \times n$ matrix whose i -th diagonal entry is x_i . The notation $x \geq 0$ denotes that the vector x belongs to the nonnegative orthant, and for such a vector, \sqrt{x} denotes the vector in \mathbb{R}^n whose i -th entry is $\sqrt{x_i}$. For two matrices $X, Y \in \mathbb{R}^{m \times n}$, we denote by $X \cdot Y$ the matrix in $\mathbb{R}^{m \times n}$ whose (i, j) -th entry is $X_{ij}Y_{ij}$. For vectors $x, y \in \mathbb{R}^n$, the notation y_x (sometimes $(y)_x$ if there is room for confusion with other indices) denotes the vector containing the entries of y where x_i is nonzero in the same order (the length of y_x is hence equal to the number of nonzero entries in x). Similarly, for a vector $x \in \mathbb{R}^n$ and a matrix $Y \in \mathbb{R}^{n \times n}$, the notation Y_x (sometimes $(Y)_x$ if there is room for confusion with other indices) denotes the principal submatrix of Y consisting of rows and columns of Y whose indices correspond to indices of nonzero entries of x . The notation I (resp. J) refers to the identity matrix (resp. the matrix of all ones); the dimension will be clear from context. For a symmetric matrix $M \in \mathbb{R}^{n \times n}$, the notation $M \succeq 0$ (resp. $M \succ 0$) denotes that M is *positive semidefinite* (resp. *positive definite*), i.e. that it has nonnegative (resp. positive) eigenvalues. As mentioned already, we say that M is *copositive* if $x^T M x \geq 0, \forall x \geq 0$. The

simplex in \mathbb{R}^n is denoted by $\Delta_n := \{x \in \mathbb{R}^n \mid x \geq 0, \sum_{i=1}^n x_i = 1\}$. Finally, for a scalar c , the notation $\lceil c \rceil$ denotes the ceiling of c , i.e. the smallest integer greater than or equal to c .

We recall that a *form* is a homogeneous polynomial; i.e. a polynomial whose monomials all have the same degree. A form $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be *nonnegative* if $p(x) \geq 0, \forall x \in \mathbb{R}^n$, and *positive definite* if $p(x) > 0, \forall x \neq 0$. A *critical point* of a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a point $x \in \mathbb{R}^n$ at which the gradient $\nabla f(x)$ is zero. A *second-order point* of a twice-differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a critical point x at which the Hessian matrix $\nabla^2 f(x)$ is positive semidefinite.

All graphs in this chapter are undirected, unweighted, and have no self-loops. The *adjacency matrix* of a graph G on n vertices is the $n \times n$ symmetric matrix whose (i, j) -th entry equals one if vertices i and j share an edge in G and zero otherwise. The *complement* of a graph G , denoted by \bar{G} , is the graph with the same vertex set as G and such that two distinct vertices are adjacent if and only if they are not adjacent in G . An *induced subgraph* of G is a graph containing a subset of the vertices of G and all edges connecting pairs of vertices in that subset.

3.2 The Main Result

3.2.1 Complexity of Deciding Existence of Local minima

To show that a polynomial-time algorithm for finding a local minimum of a quadratic function over a polytope (i.e., a bounded polyhedron) implies $P = NP$, we show as an intermediate step that it is NP-hard to decide whether a quadratic program with an unbounded feasible set has a local minimum (Theorem 3.2.5). To achieve this intermediate step, we first establish the following hardness result.

Theorem 3.2.1. *It is strongly NP-hard to decide if a degree-4 polynomial has a local minimum.*

We will prove this theorem by presenting a polynomial-time reduction from the STABLE-SET problem, which is known to be (strongly) NP-hard [39]. Recall that in the STABLESET problem, we are given as input a graph G on n vertices and a positive integer $r \leq n$. We are then asked to decide whether G has a *stable set* of size r , i.e. a set of r pairwise non-adjacent vertices. We denote the size of the largest stable set in a graph G by the standard notation $\alpha(G)$. We also recall that a *clique* in a graph G is a set of pairwise adjacent vertices. The size of the largest clique in G is denoted by $\omega(G)$. The following theorem of Motzkin and Straus [74] relates $\omega(G)$ to the optimal value of a quadratic program.

Theorem 3.2.2 ([74]). *Let G be a graph on n vertices with adjacency matrix A and clique number ω . The optimal value of the quadratic program*

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & x^T A x \\ \text{subject to} \quad & x \geq 0, \\ & \sum_{i=1}^n x_i = 1 \end{aligned} \tag{3.2}$$

is $1 - \frac{1}{\omega}$.

For a scalar k and a symmetric matrix A (which will always be an adjacency matrix), the following notation will be used repeatedly in our proofs:

$$M_{A,k} := kA + kI - J, \tag{3.3}$$

$$q_{A,k}(x) := x^T M_{A,k} x, \tag{3.4}$$

and

$$p_{A,k}(x) := (x^2)^T M_{A,k} x^2. \tag{3.5}$$

Note that nonnegativity of the quadratic form $q_{A,k}$ over the nonnegative orthant is equivalent to (global) nonnegativity of the quartic form $p_{A,k}$ and to copositivity of the matrix $M_{A,k}$.

The following corollary of Theorem 3.2.2 will be of more direct relevance to our proofs. The first statement in the corollary has been observed e.g. by de Klerk and Pashechnik [35], but its proof is included here for completeness. The second statement, which will also be needed in the proof of Theorem 3.2.1, follows straightforwardly.

Corollary 3.2.3. *For a scalar $k > 0$ and a graph G with adjacency matrix A , the matrix $M_{A,k}$ in (3.3) is copositive if and only if $\alpha(G) \leq k$. Furthermore, if $\alpha(G) < k$, the quartic form $p_{A,k}$ in (3.5) is positive definite.²*

Proof. First observe that $\alpha(G) = \omega(\bar{G})$, and that the adjacency matrix of \bar{G} is $J - A - I$. Thus from Theorem 3.2.2, the maximum value of $x^T(J - A - I)x$ over Δ_n is $1 - \frac{1}{\alpha(G)}$, and hence the minimum value of $x^T(A + I)x$ over Δ_n is $\frac{1}{\alpha(G)}$. Therefore, for any $k > 0$, $\alpha(G) \leq k$ if and only if $x^T(A + I)x \geq \frac{1}{k}$ for all $x \in \Delta_n$, which holds if and only if $x^T(k(A + I) - J)x \geq 0$ for all $x \in \Delta_n$. The first statement of the corollary then follows from the homogeneity of $x^T(k(A + I) - J)x$.

To show that $p_{A,k}$ is positive definite when $\alpha(G) < k$, observe that

$$kA + kI - J = (\alpha(G)(A + I) - J) + (k - \alpha(G))(A + I).$$

Considering the two terms on the right separately, we observe that $(x^2)^T(\alpha(G)(A + I) - J)x^2$ (i.e., $p_{A,\alpha(G)}$) is nonnegative since $M_{A,\alpha(G)}$ is copositive, and that $(x^2)^T(k - \alpha(G))(A + I)x^2$ is positive definite. Therefore, their sum $(x^2)^T(kA + kI - J)x^2$ is positive definite. \square

We now present the proof of Theorem 3.2.1. While the statement of the theorem is given for degree-4 polynomials, it is straightforward to extend the result to higher-degree polynomials. We note that degree four is the smallest degree for which deciding existence of local minima is intractable. For degree-3 polynomials, it turns out that this question can be answered by solving semidefinite programs of polynomial size [7].

²The converse of this statement also holds, but we do not need it for the proof of Theorem 3.2.1.

Proof (of Theorem 3.2.1). We present a polynomial-time reduction from the STABLESET problem. Let a graph G on n vertices with adjacency matrix A and a positive integer $r \leq n$ be given. We show that G has a stable set of size r if and only if the quartic form $p_{A,r-0.5}$ defined in (3.5) has no local minimum. This is a consequence of the following more general fact that we prove below: For a noninteger scalar k , the quartic form $p_{A,k}$ has no local minimum if and only if $\alpha(G) \geq k$.

We first observe that if $\alpha(G) < k$, then $p_{A,k}$ has a local minimum. Indeed, recall from the second claim of Corollary 3.2.3 that under this assumption, $p_{A,k}$ is positive definite. Since $p_{A,k}$ vanishes at the origin, it follows that the origin is a local minimum. Suppose now that $\alpha(G) \geq k$. Since k is noninteger, this implies that $\alpha(G) > k$. We show that in this case, $p_{A,k}$ has no local minimum by showing that the origin must be the only second-order point of $p_{A,k}$. Since any local minimum of a polynomial is a second-order point, only the origin can be a candidate local minimum for $p_{A,k}$. However, by the first claim of Corollary 3.2.3, the matrix $M_{A,k}$ is not copositive and hence $p_{A,k}$ is not nonnegative. As $p_{A,k}$ is homoeogenous, this implies that $p_{A,k}$ takes negative values arbitrarily close to the origin, ruling out the possibility of the origin being a local minimum.

To show that when $\alpha(G) > k$, the origin is the only second-order point of $p_{A,k}$, we compute the gradient and Hessian of $p_{A,k}$. We have

$$\nabla p_{A,k}(x) = 4x \cdot M_{A,k}x^2,$$

and

$$\nabla^2 p_{A,k}(x) = 8M_{A,k} \cdot xx^T + 4diag(M_{A,k}x^2).$$

Suppose for the sake of contradiction that $p_{A,k}$ has a nonzero second-order point \bar{x} . Since \bar{x} is a critical point, $\nabla p_{A,k}(\bar{x}) = 0$ and thus $(M_{A,k}\bar{x}^2)_{\bar{x}} = 0$. It then follows that

$$(\nabla^2 p_{A,k}(\bar{x}))_{\bar{x}} = 8(M_{A,k})_{\bar{x}} \cdot \bar{x}_{\bar{x}}\bar{x}_{\bar{x}}^T.$$

Because $\nabla^2 p_{A,k}(\bar{x}) \succeq 0$ and thus all its principal submatrices are positive semidefinite, we have $8(M_{A,k})_{\bar{x}} \cdot \bar{x}_{\bar{x}} \bar{x}_{\bar{x}}^T \succeq 0$. Since

$$(M_{A,k})_{\bar{x}} \cdot \bar{x}_{\bar{x}} \bar{x}_{\bar{x}}^T = \text{diag}(\bar{x}_{\bar{x}})(M_{A,k})_{\bar{x}} \text{diag}(\bar{x}_{\bar{x}}),$$

and since $\text{diag}(\bar{x}_{\bar{x}})$ is an invertible matrix, it follows that $(M_{A,k})_{\bar{x}} \succeq 0$.

We now consider the induced subgraph $G_{\bar{x}}$ of G with vertices corresponding to the indices of the nonzero entries of \bar{x} . Note that the adjacency matrix of $G_{\bar{x}}$ is $A_{\bar{x}}$. Furthermore, observe that $M_{A_{\bar{x}},k} = (M_{A,k})_{\bar{x}}$, and therefore $M_{A_{\bar{x}},k}$ is positive semidefinite and thus copositive. We conclude from the first claim of Corollary 3.2.3 that $\alpha(G_{\bar{x}}) \leq k$. We now claim that

$$M_{A_{\bar{x}},k} \bar{x}_{\bar{x}}^2 = (M_{A,k})_{\bar{x}} \bar{x}_{\bar{x}}^2 = (M_{A,k} \bar{x}^2)_{\bar{x}} = 0.$$

The first equality follows from that $M_{A_{\bar{x}},k} = (M_{A,k})_{\bar{x}}$, the second from that the indices of the nonzero entries of \bar{x} are the same as those of \bar{x}^2 , and the third from that $\nabla p_{A,k}(\bar{x}) = 0$. Hence, $p_{A_{\bar{x}},k}(\bar{x}_{\bar{x}}) = 0$. Since $\bar{x}_{\bar{x}}$ is nonzero, $p_{A_{\bar{x}},k}$ is not positive definite. By the second claim of Corollary 3.2.3, we must have $\alpha(G_{\bar{x}}) \geq k$. Therefore, $\alpha(G_{\bar{x}}) = k$. However, because k was assumed to be noninteger, we have a contradiction. \square

It turns out that the proof of Theorem 3.2.1 also shows that it is NP-hard to decide if a quartic polynomial has a strict local minimum. Recall that a *strict local minimizer* of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ over a set $\Omega \subseteq \mathbb{R}^n$ is a point $\bar{x} \in \Omega$ for which there exists a scalar $\epsilon > 0$ such that $p(\bar{x}) < p(x)$ for all $x \in \Omega \setminus \bar{x}$ with $\|x - \bar{x}\| \leq \epsilon$.

Corollary 3.2.4. *It is strongly NP-hard to decide if a degree-4 polynomial has a strict local minimum.*

Proof. Observe from the proof of Theorem 3.2.1 that for a graph G and a noninteger scalar k , the quartic form $p_{A,k}$ has a local minimum if and only if $\alpha(G) < k$. In the case where

$p_{A,k}$ does have a local minimum, we showed that $p_{A,k}$ is positive definite, and thus the local minimum (the origin) must be a strict local minimum. \square

We now turn our attention to local minima of quadratic programs.

Theorem 3.2.5. *It is strongly NP-hard to decide if a quadratic function has a local minimum over a polyhedron. The same is true for deciding if a quadratic function has a strict local minimum over a polyhedron.*

Proof. We present a polynomial-time reduction from the STABLESET problem to the problem of deciding if a quadratic function has a local minimum over a polyhedron. The reader can check that same reduction is valid for the case of strict local minima.

Let a graph G on n vertices with adjacency matrix A and a positive integer $r \leq n$ be given. Let $k = r - 0.5$, $q_{A,k}$ be the quadratic form defined in (3.4), and consider the optimization problem

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & q_{A,k}(x) \\ \text{subject to} \quad & x \geq 0. \end{aligned} \tag{3.6}$$

We show that a point $x \in \mathbb{R}^n$ is a local minimum of (3.6) if and only if \sqrt{x} is a local minimum of the quartic form $p_{A,k}$ defined in (3.5). By the arguments in the proof of Theorem 3.2.1, we would have that (3.6) has no local minimum if and only if G has a stable set of size r .

Indeed, if x is not a local minimum of (3.6), there exists a sequence $\{y_j\} \subseteq \mathbb{R}^n$ with $y_j \rightarrow x$ and such that for all j , $y_j \geq 0$ and $q_{A,k}(y_j) < q_{A,k}(x)$. The sequence $\{\sqrt{y_j}\}$ would then satisfy $p_{A,k}(\sqrt{y_j}) < p_{A,k}(\sqrt{x})$ and $\sqrt{y_j} \rightarrow \sqrt{x}$, proving that \sqrt{x} is not a local minimum of $p_{A,k}$. Similarly, if x is not a local minimum of $p_{A,k}$, there exists a sequence $\{z_j\} \subseteq \mathbb{R}^n$ such that $z_j \rightarrow x$ and $p_{A,k}(z_j) < p_{A,k}(x)$ for all j . The sequence $\{z_j^2\}$ would then prove that x^2 is not a local minimum of (3.6). \square

3.2.2 Complexity of Finding a Local minimum of a Quadratic Function Over a Polytope

We now address the original question posed by Pardalos and Vavasis concerning the complexity of finding a local minimum of a quadratic program with a compact feasible set. Note again that if the feasible set is compact, the existence of a local minimum is guaranteed. In fact, there will always be a local minimum that has rational entries with polynomial bitsize [109].

Theorem 3.2.6. *If there is a polynomial-time algorithm that finds a point within Euclidean distance c^n (for any constant $c \geq 0$) of a local minimum of an n -variate quadratic function over a polytope, then $P = NP$.*

Proof. Fix any constant $c \geq 0$. We show that if an algorithm could take as input a quadratic program with a bounded feasible set and in polynomial time return a point within distance c^n of any local minimum, then this algorithm would solve the the STABLESET problem in polynomial time.

Let a graph G on n vertices with adjacency matrix A and a positive integer $r \leq n$ be given. Let $k = r - 0.5$, $q_{A,k}$ be the quadratic form defined in (3.4), and consider the quadratic program³

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & q_{A,k}(x) \\ \text{subject to} \quad & x \geq 0, \\ & \sum_{i=1}^n x_i \leq 3c^n \sqrt{n}. \end{aligned} \tag{3.7}$$

Note that the feasible set of this problem is bounded. Moreover, the number of bits required to write down this quadratic program is polynomial in n . This is because the scalar $3c^n \sqrt{n}$ takes $2 + n \lceil \log_2(c+1) \rceil + \frac{1}{2} \lceil \log_2(n+1) \rceil$ bits to write down, and the remaining $O(n^2)$ numbers

³Without loss of generality, we suppose that c is rational. If c is irrational, one can e.g. replace it in (3.7) with $\lceil c \rceil$.

in the problem data are bounded in magnitude by n , so they each take $O(\log_2(n))$ bits to write down.

We will show that if $\alpha(G) < k$, the origin is the unique local minimum of (3.7), and that if $\alpha(G) \geq k$ (equivalently $\alpha(G) > k$), any local minimum \bar{x} of (3.7) satisfies $\sum_{i=1}^n \bar{x}_i = 3c^n \sqrt{n}$. Since the (Euclidean) distance from the origin to the hyperplane $\{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 3c^n \sqrt{n}\}$ is $3c^n$, there is no point that is within distance c^n of both the origin and this hyperplane. Thus, the graph G has no stable set of size r (or equivalently $\alpha(G) < k$) if and only if the Euclidean norm of all points within distance c^n of any local minimum of (3.7) is less than or equal to c^n .

To see why $\alpha(G) < k$ implies that the origin is the unique local minimum of (3.7), recall from the second claim of Corollary 3.2.3 that the quartic form $p_{A,k}$ defined in (3.5) must be positive definite. Thus, for any nonzero vector $x \geq 0$, we have $q_{A,k}(x) > 0$. This implies that the origin is a local minimum of (3.7). Moreover, since $q_{A,k}$ is homogeneous, we have that no other feasible point can be a local minimum. Indeed, for any nonzero vector $x \geq 0$ and any nonnegative scalar $\epsilon < 1$, $q_{A,k}(\epsilon x) < q_{A,k}(x)$.

To see why when $\alpha(G) > k$, the last constraint of (3.7) must be tight at all local minima, recall from the proof of Theorem 3.2.5 that when $\alpha(G) > k$, the optimization problem in (3.6) has no local minimum. Therefore, for any vector x that is feasible to (3.7) and satisfies $\sum_{i=1}^n x_i < 3c^n \sqrt{n}$, there exists a sequence $\{y_i\} \subseteq \mathbb{R}^n$ with $y_i \rightarrow x$, and satisfying

$$y_i \geq 0, \sum_{i=1}^n y_i < 3c^n \sqrt{n}, q_{A,k}(y_i) < q_{A,k}(x), \forall i.$$

As the points y_i are feasible to (3.7), any vector x satisfying $\sum_{i=1}^n x_i < 3c^n \sqrt{n}$ cannot be a local minimum of (3.7). Thus, if $\alpha(G) > k$, any local minimum \bar{x} of (3.7) satisfies $\sum_{i=1}^n \bar{x}_i = 3c^n \sqrt{n}$. □

By replacing the quantity $3c^n \sqrt{n}$ in the proof of Theorem 3.2.6 with $3n^{c+0.5}$ and $2n$ respectively, we get the following two corollaries.

Corollary 3.2.7. *If there is a pseudo-polynomial time algorithm that finds a point within Euclidean distance n^c (for any constant $c \geq 0$) of a local minimum of an n -variate quadratic function over a polytope, then $P = NP$.*

Corollary 3.2.8. *If there is a polynomial-time algorithm that finds a point within Euclidean distance $\epsilon\sqrt{n}$ (for any constant $\epsilon \in [0, 1)$) of a local minimum of a restricted set of quadratic programs over n variables whose numerical data are integers bounded in magnitude by $2n$, then $P = NP$.*

In [87], Pardalos and Vavasis also propose two follow-up questions about quadratic programs with compact feasible sets. The first is about the complexity of finding a “KKT point”. As is, our proof does not have any implications for this question since the origin is always a KKT point of the quadratic programs that arise from our reductions. The second question asks whether finding a local minimum is easier in the special case where the problem only has one local minimum (which is thus also the global minimum). Related to this question, we can prove the following claim.

Theorem 3.2.9. *If there is a polynomial-time algorithm which decides whether a quadratic program with a bounded feasible set has a unique local minimum, and if so returns this minimum⁴, then $P=NP$.*

Proof. Suppose there was such an algorithm (call it Algorithm U). We show that Algorithm U would solve the STABLESET problem in polynomial time. Let a graph G on n vertices with adjacency matrix A and a positive integer $r \leq n$ be given, and input the quadratic program (3.7), with $k = r - 0.5$, into Algorithm U. Observe from the proof of Theorem 3.2.6 that there are three possibilities for this quadratic program: (i) the origin is the unique local minimum, (ii) there is a unique local minimum and it is on the hyperplane $\{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 3c^n \sqrt{n}\}$, and (iii) there are multiple local minima and they are all on the hyperplane $\{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 3c^n \sqrt{n}\}$. Case (i) indicates that $\alpha(G) < k$, and the

⁴This unique local (and therefore global) minimum is guaranteed to have rational entries with polynomial bitsize; see [109].

output of Algorithm U in this case would be the origin. Cases (ii) and (iii) both indicate that $\alpha(G) > k$. The output of Algorithm U is a point away from the origin in case (ii), and the declaration that the local minimum is not unique in case (iii). Thus Algorithm U would reveal which case we are in, and that would allow us to decide if G has a stable set of size r in polynomial time. \square

To conclude, we have established intractability of several problems related to local minima of quadratic programs. We hope our results motivate more research on identifying classes of quadratic programs where local minima can be found more efficiently than global minima. One interesting example is the case of the concave knapsack problem, where Moré and Vavasis [72] show that a local minimum can be found in polynomial time even though, unless $P=NP$, a global minimum cannot.

Chapter 4

On Attainment of the Optimal Value in Polynomial Optimization

4.1 Introduction

In this chapter, we again consider problems of the form

$$\begin{aligned} \inf_x \quad & p(x) \\ \text{subject to} \quad & q_i(x) \geq 0, \forall i \in \{1, \dots, m\}, \end{aligned} \tag{4.1}$$

where p, q_i are polynomial functions, and address the problem of testing whether the optimal value is *attained*, provided that the optimal value p^* is finite. More formally, does there exist a feasible point x^* such that $p(x^*) = p^*$? Such a point x^* will be termed an *optimal solution*. For this chapter, we will refer to sets of the type $\{x \in \mathbb{R}^n \mid q_i(x) \geq 0, \forall i \in \{1, \dots, m\}\}$ as *closed basic semialgebraic sets*.

Existence of optimal solutions is a fundamental question in optimization and its study has a long history, dating back to the nineteenth century with the extreme value theorem of

Bolzano and Weierstrass.¹ The question of testing attainment of the optimal value for POPs has appeared in the literature explicitly. For example, Nie, Demmel, and Sturmfels describe an algorithm for globally solving an unconstrained POP which requires as an assumption that the optimal value be attained [84]. This leads them to make the following remark in their conclusion section:

“This assumption is non-trivial, and we do not address the (important and difficult) question of how to verify that a given polynomial $f(x)$ has this property.”

Prior literature on existence of optimal solutions to POPs has focused on identifying cases where existence is always guaranteed. The best-known result here is the case of linear programming (i.e., when the degrees of p and q_i are one). In this case, the optimal value of the problem is always attained. This result was extended by Frank and Wolfe to the case where p is quadratic and the polynomials q_i are linear [38]. Consequently, results concerning attainment of the optimal value are sometimes referred to as “Frank-Wolfe type” theorems in the literature [17, 68]. Andronov et al. showed that the same statement holds again when p is cubic (and the polynomials q_i are linear) [10].

Our results in this chapter show that in all other cases, it is strongly NP-hard to determine whether a polynomial optimization problem attains its optimal value. This implies that unless $P=NP$, there is no polynomial-time (or even pseudo-polynomial time) algorithm for checking this property. Nevertheless, it follows from the Tarski-Seidenberg quantifier elimination theory [100, 105] that this problem is decidable, i.e., can be solved in finite time. There are also probabilistic algorithms that test for attainment of the optimal value of a POP [43, 44], but their complexities are exponential in the number of variables.

In this chapter, we also study the complexity of testing several well-known sufficient conditions for attainment of the optimal value (see Section 4.1.1 below). One sufficient condition that we do not consider but that is worth noting is for the polynomials $p, -q_1, \dots, -q_m$ to

¹To remove possible confusion, we emphasize that our focus in this chapter is not on the complexity of testing feasibility or unboundedness of problem (4.1), which have already been studied extensively. On the contrary, all optimization problems that we consider are by construction feasible and bounded below.

all be convex (see [17] for a proof, [68] for the special case where p and q_i are quadratics, and [20] for other extensions). The reason we exclude this sufficient condition from our study is that the complexity of checking convexity of polynomials has already been analyzed in [3].

4.1.1 Organization and Contributions of the Chapter

As mentioned before, this chapter concerns itself with the complexity of testing attainment of the optimal value of a polynomial optimization problem. More specifically, we show in Section 4.2 that it is strongly NP-hard to test attainment when the objective function has degree 4, even in absence of any constraints (Theorem 4.2.1), and when the constraints are of degree 2, even when the objective is linear (Theorem 4.2.2).

In Section 4.3, we show that several well-known sufficient conditions for attainment of the optimal value in a POP are also strongly NP-hard to test. These include coercivity of the objective function (Theorem 4.3.1), closedness of a bounded feasible set (Theorem 4.3.2 and Remark 4.3.1), boundedness of a closed feasible set (Corollary 4.3.3), a robust analogue of compactness known as stable compactness (Corollary 4.3.5), and an algebraic certificate of compactness known as the Archimedean property (Theorem 4.3.8). The latter property is of independent interest to the convergence of the Lasserre hierarchy, as discussed in Section 4.3.2.

In Section 4.4, we give semidefinite programming (SDP) based hierarchies for testing compactness of the feasible set and coercivity of the objective function of a POP (Propositions 4.4.1 and 4.4.2). The hierarchy for compactness comes from a straightforward application of Stengle's Positivstellensatz (cf. Theorem 4.3.7), but the one for coercivity requires us to develop a new characterization of coercive polynomials (Theorem 4.4.3). We end the chapter in Section 4.5 with a summary and some brief concluding remarks.

4.2 NP-hardness of Testing Attainment of the Optimal Value

In this section, we show that testing attainment of the optimal value of a polynomial optimization problem is NP-hard. Throughout this chapter, when we study complexity questions around problem (4.1), we fix the degrees of all polynomials involved and think of the number of variables and the coefficients of these polynomials as input. Since we are working in the Turing model of computation, all the coefficients are rational numbers and the input size can be taken to be the total number of bits needed to represent the numerators and denominators of these coefficients.

Our proofs of hardness are based on reductions from ONE-IN-THREE 3SAT which is known to be NP-hard [99]. Recall that in ONE-IN-THREE 3SAT, we are given a 3SAT instance (i.e., a collection of clauses, where each clause consists of exactly three literals, and each literal is either a variable or its negation) and we are asked to decide whether there exists a $\{0, 1\}$ assignment to the variables that makes the expression true with the additional property that each clause has *exactly* one true literal.

Theorem 4.2.1. *Testing whether a degree-4 polynomial attains its unconstrained infimum is strongly NP-hard.*

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables x_1, \dots, x_n , and k clauses. Let $s_\phi(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as

$$s_\phi(x) = \sum_{i=1}^k (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)^2 + \sum_{i=1}^n (1 - x_i^2)^2, \quad (4.2)$$

where $\phi_{it} = x_j$ if the t -th literal in the i -th clause is x_j , and $\phi_{it} = -x_j$ if it is $\neg x_j$ (i.e., the negation of x_j). Now, let

$$p_\phi(x, y, z, \lambda) := \lambda^2 s_\phi(x) + (1 - \lambda)^2 (y^2 + (yz - 1)^2), \quad (4.3)$$

where $y, z, \lambda \in \mathbb{R}$. We show that p_ϕ achieves its infimum if and only if ϕ is satisfiable. Note that the reduction is polynomial in length and the coefficients of p_ϕ are at most a constant multiple of $n + k$ in absolute value.

If ϕ has a satisfying assignment, then for any y and z , letting $\lambda = 1$, $x_i = 1$ if the variable is true in that assignment and $x_i = -1$ if it is false, results in a zero of p_ϕ . As p_ϕ is a sum of squares and hence nonnegative, we have shown that it achieves its infimum.

Now suppose that ϕ is not satisfiable. We will show that p_ϕ is positive everywhere but gets arbitrarily close to zero. To see the latter claim, simply set $\lambda = 0$, $z = \frac{1}{y}$, and let $y \rightarrow 0$. To see the former claim, suppose for the sake of contradiction that p_ϕ has a zero. Since $y^2 + (yz - 1)^2$ is always positive, we must have $\lambda = 1$ in order for the second term to be zero. Then, in order for the whole expression to be zero, we must also have that $s_\phi(x)$ must vanish at some x . But any zero of s_ϕ must have each $x \in \{-1, 1\}^n$, due to the second term of s_ϕ . However, because the instance ϕ is not satisfiable, for any such x , there exists $i \in \{1, \dots, k\}$ such that $\phi_{i1} + \phi_{i2} + \phi_{i3} + 1 \neq 0$, as there must be a clause where not exactly one literal is set to one. This means that s_ϕ is positive everywhere, which is a contradiction.

We have thus shown that testing attainment of the optimal value is NP-hard for unconstrained POPs where the objective is a polynomial of degree 6. In the interest of minimality, we now extend the proof to apply to an objective function of degree 4. To do this, we first introduce $n + 1$ new variables χ_1, \dots, χ_n and w . We replace every occurrence of the product λx_i in $\lambda^2 s_\phi$ with the variable χ_i . For example, the term $\lambda^2 x_1 x_2$ would become $\chi_1 \chi_2$. Let $\hat{s}_\phi(x, \chi, \lambda)$ denote this transformation on $\lambda^2 s_\phi(x)$. Note that $\hat{s}_\phi(x, \chi, \lambda)$ is now a quartic polynomial. Now consider the quartic polynomial (whose coefficients are again at most a constant multiple of $n + k$ in absolute value)

$$\hat{p}_\phi(x, y, z, \lambda, \chi, w) = \hat{s}_\phi(x, \chi, \lambda) + (1 - \lambda)^2 (y^2 + (w - 1)^2) + (w - yz)^2 + \sum_{i=1}^n (\chi_i - \lambda_i x_i)^2. \quad (4.4)$$

Observe that \hat{p}_ϕ is a sum of squares as \hat{s}_ϕ can be verified to be a sum of squares by bringing λ inside every squared term of s_ϕ . Hence, \hat{p}_ϕ is nonnegative. Furthermore, its infimum is still zero, as the choice of variables $\lambda = 0, w = 1, \chi = 0, x$ arbitrary, $z = \frac{1}{y}$, and letting $y \rightarrow 0$ will result in arbitrarily small values of \hat{p}_ϕ . Now it remains to show that this polynomial will have a zero if and only if p_ϕ in (4.3) has a zero. Observe that if (x, y, z, λ) is a zero of p_ϕ , then $(x, y, z, \lambda, \lambda x, yz)$ is a zero of \hat{p}_ϕ . Conversely, if $(x, y, z, \lambda, \chi, w)$ is a zero of \hat{p}_ϕ , then (x, y, z, λ) is a zero of p_ϕ . \square

Remark 4.2.1. Because we use the ideas behind this reduction repeatedly in the remainder of this chapter, we refer to the quartic polynomial defined in (4.2) as s_ϕ throughout. The same convention for ϕ_{it} relating the literals of ϕ to the variables x will be assumed as well.

We next show that testing attainment of the optimal value of a POP is NP-hard when the objective function is linear and the constraints are quadratic. Together with the previously-known Frank-Wolfe type theorems which we reviewed in the introduction, Theorems 4.2.1 and 4.2.2 characterize the complexity of testing attainment of the optimal value in polynomial optimization problems of any given degree. Indeed, our reductions can trivially be extended to the case where the constraints or the objective have higher degrees. For example to increase the degree of the constraints to some positive integer d , one can introduce a new variable γ along with the trivial constraint $\gamma^d = 0$. To increase the degree of the objective from four to a higher degree $2d$, one can again introduce a new variable γ and add the term γ^{2d} to the objective function.

Theorem 4.2.2. *Testing whether a degree-1 polynomial attains its infimum on a feasible set defined by degree-2 inequalities is strongly NP-hard.*

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses. Define the following POP, with $x, \chi \in \mathbb{R}^n$ and $\lambda, y, z, w, \gamma, \zeta, \psi \in \mathbb{R}$:

$$\min_{x, \chi, \lambda, y, z, w, \gamma, \zeta, \psi} \quad \gamma \quad (4.5)$$

$$\text{subject to} \quad \gamma \geq \lambda \sum_{i=1}^n \chi_i + (1 - \lambda)(\psi + \zeta) \quad (4.6)$$

$$1 - x_i^2 = 0, \quad \forall i \in \{1, \dots, n\}, \quad (4.7)$$

$$\chi_i = (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)^2, \quad \forall i \in \{1, \dots, k\}, \quad (4.8)$$

$$\psi = y^2, \quad (4.9)$$

$$yz = w, \quad (4.10)$$

$$\zeta = (w - 1)^2, \quad (4.11)$$

$$\lambda(1 - \lambda) = 0. \quad (4.12)$$

We show that the infimum of this POP is attained if and only if ϕ is satisfiable. Note first that the objective value is always nonnegative because of (4.6) and in view of (4.8), (4.9), (4.11), and (4.12). Observe that if ϕ has a satisfying assignment, then letting $x_i = 1$ if the variable is true in that assignment and $x_i = -1$ if it is false, along with $\lambda = 1$, y and z arbitrary, $\psi = y^2$, $w = yz$, and $\zeta = (w - 1)^2$, results in a feasible solution with an objective value of 0.

If ϕ is not satisfiable, the objective value can be made arbitrarily close to zero by taking an arbitrary $x \in \{-1, 1\}^n$, χ_i accordingly to satisfy (4.8), $\lambda = 0$, $\psi = y^2$, $z = \frac{1}{y}$, $w = 1$, $\zeta = 0$, and letting $y \rightarrow 0$. Suppose for the sake of contradiction that there exists a feasible solution to the POP with $\gamma = 0$. As argued before, because of the constraints (4.8), (4.9), (4.11), and (4.12), $\lambda \sum_{i=1}^n \chi_i + (1 - \lambda)(\psi + \zeta)$ is always nonnegative, and so for γ to be exactly zero, we need to have

$$\lambda \sum_{i=1}^n \chi_i + (1 - \lambda)(\psi + \zeta) = 0.$$

From (4.12), either $\lambda = 0$ or $\lambda = 1$. If $\lambda = 1$, then we must have $\chi_i = 0, \forall i = 1, \dots, n$, which is not possible as ϕ is not satisfiable. If $\lambda = 0$, then we must have $\psi + \zeta = y^2 + (yz - 1)^2 = 0$, which cannot happen as this would require $y = 0$ and $yz = 1$ concurrently. \square

4.3 NP-hardness of Testing Sufficient Conditions for Attainment

Arguably, the two best-known sufficient conditions under which problem (4.1) attains its optimal value are *compactness* of the feasible set and *coercivity* of the objective function. In this section, we show that both of these properties are NP-hard to test for POPs of low degree. We also prove that certain stronger conditions, namely the *Archimedean property* of the quadratic module associated with the constraints and *stable compactness* of the feasible set, are NP-hard to test.

4.3.1 Coercivity of the Objective Function

A function $p : \mathbb{R}^n \rightarrow \mathbb{R}$ is *coercive* if for every sequence $\{x_k\}$ such that $\|x_k\| \rightarrow \infty$, we have $p(x_k) \rightarrow \infty$. It is well known that a continuous coercive function achieves its infimum on a closed set (see, e.g., Appendix A.2 of [19]). This is because all sublevel sets of continuous coercive functions are compact.

Theorem 4.3.1. *Testing whether a degree-4 polynomial is coercive is strongly NP-hard.*

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses, and the associated quartic polynomial $s_\phi(x)$ as in (4.2). Let $s_{\phi h} : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ be the homogenization of this polynomial:

$$s_{\phi h}(x_0, x) := x_0^4 s_\phi\left(\frac{x}{x_0}\right) = \sum_{i=1}^k x_0^2 (\phi_{i1} + \phi_{i2} + \phi_{i3} + x_0)^2 + \sum_{i=1}^n (x_0^2 - x_i^2)^2. \quad (4.13)$$

By construction, $s_{\phi h}$ is a homogeneous polynomial of degree 4. We show that $s_{\phi h}$ is coercive if and only if ϕ is not satisfiable.

Suppose first that the instance ϕ has a satisfying assignment $\hat{x} \in \{-1, 1\}^n$. Then it is easy to see that $s_{\phi h}(1, \hat{x}) = 0$. As $s_{\phi h}$ is homogeneous, $s_{\phi h}(\alpha, \alpha\hat{x}) = 0$ for all α , showing that $s_{\phi h}$ is not coercive.

Now suppose that ϕ is not satisfiable. We show that $s_{\phi h}$ is positive definite (i.e., $s_{\phi h}(x_0, x) > 0$ for all $(x_0, x) \neq (0, 0)$). This would then imply that $s_{\phi h}$ is coercive as

$$\begin{aligned} s_{\phi h}(x_0, x) &= \|(x_0, x)^T\|^4 \cdot s_{\phi h}\left(\frac{(x_0, x)}{\|(x_0, x)^T\|}\right) \\ &\geq \mu \|(x_0, x)^T\|^4, \end{aligned}$$

where $\mu > 0$ is defined as the minimum of $s_{\phi h}$ on the unit sphere:

$$\mu = \min_{(x_0, x) \in S^n} s_{\phi h}(x_0, x).$$

Suppose that $s_{\phi h}$ was not positive definite. Then there exists a point $(\hat{x}_0, \hat{x}) \neq (0, 0)$ such that $s_{\phi h}(\hat{x}_0, \hat{x}) = 0$. First observe \hat{x}_0 cannot be zero due to the $(x_0 - x_i)^2$ terms in (4.13). As $\hat{x}_0 \neq 0$, then, by homogeneity, the point $(1, \frac{\hat{x}}{\hat{x}_0})$ is a zero of $s_{\phi h}$ as well. This however implies that $s_{\phi}(\hat{x}) = 0$, which we have previously argued (cf. the proof of Theorem 4.2.1) is equivalent to satisfiability of ϕ , hence a contradiction. \square

We remark that the above hardness result is minimal in the degree as odd-degree polynomials are never coercive and a quadratic polynomial $x^T Q x + b^T x + c$ is coercive if and only if the matrix Q is positive definite, a property that can be checked in polynomial time (e.g., by checking positivity of the leading principal minors of Q).

4.3.2 Closedness and Boundedness of the Feasible Set

The well-known Bolzano-Weierstrass extreme value theorem states that the infimum of a continuous function on a compact (i.e., closed and bounded) set is attained. In this section, we show that testing closedness or boundedness of a basic semialgebraic set defined by degree-2 inequalities is NP-hard. Once again, these hardness results are minimal in degree since these properties can be tested in polynomial time for sets defined by affine inequalities, as we describe next.

To check boundedness of a set $P := \{x \in \mathbb{R}^n \mid a_i^T x \geq b_i, i = 1, \dots, m\}$ defined by affine inequalities, one can first check that P is nonempty, and if it is, for each i minimize and maximize x_i over P . Note that P is unbounded if and only if at least one of these $2n$ linear programs is unbounded, which can be certified e.g. by detecting infeasibility of the corresponding dual problem. Thus, boundedness of P can be tested by solving $2n + 1$ linear programming feasibility problems, which can be done in polynomial time.

To check closedness of a set $P := \{x \in \mathbb{R}^n \mid a_i^T x \geq b_i, i = 1, \dots, m, c_j^T x > d_j, j = 1, \dots, r\}$, one can for each j minimize $c_j^T x$ over $\{x \in \mathbb{R}^n \mid a_i^T x \geq b_i, i = 1, \dots, m\}$ and declare that P is closed if and only if all of the respective optimal values are greater than d_j . Thus, closedness of P can be tested by solving r linear programs, which can be done in polynomial time.

Theorem 4.3.2. *Given a set of quadratic polynomials $q_i, i = 1, \dots, m, t_j, j = 1, \dots, r$, it is strongly NP-hard to test whether the basic semialgebraic set*

$$\{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m, t_j(x) > 0, j = 1, \dots, r\}$$

*is closed*².

²Note that m is not fixed in this statement or in Corollaries 4.3.3 and 4.3.5 below.

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses. Let ϕ_{ij} be as in the proof of Theorem 4.2.1 and consider the set

$$S_\phi = \{(x, y) \in \mathbb{R}^{n+1} \mid (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y = 0, i = 1, \dots, k, 1 - x_j^2 = 0, j = 1, \dots, n, y < 1\}. \quad (4.14)$$

We show that S_ϕ is closed if and only if the instance ϕ is not satisfiable. To see this, first note that we can rewrite S_ϕ as

$$S_\phi = \left\{ \{-1, 1\}^n \times \{0\} \right\} \cup \left\{ \{x \in \mathbb{R}^n \mid s_\phi(x) = 0\} \times \{y \in \mathbb{R} \mid y < 1\} \right\},$$

where s_ϕ is as in the proof of Theorem 4.2.1. If ϕ is not satisfiable, then $S_\phi = \{-1, 1\}^n \times \{0\}$, which is closed. If ϕ is satisfiable, then $\{x \in \mathbb{R}^n \mid s_\phi(x) = 0\}$ is nonempty and

$$\{x \in \mathbb{R}^n \mid s_\phi(x) = 0\} \times \{y \in \mathbb{R} \mid y < 1\}$$

is not closed and not a subset of $\{-1, 1\}^n \times \{0\}$. This implies that S_ϕ is not closed. \square

Remark 4.3.1. We note that the problem of testing closedness of a basic semialgebraic set remains NP-hard even if one has a promise that the set is bounded. Indeed, one can add the constraint $y \geq -1$ to the set S_ϕ in (4.14) to make it bounded and this does not change the previous proof.

Corollary 4.3.3. *Given a set of quadratic polynomials $q_i, i = 1, \dots, m$, it is strongly NP-hard to test whether the set*

$$\{x \in \mathbb{R}^n \mid q_i(x) = 0, i = 1, \dots, m\}$$

is bounded.

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses. Let ϕ_{ij} be as in the proof of Theorem 4.2.1 and consider the set

$$S = \{(x, y) \in \mathbb{R}^{n+1} \mid (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y = 0, i = 1, \dots, k, 1 - x_j^2 = 0, j = 1, \dots, n\}.$$

This set is bounded if and only if ϕ is not satisfiable. One can see this by following the proof of Theorem 4.3.2 and observing that y will be unbounded in the satisfiable case, and only 0 otherwise. \square

Note that it follows immediately from either of the results above that testing compactness of a basic semialgebraic set is NP-hard. We end this subsection by establishing the same hardness result for a sufficient condition for compactness that has featured in the literature on polynomial optimization (see, e.g., [69], [84, Section 7]).

Definition 4.3.4. *A closed basic semialgebraic set $S = \{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m\}$ is stably compact if there exists $\epsilon > 0$ such that the set $\{x \in \mathbb{R}^n \mid \delta_i(x) + q_i(x) \geq 0, i = 1, \dots, m\}$ is compact for any set of polynomials δ_i having degree at most that of q_i and coefficients at most ϵ in absolute value.*

Intuitively, a closed basic semialgebraic set is stably compact if it remains compact under small perturbations of the coefficients of its defining polynomials. A stably compact set is clearly compact, though the converse is not true as shown by the set

$$S = \{(x_1, x_2) \in \mathbb{R}^2 \mid (x_1 - x_2)^4 + (x_1 + x_2)^2 \leq 1\}.$$

Indeed, this set is contained inside the unit disk, but for $\epsilon > 0$, the set

$$S_\epsilon = \{(x_1, x_2) \in \mathbb{R}^2 \mid (x_1 - x_2)^4 - \epsilon x_1^4 + (x_1 + x_2)^2 \leq 1\}$$

is unbounded as its defining polynomial tends to $-\infty$ along the line $x_1 = x_2$.

Section 5 of [69] shows that the set S in Definition 4.3.4 is stably compact if and only if the function

$$q(x) = \max_{i,j} \{-q_{ij}(x)\}$$

is positive on the unit sphere. Here, $q_{ij}(x)$ is a homogenous polynomial that contains all terms of degree j in $q_i(x)$. Perhaps because of this characterization, the same section in [69] remarks that “stable compactness is easier to check than compactness”, though as far as polynomial-time checkability is concerned, we show that the situation is no better.

Corollary 4.3.5. *Given a set of quadratic polynomials $q_i, i = 1, \dots, m$, it is strongly NP-hard to test whether the set*

$$\{x \in \mathbb{R}^n \mid q_i(x) = 0, i = 1, \dots, m\}$$

is stably compact.

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses. Let ϕ_{ij} be as in the proof of Theorem 4.2.1 and consider the set

$$T_\phi = \{(x_0, x) \in \mathbb{R}^{n+1} \mid (\phi_{i1} + \phi_{i2} + \phi_{i3} + x_0)^2 = 0, i = 1, \dots, k, x_0^2 - x_j^2 = 0, j = 1, \dots, n\}.$$

We show that the function

$$q_\phi(x_0, x) = \max_{i=1, \dots, k, j=1, \dots, n} \{-(\phi_{i1} + \phi_{i2} + \phi_{i3} + x_0)^2, (\phi_{i1} + \phi_{i2} + \phi_{i3} + x_0)^2, x_0^2 - x_j^2, x_j^2 - x_0^2\}$$

is positive on the unit sphere if and only if the instance ϕ is not satisfiable. Suppose first that ϕ is not satisfiable and assume for the sake of contradiction that there is a point (x_0, x) on the sphere such that $q_\phi(x_0, x) = 0$. This implies that $\phi_{i1} + \phi_{i2} + \phi_{i3} + x_0 = 0, \forall i = 1, \dots, k$ and $x_0^2 = x_j^2, \forall j = 1, \dots, n$. Hence, $x_0 \neq 0$ and $\frac{x}{x_0}$ is a satisfying assignment to ϕ , which is a contradiction. Suppose now that ϕ has a satisfying assignment $\hat{x} \in \{-1, 1\}^n$. Then it is

easy to check that

$$q_\phi \left(\frac{(1, \hat{x})}{\|(1, \hat{x})^T\|} \right) = 0.$$

□

The Archimedean Property

An algebraic notion closely related to compactness is the so-called Archimedean property. This notion has frequently appeared in recent literature at the interface of algebraic geometry and polynomial optimization. The Archimedean property is the assumption needed for the statement of Putinar's Positivstellensatz [92] and convergence of the Lasserre hierarchy [61]. In this subsection, we recall the definition of the Archimedean property and study the complexity of checking it. To our knowledge, the only previous result in this direction is that testing the Archimedean property is decidable [111, Section 3.3].

Recall that the *quadratic module* associated with a set of polynomials q_1, \dots, q_m is the set of polynomials that can be written as

$$\sigma_0(x) + \sum_{i=1}^m \sigma_i(x)q_i(x),$$

where $\sigma_0, \dots, \sigma_m$ are sum of squares polynomials.

Definition 4.3.6. *A quadratic module Q is Archimedean if there exists a scalar $R > 0$ such that $R - \sum_{i=1}^n x_i^2 \in Q$.*

Several equivalent characterizations of this property can be found in [63, Theorem 3.17]. Note that a set $\{x \in \mathbb{R}^n \mid q_i(x) \geq 0\}$ for which the quadratic module associated with the polynomials $\{q_i\}$ is Archimedean is compact. However, the converse is not true. For example, for $n > 1$, the sets

$$\left\{ x \in \mathbb{R}^n \mid x_1 - \frac{1}{2} \geq 0, \dots, x_n - \frac{1}{2} \geq 0, 1 - \prod_{i=1}^n x_i \geq 0 \right\}$$

are compact but not Archimedean; see [63], [91] for a proof of the latter claim. Hence, hardness of testing the Archimedean property does not follow from hardness of testing compactness.

As mentioned previously, the Archimedean property has received recent attention in the optimization community due to its connection to the Lasserre hierarchy. Indeed, under the assumption that the quadratic module associated with the defining polynomials of the feasible set of (4.1) is Archimedean, the Lasserre hierarchy [61] produces a sequence of SDP-based lower bounds that converge to the optimal value of the POP. Moreover, Nie has shown [82] that under the Archimedean assumption, convergence happens in a finite number of rounds generically. One way to ensure the Archimedean property—assuming that we know that our feasible set is contained in a ball of radius R —is to add the redundant constraint $R^2 \geq \sum_{i=1}^n x_i^2$ to the constraints of (4.1). This approach however increases the size of the SDP that needs to be solved at each level of the hierarchy. Moreover, such a scalar R may not be readily available for some applications.

Our proof of NP-hardness of testing the Archimedean property will be based on showing that the specific sets that arise from the proof of Corollary 4.3.3 are compact if and only if their corresponding quadratic modules are Archimedean. Our proof technique will use the Stengle’s Positivstellensatz, which we recall next.

Theorem 4.3.7 (Stengle’s Positivstellensatz [103]). *A basic semialgebraic set*

$$\mathcal{S} := \{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m, r_j(x) = 0, j = 1, \dots, k\}$$

is empty if and only if there exist sos polynomials σ_{c_1, \dots, c_m} and polynomials t_i such that

$$-1 = \sum_{j=1}^k t_j r_j + \sum_{c_1, \dots, c_m \in \{0,1\}^m} \sigma_{c_1, \dots, c_m}(x) \prod_{i=1}^m q_i(x)^{c_i}.$$

Remark 4.3.2. Note that if only equality constraints are considered, the second term on the right hand side is a single sos polynomial $\sigma_{0,\dots,0}$. In the next theorem, we only need this special case, which is also known as the Real Nullstellensatz [59].

Theorem 4.3.8. *Given a set of quadratic polynomials q_1, \dots, q_m , it is strongly NP-hard to test whether their quadratic module has the Archimedean property.*

Proof. Consider a ONE-IN-THREE 3SAT instance ϕ with n variables and k clauses. Let ϕ_{ij} be as in the proof of Theorem 4.2.1 and consider the set of quadratic polynomials

$$\{(\phi_{i1} + \phi_{i2} + \phi_{i3})y, -(\phi_{i1} + \phi_{i2} + \phi_{i3})y, i = 1, \dots, k; 1 - x_j^2, x_j^2 - 1, j = 1, \dots, n\}.$$

We show that the quadratic module associated with these polynomials is Archimedean if and only if ϕ is not satisfiable. First observe that if ϕ is satisfiable, then the quadratic module cannot be Archimedean as the set

$$S = \{(x, y) \in \mathbb{R}^{n+1} \mid (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y = 0, i = 1, \dots, k, 1 - x_j^2 = 0, j = 1, \dots, n\}$$

is not compact (see the proof of Corollary 4.3.3).

Now suppose that the instance ϕ is not satisfiable. We need to show that for some scalar $R > 0$ and some sos polynomials $\sigma_0, \sigma_1, \dots, \sigma_k, \hat{\sigma}_1, \dots, \hat{\sigma}_k, \tau_1, \dots, \tau_n, \hat{\tau}_1, \dots, \hat{\tau}_n$, we have

$$\begin{aligned} R - \sum_{i=1}^n x_i^2 - y^2 &= \sigma_0(x, y) + \sum_{i=1}^k \sigma_i(x, y)(\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y \\ &\quad + \sum_{i=1}^k \hat{\sigma}_i(x, y)(-\phi_{i1} - \phi_{i2} - \phi_{i3} - 1)y + \sum_{j=1}^n \tau_j(x, y)(1 - x_j^2) + \sum_{j=1}^n \hat{\tau}_j(x, y)(x_j^2 - 1). \end{aligned}$$

Since any polynomial can be written as the difference of two sos polynomials (see, e.g., [4, Lemma 1]), this is equivalent to existence of a scalar $R > 0$, an sos polynomial σ_0 , and some

polynomials $v_1, \dots, v_k, t_1, \dots, t_n$ such that

$$R - \sum_{i=1}^n x_i^2 - y^2 = \sigma_0(x, y) + \sum_{i=1}^k v_i(x, y)(\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y + \sum_{j=1}^n t_j(x, y)(1 - x_j^2). \quad (4.15)$$

First, note that

$$n - \sum_{i=1}^n x_i^2 = \sum_{i=1}^n (1 - x_i^2). \quad (4.16)$$

Secondly, as ϕ is not satisfiable, we know that the set

$$\{x \in \mathbb{R}^n \mid 1 - x_j^2 = 0, j = 1, \dots, n, \phi_{i1} + \phi_{i2} + \phi_{i3} + 1 = 0, i = 1, \dots, k\}$$

is empty. From Stengle's Positivstellensatz, it follows that there exist an sos polynomial $\tilde{\sigma}_0$ and some polynomials $\tilde{v}_1, \dots, \tilde{v}_k, \tilde{t}_1, \dots, \tilde{t}_n$ such that

$$-1 = \tilde{\sigma}_0(x) + \sum_{i=1}^k \tilde{v}_i(x)(\phi_{i1} + \phi_{i2} + \phi_{i3} + 1) + \sum_{j=1}^n \tilde{t}_j(x)(1 - x_j^2).$$

Multiplying this identity on either side by y^2 , we obtain:

$$-y^2 = y^2 \tilde{\sigma}_0(x) + \sum_{i=1}^k \tilde{v}_i(x)y \cdot (\phi_{i1} + \phi_{i2} + \phi_{i3} + 1)y + \sum_{j=1}^n \tilde{t}_j(x)y^2(1 - x_j^2). \quad (4.17)$$

Note that if we sum (4.16) and (4.17) and take $R = n$, $\sigma_0(x, y) = y^2 \tilde{\sigma}_0(x)$, $v_i(x, y) = y \tilde{v}_i(x)$ for all $i = 1, \dots, k$, and $t_j(x, y) = y^2 \cdot \tilde{t}_j(x) + 1$ for all $j = 1, \dots, n$, we recover (4.15). □

4.4 Algorithms for Testing Attainment of the Optimal Value

In this section, we give a hierarchy of sufficient conditions for compactness of a closed basic semialgebraic set, and a hierarchy of sufficient conditions for coercivity of a polynomial. These hierarchies are amenable to semidefinite programming (SDP) as they all involve, in one way or another, a search over the set of sum of squares polynomials. The connection between SDP and sos polynomials is well known: recall that a polynomial σ of degree $2d$ is sos if and only if there exists a symmetric positive semidefinite matrix Q such that $\sigma(x) = z(x)^T Q z(x)$ for all x , where $z(x)$ here is the standard vector of monomials of degree up to d in the variables x (see, e.g. [88]).

The hierarchies that we present are such that if the property in question (i.e., compactness or coercivity) is satisfied on an input instance, then some level of the SDP hierarchy will be feasible and provide a certificate that the property is satisfied. The test for compactness is a straightforward application of Stengle's Positivstellensatz, but the test for coercivity requires a new characterization of this property, which we give in Theorem 4.4.3.

4.4.1 Compactness of the Feasible Set

Consider a closed basic semialgebraic set

$$\mathcal{S} := \{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m\},$$

where the polynomials q_i have integer coefficients³ and are of degree at most d . A result of Basu and Roy [16, Theorem 3] implies that if \mathcal{S} is bounded, then it must be contained in a

³If some of the coefficients of the polynomials q_i are rational (but not integer), we can make them integers by clearing denominators without changing the set \mathcal{S} .

ball of radius

$$R^* := \sqrt{n} \left((2d+1)(2d)^{n-1} + 1 \right) 2^{(2d+1)(2d)^{n-1}(2nd+2)(2\tau + \text{bit}((2d+1)(2d)^{n-1}) + (n+1)\text{bit}(d+1) + \text{bit}(m))}, \quad (4.18)$$

where τ is the largest bitsize of any coefficient of any q_i , and $\text{bit}(\eta)$ denotes the bitsize of η .

With this result in mind, the following proposition is an immediate consequence of Stengle's Positivstellensatz (c.f. Theorem 4.3.7) after noting that the set \mathcal{S} is bounded if and only if the set

$$\{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m, \sum_{i=1}^n x_i^2 \geq R^* + 1\}$$

is empty.

Proposition 4.4.1. *Consider a closed basic semialgebraic set $\mathcal{S} := \{x \in \mathbb{R}^n \mid q_i(x) \geq 0, i = 1, \dots, m\}$, where the polynomials q_i have integer coefficients and are of degree at most d . Let R^* be as in (4.18) and let $q_0(x) = \sum_{i=1}^n x_i^2 - R^* - 1$. Then the set \mathcal{S} is compact if and only if there exist sos polynomials σ_{h_0, \dots, h_m} such that*

$$-1 = \sum_{h_0, \dots, h_m \in \{0,1\}^{m+1}} \sigma_{h_0, \dots, h_m}(x) \prod_{i=0}^m q_i(x)^{h_i}.$$

This proposition naturally yields the following semidefinite programming-based hierarchy indexed by a nonnegative integer r :

$$\begin{aligned} & \min_{\sigma_{h_0, \dots, h_m}} && 0 \\ \text{subject to} & && -1 = \sum_{h_0, \dots, h_m \in \{0,1\}^{m+1}} \sigma_{h_0, \dots, h_m}(x) \prod_{i=0}^m q_i(x)^{h_i}, \\ & && \sigma_{h_0, \dots, h_m} \text{ is sos and has degree } \leq 2r. \end{aligned} \quad (4.19)$$

Note that for a fixed level r , one is solving a semidefinite program whose size is polynomial in the description of \mathcal{S} . If for some r the SDP is feasible, then we have an algebraic certificate of compactness of the set \mathcal{S} . Conversely, as Proposition 4.4.1 implies, if \mathcal{S} is compact, then

the above SDP will be feasible for some level r^* . One can upper bound r^* by a function of n, m , and d only using the main theorem of [66]. This bound is however very large and mainly of theoretical interest.

4.4.2 Coercivity of the Objective Function

It is well known that the infimum of a continuous coercive function over a closed set is attained. This property has been widely studied, even in the case of polynomial functions; see e.g. [14, 13, 53]. A simple sufficient condition for coercivity of a polynomial p is for its terms of highest order to form a positive definite (homogeneous) polynomial; see, e.g., [53, Lemma 4.1]. One can give a hierarchy of SDPs to check for this condition as is done in [53, Section 4.2]. However, this condition is sufficient but not necessary for coercivity. For example, the polynomial $x_1^4 + x_2^2$ is coercive, but its top homogeneous component is not positive definite. Theorem 4.4.3 below gives a necessary and sufficient condition for a polynomial to be coercive which lends itself again to an SDP hierarchy. To start, we need the following proposition, whose proof is straightforward and thus omitted.

Proposition 4.4.2. *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is coercive if and only if the sets*

$$\mathcal{S}_\gamma := \{x \in \mathbb{R}^n \mid f(x) \leq \gamma\}$$

are bounded for all $\gamma \in \mathbb{R}$.

A polynomial p is said to be *s-coercive* if $p(x)/\|x\|^s$ is coercive. The *order of coercivity* of p is the supremum over $s \geq 0$ for which p is *s-coercive*. It is known that the order of coercivity of a coercive polynomial is always positive [13, 42].

Theorem 4.4.3. *A polynomial p is coercive if and only if there exist an even integer $c > 0$ and a scalar $k \geq 0$ such that for all $\gamma \in \mathbb{R}$, the γ -sublevel set of p is contained within a ball of radius $\gamma^c + k$.*

Proof. The “if” direction follows immediately from the fact that each γ -sublevel set is bounded. For the converse, suppose that p is coercive and denote its order of coercivity by $q > 0$. Then, from Observation 2 of [13], we get that there exists a scalar $M \geq 0$ such that

$$\|x\| \geq M \Rightarrow p(x) \geq \|x\|^q, \quad (4.20)$$

or equivalently $\|x\| \leq p(x)^{\frac{1}{q}}$. Now consider the function $R_p : \mathbb{R} \rightarrow \mathbb{R}$ which is defined as

$$R_p(\gamma) := \max_{p(x) \leq \gamma} \|x\|,$$

i.e. the radius of the γ -sublevel set of p . We note two relevant properties of this function.

- The function $R_p(\gamma)$ is nondecreasing. This is because the γ -sublevel set of p is a subset of the $(\gamma + \epsilon)$ -sublevel set of p for any $\epsilon > 0$.
- Let $m = \inf\{\gamma \mid R_p(\gamma) \geq M\}$. We claim that

$$m \geq M^q. \quad (4.21)$$

Suppose for the sake of contradiction that we had $m < M^q$. By the definition of m , there exists $\bar{\gamma} \in (m, M^q)$ such that $R_p(\bar{\gamma}) \geq M$. This means that there exists $\bar{x} \in \mathbb{R}^n$ such that $p(\bar{x}) \leq \bar{\gamma} < M^q$ and $\|\bar{x}\| \geq M$. From (4.20) we then have $p(\bar{x}) \geq \|\bar{x}\|^q \geq M^q$, which is a contradiction.

We now claim that $R_p(\gamma) \leq \gamma^{\frac{1}{q}}$ for all $\gamma > m$. Suppose for the sake of contradiction that there exists $\gamma_0 > m$ such that $R_p(\gamma_0) > \gamma_0^{\frac{1}{q}}$. This means that there exists $x_0 \in \mathbb{R}^n$ such that $\|x_0\| > \gamma_0^{\frac{1}{q}}$ but $p(x_0) \leq \gamma_0$.

Consider first the case where $p(x_0) \geq m$. Since $\gamma_0 > m$, we have

$$\|x_0\| > \gamma_0^{1/q} > m^{1/q} \geq M,$$

where the last inequality follows from (4.21). It follows from (4.20) that $p(x_0) \geq \|x_0\|^q > \gamma_0$ which is a contradiction.

Now consider the case where $p(x_0) < m$. By definition of m , we have $R_p(p(x_0)) < M$, and so $\|x_0\| < M$. Furthermore, since

$$\gamma_0 > m \stackrel{(4.21)}{\geq} M^q,$$

we have $M < \gamma_0^{1/q}$, which gives $\|x_0\| < M < \gamma_0^{1/q}$. This contradicts our previous assumption that $\|x_0\| > \gamma_0^{1/q}$.

If we let c be the smallest even integer greater than $1/q$, we have shown that

$$R_p(\gamma) \leq \gamma^{1/q} \leq \gamma^c$$

on the set $\gamma > m$. Finally, if we let $k = R_p(m)$, by monotonicity of R_p , we get that $R_p(\gamma) \leq \gamma^c + k$, for all γ .

□

Remark 4.4.1. One can easily show now that for any coercive polynomial p , there exist an integer $c' > 0$ and a scalar $k' \geq 0$ (possibly differing from the scalars c and k given in the proof of Theorem 4.4.3) such that

$$R_p^2(\gamma) < \gamma^{2c'} + k'.$$

For the following hierarchy it will be easier to work with this form.

In view of the above remark, observe that coercivity of a polynomial p is equivalent to existence of an integer $c' > 0$ and a scalar $k' \geq 0$ such that the set

$$\left\{ (\gamma, x) \in \mathbb{R}^{n+1} \mid p(x) \leq \gamma, \sum_{i=1}^n x_i^2 \geq \gamma^{2c'} + k' \right\} \quad (4.22)$$

is empty. This formulation naturally leads to the following SDP hierarchy indexed by a positive integer r .

Proposition 4.4.4. *A polynomial p of degree d is coercive if and only if for some integer $r \geq 1$, the following SDP is feasible:*

$$\begin{aligned}
& \min_{\sigma_0, \dots, \sigma_3} && 0 \\
\text{subject to} &&& -1 = \sigma_0(x, \gamma) + \sigma_1(x, \gamma)(\gamma - p(x)) + \sigma_2(x, \gamma) \left(\sum_{i=1}^n x_i^2 - \gamma^{2r} - 2^r \right) \\
&&& + \sigma_3(x, \gamma)(\gamma - p(x)) \left(\sum_{i=1}^n x_i^2 - \gamma^{2r} - 2^r \right), \\
&&& \sigma_0 \text{ is sos and of degree } \leq 4r, \\
&&& \sigma_1 \text{ is sos and of degree } \leq \max\{4r - d, 0\}, \\
&&& \sigma_2 \text{ is sos and of degree } \leq 2r, \\
&&& \sigma_3 \text{ is sos and of degree } \leq \max\{2r - d, 0\}.
\end{aligned} \tag{4.23}$$

Proof. If the SDP in (4.23) is feasible for some r , then the set

$$\left\{ (\gamma, x) \in \mathbb{R}^{n+1} \mid p(x) \leq \gamma, \sum_{i=1}^n x_i^2 \geq \gamma^{2r} + 2^r \right\} \tag{4.24}$$

must be empty. Indeed, if this was not the case, a feasible (γ, x) pair would make the right hand side of the equality constraint of (4.23) nonnegative, while the left hand side is negative. As the set in (4.24) is empty, then for all γ , the γ -sublevel set of p is contained within a ball of radius $\sqrt{\gamma^{2r} + 2^r}$ and thus p is coercive.

To show the converse, suppose that p is coercive. Then we know from Theorem 4.4.3 and Remark 4.4.1 that there exist an integer $c' > 0$ and a scalar $k' \geq 0$ such that the set in (4.22) is empty. From Stengle's Positivstellensatz (c.f. Theorem 4.3.7), there exist an even

nonnegative integer \hat{r} and sos polynomials $\sigma'_0, \dots, \sigma'_3$ of degree at most \hat{r} such that

$$\begin{aligned}
-1 = & \sigma'_0(x, \gamma) + \sigma'_1(x, \gamma)(\gamma - p(x)) + \sigma'_2(x, \gamma) \left(\sum_{i=1}^n x_i^2 - \gamma^{2c'} - k' \right) \\
& + \sigma'_3(x, \gamma)(\gamma - p(x)) \left(\sum_{i=1}^n x_i^2 - \gamma^{2c'} - k' \right).
\end{aligned} \tag{4.25}$$

Let $r^* = \lceil \max\{c', \log_2(k' + 1), \frac{\hat{r}+d}{2}\} \rceil$. We show that the SDP in (4.23) is feasible for $r = r^*$ by showing that the polynomials

$$\sigma_0(x, \gamma) = \sigma'_0(x, \gamma) + \sigma'_2(x, \gamma)(\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k'),$$

$$\sigma_1(x, \gamma) = \sigma'_1(x, \gamma) + \sigma'_3(x, \gamma)(\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k'),$$

$$\sigma_2(x, \gamma) = \sigma'_2(x, \gamma),$$

$$\sigma_3(x, \gamma) = \sigma'_3(x, \gamma)$$

are a feasible solution to the problem. First, note that

$$4r^* - d \geq 2r^* - d \geq \hat{r} \geq 0,$$

and so σ_0 is of degree at most $\hat{r} + 2r^* \leq 4r^*$, σ_1 is of degree at most $\hat{r} + 2r^* \leq \max\{4r^* - d, 0\}$, σ_2 is of degree at most $\hat{r} \leq 2r^*$, and σ_3 is of degree at most $\hat{r} \leq \max\{2r^* - d, 0\}$. Furthermore, these polynomials are sums of squares. To see this, note that $\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k'$ is nonnegative as $r^* \geq c'$ and $2^{r^*} \geq k' + 1$. As any nonnegative univariate polynomial is a sum of squares (see, e.g., [21]), it follows that $\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k'$ is a sum of squares. Combining this with the facts that $\sigma'_0, \dots, \sigma'_3$ are sums of squares, and products and sums of sos polynomials are

so again, we get that $\sigma_0, \dots, \sigma_3$ are sos. Finally, the identity

$$\begin{aligned}
-1 &= \left(\sigma'_0(x, \gamma) + \sigma'_2(x, y)(\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k') \right) \\
&\quad + \left(\sigma'_1(x, \gamma) + \sigma'_3(x, \gamma)(\gamma^{2r^*} - \gamma^{2c'} + 2^{r^*} - k') \right) (\gamma - p(x)) \\
&\quad + \sigma'_2(x, \gamma) \left(\sum_{i=1}^n x_i^2 - \gamma^{2r^*} - 2^{r^*} \right) + \sigma'_3(x, \gamma) (\gamma - p(x)) \left(\sum_{i=1}^n x_i^2 - \gamma^{2r^*} - 2^{r^*} \right)
\end{aligned}$$

holds by a simple rewriting of (4.25). □

As an illustration, we revisit the simple example $p(x) = x_1^4 + x_2^2$, whose top homogeneous component is not positive definite. The hierarchy in Proposition 4.4.4 with $r = 1$ gives an automated algebraic proof of coercivity of p in terms of the following identity:

$$-1 = \left(\frac{2}{3} \left(x_1^2 - \frac{1}{2} \right)^2 + \frac{2}{3} \left(\gamma - \frac{1}{2} \right)^2 \right) + \frac{2}{3} (\gamma - x_1^4 - x_2^2) + \frac{2}{3} (x_1^2 + x_2^2 - \gamma^2 - 2). \quad (4.26)$$

Note that this is a certificate that the γ -sublevel set of p is contained in a ball of radius $\sqrt{\gamma^2 + 2}$.

Remark 4.4.2. From a theoretical perspective, our developments so far show that coercivity of multivariate polynomials is a decidable property as it can be checked by solving a finite number of SDP feasibility problems (each of which can be done in finite time [90]). Indeed, given a polynomial p , one can think of running two programs in parallel. The first one solves the SDPs in Proposition 4.4.4 for increasing values of r . The second uses Proposition 4.4.1 and its degree bound to test whether the β -sublevel set p is compact, starting from $\beta = 1$, and doubling β in each iteration. On every input polynomial p whose coercivity is in question, either the first program halts with a yes answer or the second program halts with a no answer. We stress that this remark is of theoretical interest only, as the value of our contribution is really in providing proofs of coercivity, not proofs of non-coercivity. Moreover, coercivity

can alternatively be decided in finite time by applying the quantifier elimination theory of Tarski and Seidenberg [105, 100] to the characterization in Proposition 4.4.2.

4.5 Summary and Conclusions

We studied the complexity of checking existence of optimal solutions in mathematical programs (given as minimization problems) that are feasible and lower bounded. We showed that unless $P=NP$, this decision problem does not have a polynomial time (or even pseudo-polynomial time) algorithm when the constraints and the objective function are defined by polynomials of low degree. More precisely, this claim holds if the constraints are defined by quadratic polynomials (and the objective has degree as low as one) or if the objective function is a quartic polynomial (even in absence of any constraints). For polynomial optimization problems with linear constraints and objective function of degrees 1,2, or 3, previous results imply that feasibility and lower boundedness always guarantee existence of an optimal solution.

We also showed, again for low-degree polynomial optimization problems, that several well-known sufficient conditions for existence of optimal solutions are NP-hard to check. These were coercivity of the objective function, closedness of the feasible set (even when bounded), boundedness of the feasible set (even when closed), an algebraic certificate of compactness known as the Archimedean property, and a robust analogue of compactness known as stable compactness.

Our negative results should by no means deter researchers from studying algorithms that can efficiently check existence of optimal solutions—or, for that matter, any of the other properties mentioned above such as compactness and coercivity—on special instances. On the contrary, our results shed light on the intricacies that can arise when studying these properties and calibrate the expectations of an algorithm designer. Hopefully, they will even motivate further research in identifying problem structures (e.g., based on the Newton

polytope of the objective and/or constraints) for which checking these properties becomes more tractable, or efficient algorithms that can test useful sufficient conditions that imply these properties.

In the latter direction, we argued that sum of squares techniques could be a natural tool for certifying compactness of basic semialgebraic sets via semidefinite programming. By deriving a new characterization of coercive polynomials, we showed that the same statement also applies to the task of certifying coercivity. This final contribution motivates a problem that we leave for our future research. While coercivity (i.e., boundedness of all sublevel sets) of a polynomial objective function guarantees existence of optimal solutions to a feasible POP, the same guarantee can be made from the weaker requirement that *some* sublevel set of the objective be bounded and have a non-empty intersection with the feasible set. It is not difficult to show that this property is also NP-hard to check. However, it would be useful to derive a hierarchy of sufficient conditions for it, where each level can be efficiently tested (perhaps again via SDP), and such that if the property was satisfied, then a level of the hierarchy would hold.

Chapter 5

Semidefinite Programming

Relaxations for Nash Equilibria in Bimatrix Games

5.1 Introduction

A bimatrix game is a game between two players (referred to in this chapter as players A and B) defined by a pair of $m \times n$ payoff matrices A and B . Let Δ_m and Δ_n denote the m -dimensional and n -dimensional simplices

$$\Delta_m = \{x \in \mathbb{R}^m \mid x_i \geq 0, \forall i, \sum_{i=1}^m x_i = 1\}, \Delta_n = \{y \in \mathbb{R}^n \mid y_i \geq 0, \forall i, \sum_{i=1}^n y_i = 1\}.$$

These form the strategy spaces of player A and player B respectively. For a strategy pair $(x, y) \in \Delta_m \times \Delta_n$, the payoff received by player A (resp. player B) is $x^T A y$ (resp. $x^T B y$). In particular, if the players pick vertices i and j of their respective simplices (also called pure strategies), their payoffs will be $A_{i,j}$ and $B_{i,j}$. One of the prevailing solution concepts for bimatrix games is the notion of *Nash equilibrium*. At such an equilibrium, the players are playing mutual best responses, i.e., a payoff maximizing strategy against the opposing

player's strategy. In our notation, a Nash equilibrium for the game (A, B) is a pair of strategies $(x^*, y^*) \in \Delta_m \times \Delta_n$ such that

$$x^{*T} A y^* \geq x^T A y^*, \forall x \in \Delta_m,$$

and

$$x^{*T} B y^* \geq x^{*T} B y, \forall y \in \Delta_n.^1$$

Nash [77] proved that for any bimatrix game, such pairs of strategies exist (in fact his result more generally applies to games with a finite number of players and a finite number of pure strategies). While existence of these equilibria is guaranteed, finding them is believed to be a computationally intractable problem. More precisely, a result of [31] implies that computing Nash equilibria is PPAD-complete (see [31] for a definition) even when the number of players is 3. This result was later improved by [24] who showed the same hardness result for bimatrix games.

These results motivate the notion of an approximate Nash equilibrium, a solution concept in which players receive payoffs “close” to their best response payoffs. More precisely, a pair of strategies $(x^*, y^*) \in \Delta_m \times \Delta_n$ is an (*additive*) ϵ -Nash equilibrium for the game (A, B) if

$$x^{*T} A y^* \geq x^T A y^* - \epsilon, \forall x \in \Delta_m,$$

and

$$x^{*T} B y^* \geq x^{*T} B y - \epsilon, \forall y \in \Delta_n.^2$$

Note that when $\epsilon = 0$, (x^*, y^*) form an exact Nash equilibrium, and hence it is of interest to find ϵ -Nash equilibria with ϵ small. Unfortunately, approximation of Nash equilibria has also proved to be computationally difficult. [25] have shown that, unless $\text{PPAD} \subseteq \text{P}$, there cannot

¹In this chapter we assume that all entries of A and B are between 0 and 1, and argue at the beginning of Section 5.2 why this is without loss of generality for the purpose of computing Nash equilibria.

²There are also other important notions of approximate Nash equilibria, such as ϵ -well-supported Nash equilibria [37] and relative approximate Nash equilibria [30] which are not considered in this chapter.

be a fully polynomial-time approximation scheme for computing Nash equilibria in bimatrix games. There have, however, been a series of constant factor approximation algorithms for this problem ([33, 32, 58, 107]), with the current best producing a .3393 approximation via an algorithm by [107].

We remark that there are exponential-time algorithms for computing Nash equilibria, such as the Lemke-Howson algorithm ([64, 98]). There are also certain subclasses of the problem which can be solved in polynomial time, the most notable example being the case of zero-sum games (i.e. when $B = -A$). This problem was shown to be solvable via linear programming by [29], and later shown to be polynomially equivalent to linear programming by [1]. Aside from computation of Nash equilibria, there are a number of related decision questions which are of economic interest but unfortunately NP-hard. Examples include deciding whether a player's payoff exceeds a certain threshold in some Nash equilibrium, deciding whether a game has a unique Nash equilibrium, or testing whether there exists a Nash equilibrium where a particular set of strategies is not played ([40, 27]).

Our focus in this chapter is on understanding the power of semidefinite programming³ (SDP) for finding approximate Nash equilibria in bimatrix games or providing certificates for related decision questions. The goal is not to develop a competitive solver, but rather to analyze the algorithmic power of SDP when applied to basic problems around computation of Nash equilibria. Semidefinite programming relaxations have been analyzed in depth in areas such as combinatorial optimization ([41], [67]) and systems theory ([23]), but not to such an extent in game theory. To our knowledge, the appearance of SDP in the game theory literature includes the work of [102] for exchangeable equilibria in symmetric games, of [89] on zero-sum polynomial games, of [101] for zero-sum stochastic games, and of [60] for semialgebraic min-max problems in static and dynamic games.

³The unfamiliar reader is referred to [108] for the theory of SDPs and a description of polynomial-time algorithms for them based on interior point methods.

5.1.1 Organization and Contributions of the chapter

In Section 5.2, we formulate the problem of finding a Nash equilibrium in a bimatrix game as a nonconvex quadratically constrained quadratic program and pose a natural SDP relaxation for it. In Section 5.3, we show that our SDP is exact when the game is strictly competitive (see Definition 5.3.3). In Section 5.4, we design two continuous but nonconvex objective functions for our SDP whose global minima coincide with rank-1 solutions. We provide a heuristic based on iterative linearization for minimizing both objective functions. We show empirically that these approaches produce ϵ very close to zero (on average in the order of 10^{-3}). In Section 5.5, we establish a number of bounds on the quality of the approximate Nash equilibria that can be read off of feasible solutions to our SDP. In Theorems 5.5.5, 5.5.6, and 5.5.8, we show that when the SDP returns solutions which are “close” to rank-1, the resulting strategies have small ϵ . We then present an improved analysis in the rank-2 case which shows how one can recover a $\frac{5}{11}$ -Nash equilibrium from the SDP solution (Theorem 5.5.10). We further prove that for symmetric games (i.e., when $B = A^T$), a $\frac{1}{3}$ -Nash equilibrium can be recovered in the rank-2 case (Theorem 5.5.17). We do not currently know of a polynomial-time algorithm for finding rank-2 solutions to our SDP. If such an algorithm were found, it would, together with our analysis, improve the best known approximation bound for symmetric games. In Section 5.6, we show how our SDP formulation can be used to provide certificates for certain (NP-hard) questions of economic interest about Nash equilibria in symmetric games. These are the problems of testing whether the maximum welfare achievable under any symmetric Nash equilibrium exceeds some threshold, and whether a set of strategies is played in every symmetric Nash equilibrium. In Section 5.7, we show that the SDP analyzed in this chapter dominates the first level of the Lasserre hierarchy (Proposition 5.7.1). Some directions for future research are discussed in Section 5.8. The four appendices of the chapter add some numerical and technical details.

5.1.2 Notation

We establish some notation that will be used throughout the chapter. The symbol Δ_k denotes the k -dimensional simplex. For a matrix A , the notation A_i , refers to its i -th row, and $A_{.j}$ refers to its j -th column. The notation e_i refers to the elementary vector $(0, \dots, 0, 1, 0, \dots, 0)^T$ with the 1 being in position i , 0_m refers to the m -dimensional vector of zero's, 1_m refers to the m -dimensional vector of one's, and $J_{m \times n}$ refers to the $m \times n$ matrix of one's. The notation $A \succeq 0$ (resp $A \geq 0$) denotes that a matrix A is positive semidefinite (resp. elementwise nonnegative), $S^{k \times k}$ denotes the set of symmetric $k \times k$ matrices, and $\text{Tr}(A)$ denotes the trace of a matrix A , i.e., the sum of its diagonal elements. For two matrices A and B , $A \succeq B$ denotes that $A - B$ is positive semidefinite and $A \otimes B$ denotes their Kronecker product. Finally, for a vector v , $\text{diag}(v)$ denotes the diagonal matrix with v on its diagonal. For a square matrix M , $\text{diag}(M)$ denotes the vector containing its diagonal entries.

5.2 The Formulation of our SDP Relaxation

In this section we present an SDP relaxation for the problem of finding Nash equilibria in bimatrix games. This is done after a straightforward reformulation of the problem as a nonconvex quadratically constrained quadratic program. We also assume that all entries of the payoff matrices A and B are between 0 and 1. This can be done without loss of generality because Nash equilibria are invariant under certain affine transformations in the payoffs. In particular, the games (A, B) and $(cA + dJ_{m \times n}, eB + fJ_{m \times n})$ have the same Nash equilibria

for any scalars c, d, e , and f , with c and e positive. This is because

$$\begin{aligned}
& x^{*T}Ay \geq x^T Ay \\
& \Leftrightarrow c(x^{*T}Ay^*) + d \geq c(x^T Ay^*) + d \\
& \Leftrightarrow c(x^{*T}Ay^*) + d(x^{*T}J_{m \times n}y^*) \geq c(x^T Ay^*) + d(x^T J_{m \times n}y^*) \\
& \Leftrightarrow x^{*T}(cA + dJ_{m \times n})y^* \geq x^T(cA + dJ_{m \times n})y
\end{aligned}$$

Identical reasoning applies for player B.

5.2.1 Nash Equilibria as Solutions to Quadratic Programs

Recall the definition of a Nash equilibrium from Section 5.1. An equivalent characterization is that a strategy pair $(x^*, y^*) \in \Delta_m \times \Delta_n$ is a Nash equilibrium for the game (A, B) if and only if

$$\begin{aligned}
x^{*T}Ay^* & \geq e_i^T Ay^*, \forall i \in \{1, \dots, m\}, \\
x^{*T}By^* & \geq x^{*T}Be_i, \forall i \in \{1, \dots, n\}.
\end{aligned} \tag{5.1}$$

The equivalence can be seen by noting that because the payoff from playing any mixed strategy is a convex combination of payoffs from playing pure strategies, there is always a pure strategy best response to the other player's strategy.

We now treat the Nash problem as the following quadratic programming (QP) feasibility problem:

$$\begin{aligned}
& \min_{x \in \mathbb{R}^m, y \in \mathbb{R}^n} && 0 \\
& \text{subject to} && x^T A y \geq e_i^T A y, \forall i \in \{1, \dots, m\}, \\
& && x^T B y \geq x^T B e_j, \forall j \in \{1, \dots, n\}, \\
& && x_i \geq 0, \forall i \in \{1, \dots, m\}, \\
& && y_i \geq 0, \forall j \in \{1, \dots, n\}, \\
& && \sum_{i=1}^m x_i = 1, \\
& && \sum_{i=1}^n y_i = 1.
\end{aligned} \tag{5.2}$$

Similarly, a pair of strategies $x^* \in \Delta_m$ and $y^* \in \Delta_n$ form an ϵ -Nash equilibrium for the game (A, B) if and only if

$$\begin{aligned}
x^{*T} A y^* &\geq e_i^T A y^* - \epsilon, \forall i \in \{1, \dots, m\}, \\
x^{*T} B y^* &\geq x^{*T} B e_i - \epsilon, \forall i \in \{1, \dots, n\}.
\end{aligned}$$

Observe that any pair of simplex vectors (x, y) is an ϵ -Nash equilibrium for the game (A, B) for any ϵ that satisfies

$$\epsilon \geq \max\{\max_i e_i^T A y - x^T A y, \max_i x^T B e_i - x^T B y\}.$$

We use the following notation throughout the chapter:

- $\epsilon_A(x, y) := \max_i e_i^T A y - x^T A y,$
- $\epsilon_B(x, y) := \max_i x^T B e_i - x^T B y,$
- $\epsilon(x, y) := \max\{\epsilon_A(x, y), \epsilon_B(x, y)\},$

and the function parameters are later omitted if they are clear from the context.

5.2.2 SDP Relaxation

The QP formulation in (5.2) lends itself to a natural SDP relaxation. We define a matrix

$$\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix},$$

and an augmented matrix

$$\mathcal{M}' := \begin{bmatrix} X & P & x \\ Z & Y & y \\ x & y & 1 \end{bmatrix},$$

with $X \in S^{m \times m}$, $Z \in \mathbb{R}^{n \times m}$, $Y \in S^{n \times n}$, $x \in \mathbb{R}^m$, $y \in \mathbb{R}^n$ and $P = Z^T$.

The SDP relaxation can then be expressed as

$$\min_{\mathcal{M}' \in \mathbb{S}^{m+n+1, m+n+1}} 0 \quad (\text{SDP1})$$

$$\text{subject to} \quad \text{Tr}(AZ) \geq e_i^T A y, \forall i \in \{1, \dots, m\}, \quad (5.3)$$

$$\text{Tr}(BZ) \geq x^T B e_j, \forall j \in \{1, \dots, n\}, \quad (5.4)$$

$$\sum_{i=1}^m x_i = 1, \quad (5.5)$$

$$\sum_{i=1}^n y_i = 1, \quad (5.6)$$

$$\mathcal{M}' \succeq 0, \quad (5.7)$$

$$\mathcal{M}'_{m+n+1, m+n+1} = 1, \quad (5.8)$$

$$\mathcal{M}' \succeq 0. \quad (5.9)$$

We refer to the constraints (5.3) and (5.4) as the relaxed Nash constraints and the constraints (5.5) and (5.6) as the unity constraints. This SDP is motivated by the following observation.

Proposition 5.2.1. *Let \mathcal{M}' be any rank-1 feasible solution to SDP1. Then the vectors x and y from its last column constitute a Nash equilibrium for the game (A, B) .*

Proof. We know that x and y are in the simplex from the constraints (5.5), (5.6), and (5.7).

If the matrix \mathcal{M}' is rank-1, then it takes the form

$$\begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T. \quad (5.10)$$

Then, from the relaxed Nash constraints we have that

$$\begin{aligned} e_i^T A y &\leq \text{Tr}(AZ) = \text{Tr}(Ayx^T) = \text{Tr}(x^T A y) = x^T A y, \\ x^T A e_i &\leq \text{Tr}(BZ) = \text{Tr}(Byx^T) = \text{Tr}(x^T B y) = x^T B y. \end{aligned}$$

The claim now follows from the characterization given in (5.1). □

Remark 5.2.1. Because a Nash equilibrium always exists, there will always be a matrix of the form (5.10) which is feasible to SDP1. Thus we can disregard any concerns about SDP1 being feasible, even when we add valid inequalities to it in Section 5.2.3.

Remark 5.2.2. It is intuitive to note that the submatrix $P = Z^T$ of the matrix \mathcal{M}' corresponds to a probability distribution over the strategies, and that seeking a rank-1 solution to our SDP can be interpreted as making P a product distribution.

The following theorem shows that SDP1 is a weak relaxation and stresses the necessity of additional valid constraints.

Theorem 5.2.2. Consider a bimatrix game with payoff matrices bounded in $[0, 1]$. Then for

any two vectors $x \in \Delta_m$ and $y \in \Delta_n$, there exists a feasible solution \mathcal{M}' to SDP1 with $\begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$ as its last column.

Proof. Consider any $x, y, \gamma > 0$, and the matrix

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T + \begin{bmatrix} \gamma J_{m+n, m+n} & 0_{m+n} \\ 0_{m+n}^T & 0 \end{bmatrix}.$$

This matrix is the sum of two nonnegative psd matrices and is hence nonnegative and psd. By assumption x and y are in the simplex, and so constraints (5.5) – (5.9) of SDP1 are satisfied. To check that constraints (5.3) and (5.4) hold, note that since A and B are nonnegative, as long as the matrices A and B are not the zero matrices, the quantities $\text{Tr}(AZ)$ and $\text{Tr}(BZ)$ will become arbitrarily large as γ increases. Since $e_i^T A y$ and $x^T B e_i$ are bounded by 1 by assumption, we will have that constraints (5.3) and (5.4) hold for γ large enough. In the case where A or B is the zero matrix, the Nash constraints are trivially satisfied for the respective player. \square

5.2.3 Valid Inequalities

In this subsection, we introduce a number of valid inequalities to improve upon the SDP relaxation in SDP1. These inequalities are justified by being valid if the matrix returned by the SDP is rank-1. The terminology we introduce here to refer to these constraints is used throughout the chapter. Constraints (5.11) and (5.12) will be referred to as the *row inequalities*, and (5.13) and (5.14) will be referred to as the *correlated equilibrium inequalities*.

Proposition 5.2.3. Any rank-1 solution \mathcal{M}' to SDP1 must satisfy the following:

$$\sum_{j=1}^m X_{i,j} = \sum_{j=1}^n P_{i,j} = x_i, \forall i \in \{1, \dots, m\}, \quad (5.11)$$

$$\sum_{j=1}^n Y_{i,j} = \sum_{j=1}^m Z_{i,j} = y_i, \forall i \in \{1, \dots, n\}. \quad (5.12)$$

$$\sum_{j=1}^n A_{i,j} P_{i,j} \geq \sum_{j=1}^n A_{k,j} P_{i,j}, \forall i, k \in \{1, \dots, m\}, \quad (5.13)$$

$$\sum_{j=1}^m B_{j,i} P_{j,i} \geq \sum_{j=1}^m B_{j,k} P_{j,i}, \forall i, k \in \{1, \dots, n\}. \quad (5.14)$$

Proof. Recall from (5.10) that if \mathcal{M}' is rank-1, it is of the form

$$\begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T.$$

To show (5.11), observe that

$$\sum_{j=1}^m X_{i,j} = \sum_{j=1}^m x_i x_j = x_i, \forall i \in \{1, \dots, m\}.$$

An identical argument works for the remaining matrices $P, Z,$ and Y . To show (5.13) and (5.14), observe that a pair (x, y) is a Nash equilibrium if and only if

$$\forall i, x_i > 0 \Rightarrow e_i^T Ay = x^T Ay = \max_i e_i^T Ay,$$

$$\forall i, y_i > 0 \Rightarrow x^T Be_i = x^T By = \max_i x^T Be_i.$$

This is because the Nash conditions require that $x^T Ay$, a convex combination of $e_i^T Ay$, be at least $e_i^T Ay$ for all i . Indeed, if $x_i > 0$ but $e_i^T Ay < x^T Ay$, the convex combination must be less than $\max_i e_i^T Ay$.

For each i such that $x_i = 0$ or $y_i = 0$, inequalities (5.13) and (5.14) reduce to $0 \geq 0$, so we only need to consider strategies played with positive probability. Observe that if \mathcal{M}' is rank-1, then

$$\begin{aligned} \sum_{j=1}^n A_{i,j} P_{i,j} &= x_i \sum_{j=1}^n A_{i,j} y_j = x_i e_i^T A y \geq x_i e_k^T A y = \sum_{j=1}^n A_{k,j} P_{i,j}, \forall i, k \\ \sum_{j=1}^m B_{j,i} P_{j,i} &= y_i \sum_{j=1}^m B_{j,i} x_j = y_i x^T B e_i \geq y_i x^T B e_k = \sum_{j=1}^m B_{j,i} P_{j,k}, \forall i, k. \end{aligned}$$

□

Remark 5.2.3. There are two ways to interpret the inequalities in (5.13) and (5.14): the first is as a relaxation of the constraint $x_i(e_i^T A y - e_j^T A y) \geq 0, \forall i, j$, which must hold since any strategy played with positive probability must give the best response payoff. The other interpretation is to have the distribution over outcomes defined by P be a correlated equilibrium [11]. This can be imposed by a set of linear constraints on the entries of P as explained next.

Suppose the players have access to a public randomization device which prescribes a pure strategy to each of them (unknown to the other player). The distribution over the assignments can be given by a matrix P , where $P_{i,j}$ is the probability that strategy i is assigned to player A and strategy j is assigned to player B. This distribution is a correlated equilibrium if both players have no incentive to deviate from the strategy prescribed, that is, if the prescribed pure strategies a and b satisfy

$$\begin{aligned} \sum_{j=1}^n A_{i,j} \text{Prob}(b = j | a = i) &\geq \sum_{j=1}^n A_{k,j} \text{Prob}(b = j | a = i), \\ \sum_{i=1}^m B_{i,j} \text{Prob}(a = i | b = j) &\geq \sum_{i=1}^m B_{i,k} \text{Prob}(a = i | b = j). \end{aligned}$$

If we interpret the P submatrix in our SDP as the distribution over the assignments by the public device, then because of our row constraints, $\text{Prob}(b = j | a = i) = \frac{P_{i,j}}{x_i}$ whenever

$x_i \neq 0$ (otherwise the above inequalities are trivial). Similarly, $P(a = i|b = j) = \frac{P_{i,j}}{y_j}$ for nonzero y_j . Observe now that the above two inequalities imply (5.13) and (5.14). Finally, note that every Nash equilibrium generates a correlated equilibrium, since if P is a product distribution given by xy^T , then $Prob(b = j|a = i) = y_j$ and $P(a = i|b = j) = x_i$.

Implied Inequalities

In addition to those explicitly mentioned in the previous section, there are other natural valid inequalities which are omitted because they are implied by the ones we have already proposed. We give two examples of such inequalities in the next proposition. We refer to the constraints in (5.15) below as the *distribution constraints*. The constraints in (5.16) are the familiar McCormick inequalities [70] for box-constrained quadratic programming.

Proposition 5.2.4. Let $z := \begin{bmatrix} x \\ y \end{bmatrix}$. Any rank-1 solution \mathcal{M}' to SDP1 must satisfy the following:

$$\sum_{i=1}^m \sum_{j=1}^m X_{i,j} = \sum_{i=1}^n \sum_{j=1}^m Z_{i,j} = \sum_{i=1}^n \sum_{j=1}^n Y_{i,j} = 1. \quad (5.15)$$

$$\mathcal{M}_{i,j} \leq z_i, \forall i, j \in \{1, \dots, m+n\}, \quad (5.16)$$

$$\mathcal{M}_{i,j} + 1 \geq z_i + z_j, \forall i, j \in \{1, \dots, m+n\}.$$

Proof. The distribution constraints follow immediately from the row constraints (5.11) and (5.12), along with the unity constraints (5.5) and (5.6).

The first McCormick inequality is immediate as a consequence of (5.11) and (5.12), as all entries of \mathcal{M} are nonnegative. To see why the second inequality holds, consider whichever submatrix X, Y, P , or Z that contains $\mathcal{M}_{i,j}$. Suppose that this submatrix is, e.g., P . Then, since P is nonnegative,

$$0 \leq \sum_{k=1, k \neq i}^m \sum_{l=1, l \neq j}^n P_{k,l} \stackrel{(5.11)}{=} \sum_{k=1, k \neq i}^m (x_k - P_{k,j}) \stackrel{(5.12)}{=} (1 - x_i) - (y_j - P_{i,j}) = P_{i,j} + 1 - x_i - y_j.$$

The same argument holds for the other submatrices, and this concludes the proof. \square

5.2.4 Simplifying our SDP

We observe that the row constraints (5.11) and (5.12) along with the correlated equilibrium constraints (5.13) and (5.14) imply the relaxed Nash constraints (5.3) and (5.4). Indeed, if we fix an index $k \in \{1, \dots, m\}$, then

$$\text{Tr}(AZ) = \sum_{i=1}^m \sum_{j=1}^n A_{i,j} P_{i,j} \stackrel{(5.13)}{\geq} \sum_{i=1}^m \sum_{j=1}^n A_{k,j} P_{i,j} \geq \sum_{j=1}^n A_{k,j} \left(\sum_{i=1}^m P_{i,j} \right) \stackrel{(5.12), P=Z^T}{\geq} \sum_{j=1}^n A_{k,j} y_j = e_k^T A y.$$

The proof for player B proceeds identically. Then, after collecting the valid inequalities and removing the relaxed Nash constraints, we arrive at an SDP given by

$$\begin{aligned} & \min_{\mathcal{M}' \in S^{(m+n+1) \times (m+n+1)}} && 0 && \text{(SDP1')} \\ & \text{subject to} && (5.5) - (5.9), (5.11) - (5.14). \end{aligned}$$

We make the observation that the last row and column of \mathcal{M}' can be removed from this SDP, that is, there is a one-to-one correspondence between solutions to SDP1' and those to

the following SDP (where $\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix}$, with $P = Z^T$):

$$\min_{\mathcal{M} \in S^{(m+n) \times (m+n)}} 0 \quad (\text{SDP2})$$

$$\text{subject to } \mathcal{M} \succeq 0, \quad (5.17)$$

$$\mathcal{M} \geq 0, \quad (5.18)$$

$$\sum_{i=1}^n \sum_{j=1}^n P_{i,j} = 1, \quad (5.19)$$

$$\sum_{j=1}^m X_{i,j} = \sum_{j=1}^n P_{i,j}, \forall i \in \{1, \dots, m\}, \quad (5.20)$$

$$\sum_{j=1}^n Y_{i,j} = \sum_{j=1}^m Z_{i,j}, \forall i \in \{1, \dots, n\}, \quad (5.21)$$

$$\sum_{j=1}^n A_{i,j} P_{i,j} \geq \sum_{j=1}^n A_{k,j} P_{i,j}, \forall i, k \in \{1, \dots, m\}, \quad (5.22)$$

$$\sum_{j=1}^m B_{j,i} P_{j,i} \geq \sum_{j=1}^m B_{j,k} P_{j,i}, \forall i, k \in \{1, \dots, n\}. \quad (5.23)$$

Indeed, it is readily verified that the submatrix \mathcal{M} from any feasible solution \mathcal{M}' to SDP1' is feasible to SDP2. Conversely, let \mathcal{M} be any feasible matrix to SDP2. Consider an eigendecomposition $\mathcal{M} = \sum_{i=1}^k \lambda_i v_i v_i^T$ and let $\begin{bmatrix} x \\ y \end{bmatrix} := \mathcal{M}^{\frac{1}{2}}$. Then the matrix

$$\mathcal{M}' := \begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix} = \sum_{i=1}^k \lambda_i \begin{bmatrix} v_i \\ 1_{m+n}^T v_i / 2 \end{bmatrix} \begin{bmatrix} v_i \\ 1_{m+n}^T v_i / 2 \end{bmatrix}^T \quad (5.24)$$

is easily seen to be feasible to SDP1'.

Given any feasible solution \mathcal{M} to SDP2, observe that the submatrix P is a correlated equilibrium. We take our candidate approximate Nash equilibrium to be the pair $x = P1_n$

and $y = P^T \mathbf{1}_m$. If the correlated equilibrium P is rank-1, then the pair (x, y) so defined constitutes an exact Nash equilibrium. In Section 5.4, we will add certain objective functions to SDP2 with the interpretation of searching for low-rank correlated equilibria.

5.3 Exactness for Strictly Competitive Games

In this section, we show that SDP1 recovers a Nash equilibrium for any zero-sum game, and that SDP2 recovers a Nash equilibrium for any strictly competitive game (see Definition 5.3.3 below). Both these notions represent games where the two players are in direct competition, but strictly competitive games are more general, and for example, allow both players to have nonnegative payoff matrices. These classes of games are solvable in polynomial time via linear programming. Nonetheless, it is reassuring to know that our SDPs recover these important special cases.

Definition 5.3.1. *A zero-sum game is a game in which the payoff matrices satisfy $A = -B$.*

Theorem 5.3.2. *For a zero-sum game, the vectors x and y from the last column of any feasible solution \mathcal{M}' to SDP1 constitute a Nash equilibrium.*

Proof. Recall that the relaxed Nash constraints (5.3) and (5.4) read

$$\text{Tr}(AZ) \geq e_i^T A y, \forall i \in \{1, \dots, m\},$$

$$\text{Tr}(BZ) \geq x^T B e_j, \forall j \in \{1, \dots, n\}.$$

Since $B = -A$, the latter statement is equivalent to

$$\text{Tr}(AZ) \leq x^T A e_j, \forall j \in \{1, \dots, n\}.$$

In conjunction these imply

$$e_i^T Ay \leq \text{Tr}(AZ) \leq x^T Ae_j, \forall i \in \{1, \dots, m\}, j \in \{1, \dots, n\}. \quad (5.25)$$

We claim that any pair $x \in \Delta_m$ and $y \in \Delta_n$ which satisfies the above condition is a Nash equilibrium. To see that $x^T Ay \geq e_i^T Ay, \forall i \in \{1, \dots, m\}$, observe that $x^T Ay$ is a convex combination of $x^T Ae_j$, which are at least $e_i^T Ay$ by (5.25). To see that $x^T By \geq x^T Be_j \Leftrightarrow x^T Ay \leq x^T Ae_j, \forall j \in \{1, \dots, n\}$, observe that $x^T Ay$ is a convex combination of $e_i^T Ay$, which are at most $x^T Ae_j$ by (5.25).

□

Definition 5.3.3. A game (A, B) is strictly competitive if for all $x, x' \in \Delta_m, y, y' \in \Delta_n$, $x^T Ay - x'^T Ay'$ and $x'^T By' - x^T By$ have the same sign.

The interpretation of this definition is that if one player benefits from changing from one outcome to another, the other player must suffer. Adler, Daskalakis, and Papadimitriou show in [2] that the following much simpler characterization is equivalent.

Theorem 5.3.4 (Theorem 1 of [2]). A game is strictly competitive if and only if there exist scalars c, d, e , and f , with $c > 0, e > 0$, such that $cA + dJ_{m \times n} = -eB + fJ_{m \times n}$.

One can easily show that there exist strictly competitive games for which not all feasible solutions to SDP1 have Nash equilibria as their last columns (see Theorem 5.2.2). However, we show that this is the case for SDP2.

Theorem 5.3.5. For a strictly competitive game, the vectors $x := P1_n$ and $y := P^T 1_m$ from any feasible solution \mathcal{M} to SDP2 constitute a Nash equilibrium.

To prove Theorem 5.3.5 we need the following lemma, which shows that feasibility of a matrix \mathcal{M} in SDP2 is invariant under certain transformations of A and B .

Lemma 5.3.6. *Let c, d, e , and f be any set of scalars with $c > 0$ and $e > 0$. If a matrix \mathcal{M} is feasible to SDP2 with input payoff matrices A and B , then it is also feasible to SDP2 with input matrices $cA + dJ_{m \times n}$ and $eB + fJ_{m \times n}$.*

Proof. It suffices to check that constraints (5.22) and (5.23) of SDP2 still hold, as only the correlated equilibrium constraints use the matrices A and B . We only show that constraint (5.22) still holds because the argument for constraint (5.23) is identical.

Note from the definition of x that for each $i \in \{1, \dots, m\}$, $x_i = \sum_{j=1}^n (J_{m \times n})_{i,j} P_{i,j}$. To check that the correlated equilibrium constraints hold, observe that for scalars $c > 0, d$, and for all $i, k \in \{1, \dots, m\}$,

$$\begin{aligned} & \sum_{j=1}^n A_{i,j} P_{i,j} \geq \sum_{j=1}^n A_{k,j} P_{i,j} \\ \Leftrightarrow & c \sum_{j=1}^n A_{i,j} P_{i,j} + d \sum_{j=1}^n P_{i,j} \geq c \sum_{j=1}^n A_{k,j} P_{i,j} + d \sum_{j=1}^n P_{i,j} \\ \Leftrightarrow & c \sum_{j=1}^n A_{i,j} P_{i,j} + d \sum_{j=1}^n (J_{m \times n})_{i,j} P_{i,j} \geq c \sum_{j=1}^n A_{k,j} P_{i,j} + d \sum_{j=1}^n (J_{m \times n})_{k,j} P_{i,j} \\ \Leftrightarrow & \sum_{j=1}^n (cA_{i,j} + dJ_{m \times n})_{k,j} P_{i,j} \geq \sum_{j=1}^n (cA_{i,j} + dJ_{m \times n})_{k,j} P_{i,j}. \end{aligned}$$

□

Proof. Proof (of Theorem 5.3.5). Let A and B be the payoff matrices of the given strictly competitive game and let \mathcal{M} be a feasible solution to SDP2. Since the game is strictly competitive, we know from Theorem 5.3.4 that $cA + dJ_{m \times n} = -eB + fJ_{m \times n}$ for some scalars $c > 0, e > 0, d, f$. Consider a new game with input matrices $\tilde{A} = cA + dJ_{m \times n}$ and $\tilde{B} = eB - fJ_{m \times n}$. By Lemma 5.3.6, \mathcal{M} is still feasible to SDP2 with input matrices \tilde{A}

and \tilde{B} . By the arguments in Section 5.2.4, the matrix $\mathcal{M}' := \begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix}$ is feasible to SDP1', and hence also to SDP1. Now notice that since $\tilde{A} = -\tilde{B}$, Theorem 5.3.2 implies

that the vectors x and y in the last column form a Nash equilibrium to the game (\tilde{A}, \tilde{B}) . Finally recall from the arguments at the beginning of Section 5.2 that Nash equilibria are invariant to scaling and shifting of the payoff matrices, and hence (x, y) is a Nash equilibrium to the game (A, B) . \square

5.4 Algorithms for Lowering Rank

In this section, we present heuristics which aim to find low-rank solutions to SDP2 and present some empirical results. Recall that our SDP2 in Section 5.2.4 did not have an objective function. Hence, we can encourage low-rank solutions by choosing certain objective functions, in particular the trace of the matrix \mathcal{M} , which is a general heuristic for minimizing the rank of symmetric matrices [96, 36]. This simple objective function is already guaranteed to produce a rank-1 solution in the case of strictly competitive games (see Proposition 5.4.1 below). For general games, however, one can design better objective functions in an iterative fashion (see Section 5.4.1).

Notational Remark: For the remainder of this section, we will use the shorthand $x := P1_n$ and $y := P^T 1_m$, where P is the upper right submatrix of a feasible solution \mathcal{M} to SDP2.

Proposition 5.4.1. *For a strictly competitive game, any optimal solution to SDP2 with $\text{Tr}(\mathcal{M})$ as the objective function must be rank-1.*

Proof. Let

$$\mathcal{M} := \begin{bmatrix} X & P \\ P^T & Y \end{bmatrix}$$

be a feasible solution to SDP2. In the case of strictly competitive games, from Theorem 5.3.5 we know that that (x, y) is a Nash equilibrium. Then because the matrix \mathcal{M} is psd, from (5.24) and an application of the Schur complement (see, e.g. [22, Sect. A.5.5])

to $\begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix}$, we have that $\mathcal{M} \succeq \begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$. Hence, $\mathcal{M} = \begin{bmatrix} xx^T & xy^T \\ yx^T & yy^T \end{bmatrix} + \mathcal{P}$ for some psd matrix \mathcal{P} and the Nash equilibrium (x, y) . Given this expression, the objective function $\text{Tr}(\mathcal{M})$ is then $x^T x + y^T y + \text{Tr}(\mathcal{P})$. As (x, y) is a Nash equilibrium, the choice of $\mathcal{P} = 0$ results in a feasible solution. Since the zero matrix has the minimum possible trace among all psd matrices, the solution will be the rank-1 matrix $\begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$. \square

Remark 5.4.1. If the row constraints and the nonnegativity constraints on X and Y are removed from SDP2, then this SDP with $\text{Tr}(\mathcal{M})$ as the objective function can be interpreted as searching for a minimum-rank correlated equilibrium P via the nuclear norm relaxation; see [96, Section 2].

5.4.1 Linearization Algorithms

The algorithms we present in this section for minimizing the rank of the matrix \mathcal{M} in SDP2 are based on iterative linearization of certain nonconvex objective functions. Motivated by the next proposition, we design two continuous (nonconvex) objective functions that, if minimized exactly, would guarantee rank-1 solutions. We will then linearize these functions iteratively.

Proposition 5.4.2. *Let the matrices X and Y and vectors $x := P1_n$ and $y := P^T 1_m$ be taken from a feasible solution to SDP2. Then the matrix \mathcal{M} is rank-1 if and only if $X_{i,i} = x_i^2$ and $Y_{i,i} = y_i^2$ for all i .*

Proof. Note that if \mathcal{M} is rank-1, then it can be written as zz^T for some $z \in \mathbb{R}^{m+n}$. The i -th diagonal entry in the X submatrix will then be equal to

$$z_i^2 \stackrel{(5.15)}{=} \frac{1}{4} z_i^2 (1_{m+n}^T z z^T 1_{m+n}) = \left(\frac{1}{2} \mathcal{M}_{i,1_{m+n}}\right)^2 \stackrel{(5.11)}{=} (P_{i,1_n})^2 = x_i^2,$$

where the second equality holds because \mathcal{M}_i ,—the i -th row of \mathcal{M} —is $z_i z^T$. An analogous statement holds for the diagonal entries of Y , and hence the condition is necessary.

To show sufficiency, let $z := \begin{bmatrix} x \\ y \end{bmatrix}$. Since \mathcal{M} is psd, we have that $\mathcal{M}_{i,j} \leq \sqrt{\mathcal{M}_{i,i}\mathcal{M}_{j,j}}$, which implies $\mathcal{M}_{i,j} \leq z_i z_j$ by the assumption of the proposition. Recall from the distribution constraint (5.15) that $\sum_{i=1}^{m+n} \sum_{j=1}^{m+n} \mathcal{M}_{i,j} = 4$. Further, the same constraint along with the definitions of x and y imply that $\sum_{i=1}^{m+n} z_i = 2$, which means that $\sum_{i=1}^{m+n} \sum_{j=1}^{m+n} z_i z_j = 4$. Hence in order to have the equality

$$4 = \sum_{i=1}^{m+n} \sum_{j=1}^{m+n} \mathcal{M}_{i,j} \leq \sum_{i=1}^{m+n} \sum_{j=1}^{m+n} z_i z_j = 4,$$

we must have $\mathcal{M}_{i,j} = z_i z_j$ for each i and j . Consequently \mathcal{M} is rank-1. \square

We focus now on two nonconvex objectives that as a consequence of the above proposition would return rank-1 solutions:

Proposition 5.4.3. *All optimal solutions to SDP2 with the objective function $\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}}$ or $\text{Tr}(\mathcal{M}) - x^T x - y^T y$ are rank-1.*

Proof. We show that each of these objectives has a specific lower bound which is achieved if and only if the matrix is rank-1.

Observe that since $\mathcal{M} \succeq \begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$, we have $\sqrt{X_{i,i}} \geq x_i$ and $\sqrt{Y_{i,i}} \geq y_i$, and hence

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}} \geq \sum_{i=1}^m x_i + \sum_{i=1}^n y_i = 2.$$

Further note that

$$\text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} \geq \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} = 0.$$

We can see that the lower bounds are achieved if and only if $X_{i,i} = x_i^2$ and $Y_{i,i} = y_i^2$ for all i , which by Proposition 5.4.2 happens if and only if \mathcal{M} is rank-1. \square

We refer to our two objective functions in Proposition 5.4.3 as the “*square root objective*” and the “*diagonal gap objective*” respectively. While these are both nonconvex, we will attempt to iteratively minimize them by linearizing them through a first order Taylor expansion. For example, at iteration k of the algorithm,

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}} \simeq \sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k-1)}} + \frac{1}{2\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} (\mathcal{M}_{i,i}^{(k)} - \mathcal{M}_{i,i}^{(k-1)}).$$

Note that for the purposes of minimization, this reduces to minimizing $\sum_{i=1}^{m+n} \frac{1}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} \mathcal{M}_{i,i}^{(k)}$.

In similar fashion, for the second objective function, at iteration k we can make the approximation

$$\text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} \simeq \text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \left(\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \right).$$

Once again, for the purposes of minimization this reduces to minimizing $\text{Tr}(\mathcal{M}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)}$. This approach then leads to the following two algorithms.⁴

⁴An algorithm similar to Algorithm 4 is used in [52].

Algorithm 3 Square Root Minimization Algorithm

- 1: Let $x^{(0)} = 1_m, y^{(0)} = 1_n, k = 1$.
 - 2: **while** !convergence **do**
 - 3: Solve SDP2 with $\sum_{i=1}^m \frac{1}{\sqrt{x_i^{(k-1)}}} X_{i,i} + \sum_{i=1}^n \frac{1}{\sqrt{y_i^{(k-1)}}} Y_{i,i}$ as the objective, and let \mathcal{M}^* be an optimal solution.
 - 4: Let $x^{(k)} = \text{diag}(X^*), y^{(k)} = \text{diag}(Y^*)$.
 - 5: Let $k = k + 1$.
 - 6: **end while**
-

Algorithm 4 Diagonal Gap Minimization Algorithm

- 1: Let $x^{(0)} = 0_m, y^{(0)} = 0_n, k = 1$.
 - 2: **while** !convergence **do**
 - 3: Solve SDP2 with $\text{Tr}(X) + \text{Tr}(Y) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)}$ as the objective, and let \mathcal{M}^* be an optimal solution.
 - 4: Let $x^{(k)} = P^* 1_n, y^{(k)} = P^{*T} 1_m$.
 - 5: Let $k = k + 1$.
 - 6: **end while**
-

Remark 5.4.2. Note that the first iteration of both algorithms uses the nuclear norm (i.e. trace) of \mathcal{M} as the objective.

The square root algorithm has the following property.

Theorem 5.4.4. *Let $\mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \dots$ be the sequence of optimal matrices obtained from the square root algorithm. Then the sequence*

$$\left\{ \sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}} \right\} \quad (5.26)$$

is nonincreasing and is lower bounded by two. If it reaches two at some iteration t , then the matrix $\mathcal{M}^{(t)}$ is rank-1.

Proof. Observe that for any $k > 1$,

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}} \leq \frac{1}{2} \sum_{i=1}^{m+n} \left(\frac{\mathcal{M}_{i,i}^{(k)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} + \sqrt{\mathcal{M}_{i,i}^{(k-1)}} \right) \leq \frac{1}{2} \sum_{i=1}^{m+n} \left(\frac{\mathcal{M}_{i,i}^{(k-1)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} + \sqrt{\mathcal{M}_{i,i}^{(k-1)}} \right) = \sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k-1)}},$$

where the first inequality follows from the arithmetic-mean-geometric-mean inequality, and the second follows from that $\mathcal{M}_{i,i}^{(k)}$ is chosen to minimize $\sum_{i=1}^{m+n} \frac{\mathcal{M}_{i,i}^{(k)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}}$ and hence achieves a no larger value than the feasible solution $\mathcal{M}^{(k-1)}$. This shows that the sequence is nonincreasing.

The proof of Proposition 5.4.3 already shows that the sequence is lower bounded by two, and Proposition 5.4.3 itself shows that reaching two is sufficient to have the matrix be rank-1. \square

The diagonal gap algorithm has the following property.

Theorem 5.4.5. *Let $\mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \dots$ be the sequence of optimal matrices obtained from the diagonal gap algorithm. Then the sequence*

$$\left\{ \text{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} \right\} \quad (5.27)$$

is nonincreasing and is lower bounded by zero. If it reaches zero at some iteration t , then the matrix $\mathcal{M}^{(t)}$ is rank-1.

Proof. Observe that

$$\begin{aligned} & \text{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} \\ & \leq \text{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} + \left(\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \right)^T \left(\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \right) \end{aligned}$$

$$\begin{aligned}
&= \text{Tr}(\mathcal{M}^{(k)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} + \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \\
&\leq \text{Tr}(\mathcal{M}^{(k-1)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} + \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \\
&= \text{Tr}(\mathcal{M}^{(k-1)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)},
\end{aligned}$$

where the second inequality follows from that $\mathcal{M}^{(k)}$ is chosen to minimize

$$\text{Tr}(\mathcal{M}^{(k-1)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)}$$

and hence achieves a no larger value than the feasible solution $\mathcal{M}^{(k-1)}$. This shows that the sequence is nonincreasing.

The proof of Proposition 5.4.3 already shows that the sequence is lower bounded by zero, and Proposition 5.4.3 itself shows that reaching zero is sufficient to have the matrix be rank-1. \square

We also invite the reader to also see Theorem 5.5.6 in the next section which relates the objective value of the diagonal gap minimization algorithm and the quality of approximate Nash equilibria that the algorithm produces.

5.4.2 Numerical Experiments

We tested Algorithms 3 and 4 on games coming from 100 randomly generated payoff matrices with entries bounded in $[0, 1]$ of varying sizes. Below is a table of statistics for 20×20

matrices; the data for the rest of the sizes can be found in Appendix A.1.⁵ We can see that our algorithms return approximate Nash equilibria with fairly low ϵ (recall the definition from Section 5.2.1). We ran 20 iterations of each algorithm on each game. Using the SDP solver of [73], each iteration takes on average under 4 seconds to solve on a standard personal machine with a 3.4 GHz processor and 16 GB of memory.

Table 5.1: Statistics on ϵ for 20×20 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0198	0.0046	0.0039	0.0034
Diagonal Gap	0.0159	0.0032	0.0024	0.0032

The histograms below show the effect of increasing the number of iterations on lowering ϵ on 20×20 games. For both algorithms, there was a clear improvement of the ϵ by increasing the number of iterations.

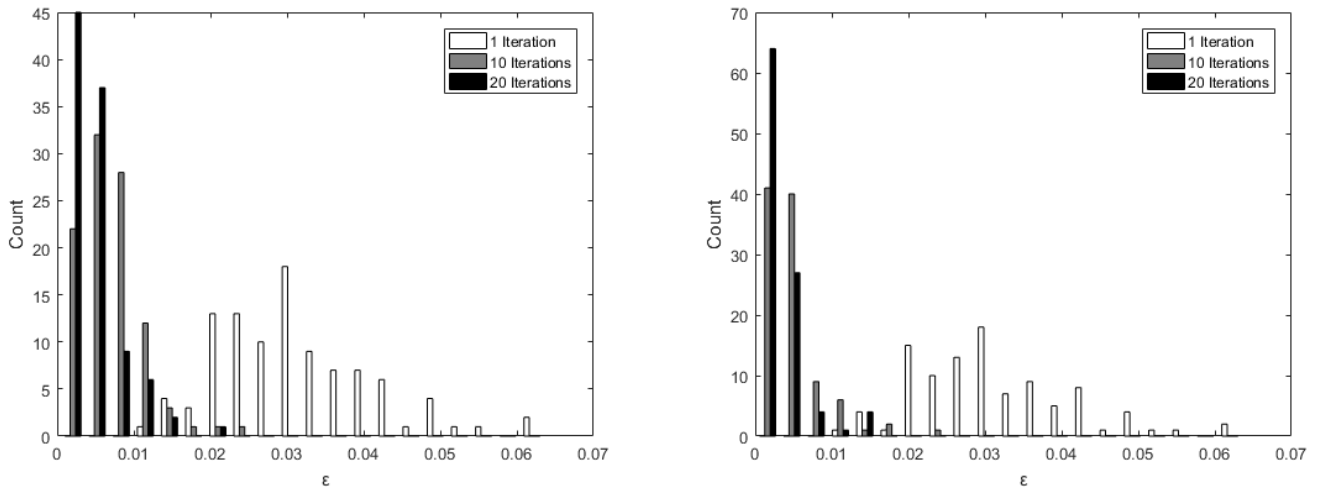


Figure 5.1: Distribution of ϵ over numbers of iterations for the square root algorithm (left) and the diagonal gap algorithm (right).

⁵The code and instance data that produced these results is publicly available at https://github.com/jeffreyzhang92/SDP_Nash. The function `nash.m` computes an approximate Nash equilibrium using one of our two algorithms as specified by the user.

5.5 Bounds on ϵ for General Games

Since the problem of computing a Nash equilibrium to an arbitrary bimatrix game is PPAD-complete, it is unlikely that one can find rank-1 solutions to this SDP in polynomial time. In Section 5.4, we designed objective functions (such as variations of the nuclear norm) that empirically do very well in finding low-rank solutions to SDP2. Nevertheless, it is of interest to know if the solution returned by SDP2 is not rank-1, whether one can recover an ϵ -Nash equilibrium from it and have a guarantee on ϵ . Our goal in this section is to study this question.

Notational Remark: Recall our notation for the matrix

$$\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix}.$$

Throughout this section, any matrices $X, Z, P = Z^T$ and Y are assumed to be taken from a feasible solution to SDP2. Furthermore, x and y will be $P1_n$ and P^T1_m respectively.

The ultimate results of this section are the theorems in Sections 5.5.2 and 5.5.3. To work towards them, we need a number of preliminary lemmas which we present in Section 5.5.1.

5.5.1 Lemmas Towards Bounds on ϵ

We first observe the following connection between the approximate payoffs $\text{Tr}(AZ)$ and $\text{Tr}(BZ)$, and $\epsilon(x, y)$, as defined in Section 5.2.1.

Lemma 5.5.1. *Consider any feasible solution to SDP2. Then*

$$\epsilon(x, y) \leq \max\{\text{Tr}(AZ) - x^T Ay, \text{Tr}(BZ) - x^T By\}.$$

Proof. Recall from the argument at the beginning of Section 5.2.4 that constraints (5.13) and (5.14) imply $\text{Tr}(AZ) \geq e_i^T Ay$ and $\text{Tr}(BZ) \geq x^T B e_i$ for all i . Hence, we have $\epsilon_A \leq \text{Tr}(AZ) - x^T Ay$ and $\epsilon_B \leq \text{Tr}(BZ) - x^T By$. \square

We thus are interested in the difference of the two matrices $P = Z^T$ and xy^T . These two matrices can be interpreted as two different probability distributions over the strategy outcomes. The matrix P is the probability distribution from the SDP which generates the approximate payoffs $\text{Tr}(AZ)$ and $\text{Tr}(BZ)$, while xy^T is the product distribution that would have resulted if the matrix had been rank-1. We will see that the difference of these distributions is key in studying the ϵ which results from SDP2. Hence, we first take steps to represent this difference.

Lemma 5.5.2. *Consider any feasible matrix \mathcal{M} to SDP2 with an eigendecomposition*

$$\mathcal{M} = \sum_{i=1}^k \lambda_i v_i v_i^T =: \sum_{i=1}^k \lambda_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T, \quad (5.28)$$

so that the eigenvectors $v_i \in \mathbb{R}^{m+n}$ are partitioned into vectors $a_i \in \mathbb{R}^m$ and $b_i \in \mathbb{R}^n$. Then for all i , $\sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$.

Proof. We know from (5.19), (5.20), and (5.21) that

$$\sum_{i=1}^k \lambda_i 1_m^T a_i a_i^T 1_m \stackrel{(5.19),(5.20)}{=} 1, \quad (5.29)$$

$$\sum_{i=1}^k \lambda_i 1_m^T a_i b_i^T 1_n \stackrel{(5.19)}{=} 1, \quad (5.30)$$

$$\sum_{i=1}^k \lambda_i 1_n^T b_i a_i^T 1_m \stackrel{(5.19)}{=} 1, \quad (5.31)$$

$$\sum_{i=1}^k \lambda_i 1_n^T b_i b_i^T 1_n \stackrel{(5.19),(5.21)}{=} 1. \quad (5.32)$$

Then by subtracting terms we have

$$(5.29) - (5.30) = \sum_{i=1}^k \lambda_i 1_m^T a_i (a_i^T 1_m - b_i^T 1_n) = 0, \quad (5.33)$$

$$(5.31) - (5.32) = \sum_{i=1}^k \lambda_i 1_n^T b_i (a_i^T 1_m - b_i^T 1_n) = 0. \quad (5.34)$$

By subtracting again these imply

$$(5.33) - (5.34) = \sum_{i=1}^k \lambda_i (1_m^T a_i - 1_n^T b_i)^2 = 0. \quad (5.35)$$

As all λ_i are nonnegative due to positive semidefiniteness of \mathcal{M} , the only way for this equality to hold is to have $1_m^T a_i = 1_n^T b_i, \forall i$. This is equivalent to the statement of the claim. \square

From Lemma 5.5.2, we can let $s_i := \sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$, and furthermore we assume without loss of generality that each s_i is nonnegative. Note that from the definition of x we have

$$x_i = \sum_{j=1}^m P_{ij} = \sum_{l=1}^k \sum_{j=1}^m \lambda_l (a_l)_i (b_l)_j = \sum_{j=1}^k \lambda_j s_j (a_l)_i. \quad (5.36)$$

Hence,

$$x = \sum_{i=1}^k \lambda_i s_i a_i. \quad (5.37)$$

Similarly,

$$y = \sum_{i=1}^k \lambda_i s_i b_i. \quad (5.38)$$

Finally note from the distribution constraint (5.15) that this implies

$$\sum_{i=1}^k \lambda_i s_i^2 = 1. \quad (5.39)$$

Lemma 5.5.3. *Let*

$$\mathcal{M} = \sum_{i=1}^k \lambda_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T,$$

be a feasible solution to SDP2, such that the eigenvectors of \mathcal{M} are partitioned into a_i and b_i with $\sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j = s_i, \forall i$. Then

$$P - xy^T = \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T.$$

Proof. Using equations (5.37) and (5.38) we can write

$$\begin{aligned}
P - xy^T &= \sum_{i=1}^k \lambda_i a_i b_i^T - \left(\sum_{i=1}^k \lambda_i s_i a_i \right) \left(\sum_{j=1}^k \lambda_j s_j b_j \right)^T \\
&= \sum_{i=1}^k \lambda_i a_i \left(b_i - s_i \sum_{j=1}^k \lambda_j s_j b_j \right)^T \\
&\stackrel{(5.39)}{=} \sum_{i=1}^k \lambda_i a_i \left(\sum_{j=1}^k \lambda_j s_j^2 b_i - s_i \sum_{j=1}^k \lambda_j s_j b_j \right)^T \\
&= \sum_{i=1}^k \sum_{j=1}^k \lambda_i \lambda_j a_i s_j (s_j b_i - s_i b_j)^T \\
&= \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j a_i - s_i a_j) (s_j b_i - s_i b_j)^T,
\end{aligned}$$

where the last line follows from observing that terms where i and j are switched can be combined. □

We can relate ϵ and $P - xy^T$ with the following lemma.

Lemma 5.5.4. *Let the matrix P and the vectors $x := P1_n$ and $y := P^T 1_m$ come from any feasible solution to SDP2. Then*

$$\epsilon \leq \frac{\|P - xy^T\|_1}{2},$$

where $\|\cdot\|_1$ here denotes the entrywise L_1 norm, i.e., the sum of the absolute values of the entries of the matrix.

Proof. Let $D := P - xy^T$. From Lemma 5.5.1,

$$\epsilon_A \leq \text{Tr}(AZ) - x^T Ay = \text{Tr}(A(Z - yx^T)).$$

If we then hold D fixed and restrict that A has entries bounded in $[0,1]$, the quantity $\text{Tr}(AD^T)$ is maximized when

$$A_{i,j} = \begin{cases} 1 & D_{i,j} \geq 0 \\ 0 & D_{i,j} < 0 \end{cases}.$$

The resulting quantity $\text{Tr}(AD^T)$ will then be the sum of all nonnegative elements of D . Since the sum of all elements in D is zero, this quantity will be equal to $\frac{1}{2}\|D\|_1$.

The proof for ϵ_B is identical, and the result follows from that ϵ is the maximum of ϵ_A and ϵ_B . □

5.5.2 Bounds on ϵ

We provide a number of bounds on $\epsilon(x,y)$ for $x := P1_n$ and $y := P^T1_m$ coming from any feasible solution to SDP2. Our first two theorems roughly state that solutions which are “close” to rank-1 provide small ϵ .

Theorem 5.5.5. *Consider any feasible solution \mathcal{M} to SDP2. Suppose \mathcal{M} is rank- k and its eigenvalues are $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0$. Then x and y constitute an ϵ -NE to the game (A, B) with $\epsilon \leq \frac{m+n}{2} \sum_{i=2}^k \lambda_i$.*

Proof. By the Perron Frobenius theorem (see e.g. [71, Chapter 8.3]), the eigenvector corresponding to λ_1 can be assumed to be nonnegative, and hence

$$s_1 = \|a_1\|_1 = \|b_1\|_1. \tag{5.40}$$

We further note that for all i , since $\begin{bmatrix} a_i \\ b_i \end{bmatrix}$ is a vector of length $m+n$ with 2-norm equal to 1, we must have

$$\left\| \begin{bmatrix} a_i \\ b_i \end{bmatrix} \right\|_1 \leq \sqrt{m+n}. \tag{5.41}$$

Since s_i is the sum of the elements of a_i and b_i , we know that

$$s_i \leq \min\{\|a_i\|_1, \|b_i\|_1\} \leq \frac{\sqrt{m+n}}{2}. \quad (5.42)$$

This then gives us

$$s_i^2 \leq \|a_i\|_1 \|b_i\|_1 \leq \frac{m+n}{4}, \quad (5.43)$$

with the first inequality following from (5.42) and the second from (5.41). Finally note that a consequence of the nonnegativity of $\|\cdot\|_1$ and (5.41) is that for all i, j ,

$$\|a_i\|_1 \|b_j\|_1 + \|b_i\|_1 \|a_j\|_1 \leq (\|a_i\|_1 + \|b_i\|_1)(\|a_j\|_1 + \|b_j\|_1) = \left\| \begin{bmatrix} a_i \\ b_i \end{bmatrix} \right\|_1 \left\| \begin{bmatrix} a_j \\ b_j \end{bmatrix} \right\|_1 \stackrel{(5.41)}{\leq} m+n. \quad (5.44)$$

Now we let $D := P - xy^T$ and upper bound $\frac{1}{2}\|D\|_1$ using Lemma 5.5.3.

$$\begin{aligned} \frac{1}{2}\|D\|_1 &= \frac{1}{2} \left\| \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T \right\|_1 \\ &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \|\lambda_i \lambda_j (s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T\|_1 \\ &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j \|s_j a_i - s_i a_j\|_1 \|s_j b_i - s_i b_j\|_1 \\ &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j \|a_i\|_1 + s_i \|a_j\|_1)(s_j \|b_i\|_1 + s_i \|b_j\|_1) \quad (5.45) \\ &\stackrel{(5.40), (5.43)}{\leq} \frac{1}{2} \sum_{j=2}^k \lambda_1 s_1^2 \lambda_j (s_j + \|a_j\|_1)(s_j + \|b_j\|_1) \\ &\quad + \frac{1}{2} \sum_{i=2}^k \sum_{j>i}^k \lambda_i \lambda_j \left(s_j^2 \frac{m+n}{4} + s_i^2 \frac{m+n}{4} + s_i s_j \|a_i\|_1 \|b_j\|_1 + s_i s_j \|a_j\|_1 \|b_i\|_1 \right) \\ &\stackrel{(5.41), (5.44), (5.42)}{\leq} \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i \\ &\quad + \frac{1}{2} \sum_{i=2}^k \sum_{j>i}^k \lambda_i \lambda_j \frac{m+n}{4} (s_i^2 + s_j^2) + \lambda_i \lambda_j s_i s_j (m+n) \end{aligned}$$

$$\begin{aligned}
&\stackrel{\text{AMGM}^6}{\leq} \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{m+n}{2} \sum_{i=2}^k \sum_{j>i}^k \lambda_i \lambda_j \left(\frac{s_i^2 + s_j^2}{4} + \frac{s_i^2 + s_j^2}{2} \right) \\
&= \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{3(m+n)}{8} \sum_{i=2}^k \sum_{j>i}^k \lambda_i \lambda_j (s_i^2 + s_j^2) \\
&= \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{3(m+n)}{8} \left(\sum_{i=2}^k \lambda_i s_i^2 \sum_{j>i}^k \lambda_j + \sum_{i=2}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2 \right) \\
&= \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{3(m+n)}{8} \left(\sum_{j=2}^k \lambda_j \sum_{2 \leq i < j} \lambda_i s_i^2 + \sum_{i=2}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2 \right) \\
&\leq \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{3(m+n)}{8} \left(\sum_{j=2}^k \lambda_j s_j^2 \right) \sum_{i=2}^k \lambda_i \\
&\stackrel{(5.39)}{=} \frac{m+n}{2} \lambda_1 s_1^2 \sum_{i=2}^k \lambda_i + \frac{3(m+n)}{8} (1 - \lambda_1 s_1^2) \sum_{i=2}^k \lambda_i \\
&= \frac{m+n}{8} (3 + \lambda_1 s_1^2) \sum_{i=2}^k \lambda_i \\
&\stackrel{(5.39)}{\leq} \frac{m+n}{2} \sum_{i=2}^k \lambda_i.
\end{aligned}$$

□

The following theorem quantifies how making the objective of the diagonal gap algorithm from Section 5.4 small makes ϵ small. The proof is similar to the proof of Theorem 5.5.5.

Theorem 5.5.6. *Let \mathcal{M} be a feasible solution to SDP2. Then, x and y constitute an ϵ -NE to the game (A, B) with $\epsilon \leq \frac{3(m+n)}{8} (\text{Tr}(\mathcal{M}) - x^T x - y^T y)$.*

Proof. Let \mathcal{M} be rank- k with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0$ and eigenvectors v_1, \dots, v_k

partitioned as in Lemma 5.5.2 so that $v_i = \begin{bmatrix} a_i \\ b_i \end{bmatrix}$ with $\sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$ for $i = 1, \dots, k$.

Let $s_i := \sum_{j=1}^m (a_i)_j$. Then we have $\text{Tr}(\mathcal{M}) = \sum_{i=1}^k \lambda_i$, and

$$x^T x + y^T y \stackrel{(5.37), (5.38)}{=} \left(\sum_{i=1}^k \lambda_i s_i v_i \right)^T \left(\sum_{i=1}^k \lambda_i s_i v_i \right) = \sum_{i=1}^k \lambda_i^2 s_i^2. \quad (5.46)$$

⁶AMGM is used to denote the arithmetic-mean-geometric-mean inequality.

We now get the following chain of inequalities (the first one follows from Lemma 5.5.4 and inequality (5.45)):

$$\begin{aligned}
\epsilon &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j \|a_i\|_1 + s_i \|a_j\|_1) (s_j \|b_i\|_1 + s_i \|b_j\|_1) \\
&\stackrel{(5.40),(5.43)}{\leq} \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j \left(s_j^2 \frac{m+n}{4} + s_i^2 \frac{m+n}{4} + s_i s_j \|a_i\|_1 \|b_j\|_1 + s_i s_j \|a_j\|_1 \|b_i\|_1 \right) \\
&\stackrel{(5.44)}{\leq} \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j \frac{m+n}{4} (s_i^2 + s_j^2) + \lambda_i \lambda_j s_i s_j (m+n) \\
&\stackrel{AMGM}{\leq} \frac{m+n}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j \left(\frac{s_i^2 + s_j^2}{4} + \frac{s_i^2 + s_j^2}{2} \right) \\
&= \frac{3(m+n)}{8} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_i^2 + s_j^2) \\
&= \frac{3(m+n)}{8} \left(\sum_{i=1}^k \lambda_i s_i^2 \sum_{j>i}^k \lambda_j + \sum_{i=1}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2 \right) \\
&= \frac{3(m+n)}{8} \left(\sum_{j=1}^k \lambda_j \sum_{1 \leq i < j} \lambda_i s_i^2 + \sum_{i=1}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2 \right) \\
&= \frac{3(m+n)}{8} \left(\sum_{i=1}^k \lambda_i \sum_{j \neq i} \lambda_j s_j^2 \right) \\
&\stackrel{(5.39)}{=} \frac{3(m+n)}{8} \left(\sum_{i=1}^k \lambda_i (1 - \lambda_i s_i^2) \right) \\
&= \frac{3(m+n)}{8} \left(\sum_{i=1}^k \lambda_i - \sum_{i=1}^k \lambda_i^2 s_i^2 \right) \stackrel{(5.46)}{=} \frac{3(m+n)}{8} (\text{Tr}(\mathcal{M}) - x^T x - y^T y).
\end{aligned}$$

□

We now give a bound on ϵ which is dependent on the nonnegative rank of the matrix returned by SDP2. Our analysis will also be useful for the next subsection. To begin, we first recall the definition of the nonnegative rank.

Definition 5.5.7. *The nonnegative rank of a (nonnegative) $m \times n$ matrix M is the smallest k for which there exist a nonnegative $m \times k$ matrix U and a nonnegative $n \times k$ matrix V*

such that $M = UV^T$. Such a decomposition is called a nonnegative matrix factorization of M .

Theorem 5.5.8. *Consider the matrix P from any feasible solution to SDP2. Suppose its nonnegative rank is k . Then $x := P1_n$ and $y := P^T1_m$ constitute an ϵ -NE to the game (A, B) with $\epsilon \leq 1 - \frac{1}{k}$.*

Proof. Since P has nonnegative rank k and its entries sum up to 1, we can write $P = \sum_{i=1}^k \sigma_i a_i b_i^T$, where $a_i \in \Delta_m, b_i \in \Delta_n$, and $\sum_{i=1}^k \sigma_i = 1$. From Lemma 5.5.4 and inequality (5.45) (keeping in mind that $s_i = 1, \forall i$) we have

$$\begin{aligned} \epsilon &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \sigma_i \sigma_j (\|a_i\|_1 + \|a_j\|_1) (\|b_i\|_1 + \|b_j\|_1) \\ &\leq 2 \sum_{i=1}^k \sum_{j>i}^k \sigma_i \sigma_j \\ &= 2 \left(\frac{1}{2} \left(\sum_{i=1}^k \sigma_i \sum_{j=1}^k \sigma_j - \sum_{i=1}^k \sigma_i^2 \right) \right) \\ &= 1 - \sum_{i=1}^k \sigma_i^2 \\ &\leq 1 - \frac{1}{k}, \end{aligned}$$

where the last line follows from the fact that $\|v\|_2^2 \geq \frac{1}{k}$ for any vector $v \in \Delta_k$. □

5.5.3 Bounds on ϵ in the Rank-2 Case

We now provide a number of bounds on $\epsilon(x, y)$ with $x := P1_n$ and $y := P^T1_m$ which hold for rank-2 feasible solutions \mathcal{M} to SDP2 (note that P will have rank at most 2 in this case). This is motivated by our ability to show stronger (constant) bounds in this case, and the fact that we often recover rank-2 (or rank-1) solutions with our algorithms in Section 5.4. Furthermore, our analysis will use the special property that a rank-2 nonnegative matrix will have nonnegative rank also equal to two, and that a nonnegative factorization of it can

be computed in polynomial time (see, e.g., Section 4 of [26]). We begin with the following observation, which follows from Theorem 5.5.8 when $k = 2$.

Corollary 5.5.9. *If the matrix P from a feasible solution to SDP2 is rank-2, then x and y constitute a $\frac{1}{2}$ -NE.*

We now show how this pair of strategies can be refined.

Theorem 5.5.10. *If the matrix P from a feasible solution to SDP2 is rank-2, then either x and y constitute a $\frac{5}{11}$ -NE, or a $\frac{5}{11}$ -NE can be recovered from P in polynomial time.*

Proof. We consider 3 cases, depending on whether $\epsilon_A(x, y)$ and $\epsilon_B(x, y)$ are greater than or less than .4. If $\epsilon_A \leq .4, \epsilon_B \leq .4$, then (x, y) is already a .4-Nash equilibrium. Now consider the case when $\epsilon_A \geq .4, \epsilon_B \geq .4$. Since $\epsilon_A \leq \text{Tr}(A(P - xy^T)^T)$ and $\epsilon_B \leq \text{Tr}(B(P - xy^T)^T)$ as seen in the proof of Lemma 5.5.1, we have, reusing the notation in the proof of Theorem 5.5.8,

$$\sigma_1 \sigma_2 (a_1 - a_2)^T A (b_1 - b_2) \geq .4, \sigma_1 \sigma_2 (a_1 - a_2)^T B (b_1 - b_2) \geq .4.$$

Since A, a_1, a_2, b_1 , and b_2 are all nonnegative and $\sigma_1 \sigma_2 \leq \frac{1}{4}$,

$$a_1^T A b_1 + a_2^T A b_2 \geq (a_1 - a_2)^T A (b_1 - b_2) \geq 1.6,$$

and the same inequalities hold for for player B. In particular, since A and B have entries bounded in $[0,1]$ and a_1, a_2, b_1 , and b_2 are simplex vectors, all the quantities $a_1^T A b_1, a_2^T A b_2, a_1^T B b_1$, and $a_2^T B b_2$ are at most 1, and consequently at least .6. Hence (a_1, a_2) and (a_2, b_2) are both .4-Nash equilibria.

Now suppose that (x, y) is a .4-NE for one player (without loss of generality player A) but not for the other (without loss of generality player B). Then $\epsilon_A \leq .4$, and $\epsilon_B \geq .4$. Let y^* be a best response for player B to x , and let $p = \frac{1}{1 + \epsilon_B - \epsilon_A}$. Consider the strategy profile $(\tilde{x}, \tilde{y}) := (x, py + (1 - p)y^*)$. This can be interpreted as the outcome (x, y) occurring with probability p , and the outcome (x, y^*) happening with probability $1 - p$. In the first case,

player A will have $\epsilon_A(x, y) = \epsilon_A$ and player B will have $\epsilon_B(x, y) = \epsilon_B$. In the second outcome, player A will have $\epsilon_A(x, y^*)$ at most 1, while player B will have $\epsilon_B(x, y^*) = 0$. Then under this strategy profile, both players have the same upper bound for ϵ , which equals $\epsilon_{BP} = \frac{\epsilon_B}{1 + \epsilon_B - \epsilon_A}$. To find the worst case for this value, let $\epsilon_B = .5$ (note from Theorem 5.5.9 that $\epsilon_B \leq \frac{1}{2}$) and $\epsilon_A = .4$, and this will return $\epsilon = \frac{5}{11}$.

□

We now show a stronger result in the case of symmetric games.

Definition 5.5.11. *A symmetric game is a game in which the payoff matrices A and B satisfy $B = A^T$.*

Definition 5.5.12. *A Nash equilibrium strategy (x, y) is said to be symmetric if $x = y$.*

Theorem 5.5.13 (see Theorem 2 in [77]). *Every symmetric bimatrix game has a symmetric Nash equilibrium.*

For the proof of Theorem 5.5.17 below we modify SDP2 so that we are seeking a symmetric solution. We also need a more specialized notion of the nonnegative rank.

Definition 5.5.14. *A matrix M is completely positive (CP) if it admits a decomposition $M = UU^T$ for some nonnegative matrix U .*

Definition 5.5.15. *The CP-rank of an $n \times n$ CP matrix M is the smallest k for which there exists a nonnegative $n \times k$ matrix U such that $M = UU^T$.*

Theorem 5.5.16 (see e.g. Theorem 2.1 in [18]). *A rank-2, nonnegative, and positive semidefinite matrix is CP and has CP-rank 2.*

It is also known (see e.g., Section 4 in [55]) that the CP factorization of a rank-2 CP matrix can be found to arbitrary accuracy in polynomial time.

Theorem 5.5.17. *Suppose the constraint $P \succeq 0$ is added to SDP2. Then if in a feasible solution to this new SDP the matrix P is rank-2, either x and y constitute a symmetric $\frac{1}{3}$ -NE, or a symmetric $\frac{1}{3}$ -NE can be recovered from P in polynomial time.*

Proof. If (x, y) is already a symmetric $\frac{1}{3}$ -NE, then the claim is established. Now suppose that (x, y) does not constitute a $\frac{1}{3}$ -Nash equilibrium. Similarly as in the proof of Theorem 5.5.8, we can decompose P into $\sum_{i=1}^2 \sigma_i a_i a_i^T$, where $\sum_{i=1}^2 \sigma_i = 1$ and each a_i is a vector on the unit simplex. Then we have

$$\sigma_1 \sigma_2 (a_1 - a_2)^T A (a_1 - a_2) \geq \frac{1}{3}.$$

Since A, a_1 , and a_2 are all nonnegative, and $\sigma_1 \sigma_2 \leq \frac{1}{4}$, we get

$$a_1^T A a_1 + a_2^T A a_2 \geq (a_1 - a_2)^T A (a_1 - a_2) \geq \frac{4}{3}.$$

In particular, at least one of $a_1^T A a_1$ and $a_2^T A a_2$ is at least $\frac{2}{3}$. Since the maximum possible payoff is 1, at least one of (a_1, a_1) and (a_2, a_2) is a (symmetric) $\frac{1}{3}$ -Nash equilibrium. \square

Remark 5.5.1. For symmetric games, instead of the construction stated in Theorem 5.5.17, one can simply optimize over a smaller $m \times m$ matrix (note $m = n$). This is the relaxed version of exchangeable equilibria [102], with the completely positive constraint relaxed to a psd constraint.

Remark 5.5.2. The statements of Corollary 5.5.9, and Theorem 5.5.10, and Theorem 5.5.17 hold for any rank-2 correlated equilibrium. Indeed, given any rank-2 (equivalently, nonnegative-rank-2) correlated equilibrium P , one can complete it to a (rank-2) feasible solution to SDP2 as follows. Let $P = \sum_{i=1}^2 \sigma_i a_i b_i^T$, where $a_i \in \Delta_m, b_i \in \Delta_n$, and $\sigma_1 + \sigma_2 = 1$.

It is easy to check that

$$\mathcal{M} := \sum_{i=1}^2 \sigma_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T$$

is feasible to SDP2.

5.6 Bounding Payoffs and Strategy Exclusion in Symmetric Games

In addition to finding ϵ -additive Nash equilibria, our SDP approach can be used to answer certain questions of economic interest about Nash equilibria without actually computing them. For instance, economists often would like to know the maximum welfare (sum of the two players' payoffs) achievable under any Nash equilibrium, or whether there exists a Nash equilibrium in which a given subset of strategies (corresponding, e.g., to undesirable behavior) is not played. Both these questions are NP-hard for bimatrix games [40], even when the game is symmetric and only symmetric equilibria are considered [28]. In this section, we consider these two problems in the symmetric setting and compare the performance of our SDP approach to an LP approach which searches over symmetric correlated equilibria. For general equilibria, it turns out that for these two specific questions, our SDP approach is equivalent to an LP that searches over correlated equilibria.

5.6.1 Bounding Payoffs

When designing policies that are subject to game theoretic behavior by agents, economists would often like to find one with a good socially optimal outcome, which usually corresponds to an equilibrium giving the maximum welfare. Hence, given a game, it is of interest to know the highest achievable welfare under any Nash equilibrium. For symmetric games, symmetric equilibria are of particular interest as they reflect the notion that identical agents should behave similarly given identical options.

Note that the maximum welfare of a symmetric game under any symmetric Nash equilibrium is equal to the optimal value of the following quadratic program:

$$\begin{aligned} \max_{x \in \Delta_m} \quad & 2x^T Ax \\ \text{subject to} \quad & x^T Ax \geq e_i^T Ax, \forall i \in \{1, \dots, m\}. \end{aligned} \tag{5.47}$$

One can find an upper bound on this number by solving an LP which searches over symmetric correlated equilibria:

$$\max_{P \in \mathbb{S}^{m,m}} \text{Tr}(AP^T) \tag{LP1}$$

$$\text{subject to} \quad \sum_{i=1}^m \sum_{j=1}^m P_{i,j} = 1 \tag{5.48}$$

$$\sum_{j=1}^m A_{i,j} P_{i,j} \geq \sum_{j=1}^m A_{k,j} P_{i,j}, \forall i, k \in \{1, \dots, m\}, \tag{5.49}$$

$$P \succeq 0. \tag{5.50}$$

A potentially better upper bound on the maximum welfare can be obtained from a version of SDP2 adapted to this specific problem:

$$\max_{P \in \mathbb{S}^{m,m}} \text{Tr}(AP^T) \tag{SDP3}$$

$$\text{subject to} \quad (5.48), (5.49), (5.50)$$

$$P \succeq 0.$$

To test the quality of these upper bounds, we tested this LP and SDP on a random sample of one hundred 5×5 and 10×10 games⁷. The resulting upper bounds are in Figure 5.2, which shows that the bound returned by SDP3 was exact in a large number of the experiments.⁸

⁷The matrix A in each game was randomly generated with diagonal entries uniform and independent in $[0, .5]$ and off-diagonal entries uniform and independent in $[0, 1]$.

⁸The computation of the exact maximum payoffs was done with the `1rsnash` software [12], which computes all extreme Nash equilibria. For a definition of extreme Nash equilibria and for understanding why it is sufficient for us to compare against extreme Nash equilibria (both in Section 5.6.1 and in Section 5.6.2), see Appendix A.3. The computation of the SDP upper bound has been implemented in the file `nashbound.m`, which is publicly available at https://github.com/jeffreyzhang92/SDP_Nash along with the instance data. This file more generally computes an SDP-based lower bound on the minimum of an input quadratic function over the set of Nash equilibria of a bimatrix game. The file also takes as an argument whether one wishes to only consider symmetric equilibria when the game is symmetric.




Figure 5.2: The quality of the upper bound on the maximum welfare obtained by LP1 and SDP3 on 100 5×5 games (left) and 100 10×10 games (right).

5.6.2 Strategy Exclusion

The strategy exclusion problem asks, given a subset of strategies $\mathcal{S} = (\mathcal{S}_x, \mathcal{S}_y)$, with $\mathcal{S}_x \subseteq \{1, \dots, m\}$ and $\mathcal{S}_y \subseteq \{1, \dots, n\}$, is there a Nash equilibrium in which no strategy in \mathcal{S} is played with positive probability. We will call a set \mathcal{S} “persistent” if the answer to this question is negative, i.e. at least one strategy in \mathcal{S} is played with positive probability in every Nash equilibrium. One application of the strategy exclusion problem is to understand whether certain strategies can be discouraged in the design of a game, such as reckless behavior in a game of chicken or defecting in a game of prisoner’s dilemma. In these particular examples these strategy sets are persistent and cannot be discouraged.

As in the previous subsection, we consider the strategy exclusion problem for symmetric strategies in symmetric games (such as the aforementioned games of chicken and prisoner's dilemma). A quadratic program which addresses this problem is as follows:

$$\begin{aligned} \min_{x \in \Delta_m} \quad & \sum_{i \in \mathcal{S}_x} x_i \\ \text{subject to} \quad & x^T A x \geq e_i^T A x, \forall i \in \{1, \dots, m\}. \end{aligned} \tag{5.51}$$

Observe that by design, \mathcal{S} is persistent if and only if this quadratic program has a positive optimal value. As in the previous subsection, an LP relaxation of this problem which searches over symmetric correlated equilibria is given by

$$\begin{aligned} \min_{P \in \mathbb{S}^{m,m}} \quad & \sum_{i \in \mathcal{S}_x} \sum_{j=1}^m P_{ij} \\ \text{subject to} \quad & (5.48), (5.49), (5.50). \end{aligned} \tag{LP2}$$

The SDP relaxation that we propose for the strategy exclusion problem is the following:

$$\begin{aligned} \min_{P \in \mathbb{S}^{m,m}} \quad & \sum_{i \in \mathcal{S}_x} \sum_{j=1}^m P_{ij} \\ \text{subject to} \quad & (5.48), (5.49), (5.50) \\ & P \succeq 0. \end{aligned} \tag{SDP4}$$

Our approach would be to declare that the strategy set \mathcal{S}_x is persistent if and only if SDP4 has a positive optimal value.

Note that since the optimal value of SDP4 is a lower bound for that of (5.51), SDP4 carries over the property that if a set \mathcal{S} is not persistent, then the SDP for sure returns zero. Thus, when using SDP4 on a set which is not persistent, our algorithm will always be correct. However, this is not necessarily the case for a persistent set. While we can be certain that a set is persistent if SDP4 returns a positive optimal value (again, because the

optimal value of SDP4 is a lower bound for that of (5.51)), there is still the possibility that for a persistent set SDP4 will have optimal value zero. The same arguments hold for the optimal value of LP2.

To test the performance of LP2 and SDP4, we generated 100 random games of size 5×5 and 10×10 and computed all their symmetric extreme Nash equilibria⁹. We then, for every strategy set \mathcal{S} of cardinality one and two, checked whether that set of strategies was persistent, first by checking among the extreme Nash equilibria, then through LP2 and SDP4. The results are presented in Tables 5.2 and 5.3. As can be seen, SDP4 was quite effective for the strategy exclusion problem.

Table 5.2: Performance of LP2 and SDP4 on 5×5 games

$ \mathcal{S} $	1	2
Number of total sets	500	1000
Number of persistent sets	245	748
Persistent sets certified (LP2)	177 (72.2%)	661 (88.7%)
Persistent sets certified (SDP4)	245 (100%)	748 (100%)

Table 5.3: Performance of LP2 and SDP4 on 10×10 games

$ \mathcal{S} $	1	2
Number of total sets	1000	4500
Number of persistent sets	326	2383
Persistent sets certified (LP2)	39 (12.0%)	630 (26.4%)
Persistent sets certified (SDP4)	318 (97.5%)	2368 (99.4%)

5.7 Connection to the Sum of Squares/Lasserre Hierarchy

In this section, we clarify the connection of the SDPs we have proposed in this chapter to those arising in the sum of squares/Lasserre hierarchy. We start by briefly reviewing this hierarchy.

⁹The exact computation of the exact Nash equilibria was done again with the `lrsnash` software [12], which computes extreme Nash equilibria. To understand why this suffices for our purposes see Appendix A.3.

5.7.1 Sum of Squares/Lasserre Hierarchy

The sum of squares/Lasserre hierarchy¹⁰ gives a recipe for constructing a sequence of SDPs whose optimal values converge to the optimal value of a given polynomial optimization problem. Recall that for a POP of the form

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & p(x) \\ \text{subject to} \quad & q_i(x) \geq 0, \forall i \in \{1, \dots, m\}, \end{aligned} \tag{5.52}$$

where p, q_i are polynomial functions, k -th level of the Lasserre hierarchy is given by

$$\begin{aligned} \gamma_{sos}^k := \max_{\gamma, \sigma_i} \quad & \gamma \\ \text{subject to} \quad & p(x) - \gamma = \sigma_0(x) + \sum_{i=1}^m \sigma_i(x)q_i(x), \\ & \sigma_i \text{ is sos, } \forall i \in \{0, \dots, m\}, \\ & \sigma_0, g_i \sigma_i \text{ have degree at most } 2k, \forall i \in \{1, \dots, m\}. \end{aligned} \tag{5.53}$$

Recall that any fixed level of this hierarchy gives an SDP of size polynomial in n and that, if the quadratic module generated by $\{x \in \mathbb{R}^n | g_i(x) \geq 0\}$ is Archimedean (see, e.g. [63] for definition), then $\lim_{k \rightarrow \infty} \gamma_{sos}^k = p^*$, where p^* is the optimal value of the pop in (5.52). The latter statement is a consequence of Putinar's positivstellensatz (see, e.g. [92], [61]).

5.7.2 The Lasserre Hierarchy and SDP1

One can show, e.g. via the arguments in [62], that the feasible sets of the SDPs dual to the SDPs underlying the hierarchy we summarized above produce an arbitrarily tight outer approximation to the convex hull of the set of Nash equilibria of any game. The downside of this approach, however, is that the higher levels of the hierarchy can get expensive very

¹⁰The unfamiliar reader is referred to [61, 88, 63] for an introduction to this hierarchy and the related theory of moment relaxations.

quickly. This is why the approach we took in this chapter was instead to improve the first level of the hierarchy. The next proposition formalizes this connection.

Proposition 5.7.1. *Consider the problem of minimizing any quadratic objective function over the set of Nash equilibria of a bimatrix game. Then, SDP1 (and hence SDP2) gives a lower bound on this problem which is no worse than that produced by the first level of the Lasserre hierarchy.*

Proof. To prove this proposition we show that the first level of the Lasserre hierarchy is dual to a weakened version of SDP1.

Explicit parametrization of first level of the Lasserre hierarchy. Consider the formulation of the Lasserre hierarchy in (5.53) with $k = 1$. Suppose we are minimizing a quadratic function

$$f(x, y) = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T \mathcal{C} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

over the set of Nash equilibria as described by the linear and quadratic constraints in (5.2).

If we apply the first level of the Lasserre hierarchy to this particular pop, we get

$$\begin{aligned}
& \max_{Q, \alpha, \chi, \beta, \psi, \eta} \quad \gamma \\
\text{subject to} \quad & \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T \mathcal{C} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} - \gamma = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T Q \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} + \sum_{i=1}^m \alpha_i (x^T A y - e_i^T A y) \\
& \quad + \sum_{i=1}^n \beta_i (x^T B y - x^T B e_i) \\
& \quad + \sum_{i=1}^m \chi_i x_i + \sum_{i=1}^n \psi_i y_i \\
& \quad + \eta_1 \left(\sum_{i=1}^m x_i - 1 \right) + \eta_2 \left(\sum_{i=1}^n y_i - 1 \right), \\
& \quad Q \succeq 0, \\
& \quad \alpha, \chi, \beta, \psi \geq 0,
\end{aligned} \tag{5.54}$$

where $Q \in \mathbb{S}^{m+n+1 \times m+n+1}$, $\alpha, \chi \in \mathbb{R}^m$, $\beta, \psi \in \mathbb{R}^n$, $\eta \in \mathbb{R}^2$.

By matching coefficients of the two quadratic functions on the left and right hand sides of (5.54), this SDP can be written as

$$\begin{aligned}
& \max_{\gamma, \alpha, \beta, \chi, \psi, \eta} \quad \gamma \\
\text{subject to} \quad & \mathcal{H} \succeq 0, \\
& \quad \alpha, \beta, \chi, \psi \geq 0,
\end{aligned} \tag{5.55}$$

where

$$\mathcal{H} := \frac{1}{2} \begin{bmatrix} 0 & (-\sum_{i=1}^m \alpha_i)A + (-\sum_{i=1}^n \beta_i)B & \sum_{i=1}^n \beta_i B_{:,i} - \chi - \eta_1 \mathbf{1}_m \\ (-\sum_{i=1}^m \alpha_i)A + (-\sum_{i=1}^n \beta_i)B & 0 & \sum_{i=1}^m \alpha_i A_{:,i}^T - \psi - \eta_2 \mathbf{1}_n \\ \sum_{i=1}^n \beta_i B_{:,i}^T - \chi^T - \eta_1 \mathbf{1}_m^T & \sum_{i=1}^m \alpha_i A_{:,i} - \psi^T - \eta_2 \mathbf{1}_n^T & 2\eta_1 + 2\eta_2 - 2\gamma \end{bmatrix} + \mathcal{C}. \tag{5.56}$$

Dual of a weakened version of SDP1. With this formulation in mind, let us consider a weakened version of SDP1 with only the relaxed Nash constraints, unity constraints, and nonnegativity constraints on x and y in the last column (i.e., the nonnegativity constraint is not applied to the entire matrix). Let the objective be $\text{Tr}(C\mathcal{M}')$. To write this new SDP in standard form, let

$$\mathcal{A}_i := \frac{1}{2} \begin{bmatrix} 0 & A & 0 \\ A^T & 0 & -A_i^T \\ 0 & -A_i & 0 \end{bmatrix}, \mathcal{B}_i := \frac{1}{2} \begin{bmatrix} 0 & B & -B_{,i} \\ B^T & 0 & 0 \\ -B_{,i}^T & 0 & 0 \end{bmatrix},$$

$$\mathcal{S}_1 := \frac{1}{2} \begin{bmatrix} 0 & 0 & 1_m \\ 0 & 0 & 0 \\ 1_m^T & 0 & -2 \end{bmatrix}, \mathcal{S}_2 := \frac{1}{2} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1_n \\ 0 & 1_n^T & -2 \end{bmatrix}.$$

Let \mathcal{N}_i be the matrix with all zeros except a $\frac{1}{2}$ at entry $(i, m+n+1)$ and $(m+n+1, i)$ (or a 1 if $i = m+n+1$).

Then this SDP can be written as

$$\min_{\mathcal{M}'} \quad \text{Tr}(C\mathcal{M}') \quad (\text{SDP0})$$

$$\text{subject to} \quad \mathcal{M}' \succeq 0, \quad (5.57)$$

$$\text{Tr}(\mathcal{N}_i\mathcal{M}') \geq 0, \forall i \in \{1, \dots, m+n\}, \quad (5.58)$$

$$\text{Tr}(\mathcal{A}_i\mathcal{M}') \geq 0, \forall i \in \{1, \dots, m\}, \quad (5.59)$$

$$\text{Tr}(\mathcal{B}_i\mathcal{M}') \geq 0, \forall i \in \{1, \dots, n\}, \quad (5.60)$$

$$\text{Tr}(\mathcal{S}_1\mathcal{M}') = 0, \quad (5.61)$$

$$\text{Tr}(\mathcal{S}_2\mathcal{M}') = 0, \quad (5.62)$$

$$\text{Tr}(\mathcal{N}_{m+n+1}) = 1. \quad (5.63)$$

We now create dual variables for each constraint; we choose α_i and β_i for the relaxed Nash constraints (5.59) and (5.60), η_1 and η_2 for the unity constraints (5.61) and (5.62), χ for the nonnegativity of x (5.58), ψ for the nonnegativity of y (5.58), and γ for the final constraint on the corner (5.63). These variables are chosen to coincide with those used in the parametrization of the first level of the Lasserre hierarchy, as can be seen more clearly below.

We then write the dual of the above SDP as

$$\begin{aligned} & \max_{\alpha, \beta, \lambda, \gamma} \quad \gamma \\ \text{subject to} \quad & \sum_{i=1}^m \alpha_i \mathcal{A}_i + \sum_{i=1}^n \beta_i \mathcal{B}_i + \sum_{i=1}^2 \eta_i \mathcal{S}_i + \sum_{i=1}^m \mathcal{N}_{i+n} \chi_i + \sum_{i=1}^n \mathcal{N}_i \psi_i + \gamma \mathcal{N}_{m+n+1} \preceq \mathcal{C}, \\ & \alpha, \beta, \chi, \psi \geq 0. \end{aligned}$$

which can be rewritten as

$$\begin{aligned} & \max_{\alpha, \beta, \chi, \psi, \gamma} \quad \gamma \\ \text{subject to} \quad & \mathcal{G} \succeq 0, \\ & \alpha, \beta, \chi, \psi \geq 0, \end{aligned} \tag{5.64}$$

where

$$\mathcal{G} := \frac{1}{2} \begin{bmatrix} 0 & (-\sum_{i=1}^m \alpha_i)A + (-\sum_{i=1}^n \beta_i)B & \sum_{i=1}^n \beta_i B_{,i} - \chi - \eta_1 \mathbf{1}_m \\ (-\sum_{i=1}^m \alpha_i)A + (-\sum_{i=1}^n \beta_i)B & 0 & \sum_{i=1}^m \alpha_i A_{,i}^T - \psi - \eta_2 \mathbf{1}_n \\ \sum_{i=1}^n \beta_i B_{,i}^T - \chi^T - \eta_1 \mathbf{1}_m^T & \sum_{i=1}^m \alpha_i A_{,i} - \psi^T - \eta_2 \mathbf{1}_n^T & 2\eta_1 + 2\eta_2 - 2\gamma \end{bmatrix} + \mathcal{C}.$$

We can now see that the matrix \mathcal{G} coincides with the matrix \mathcal{H} in the SDP (5.55). Then we have

$$(5.54)^{opt} = (5.55)^{opt} = (5.64)^{opt} \leq SDP0^{opt} \leq SDP1^{opt},$$

where the first inequality follows from weak duality, and the second follows from that the constraints of SDP0 are a subset of the constraints of SDP1. \square

Remark 5.7.1. The Lasserre hierarchy can be viewed in each step as a pair of primal-dual SDPs: the sum of squares formulation which we have just presented, and a moment formulation which is dual to the sos formulation [61]. All our SDPs in this chapter can be viewed more directly as an improvement upon the moment formulation.

Remark 5.7.2. One can see, either by inspection or as an implication of the proof of Theorem 5.2.2, that in the case where the objective function corresponds to maximizing player A's and/or B's payoffs¹¹, SDPs (5.55) and (5.64) are infeasible. This means that for such problems the first level of the Lasserre hierarchy gives an upper bound of $+\infty$ on the maximum payoff. On the other hand, the additional valid inequalities in SDP2 guarantee that the resulting bound is always finite.

5.8 Future Work

Our work leaves many avenues of further research. Are there other interesting subclasses of games (besides strictly competitive games) for which our SDP is guaranteed to recover an exact Nash equilibrium? Can the guarantees on ϵ in Section 5.5 be improved in the rank-2 case (or the general case) by improving our analysis? Is there a polynomial time algorithm that is guaranteed to find a rank-2 solution to SDP2? Such an algorithm, together with our analysis, would improve the best known approximation bound for symmetric games (see Theorem 5.5.17). Can this bound be extended to general games? We show in Appendix A.4 that some natural approaches based on symmetrization of games do not immediately lead to a positive answer to this question. Can SDPs in a higher level of the Lasserre hierarchy be used to achieve better ϵ guarantees? What are systematic ways of adding valid inequalities

¹¹This would be the case, for example, in the maximum social welfare problem of Section 5.6.1, where the matrix of the quadratic form in the objective function is given by

$$\mathcal{C} = \begin{bmatrix} 0 & -A - B & 0 \\ -A - B & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

to these higher-order SDPs by exploiting the structure of the Nash equilibrium problem? For example, since any strategy played with positive probability must give the same payoff, one can add a relaxed version of the cubic constraints

$$x_i x_j (e_i^T A y - e_j^T A y) = 0, \forall i, j \in \{1, \dots, m\}$$

to the SDP underlying the second level of the Lasserre hierarchy. What are other valid inequalities for the second level? Finally, our algorithms were specifically designed for two-player one-shot games. This leaves open the design and analysis of semidefinite relaxations for repeated games or games with more than two players.

Appendix A

Appendices for Nash Equilibria

A.1 Statistics on ϵ from Algorithms in Section 5.4

Below are statistics for the ϵ recovered in 100 random games of varying sizes using the algorithms of Section 5.4.

Table A.1: Statistics on ϵ for 5×5 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0702	0.0040	0.0004	0.0099
Diagonal Gap	0.0448	0.0027	0	0.0061

Table A.2: Statistics on ϵ for 10×5 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0327	0.0044	0.0021	0.0064
Diagonal Gap	0.0267	0.0033	0.0006	0.0053

Table A.3: Statistics on ϵ for 10×10 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0373	0.0058	0.0039	0.0065
Diagonal Gap	0.0266	0.0043	0.0026	0.0051

Table A.4: Statistics on ϵ for 15×10 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0206	0.0050	0.0034	0.0045
Diagonal Gap	0.0212	0.0038	0.0025	0.0039

Table A.5: Statistics on ϵ for 15×15 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0169	0.0051	0.0042	0.0039
Diagonal Gap	0.0159	0.0038	0.0029	0.0034

Table A.6: Statistics on ϵ for 20×15 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0152	0.0046	0.0035	0.0036
Diagonal Gap	0.0119	0.0032	0.0022	0.0027

Table A.7: Statistics on ϵ for 20×20 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0198	0.0046	0.0039	0.0034
Diagonal Gap	0.0159	0.0032	0.0024	0.0032

A.2 Comparison with an SDP Approach from [60]

In this section, at the request of a referee, we compare the first level of the SDP hierarchy given in [60, Section 4] to SDP2 using $\text{Tr}(M)$ as the objective function on 100 randomly generated games for each size given in the tables below. The first level of the hierarchy in [60] optimizes over a matrix which is slightly bigger than the one in SDP2, though it has a number of constraints linear in the size of the game considered, as opposed to the quadratic number in SDP2. We remark that the approach in [60] is applicable more generally to many other problems, including several in game theory.

The scalar ϵ reported in Table A.8 is computed using the strategies (x, y) extracted from the first row of the optimal matrix M_1 as described in Section 4.1 of [60]. The scalar ϵ reported in Table A.9 is computed using $x = P1_n$ and $y = P^T 1_m$ from the optimal solution to SDP2 with $\text{Tr}(\mathcal{M})$ as the objective function.

Table A.8: Statistics on ϵ for first level of the hierarchy in [60].

	5×5	10×5	10×10	15×10	15×15	20×15	20×20
Max	0.3357	0.3304	0.2557	0.2189	0.1987	0.1837	0.1828
Mean	0.1883	0.1889	0.1513	0.1446	0.1262	0.1217	0.1087
Median	0.1803	0.1865	0.1452	0.1418	0.1271	0.1208	0.1070

Table A.9: Statistics on ϵ for SDP2 with $\text{Tr}(M)$ as the objective function.

	5×5	10×5	10×10	15×10	15×15	20×15	20×20
Max	0.1581	0.1589	0.115	0.1335	0.0878	0.082	0.0619
Mean	0.0219	0.0332	0.0405	0.04	0.0366	0.0356	0.0298
Median	0.0046	0.0233	0.036	0.0346	0.0345	0.0325	0.0293

We also ran the second level of the hierarchy in [60] on the same 100 5×5 games. The maximum ϵ observed was .3362, while the mean was .1880 and the median was .1800. The size of the variable matrix that needs to be positive semidefinite for this level is 78×78 .

A.3 Lemmas for Extreme Nash Equilibria

The results reported in Section 5.6 were found using the `1rsnash` [12] software which computes extreme Nash equilibria (see definition below). In particular the true maximum welfare and the persistent strategy sets were found in relation to extreme symmetric Nash equilibria only. We show in this appendix why this is sufficient for the claims we made about *all* symmetric Nash equilibria. We prove a more general statement below about general games and general Nash equilibria since this could be of potential independent interest. The proof for symmetric games is identical once the strategies considered are restricted to be symmetric.

Definition A.3.1. *An extreme Nash equilibrium is a Nash equilibrium which cannot be expressed as a convex combination of other Nash equilibria.*

Lemma A.3.2. *All Nash equilibria are convex combinations of extreme Nash equilibria.*

Proof. It suffices to show that any extreme point of the convex hull of the set of Nash equilibria must be an extreme Nash equilibrium, as any point in a compact convex set can be written as a convex combination of its extreme points. Note that this convex hull contains three types of points: extreme Nash equilibria, Nash equilibria which are not extreme, and convex combinations of Nash equilibria which are not Nash equilibria. The claim then follows because any extreme point of the convex hull cannot be of the second or third type, as they can be written as convex combinations of other points in the hull. \square

The next lemma shows that checking extreme Nash equilibria are sufficient for the maximum welfare problem.

Lemma A.3.3. *For any bimatrix game, there exists an extreme Nash equilibrium giving the maximum welfare among all Nash equilibria.*

Proof. Consider any Nash equilibrium (\tilde{x}, \tilde{y}) , and let it be written as $\begin{bmatrix} \tilde{x} \\ \tilde{y} \end{bmatrix} = \sum_{i=1}^r \lambda_i \begin{bmatrix} x^i \\ y^i \end{bmatrix}$ for some set of extreme Nash equilibria $\begin{bmatrix} x^1 \\ y^1 \end{bmatrix}, \dots, \begin{bmatrix} x^r \\ y^r \end{bmatrix}$ and $\lambda \in \Delta_r$. Observe that for any i, j ,

$$x^{iT} A y^j \leq x^{jT} A y^j, x^{iT} B y^j \leq x^{iT} B y^i, \quad (\text{A.1})$$

from the definition of a Nash equilibrium. Now note that

$$\begin{aligned} \tilde{x}^T (A + B) \tilde{y} &= \left(\sum_{i=1}^r \lambda_i x^i \right)^T (A + B) \left(\sum_{i=1}^r \lambda_i y^i \right) \\ &= \sum_{i=1}^r \sum_{j=1}^r \lambda_i \lambda_j x^{iT} (A + B) y^j \\ &= \sum_{i=1}^r \sum_{j=1}^r \lambda_i \lambda_j x^{iT} A y^j + \sum_{i=1}^r \sum_{j=1}^r \lambda_i \lambda_j x^{iT} B y^j \\ &\stackrel{(\text{A.1})}{\leq} \sum_{i=1}^r \sum_{j=1}^r \lambda_i \lambda_j x^{iT} A y^j + \sum_{i=1}^r \sum_{j=1}^r \lambda_i \lambda_j x^{iT} B y^i \\ &= \sum_{i=1}^r \lambda_i x^{iT} A y^i + \sum_{i=1}^r \lambda_i x^{iT} B y^i \\ &= \sum_{i=1}^r \lambda_i x^{iT} (A + B) y^i. \end{aligned}$$

In particular, since each (x^i, y^i) is an extreme Nash equilibrium, this tells us for any Nash equilibrium (\tilde{x}, \tilde{y}) there must be an extreme Nash equilibrium which has at least as much welfare. \square

Similarly for the results for persistent sets in Section 5.6.2, there is no loss in restricting attention to extreme Nash equilibria.

Lemma A.3.4. *For a given strategy set \mathcal{S} , if every extreme Nash equilibrium plays at least one strategy in \mathcal{S} with positive probability, then every Nash equilibrium plays at least one strategy in \mathcal{S} with positive probability.*

Proof. Let \mathcal{S} be a persistent set of strategies. Since all Nash equilibria are composed of nonnegative entries, and every extreme Nash equilibrium has positive probability on some entry in \mathcal{S} , any convex combination of extreme Nash equilibria must have positive probability on some entry in \mathcal{S} . \square

A.4 A Note on Reductions from General Games to Symmetric Games

An anonymous referee asked us if our guarantees for symmetric games transfer over to general games by symmetrization. Indeed, there are reductions in the literature that take a general game, construct a symmetric game from it, and relate the Nash equilibria of the original game to symmetric Nash equilibria of its symmetrized version. In this Appendix, we review two well-known reductions of this type, which are shown in [45] and [54], and show that the quality of approximate Nash equilibria can differ greatly between the two games. We hope that our examples can be of independent interest.

A.4.1 The Reduction of [45]

Consider a game (A, B) with $A, B > 0$ and a Nash equilibrium (x^*, y^*) of it with payoffs $p_A := x^{*T} A y^*$ and $p_B := x^{*T} B y^*$. Then the symmetric game (S_{AB}, S_{AB}^T) with

$$S_{AB} := \begin{bmatrix} 0 & A \\ B^T & 0 \end{bmatrix}$$

admits a symmetric Nash equilibrium in which both players play $\begin{bmatrix} \frac{p_A}{p_A+p_B} x^* \\ \frac{p_B}{p_A+p_B} y^* \end{bmatrix}$. In the reverse

direction, any symmetric equilibrium $\left(\begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix} \right)$ of (S_{AB}, S_{AB}^T) yields a Nash equilibrium $(\frac{x}{1_m^T}, \frac{y}{1_n^T})$ to the original game (A, B) .

To demonstrate that high-quality approximate Nash equilibria in the symmetrized game can map to low-quality approximate Nash equilibria in the original game, consider the game given by $(A, B) = \left(\begin{bmatrix} \epsilon & 0 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} \epsilon^2 & 0 \\ 0 & 1 \end{bmatrix} \right)$ for some $\epsilon > 0$. The symmetric strategy

$$\left(\begin{bmatrix} \frac{1}{1+\epsilon} \\ 0 \\ \frac{\epsilon}{1+\epsilon} \\ 0 \end{bmatrix}, \begin{bmatrix} \frac{1}{1+\epsilon} \\ 0 \\ \frac{\epsilon}{1+\epsilon} \\ 0 \end{bmatrix} \right)$$

is an $\epsilon \frac{1-\epsilon}{1+\epsilon}$ -NE for (S_{AB}, S_{AB}^T) , but the strategy pair $\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right)$ is a $(1-\epsilon)$ -NE for (A, B) .

A.4.2 The Reduction of [54]

Consider a game (A, B) with $A > 0, B < 0$ and a Nash equilibrium (x^*, y^*) of it with payoffs $p_A := x^{*T} A y^*$ and $p_B := x^{*T} B y^*$. Then the symmetric game (S_{AB}, S_{AB}^T) with

$$S_{AB} := \begin{bmatrix} 0_{m \times m} & A & -1_m \\ B^T & 0_{n \times n} & 1_n \\ 1_m^T & -1_n^T & 0 \end{bmatrix}$$

admits a symmetric Nash equilibrium in which both players play

$$\begin{bmatrix} \frac{x^*}{2-p_B} \\ \frac{y^*}{2+p_A} \\ 1 - \frac{1}{2-p_B} - \frac{1}{2+p_A} \end{bmatrix}.$$

In the reverse direction, any symmetric equilibrium $\left(\begin{bmatrix} x \\ y \\ z \end{bmatrix}, \begin{bmatrix} x \\ y \\ z \end{bmatrix} \right)$ of (S_{AB}, S_{AB}^T) yields a Nash equilibrium $(\frac{x}{1_m^T}, \frac{y}{1_n^T})$ to the original game (A, B) . This reduction has some advantages over the previous one (see [54, Section 1]).

To demonstrate that high-quality approximate Nash equilibria in the new symmetrized game can again map to low-quality approximate Nash equilibria in the original game, consider the game given by $(A, B) = \left(\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 0 & 0 \end{bmatrix} \right)$. Let $\epsilon \in (0, \frac{1}{2})$. The symmetric strategy

$$\left(\begin{bmatrix} \epsilon \\ 0 \\ 1 - \epsilon \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \epsilon \\ 0 \\ 1 - \epsilon \\ 0 \\ 0 \end{bmatrix} \right)$$

is an $\frac{\epsilon}{2}(1 - \epsilon)$ -NE¹ for (S_{AB}, S_{AB}^T) , but the strategy pair $\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right)$ is a 1-NE for (A, B) .

¹Note that approximation factor is halved since the range of the entries of the payoff matrix in the symmetrized game is $[-1, 1]$.

Bibliography

- [1] Ilan Adler. The equivalence of linear programs and zero-sum games. *International Journal of Game Theory*, 42(1):165–177, 2013.
- [2] Ilan Adler, Constantinos Daskalakis, and Christos H Papadimitriou. A note on strictly competitive games. In *International Workshop on Internet and Network Economics*, pages 471–474. Springer, 2009.
- [3] A. A. Ahmadi, A. Olshevsky, P. A. Parrilo, and J. N. Tsitsiklis. NP-hardness of deciding convexity of quartic polynomials and related problems. *Mathematical Programming*, 137(1-2):453–476, 2013.
- [4] Amir Ali Ahmadi and Georgina Hall. DC decomposition of nonconvex polynomials with algebraic techniques. *Mathematical Programming*, pages 1–26, 2015.
- [5] Amir Ali Ahmadi and Georgina Hall. On the complexity of detecting convexity over a box. *Mathematical Programming*, pages 1–15, 2019.
- [6] Amir Ali Ahmadi and Anirudha Majumdar. Some applications of polynomial optimization in operations research and real-time decision making. *Optimization Letters*, 10(4):709–729, 2016.
- [7] Amir Ali Ahmadi and Jeffrey Zhang. Complexity aspects of local minima and related notions. *In Preparation*.
- [8] Farid Alizadeh. Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization*, 5(1):13–51, 1995.
- [9] Animashree Anandkumar and Rong Ge. Efficient approaches for escaping higher order saddle points in non-convex optimization. In *Conference on learning theory*, pages 81–102, 2016.
- [10] VG Andronov, EG Belousov, and VM Shironin. On solvability of the problem of polynomial programming. *Izvestija Akadem. Nauk SSSR, Tekhnicheskaja Kibernetika*, 4:194–197, 1982.
- [11] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

- [12] David Avis, Gabriel D Rosenberg, Rahul Savani, and Bernhard Von Stengel. Enumeration of Nash equilibria for two-player games. *Economic Theory*, 42(1):9–37, 2010.
- [13] Tomáš Bajbar and Sönke Behrends. How fast do coercive polynomials grow? Technical report, Institut für Numerische und Angewandte Mathematik, Georg-August-Universität Göttingen, 2017.
- [14] Tomas Bajbar and Oliver Stein. Coercive polynomials and their Newton polytopes. *SIAM Journal on Optimization*, 25(3):1542–1570, 2015.
- [15] Erwin H Bareiss. Sylvester’s identity and multistep integer-preserving Gaussian elimination. *Mathematics of computation*, 22(103):565–578, 1968.
- [16] Saugata Basu and Marie-Françoise Roy. Bounding the radii of balls meeting every connected component of semi-algebraic sets. *Journal of Symbolic Computation*, 45(12):1270–1279, 2010.
- [17] Evgeny G Belousov and Diethard Klatte. A Frank–Wolfe type theorem for convex polynomial programs. *Computational Optimization and Applications*, 22(1):37–48, 2002.
- [18] Abraham Berman and Naomi Shaked-Monderer. *Completely positive matrices*. World Scientific, 2003.
- [19] Dimitri P Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [20] Dimitri P Bertsekas and Paul Tseng. Set intersection theorems and existence of optimal solutions. *Mathematical Programming*, 110(2):287–314, 2007.
- [21] Grigoriy Blekherman, Pablo A Parrilo, and Rekha R Thomas. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM, 2012.
- [22] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [23] Stephen P Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. SIAM, 1994.
- [24] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*, volume 6, page 47th, 2006.
- [25] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS’06)*, pages 603–612. IEEE, 2006.
- [26] Joel E Cohen and Uriel G Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993.
- [27] Vincent Conitzer and Tuomas Sandholm. Complexity results about nash equilibria. *arXiv preprint cs/0205074*, 2002.

- [28] Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.
- [29] George B Dantzig. A proof of the equivalence of the programming problem and the game problem. *Activity analysis of production and allocation*, 13:330–338, 1951.
- [30] Constantinos Daskalakis. On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms (TALG)*, 9(3):23, 2013.
- [31] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [32] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. A note on approximate Nash equilibria. In *International Workshop on Internet and Network Economics*, pages 297–306. Springer, 2006.
- [33] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. Progress in approximate Nash equilibria. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 355–358. ACM, 2007.
- [34] Etienne De Klerk. *Aspects of semidefinite programming: interior point algorithms and selected applications*, volume 65. Springer Science & Business Media, 2006.
- [35] Etienne De Klerk and Dmitrii V Pasechnik. Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002.
- [36] Maryam Fazel. *Matrix rank minimization with applications*. PhD thesis, PhD thesis, Stanford University, 2002.
- [37] John Fearnley, Paul W Goldberg, Rahul Savani, and Troels Bjerre Sørensen. Approximate well-supported Nash equilibria below two-thirds. *Algorithmica*, 76(2):297–319, 2016.
- [38] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics (NRL)*, 3(1-2):95–110, 1956.
- [39] Michael R Garey and David S Johnson. *Computers and Intractability*, volume 29. WH Freeman New York, 2002.
- [40] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [41] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [42] Evgenii Alekseevich Gorin. Asymptotic properties of polynomials and algebraic functions of several variables. *Russian mathematical surveys*, 16(1):93–119, 1961.

- [43] Aurélien Greuet and Mohab Safey El Din. Deciding reachability of the infimum of a multivariate polynomial. In *Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation*, pages 131–138. ACM, 2011.
- [44] Aurélien Greuet and Mohab Safey El Din. Probabilistic algorithm for polynomial optimization over a real algebraic set. *SIAM Journal on Optimization*, 24(3):1313–1343, 2014.
- [45] JH Griesmer, AJ Hoffman, and A Robinson. On symmetric bimatrix games. *IBM Research Paper RC-959. IBM Corp, Thomas J Watson Research Center, Yorktown Heights, New York*, 1963.
- [46] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2. Springer Science & Business Media, 2012.
- [47] J William Helton and Jiawang Nie. Semidefinite representation of convex sets. *Mathematical Programming*, 122(1):21–64, 2010.
- [48] Lane A Hemaspaandra and Ryan Williams. SIGACT News Complexity Theory Column 76: An atypical survey of typical-case heuristic algorithms. *ACM SIGACT News*, 43(4):70–89, 2012.
- [49] David Hilbert. Über die darstellung definiter formen als summe von formenquadraten. *Mathematische Annalen*, 32(3):342–350, 1888.
- [50] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- [51] M Huneault and FD Galiana. A survey of the optimal power flow literature. *IEEE Transactions on Power Systems*, 6(2):762–770, 1991.
- [52] Soichi Ibaraki and Masayoshi Tomizuka. Rank minimization approach for solving BMI problems with random search. In *American Control Conference, 2001. Proceedings of the 2001*, volume 3, pages 1870–1875. IEEE, 2001.
- [53] Vaithilingam Jeyakumar, Jean B Lasserre, and Guoyin Li. On polynomial optimization over non-compact semi-algebraic sets. *Journal of Optimization Theory and Applications*, 163(3):707–718, 2014.
- [54] AP Jurg, MJM Jansen, Jos AM Potters, and SH Tijs. A symmetrization for finite two-person games. *Zeitschrift für Operations Research*, 36(2):111–123, 1992.
- [55] Vassilis Kalofolias and Efstratios Gallopoulos. Computing symmetric nonnegative rank factorizations. *Linear Algebra and its Applications*, 436(2):421–435, 2012.
- [56] Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pages 302–311, 1984.

- [57] Leonid Genrikhovich Khachiyan. A polynomial algorithm in linear programming. In *Doklady Akademii Nauk*, volume 244, pages 1093–1096. Russian Academy of Sciences, 1979.
- [58] Spyros C Kontogiannis, Panagiota N Panagopoulou, and Paul G Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *International Workshop on Internet and Network Economics*, pages 286–296. Springer, 2006.
- [59] Jean-Louis Krivine. Anneaux préordonnés. *Journal d'analyse mathématique*, 12(1):307–326, 1964.
- [60] Rida Laraki and Jean B Lasserre. Semidefinite programming for min–max problems and games. *Mathematical Programming*, 131(1-2):305–332, 2012.
- [61] Jean B Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [62] Jean B Lasserre. Convex sets with semidefinite representation. *Mathematical programming*, 120(2):457–477, 2009.
- [63] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, pages 157–270. Springer, 2009.
- [64] Carlton E Lemke and Joseph T Howson, Jr. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423, 1964.
- [65] Eliane Maria Loiola, Nair Maria Maia de Abreu, Paulo Oswaldo Boaventura-Netto, Peter Hahn, and Tania Querido. A survey for the quadratic assignment problem. *European journal of operational research*, 176(2):657–690, 2007.
- [66] Henri Lombardi, Daniel Perrucci, and Marie-Françoise Roy. An elementary recursive bound for effective Positivstellensatz and Hilbert’s 17th problem. *Available at arXiv:1404.2338*, 2014.
- [67] László Lovász. On the shannon capacity of a graph. *IEEE Transactions on Information theory*, 25(1):1–7, 1979.
- [68] Z.-Q. Luo and S. Zhang. On extensions of the Frank-Wolfe theorems. *Computational Optimization and Applications*, 13(1-3):87–110, 1999.
- [69] Murray Marshall. Optimization of polynomial functions. *Canadian Mathematical Bulletin*, 46(4):575–587, 2003.
- [70] Garth P McCormick. Computability of global solutions to factorable nonconvex programs: Part I : Convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976.
- [71] Carl D Meyer. *Matrix analysis and applied linear algebra*, volume 2. SIAM, 2000.

- [72] Jorge J Moré and Stephen A Vavasis. On the solution of concave knapsack problems. *Mathematical programming*, 49(1-3):397–411, 1990.
- [73] MOSEK. *MOSEK reference manual*, 2013. Version 7. Latest version available at <http://www.mosek.com/>.
- [74] Theodore S Motzkin and Ernst G Straus. Maxima for graphs and a new proof of a theorem of turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.
- [75] Katta G Murty and Santosh N Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.
- [76] Katta G Murty and Feng-Tien Yu. *Linear Complementarity, Linear and Nonlinear Programming*, volume 3. Berlin: Heldermann, 1988.
- [77] John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- [78] Yurii Nesterov. Implementable tensor methods in unconstrained convex optimization. *Mathematical Programming*, pages 1–27, 2019.
- [79] Yurii Nesterov et al. Random walk in a simplex and quadratic optimization over convex polytopes. Technical report, CORE, 2003.
- [80] Yurii Nesterov and Arkadii Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*, volume 13. SIAM, 1994.
- [81] Tim Netzer. On semidefinite representations of non-closed sets. *Linear algebra and its applications*, 432(12):3072–3078, 2010.
- [82] Jiawang Nie. Optimality conditions and finite convergence of the Lasserre hierarchy. *Mathematical programming*, 146(1-2):97–121, 2014.
- [83] Jiawang Nie. The hierarchy of local minimums in polynomial optimization. *Mathematical Programming*, 151(2):555–583, 2015.
- [84] Jiawang Nie, James Demmel, and Bernd Sturmfels. Minimizing polynomials via sum of squares over the gradient ideal. *Mathematical Programming*, 106(3):587–606, 2006.
- [85] Panos M Pardalos and Georg Schnitger. Checking local optimality in constrained quadratic programming is np-hard. *Operations Research Letters*, 7(1):33–35, 1988.
- [86] Panos M Pardalos and Stephen A Vavasis. Quadratic programming with one negative eigenvalue is np-hard. *Journal of Global optimization*, 1(1):15–22, 1991.
- [87] Panos M Pardalos and Stephen A Vavasis. Open questions in complexity theory for numerical optimization. *Mathematical Programming*, 57(1-3):337–339, 1992.
- [88] Pablo A Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, 2003.

- [89] Pablo A Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2855–2860. IEEE, 2006.
- [90] L. Porkolab and L. Khachiyan. On the complexity of semidefinite programs. *Journal of Global Optimization*, 10(4):351–365, 1997.
- [91] Alexander Prestel and Charles N Delzell. Positive Polynomials: from Hilbert’s 17th Problem to Real Algebra. *Springer Monographs in Mathematics*. Springer, Berlin, Germany, 2001.
- [92] Mihai Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.
- [93] Motakuri Ramana and AJ Goldman. Some geometric results in semidefinite programming. *Journal of Global Optimization*, 7(1):33–50, 1995.
- [94] Motakuri V Ramana. An exact duality theory for semidefinite programming and its complexity implications. *Mathematical Programming*, 77(1):129–162, 1997.
- [95] Motakuri Venkata Ramana. *An algorithmic analysis of multiquadratic and semidefinite programming problems*. PhD thesis, Citeseer, 1993.
- [96] Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [97] R Tyrrell Rockafellar. *Convex Analysis*, volume 28. Princeton University Press, 1970.
- [98] Rahul Savani and Bernhard Stengel. Hard-to-solve bimatrix games. *Econometrica*, 74(2):397–429, 2006.
- [99] Thomas J Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226. ACM, 1978.
- [100] A. Seidenberg. A new decision method for elementary algebra. *Annals of Mathematics*, pages 365–374, 1954.
- [101] Parikshit Shah and Pablo A Parrilo. Polynomial stochastic games via sum of squares optimization. In *Decision and Control, 2007 46th IEEE Conference on*, pages 745–750. IEEE, 2007.
- [102] Noah D Stein. *Exchangeable equilibria*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [103] Gilbert Stengle. A Nullstellensatz and a Positivstellensatz in semialgebraic geometry. *Mathematische Annalen*, 207(2):87–97, 1974.
- [104] Johan AK Suykens and Joos Vandewalle. Least squares support vector machine classifiers. *Neural processing letters*, 9(3):293–300, 1999.

- [105] Alfred Tarski. A decision method for elementary algebra and geometry. In *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pages 24–84. Springer, 1998.
- [106] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [107] Haralampos Tsaknakis and Paul G Spirakis. An optimization approach for approximate Nash equilibria. In *International Workshop on Web and Internet Economics*, pages 42–56. Springer, 2007.
- [108] Lieven Vandenberghhe and Stephen Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [109] Stephen A Vavasis. Quadratic programming is in np. *Information Processing Letters*, 36(2):73–77, 1990.
- [110] Cynthia Vinzant. What is... a spectrahedron? *Notices Amer. Math. Soc*, 61(5):492–494, 2014.
- [111] Sven Wagner. *Archimedean quadratic modules: a decision problem for real multivariate polynomials*. PhD thesis, University of Konstanz, 2009.