Simple Complexity from Imitation Games

Andrew McLennan* and Rabee Tourky†

February 21, 2005

Abstract

We give simple proofs of refinements of the complexity results of Gilboa and Zemel (1989), and we derive additional results of this sort. Our constructions employ imitation games, which are two person games in which both players have the same sets of pure strategies and the second player wishes to play the same pure strategy as the first player.

1 Introduction

NASH is the problem of computing a Nash equilibrium of a two person finite game. Whether NASH has a polynomial-time algorithm is not known, and this problem has recently been described by Papadimitriou (2001) as (along with factoring) "the most important concrete open question on the boundary of **P** today." Gilboa and Zemel (1989) have shown that several closely related problems are **NP**-complete.

An *imitation game* is a finite two player game in which the two agents' sets of pure strategies are the same and the second player's payoff is either 1 or 0 according to whether she chooses the same pure strategy as the first player. McLennan and Tourky (2004) recently initiated the study of these games. Among their findings are results showing that, in spite of the simple structure of imitation games, they capture the computational complexity of general two player games. This paper provides simple proofs of slight refinements of the results of Gilboa and Zemel (1989) by constructing imitation games whose Nash equilibria mirror a known **NP**-complete problem. In addition to strengthening those results and simplifying their presentation, the arguments given here provide further confirmation of the usefulness of imitation games in the study of complexity issues related to two player games.

We now review some relevant concepts of theoretical computer science. A *decision problem* is a computational task whose desired output is either YES or NO. Such a problem is in the class **NP** if there is a "nondeterministic" computational procedure whose running time is bounded by a polynomial function of the size of the input and which can verify that the desired ouput is YES when that is the case. Roughly, a nondeterministic computation is modelled by a Turing machine for which some states have two successor states, so that for

*Department of Economics, The University of Minnesota, 1035 Management and Economics, 271 19th Ave. S. Minneapolis, MN 55455. The hospitality of the Economic Theory Center of the University of Melbourne is gratefully acknowledged.

†Department of Economics, The University of Melbourne, Victoria 3010, Australia.

each input a tree of states is generated. One may imagine a technology in which all paths to leafs of the tree are followed simultaneously. Alternatively, one may imagine that the successor node is chosen randomly at states where branching occurs, in which case (provided that both allowed successor states always have positive probability) there is a positive probability of following a branch that verifies that the desired output is YES.

An example of such a problem is CLIQUE: given a simple undirected graph[1] $G = (V, E)$ and a positive integer k, determine whether G contains a clique of size k. The process of determining whether a set of k vertices is a clique is a computation whose running time is bounded by a polynomial function of the size of G (more specifically, it requires $k(k-1)/2$ queries) so generating such a set randomly has a positive probability of producing a *certificate* that the correct output is YES when this is the case. In fact any problem in **NP** has a checking procedure that allows one to pass from a certificate (for instance suitable branching choices of a nondeterministic algorithm) to a verification that that the desired answer is YES, when this is the case, in polynomial time. The existence of such a checking procedure implies that a problem is in **NP**, since guessing an element of the space of certificates is a suitable nondeteministic procedure, and in fact **NP** is most commonly defined in terms of such checking procedures.

The class **P** of decision problems that have polynomial time algorithms is a subclass of **NP**. (Deterministic computational procedures are a subclass of the class of nondeterministic procedures.) It is generally believed that the containment is strict, but this has not been proved, and is a preeminent open problem, not just of computer science, but of all of mathematics.

A decision problem P is **NP**-*hard* if, for any problem Q in **NP**, there is a polynomial-time procedure for transforming an instance of Q into an instance of P that has the same desired output. If P is both **NP**-hard and in **NP** itself, then P is said to be ***NP**-complete*. A polynomial-time algorithm for any **NP**-hard problem would prove that **NP** = **P**, but of course no such algorithm has ever been found. CLIQUE is one of the many problems that are known to be **NP**-complete.

For each problem P in **NP** there is a related problem for which the desired output is either NO, if that is the correct output for P, or a certificate allowing polynomial time verification that the correct output for P is YES. The class of such derived problems is **FNP**. The subclass **TFNP** (Megiddo and Papadimitriou (1991), Papadimitriou (1994)) is the class of problems in **FNP** for which the desired output of the associated decision problem is always YES. An important example is the problem of factoring an integer as a product of primes. (The associated decision problem is "determine whether the integer has a prime factorization.")

The *support* of a mixed strategy profile is the set of pure strategies that are assigned positive probability. The problem of determining whether a two player game has a Nash equilibrium with a particular support is in **P**, being a straightforward matter of linear algebra. The decision problem "does the game

[1]A *simple undirected graph* is a pair $G = (V, E)$ in which V is a finite set of *vertices* and E is a set of unordered pairs of elements of distinct elements of V. Elements of E are called *edges*. If u and v are distinct elements of V, the corresponding unordered pair is denoted by uv or vu. If $uv \in E$, then u and v are said to be *neighbors*. A set $C \subset V$ of vertices is a *clique* if $uv \in E$ for all distinct $u, v \in C$.

have a Nash equilibrium" is consequently in **NP**, since guessing a support has a positive probability of producing an equilibrium in polynomial time, and the fact that every finite game has a Nash equilibrium implies that NASH is in **TFNP**.

Megiddo and Papadimitriou (1991) and Papadimitriou (1994) argue that **TFNP** is unlikely to contain any problems that are complete for **TFNP**. Thus the most promising approach to the study of **TFNP** would seem to be to identify problems that are representative of important subclasses. NASH is one such problem, since it is a simple instance of a large class of problems that are special cases or reformulations of the problem of computing a fixed point under the hypotheses of the Kakutani fixed point theorem. (The arguments in McLennan and Tourky (2004) may be construed as evidence that these problems are "not much more complex" than NASH.)

Although Gilboa and Zemel (1989) do not resolve the status of NASH, they provide a host of **NP**-hard problems that are (at least on superficial inspection) quite similar to NASH, suggesting that the "distance" from NASH to **NP**-complete problems is less than the distance from **P** to NASH. Their results are established by constructing games with one "trivial equilibrium" whose other equilibria mirror a known **NP**-complete problem.

Recently Conitzer and Sandholm (2003), Blum and Toth (2004), and Codenotti and Štefanovič (2004) have provided alternative proofs of these results by constructing games whose equilibria mirror SAT, which is the following decision problem: given a compound proposition in conjunctive-disjunctive form—that is, a conjunction of disjunctions of primitive propositions $P, Q, \dots, R$ and their negations—does there exist a vector of truth values for $P, Q, \dots, R$ such that the compound proposition is true. For example, $(P \vee Q) \wedge (\neg Q \vee R) \wedge (\neg P \vee \neg R)$ is satisfied by the vector of truth values (T, F, F). Like CLIQUE, SAT is **NP**-complete. Roughly, their constructions pass from a proposition in conjunctive-disjunctive form to a two player game in which there is a trivial Nash equilibrium, and other equilibria correspond to vectors of truth values satisfying the formula.

In Gilboa and Zemel (1989), and in the recent papers using reduction to SAT, the games are surprisingly complicated, with detailed strategic incentives whose implications are not obvious. The purpose of this note is to give simple proofs of similar results using imitation games. As in Gilboa and Zemel, we begin with an instance of CLIQUE and construct a game whose equilibria correspond to maximal cliques. Relative to their constructions, ours are simple and direct.

Blum and Toth (2004) point out that the games employed by Gilboa and Zemel (1989) and Conitzer and Sandholm (2003) have equilibria that are not *regular* in the sense of Harsanyi (1973). That is, the equilibrium either fails to be *strict* (some pure strategy is assigned probability zero even though it yields the equilibrium expected utility) or the system of equations expressing indifference, for each agent, between all pure strategies in the support of the agent's mixed strategy, has a singular Jacobean. The games used by Blum and Toth (2004) have only regular equilibria, so they show that the **NP**-completeness results hold even if attention is restricted to games that are non-degenerate in this sense. We do not know whether our games are non-degenerate, but the equilibria that figure in our argument are regular. Non-degenerate games are dense (Harsanyi (1973)) and regular equilibria are preserved under small perturbations of the payoffs, so our results also confirm the Blum-Toth result because games near the ones we use are non-degenerate and have equilibria of the sort required by the analysis.

As with the Blum-Toth proof, the other recent proofs refine the Gilboa-Zemel result by showing that the problems remain **NP**-complete if attention is restricted to games with particular features. Conitzer and Sandholm (2003) emphasize that their games are symmetric, and the payoffs in the games employed by Codenotti and Štefanovič (2004) are all either 0 or 1. The statement of our main result, Theorem 3, emphasizes that the imitation game aspect, but it will be evident that the games are quite special even within that class.

In the next section we describe a relationship between Nash equilibria of imitation games and symmetric equilibria of symmetric games. Section 3 describes a relationship between the Nash equilibria of an imitation game and the stationary points (Karush–Kuhn–Tucker points) of an associated quadratic programming problem. It also introduces an imitation game based on the elegant Motzkin-Strauss formulation of the maximum clique problem as a quadratic program (see Motzkin and Straus (1965) and more recently Bomze (1998)) and relates the equilibria of this game to CLIQUE, and to the quadratic programming problem. Section 4 establishes the complexity results.

2 Nash Equilibria of Imitation Games

For an integer $k \geq 1$ let

$$\Delta^k := \{\, \rho \in \mathbb{R}^k_{\geq} : \rho_1 + \cdots + \rho_k = 1 \,\}$$

be the $(k-1)$-dimensional simplex. For $\rho \in \Delta^k$ let

$$\operatorname{supp} \rho := \{\, i = 1, \ldots, k : \rho_i > 0 \,\}$$

be the *support* of ρ. For $S \subset \{1, \ldots, k\}$ let $\beta_S \in \Delta^k$ be the *uniform distribution* on S: the i^{th} component of β_S is $1/|S|$ or 0 according to whether $i \in S$. For $i = 1, \ldots, k$ let $\mathbf{e}_i$ be the i^{th} standard unit basis vector. Usually $\mathbf{e}_i$ will occur in its role as $\beta_{\{i\}}$, but it will also have other applications.

We represent a two person game with a pair (M, N) of $m \times n$ matrices with real entries. A *Nash equilibrium* of (M, N) is a pair $(\sigma, \tau) \in \Delta^m \times \Delta^n$ such that $\sigma^T M \tau \geq \tilde{\sigma}^T M \tau$ for all $\tilde{\sigma} \in \Delta^m$ and $\sigma^T N \tau \geq \sigma^T N \tilde{\tau}$ for all $\tilde{\tau} \in \Delta^n$.

The game (M, N) is an *imitation game* if $n = m$ and N is the $m \times m$ identity matrix I, in which case the row player is called the *mover* and the column player is called the *imitator*. An *I-equilibrium* of (M, I) is a $\tau \in \Delta^m$ such that

$$\operatorname{supp} \tau \subset \operatorname{argmax}_{i=1,\ldots,m} \mathbf{e}_i^T M \tau.$$

The game (M, N) is *symmetric* if $n = m$ and $N = M^T$. A *symmetric equilibrium* of (M, M^T) is a $\tau \in \Delta^m$ such that (τ, τ) is a Nash equilibrium.

The definition of an I-equilibrium is relatively simple, involving half of the conditions involved in the definition of Nash equilibrium. The following result shows that this concept is an adequate proxy in connection with an imitation game. In addition, it will imply that each of our complexity results has an immediate corollary concerning symmetric equilibria of symmetric games.

Proposition 1. *For an $m \times m$ matrix M and $\tau \in \Delta^m$ the following are equivalent:*

(a) τ is an I-equilibrium of the imitation game (M, I);

(b) $(\beta_{\operatorname{supp}\tau}, \tau)$ is a Nash equilibrium of the imitation game (M, I);

(c) τ is a symmetric equilibrium of the symmetric game (M, M^T).

Proof. If τ is an I-equilibrium of (M, I), then, by virtue of the definition, $\beta_{\operatorname{supp}\tau}$ is a best response to τ for the mover, and τ is clearly a best response to $\beta_{\operatorname{supp}\tau}$ for the imitator. Thus (a) implies (b). If $\beta_{\operatorname{supp}\tau}$ is a best response to τ for the mover in (M, I), then τ is a best response to itself for the first agent in (M, M^T), and, by symmetry, for the second agent as well. Thus (b) implies (c). If τ is a symmetric equilibrium of (M, M^T) then, because τ is a best response for the first agent, τ is an I-equilibrium of (M, I). Thus (c) implies (a). ∎

3 Quadratic Programming

A *stationary point* of the quadratic programming problem $\max_{\tau\in\Delta^m} \tau^T M\tau$ is a $\tau \in \Delta^m$ such that

$$0 \geq \frac{d}{dt}[(\tau + t\nu)^T M(\tau + t\nu)]\big|_{t=0}$$

for all $\nu \in \mathbb{R}^m$ such that $\tau + t\nu \in \Delta^m$ for small $t > 0$, i.e., $\nu_1 + \cdots + \nu_m = 0$ and $\nu_i \geq 0$ for each i such that $\tau_i = 0$.

Proposition 2. *For a symmetric $m \times m$ matrix M, τ is a stationary point of the quadratic programming problem $\max_{\tau\in\Delta^m} \tau^T M\tau$ if and only if it is an I-equilibrium of the imitation game (M, I).*

Proof. Since M is symmetric, for any $\nu \in \mathbb{R}^m$ we have

$$\frac{d}{dt}[(\tau + t\nu)^T M(\tau + t\nu)]\big|_{t=0} = \nu^T M\tau + \tau^T M\nu = 2\nu^T M\tau.$$

If τ is an I-equilibrium of (M, I), then it is a symmetric equilibrium of the symmetric game (M, M), so $\tau^T M\tau = \max_{\sigma\in\Delta^m} \sigma^T M\tau$ and, for any ν as above, $\tau^T M\tau \geq (\tau + t\nu)^T M\tau$ for small $t \geq 0$, whence $\nu^T M\tau \leq 0$. Conversely, if this inequality holds for all ν as above, then $(\sigma^T - \tau^T)M\tau \leq 0$ for all $\sigma \in \Delta^m$, so that τ is a symmetric equilibrium of (M, M) and an I-equilibrium of (M, I). ∎

We now fix a simple undirected graph $G = (V, E)$ with $V = \{1, \ldots, m\}$. For the remainder of the section M will denote the symmetric $m \times m$ matrix with entries:

$$m_{ij} := \begin{cases} 0, & \text{if } ij \in \mathcal{E}; \\ -1, & \text{if } i = j; \\ -2, & \text{otherwise}. \end{cases}$$

A *maximal clique* for G is a clique $C \subset V$ that is not a proper subset of another clique.

Lemma 1. *If C is a maximal clique, then β_C is an I-equilibrium of the imitation game (M, I).*

Proof. It suffices to show that β_C is a best response to itself in the symmetric game (M, M), hence a symmetric equilibrium. If $i \in C$, then $\mathbf{e}_i^T M \beta_C = -1/|C|$, and if $i \notin C$, then $ij \in E$ for at most $|C| - 1$ vertices $j \in C$, so that $\mathbf{e}_i^T M \beta_C \leq -2/|C|$. ■

When C is a maximal clique, (β_C, β_C) is a regular equilibrium of (M, I). Specifically, strictness for the mover was verified in the proof above, strictness for the imitator is obvious, and one may easily compute that the Jacobean of the system of equations expressing each agent's indifference between all elements of the support of her mixed strategy is a diagonal matrix, hence nonsingular. It will be easy to see that the related equilibria appearing in the proof of Theorem 3 are also regular, so that (as we explained in the Introduction) there are nearby non-degenerate games that also support the analysis.

Lemma 2. *If τ is a local maximum of $\tau^T M \tau$, then $\tau = \beta_C$ for some maximal clique C.*

Proof. Suppose by way of contradiction that $\operatorname{supp} \tau$ is not a clique. Pick $i, j \in \operatorname{supp} \tau$ such that $ij \notin E$, and let $\nu := \mathbf{e}_i - \mathbf{e}_j$. For $\sigma = \tau + t\nu$ we have

$$\sigma^T M \sigma = \tau^T M \tau + 2t\tau^T M \nu + t^2 \nu^T M \nu \, .$$

Since $\nu^T M \nu = 2$, $\sigma^T M \sigma > \tau^T M \tau$ when t has the same sign as $\tau^T M \nu$, contradicting local maximality.

Therefore $\operatorname{supp} \tau$ is a clique. Let C be a maximal clique containing it. Elementary algebra gives

$$\tau^T M \tau = -\sum_{i \in C} \tau_i^2 = -\tfrac{1}{|C|} - \sum_{i \in C} \left(\tau_i - \tfrac{1}{|C|}\right)^2 .$$

We must have $\tau = \beta_C$, since otherwise this quantity could be increased locally by moving along the line segment from τ to β_C. ■

A *maximum clique* of G is a clique that does not have fewer elements than some other clique.

Proposition 3. *The solutions of $\max_{\tau \in \Delta^m} \tau^T M \tau$ are the uniform distributions β_C on maximum cliques C.*

Proof. By compactness the set of solutions is nonempty. It is a subset of the set of local maximizers, which is contained in the set of β_C for maximal cliques C. For any clique C we have $\beta_C^T M \beta_C = -1/|C|$. ■

4 Complexity Results

We are now ready to state and prove our complexity results. Problems (a), (g), (i), (j), and (k) below are specializations to imitation games of problems shown by GZ to be **NP**-complete[2]. Our findings are similar to those of CS insofar as they strengthen several of the results of GZ by showing that the relevant decision problems remain **NP**-complete if their scope is restricted, to symmetric games in the case of GZ and to imitation games (M, I) with symmetric M here.

[2]We have not found a way to use our methods to prove the following result of GZ: deciding, for a given game and integer k, whether there is an equilibrium in which each player uses no more than k pure strategies, is **NP**-complete.

Theorem 3. *The following decisions problems, whose given data are a symmetric matrix M, an integer k, and $\alpha \in \mathbb{R}$, are **NP**-complete:*

(a) **(Max payoff of mover)** *Does the imitation game (M, I) have an equilibrium in which the expected payoff of the mover is greater than or equal to α?*

(b) **(Max payoff of two players)** *Does the imitation game (M, I) have an equilibrium in which the expected payoff of each player is greater than or equal to α?*

(c) **(Min probability of all pure strategies)** *Does the imitation game (M, I) have an I-equilibrium τ in which all probabilities τ_i are less than or equal to α?*

(d) **(Min probability of a given pure strategy)** *Does the imitation game (M, I) have an equilibrium in which a given pure strategy is played with probability less than or equal to α?*

(e) **(Max cardinality of max probability)** *Does the imitation game (M, I) have an equilibrium in which the set of pure strategies played with maximum probability has at least k elements?*

(f) **(Max cardinality of equal probability)** *Does the imitation game (M, I) have an equilibrium in which the set of pure strategies played with the same probability as a given pure strategy has at least k elements?*

(g) **(Uniqueness)** *Does the imitation game (M, I) have more than one Nash equilibrium?*

(h) **(Containing Subset)** *Does the imitation game (M, I) have an equilibrium in which a particular pure strategy is played with positive probability?*

(i) **(In a Subset)** *Does the imitation game (M, I) have an equilibrium in which a particular pure strategy is not played with positive probability?*

(j) **(Mixed Strategy)** *Does the imitation game (M, I) have an equilibrium in which the mover does not play a pure strategy?*

Proof. To show that any one of these problems is in **NP** it suffices to observe that guessing a support (or two supports in the case of (g)) then computing the equilibria with that support, constitutes a nondeterministic polynomial time algorithm that has a positive probability of proving that the answer is YES when that is the case. The remainder of the proof shows that each of these problems is **NP**-hard.

Henceforth we will assume that M is derived from a graph $G = (V, E)$ as described in the last section. Combining Propositions 1, 2, and 3, (M, I) has an I-equilibrium τ in which the mover's expected payoff $\tau^T M \tau$ is at least $-1/k$ if and only if G has a clique with k elements, so the **NP**-hardness of (a) and (b) follows from the **NP**-hardness of CLIQUE.

For the remainder of the proof we will assume that $im \in E$ for all $i = 1, \ldots, m-1$. (Of course the restriction of CLIQUE to graphs satisfying this condition is still **NP**-hard.) Let τ be an I-equilibrium of (M, I). Then $\tau_m > 0$, since otherwise we would have $\mathbf{e}_m^T M\tau = 0 > \tau^T M\tau$. In addition, $\tau_i \leq \tau_m$ for all i, since when $\tau_i > 0$ we have

$$-\tau_i \geq \mathbf{e}_i^T M\tau = \mathbf{e}_m^T M\tau = -\tau_m = \tau^T M\tau.$$

This computation also implies that $\tau_i = \tau_m$ if and only if i is a neighbor of every element of the support of τ. It follows that for this M, $\alpha = 1/k$, and m as the particular pure strategy, each of the problems (c)-(f) is equivalent to whether G has a k-clique. Thus (c)-(f) are **NP**-hard.

We now employ the "trick" of Gilboa and Zemel (1989), which is also used in Conitzer and Sandholm (2003), of appending an extra strategy for each agent. For $\gamma, \delta \in \mathbb{R}$ let

$$M_{\gamma,\delta} := \begin{pmatrix} m_{11} & m_{1m} & \cdots & m_{1m} & -\gamma \\ m_{12} & m_{1m} & \cdots & m_{1m} & -\gamma \\ \vdots & \vdots & & \vdots & -\gamma \\ m_{1m} & m_{2m} & \cdots & m_{mm} & -\gamma \\ -\gamma & -\gamma & \cdots & -\gamma & -\gamma+\delta \end{pmatrix},$$

and let $M_\gamma := M_{\gamma,0}$. Clearly $\mathbf{e}_{m+1}$ is an I-equilibrium of the imitation game (M_γ, I). Suppose that τ is an equilibrium of (M_γ, I) with $1 > \tau_{m+1}$. If $\tau^* \in \mathbb{R}^m$ is the mixed strategy for (M, I) whose i-th component is $\tau_i^* := \frac{\tau_i}{(1-\tau_{m+1})}$, then τ^* is an I-equilibrium of (M, I), and (since m is a best response to τ) $\tau_m^* \leq \gamma$. Conversely, if τ^* is an I-equilibrium of (M, I) with $\tau_m^* \leq \gamma$, then $(\tau^*, 0)$ is an I-equilibrium of (M_γ, I). So, there is an I-equilibrium τ of (M_γ, I) with $\tau_{m+1} < 1$ if and only if there exists an equilibrium τ of (M, I) in which $\tau_m \leq \gamma$. As we saw above, when $\gamma = 1/k$ this will be the case if and only if there is a k-clique. Letting m and $m+1$ be the particular pure strategies in (h) and (i) respectively, the **NP**-hardness of (h)-(k) follows.

This argument also establishes the **NP**-hardness of the problem of deciding whether an imitation game has more that one I-equilibrium, but in the game (M_γ, I) there are many Nash equilibria in which the imitator plays $\mathbf{e}_{m+1}$. To establish that (g) is **NP**-hard one may apply the line of reasoning above to the game $(M_{\gamma,\delta}, I)$ with $-\frac{1}{k-1} < -\gamma < -\gamma + \delta < \frac{1}{k}$. ■

In spite of the similarity between the results of CS and the following consequences (via Proposition 1) of Theorem 3, we do not see a simple sense in which these results are implied by theirs or vice versa. Concretely, questions concerning Nash equilibria of symmetric games seem, on superficial inspection, to be similar to questions concerning symmetric equilibria of symmetric games, but in fact the two sorts of questions are quite distinct.

Corollary 4. *The following decisions problems, whose given data are a symmetric matrix M, an integer k, and $\alpha \in \mathbb{R}$, are NP-complete:*

(a) Does a symmetric game (M, M) have a symmetric equilibrium in which the expected payoff of each player is greater than or equal to α?

(b) Does the symmetric game (M, M) have a symmetric equilibrium τ in which all probabilities τ_i are less than or equal to α?

(c) *Does the symmetric game (M, M) have a symmetric equilibrium in which a particular pure strategy is played with probability less than or equal to α?*

(d) *Does the symmetric game (M, M) have a symmetric equilibrium in which the cardinality of the set of pure strategies played with maximum probability is greater than or equal to k?*

(e) *Does the symmetric game (M, M) have a symmetric equilibrium in which the cardinality of the set of pure strategies played with the same probability as a particular pure strategy is greater than or equal to k?*

(f) *Does the symmetric game (M, M) have more than one symmetric equilibrium?*

(g) *Does the symmetric game (M, M) have a symmetric equilibrium in which a particular pure strategy is played with positive probability?*

(h) *Does the symmetric game (M, M) have a symmetric equilibrium in which a particular pure strategy is not played with positive probability?*

(i) *Does the symmetric game (M, M) have a symmetric equilibrium in which at least two pure strategies are played with positive probability?*

(j) *Does the symmetric game (M, M) have a symmetric equilibrium that is not pure.*

References

A. Blum and B. Toth. NP-hardness involving Nash equilibria in a non-degenerate setting. unpublished, 2004.

I. M. Bomze. On standard quadratic optimization problems. *J. Global Optim.*, 13(4):369–387, 1998. Workshop on Global Optimization (Trier, 1997).

B. Codenotti and D. Štefanovič. On the computational complexity of Nash equilibria for $(0, 1)$ bimatrix games. unpublished, 2004.

V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 765–771, 2003.

I. Gilboa and E. Zemel. Nash and correlated equilibria: some complexity considerations. *Games Econom. Behav.*, 1(1):80–93, 1989.

J. Harsanyi. Oddness of the number of equilibrium points: A new proof. *J. of Game Theory*, 2:235–250, 1973.

A. McLennan and R. Tourky. From imitation games to Kakutani. unpublished, 2004.

N. Megiddo and C. H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theoret. Comput. Sci.*, 81(2, Algorithms Automat. Complexity Games):317–324, 1991.

T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canad. J. Math.*, 17:533–540, 1965.

C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48 (3):498–532, 1994.

C. H. Papadimitriou. Algorithms, games and the internet. In *Annual ACM Symposium on the Theory of Computing*, pages 749–253, 2001.