

The Exact Computational Complexity of Evolutionarily Stable Strategies

Vincent Conitzer

Department of Computer Science
Duke University
Durham, NC, USA
`conitzer@cs.duke.edu`

Abstract. While the computational complexity of many game-theoretic solution concepts, notably Nash equilibrium, has now been settled, the question of determining the exact complexity of computing an evolutionarily stable strategy has resisted solution since attention was drawn to it in 2004. In this paper, I settle this question by proving that deciding the existence of an evolutionarily stable strategy is Σ_2^P -complete.

Key words: Algorithmic game theory, equilibrium computation, evolutionarily stable strategies

1 Introduction

Game theory provides ways of formally representing strategic interactions between multiple players, as well as a variety of *solution concepts* for the resulting games. The best-known solution concept is that of Nash equilibrium [Nash, 1950], where each player plays a best response to all the other players' strategies. The computational complexity of, given a game in strategic form, computing a (any) Nash equilibrium, remained open for a long time and was accorded significant importance [Papadimitriou, 2001]. An elegant algorithm for the two-player case, the Lemke-Howson algorithm [Lemke and Howson, 1964], was proved to require exponential time on some game families [Savani and von Stengel, 2006]. Finally, in a breakthrough series of papers, the problem was established to be PPAD-complete, even in the two-player case [Daskalakis *et al.*, 2009; Chen *et al.*, 2009].¹

Not all Nash equilibria are created equal; for example, one can Pareto-dominate another. Moreover, generally, the set of Nash equilibria does not satisfy *interchangeability*. That is, if player 1 plays her strategy from one Nash equilibrium, and player 2 plays his strategy from another Nash equilibrium, the result is not guaranteed to be a Nash equilibrium. This leads to the dreaded *equilibrium selection* problem: if one plays a game for the first time, how is one to know according to which equilibrium to play? This problem is arguably exacerbated

¹ Depending on the precise formulation, the problem can actually be FIXP-complete for more than 2 players [Etessami and Yannakakis, 2010].

by the fact that determining whether equilibria with particular properties, such as placing probability on a particular pure strategy or having at least a certain level of social welfare, exist is NP-complete in two-player games (and associated optimization problems are inapproximable unless $P=NP$) [Gilboa and Zemel, 1989; Conitzer and Sandholm, 2008]. In any case, equilibria are often seen as a state to which play could reasonably converge, rather than an outcome that can necessarily be arrived at immediately by deduction. Many other solution concepts have been studied from a computational perspective, including refinements of Nash equilibrium [Hansen *et al.*, 2010; Sørensen, 2012], coarsenings of Nash equilibrium (such as correlated equilibrium [Papadimitriou and Roughgarden, 2008; Jiang and Leyton-Brown, 2013] and equilibria of repeated games [Littman and Stone, 2005; Borgs *et al.*, 2010; Kontogiannis and Spirakis, 2008]), and incomparable concepts such as Stackelberg equilibrium [Conitzer and Sandholm, 2006; von Stengel and Zamir, 2010].

In this paper, we consider the concept of *evolutionarily stable strategies*, a solution concept for symmetric games with two players. s will denote a pure strategy and σ a mixed strategy, where $\sigma(s)$ denotes the probability that mixed strategy σ places on pure strategy s . $u(s, s')$ is the utility that a player playing s obtains when playing against a player playing s' , and $u(\sigma, \sigma') = \sum_{s, s'} \sigma(s)\sigma'(s')u(s, s')$ is the natural extension to mixed strategies.

Definition 1 (Price and Smith [1973]) *Given a symmetric two-player game, a mixed strategy σ is said to be an evolutionarily stable strategy (ESS) if both of the following properties hold.*

1. (*symmetric Nash equilibrium property*) For any mixed strategy σ' , we have $u(\sigma, \sigma) \geq u(\sigma', \sigma)$.
2. For any mixed strategy σ' ($\sigma' \neq \sigma$) for which $u(\sigma, \sigma) = u(\sigma', \sigma)$, we have $u(\sigma, \sigma') > u(\sigma', \sigma')$.

The intuition behind this definition is that a population of players playing σ cannot be successfully “invaded” by a small population of players playing some $\sigma' \neq \sigma$, because they will perform *strictly* worse than the players playing σ and therefore they will shrink as a fraction of the population. They perform strictly worse either because (1) $u(\sigma, \sigma) > u(\sigma', \sigma)$, and because σ has dominant presence in the population this outweighs performance against σ' ; or because (2) $u(\sigma, \sigma) = u(\sigma', \sigma)$ so the second-order effect of performance against σ' becomes significant, but in fact σ' performs worse against itself than σ performs against it, that is, $u(\sigma, \sigma') > u(\sigma', \sigma')$.

Example (Hawk-Dove game [Price and Smith, 1973]). Consider the following symmetric two-player game:

	Dove	Hawk
Dove	1,1	0,2
Hawk	2,0	-1,-1

The unique symmetric Nash equilibrium σ of this game is 50% Dove, 50% Hawk. For any σ' , we have $u(\sigma, \sigma) = u(\sigma', \sigma) = 1/2$. That is, everything is a best response to σ . We also have $u(\sigma, \sigma') = 1.5\sigma'(\text{Dove}) - 0.5\sigma'(\text{Hawk}) = 2\sigma'(\text{Dove}) - 0.5$, and $u(\sigma', \sigma') = 1\sigma'(\text{Dove})^2 + 2\sigma'(\text{Hawk})\sigma'(\text{Dove}) + 0\sigma'(\text{Dove})\sigma'(\text{Hawk}) - 1\sigma'(\text{Hawk})^2 = -2\sigma'(\text{Dove})^2 + 4\sigma'(\text{Dove}) - 1$. The difference between the former and the latter expression is $2\sigma'(\text{Dove})^2 - 2\sigma'(\text{Dove}) + 0.5 = 2(\sigma'(\text{Dove}) - 0.5)^2$. The latter is clearly positive for all $\sigma' \neq \sigma$, implying that σ is an ESS.

Intuitively, the problem of computing an ESS appears significantly harder than that of computing a Nash equilibrium, or even a Nash equilibrium with a simple additional property such as those described earlier. In the latter type of problem, while it may be difficult to find the solution, once found, it is straightforward to verify that it is in fact a Nash equilibrium (with the desired simple property). This is not so for the notion of ESS: given a candidate strategy, it does not appear straightforward to figure out whether there exists a strategy that successfully invades it. However, appearances can be deceiving; perhaps there is a not entirely obvious, but nevertheless fast and elegant way of checking whether such an invading strategy exists. Even if not, it is not immediately clear whether this makes the problem of *finding* an ESS genuinely harder. Computational complexity provides the natural toolkit for answering these questions.

The complexity of computing whether a game has an evolutionarily stable strategy (for an overview, see Chapter 29 of the Algorithmic Game Theory book [Suri, 2007]) was first studied by Etessami and Lochbihler [2008], who proved that the problem is both NP-hard and coNP-hard, as well as that the problem is contained in Σ_2^P (the class of decision problems that can be solved in nondeterministic polynomial time when given access to an NP oracle). Nisan [2006] subsequently² proved the stronger hardness result that the problem is co D^P -hard. He also observed that it follows from his reduction that the problem of determining whether a given strategy is an ESS is coNP-hard (and Etessami and Lochbihler [2008] then pointed out that this also follows from their reduction). Etessami and Lochbihler [2008] also showed that the problem of determining the existence of a *regular* ESS is NP-complete. As was pointed out in both papers, all of this still leaves the main question of the exact complexity of the general ESS problem open. In this paper, this is settled: the problem is in fact Σ_2^P -complete.

The proof is structured as follows. Lemma 1 shows that the slightly more general problem of determining whether an ESS exists whose support is restricted to a subset of the strategies is Σ_2^P -hard. This is the main part of the proof. Then, Lemma 2 points out that if two pure strategies are exact duplicates, neither of them can occur in the support of any ESS. By this, we can disallow selected strategies from taking part in any ESS simply by duplicating them. Combining this with the first result, we arrive at the main result, Theorem 1.

One may well complain that Lemma 2 is a bit of a cheat; perhaps we should just consider duplicate strategies to be “the same” strategy and merge them

² An early version of Etessami and Lochbihler [2008] appeared in 2004.




Fig. 1. An example MINMAX-CLIQUE instance (with $k = 2$), for which the answer is “no.”

back into one. As the reader probably suspects, such a hasty and limited patch will not avoid the hardness result. Even something a little more thorough, such as iterated elimination of very weakly dominated strategies (in some order), will not suffice: in Appendix A I show, with additional analysis and modifications, that the result holds even in games where each pure strategy is the unique best response to some mixed strategy.

2 Hardness with Restricted Support

Definition 2 In *ESS-RESTRICTED-SUPPORT*, we are given a symmetric two-player normal-form game G with strategies S , and a subset $T \subseteq S$. We are asked whether there exists an evolutionarily stable strategy of G that places positive probability only on strategies in T (but not necessarily on all strategies in T).

Definition 3 (MINMAX-CLIQUE) We are given a graph $G = (V, E)$, sets I and J , a partition of V into subsets V_{ij} for $i \in I$ and $j \in J$, and a number k . We are asked whether it is the case that for every function $t : I \rightarrow J$, there is a clique of size (at least) k in the subgraph induced on $\bigcup_{i \in I} V_{i,t(i)}$. (Without loss of generality, we will require $k > 1$.)

Example. Figure 1 shows a tiny MINMAX-CLIQUE instance (let $k = 2$). The answer to this instance is “no” because for $t(1) = 2, t(2) = 1$, the graph induced on $\bigcup_{i \in I} V_{i,t(i)} = V_{12} \cup V_{21} = \{v_{12}, v_{21}\}$ has no clique of size at least 2.

Recall that $\Pi_2^P = \text{co}\Sigma_2^P$.

Known Theorem 1 ([Ko and Lin, 1995]) *MINMAX-CLIQUE is Π_2^P -complete.*

Lemma 1 *ESS-RESTRICTED-SUPPORT is Σ_2^P -hard.*

Proof: We reduce from the complement of MINMAX-CLIQUE. That is, we show how to transform any instance of MINMAX-CLIQUE into a symmetric two-player normal-form game with a distinguished subset T of its strategies, so that this game has an ESS with support in T if and only if the answer to the MINMAX-CLIQUE instance is “no.”

The Reduction. For every $i \in I$ and every $j \in J$, create a strategy s_{ij} . For every $v \in V$, create a strategy s_v . Finally, create a single additional strategy s_0 .

- For all $i \in I$ and $j \in J$, $u(s_{ij}, s_{ij}) = 1$.
- For all $i \in I$ and $j, j' \in J$ with $j \neq j'$, $u(s_{ij}, s_{ij'}) = 0$.
- For all $i, i' \in I$ with $i \neq i'$ and $j, j' \in J$, $u(s_{ij}, s_{i'j'}) = 2$.
- For all $i \in I$, $j \in J$, and $v \in V$, $u(s_{ij}, s_v) = 2 - 1/|I|$.
- For all $i \in I$ and $j \in J$, $u(s_{ij}, s_0) = 2 - 1/|I|$.
- For all $i \in I$, $j \in J$, and $v \in V_{ij}$, $u(s_v, s_{ij}) = 2 - 1/|I|$.
- For all $i \in I$, $j, j' \in J$ with $j \neq j'$, and $v \in V_{ij}$, $u(s_v, s_{ij'}) = 0$.
- For all $i, i' \in I$ with $i \neq i'$, $j, j' \in J$, and $v \in V_{ij}$, $u(s_v, s_{i'j'}) = 2 - 1/|I|$.
- For all $v \in V$, $u(s_v, s_v) = 0$.
- For all $v, v' \in V$ with $v \neq v'$ where $(v, v') \notin E$, $u(s_v, s_{v'}) = 0$.
- For all $v, v' \in V$ with $v \neq v'$ where $(v, v') \in E$, $u(s_v, s_{v'}) = (k/(k-1))(2 - 1/|I|)$.
- For all $v \in V$, $u(s_v, s_0) = 0$.
- For all $i \in I$ and $j \in J$, $u(s_0, s_{ij}) = 2 - 1/|I|$.
- For all $v \in V$, $u(s_0, s_v) = 0$.
- $u(s_0, s_0) = 0$.

We are asked whether there exists an ESS that places positive probability only on strategies s_{ij} with $i \in I$ and $j \in J$. That is, $T = \{s_{ij} : i \in I, j \in J\}$.

Example. Consider again the MINMAX-CLIQUE instance from Figure 1. The game to which the reduction maps this instance is:

	s_{11}	s_{12}	s_{21}	s_{22}	$s_{v_{11}}$	$s_{v_{12}}$	$s_{v_{21}}$	$s_{v_{22}}$	s_0
s_{11}	1	0	2	2	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{12}	0	1	2	2	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{21}	2	2	1	0	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{22}	2	2	0	1	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
$s_{v_{11}}$	$3/2$	0	$3/2$	$3/2$	0	0	3	3	0
$s_{v_{12}}$	0	$3/2$	$3/2$	$3/2$	0	0	0	3	0
$s_{v_{21}}$	$3/2$	$3/2$	$3/2$	0	3	0	0	0	0
$s_{v_{22}}$	$3/2$	$3/2$	0	$3/2$	3	3	0	0	0
s_0	$3/2$	$3/2$	$3/2$	$3/2$	0	0	0	0	0

It has an ESS σ with weight $1/2$ on each of s_{12} and s_{21} . In contrast, (for example) σ' with weight $1/2$ on each of s_{11} and s_{21} is invaded by the strategy σ'' with weight $1/2$ on each of $s_{v_{11}}$ and $s_{v_{21}}$, because $u(\sigma'', \sigma') = u(\sigma', \sigma') = 3/2$ and $u(\sigma'', \sigma'') = u(\sigma', \sigma'') = 3/2$.

Proof of equivalence. Suppose there exists a function $t : I \rightarrow J$ such that every clique in the subgraph induced on $\bigcup_{i \in I} V_{i,t(i)}$ has size strictly less than k . We will show that the mixed strategy σ that places probability $1/|I|$ on $s_{i,t(i)}$ for each $i \in I$ (and 0 everywhere else) is an ESS.

First, we show that σ is a best response against itself. For any s_{ij} in the support of σ , we have $u(s_{ij}, \sigma) = (1/|I|) \cdot 1 + (1 - 1/|I|) \cdot 2 = 2 - 1/|I|$, and hence we also have $u(\sigma, \sigma) = 2 - 1/|I|$. For s_{ij} not in the support of σ , we have $u(s_{ij}, \sigma) = (1/|I|) \cdot 0 + (1 - 1/|I|) \cdot 2 = 2 - 2/|I| < 2 - 1/|I|$. For all $i \in I$, for all $v \in V_{i,t(i)}$, we have $u(s_v, \sigma) = (1/|I|) \cdot (2 - 1/|I|) + (1 - 1/|I|) \cdot (2 - 1/|I|) = 2 - 1/|I|$. For all $i \in I$, $j \in J$ with $j \neq t(i)$, and $v \in V_{ij}$, we have $u(s_v, \sigma) = (1/|I|) \cdot 0 + (1 - 1/|I|) \cdot (2 - 1/|I|) = (1 - 1/|I|)(2 - 1/|I|) < 2 - 1/|I|$. Finally, $u(s_0, \sigma) = 2 - 1/|I|$. So σ is a best response to itself.

It follows that if there were a strategy $\sigma' \neq \sigma$ that could successfully invade σ , then σ' must put probability only on best responses to σ . Based on the calculations in the previous paragraph, these best responses are s_0 , and, for any i , $s_{i,t(i)}$ and, for all $v \in V_{i,t(i)}$, s_v . The expected utility of σ against any of these is $2 - 1/|I|$ (in particular, for any i , we have $u(\sigma, s_{i,t(i)}) = (1/|I|) \cdot 1 + (1 - 1/|I|) \cdot 2 = 2 - 1/|I|$). Hence, $u(\sigma, \sigma') = 2 - 1/|I|$, and to successfully invade, σ' must attain $u(\sigma', \sigma') \geq 2 - 1/|I|$.

We can write $\sigma' = p_0 s_0 + p_1 \sigma'_1 + p_2 \sigma'_2$, where $p_0 + p_1 + p_2 = 1$, σ'_1 only puts positive probability on the $s_{i,t(i)}$ strategies, and σ'_2 only puts positive probability on the s_v strategies with $v \in V_{i,t(i)}$. The strategy that results from conditioning σ' on σ'_1 not being played may be written as $(p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2$, and thus we may write $u(\sigma', \sigma') = p_1^2 u(\sigma'_1, \sigma'_1) + p_1(p_0 + p_2)u(\sigma'_1, (p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2) + (p_0 + p_2)p_1 u((p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2, \sigma'_1) + (p_0 + p_2)^2 u((p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2, (p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2)$. Now, if we shift probability mass from s_0 to σ'_2 , i.e., we decrease p_0 and increase p_2 by the same amount, this will not affect any of the coefficients in the previous expression; it will not affect any of $u(\sigma'_1, \sigma'_1)$, $u(\sigma'_1, (p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2)$ (because $u(s_{ij}, v) = u(s_{ij}, s_0) = 2 - 1/|I|$), and $u((p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2, \sigma'_1)$ (because $u(s_0, s_{ij}) = u(s_v, s_{ij}) = 2 - 1/|I|$ when $v \in V_{ij}$ or $v \in V_{i'j'}$ with $i' \neq i$); and it will not decrease $u((p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2, (p_0/(p_0 + p_2))s_0 + (p_2/(p_0 + p_2))\sigma'_2)$ (because for any $v \in V$, $u(s_0, s_0) = u(s_0, s_v) = u(s_v, s_0) = 0$). Therefore, we may assume without loss of generality that $p_0 = 0$, and hence $\sigma' = p_1 \sigma'_1 + p_2 \sigma'_2$.

It follows that we can write $u(\sigma', \sigma') = p_1^2 u(\sigma'_1, \sigma'_1) + p_1 p_2 u(\sigma'_1, \sigma'_2) + p_2 p_1 u(\sigma'_2, \sigma'_1) + p_2^2 u(\sigma'_2, \sigma'_2)$. We first note that $u(\sigma'_1, \sigma'_1)$ can be at most $2 - 1/|I|$. Specifically, $u(\sigma'_1, \sigma'_1) = (\sum_i \sigma'_1(s_{i,t(i)})^2) \cdot 1 + (1 - \sum_i \sigma'_1(s_{i,t(i)})^2) \cdot 2$, and this expression is uniquely maximized by setting each $\sigma'_1(s_{i,t(i)})$ to $1/|I|$. $u(\sigma'_1, \sigma'_2)$ is easily seen to also be $2 - 1/|I|$, and $u(\sigma'_2, \sigma'_1)$ is easily seen to be at most $2 - 1/|I|$ (in fact, it is exactly that). Thus, to obtain $u(\sigma', \sigma') \geq 2 - 1/|I|$, we must

have either $p_1 = 1$ or $u(\sigma'_2, \sigma'_2) \geq 2 - 1/|I|$. However, in the former case, we would require $u(\sigma'_1, \sigma'_1) = 2 - 1/|I|$, which can only be attained by setting each $\sigma'_1(s_{i,t(i)})$ to $1/|I|$ —but this would result in $\sigma' = \sigma$. Thus, we can conclude $u(\sigma'_2, \sigma'_2) \geq 2 - 1/|I|$. But then σ'_2 would also successfully invade σ . Hence, we can assume without loss of generality that $\sigma' = \sigma'_2$, i.e., $p_0 = p_1 = 0$ and $p_2 = 1$.

That is, we can assume that σ' only places positive probability on strategies s_v with $v \in \bigcup_{i \in I} V_{i,t(i)}$. For any $v, v' \in V$, we have $u(s_v, s_{v'}) = u(s_{v'}, s_v)$. Specifically, $u(s_v, s_{v'}) = u(s_{v'}, s_v) = (k/(k-1))(2 - 1/|I|)$ if $v \neq v'$ and $(v, v') \in E$, and $u(s_v, s_{v'}) = u(s_{v'}, s_v) = 0$ otherwise. Now, suppose that $\sigma'(s_v) > 0$ and $\sigma'(s_{v'}) > 0$ for $v \neq v'$ with $(v, v') \notin E$. We can write $\sigma' = p_0\sigma'' + p_1s_v + p_2s_{v'}$, where $p_0, p_1 = \sigma'(s_v)$, and $p_2 = \sigma'(s_{v'})$ sum to 1. We have $u(\sigma', \sigma') = p_0^2u(\sigma'', \sigma'') + 2p_0p_1u(\sigma'', s_v) + 2p_0p_2u(\sigma'', s_{v'})$ (because $u(s_v, s_v) = u(s_{v'}, s_{v'}) = u(s_v, s_{v'}) = 0$). Suppose, without loss of generality, that $u(\sigma'', s_v) \geq u(\sigma'', s_{v'})$. Then, if we shift all the mass from $s_{v'}$ to s_v (so that the mass on the latter becomes $p_1 + p_2$), this can only increase $u(\sigma', \sigma')$, and it reduces the size of the support of σ' by 1. By repeated application, we can assume without loss of generality that the support of σ' corresponds to a clique of the induced subgraph on $\bigcup_{i \in I} V_{i,t(i)}$. We know this clique has size c where $c < k$. $u(\sigma', \sigma')$ is maximized if σ' randomizes uniformly over its support, in which case $u(\sigma', \sigma') = ((c-1)/c)(k/(k-1))(2 - 1/|I|) < ((k-1)/k)(k/(k-1))(2 - 1/|I|) = 2 - 1/|I|$. But this contradicts that σ' would successfully invade σ . It follows that σ is indeed an ESS.

Conversely, suppose that there exists an ESS σ that places positive probability only on strategies s_{ij} with $i \in I$ and $j \in J$. We must have $u(\sigma, \sigma) \geq 2 - 1/|I|$, because otherwise s_0 would be a better response to σ . First suppose that for every $i \in I$, there is at most one $j \in J$ such that σ places positive probability on s_{ij} (we will shortly show that this must be the case). Let $t(i)$ denote the $j \in J$ such that $\sigma(s_{ij}) > 0$ (if there is no such j for some i , then choose an arbitrary j to equal $t(i)$). Then, $u(\sigma, \sigma)$ is uniquely maximized by setting $\sigma(s_{i,t(i)}) = 1/|I|$ for all $i \in I$, resulting in $u(\sigma, \sigma) = (1/|I|) \cdot 1 + (1 - 1/|I|) \cdot 2 = 2 - 1/|I|$. Hence, this is the only way to ensure that $u(\sigma, \sigma) \geq 2 - 1/|I|$, under the assumption that for every $i \in I$, there is at most one $j \in J$ such that σ places positive probability on s_{ij} .

Now, let us consider the case where there exists an $i \in I$ such that there exist $j, j' \in J$ with $j \neq j'$, $\sigma(s_{ij}) > 0$, and $\sigma(s_{ij'}) > 0$, to show that such a strategy cannot obtain a utility of $2 - 1/|I|$ or more against itself. We can write $\sigma = p_0\sigma' + p_1s_{ij} + p_2s_{ij'}$, where σ' places probability zero on s_{ij} and $s_{ij'}$. We observe that $u(\sigma', s_{ij}) = u(s_{ij}, \sigma')$ and $u(\sigma', s_{ij'}) = u(s_{ij'}, \sigma')$, because when the game is restricted to these strategies, each player always gets the same payoff as the other player. Moreover, $u(\sigma', s_{ij}) = u(\sigma', s_{ij'})$, because σ' does not place positive probability on either s_{ij} or $s_{ij'}$. Hence, we have that $u(\sigma, \sigma) = p_0^2u(\sigma', \sigma') + 2p_0(p_1 + p_2)u(\sigma', s_{ij}) + p_1^2 + p_2^2$. But then, if we shift all the mass from $s_{ij'}$ to s_{ij} (so that the mass on the latter becomes $p_1 + p_2$) to obtain strategy σ'' , it follows that $u(\sigma'', \sigma'') > u(\sigma, \sigma)$. By repeated application, we can find a strategy σ''' such that $u(\sigma''', \sigma''') > u(\sigma, \sigma)$ and for every $i \in I$, there is

at most one $j \in J$ such that σ''' places positive probability on s_{ij} . Because we showed previously that the latter type of strategy can obtain expected utility at most $2 - 1/|I|$ against itself, it follows that it is in fact the *only* type of strategy (among those that randomize only over the s_{ij} strategies) that can obtain expected utility $2 - 1/|I|$ against itself. Hence, we can conclude that the ESS σ must have, for each $i \in I$, exactly one $j \in J$ (to which we will refer as $t(i)$) such that $\sigma(s_{i,t(i)}) = 1/|I|$, and that σ places probability 0 on every other strategy.

Finally, suppose, for the sake of contradiction, that there exists a clique of size k in the induced subgraph on $\bigcup_{i \in I} V_{i,t(i)}$. Consider the strategy σ' that places probability $1/k$ on each of the corresponding strategies s_v . We have that $u(\sigma, \sigma) = u(\sigma, \sigma') = u(\sigma', \sigma) = 2 - 1/|I|$. Moreover, $u(\sigma', \sigma') = (1/k) \cdot 0 + ((k - 1)/k) \cdot (k/(k - 1))(2 - 1/|I|) = 2 - 1/|I|$. It follows that σ' successfully invades σ —but this contradicts σ being an ESS. It follows, then, that t is such that every clique in the induced graph on $\bigcup_{i \in I} V_{i,t(i)}$ has size strictly less than k . ■

3 Hardness without Restricted Support

Lemma 2 (No duplicates in ESS) *Suppose that strategies s_1 and s_2 ($s_1 \neq s_2$) are duplicates, i.e., for all s , $u(s_1, s) = u(s_2, s)$.³ Then no ESS places positive probability on s_1 or s_2 .*

Proof: For the sake of contradiction, suppose σ is an ESS that places positive probability on s_1 or s_2 (or both). Then, let $\sigma' \neq \sigma$ be identical to σ with the exception that $\sigma'(s_1) \neq \sigma(s_1)$ and $\sigma'(s_2) \neq \sigma(s_2)$ (but it must be that $\sigma'(s_1) + \sigma'(s_2) = \sigma(s_1) + \sigma(s_2)$). That is, σ' redistributes some mass between s_1 and s_2 . Then, σ cannot repel σ' , because $u(\sigma, \sigma) = u(\sigma', \sigma)$ and $u(\sigma, \sigma') = u(\sigma', \sigma')$. ■

Definition 4 *In ESS, we are given a symmetric two-player normal-form game G . We are asked whether there exists an evolutionarily stable strategy of G .*

Theorem 1 *ESS is Σ_2^P -complete.*

Proof: Etessami and Lochbihler [2008] proved membership in Σ_2^P . We prove hardness by reduction from ESS-RESTRICTED-SUPPORT, which is hard by Lemma 1. Given the game G with strategies S and subset of strategies $T \subseteq S$ that can receive positive probability, construct a modified game G' by duplicating all the strategies in $S \setminus T$. (At this point, for duplicate strategies s_1 and s_2 , we require $u(s, s_1) = u(s, s_2)$ as well as $u(s_1, s) = u(s_2, s)$.) If G has an ESS σ that places positive probability only on strategies in T , this will still be an ESS in G' , because any strategy that uses the new duplicate strategies will still be repelled,

³ It is fine to require $u(s, s_1) = u(s, s_2)$ as well, and we will do so in the proof of Theorem 1, but it is not necessary for this lemma to hold.

just as its equivalent strategy that does not use the new duplicates was repelled in the original game. (Here, it should be noted that the equivalent strategy in the original game cannot turn out to be σ , because σ does not put any probability on a strategy that is duplicated.) On the other hand, if G' has an ESS, then by Lemma 2, this ESS can place positive probability only on strategies in T . This ESS will still be an ESS in G (all of whose strategies also exist in G'), and naturally it will still place positive probability only on strategies in T . ■

A Hardness without duplication

In this appendix, it is shown that with some additional analysis and modifications, the result holds even in games where each pure strategy is the unique best response to some mixed strategy. That is, the hardness is not simply an artifact of the introduction of duplicate or otherwise redundant strategies.

Definition 5 *In the MINMAX-CLIQUE problem, say vertex v dominates vertex v' if they are in the same partition element V_{ij} , there is no edge between them, and the set of neighbors of v is a superset (not necessarily strict) of the set of neighbors of v' .*

Lemma 3 *Removing a dominated vertex does not change the answer to a MINMAX-CLIQUE instance.*

Proof: In any clique in which dominated vertex v' participates (and therefore its dominator v does not), v can participate in its stead. ■

Modified Lemma 1 *ESS-RESTRICTED-SUPPORT is Σ_2^P -hard, even if every pure strategy is the unique best response to some mixed strategy.*

Proof: We use the same reduction as in the proof of Lemma 1. We restrict our attention to instances of the MINMAX-CLIQUE problem where $|I| \geq 2$, $|J| \geq 2$, there are no dominated vertices, and every vertex is part of at least one edge. Clearly, the problem remains Π_2^P -complete when restricting attention to these instances. For the games resulting from these restricted instances, we show that every pure strategy is the unique best response to some mixed strategy. Specifically:

- s_{ij} is the unique best response to the strategy that distributes $1 - \epsilon$ mass uniformly over the $s_{i'j'}$ with $i' \neq i$, and ϵ mass uniformly over the $s_{ij'}$ with $j' \neq j$. (This is because only pure strategies $s_{ij'}$ will get a utility of 2 against the part with mass $1 - \epsilon$, and among these only s_{ij} will get a utility of 1 against the part with mass ϵ .)
- s_v (with $v \in V_{ij}$) is the unique best response to the strategy that places $(1/|I|)(1 - \epsilon)$ probability on s_{ij} and $(1/(|I||J|))(1 - \epsilon)$ probability on every $s_{i'j'}$ with $i' \neq i$, and that distributes the remaining ϵ mass uniformly over

- the vertex strategies corresponding to neighbors of v . (This is because s_v obtains an expected utility of $2 - 1/|I|$ against the part with mass $1 - \epsilon$, and an expected utility of $(k/(k-1))(2 - 1/|I|)$ against the part with mass ϵ ; strategies $s_{v'}$ with $v' \notin V_{ij}$ obtain utility strictly less than $2 - 1/|I|$ against the part with mass $1 - \epsilon$; and strategies $s_{i''j''}$, s_0 , and $s_{v'}$ with $v' \in V_{ij}$ obtain utility at most $2 - 1/|I|$ against the part with mass $1 - \epsilon$, and an expected utility of strictly less than $(k/(k-1))(2 - 1/|I|)$ against the part with mass ϵ . (In the case of $s_{v'}$ with $v' \in V_{ij}$, this is because by assumption, v' does not dominate v , so either v has a neighbor that v' does not have, which gets positive probability and against which $s_{v'}$ gets a utility of 0; or, there is an edge between v and v' , so that $s_{v'}$ gets positive probability and $s_{v'}$ gets utility 0 against itself.))
- s_0 is the unique best response to the strategy that randomizes uniformly over the s_{ij} . (This is because it obtains utility $2 - 1/|I|$ against that strategy, and all the other pure strategies obtain utility strictly less against that strategy, due to getting utility 0 against at least one pure strategy in its support.)

■

The following lemma is a generalization of Lemma 2.

Modified Lemma 2 *Suppose that subset $S' \subseteq S$ satisfies:*

- *for all $s \in S \setminus S'$ and $s', s'' \in S'$, we have $u(s', s) = u(s'', s)$ (that is, strategies in S' are interchangeable when they face a strategy outside S');⁴ and*
- *the restricted game where players must choose from S' has no ESS.*

Then no ESS of the full game places positive probability on any strategy in S' .

Proof: Consider a strategy σ that places positive probability on S' . We can write $\sigma = p_1\sigma_1 + p_2\sigma_2$, where $p_1 + p_2 = 1$, σ_1 places positive probability only on $S \setminus S'$, and σ_2 places positive probability only on S' . Because no ESS exists in the game restricted to S' , there must be a strategy σ'_2 (with $\sigma'_2 \neq \sigma_2$) whose support is contained in S' that successfully invades σ_2 , so either (1) $u(\sigma'_2, \sigma_2) > u(\sigma_2, \sigma_2)$ or (2) $u(\sigma'_2, \sigma_2) = u(\sigma_2, \sigma_2)$ and $u(\sigma'_2, \sigma'_2) \geq u(\sigma_2, \sigma'_2)$. Now consider the strategy $\sigma' = p_1\sigma_1 + p_2\sigma'_2$; we will show that it successfully invades σ . This is because $u(\sigma', \sigma) = p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma_2) + p_2p_1u(\sigma'_2, \sigma_1) + p_2^2u(\sigma'_2, \sigma_2) = p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma_2) + p_2p_1u(\sigma_2, \sigma_1) + p_2^2u(\sigma'_2, \sigma_2) \geq p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma_2) + p_2p_1u(\sigma_2, \sigma_1) + p_2^2u(\sigma_2, \sigma_2) = u(\sigma, \sigma)$, where the second equality follows from the property assumed in the lemma. If case (1) above holds, then the inequality is strict and σ is not a best response against itself. If case (2) holds, then we have equality; moreover, $u(\sigma', \sigma') = p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma'_2) + p_2p_1u(\sigma'_2, \sigma_1) + p_2^2u(\sigma'_2, \sigma'_2) = p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma'_2) + p_2p_1u(\sigma_2, \sigma_1) + p_2^2u(\sigma'_2, \sigma'_2) \geq p_1^2u(\sigma_1, \sigma_1) + p_1p_2u(\sigma_1, \sigma'_2) + p_2p_1u(\sigma_2, \sigma_1) + p_2^2u(\sigma_2, \sigma'_2) = u(\sigma, \sigma')$, where the second equality follows from the property assumed in the lemma. So in this case too, σ' successfully invades σ . ■

⁴ Again, it is fine to require $u(s, s') = u(s, s'')$ as well, and we will do so in the proof of Modified Theorem 1, but it is not necessary for the lemma to hold.

Modified Theorem 1 *ESS is Σ_2^P -complete, even if every pure strategy is the unique best response to some mixed strategy.*

Proof: Again, Etessami and Lochbihler [2008] proved membership in Σ_2^P . For hardness, we use a similar proof strategy as in Theorem 1, again reducing from ESS-RESTRICTED-SUPPORT, which is hard even if every pure strategy is the unique best response to some mixed strategy, by Modified Lemma 1. Given the game G with strategies S and subset of strategies $T \subseteq S$ that can receive positive probability, construct a modified game G' by replacing each pure strategy $s \in S \setminus T$ by three new pure strategies, s^1, s^2, s^3 . For each $s' \notin \{s^1, s^2, s^3\}$, we will have $u(s^i, s') = u(s, s')$ (the utility of the original s) and $u(s', s^i) = u(s', s)$ for all $i \in \{1, 2, 3\}$; for all $i, j \in \{1, 2, 3\}$, we will have $u(s^i, s^j) = u(s, s) + \rho(i, j)$, where ρ gives the payoffs of rock-paper-scissors (with -1 for a loss, 0 for a tie, and 1 for a win).

If G has an ESS that places positive probabilities only on strategies in T , this will still be an ESS in G' because any strategy σ' that uses new strategies s^i will still be repelled, just as the corresponding strategy σ'' that put the mass on the corresponding original strategies s (i.e., $\sigma''(s) = \sigma'(s^1) + \sigma'(s^2) + \sigma'(s^3)$) was repelled in the original game. (Unlike in the proof of the original Theorem 1, here it is perhaps not immediately obvious that $u(\sigma'', \sigma'') = u(\sigma', \sigma')$, because the right-hand side involves additional terms involving ρ . But ρ is a symmetric zero-sum game, and any strategy results in an expected utility of 0 against itself in such a game.) On the other hand, if G' has an ESS, then by Modified Lemma 2 (letting $S' = \{s^1, s^2, s^3\}$ and using the fact that rock-paper-scissors has no ESS), this ESS can place positive probability only on strategies in T . This ESS will still be an ESS in G (for any potentially invading strategy in G there would be an equivalent such strategy in G' , for example replacing s by s^1 as needed), and naturally it will still place positive probability only on strategies in T .

Finally it remains to be shown that in G' each pure strategy is the unique best response to some mixed strategy, using the fact that this is the case for G . For a pure strategy in T , we can simply use the same mixed strategy as we use for that pure strategy in G , replacing mass placed on each $s \notin T$ in G with a uniform mixture over s^1, s^2, s^3 where needed. (By using a uniform mixture, we guarantee that each s^i obtains the same expected utility against the mixed strategy as the corresponding s strategy in G .) For a pure strategy $s^i \notin T$, we cannot simply use the same mixed strategy as we use for the corresponding s in G (with the same uniform mixture trick), because s^1, s^2, s^3 would all be equally good responses. But because these three would be the *only* best responses, we can mix in a sufficiently small amount of $s^{i+1 \pmod{3}}$ (where i beats $i+1 \pmod{3}$) in ρ to make s^i the unique best response. ■

Acknowledgments

I thank ARO and NSF for support under grants W911NF-12-1-0550, W911NF-11-1-0332, IIS-0953756, and CCF-1101659. The compendium by Schaefer and Umans [2008] guided me to the MINMAX-CLIQUE problem.

Bibliography

- Christian Borgs, Jennifer Chayes, Nicole Immorlica, Adam Tauman Kalai, Vahab Mirrokni, and Christos Papadimitriou. The myth of the Folk Theorem. *Games and Economic Behavior*, 70(1):34–43, 2010.
- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 82–90, Ann Arbor, MI, USA, 2006.
- Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.
- Constantinos Daskalakis, Paul Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- Kousha Etessami and Andreas Lochbihler. The computational complexity of evolutionarily stable strategies. *International Journal of Game Theory*, 37(1):93–113, 2008. An earlier version was made available as ECCO tech report TR04-055, 2004.
- Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.
- Kristoffer Arnsfelt Hansen, Peter Bro Miltersen, and Troels Bjerre Sørensen. The computational complexity of trembling hand perfection and other equilibrium refinements. In *Proceedings of the Third Symposium on Algorithmic Game Theory (SAGT-10)*, pages 198–209, Athens, Greece, 2010.
- Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 2013.
- Ker-I Ko and Chih-Long Lin. On the complexity of min-max optimization problems and their approximation. In D.-Z. Du and P. M. Pardalos, editors, *Minimax and Applications*, pages 219–239. Kluwer Academic Publishers, Boston, 1995.
- Spyros C. Kontogiannis and Paul G. Spirakis. Equilibrium points in fear of correlated threats. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 210–221, Shanghai, China, 2008.
- Carlton Lemke and Joseph Howson. Equilibrium points of bimatrix games. *Journal of the Society of Industrial and Applied Mathematics*, 12:413–423, 1964.
- Michael L. Littman and Peter Stone. A polynomial-time Nash equilibrium algorithm for repeated games. *Decision Support Systems*, 39:55–66, 2005.
- John Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.

- Noam Nisan. A note on the computational hardness of evolutionary stable strategies. *Electronic Colloquium on Computational Complexity (ECCC)*, 13(076), 2006.
- Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), 2008.
- Christos H. Papadimitriou. Algorithms, games and the Internet. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
- George Price and John Maynard Smith. The logic of animal conflict. *Nature*, 246:15–18, 1973.
- Rahul Savani and Bernhard von Stengel. Hard-to-solve bimatrix games. *Econometrica*, 74:397–429, 2006.
- Marcus Schaefer and Christopher Umans. Completeness in the polynomial-time hierarchy: A compendium, 2008.
- Troels Bjerre Sørensen. Computing a proper equilibrium of a bimatrix game. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 916–928, Valencia, Spain, 2012.
- Siddharth Suri. Computational evolutionary game theory. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 29. Cambridge University Press, 2007.
- Bernhard von Stengel and Shmuel Zamir. Leadership games with convex strategy sets. *Games and Economic Behavior*, 69:446–457, 2010.