A Generalized Strategy Eliminability Criterion and Computational Methods for Applying It*

Vincent Conitzer
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
conitzer@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
sandholm@cs.cmu.edu

ABSTRACT

We define a generalized strategy eliminability criterion for bimatrix games that considers whether a given strategy is eliminable relative to given dominator & eliminee subsets of the players' strategies. We show that this definition spans a spectrum of eliminability criteria from strict dominance (when the sets are as small as possible) to Nash equilibrium (when the sets are as large as possible). We show that checking whether a strategy is eliminable according to this criterion is coNP-complete (both when all the sets are as large as possible and when the dominator sets each have size 1). We then give an alternative definition of the eliminability criterion and show that it is equivalent using the Minimax Theorem. We show how this alternative definition can be translated into a mixed integer program of polynomial size with a number of (binary) integer variables equal to the sum of the sizes of the eliminee sets, implying that checking whether a strategy is eliminable according to the criterion can be done in polynomial time, given that the eliminee sets are small. Finally, we study using the criterion for iterated elimination of strategies.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

General Terms

Algorithms, Economics, Theory

Keywords

Game Theory, Solution Concepts, (Iterated) Elimination of Strategies, Nash Equilibrium, Dominance

*This material is based upon work supported by the National Science Foundation under ITR grants IIS-0121678 and IIS-0427858, and a Sloan Fellowship. An earlier version of the paper appeared at AAAI-05.

EC'06 Workshop on Alternative Solution Concepts in Mechanism Design. June 12, 2006, Ann Arbor, Michigan, USA.
.

1. INTRODUCTION

Solving general-sum games is a topic of growing interest in computer science. To solve such games, the concept of (iterated) dominance is often too strong: it cannot eliminate enough strategies. But, if possible, we would like a stronger argument for eliminating a strategy than (mixed-strategy) Nash equilibrium. Similarly, in mechanism design (where one gets to create the game), implementation in dominant strategies is often excessively restrictive, but implementation in (Bayes-)Nash equilibrium may not be sufficiently strong for the designer's purposes. Hence, it is desirable to have eliminability criteria that are *between* these concepts in strength. In this paper, we will introduce such a criterion. The criterion we introduce considers whether a given strategy is eliminable relative to given dominator & eliminee subsets of the players' strategies. The criterion spans an entire *spectrum* of strength between Nash equilibrium and strict dominance (in terms of which strategies it can eliminate), and in the extremes can be made to coincide with either of these two concepts, depending on how the dominator & eliminee sets are set. It can also be used for iterated elimination of strategies.

An important question to ask of any solution concept is how efficiently a solution can be *computed*. A recent sequence of papers [8, 4, 9, 5] shows that the problem of finding a Nash equilibrium (even in the two-player case) is complete for the class PPAD. (In contrast, *approximate* Nash equilibria can be found in quasi-polynomial time [19]. Also, Nash equilibria can be found in polynomial time for average-payoff repeated games [20].) The best-known algorithm for finding a Nash equilibrium, the *Lemke-Howson* algorithm [17], has a worst-case exponential running time [27], and methods based on exhaustively searching through the space of the mixed strategies' supports fare comparatively well for many classes of games [26]. It is also known that finding Nash equilibria *with certain additional properties* (for example, the social-welfare maximizing Nash equilibrium) is NP-complete [12, 6]. The computational complexity of dominance and iterated dominance has been studied as well [16, 11, 7]. In this paper, we will study the computational complexity of applying the new eliminability criterion, and provide a mixed integer programming approach for it.

Throughout, we focus on two-player games only. The eliminability criterion itself can be generalized to more players, but the computational tools we introduce do not straightforwardly generalize to more players. Moreover, we focus only on normal-form games (rather than make use of structured representations of games [14, 18, 2, 13]).

One of the benefits of the new criterion is that when a strategy cannot be eliminated by dominance (but it can be eliminated by the Nash equilibrium concept), the new criterion may provide a stronger argument than Nash equilibrium for eliminating the strategy, by using dominator & eliminee sets smaller than the entire strategy set. To get the strongest possible argument for eliminating a strategy, the dominator & eliminee sets should be chosen to be as small as possible while still having the strategy be eliminable relative to these sets.[1] Iterated elimination of strategies using the new criterion is also possible, and again, to get the strongest possible argument for eliminating a strategy, the sequence of eliminations leading up to it should use dominator & eliminee sets that are as small as possible.[2]

As another benefit, the algorithm that we provide for checking whether a strategy is eliminable according to the new criterion can also be used as a subroutine in the computation of Nash equilibria. Specifically, any strategy that is eliminable (even using iterated elimination) according to the criterion is guaranteed not to occur in any Nash equilibrium. Current state-of-the-art algorithms for computing Nash equilibria already use a subroutine that eliminates (conditionally) dominated strategies [26]. Because the new criterion can eliminate more strategies than dominance, the algorithm we provide may speed up the computation of Nash equilibria. (For purposes of speed, it is probably desirable to only apply special cases of the criterion that can be computed fast—in particular, as we will show, eliminability according to the criterion can be computed fast when the eliminee sets are small. Even these special cases are more powerful than dominance.)

2. A MOTIVATING EXAMPLE

Because the definition of the new eliminability criterion is complex, we will first illustrate it with an example. Consider the following (partially specified) game.

	σ_c^1	σ_c^2	σ_c^3	σ_c^4
σ_r^1	?, ?	?, 2	?, 0	?, 0
σ_r^2	2, ?	2, 2	2, 0	2, 0
σ_r^3	0, ?	0, 2	3, 0	0, 3
σ_r^4	0, ?	0, 2	0, 3	3, 0

A quick look at this game reveals that strategies σ_r^3 and σ_r^4 are both *almost* dominated by σ_r^2—but they perform better than σ_r^2 against σ_c^3 and σ_c^4, respectively. Similarly, strategies σ_c^3 and σ_c^4 are both almost dominated by σ_c^2—but they perform better than σ_c^2 against σ_r^4 and σ_r^3, respectively. So we are unable to eliminate any strategies using (even weak) dominance.

Now consider the following reasoning. In order for it to be worthwhile for the row player to ever play σ_r^3 rather than σ_r^2, the column player should play σ_c^3 at least $\frac{2}{3}$ of the time. (If it is exactly $\frac{2}{3}$, then switching from σ_r^2 to σ_r^3 will cost the row player 2 exactly $\frac{1}{3}$ of the time, but the row player will gain 1 exactly $\frac{2}{3}$ of the time, so the expected benefit is 0.) But, similarly, in order for it to be worthwhile for the column player to ever play σ_c^3, the row player should play σ_r^4 at least $\frac{2}{3}$ of the time. But again, in order for it to be worthwhile for the row player to ever play σ_r^4, the column player should play σ_c^4 at least $\frac{2}{3}$ of the time. Thus, if both the row and the column player accurately assess the probabilities that the other places on these strategies, and their strategies are rational with respect to these assessments (as would be the case in a Nash equilibrium), then, if the row player puts positive probability on σ_r^3, by the previous reasoning, the column player should be playing σ_c^3 at least $\frac{2}{3}$ of the time, and σ_c^4 at least $\frac{2}{3}$ of the time. Of course, this is impossible; so, in a sense, the row player should not play σ_r^3.

It may appear that all we have shown is that σ_r^3 is not played in any Nash equilibrium. But, to some extent, our argument for not playing σ_r^3 did not make use of the full elimination power of the Nash equilibrium concept. Most notably, we only reasoned about a small part of the game: we never mentioned strategies σ_r^1 and σ_c^1, and we did not even specify most of the utilities for these strategies. (It is easy to extend this example so that the argument only uses an arbitrarily small fraction of the strategies and of the utilities in the matrix, for instance by adding many copies of σ_r^1 and σ_c^1.) The locality of the reasoning that we did is more akin to the notion of dominance, which is perhaps the extreme case of local reasoning about eliminability—only two strategies are mentioned in it. So, in this sense, the argument for eliminating σ_r^3 is somewhere between dominance and Nash equilibrium in strength.

3. DEFINITION OF THE ELIMINABILITY CRITERION

We are now ready to give the formal definition of the generalized eliminability criterion. To make the definition a bit simpler, we define its negation—when a strategy is *not* eliminable relative to certain sets of strategies. Also, we only define when one of the *row player's* strategies is eliminable, but of course the definition is analogous for the column player.

The definition, which considers when a strategy e_r^* is eliminable relative to subsets D_r, E_r of the row player's pure strategies (with $e_r^* \in E_r$) and subsets D_c, E_c of the column player's pure strategies, can be stated informally as follows. To protect e_r^* from elimination, we should be able to specify the probabilities that the players' mixed strategies place on the E_i sets in such a way that 1) e_r^* receives nonzero probability, and 2) for every pure strategy e_i that receives nonzero probability, for every mixed strategy d_i using only strategies in D_i, it is conceivable that player $-i$'s mixed strategy[3] is completed so that e_i is no worse than d_i.[4] The formal definition follows.

DEFINITION 1. *Given a two-player game in normal form, subsets D_r, E_r of the row player's pure strategies Σ_r, subsets D_c, E_c of the column player's pure strategies Σ_c, and a distinguished strategy $e_r^* \in E_r$, we say that e_r^* is* not eliminable relative to *D_r, E_r, D_c, E_c, if there exist functions (partial mixed strategies) $p_r : E_r \to [0,1]$ and $p_c : E_c \to [0,1]$ with $p_r(e_r^*) > 0$, $\sum_{e_r \in E_r} p_r(e_r) \leq 1$, and $\sum_{e_c \in E_c} p_c(e_c) \leq 1$, such that the following holds. For both $i \in \{r, c\}$, for any $e_i \in E_i$ with $p_i(e_i) > 0$, for any mixed strategy d_i placing positive probability only on strategies in D_i, there is some pure strategy $\sigma_{-i} \in \Sigma_{-i} - E_{-i}$ such that (letting*

[1]There may be multiple minimal vectors of dominator & eliminee sets relative to which the strategy is eliminable; in this paper, we will not attempt to settle which of these minimal vectors, if any, constitutes the most powerful argument for eliminating the strategy.

[2]Here, there may also be a tradeoff with the length of the elimination path. For example, there may be a path of several eliminations using dominator & eliminee sets that are small, as well as a single elimination using dominator & eliminee sets that are large, both of which eliminate a given strategy. (In fact, we will *always* be confronted with this situation, as Corollary 3 will show.) Again, in this paper, we will not attempt to settle which argument for eliminating the strategy is stronger.

[3]As is common in the game theory literature, $-i$ denotes "the player other than i."

[4]This description may sound similar to the concept of *rationalizability*. However, in two-player games (the subject of this paper), rationalizability is known to coincide with iterated strict dominance [25].

$p_{-i} \diamond \sigma_{-i}$ *denote the mixed strategy that results from placing the remaining probability* $1 - \sum_{e_{-i} \in E_{-i}} p_{-i}(e_{-i})$ *that is not used by the partial mixed strategy* p_{-i} *on* σ_{-i}*), we have:* $u_i(e_i, p_{-i} \diamond \sigma_{-i}) \geq u_i(d_i, p_{-i} \diamond \sigma_{-i})$*. (If* p_{-i} *already uses up all the probability, we simply have* $u_i(e_i, p_{-i}) \geq u_i(d_i, p_{-i})$*—no* σ_{-i} *needs to be chosen.)*[5]

In the example from the previous subsection, we can set $D_r = \{\sigma_r^2\}$, $D_c = \{\sigma_c^2\}$, $E_r = \{\sigma_r^3, \sigma_r^4\}$, $E_c = \{\sigma_c^3, \sigma_c^4\}$, and $e_r^* = \sigma_r^3$. Then, by the reasoning that we did, it is impossible to set p_r and p_c so that the conditions are satisfied, and hence σ_r^3 is eliminable relative to these sets.

4. THE SPECTRUM OF STRENGTH

In this section we show that the generalized eliminability criterion we defined in in the previous section spans a spectrum of strength all the way from Nash equilibrium (when the sets D_r, E_r, D_c, E_c are chosen as large as possible), to strict dominance (when the sets are chosen as small as possible). First, we show that the criterion is monotonically increasing, in the sense that the larger we make the sets D_r, E_r, D_c, E_c, the more strategies are eliminable.

PROPOSITION 1. *If* e_r^* *is eliminable relative to* $D_r^1, E_r^1, D_c^1, E_c^1$*, and* $D_r^1 \subseteq D_r^2, E_r^1 \subseteq E_r^2, D_c^1 \subseteq D_c^2, E_c^1 \subseteq E_c^2$*, then* e_r^* *is eliminable relative to* $D_r^2, E_r^2, D_c^2, E_c^2$*.*

PROOF. We will prove this by showing that if e_r^* is not eliminable relative to $D_r^2, E_r^2, D_c^2, E_c^2$, then e_r^* is not eliminable relative to $D_r^1, E_r^1, D_c^1, E_c^1$. It is straightforward that making the D_i sets smaller only weakens the condition on strategies e_i with $p_i(e_i) > 0$ in Definition 1. Hence, if e_r^* is not eliminable relative to $D_r^2, E_r^2, D_c^2, E_c^2$, then e_r^* is not eliminable relative to $D_r^1, E_r^2, D_c^1, E_c^2$. All that remains to show is that making the E_i sets smaller will not make e_r^* eliminable. To show this, we first observe that, if in its last step Definition 1 allowed for distributing the remaining probability arbitrarily over the strategies in $\Sigma_{-i} - E_{-i}$ (rather than requiring a single one of these strategies to receive all the remaining probability), this would not change the definition, because we might as well place all the remaining probability on the strategy $\sigma_{-i} \in \Sigma_{-i} - E_{-i}$ that maximizes $u_i(e_i, \sigma_{-i}) - u_i(d_i, \sigma_{-i})$. Now, let p_r and p_c be partial mixed strategies over E_r^2 and E_c^2 that prove that e_r^* is not eliminable relative to $D_r^1, E_r^2, D_c^1, E_c^2$. Then, to show that e_r^* is not eliminable relative to $D_r^1, E_r^1, D_c^1, E_c^1$, use the partial mixed strategies p'_r and p'_c, which are simply the restrictions of p_r and p_c to E_r^1 and E_c^1, respectively. For any $e_i \in E_i^1$ with $p'_i(e_i) > 0$ and for any mixed strategy d_i over D_i^1, we know that there exists some $\sigma_{-i} \in \Sigma_{-i} - E_{-i}^2$ such that $u_i(e_i, p_{-i} \diamond \sigma_{-i}) \geq u_i(d_i, p_{-i} \diamond \sigma_{-i})$ (because the p_i prove that e_r^* is not eliminable relative to $D_r^1, E_r^2, D_c^1, E_c^2$). But, the distribution $p_{-i} \diamond \sigma_{-i}$ is a legitimate completion of the partial mixed strategy p'_{-i} as well (albeit one that distributes the remaining probability over multiple strategies), and hence the p'_i prove that e_r^* is not eliminable relative to $D_r^1, E_r^1, D_c^1, E_c^1$. □

Next, we show that the Nash equilibrium concept is weaker[6] than our generalized eliminability criterion—in the sense that the generalized criterion can never eliminate a strategy that is in some Nash equilibrium. So, if a strategy can be eliminated by the generalized criterion, it can be eliminated by the Nash equilibrium concept.

PROPOSITION 2. *If there is some Nash equilibrium that places positive probability on pure strategy* σ_r^**, then* σ_r^* *is not eliminable relative to any* D_r, E_r, D_c, E_c*.*

PROOF. Let σ'_r be the row player's (mixed) strategy in the Nash equilibrium (which places positive probability on σ_r^*), and let σ'_c be the column player's (mixed) strategy in the Nash equilibrium. For any D_r, E_r, D_c, E_c with $\sigma_r^* \in E_r$, to prove that σ_r^* is not eliminable relative to these sets, simply let p_r coincide with σ'_r on E_r—that is, let p_r be the probabilities that the row player places on the strategies in E_r in the equilibrium. (Thus, $p_r(\sigma_r^*) > 0$). Similarly, let p_c coincide with σ'_c on E_c. We will prove that the condition on strategies with positive probability is satisfied for the row player; the case of the column player follows by symmetry. For any $e_r \in E_r$ with $p_r(e_r) > 0$, for any mixed strategy d_r, we have $u_r(e_r, \sigma'_c) - u_r(d_r, \sigma'_c) \geq 0$, by the Nash equilibrium condition. Now, let pure strategy $\sigma_c \in \arg\max_{\sigma \in \Sigma_c - E_c}(u_r(e_r, p_c \diamond \sigma) - u_r(d_r, p_c \diamond \sigma))$. Then we must have $u_r(e_r, p_c \diamond \sigma_c) - u_r(d_r, p_c \diamond \sigma_c) \geq u_r(e_r, \sigma'_c) - u_r(d_r, \sigma'_c) \geq 0$ (because $p_c \diamond \sigma_c$ and σ'_c coincide on E_c, and for the former, the remainder of the distribution is chosen to maximize this expression). It follows that σ_r^* is not eliminable relative to any D_r, E_r, D_c, E_c. □

We next show that by choosing the sets D_r, E_r, D_c, E_c as large as possible, we can make the generalized eliminability criterion coincide with the Nash equilibrium concept.[7]

PROPOSITION 3. *Let* $D_r = E_r = \Sigma_r$ *and* $D_c = E_c = \Sigma_c$*. Then* e_r^* *is eliminable relative to these sets if and only if there is no Nash equilibrium that places positive probability on* e_r^**.*

PROOF. The "only if" direction follows from Proposition 2. For the "if" direction, suppose e_r^* is not eliminable relative to $D_r = E_r = \Sigma_r$ and $D_c = E_c = \Sigma_c$. The partial distributions p_r and p_c with $p_r(e_r^*) > 0$ that show that e_r^* is not eliminable must use up all the probability (the probabilities must sum to one), because there are no strategies outside $E_c = \Sigma_c$ and $E_r = \Sigma_r$ to place any remaining probability on. Hence, we must have, for any strategy $e_r \in E_r = \Sigma_r$ with $p_r(e_r) > 0$, that for any mixed strategy d_r, $u_r(e_r, p_c) \geq u_r(d_r, p_c)$ (and the same for the column player). But these are precisely the conditions for p_r and p_c to constitute a Nash equilibrium. It follows that there is a Nash equilibrium with positive probability on e_r^*. □

Moving to the other side of the spectrum, we now show that the concept of strict dominance is stronger than the generalized eliminability criterion—in the sense that the generalized eliminability criterion can always eliminate a strictly dominated strategy (as long as the dominating strategy is in D_r).

PROPOSITION 4. *If pure strategy* σ_r^* *is strictly dominated by some mixed strategy* d_r*, then* σ_r^* *is eliminable relative to any* D_r, E_r, D_c, E_c *such that 1)* $\sigma_r^* \in E_r$*, and 2) all the pure strategies on which* d_r *places positive probability are in* D_r*.*

PROOF. To show that σ_r^* is not eliminable relative to these sets, we must set $p_r(\sigma_r^*) > 0$, and thus we must demonstrate that for

[5]We need to make this case explicit for the case $E_{-i} = \Sigma_{-i}$.

[6]When discussing elimination of strategies, it is tempting to say that the stronger criterion is the one that can eliminate more strategies. However, when discussing solution concepts, the convention is that the stronger concept is the one that implies the other. Therefore, the criterion that can eliminate fewer strategies is actually the stronger one. For example, strict dominance is stronger than weak dominance, even though weak dominance can eliminate more strategies.

[7]Unlike Nash equilibrium, the generalized eliminability criterion does not discuss what probabilities should be placed on strategies that are not eliminated, so it only "coincides" with Nash equilibrium in terms of what it can eliminate.

some pure strategy $\sigma_c \in \Sigma_c - E_c$, $u_r(\sigma_r^*, p_c \diamond \sigma_c) \geq u_r(d_r, p_c \diamond \sigma_c)$ (or, if all the probability is used up, $u_r(\sigma_r^*, p_c) \geq u_r(d_r, p_c)$), because d_r only places positive probability on strategies in D_r. But this is impossible, because by strict dominance, $u_r(\sigma_r^*, \sigma_c) < u_r(d_r, \sigma_c)$ for any mixed strategy σ_c. □

Finally, we show that by choosing the sets E_r, E_c as small as possible, we can make the generalized eliminability criterion coincide with the strict dominance concept.

PROPOSITION 5. *Let $E_c = \{\}$ and $E_r = \{e_r\}$. Then e_r is eliminable relative to D_r, E_r, D_c, E_c if and only if it is strictly dominated by some mixed strategy that places positive probability only on elements of D_r.*

PROOF. The "if" direction follows from Proposition 4. For the "only if" direction, suppose that e_r is eliminable relative to these sets. That means that there exists a mixed strategy d_r that places positive probability only on strategies in D_r such that for any pure strategy $\sigma_c \in \Sigma_c - E_c = \Sigma_c$, $u(e_r, \sigma_c) < u(d_r, \sigma_c)$ (because $E_c = \{\}$ and $E_r = \{e_r\}$, this is the only way in which an attempt to prove that e_r is not eliminable could fail). But this is precisely the condition for d_r to strictly dominate e_r. □

We are now ready to turn to computational aspects of the new eliminability criterion.

5. APPLYING THE NEW ELIMINABILITY CRITERION CAN BE COMPUTATIONALLY HARD

In this section, we demonstrate that applying the eliminability criterion can be computationally hard, in the sense of worst-case complexity.[8] We show that applying the eliminability criterion is coNP-complete in two key special cases (subclasses of the problem). The first case is the one in which the D_r, E_r, D_c, E_c sets are set to be as large as possible. Here, the hardness follows directly from Proposition 3 and a known hardness result on computing Nash equilibria [12, 6].

THEOREM 1. *Deciding whether a given strategy is eliminable relative to $D_r = E_r = \Sigma_r$ and $D_c = E_c = \Sigma_c$ is coNP-complete, even when the game is symmetric.*

PROOF. By Proposition 3, this is the converse of asking whether there exists a Nash equilibrium with positive probability on the given strategy. This is NP-complete [12, 6]. □

While this shows that the eliminability criterion is, in general, computationally hard to apply, we may wonder if there are special cases in which it is computationally easy to apply. Natural special cases to look at include those in which some of the sets D_r, E_r, D_c, E_c are small. The next theorem shows that applying the eliminability criterion remains coNP-complete even when $|D_r| = |D_c| = 1$.

THEOREM 2. *Deciding whether a given strategy is eliminable relative to given D_r, E_r, D_c, E_c is coNP-complete, even when $|D_r| = |D_c| = 1$.*

PROOF. We will show later (Corollary 1) that the problem is in coNP. To show that the problem is coNP-hard, we reduce an arbitrary KNAPSACK instance (given by m cost-value pairs (c_i, v_i), a cost constraint C and a value target V; we assume without loss of generality that $C = 1 - \epsilon$, for some ϵ small enough that it is impossible for a subset of the c_i to sum to a value strictly between C and 1,[9] that $c_i > 0$ for all i, and that $\sum_{i=1}^{m} v_i \leq 1$) to the following eliminability question. Let the game be as follows. The row player has $m+2$ distinct pure strategies: $e_r^1, e_r^2, \ldots, e_r^m, e_r^*, d_r$ (where $E_r = \{e_r^1, e_r^2, \ldots, e_r^m, e_r^*\}$ and $D_r = \{d_r\}$). The column player has $m+1$ distinct pure strategies: $e_c^1, e_c^2, \ldots, e_c^m, d_c$ (where $E_c = \{e_c^1, e_c^2, \ldots, e_c^m\}$ and $D_c = \{d_c\}$). Let the utilities be as follows:

- $u_r(e_r^i, e_c^j) = 1$ for all $i \neq j$;
- $u_r(e_r^i, e_c^i) = 1 - \frac{1}{v_i}$ for all i;
- $u_r(e_r^i, d_c) = 1$ for all i;
- $u_r(e_r^*, e_c^i) = \frac{1}{V} - 1$ for all i;
- $u_r(e_r^*, d_c) = -1$;
- $u_r(d_r, e_c^i) = 0$ for all i;
- $u_r(d_r, d_c) = 0$;
- $u_c(e_r^i, e_c^j) = 0$ for all $i \neq j$;
- $u_c(e_r^i, e_c^i) = \frac{1}{c_i}$ for all i;
- $u_c(e_r^i, d_c) = 1$ for all i;
- $u_c(e_r^*, e_c^i) = 0$ for all i;
- $u_c(e_r^*, d_c) = 1$;
- $u_c(d_r, e_c^i) = 0$ for all i;
- $u_c(d_r, d_c) = 1$.

Thus, the matrix is as follows:

	e_c^1	e_c^2	$\cdots$	e_c^m	d_c
e_r^1	$1-\frac{1}{v_1}, \frac{1}{c_1}$	$1, 0$	$\cdots$	$1, 0$	$1, 1$
e_r^2	$1, 0$	$1-\frac{1}{v_2}, \frac{1}{c_2}$	$\cdots$	$1, 0$	$1, 1$
$\vdots$					
e_r^m	$1, 0$	$1, 0$	$\cdots$	$1-\frac{1}{v_m}, \frac{1}{c_m}$	$1, 1$
e_r^*	$\frac{1}{V}-1, 0$	$\frac{1}{V}-1, 0$	$\cdots$	$\frac{1}{V}-1, 0$	$-1, 1$
d_r	$0, 0$	$0, 0$	$\cdots$	$0, 0$	$0, 1$

We now show that e_r^* is eliminable relative to D_r, E_r, D_c, E_c if and only if there is no solution to the KNAPSACK instance.

First suppose there is a solution to the KNAPSACK instance. Then, for every i such that (c_i, v_i) is included in the KNAPSACK solution, let $p_r(e_r^i) = c_i$; for every i such that (c_i, v_i) is not included in the KNAPSACK solution, let $p_r(e_r^i) = 0$. Also, let $p_r(e_r^*) = 1 - \sum_{i=1}^{m} p_r(e_r^i)$. (We note that $\sum_{i=1}^{m} p_r(e_r^i) \leq C = 1 - \epsilon$, so that $p_r(e_r^*) \geq \epsilon > 0$.) Also, for every i such that (c_i, v_i) is included in the KNAPSACK solution, let $p_c(e_c^i) = v_i$. We now show that p_r and p_c satisfy the conditions of Definition 1. If the column player places the remaining probability on d_c, then the utility for the row player of playing any e_r^i with $p_r(e_r^i) > 0$ is $1 - \frac{v_i}{v_i} = 0$;

[8] Because we only show hardness in the worst case, it is possible that many (or even most) instances are in fact easy to solve.

[9] Because we may assume that the c_i and C are all integers divided by some number K, it is sufficient if $\epsilon < \frac{1}{K}$.

the utility of playing e_r^* is $-1+\frac{1}{V}\sum_{i=1}^{m} p_c(e_c^i) \geq -1+\frac{V}{V}=0$; and the utility of playing d_r is also 0. Thus, the condition is satisfied for all elements of E_r that have positive probability. As for E_c, we note that all of the row player's probability has already been used up. The utility of playing any e_c^i with $p_c(e_c^i) > 0$ is $\frac{c_i}{c_i}=1$, whereas the utility for playing d_c is also 1. Thus, the condition is satisfied for all elements of E_c that have positive probability. It follows that p_r and p_c satisfy the conditions of Definition 1 and e_r^* is not eliminable relative to D_r, E_r, D_c, E_c.

Now suppose that e_r^* is not eliminable relative to D_r, E_r, D_c, E_c. Let p_r and p_c be partial mixed strategies on E_r and E_c satisfying the conditions of Definition 1. We must have that $p_r(e_r^*) > 0$. The utility for the row player of playing e_r^* is $-1+\frac{1}{V}\sum_{i=1}^{m} p_c(e_c^i)$, which must be at least 0 (the utility of playing d_r); hence $\sum_{i=1}^{m} p_c(e_c^i) \geq V$. The utility for the column player of playing e_c^i is $\frac{p_r(e_r^i)}{c_i}$, which must be at least 1 (the utility of playing d_c) if $p_c(e_c^i) > 0$; hence $p_r(e_r^i) \geq c_i$ if $p_c(e_c^i) > 0$. Finally, the utility for the row player of playing e_r^i is $1-\frac{p_c(e_c^i)}{v_i}$, which must be at least 0 (the utility of playing d_r) if $p_r(e_r^i) > 0$; hence $p_c(e_c^i) \leq v_i$ if $p_r(e_r^i) > 0$. Because we must have $p_r(e_r^i) \geq c_i > 0$ if $p_c(e_c^i) > 0$, it follows that we must always have $p_c(e_c^i) \leq v_i$. Let $S = \{i : p_c(e_c^i) > 0\}$. We must have $\sum_{i \in S} v_i \geq \sum_{i \in S} p_c(e_c^i) \geq V$. Also, we must have $\sum_{i \in S} c_i \leq \sum_{i \in S} p_r(e_r^i) < 1$ (because we must have $p_r(e_r^*) > 0$). Because it is impossible that $C < \sum_{i \in S} c_i < 1$, it follows that $\sum_{i \in S} c_i \leq C$. But then, S is a solution to the KNAPSACK instance. □

However, we will show later that the eliminability criterion can be applied in polynomial time if the E_i sets are small (regardless of the size of the D_i sets). To do so, we first need to introduce an alternative version of the definition.

6. AN ALTERNATIVE, EQUIVALENT DEFINITION OF THE ELIMINABILITY CRITERION

In this section, we will give an alternative definition of eliminability, and we will show it is equivalent to the one presented in Definition 1. While the alternative definition is slightly less intuitve than the original one, it is easier to work with computationally, as we will show in the next section. Informally, the alternative definition differs from the original one as follows: in the alternative definition, the completion of player $-i$'s mixed strategy has to be chosen *before* player i's strategy d_i is chosen (but after player i's strategy e_i with $p_i(e_i) > 0$ is chosen). The formal definition follows.

DEFINITION 2. *Given a two-player game in normal form, subsets D_r, E_r of the row player's pure strategies Σ_r, subsets D_c, E_c of the column player's pure strategies Σ_c, and a distinguished strategy $e_r^* \in E_r$, we say that e_r^* is* not eliminable relative to D_r, E_r, D_c, *E_c, if there exist functions (partial mixed strategies) $p_r : E_r \to [0,1]$ and $p_c : E_c \to [0,1]$ with $p_r(e_r^*) > 0$, $\sum_{e_r \in E_r} p_r(e_r) \leq 1$, and $\sum_{e_c \in E_c} p_c(e_c) \leq 1$, such that the following holds. For both $i \in \{r, c\}$, for any $e_i \in E_i$ with $p_i(e_i) > 0$, there exists some completion of the probability distribution over $-i$'s strategies, given by $p_{-i}^{e_i} : \Sigma_{-i} \to [0,1]$ (with $p_{-i}^{e_i}(e_{-i}) = p_{-i}(e_{-i})$ for all $e_{-i} \in E_{-i}$, and $\sum_{\sigma_{-i} \in \Sigma_{-i}} p_{-i}^{e_i}(\sigma_{-i}) = 1$), such that for any pure strategy $d_i \in D_i$, we have $u_i(e_i, p_{-i}^{e_i}) \geq u_i(d_i, p_{-i}^{e_i})$.*

We now show that the two definitions are equivalent.

THEOREM 3. *The notions of eliminability put forward in Definitions 1 and 2 are equivalent. That is, e_r^* is eliminable relative to D_r, E_r, D_c, E_c according to Definition 1 if and only if e_r^* is eliminable relative to (the same) D_r, E_r, D_c, E_c according to Definition 2.*

PROOF. The definitions are identical up to the condition that each strategy with positive probability (each $e_r \in E_r$ with $p_r(e_r) > 0$ and each $e_c \in E_c$ with $p_c(e_c) > 0$) must satisfy. We will show that these conditions are equivalent across the two definitions, thereby showing that the definitions are equivalent.

To show that the conditions are equivalent, we introduce another, zero-sum game that is a function of the original game, the sets D_r, E_r, D_c, E_c, the chosen partial probability distributions p_r and p_c, and the strategy e_i for which we are checking whether the conditions are satisfied. (Without loss of generality, assume that we are checking it for some strategy $e_r \in E_r$ with $p_r(e_r) > 0$.)

The zero-sum game has two players, 1 and 2 (not to be confused with the row and column players of the original game). Player 1 chooses some $d_r \in D_r$, and player 2 chooses some $\sigma_c \in \Sigma_c - E_c$. The utility to player 1 is $u_r(d_r, p_c \diamond \sigma_c) - u_r(e_r, p_c \diamond \sigma_c)$ (and the utility to player 2 is the negative of this). (We assume without loss of generality that p_c does not already use up all the probability, because in this case the conditions are trivially equivalent across the two definitions.)

First, suppose that player 1 must declare her probability distribution (mixed strategy) over D_r first, after which player 2 best-responds. Then, letting $\Delta(X)$ denote the set of probability distributions over set X, player 1 will receive $\max_{\delta_r \in \Delta(D_r)} \min_{\sigma_c \in \Sigma_c - E_c} \sum_{d_r \in D_r} \delta_r(d_r)(u_r(d_r, p_c \diamond \sigma_c) - u_r(e_r, p_c \diamond \sigma_c)) = \max_{\delta_r \in \Delta(D_r)} \min_{\sigma_c \in \Sigma_c - E_c} u_r(\delta_r, p_c \diamond \sigma_c) - u_r(e_r, p_c \diamond \sigma_c)$. This expression is at most 0 if and only if the condition in Definition 1 is satisfied.

Second, suppose that player 2 must declare his probability distribution (mixed strategy) over $\Sigma_c - E_c$ first, after which player 1 best-responds. Then, player 1 will receive $\min_{\delta_c \in \Delta(\Sigma_c - E_c)} \max_{d_r \in D_r} \sum_{\sigma_c \in \Sigma_c - E_c} \delta_c(\sigma_c)(u_r(d_r, p_c \diamond \sigma_c) - u_r(e_r, p_c \diamond \sigma_c)) = \min_{\delta_c \in \Delta(\Sigma_c - E_c)} \max_{d_r \in D_r} \sum_{e_c \in E_c} p_c(e_c)(u_r(d_r, e_c) - u_r(e_r, e_c)) + \sum_{\sigma_c \in \Sigma_c - E_c} (1 - \sum_{e_c \in E_c} p_c(e_c))\delta_c(\sigma_c)(u_r(d_r, \sigma_c) - u_r(e_r, \sigma_c)) = \min_{\delta_c \in \Delta(\Sigma_c - E_c)} \max_{d_r \in D_r} u_r(d_r, p_c \diamond \delta_c) - u_r(e_r, p_c \diamond \delta_c)$. This expression is at most 0 if and only if the condition in Definition 2 is satisfied.

However, by the Minimax Theorem [28], the two expressions must have the same value, and hence the two conditions are equivalent. □

Informally, the reason that Definition 2 is easier to work with computationally is that all of the continuous variables (the values of the functions $p_r, p_c, p_c^{e_r}, p_r^{e_c}$) are set by the party that is trying to prove that the strategy is not eliminable; whereas in Definition 1, some of the continuous variables (the probabilities defining the mixed strategies d_r, d_c) are set by the party trying to refute the proof that the strategy is not eliminable. This will become more precise in the next section.

7. A MIXED INTEGER PROGRAMMING APPROACH

In this section, we show how to translate Definition 2 into a mixed integer program that determines whether a given strategy e_r^* is eliminable relative to given sets D_r, E_r, D_c, E_c. The variables in the program, which are all restricted to be nonnegative, are the $p_i(e_i)$ for all $e_i \in E_i$; the $p_i^{e_{-i}}(\sigma_i)$ for all $e_{-i} \in E_{-i}$ and all $\sigma_i \in \Sigma_i - E_i$; and *binary* indicator variables $b_i(e_i)$ for all $e_i \in E_i$ which can be set to zero if and only if $p_i(e_i) = 0$. The program is the following:

maximize $p_r(e_r^*)$ **subject to**

(probability constraints): for both $i \in \{r, c\}$, for all $e_i \in E_i$, $\sum_{e_{-i} \in E_{-i}} p_{-i}(e_{-i}) + \sum_{\sigma_{-i} \in \Sigma_{-i} - E_{-i}} p_{-i}^{e_i}(\sigma_{-i}) = 1$

(binary constraints): for both $i \in \{r, c\}$, for all $e_i \in E_i$, $p_i(e_i) \leq b_i(e_i)$

(main constraints): for both $i \in \{r, c\}$, for all $e_i \in E_i$ and all $d_i \in D_i$, $\sum_{e_{-i} \in E_{-i}} p_{-i}(e_{-i})(u_i(e_i, e_{-i}) - u_i(d_i, e_{-i})) + \sum_{\sigma_{-i} \in \Sigma_{-i} - E_{-i}} p_{-i}^{e_i}(\sigma_{-i})(u_i(e_i, \sigma_{-i}) - u_i(d_i, \sigma_{-i})) \geq (b_i(e_i) - 1)U_i$

In this program, the constant U_i is the maximum difference between two different utilities that player i may receive in the game, that is, $U_i = \max_{\sigma_r, \sigma_r' \in \Sigma_r, \sigma_c, \sigma_c' \in \Sigma_c} u_i(\sigma_r, \sigma_c) - u_i(\sigma_r', \sigma_c')$.

THEOREM 4. *The mixed integer program has a solution with objective value greater than zero if and only if e_r^* is not eliminable relative to D_r, E_r, D_c, E_c.*

PROOF. For any $e_i \in E_i$ with $p_i(e_i) > 0$, $b_i(e_i)$ must be 1, and thus the corresponding main constraints become: for any $d_i \in D_i$, $\sum_{e_{-i} \in E_{-i}} p_{-i}(e_{-i})(u_i(e_i, e_{-i}) - u_i(d_i, e_{-i})) + \sum_{\sigma_{-i} \in \Sigma_{-i} - E_{-i}} p_{-i}^{e_i}(\sigma_{-i})(u_i(e_i, \sigma_{-i}) - u_i(d_i, \sigma_{-i})) \geq 0$. These are equivalent to the constraints given on strategies $e_i \in E_i$ with $p_i(e_i) > 0$ in Definition 2. On the other hand, for any $e_i \in E_i$ with $p_i(e_i) = 0$, $b_i(e_i)$ can be set to 0, in which case the constraints become: for any $d_i \in D_i$, $\sum_{e_{-i} \in E_{-i}} p_{-i}(e_{-i})(u_i(e_i, e_{-i}) - u_i(d_i, e_{-i})) + \sum_{\sigma_{-i} \in \Sigma_{-i} - E_{-i}} p_{-i}^{e_i}(\sigma_{-i})(u_i(e_i, \sigma_{-i}) - u_i(d_i, \sigma_{-i})) \geq -U_i$. Because the probabilities in each of these constraints must sum to one by the probability constraints, and U_i is the maximum difference between two different utilities that player i may receive in the game, these constraints are vacuous. Therefore the main constraints correspond exactly to those in Definition 2. □

We obtain the following corollaries:

COROLLARY 1. *Checking whether a given strategy can be eliminated relative to given D_r, E_r, D_c, E_c is in coNP.*

PROOF. To see whether the strategy can be protected from elimination, we can nondeterministically choose the values for the binary variables $b_r(e_r)$ and $b_c(e_c)$. After this, only a linear program remains to be solved, which can be done in polynomial time [15]. □

COROLLARY 2. *Using the mixed integer program above, the time required to check whether a given strategy can be eliminated relative to given D_r, E_r, D_c, E_c is exponential only in $|E_r| + |E_c|$ (and not in $|D_r|, |D_c|, |\Sigma_r|$, or $|\Sigma_c|$).*

PROOF. Any mixed integer program whose only integer variables are binary variables can be solved in time exponential only in its number of binary variables (for example, by searching over all settings of its binary variables and solving the remaining linear program in each case). The number of binary variables in this program is $|E_r| + |E_c|$. □

8. ITERATED ELIMINATION

In this section, we study what happens when we eliminate strategies *iteratively* using the new criterion. The criterion can be iteratively applied by removing an eliminated strategy from the game, and subsequently checking for new eliminabilities in the game with the strategy removed, *etc.* (as in the more elementary, conventional notion of iterated dominance). First, we show that this procedure is, in a sense, sound.

THEOREM 5. *Iterated elimination according to the generalized criterion will never remove a strategy that is played with positive probability in some Nash equilibrium of the original game.*

PROOF. We will prove this by induction on the elimination round (that is, the number of strategies eliminated so far). The claim is true for the first round by Proposition 2. Now suppose it is true up to and including round k; we must show it is true for round $k + 1$. Suppose that the claim is false for round $k + 1$, that is, there exists some game G and some pure strategy σ such that 1) σ is played with positive probability in some Nash equilibrium of G, and 2) using k elimination rounds, G can be reduced to G^{k+1}, in which σ is eliminable. Now consider the game G^k which preceded G^{k+1} in the elimination sequence, that is, the game obtained by undoing the last elimination before G^{k+1}. Also, let σ' be the strategy removed from G^k to obtain G^{k+1}. Now, in G^k, σ cannot be eliminated by the induction assumption. However, by Proposition 3, any strategy that is not played with positive probability in any Nash equilibrium can be eliminated, so it follows that there is some Nash equilibrium of G^k in which σ is played with positive probability. Moreover, this Nash equilibrium cannot place positive probability on σ' (because otherwise, by Proposition 2, we would not be able to eliminate it). But then, this Nash equilibrium must also be a Nash equilibrium of G^{k+1}: it does not place any probability on strategies that are not in G^{k+1}, and the set of strategies that the players can switch to in G^{k+1} is a subset of those in G^k. Hence, by Proposition 2, we cannot eliminate σ from G^{k+1}, and we have achieved the desired contradiction. □

Because (the single-round version of) the eliminability criterion extends all the way to Nash equilibrium by Proposition 3, we get the following corollary.

COROLLARY 3. *Any strategy that can be eliminated using iterated elimination can also be eliminated in a single round (that is, without iterated application of the criterion).*

PROOF. By Proposition 3, all strategies that are not played with positive probability in any Nash equilibrium can be eliminated in a single round; but by Theorem 5, this is the only type of strategy that iterated elimination can eliminate. □

Interestingly, iterated elimination is in a sense incomplete:

PROPOSITION 6. *Removing an eliminated strategy from a game sometimes decreases the set of strategies that can be eliminated.*

PROOF. Consider the following game:

	L	M	R
U	2, 2	0, 1	0, 5
D	1, 0	1, 1	1, 0

The unique Nash equilibrium of this game is (D, M), for the following reasons. In order for it to be worthwhile for the row player to play U with positive probability, the column player should play L with probability at least $1/2$. But, in order for it to be worthwhile for the column player to play L with positive probability (rather than M), the row player should play U with probability at least $1/2$. However, if the row player plays U with probability at least $1/2$, then the column player's unique best response is to play R. Hence, the row player must play D in any Nash equilibrium, and the unique best response to D is M.

Thus, by Proposition 3, all strategies besides D and M can be eliminated. In particular, R can be eliminated. However, if we remove R from the game, the remaining game is:

	L	M
U	2, 2	0, 1
D	1, 0	1, 1

In this game, (U, L) is also a Nash equilibrium, and hence U and L can no longer be eliminated, by Proposition 2. □

This example highlights an interesting issue with respect to using this eliminability criterion as a preprocessing step in the computation of Nash equilibria: it does not suffice to simply throw out eliminated strategies and compute a Nash equilibrium for the remaining game. Rather, we need to use the criterion more carefully: if we know that a strategy is eliminable according to the criterion we can restrict our attention to supports for the player that do not include this strategy.

The example also directly implies that iterated elimination according to the generalized criterion is path-dependent (the choice of which strategy to remove first affects which strategies can/will be removed later). The same phenomenon occurs with iterated weak dominance (one strategy weakly dominates another if the former always does at least as well as the latter, and in at least one case, strictly better). There is a sizeable literature on path (in)dependence for various notions of dominance [10, 3, 24, 21, 22, 1].

In light of these results, it may appear that there is not much reason to do iterated elimination using the new criterion, because it never increases and sometimes even decreases the set of strategies that we can eliminate. However, we need to keep in mind that Theorem 5, Corollary 3, and Proposition 6 do not pose any restrictions on the sets D_r, E_r, D_c, E_c, and therefore (by Propositions 2 and 3) are effectively results about iteratively removing strategies based on whether they are played in a Nash equilibrium. However, the new criterion is more informative and useful when there are restrictions on the sets D_r, E_r, D_c, E_c. Of particular interest is the restriction $|E_r| + |E_c| \leq k$, because by Corollary 2 this quantity determines the (worst-case) runtime of the mixed integer programming approach that we presented in the previous section. Under this restriction, it turns out that iterated elimination can eliminate strategies that single-round elimination cannot.

PROPOSITION 7. *Under a restriction of the form $|E_r|+|E_c| \leq k$, iterated elimination can eliminate strategies that single-round elimination cannot (even when $k = 1$).*

PROOF. By Proposition 5, when $k = 1$ the eliminability criterion coincides with strict dominance (and hence iterated application of the criterion coincides with iterated strict dominance). So, consider the following game:

	L	R
U	1, 0	1, 1
D	0, 1	0, 0

Strict dominance cannot eliminate L, but iterated strict dominance (which can remove D first) can eliminate L. □

Of course, even under this (or any other) restriction iterated elimination remains sound in the sense of Theorem 5. Therefore, one sensible approach to eliminating strategies is the following. Iteratively apply the eliminability criterion (with whatever restrictions are desired to increase the strength of the argument, or are necessary to make it computationally manageable, such as $|E_r|+|E_c| \leq k$), removing each eliminated strategy, until the process gets stuck. Then, start again with the original game, and take a different path of iterated elimination (which may eliminate strategies that could no longer be eliminated after the first path of elimination, as described in Proposition 6), until the process gets stuck—*etc.* In the end, any strategy that was eliminated in any one of the elimination paths can be considered "eliminated", and this is safe by Theorem 5.[10]

Interestingly, here the analogy with iterated weak dominance breaks down. Because there is no soundness theorem such as Theorem 5 for iterated weak dominance, considering all the strategies that are eliminated in some iterated weak dominance elimination path to be simultaneously "eliminated" can lead to senseless results. Consider for example the following game:

	L	M	R
U	1, 1	0, 0	1, 0
D	1, 1	1, 0	0, 0

U can be eliminated by removing R first, and D can be eliminated by removing M first—but these are the row player's only strategies, so considering both of them to be eliminated makes little sense.

9. CONCLUSIONS

We defined a generalized eliminability criterion for bimatrix games that considers whether a given strategy is eliminable relative to given dominator & eliminee subsets of the players' strategies. We showed that this definition spans a spectrum of eliminability criteria from strict dominance (when the sets are as small as possible) to Nash equilibrium (when the sets are as large as possible). Thus, eliminating a strategy relative to dominator & eliminee sets of intermediate size can provide a stronger argument for eliminating a strategy than Nash equilibrium, even when the strategy cannot be eliminated by (iterated) dominance. We showed that checking whether a strategy is eliminable according to this criterion is coNP-complete (both when all the sets are as large as possible and when the dominator sets each have size 1). We then gave an alternative definition of the eliminability criterion and showed that it is equivalent using the Minimax Theorem. We showed how this alternative definition can be translated into a mixed integer program of polynomial size with a number of (binary) integer variables equal to the sum of the sizes of the eliminee sets, implying that checking whether a strategy is eliminable according to the criterion can be done in polynomial time if the eliminee sets are small. Finally, we studied using the criterion for iterated elimination of strategies.

There are numerous avenues for future research. One is to use the new eliminability criterion and the computational tools we provided for it to speed up search-based techniques for computing

[10]This procedure is reminiscent of iterative sampling.

Nash equilibria. Another avenue is to characterize the eliminability criterion at intermediate points of the spectrum. Yet another possibility is to try to find other special cases that can be computed in polynomial time. We can also experimentally analyze the runtime of the mixed integer programming approach on random games (such as those generated by GAMUT [23]). Finally, we can attempt to use the criterion as a solution concept in mechanism design.

10. REFERENCES

[1] K. R. Apt. Uniform proofs of order independence for various strategy elimination procedures. *Contributions to Theoretical Economics*, 4(1), 2004.

[2] B. Blum, C. R. Shelton, and D. Koller. A continuation method for Nash equilibria in structured games. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003.

[3] T. Borgers. Pure strategy dominance. *Econometrica*, 61(2):423–430, 1993.

[4] X. Chen and X. Deng. 3-Nash is PPAD-complete. *Electronic Colloquium on Computational Complexity*, Report No. 134, 2005.

[5] X. Chen and X. Deng. Settling the complexity of 2-player Nash equilibrium. *Electronic Colloquium on Computational Complexity*, Report No. 150, 2005.

[6] V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 765–771, Acapulco, Mexico, 2003.

[7] V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 88–97, Vancouver, Canada, 2005.

[8] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. *Electronic Colloquium on Computational Complexity*, Report No. 115, 2005.

[9] C. Daskalakis and C. Papadimitriou. Three-player games are hard. *Electronic Colloquium on Computational Complexity*, Report No. 139, 2005.

[10] I. Gilboa, E. Kalai, and E. Zemel. On the order of eliminating dominated strategies. *Operations Research Letters*, 9:85–89, 1990.

[11] I. Gilboa, E. Kalai, and E. Zemel. The complexity of eliminating dominated strategies. *Mathematics of Operation Research*, 18:553–565, 1993.

[12] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.

[13] G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: hard and easy games. In *Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 215–230, Bloomington, Indiana, USA, 2003.

[14] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.

[15] L. Khachiyan. A polynomial algorithm in linear programming. *Soviet Math. Doklady*, 20:191–194, 1979.

[16] D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7(3):103–107, 1988.

[17] C. Lemke and J. Howson. Equilibrium points of bimatrix games. *Journal of the Society of Industrial and Applied Mathematics*, 12:413–423, 1964.

[18] K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003.

[19] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 36–41, San Diego, CA, 2003.

[20] M. Littman and P. Stone. A polynomial-time Nash equilibrium algorithm for repeated games. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 48–54, San Diego, CA, 2003.

[21] L. M. Marx and J. M. Swinkels. Order independence for iterated weak dominance. *Games and Economic Behavior*, 18:219–245, 1997.

[22] L. M. Marx and J. M. Swinkels. Corrigendum, order independence for iterated weak dominance. *Games and Economic Behavior*, 31:324–329, 2000.

[23] E. Nudelman, J. Wortman, K. Leyton-Brown, and Y. Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, New York, NY, USA, 2004.

[24] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[25] D. G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–1050, 1984.

[26] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a Nash equilibrium. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 664–669, San Jose, CA, USA, 2004.

[27] R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 258–267, 2004.

[28] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1927.