Finding the Pareto Curve in Bimatrix Games is Easy

Nicola Gatti
Politecnico di Milano
Piazza Leonardo da Vinci 32
Milano, Italy
ngatti@elet.polimi.it

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

ABSTRACT

Pareto efficiency is a widely used property in solution concepts for cooperative and non–cooperative game–theoretic settings and, more generally, in multi–objective problems. However, finding or even approximating (when the objective functions are not convex) the Pareto curve is hard. Most of the literature focuses on computing concise representations to approximate the Pareto curve or on exploiting evolutionary approaches to generate approximately Pareto efficient samples of the curve. In this paper, we show that the Pareto curve of a bimatrix game can be found exactly in polynomial time and that it is composed of a polynomial number of pieces. Furthermore, each piece is a quadratic function. We use this result to provide algorithms for game-theoretic solution concepts that incorporate Pareto efficiency.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Multi–agent systems

General Terms

Algorithms, Economics

Keywords

Game theory (cooperative and non–cooperative); Pareto curve; solving games

1. INTRODUCTION

The computational study of strategic interactions is of extraordinary importance in artificial intelligence [16]. Game theory provides elegant solution concepts for such settings but they need to be accompanied by algorithms for finding such solutions in order to operationalize the theory.

Pareto efficiency [16] plays a central role in game settings with multiple objectives. Each agent is associated with one (or multiple) objective(s) and Pareto efficiency identifies the best tradeoffs among the different agents' objectives. The *Pareto curve* is commonly defined as the collection of Pareto efficient solutions. A number of solution concepts in cooperative and non–cooperative game theory are based on Pareto efficiency, prescribing the selection of specific points on the Pareto curve. For instance, in cooperative game theory, the *utilitarian bargaining solution (UBS)* [16] selects the points on the Pareto curve that maximize social welfare, the *Nash bargaining solution (NBS)* [13] selects the points on the Pareto curve that maximize the product of the agents' utilities, and the *Kalai–Smorodinsky bargaining solution (KSBS)* [11] selects the points on the Pareto curve that maintain the ratio between the maximum utility achievable by the first agent and the maximum utility achievable by the second one. In non–cooperative game theory, *strong Nash equilibrium (SNE)* strengthens the concept of *Nash equilibrium* [16] by requiring that the equilibrium is resilient to all coalitional deviations. Another way to word strong Nash equilibrium is that it is a strategy profile that is Pareto efficient for every coalition (including the grand coalition), fixing the strategies of the agents outside the coalition.

Finding the Pareto curve is typically hard as is computing solution concepts that incorporate Pareto efficiency [2, 5, 14]. Most literature focuses on approximating Pareto curve searching for concise representations [10, 12, 14], on providing results for specific cases in which objectives are convex [6], and on developing evolutionary approaches to generate approximately Pareto efficient samples [3].

In this paper we focus on bimatrix games. While the problem of verifying whether a solution is Pareto efficient has been shown to be easy [8], to our knowledge, nothing has been published about the problem of determining and characterizing the Pareto curve.

The contributions we provide in this paper are as follows.

- We study the Pareto curve of 2×2 bimatrix games, showing that there are 19 possible different cases and for each case we provide the related conditions over the parameters. The Pareto curve is, in general, non–convex, discontinuous, and piece–wise quadratic with at most 4 pieces. The curve can be computed in polynomial time.
- We study the Pareto curve of $m \times m$ bimatrix games, showing that it is composed of a polynomial (in m) number of pieces and that each piece belongs the Pareto curve of a 2×2 sub–bimatrix and is thus quadratic. The curve can be computed in time polynomial in m.
- We show that key solution concepts from cooperative game theory that incorporate Pareto efficiency in their definition—UBS, NBS, and KSBS—can be found in polynomial time in bimatrix games.
- We provide an algorithm to find a strong Nash equilibrium (SNE) [1] which calls a polynomial number of

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

times an $\mathcal{NP}$–complete oracle (thus matching the $\mathcal{NP}$-completeness of SNE finding [4, 8]). In contrast, prior algorithms for finding SNE take exponential time even if $\mathcal{P} = \mathcal{NP}$ [7, 8].

2. SETTING AND DEFINITIONS

A 2–player normal-form (aka. strategic–form) game is a tuple (N, A, U) [16] where:

- $N = \{1, 2\}$ is the set of agents (i denotes an agent),
- $A = \{A_1, A_2\}$ is the set of agents' actions and A_i is the set of agent i's actions (we denote a generic action by a, and by m_i the number of actions in A_i),
- $U = \{U_1, U_2\}$ is the set of agents' utility matrices where $U_i(a_1, a_2)$ is agent i's utility when agent 1 plays action a_1 and agent 2 plays a_2.

We denote by $\mathbf{x}_i$ the strategy (vector of probabilities) of agent i and by $x_{i,a}$ the probability with which agent i plays action $a \in A_i$. We denote by Δ_i the space of strategies over A_i, i.e., vectors $\mathbf{x}_i$ where the probabilities sum to 1. We denote by $\mathbf{x}$ the strategy profile defined as $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$. We denote by $u_i(\mathbf{x}) = \mathbf{x}_1^T U_i \mathbf{x}_2$ (where T means transpose) the expected utility of agent i as the agents' strategies vary and by $\mathbf{u}(\mathbf{x})$ the 2–dimensional vector defined as $\mathbf{u} = (u_1(\mathbf{x}), u_2(\mathbf{x}))$. We now introduce the concept of Pareto efficiency.

Definition 2.1. *A strategy profile* $\mathbf{x}$ *(analogously* $\mathbf{u}(\mathbf{x})$*) is* Pareto efficient *if there is no other strategy profile* $\mathbf{x}'$ *such that* $\mathbf{u}(\mathbf{x}') \geq \mathbf{u}(\mathbf{x})$ *and* $u_i(\mathbf{x}') > u_i(\mathbf{x})$ *for some* i.

Definition 2.2. *A strategy profile* $\mathbf{x}$ *(analogously* $\mathbf{u}(\mathbf{x})$*) is* weakly Pareto efficient *if there is no other strategy profile* $\mathbf{x}'$ *such that* $\mathbf{u}(\mathbf{x}') > \mathbf{u}(\mathbf{x})$.

We can now define the Pareto curve.

Definition 2.3. *The* Pareto curve *is the collection of* $\mathbf{u}(\mathbf{x})$ *such that* $\mathbf{x}$ *is Pareto efficient.*

Definition 2.4. *The* weak Pareto curve *is the collection of* $\mathbf{u}(\mathbf{x})$ *such that* $\mathbf{x}$ *is weakly Pareto efficient.*

Finally, we introduce the notion of Pareto dominance.

Definition 2.5. *Strategy profile* $\mathbf{x}$ *(analogously* $\mathbf{u}(\mathbf{x})$*) is* strongly Pareto dominated *by strategy profile* $\mathbf{x}'$ *(analogously* $\mathbf{u}(\mathbf{x}')$*) if* $\mathbf{u}(\mathbf{x}') > \mathbf{u}(\mathbf{x})$.

Definition 2.6. *Strategy profile* $\mathbf{x}$ *(analogously* $\mathbf{u}(\mathbf{x})$*) is* weakly Pareto dominated *by strategy profile* $\mathbf{x}'$ *(analogously* $\mathbf{u}(\mathbf{x}')$*) if* $\mathbf{u}(\mathbf{x}') \geq \mathbf{u}(\mathbf{x})$ *and* $u_i(\mathbf{x}') > u_i(\mathbf{x})$ *for some* i.

Note that the set of strategy profiles that are not weakly Pareto dominated constitute the Pareto curve, while the set of strategy profiles that are not strongly Pareto dominated constitutes the weak Pareto curve.

3. PARETO CURVE OF 2×2 GAMES

In this section we present results on the Pareto curve for 2×2 games. Results for the weak Pareto curve can be derived in a similar way. The next section will cover $m \times m$ games.

3.1 Developing analytical tools we will use later

We consider a generic 2×2 game

		agent 2	
		a_3	a_4
agent 1	a_1	$A = (A_1, A_2)$	$B = (B_1, B_2)$
	a_2	$C = (C_1, C_2)$	$D = (D_1, D_2)$

where A_i, B_i, C_i, D_i are payoffs of agent i. For simplicity, we denote by x the probability with which agent 1 plays a_1 and by y the probability with which agent 2 plays a_3. Our aim is to find a function $u_2(u_1)$ that describes the Pareto curve. To obtain that we proceed as follows. From $u_1 = u_1(x, y)$ (excluding the case $C_1 - D_1 + K_1 x = 0$, which can be treated separately), we can express y as a function of x and u_1:

$$y(x, u_1) = \frac{u_1 - D_1 + (D_1 - B_1)x}{C_1 - D_1 + K_1 x}$$

where K_i are constants (see below), under the constraints $x, y \in [0, 1]$ for each $u_1 \in [u_1^{\min}, u_1^{\max}]$ where $u_1^{\min}, u_1^{\max}$ are $u_1^{\min} = \min\{A_1, B_1, C_1, D_1\}$ and $u_1^{\max} = \max\{A_1, B_1, C_1, D_1\}$. In the general case, these constraints can be written as:

$$u_1 \in [u_1^{\min}, u_1^{\max}] \tag{1}$$

$$x \in [\underline{x}(u_1), \overline{x}(u_1)] \text{ or } \in [0, \underline{x}(u_1)] \cup [\overline{x}(u_1), 1] \tag{2}$$

where $\underline{x}(u_1), \overline{x}(u_1)$ with values from $[0, 1]$ are linear functions in u_1 whose form depends on the parameters A_1, B_1, C_1, D_1. For instance, in the case $C_1 - D_1 + K_1 x > 0$ for every $x \in [0, 1]$, $\underline{x}(u_1)$ and $\overline{x}(u_1)$ are determined to satisfy $0 \leq u_1 - D_1 + (D_1 - B_1)x$ and $1 \geq \frac{u_1 - D_1 + (D_1 - B_1)x}{C_1 - D_1 + K_1 x}$.

By substituting $y = y(x, u_1)$ in u_2, we obtain:

$$u_2(x, u_1) = \frac{D_1 C_2 - C_1 D_2 + (D_2 - C_2)u_1 + (K_5 - K_4 u_1)x_1 - K_3 x^2}{D_1 - C_1 - K_1 x}$$

In order to determine the function $u_2^*(u_1) = u_2(x^*(u_1), u_1)$ that describes the Pareto curve of the game, we aim to solve the optimization problem

$$x^*(u_1) = \arg\max_x \quad u_2(x, u_1) \qquad \text{s.t.} \quad \text{Constraint (2)}$$

In other words, we remove the dependency of $u_2(x, u_1)$ on x by substituting for x the function $x^*(u_1)$ that maximizes u_2 for each value of u_1 and satisfies Constraint (2). The function $x^*(u_1)$ can be found by studying the function u_2: we derive the derivative of u_2 w.r.t. x when u_1 is considered as a parameter and then we study how its sign changes as u_1 and as x vary according to Constraint (2). The derivative is

$$\frac{\partial u_2}{\partial x} = \frac{K_6 + K_7 - C_1 D_1 K_4 - K_2 u_1 + 2(C_1 - D_1)K_2 x + K_1 K_2 x^2}{(D_1 - C_1 - K_1 x)^2}$$

It equals zero for the following roots:

$$r_1, r_2 = \frac{C_1 - D_1}{K_1} \pm \frac{\sqrt{K_2 K_3 (B_1 C_1 - A_1 D_1 + K_1 u_1)}}{K_1 K_2}$$

Given a subrange of u_1, the maximum of u_2 can be for $x^*(u_1) = r_1$ or $x^*(u_1) = r_2$ if these roots are real values and are feasible according to Constraints 2 or it can be for some extreme value of the intervals prescribed by Constraints 2. Finally, we substitute $x = x^*(u_1)$ obtaining function $u_2^*(u_1)$ that describes the Pareto curve of the game. When $x^*(u_1) = r_1, r_2$, we have that $u_2^*(u_1)$ is a quadratic function:

$$u_2^*(u_1) = \frac{K_8 + K_9 + K_{10} - K_1 K_4 u_1}{K_1^2} \pm \frac{2\sqrt{K_2 K_3 (B_1 C_1 - A_1 D_1 + K_1 u_1)}}{K_1^2}$$

Instead, when $x^*(u_1)$ is equal to some extreme value of the intervals prescribed by Constraints 2, $u_2^*(u_1)$ is a linear function (the calculations are omitted here because they are long but trivial).

The same procedure adopted above to derive $u_2^*(u_1)$ can be used to derive $u_1^*(u_2)$. The constants K_i are as follows.

$K_1 = A_1 + D_1 - B_1 - C_1$
$K_2 = (A_1 - C_1)(B_2 - D_2) - (B_1 - D_1)(A_2 - C_2)$
$K_3 = (A_1 - B_1)(C_2 - D_2) - (C_1 - D_1)(A_1 - B_2)$
$K_4 = A_2 + D_2 - B_2 - C_2$
$K_5 = D_1 A_2 - C_1 B_2 + B_1 C_2 - 2D_1 C_2 - A_1 D_2 + 2C_1 D_2$
$K_6 = D_1(D_1(A_1 - C_2) + A_1(C_2 - D_2))$
$K_7 = C_1(C_1(B_2 - D_2) + B_1(D_2 - C_2))$
$K_8 = -(C_1 - D_1)(-D_1 A_1 + C_1 B_2) - B_1^2 C_2 - A_1^2 D_2$
$K_9 = B_1(D_1(A_2 + C_2) + C_1(-2A_2 + B_2 + C_2 - 2D_2) + A_1(C_2 + D_2))$
$K_{10} = A_1(C_1(B_2 + D_2) + D_1(A_2 - 2(B_2 + C_2) + D_2))$

In order to classify the different forms of the Pareto curve, the following properties are useful.

PROPOSITION 3.1. *The extremes (points in which u_1 is maximum or u_2 is maximum) of the Pareto curve correspond to pure strategy profiles.*

The proof is straightforward, the Pareto curve being the (non–convex) combination of the outcome payoff vectors.

Focusing on A, we have the following result.

THEOREM 3.2. *As $u_1 \to A_1^+$ (i.e., from the right), if the Pareto curve contains A and it is continuous from the right to A, then the Pareto curve is tangent to either segment $\overline{AB}$ or segment $\overline{AC}$.*

Proof. The proof follows from the application of the general envelope theorem [17]. In our specific case, that theorem reduces to distinguishing certain cases and applying the (unconstrained) envelope theorem to $u_1^*(u_2)$ and $u_2^*(u_1)$ disregarding the constraints. Specifically, the envelope theorem states that, given an optimization problem

$$f^*(z) = \max_w f(w,z),$$

where z is a parameter, we have

$$\frac{df^*(z)}{dz} = \left.\frac{\partial f(w,z)}{\partial z}\right|_{w=w^*(z)} \qquad \text{where} \quad w^*(z) = \arg\max_w f(w,z)$$

We assume the Pareto curve is continuous in A from the right. We distinguish three cases. If $B_1 > A_1$ and $A_1 > C_1$, then the derivative of the Pareto curve is $\frac{du_2^*(u_1)}{du_1}$. If $C_1 > A_1$ and $A_1 > B_1$, the derivative of the Pareto curve is $\frac{du_2^*(u_1)}{du_1}$. If $B_1 > A_1$ and $C_1 > A_1$, the derivative of the Pareto curve is $\min\left\{\frac{du_2^*(u_1)}{du_1}, \frac{du_1^*(u_2)}{du_2}\right\}$. The case $A_1 > B_1$ and $A_1 > C_1$ is not allowed when the Pareto curve is continuous in A. We now compute $\frac{du_2^*(u_1)}{du_1}$ and $\frac{du_1^*(u_2)}{du_2}$ according to the envelope theorem. We need to compute

$$\left.\frac{\partial u_2(x,u_1)}{\partial u_1}\right|_{x=x^*(u_1)} = \left.\frac{D_2 - C_2 - K_4 x}{D_1 - C_1 - K_1 x}\right|_{x=x^*(u_1)}$$

$$\left.\frac{\partial u_1(y,u_2)}{\partial u_2}\right|_{y=y^*(u_2)} = \left.\frac{D_1 - B_1 - K_4 y}{D_2 - B_2 - K_1 y}\right|_{y=y^*(u_2)}$$

and then to substitute $x = 1$ and $y = 1$ since we are studying the derivative of the Pareto curve in A. We obtain

$$\left.\frac{\partial u_2(x,u_1)}{\partial u_1}\right|_{x=1} = \frac{A_2 - B_2}{A_1 - B_1} \quad , \quad \left.\frac{\partial u_1(y,u_2)}{\partial u_2}\right|_{y=1} = \frac{A_1 - C_1}{A_2 - C_2}$$

That completes the proof. □

An analogous result holds for $u_1 \to A_1^-$ (i.e., from the left). Furthermore, we have analogous results for B, C, D: if an outcome $X \in \{A, B, C, D\}$ is on the Pareto curve and the Pareto curve is continuous to X (from the right and/or from the left), then as $u_1 \to X^{+/-}$ the Pareto curve is tangent to one of the segments whose extremes are X and an outcome in which only one agent deviates w.r.t. X. Therefore we can characterize the derivative of the Pareto curve around A, B, C, D when these points lie on such curve.

3.2 An example

We now show an example of the method discussed in the previous section. Consider the following matrices.

$$U_1 = \begin{pmatrix} 2 & 2.2 \\ 3 & 4 \end{pmatrix} \qquad U_2 = \begin{pmatrix} 4 & 2.6 \\ 1.3 & 1 \end{pmatrix}$$

Functions $u_1(x,y)$ and $u_2(x,y)$ are:

$$u_1(x,y) = (4(1-x) + 2.2x)(1-y) + (3(1-x) + 2x)y$$
$$u_2(x,y) = (1.6x + 1)(1-y) + (1.3(1-x) + 4x)y$$

Consider u_1 a parameter and derive y as function of (x, u_1):

$$y = \frac{2.1875 - 1.25u_1}{1.25 - x} + 2.25$$

Thus we can derive $u_2(x, u_1)$ as

$$u_2(x,u_1) = \frac{u_1(1.375x + 0.375) + x(4.075x - 5.825) - 2.75}{x - 1.25}$$

under the constraints

$$u_1 \in [2,4]$$
$$x \in [0,1]$$
$$y(x,u_1) \in [0,1] \longrightarrow \begin{cases} 3 - u_1 \le x \le 1 & 2 \le u_1 \le 2.2 \\ 3 - u_1 \le x \le 0.\overline{1}(20 - 5u_1) & 2.2 < u_1 \le 3 \\ 0 \le x \le 0.\overline{1}(20 - 5u_1) & 3 < u_1 \le 4 \end{cases}$$

The derivative of $u_2(x, u_1)$ w.r.t. x is

$$\frac{\partial u_2}{\partial x} = \frac{-2.09375u_1 + x(4.075x - 10.1875) + 10.0313}{(1.25 - 1.x)^2}$$

and its roots are

$$r_1 = 1.25 - 0.358401\sqrt{4u_1 - 7}$$
$$r_2 = 1.25 + 0.358401\sqrt{4u_1 - 7}$$

The function $x^*(u_1)$ that maximizes u_2 given u_1 is:

$$x^*(u_1) = \begin{cases} 3 - u_1 & 2 < u_1 \le 2.2 \\ 1.25 - 0.358401\sqrt{-7 + 4u_1} & 2.2 < u_1 \le 3 \\ 0.\overline{1}(20 - 5u_1) & 3 < u_1 \le 4 \end{cases}$$

Finally, the following function describes the Pareto curve.

$$u_2^*(u_1) = \begin{cases} 9.4 - 2.7u_1 & 2 < u_1 \le 2.2 \\ 4.3625 + 1.375u_1 - 2.92096\sqrt{-7 + 4u_1} & 2.2 < u_1 \le 3 \\ 4.55556 - 0.888889u_1 & 3 < u_1 \le 4 \end{cases}$$

At A_1^+ the Pareto curve is tangent to $\overline{AC}$, while at D_1^- the Pareto curve is tangent to $\overline{DB}$.

3.3 Classification

We now characterize all the kinds of Pareto curve that can occur in 2×2 games. We begin by defining a binary variable $\rho := ((r_1 \notin \mathbb{R}) \wedge (r_2 \notin \mathbb{R}))$ and the quantities $\overline{r} := \max\{r_1, r_2\}$ and $\underline{r} := \min\{r_1, r_2\}$.

We apply an affine transformation to achieve $A = (0,1)$ and $D = (1,0)$ so we can reduce the number of parameters in the analysis (recall that Pareto efficiency is invariant to affine transformations). We define the following conditions:

$$\phi_1 := \left(C_2 < \frac{1 - C_1}{B_1} \wedge \frac{C_1 + B_1 C_2 - 1}{C_1 - 1} \le B_2\right) \wedge \left(\forall u_1 \in [0,1] : \underline{r} \ge 1 \vee \underline{r} \le \frac{C_1 - u_1}{C_1 - 1} \vee \rho\right)$$

$$\phi_2 := \left(C_2 \ge \frac{1 - C_1}{B_1} \vee \frac{C_1 + B_1 C_2 - 1}{C_1 - 1} < B_2\right) \wedge \left(\forall u_1 \in [0,1] : \overline{r} \ge 1 \vee \overline{r} \le \frac{C_1 - u_1}{C_1 - 1} \vee \rho\right)$$

$$\psi_1 := (C_1 \le 0) \wedge \left(\forall u_1 \in [0,1] \; s.t. \; 0 < B_1 \le u_1 : \overline{r} \le \frac{C_1 - u_1}{C_1 - 1} \vee \overline{r} \ge 1 \vee \rho\right) \wedge \left(\forall u_1 \in [0,1] \; s.t. \; u_1 < B_1 < \frac{u_1 - C_1 u_1}{u_1 - C_1} : \overline{r} \le \frac{C_1 - u_1}{C_1 - 1} \vee \overline{r} \ge \frac{u_1}{B_1} \vee \rho\right) \wedge \left(\forall u_1 \in [0,1] \; s.t. \; \frac{u_1 - C_1 u_1}{u_1 - C_1} < B_1 < 1 : \overline{r} \le \frac{u_1}{B_1} \vee \overline{r} \ge \frac{C_1 - u_1}{C_1 - 1} \vee \rho\right)$$

$$\psi_2 := (C_1 > 0) \wedge \left(\forall u_1 \in [0,1] : \overline{r} \le \max\left\{0, \frac{C_1 - u_1}{C_1 - u_1}\right\} \vee \overline{r} \ge \min\left\{1, \frac{u_1}{B_1} \vee \rho\right\}\right)$$

$$\gamma_1 := (B_1 > 1 \wedge 1 - B_1 \le C_1) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ B_1 > \frac{u_1 - C_1 u_1}{u_1 - C_1} : \underline{r} \ge \frac{C_1 - u_1}{C_1 - 1} \vee \underline{r} \le \frac{u_1}{B_1} \vee \rho\right)$$
$$\gamma_2 := (B_1 \le 1 \vee 1 - B_1 > C_1) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ 0 < B_1 \le u_1 : \overline{r} \ge \frac{C_1 - u_1}{C_1 - 1} \vee \overline{r} \le 1 \vee \rho\right) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ u_1 < B_1 < \frac{C_1 - u_1}{C_1 - 1} : \overline{r} \ge \frac{C_1 - u_1}{C_1 - 1} \vee \overline{r} \le \frac{u_1}{B_1} \vee \rho\right) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ \frac{C_1 - u_1}{C_1 - 1} < B_1 < 1 : \overline{r} \ge \frac{u_1}{B_1} \vee \overline{r} \le \frac{C_1 - u_1}{C_1 - 1} \vee \rho\right)$$

We now apply a new affine transformation to achieve $A = (0,1)$ and $C = (1,0)$. We define the following conditions:

$$\beta := (\forall u_1 \in [0,1]\ s.t.\ D_1 \le u_1 < 1 : \underline{r} \ge 1 - u_1 \vee \underline{r} \le 0 \vee \rho) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ \frac{B_1}{B_1 - D_1 + 1} < u_1 < D1 : \underline{r} \ge 1 - u_1 \vee \underline{r} \le \frac{D_1 - u_1}{D_1 - B_1} \vee \rho\right) \wedge$$
$$\left(\forall u_1 \in [0,1]\ s.t.\ B_1 < u_1 < \frac{B_1}{B_1 - D_1 + 1} : \underline{r} \ge \frac{D_1 - u_1}{D_1 - B_1} \vee \underline{r} \le 1 - u_1 \vee \rho\right)$$

$$\zeta := \Big((D_1 < 0) \wedge (1 - B_1 < B_2 \le 1) \wedge (D_2 < 1) \wedge$$
$$(\forall u_1 \in [0,1]\ s.t.\ B_1 < u_1 < 1 : \underline{r} \ge 1 - u_1 \vee \underline{r} \le 0 \vee \rho)\Big) \vee$$
$$\Big((0 < D_1) \wedge ((1 - B_1 < B_2 < 1 - B_1 + D_1 \wedge D_2 \le B_1 + B_2 - D_1) \vee$$
$$(1 - B_1 + D_1 \le B_2 \le 1)) \wedge (\forall u_1 \in [0,1]\ s.t.\ B_1 < u_1 < 1 : \underline{r} \ge 1 - u_1 \vee \underline{r} \le 0 \vee \rho)\Big)$$

We will now classify the different forms of Pareto curves in 2×2 games on the basis of the number of Pareto efficient outcomes (corresponding to the number Pareto efficient pure strategy profiles). We report a graphical example of each possible form in Fig. 1. The classification is general: it is applicable to any game once affine transformations and/or transformations that switch agents and/or columns and/or rows of the bimatrix have been applied. So, for example, point A does not necessarily represent the top left entry of the matrix, as shown in the following example of transformation:

$$U_1 = \begin{pmatrix} 3 & 4 \\ 2 & 2.2 \end{pmatrix} \xrightarrow{\text{switch}} U_1' = \begin{pmatrix} 2 & 2.2 \\ 3 & 4 \end{pmatrix} \xrightarrow{\text{affine}} U_1'' = \begin{pmatrix} 0 & 0.1 \\ 0.5 & 1 \end{pmatrix}$$
$$U_2 = \begin{pmatrix} 1.3 & 1 \\ 4 & 2.8 \end{pmatrix} \xrightarrow{\text{switch}} U_2' = \begin{pmatrix} 4 & 2.8 \\ 1.3 & 1 \end{pmatrix} \xrightarrow{\text{affine}} U_2' = \begin{pmatrix} 1 & 0.6 \\ 0.1 & 0 \end{pmatrix}$$

One Pareto efficient outcome. There is 1 case.

1.1: 1–point Pareto curve. The Pareto curve is composed of only point A. This happens when A Pareto dominates all the other outcomes B, C, D (in the figure, B, C, D must be placed in the gray area): $A \ge B, C, D$.

Two Pareto efficient outcomes. There are 8 cases.

2.1: Continuous convex 1–segment Pareto curve. The Pareto curve is composed of a single segment connecting points A and B given by any possible randomization over these two points. This happens when A and B do not Pareto dominate each other and outcomes C and D are Pareto dominated by the segment (in the figure, C, D must be placed in the gray area): $B_1 \ge A_1$ and $A_2 \ge B_2$ and $\exists \lambda_1, \lambda_2 \in [0,1] : (\lambda_1 A + (1-\lambda_1) B \ge C) \wedge (\lambda_2 A + (1-\lambda_2) B \ge D)$.

2.2: Continuous non–convex 1–curve Pareto curve. The Pareto curve is composed of a single non–convex quadratic curve connecting points A and D generated by a continuous set of strategies fully randomizing over the four outcomes. This happens when A and D do not Pareto dominate each other, $B = C$, and outcomes B and C are in the triangle $AD(A_1, D_2)$ (in the figure, B and C must be in the gray area): $D_1 \ge A_1$ and $A_2 \ge D_2$ and $B = C$ and $B \in [A_1, D_1] \times [D_2, A_2]$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) D \ge B$.

2.3: Continuous non–convex 1–segment 1–curve Pareto curve. The Pareto curve is composed of a single continuous piece–wise non–convex curve connecting points A and D composed of a non–convex quadratic curve starting from A, given by a continuous set of strategies fully randomizing over the four outcomes, and of a segment ending in D, given by a continuous set of strategies randomizing over B and D. This happens when A and D do not dominate each other, B is in the triangle $AD(A_1, D_2)$ and C is after B on the line starting from A and touching B (in the figure, B must be in the light gray area while C must be in the light or dark gray areas) and satisfying condition $\neg(\psi_1 \vee \psi_2)$—this condition has no easy graphical interpretation: $D_1 \ge A_1$ and $A_2 \ge D_2$ and $B \in [A_1, D_1] \times [D_2, A_2]$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) D \ge B$ and $\exists \lambda \in [0,1] : \lambda B + (1-\lambda) D \ge C$ and $\neg(\psi_1 \vee \psi_2)$ and $\exists \lambda > 1 : C = \lambda(B - A) + A$.

2.4: Continuous non–convex 2–segment 1–curve Pareto curve. The Pareto curve is composed of a single continuous piece–wise non–convex curve connecting points A and D composed of a segment starting from A, given by a continuous set of strategies randomizing over A and B, a curve, given by a continuous set of strategies fully randomizing over the four outcomes, and, finally, of a segment ending in D, given by a continuous set of strategies randomizing over C and D. This happens when A and D do not dominate each other, B is in the triangle $AD(A_1, D_2)$ and C is placed below the segment connecting A and D (in the figure, B must be in the dark gray area while C must be in the light gray area) such that condition $\neg(\psi_1 \vee \psi_2)$ is satisfied: $D_1 \ge A_1$ and $A_2 \ge D_2$ and $B \in [A_1, D_1] \times [D_2, A_2]$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) D \ge B$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) D \ge C$ and $\neg(\psi_1 \vee \psi_2)$.

2.5: Discontinuous 1–segment 1–curve 1–point Pareto curve. The Pareto curve is discontinuous, being composed of a single continuous piece–wise non–convex curve (composed of a segment starting from A and a curve) and one point D. This happens when A and D do not dominate each other, B is in the area $[A_1, D_1] \times (-\infty, D_2]$, C is in the area $(-\infty, A_1] \times (-\infty, D_2]$ (in the figure, B must be in the dark gray area while C must be placed in the light gray area), and condition $\neg(\gamma_1 \vee \gamma_2)$ is satisfied: $D_1 \ge A_1$ and $A_2 \ge D_2$, and $B \in [A_1, D_1] \times (-\infty, D_2]$ and $C \in (-\infty, A_1] \times (-\infty, D_2]$ and $\neg(\gamma_1 \vee \gamma_2)$.

2.6: Discontinuous 1–segment 1–point Pareto curve. The Pareto curve is discontinuous, being composed of a single segment starting from A and one point D. This happens when A and D do not dominate each other, B is in the area $[A_1, D_1] \times (-\infty, D_2]$, C is in the area $(-\infty, A_1] \times (-\infty, D_2]$ (in the figure, B must be in the dark gray area while C must be in the light gray area), and condition $(\gamma_1 \vee \gamma_2)$ is satisfied: $D_1 \ge A_1$ and $A_2 \ge D_2$, and $B \in [A_1, D_1] \times (-\infty, D_2]$ and $C \in (-\infty, A_1] \times (-\infty, D_2]$ and $(\gamma_1 \vee \gamma_2)$.

2.7: Discontinuous 1–curve 2–point Pareto curve. The Pareto curve is discontinuous, being composed of a single curve and two points A and D. This happens when A and D do not dominate each other, B and C are in the area $(-\infty, A_1] \times (-\infty, D_2]$ (in the figure, B and C are in the gray area) and condition $\neg(\phi_1 \vee \phi_2)$ is satisfied: $D_1 \ge A_1$ and $A_2 \ge D_2$ and $B, C \in (-\infty, A_1] \times (-\infty, D_2]$ and $\neg(\phi_1 \vee \phi_2)$.

2.8: Discontinuous 2–point Pareto curve. The Pareto curve is discontinuous, being composed of two points A and D. This happens when A and D do not dominate each other, B and C are in the area $(-\infty, A_1] \times (-\infty, D_2]$ (in the figure,

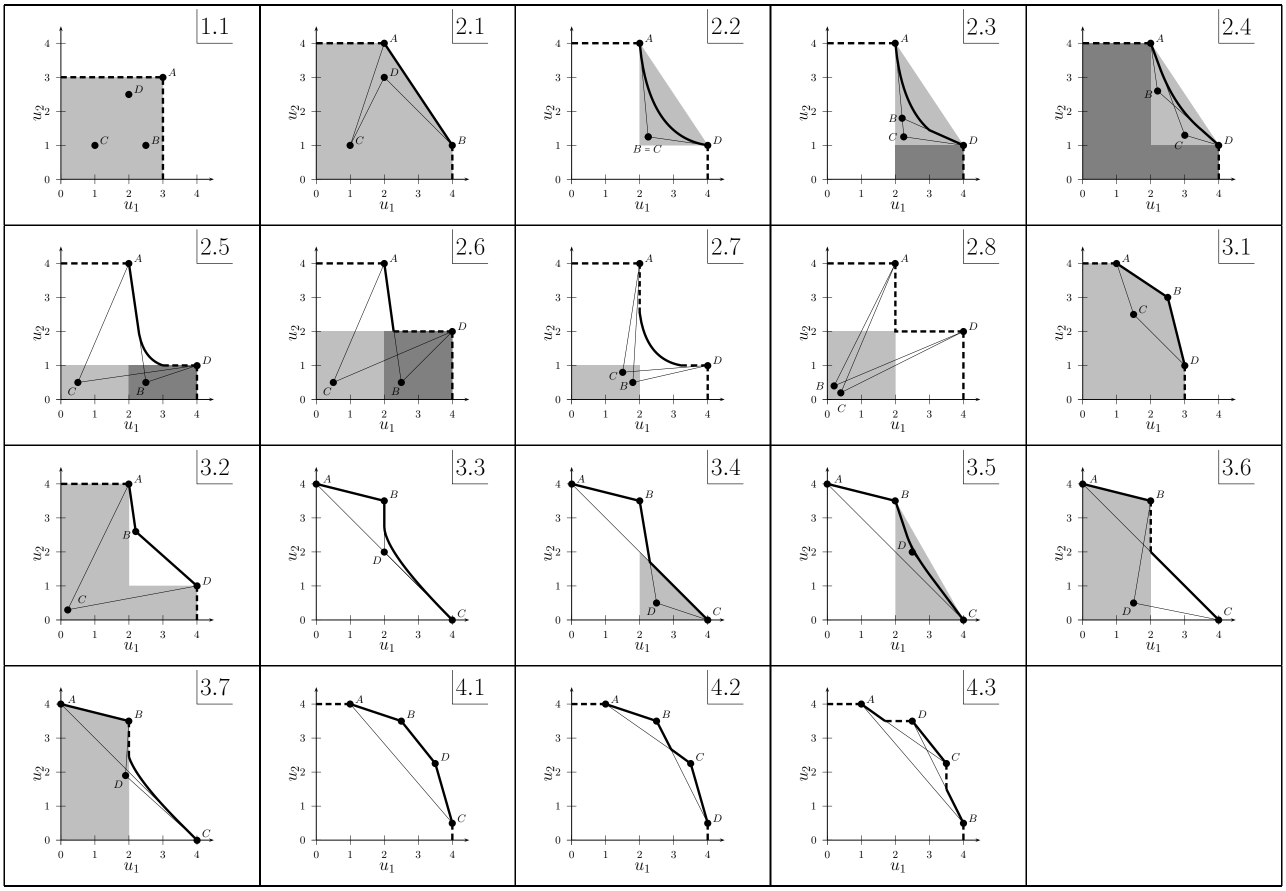

Figure 1: Classification of Pareto curve of bimatrix 2x2 games on the basis of the number of distinct pure–strategy Pareto efficient outcomes.

B and C are in the gray area) and condition $\phi_1 \vee \phi_2$ is satisfied: $D_1 \geq A_1$ and $A_2 \geq D_2$ and $B, C \in (-\infty, A_1] \times (-\infty, D_2]$ and $\phi_1 \vee \phi_2$.

Three Pareto efficient outcomes. There are 7 cases.

3.1: Continuous convex 2–segment Pareto curve. The Pareto curve is continuous, convex, and piece–wise linear. It is composed of a segment connecting A and B and a segment connecting B and D. This happens when A, B, D do not Pareto dominate each other, B is above the segment connecting A and D, and outcome C is Pareto dominated by the segments (in the figure, C must be in the gray area): $D_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq D_2$ and $\exists \lambda_1, \lambda_2 \in [0,1] : (\lambda_1 A + (1-\lambda_1)D \geq C) \wedge (B \geq \lambda_2 A + (1-\lambda_2)D)$.

3.2: Continuous non–convex 2–segment Pareto curve. The Pareto curve is continuous, non–convex, and piece–wise linear. It is composed of a segment connecting A and B and a segment connecting B and D. This happens when A, B, D do not Pareto dominate each other, B is below the segment connecting A and D, outcomes C is Pareto dominated by the segments (in the figure, C must be in the gray area) and condition $\psi_1 \vee \psi_2$ are satisfied: $D_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq D_2$ and $\exists \lambda_1, \lambda_2 \in [0,1] : (\lambda_1 A + (1-\lambda_1)D \geq C) \wedge (\lambda_2 A + (1-\lambda_2)D \geq B)$ and $\psi_1 \vee \psi_2$.

3.3: Continuous non–convex 2–segment 1–curve Pareto curve. The Pareto curve is continuous and non–convex piece–wise, composed of two segments (one connecting A and B and given by any possible randomization over these two outcomes, and one starting from B and directed to D given by a continuous set of strategies randomizing over these two outcomes) and a non–convex quadratic curve ending in C, given by a continuous set of strategies randomizing over all the four outcomes. This happens when A, B, C do not Pareto dominate each other, D is on the segment connecting A to C and with $D_1 = B_1$: $C_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq C_2$ and $\exists \lambda \in [0,1] : D = \lambda A + (1-\lambda)C$ and $D_1 = B_1$.

3.4: Continuous non–convex 3–segment Pareto curve. The Pareto curve is continuous, non–convex, and piece–wise linear. It is composed of three segments (one connecting A and B and given by any possible randomization over these two outcomes, one starting from B and directed to D given by a continuous set of strategies randomizing over these two outcomes, and one ending in C given by a continuous set of strategies randomizing over A and C). This happens when A, B, C do not Pareto dominate each other, D is below the segments connecting A to C and A to B and some additional conditions based on β are satisfied: $C_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq C_2$ and β.

3.5: Continuous non–convex 3–segment 1–curve Pareto curve. The Pareto curve is continuous and non–convex piece–wise. It is composed of three segments (one connecting A and B and given by any possible randomization over these two outcomes, one starting from B and directed to D given by a continuous set of strategies randomizing over these two outcomes, and one ending in C given by a continuous set of strategies randomizing over D and C) and one quadratic non–convex curve given by a continuous set of strategies randomizing over all the four outcomes. This happens when A, B, C do not Pareto dominate each other, D is below the segment connecting B to C and some additional conditions based on β are satisfied: $C_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq C_2$ and $\exists \lambda \in [0,1] : \lambda B + (1-\lambda) C \geq D$ and $\neg \beta$.

3.6: Discontinuous 2–segment Pareto curve. The Pareto curve is discontinuous, non–convex and piece–wise linear. It is composed of two segments (one connecting A and B and given by any possible randomization over these two outcomes, and one ending in C given by a continuous set of strategies randomizing over A and C). This happens when A, B, C do not Pareto dominate each other, D is below the segments connecting A to B and A to C and D is Pareto dominated by B and some additional conditions based on ζ are satisfied: $C_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq C_2$ and $B \geq D$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) C \geq D$ and ζ.

3.7: Discontinuous 2–segment 1–curve Pareto curve. The Pareto curve is discontinuous non–convex piece–wise, composed of two segments (one connecting A and B and given by any possible randomization over these two outcomes, and one ending in C given by a continuous set of strategies randomizing over A and C) and one quadratic non–convex curve given by a continuous set of strategies randomizing over all the four outcomes. This happens when A, B, C do not Pareto dominate each other, D is below the segments connecting A to B and A to C, and D is Pareto dominated by B, and some additional conditions based on ζ are satisfied: $C_1 \geq B_1 \geq A_1$ and $A_2 \geq B_2 \geq C_2$ and $B \geq D$ and $\exists \lambda \in [0,1] : \lambda A + (1-\lambda) C \geq D$ and $\neg \zeta$.

Four Pareto efficient outcomes. There are 3 cases.

4.1: Continuous convex 3–segment Pareto curve. The Pareto curve is continuous, convex, and piece–wise linear. It is composed of three segments (one connecting A and B and given by any possible randomization over these two outcomes, one connecting B and D and given by any possible randomization over these two outcomes, and one connecting D and C and given by any possible randomization over these two outcomes). This happens when A, B, C, D do not Pareto dominate each other and they are placed in the order A, B, D, C as u_1 increases: $C_1 \geq D_1 \geq B_1 \geq A_1$ and $C_2 \leq D_2 \leq B_2 \leq A_2$.

4.2: Continuous non–convex 4–segment Pareto curve. The Pareto curve is continuous, non–convex, and piece–wise linear. It is composed of four segments (one connecting A and B and given by any possible randomization over these two outcomes, one starting from B and given by a continuous set of strategies randomizing over B and D, one ending in C and given by a continuous set of strategies randomizing over A and C, and one connecting C and D). This happens when A, B, C, D do not Pareto dominate each other and they are placed in the order A, B, C, D as u_1 increases: $D_1 \geq C_1 \geq B_1 \geq A_1$ and $D_2 \leq C_2 \leq B_2 \leq A_2$.

4.3: Discontinuous 3–segment Pareto curve. The Pareto curve is discontinuous, non–convex, and piece–wise linear. It is composed of three segments (one starting from A and given by a continuous set of strategies randomizing over A and C, one connecting D and C and given by any possible randomization over these two outcomes, and one ending in B and given by a continuous set of strategies randomizing over D and B. This happens when A, B, C, D do not Pareto dominate each other and they are placed in the order A, D, C, B as u_1 increases: $B_1 \geq C_1 \geq D_1 \geq A_1$ and $B_2 \leq C_2 \leq D_2 \leq A_2$.

THEOREM 3.3. *The above cases constitute all the possible cases for the Pareto curve of a 2×2 bimatrix game.*

The proof is omitted since the calculations are straightforward and long. The intuition follows. For each case in Fig. 1, we applied the tools discussed in Section 3.1 with symbolic value parameters and we symbolically derived the above conditions; for each case not reported in Fig. 1, the approach is the same, but we proved that there is no feasible assignment to the parameters.

4. PARETO CURVE OF M×M GAMES

We denote by $P(U)$ the points $\mathbf{u}$ of the agents' utilities space that are Pareto efficient given the bimatrix U. On the basis of the results discussed in the previous section, in the case U is a 2x2 bimatrix, $P(U)$ can be represented as a piece–wise function with at most 4 pieces as follows:

$$P(U) = \begin{cases} f_1^U(u_1) & u_1 \in R_1^U \\ f_2^U(u_1) & u_1 \in R_2^U \\ f_3^U(u_1) & u_1 \in R_3^U \\ f_4^U(u_1) & u_1 \in R_4^U \end{cases}$$

where $\{R_j^U\}_{j \in \{1,2,3,4\}}$ is a set of (potentially open or degenerate single–point) intervals partitioning $[(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$ where $(u_2^*)^{-1}$ is the inverse function of u_2^*; f_j^U is a monotonically decreasing quadratic function or a special pair $(\varnothing, \overline{u}_2)$ denoting that the Pareto curve is not defined over such interval and that the Pareto dominator has value $u_2 = \overline{u}_2$.

EXAMPLE 4.1. *Consider Case 4.2 shown in Fig. 1. The ranges are $R_1^U = [1, 2.5]$, $R_2^U = (2.5, 2.925]$, $R_3^U = (2.925, 3.5]$, $R_4^U = (3.5, 4]$; the functions f_i^U are linear in u_1.*

EXAMPLE 4.2. *Consider Case 4.3 shown in Fig. 1. The ranges are $R_1^U = [1, 2.5]$, $R_2^U = (2.5, 2.925)$, $R_3^U = [2.925, 3.5]$, $R_4^U = (3.5, 4]$; the functions f_1^U, f_3^U, f_4^U are linear in u_1 and $f_2^U = (\varnothing, 3.5)$, the Pareto curve not being defined over R_2^U and the Pareto dominator (i.e., D) has value $u_2 = 3.5$.*

EXAMPLE 4.3. *Consider Case 3.5 shown in Fig. 1. The ranges are $R_1^U = [2, 2.25]$, $R_2^U = (2.25, 3)$, $R_3^U = (3, 4)$, $R_4^U = [4, 4]$; the function f_1^U is linear, f_2^U is quadratic, $f_3^U = (\varnothing, 1)$, $f_4^U = 1$.*

We now introduce two filtering functions. The first one, $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1), \ldots, P(U_k)\})$, returns the points of curve $P(U)$ that are not strongly Pareto dominated by any point of any Pareto curve of the collection $\{P(U_1), \ldots, P(U_k)\}$. The second one, $\mathsf{subtract}_{\text{weak}}(P(U), \{P(U_1), \ldots, P(U_k)\})$, returns the points of $P(U)$ that are not weakly Pareto dominated by any point of any Pareto curve of the collection $\{P(U_1), \ldots, P(U_k)\}$. Formally:

DEFINITION 4.4. *Given bimatrix U and a set of bimatrices $\{U_1, \ldots, U_k\}$, $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1), \ldots, P(U_k)\})$ $:= \{\mathbf{u} \in P(U) : \forall j, \nexists\, \mathbf{u}' \in P(U_j), \mathbf{u}' > \mathbf{u}\}$.*

DEFINITION 4.5. *Given bimatrix U and a set of bimatrices $\{U_1, \ldots, U_k\}$, $\mathsf{subtract}_{\text{weak}}(P(U), \{P(U_1), \ldots, P(U_k)\})$ $:= \{\mathbf{u} \in P(U) : \forall j, \nexists\, \mathbf{u}' \in P(U_j), \mathbf{u}' \geq \mathbf{u}\}$.*

We study the number of pieces returned by $\mathsf{subtract}_{\text{strong}}$ as k and the size of the bimatrix vary. (The same results can be obtained for $\mathsf{subtract}_{\text{weak}}$.)

LEMMA 4.6. *When $k = 1$ and U and U_1 are 2×2 bimatrices, the number of pieces in $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1)\})$ is no greater than 24.*

Proof. The intervals $\{R_j^U\}$ are in general different from the intervals $\{R_j^{U_1}\}$. However, by the intersection of these intervals, we can generate at most 8 intervals, say $\{R_h^{U,U_1}\}$, partitioning the compact set $[(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$ over which $P(U)$ is defined and such that at each interval we have a different pair $(f_j^U, f_{j'}^{U_1})$. Given that f_j^U and $f_{j'}^{U_1}$ are quadratic monotonically decreasing functions, there can be at most 2 intersections (computable exactly) between f_j^U and $f_{j'}^{U_1}$ over each interval R_h^{U,U_1}—excluding the degenerate case in which the number of intersections is infinite, $f_j^U = f_{j'}^{U_1}$ over the range R_h^{U,U_1}. In the case the Pareto curve of U (or U_1) is not defined over the interval, the value $\overline{u}_2$ of f_j^U (or $f_j^{U_1}$) is used to determine the intersections. We use the intersections between f_j^U and $f_{j'}^{U_1}$ to divide each interval R_h^{U,U_1} into at most 3 subintervals. This way, in each subinterval there is no intersection between the two functions. An obvious upper bound on the number of subintervals per interval R_h^{U,U_1} over which $f_j^U > f_j^{U_1}$ is 3. Thus, $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1)\})$ is composed of at most 3 subintervals per interval R_h^{U,U_1} and therefore an upper bound over the total number of pieces is $3 \times 8 = 24$. Let use remark that R_h^{U,U_1} can be degenerate single–point intervals. ◻

Lemma 4.7. *When U and U_j for every j are 2×2 bimatrices, the number of pieces in $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1), \dots, P(U_k)\})$ is at most $24k$.*

Proof. For each pair (U, U_j) we apply the same procedure adopted in the proof of Lemma 4.6, obtaining 8 intervals R_h^{U,U_j} and then dividing each of these intervals into 3 subintervals by the intersections of f_h^U and $f_{h_j}^{U_j}$. The total number of subintervals per (U, U_j) is 24. Now, different pairs (U, U_j) have in general different subintervals, but, by the intersection of these intervals, we can obtain at most $24k$ different subintervals partitioning the compact set $[(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$ over which $P(U)$ is defined and such that there is no intersection between f_h^U and any $f_{h_j}^{U_j}$. $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1), \dots, P(U_k)\})$ is composed of a subset of these subintervals and therefore the number of pieces of $\mathsf{subtract}_{\text{strong}}(P(U), \{P(U_1), \dots, P(U_k)\})$ cannot be greater than $24k$. ◻

Now we use the above lemma to prove that the Pareto curve of a 2–agent game is composed of a number of pieces that is polynomial in the size of the game. Given U, we denote by U_{a_1,a'_1,a_2,a'_2} the 2×2 sub–bimatrix of U in which agent 1 plays actions a_1 and a'_1 and agent 2 plays actions a_2 and a'_2.

Theorem 4.8. *The Pareto curve of a 2–agent game with bimatrix U of size $m\times m$ is composed of at most $24m^8$ pieces, and each piece can be described by a quadratic function.*

Proof. In [8] the authors show that, for any $m\times m$ bimatrix game U, given $\mathbf{u}$, if there exists $\mathbf{u}'$ that Pareto dominates $\mathbf{u}$, then $\mathbf{u}' \in P(U_{a_1,a'_1,a_2,a'_2})$ for some a_1, a'_1, a_2, a'_2. In other words, if $\mathbf{u} \in \overline{U}$, where $\overline{U}$ is an $l\times l$ sub–bimatrix of U with $l \geq 3$, and $\mathbf{u} \in P(U)$, then there exists some a_1, a'_1, a_2, a'_2 such that $\mathbf{u} \in P(U_{a_1,a'_1,a_2,a'_2})$. Therefore, in order to characterize the Pareto curve of a bimatrix game U, we can safely study only the Pareto curves of the 2×2 sub–bimatrices of U. Consider a 2×2 sub–bimatrix $\overline{U} = U_{\overline{a}_1,\overline{a}'_1,\overline{a}_2,\overline{a}'_2}$ of U. All the pieces in $\mathsf{subtract}_{\text{strong}}(P(\overline{U}), \{P(U_{a_1,a'_1,a_2,a'_2}) : (a_1 \neq \overline{a}_1) \vee (a'_1 \neq \overline{a}'_1) \vee (a_2 \neq \overline{a}_2) \vee (a_2 \neq \overline{a}'_2)\})$ belong to the Pareto curve $P(U)$ of the game given that their points are not Pareto dominated by the Pareto curve of any other 2×2 sub–bimatrix of U. The number of U_{a_1,a'_1,a_2,a'_2} such that $(a_1 \neq \overline{a}_1) \vee (a'_1 \neq \overline{a}'_1) \vee (a_2 \neq \overline{a}_2) \vee (a_2 \neq \overline{a}'_2)$ is at most $24m^4$. For each $\overline{U}$, there are at most $24m^4$ pieces. Therefore, the total number of pieces for all the possible $\overline{U}$—that are at most $24m^4$—is at most $24m^8$ by arguments similar to those used in the proof of Lemma 4.7. In addition, each piece of the Pareto curve of an $m\times m$ bimatrix game U is quadratic since it is a piece of a 2×2 sub–bimatrix of U. ◻

Since the computation of each subinterval in the proofs of Theorem 4.8 can be done in polynomial time in the size of the game (we only need to determine the intersection of different known intervals and solve closed–form quadratic equations), we have the following.

Corollary 4.9. *The Pareto curve of a 2–agent game can be found in polynomial time.*

5. ALGORITHMS FOR SOLUTION CONCEPTS BASED ON PARETO CURVE

We now leverage the ability to compute the Pareto curve to design algorithms for game-theoretic solution concepts that are based on Pareto efficiency.

Theorem 5.1. *The utilitarian bargaining solution (UBS) [16] of an $m\times m$ bimatrix game can be found in polynomial time.*

Proof. From Theorem 4.8 and Corollary 4.9, we can derive in polynomial time in m a number of pieces, polynomial in m, constituting the Pareto curve. To compute the UBS, for each piece, we need to compute the value of u_1 that maximizes $u_1 + u_2^*(u_1)$ under the constraint $u_1 \in [(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$. This is maximization of a quadratic single–variable function, and can be solved in closed form in constant time. Once we have found the maximum for each piece, we need to find the maximum among all of them. So, the UBS is found in $O(m^8)$ time. ◻

Theorem 5.2. *The Nash bargaining solution (NBS) [13] of an $m\times m$ bimatrix game can be found in polynomial time.*

Proof. To compute the NBS, for each piece, we need to compute the value of u_1 that maximizes $u_1 \cdot u_2^*(u_1)$ under the constraint $u_1 \in [(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$. This is maximization of a cubic single–variable function that can be solved in closed form in constant time. Once we have found the maximum for each piece, we need to find the maximum among all of them. So, the NBS is found in $O(m^8)$ time. ◻

Theorem 5.3. *The Kalai–Smorodinsky bargaining solution (KSBS) [11] of an $m\times m$ bimatrix game can be found in polynomial time (if one exists).*

Proof. To find a KSBS, for each piece, we need to find the value of u_1 such that $\frac{u_1}{u_2^*(u_1)} = \frac{u_1^{\max}}{u_2^{\max}}$ under the constraint $u_1 \in [(u_2^*)^{-1}(u_2^{\max}), u_1^{\max}]$. To do this, we need to solve the equation $u_1 = u_2^*(u_1)\frac{u_1^{\max}}{u_2^{\max}}$ that is a quadratic equation in u_1, and therefore can be solved in closed form in constant time. Hence the KSBS (if one exists) is found in time $O(m^8)$. ◻

Theorem 5.4. *Consider the following algorithm:*

1. *Enumerate all pure–strategy profiles, and for each of them verify whether it is a strong Nash equilibrium (SNE) [1]. (A polynomial–time verification routine has recently been introduced [8].)*

2. *Enumerate all the pieces of the Pareto curve that are linear (including the degenerate single–point piece), and check each of them to see whether there exists a Nash equilibrium on the curve (by using MIP Nash [15] with an additional constraint $\alpha u_1 + u_2 + q = 0$ where α and q are the parameters of the linear curve).*

On bimatrix games, this algorithm finds an SNE whenever an SNE exists, and reports that one does not exist otherwise.

Proof. The first phase of the algorithm finds all pure–strategy SNEs, if any. The second phase finds a mixed–strategy SNE, if one exists via the following reasoning. A Nash equilibrium is an SNE if and only if it is on the Pareto curve and it was proven in [9] that, for every mixed–strategy SNEs, the payoffs associated with the outcomes played with strictly positive probability are on a line. Thus, the Pareto curve in a neighborhood of mixed–strategy SNEs is linear except when the piece degenerates into a single point. Since the algorithm enumerates all these pieces, if there is an SNE, then the algorithm finds it. (We recall that MIP Nash [15] finds an NE on the Pareto curve segment if one exists.) □

Regarding the complexity of the above algorithm, one can verify whether a strategy profile is an SNE in polynomial time [8]. Therefore the first phase of the algorithm has polynomial complexity, because the number of pure strategy profiles is polynomial in the size of the game (for any constant number of agents). However, in the second phase, the complexity of deciding whether there exists a Nash equilibrium in a subspace of the utility space is $\mathcal{NP}$–complete (the reduction is the same used in [4]). By Corollary 4.9, the algorithm calls this $\mathcal{NP}$–complete oracle a polynomial number of times. Therefore the algorithm can efficiently find an SNE if there is an efficient oracle for $\mathcal{NP}$–complete problems. Recall that the complexity of deciding whether an SNE exists is $\mathcal{NP}$–complete [4, 8], so there cannot exist a polynomial–time algorithm for finding SNE unless $\mathcal{P} = \mathcal{NP}$. Our algorithm is such an algorithm if $\mathcal{P} = \mathcal{NP}$. In contrast, prior algorithms for finding SNE take exponential time even if $\mathcal{P} = \mathcal{NP}$. In [9] all the $O(4^m)$ support profiles are scanned, and in [8] an $\mathcal{NP}$–hard oracle is called potentially exponentially many times.

6. CONCLUSIONS AND FUTURE WORK

Pareto efficiency is a widely used property in solution concepts for cooperative and non–cooperative game theory problems. However, finding and even approximating the Pareto curve is hard. The literature has focused on computing concise representations that approximate the Pareto curve or on exploiting evolutionary approaches to generate approximately Pareto efficient samples of the curve. In this paper, we showed that the Pareto curve of a bimatrix game can be found exactly in polynomial time. We classified all possible forms of the Pareto curve of 2×2 bimatrix games. We then showed how the Pareto curve of a $m \times m$ bimatrix game can be derived. It is composed of a polynomial number of pieces and each piece is a quadratic function. We provided a polynomial–time algorithm for computing it. Finally, we leveraged these results to provide algorithms for computing solutions to games according to solution concepts that use Pareto efficiency in their definition.

In future research we plan to investigate whether the problem stays easy with more than two agents and in polymatrix games that exhibit similar properties to bimatrix games. One could also study whether tighter bounds on the number of pieces of the Pareto curve for bimatrix games than those we provided in this paper can be proven, and to experimentally see the forms of the Pareto curves that occur in practice in various kinds of generated game instances.

7. ACKNOWLEDGEMENTS

Tuomas Sandholm was supported by the National Science Foundation under grants IIS-0964579, IIS-1320620, and CCF-1101668.

8. REFERENCES

[1] R. Aumann. Acceptable points in games of perfect information. *PAC J MATH*, 10:381–417, 1960.

[2] H. Aziz, F. Brandt, and P. Harrenstein. Pareto optimality in coalition formation. In *SAGT*, pages 93–104, 2011.

[3] K. Bringmann, T. Friedrich, F. Neumann, and M. Wagner. Approximation-guided evolutionary multi-objective optimization. In *IJCAI*, pages 1198–1203, 2011.

[4] V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *GAME ECON BEHAV*, 63(2):621–641, 2008.

[5] T. Demuynck. The computational complexity of rationalizing pareto optimal choice behavior. *SOC CHOICE WELFARE*, pages 1–21, 2013.

[6] M. Ehrgott, L. Shao, and A. Schöbel. An approximation algorithm for convex multi–objective programming problems. *J GLOBAL OPTIM*, 50(3):397–416, 2011.

[7] N. Gatti, M. Rocco, and T. Sandholm. Algorithms for strong Nash equilibrium with more than two agents. In *AAAI*, 2013.

[8] N. Gatti, M. Rocco, and T. Sandholm. On the verification and computation of strong Nash equilibrium. In *AAMAS*, 2013.

[9] N. Gatti, M. Rocco, and T. Sandholm. Strong Nash equilibrium is in smoothed $\mathcal{P}$. In *AAAI LBT*, 2013.

[10] C. Glasser, C. Reitwiessner, H. Schmitz, and M. Witek. Approximability and hardness in multi-objective optimization. In *CiE*, pages 180–189, 2010.

[11] E. Kalai and M. Smorodinsky. Other solutions to Nash's bargaining problem. *ECONOMETRICA*, 43(3):513–518, 1975.

[12] J. Legriel, C. L. Guernic, S. Cotton, and O. Maler. Approximating the Pareto front of multi-criteria optimization problems. In *TACAS*, pages 69–83, 2010.

[13] J. Nash. The bargaining problem. *ECONOMETRICA*, 18(2):155–162, 1950.

[14] C. H. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *FOCS*, pages 86–92, 2000.

[15] T. Sandholm, A. Gilpin, and V. Conitzer. Mixed–integer programming methods for finding Nash equilibria. In *AAAI*, pages 495–501, 2005.

[16] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*. Cambridge University Press, 2008.

[17] K. Sydsaeter and P. Hammond. *Essential Mathematics for Economic Analysis*. Prentice Hall, 2008.