

Bridging the Gap Between Single and Multi Objective Games

Willem Röpke
Artificial Intelligence Lab
Vrije Universiteit Brussel
Belgium
willem.ropke@vub.be

Carla Groenland
Department of Mathematics
Universiteit Utrecht
The Netherlands
c.e.groenland@uu.nl

Roxana Rădulescu
Artificial Intelligence Lab
Vrije Universiteit Brussel
Belgium
roxana.radulescu@vub.be

Ann Nowé
Artificial Intelligence Lab
Vrije Universiteit Brussel
Belgium
ann.nowe@vub.be

Diederik M. Roijers
Artificial Intelligence Lab
Vrije Universiteit Brussel
Belgium
diederik.roijers@vub.be

ABSTRACT

A classic model to study strategic decision making in multi-agent systems is the normal-form game. This model can be generalised to allow for an infinite number of pure strategies leading to continuous games. Multi-objective normal-form games are another generalisation that model settings where players receive separate payoffs in more than one objective. We bridge the gap between the two models by providing a theoretical guarantee that a game from one setting can always be transformed to a game in the other. We extend the theoretical results to include guaranteed equivalence of Nash equilibria. The mapping makes it possible to apply algorithms from one field to the other. We demonstrate this by introducing a fictitious play algorithm for multi-objective games and subsequently applying it to two well-known continuous games. We believe the equivalence relation will lend itself to new insights by translating the theoretical guarantees from one formalism to another. Moreover, it may lead to new computational approaches for continuous games when a problem is more naturally solved in the succinct format of multi-objective games.

KEYWORDS

Game theory; Continuous game; Multi-objective; Nash equilibria

ACM Reference Format:

Willem Röpke, Carla Groenland, Roxana Rădulescu, Ann Nowé, and Diederik M. Roijers. 2023. Bridging the Gap Between Single and Multi Objective Games. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Connecting seemingly unrelated models can bridge together research communities and complement missing pieces in either setting. To that extent, we identify two well-known extensions of normal-form games and show that there are underlying equivalences which may be exploited for theoretical and algorithmic contributions. On the one hand, we consider continuous games which allow for an infinite amount of actions for players [25]. An example of such games can be found in economic models where

firms have to set a price over a continuous range which maximises their profit in a competitive environment [13]. On the other hand, we consider Multi-Objective Normal-Form Games (MONFGs) [2], which extend normal-form games by returning a vectorial payoff rather than a scalar payoff. Multi-objective games can for example be applied for scheduling household appliances in residential buildings with the objectives to minimise electricity cost while adhering as much as possible to the desired operating period [15].

Our theoretical results establish a novel equivalence relation between continuous games and MONFGs. The equivalence is shown to leave the underlying game dynamics intact, including Nash equilibria. The established connection allows one to straightforwardly transfer known results and algorithms from either model to the other, thus bridging the gap between them. From a theoretical perspective, much more is known about continuous games than MONFGs, enabling rapid advances in the latter model. Conversely, due to their succinct format, MONFGs are amenable to general algorithmic solutions which can then be applied to continuous games. The main contributions are summarised as follows¹:

- We define pure strategy equivalence between a continuous game and MONFG. This equivalence relation formalises the existence of a bijective function that maps pure strategies from the continuous game to mixed strategies in the MONFG and for which the utilities remain equal.
- We show that for every continuous game with convex strategy sets, a pure strategy equivalent MONFG can be constructed. We also show this in the other direction.
- We introduce hierarchical strategies for MONFGs and define mixed strategy equivalence between these strategies and the mixed strategies in a continuous game.
- We show that a pure strategy Nash equilibrium in a continuous game is a mixed strategy equilibrium in an equivalent MONFG. We extend this for mixed strategy equilibria in the continuous game and hierarchical equilibria in the MONFG.
- We guarantee the existence of a hierarchical Nash equilibrium in MONFGs with continuous utility functions.
- We demonstrate the algorithmic value of our equivalence notion by computing a Nash equilibrium in two continuous games, namely a polynomial game and a Bertrand price game, using a fictitious play algorithm for MONFGs.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹The supplementary material of this work can be found in [24].

2 BACKGROUND

2.1 Continuous Games

Continuous games, sometimes referred to as infinite games, extend the normal-form game model to include settings where players have a nonempty compact metric space of pure strategies, rather than a finite set. Intuitively, this means that players may have an infinite number of actions to choose from. We define a continuous game as follows [25],

Definition 2.1 (Continuous game). A continuous game is a tuple (N, \mathcal{S}, v) , where:

- N is a finite set of n players, indexed by i ;
- $\mathcal{S} = S_1 \times \dots \times S_n$, where S_i is a nonempty compact metric space of pure strategies available to player i ;
- $v = (v_1, \dots, v_n)$ where $v_i : \mathcal{S} \rightarrow \mathbb{R}$ is a real-valued and continuous utility function for player i .

The set of mixed strategies for player i in a continuous game is defined as the set of Borel probability measures on S_i and is denoted by $\mathcal{B}(S_i)$. We define the expected utility for player i for a given joint mixed strategy $\mu \in \mathcal{B}(\mathcal{S})$ as follows [4, 25],

$$v_i(\mu) = \int_{S_1 \times \dots \times S_n} v_i(s_1, \dots, s_n) d\mu_1(s_1) \dots d\mu_n(s_n). \quad (1)$$

We refer readers unfamiliar with these concepts from measure theory to a brief treatment of them in Appendix A.

2.2 Multi-Objective Normal-Form Games

Multi-Objective Normal-Form Games (MONFGs) are a generalisation of (single-objective) normal-form games to vectorial payoffs. This is formalised as follows [19].

Definition 2.2 (Multi-objective normal-form game). A (finite, n -player) multi-objective normal-form game is a tuple $(N, \mathcal{A}, \mathbf{p})$, with d objectives, where:

- N is a finite set of n players, indexed by i ;
- $\mathcal{A} = A_1 \times \dots \times A_n$, where A_i is a nonempty finite set of actions available to player i ;
- $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ where $\mathbf{p}_i : \mathcal{A} \rightarrow \mathbb{R}^d$ is the vectorial payoff function for player i .

A mixed strategy for any given player is defined as a probability distribution over their set of actions. A set of probability distributions over a finite number of points is known as a probability simplex, formally defined below.

Definition 2.3. A probability k -simplex Δ^k is a set of points for which,

$$\Delta^k = \left\{ (x_0, \dots, x_k) \in \mathbb{R}^{k+1} \mid \sum_{i=0}^k x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 0, \dots, k \right\}.$$

As such, a mixed strategy for player i is a probability distribution $\delta_i \in \Delta^{k_i}$ where $k_i = |A_i| - 1$. Let $\Delta = \Delta^{k_1} \times \dots \times \Delta^{k_n}$ be the set of joint mixed strategies. The expected payoff for player i of a mixed strategy $\delta \in \Delta$ is then naturally defined as,

$$\mathbf{p}_i(\delta) = \sum_{a \in \mathcal{A}} \mathbf{p}_i(a) \prod_{j=1}^n \delta_j(a_j). \quad (2)$$

Multi-objective decision making frequently assumes a utility-based approach [21], which is also a standard game-theoretic approach. In the MONFG model, the utility of a player is based on a particular trade-off of the various payoffs, that is, we make the additional assumption that a utility function $u_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is known for any player i . We compute the utility of a mixed strategy δ as,

$$u_i(\mathbf{p}_i(\delta)) = u_i \left(\sum_{a \in \mathcal{A}} \mathbf{p}_i(a) \prod_{j=1}^n \delta_j(a_j) \right). \quad (3)$$

We note that an alternative definition exists for the utility of mixed strategies in multi-objective games, where players first apply the utility function to the payoffs and subsequently compute their expected utility. In the multi-objective decision making literature, this method is referred to as the expected scalarised returns criterion, while the method from Eq. (2) is referred to as the scalarised expected returns criterion [9, 18].

2.3 Nash Equilibria

We consider a central solution concept in both models, namely the Nash equilibrium (NE). Informally, a Nash equilibrium is a joint strategy from which no player can unilaterally deviate and improve their utility. Definition 2.4 defines this in continuous games, while Definition 2.5 defines this in MONFGs.

Definition 2.4 (Nash equilibria in continuous games). A mixed strategy profile μ^* is a Nash equilibrium if,

$$v_i(\mu_i^*, \mu_{-i}^*) \geq v_i(\mu_i, \mu_{-i}^*),$$

for all players i and mixed strategies $\mu_i \in \mathcal{B}(S_i)$.

Definition 2.5 (Nash equilibria in multi-objective normal-form games). A mixed strategy profile δ^* is a Nash equilibrium if,

$$u_i(\delta_i^*, \delta_{-i}^*) \geq u_i(\delta_i, \delta_{-i}^*),$$

for all players i and mixed strategies $\delta_i \in \Delta^{k_i}$.

An important early result is that every continuous game must have a mixed strategy Nash equilibrium [7]. In MONFGs however, Nash equilibria are not guaranteed to exist under nonlinear utility functions [19]. In Section 3.3, we expand further on this issue and provide a novel Nash equilibrium existence result for MONFGs.

3 EQUIVALENCE RELATION

Our main contribution establishes an equivalence relation between continuous and multi-objective games. With this goal in mind, we first introduce a special game, called an identity game, which plays a crucial role in bridging the two models. Next, we prove that a bijective mapping for pure strategies in the continuous game to mixed strategies in the MONFG always exists. To complete the full equivalence, we introduce a novel concept for MONFGs, which we call hierarchical strategies. Lastly, we show that Nash equilibria are preserved in the process, allowing us to guarantee a Nash equilibrium in hierarchical strategies in MONFGs.

3.1 Identity Game

An identity game returns, as the name suggests, a payoff vector equal to the strategy it received as input. We will use such games

in Section 3.2 to prove an equivalent MONFG can be constructed for every continuous game and vice versa.

LEMMA 3.1 (IDENTITY GAME). *For any finite set of players and finite sets of pure strategies, there exists a set of payoff functions \mathbf{p} such that for each player i , $\mathbf{p}_i(\delta) = \delta$.*

PROOF SKETCH. The main idea behind the proof is to define the payoff vectors for all pure strategies as the pure strategy itself. Note that, strictly speaking, this relies on the notation of joint strategies, which we define as a vector of concatenated individual strategies. It is then possible to show, using the law of total probability, that this results in the desired property. \square

A formal proof is given in Appendix B. To illustrate the payoff mechanism in more detail, we provide an example below.

	A	B
A	(1, 0, 1, 0); (1, 0, 1, 0)	(1, 0, 0, 1); (1, 0, 0, 1)
B	(0, 1, 1, 0); (0, 1, 1, 0)	(0, 1, 0, 1); (0, 1, 0, 1)

Figure 1: The identity game for a 2-player 2-action setting.

Example 1. Consider the identity game in Fig. 1. Assume that player one plays the mixed strategy $\delta_1 = (\frac{1}{2}, \frac{1}{2})$ and player two plays the mixed strategy $\delta_2 = (\frac{1}{3}, \frac{2}{3})$. This leads to a joint strategy $\delta = (\delta_1, \delta_2) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{3}, \frac{2}{3})$. According to Lemma 3.1, the expected payoff vector should then also be $(\frac{1}{2}, \frac{1}{2}, \frac{1}{3}, \frac{2}{3})$. We verify this:

$$\begin{aligned}
\mathbf{p}_i(\delta) &= \sum_{a \in \mathcal{A}} \mathbf{p}_i(a) \prod_{j=1}^n \delta_j(a_j) \\
&= (1, 0, 1, 0) \cdot \frac{1}{6} + (1, 0, 0, 1) \cdot \frac{1}{3} + (0, 1, 1, 0) \cdot \frac{1}{6} + (0, 1, 0, 1) \cdot \frac{1}{3} \\
&= \left(\frac{1}{6}, 0, \frac{1}{6}, 0\right) + \left(\frac{1}{3}, 0, 0, \frac{1}{3}\right) + \left(0, \frac{1}{6}, \frac{1}{6}, 0\right) + \left(0, \frac{1}{3}, 0, \frac{1}{3}\right) \\
&= \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{3}, \frac{2}{3}\right) \\
&= \delta.
\end{aligned}$$

3.2 Pure Strategy Equivalence

We introduce a novel equivalence notion between continuous games and MONFGs, called Pure Strategy Equivalence (PSE). Informally, two games are pure strategy equivalent whenever the pure strategies from the continuous game can be bijectively mapped to mixed strategies in the MONFG while keeping the corresponding utilities equal. We formally define this below.

Definition 3.1 (Pure strategy equivalence). A continuous game $G_c = (N_c, \mathcal{S}, v)$ is pure strategy equivalent to a finite multi-objective normal-form game $G_m = (N_m, \mathcal{A}, \mathbf{p})$ with utility functions u when there exists a tuple of functions (π, φ) such that:

- $\pi : N_c \rightarrow N_m$ is a bijective function called the *player bijection*;
- $\varphi = \varphi_1 \times \dots \times \varphi_n$ where $\forall i \in N_c, \varphi_i : S_i \rightarrow \Delta^{k_{\pi(i)}}$ is a continuous bijective function with a continuous inverse, called the *strategy bijection*;

$$\bullet \forall i \in N_c, \forall s \in \mathcal{S}, v_i(s) = u_{\pi(i)}(\varphi(s)).$$

We first show that for every MONFG there exists a pure strategy equivalent continuous game. We note that this property has already been used [22, Lemma 1], but has not yet been explicitly described in terms of pure strategy equivalence.

THEOREM 3.2. *For every multi-objective normal-form game with continuous utility functions, there exists a pure strategy equivalent continuous game.*

PROOF. Let $G_m = (N_m, \mathcal{A}, \mathbf{p})$ be a multi-objective normal-form game and let u be the set of continuous utility functions used by players in G_m . We construct a continuous game $G_c = (N_c, \mathcal{S}, v)$ that is pure strategy equivalent to G_m .

First, let $N_c = N_m$, making the player bijection $\pi = \mathbb{1}$. To simplify the notation, we can directly substitute $\pi(i) = i$.

Recall that the set of mixed strategies for player i in G_m was defined as a simplex Δ^{k_i} . For all $S_i \in \mathcal{S}$, we define $S_i = \Delta^{k_i}$, thus satisfying the condition of being a nonempty compact metric space. As such, each strategy bijection $\varphi_i = \mathbb{1}$.

Finally, define each $v_i = u_i \circ \mathbf{p}_i$. Observe that each v_i is continuous as it is a composition of two continuous functions. As such, G_c is pure strategy equivalent to G_m . \square

The proof presented here provides an explicit construction of a pure strategy equivalent continuous game for any MONFG, which we outline in Appendix C.1 (see Algorithm 2). Moreover, as the strategy sets in both games are equal, the strategy bijection is simply the identity function. This is critical from an algorithmic viewpoint, as it ensures that the mapping may be computed efficiently. This result has two important implications. First, observe that we may also perform the construction in single-player settings, such as those studied in multi-objective planning and reinforcement learning. This suggests that algorithms designed for continuous action spaces can be used whenever the utility function of the agent is known a priori. Secondly, whenever the resulting utility functions are (twice) differentiable, the resulting game falls under the class of differentiable games, which may be solved using efficient gradient-based methods [14].

We now show a converse to Theorem 3.2, namely that every continuous game with convex strategy sets can be mapped to a pure strategy equivalent MONFG.

THEOREM 3.3. *For every continuous game whose strategy spaces are convex subsets of an Euclidean space, there exists a pure strategy equivalent multi-objective normal-form game.*

PROOF. Let $G_c = (N_c, \mathcal{S}, v)$ be a continuous game where each $S_i \in \mathcal{S}$ is a compact, convex and nonempty subset of Euclidean space. We construct a finite multi-objective normal-form game $G_m = (N_m, \mathcal{A}, \mathbf{p})$ with utility functions u that is pure strategy equivalent to G_c .

The player bijection is trivial by letting $N_m = N_c$ and thus $\pi = \mathbb{1}$. To simplify the notation, we can directly substitute $\pi(i) = i$ and refer to N_m or N_c simply as N .

It is a known property that all compact convex subsets of \mathbb{R}^k with a nonempty interior are homeomorphic to the probability k -simplex [3, Theorem 16.4]. Let us first assume that each strategy set

S_i indeed has a nonempty interior. This means that for every player i there exists a continuous bijective function f_i with a continuous inverse f_i^{-1} such that,

$$f_i : S_i \rightarrow \Delta^{k_i}. \quad (4)$$

Recall that the set of mixed strategies Δ^{k_i} for player i in G_m is defined as the probability simplex over their actions. Therefore, we may use this homeomorphism to construct the set of mixed strategies over $k_i + 1$ actions,

$$\forall i \in N : f_i(S_i) = \Delta^{k_i}. \quad (5)$$

Note that the vertices of the simplex naturally represent the actions in G_m . This ensures that each joint strategy can be bijectively mapped from G_c to G_m with

$$\forall s \in \mathcal{S} : f(s) = f_1(s_1) \times \cdots \times f_n(s_n). \quad (6)$$

We can therefore define the strategy bijection $\varphi = f$. Because f is a homeomorphism, it is by definition a continuous bijective function with a continuous inverse.

Let the payoff functions \mathbf{p} of G_m be the payoffs of the identity game from Lemma 3.1. We define the utility functions u in G_m as follows,

$$\forall i \in N, u_i = v_i \circ f_i^{-1}. \quad (7)$$

We show that this ensures that,

$$\forall i \in N, \forall s \in \mathcal{S} : v_i(s) = u_i(\mathbf{p}_i(\varphi_i(s))). \quad (8)$$

Substituting the necessary values, we get $\forall i \in N, \forall s \in \mathcal{S}$,

$$u_i(\mathbf{p}_i(\varphi_i(s))) = u_i(\mathbf{p}_i(f_i(s))) \quad (9)$$

$$= u_i(\mathbf{p}_i(\delta)) \quad (10)$$

$$= u_i(\delta) \quad (11)$$

$$= v_i \circ f_i^{-1}(\delta) \quad (12)$$

$$= v_i(s). \quad (13)$$

As such, G_c is pure strategy equivalent to G_m .

In the case where one or more S_i has an empty interior, those S_i have a nonempty interior with respect to their affine span. Then S_i is homeomorphic to the k -simplex, where $k = \dim \text{aff}(S_i)$. The remainder of the construction follows analogous to before. \square

The main idea behind the proof is to construct an MONFG using the identity game payoffs and reuse the utility functions from the continuous game. As we assumed convexity of the strategy sets, it is known that a homeomorphism between each strategy set and a probability simplex can be constructed. These simplices are the mixed strategies in the MONFG. The utility functions in the MONFG can then be defined as the composition between the utility functions of the continuous game and the inverse of each homeomorphism.

The construction defined in this proof also establishes a computational approach to transform a continuous game to an MONFG, formalised in Appendix C.1 (see Algorithm 3). However, contrary to the other direction, the strategy bijection does appear here. As such, the algorithm requires the strategy bijection to be explicitly defined which poses two distinct challenges. First, these strategy bijections may not be easily obtainable, in which case we present a standard approach that initially goes through a unit ball and subsequently to the probability simplex as a suitable first attempt in Appendix C.2. Second, the resulting functions are not guaranteed to be efficiently

computable, hence rendering the procedure intractable for some applications. We leave the application of Theorem 3.3 to such scenarios for future work.

An interesting implication of Theorem 3.3 is that it can straightforwardly be extended to continuous games with non-convex strategy sets whenever they are still homeomorphic to a simplex. However, while convexity of strategy sets is a frequently made assumption for continuous games, there can be games of interest which do not have this property. In this case, one can approximate the original continuous game with another continuous game that does satisfy the convexity requirement by, for example, taking the convex hull of each strategy set. In a subsequent stage, a pure strategy equivalent MONFG can be constructed for this approximate game, resulting in an MONFG which is approximately PSE to the original continuous game. We demonstrate this approach in Section 4.3 and show that it may still succeed in capturing the original game.

Given the constructions outlined in Theorems 3.2 and 3.3, it raises the question of whether they lead to unique descriptions of pure strategy equivalent games. The following remark shows that this is not the case. We show this formally in Appendix C.1 by constructing multiple pure strategy equivalent games.

Remark 1. A continuous game may have multiple pure strategy equivalent multi-objective normal-form games and vice versa.

An important consequence of Remark 1 is that some PSE games can be more amenable to analysis than others. For example, we have shown that for a continuous game a pure strategy equivalent MONFG can be constructed with identity game payoffs. If this continuous game can also be proven to be PSE to an MONFG with a dominated action in its payoff and monotonic utility functions, the latter game will certainly be easier to analyse. Using equivalence between games to solve a difficult game through a simpler equivalent has been studied with success before [10].

3.3 Mixed Strategy Equivalence

Given the results for pure strategy equivalence, a pressing question becomes how to handle the mixed strategies from a continuous game in an equivalent MONFG. For this purpose, we introduce a novel strategy concept in MONFGs, named *hierarchical strategies*. Hierarchical strategies allow for mixing over mixed strategies and are defined for each player i as the set of Borel probability measures on their set of mixed strategies Δ^{k_i} and denoted by $\mathcal{B}(\Delta^{k_i})$. The expected utility of a hierarchical strategy μ is then defined analogously to the way expected utility of a mixed strategy in continuous games is defined in Eq. (1), that is,

$$u_i(\mu) = \int_{\Delta^{k_1} \times \cdots \times \Delta^{k_n}} u_i(\mathbf{p}_i(\delta_1, \dots, \delta_n)) d\mu_1(\delta_1) \cdots d\mu_n(\delta_n). \quad (14)$$

We stress that the probabilities in a hierarchical strategy cannot be distributed to form an equivalent mixed strategy. This is because the utility of the resulting mixed strategy need not equal the expected utility of the hierarchical strategy, specifically when nonlinear utility functions are used. We illustrate this in Example 2.

Example 2. Consider the game in Fig. 2 and utility functions

$$u_1(x, y) = u_2(x, y) = x^2 + y^2. \quad (15)$$

	A	B
A	(3, 1); (3, 1)	(1, 3); (1, 3)
B	(1, 3); (1, 3)	(3, 1); (3, 1)

Figure 2: The game used in Example 2.

For the row player, we define two strategies, $\delta_{1,1} = (1, 0)$ and $\delta_{1,2} = (0, 1)$. For simplicity, we assume a deterministic strategy $\delta_2 = (1, 0)$ for the column player where they always play action A. For both players, the joint strategy $(\delta_{1,1}, \delta_2)$ leads to an expected payoff vector of (3, 1) and $(\delta_{1,2}, \delta_2)$ to an expected payoff vector of (1, 3). Both joint strategies individually result in a utility of 10.

Consider now the following hierarchical strategy for the row player,

$$\mu_1 = \left(P(\delta_{1,1}) = \frac{1}{2}, P(\delta_{1,2}) = \frac{1}{2} \right).$$

This hierarchical strategy denotes the fact that they will play $\delta_{1,1}$ with 50% probability and $\delta_{1,2}$ with 50% probability. As both strategies result in a utility of 10, the expected utility of μ_1 is also 10. However, distributing the probabilities in μ_1 to form a mixed strategy, $\delta_{1,3} = (\frac{1}{2}, \frac{1}{2})$, results in an expected payoff vector of (2, 2) with a utility of 8. This demonstrates that a hierarchical strategy cannot be distributed to form an equivalent mixed strategy.

While mainly used as an instrument in our analysis, hierarchical strategies may also have an applied use in multi-objective games. Concretely, hierarchical strategies are appropriate to consider when agents have to decide on a strategy which is then executed for a given period without the possibility for downstream adjustments. For example, agents joining an auction may have to commit a single mixed strategy to an automated auctioneer a priori, after which auctions are held for a number of rounds without further input from the agents.

We now define mixed strategy equivalence between a continuous game and an MONFG. Informally, this equivalence notion generalises pure strategy equivalence to relate mixed strategies from the continuous game with hierarchical strategies in the MONFG.

Definition 3.2 (Mixed strategy equivalence). Let $G_c = (N_c, \mathcal{S}, v)$ be a continuous game and $G_m = (N_m, \mathcal{A}, p)$ a finite multi-objective normal-form game with utility functions u . G_c is mixed strategy equivalent to G_m if they are pure strategy equivalent with (π, φ) and there exists a function ψ such that,

- $\psi = \psi_1 \times \dots \times \psi_n$ where $\forall i \in N_c, \psi_i : \mathcal{B}(S_i) \rightarrow \mathcal{B}(\Delta^{k_{\pi(i)}})$ is a bijective function;
- $\forall i \in N_c, \forall \mu \in \mathcal{B}(\mathcal{S}), v_i(\mu) = u_{\pi(i)}(\psi(\mu))$.

We remark that, by definition, mixed strategy equivalence implies pure strategy equivalence. Moreover, recall that the definition of mixed strategies in continuous games is similar to the definition of the set of hierarchical strategies in MONFGs. This is no coincidence and allows us to show that whenever a continuous game is pure strategy equivalent to an MONFG, it also implies mixed strategy equivalence. A formal proof of this is deferred to Appendix D.

THEOREM 3.4. *If a continuous game is pure strategy equivalent to a multi-objective normal-form game, they are also mixed strategy equivalent.*

PROOF SKETCH. By definition of pure strategy equivalence, a continuous bijective function with a continuous inverse is given which maps strategies from one game to the other. Moreover, mixed strategies and hierarchical strategies were respectively defined to be the set of all Borel probability measures over their pure strategies and mixed strategies. We can use this fact to show that a measure defined over the strategies in one set must also be a measure in the other set defined over the mapped strategies. This also implies equal utility, thereby completing the proof. \square

3.4 Mapping of Nash Equilibria

Our final theoretical contributions consider Nash equilibria in pure strategy equivalent games. Theorem 3.5 first specifies that pure strategy Nash equilibria in continuous games correspond to mixed strategy Nash equilibria in pure strategy equivalent MONFGs. Intuitively, this is clear as utilities for pure strategies in the continuous game were already guaranteed to be equal to the utilities for mixed strategies in the MONFG. Therefore, if a joint strategy cannot be improved upon unilaterally in either game, the mapped joint strategy in the related game can also not be improved upon.

THEOREM 3.5. *A pure strategy is a Nash equilibrium in a continuous game if and only if it is a mixed strategy Nash equilibrium in a pure strategy equivalent multi-objective normal-form game.*

We define a hierarchical Nash equilibrium below, such that we can construct a similar argument to show that mixed strategy NE in continuous games necessarily correspond to a hierarchical strategy NE in pure strategy equivalent MONFGs.

Definition 3.3 (Hierarchical Nash equilibria in multi-objective normal-form games). A hierarchical strategy profile μ^* is a hierarchical Nash equilibrium if,

$$u_i(\mu_i^*, \mu_{-i}^*) \geq u_i(\mu_i, \mu_{-i}^*),$$

for all players i and alternative hierarchical strategies $\mu_i \in \mathcal{B}(\Delta^{k_i})$.

THEOREM 3.6. *A mixed strategy is a Nash equilibrium in a continuous game if and only if it is a hierarchical Nash equilibrium in a pure strategy equivalent multi-objective normal-form game.*

Complete proofs for Theorems 3.5 and 3.6 are provided in Appendix E. These properties are of significant importance as they introduce algorithmic methods for computing Nash equilibria in either game model to the other. Moreover, as a consequence of Theorem 3.6 in particular, we can state the first general result for Nash equilibria to exist in MONFGs if we allow players to assume hierarchical strategies rather than limiting them to mixed strategies only. The proof for Corollary 3.7 is given below.

COROLLARY 3.7. *Every multi-objective normal-form game with continuous utility functions has a hierarchical Nash equilibrium.*

PROOF. Theorem 3.2 shows that every MONFG with continuous utility functions can be mapped to a pure strategy equivalent continuous game. Furthermore, it is known that a mixed strategy Nash equilibrium exists in every continuous game [7]. Theorem 3.6 guarantees that a mixed strategy Nash equilibrium in a continuous game is a hierarchical Nash equilibrium in the MONFG, therefore guaranteeing a hierarchical Nash equilibrium in every MONFG. \square

Algorithm 1 Multi-Objective Fictitious Play

Input: An MONFG $G = (N, \mathcal{A}, \mathbf{p})$, utility functions u and maximum timestep T

Output: A joint strategy δ

```

1:  $\delta_1 \leftarrow \frac{1}{|A_1|}$ 
2:  $h_1 \leftarrow \mathbf{0}$  ▷ The history for player 1
3:  $\delta_2 \leftarrow \frac{1}{|A_2|}$ 
4:  $h_2 \leftarrow \mathbf{0}$ 
5: for  $t \in 1, \dots, T$  do
6:    $\tilde{\delta}_2 \leftarrow \frac{h_1}{t}$  ▷ Compute an empirical mixed strategy
7:    $\tilde{\delta}_1 \leftarrow \frac{h_2}{t}$ 
8:    $\delta_1 \leftarrow \text{BESTRESPONSE}(\mathbf{p}_1, u_1, \tilde{\delta}_2)$  ▷ Compute a best response
9:    $\delta_2 \leftarrow \text{BESTRESPONSE}(\mathbf{p}_2, u_2, \tilde{\delta}_1)$ 
10:   $a_1 \leftarrow a \sim \delta_1$  ▷ Sample an action from the best response
11:   $a_2 \leftarrow a \sim \delta_2$ 
12:   $h_{1,a_2} \leftarrow h_{1,a_2} + 1$  ▷ Update the history
13:   $h_{2,a_1} \leftarrow h_{2,a_1} + 1$ 
14: end for
15: return  $\delta$ 

```

4 EMPIRICAL RESULTS

We provide empirical evidence for the provided theorems and show that it can also be applied to compute approximate equilibria when the strategy sets do not satisfy the necessary conditions for pure strategy equivalence. We adapt the well-known fictitious play algorithm from single-objective games to multi-objective games and use it to compute pure strategy Nash equilibria in continuous games. The results empirically demonstrate the applicability of our contributions and may serve as a useful template for future applications. Our implementation is available at <https://github.com/wilrop/pure-strategy-equivalence>.

4.1 Multi-Objective Fictitious Play

Fictitious play aims to learn strategies resulting in a Nash equilibrium through repeated plays of the game. While it is not guaranteed to converge in general-sum single-objective games, fictitious play and its extensions are widely used in practice. In Algorithm 1, we show an extension for fictitious play to multi-objective games. For simplicity, we consider a two-player variant but this can trivially be extended to n -player games. In each iteration of the algorithm, players calculate the empirical strategy of their opponent based on their history of play and compute a best response to this strategy. Players subsequently sample an action from their new strategy and update their histories.

Recent work has studied an adaptation of fictitious play to continuous games [6]. In their algorithm, a growing array of past strategies is kept to later compute a best response to, which imposes a significant memory requirement. A key advantage of our approach is that it only requires an array of fixed length, i.e., one entry per action, where a counter is incremented each time an action is played. The empirical mixed strategy of the opponent is then calculated by taking the relative frequency of each action. A limitation of this approach is that it can only learn pure-strategy equilibria from the continuous game.

The fictitious play algorithm shown above appears identical to the original fictitious play algorithm. The exception, however, lies in the best response computation steps. In single-objective games, this can be done efficiently by selecting the action with the highest expected returns, i.e.

$$BR(A_i, s_{-i}, p_i) = \arg \max_{a_i \in A_i} p_i(a_i, s_{-i}). \quad (16)$$

In multi-objective games, this approach can only be guaranteed to return a correct best response when employing a quasiconvex utility function [22]. In general MONFGs, the best response can be a mixed strategy and thus requires executing an optimisation subroutine to find the strategy generating the maximum utility. As a best response needs to be a global maximum, this requires the use of a global optimisation algorithm. Under specific utility functions or when approximate best responses suffice, a local optimiser could also be used.

4.2 Polynomial Game

Polynomial games are a subset of continuous games, where utility functions are guaranteed to be polynomial functions of the player strategies [25, 26]. We demonstrate that such games can also be represented as an MONFG and may be solved without employing any continuous game or polynomial game specific machinery. We cover a simple example as described by Parrilo [17].

Consider a zero-sum game where both players select a strategy from the interval $[-1, 1]$. The utility function for player one is defined as,

$$v_1(x, y) = 2xy^2 - x^2 - y, \quad (17)$$

with x the strategy selected by player one and y the strategy of player two. As the game is zero-sum, player two's utility is given by $v_2(x, y) = -v_1(x, y)$. The utility functions used in the game guarantee the existence of a unique Nash equilibrium in pure strategies where $x^* = 0.397$ and $y^* = 0.630$.

As the strategy sets are line segments, and thus are 1-simplices, a pure strategy equivalent multi-objective game is guaranteed to exist. To complete the transformation from the polynomial game to a multi-objective game, a strategy bijection $\varphi_i : S_i \rightarrow \Delta^{k_i}$ is required for each player i . The strategy bijection for both players is given by,

$$\varphi_i(s_i) = \left(\frac{s_i - s_{i,\min}}{s_{i,\max} - s_{i,\min}}, 1 - \frac{s_i - s_{i,\min}}{s_{i,\max} - s_{i,\min}} \right), \quad (18)$$

where $s_{i,\min} = -1$ and $s_{i,\max} = 1$ for both players. The inverse strategy bijection is given by,

$$\varphi_i^{-1}(\delta_i) = s_{i,\min} + \delta_{i,0} \cdot (s_{i,\max} - s_{i,\min}). \quad (19)$$

The final multi-objective game thus has two players with two actions each and corresponding identity game payoffs. Furthermore, the utility functions for both players are $u_i = v_i \circ \varphi_i^{-1}$. Because the original utility functions v_1 and v_2 guarantee a pure strategy Nash equilibrium in the continuous game, fictitious play is well suited to find the mixed strategy equilibrium in the MONFG.

We execute the fictitious play algorithm for 200 iterations on the constructed multi-objective game and repeat this for 1000 trials. Figure 3 shows the learned strategies over time, with the shaded area denoting the standard variation at that time. We illustrate




Figure 3: Learning curves for the polynomial game.

the Nash equilibrium $(0.397, 0.630)$ with dotted lines. It is clear that our algorithm learns the equilibrium after approximately 100 iterations and is able to keep improving its strategies closer to the exact equilibrium over time.

4.3 Bertrand Price Game

Theorem 3.3 states that a pure strategy equivalent MONFG is only guaranteed to exist for continuous games whose strategy spaces are convex subsets of an Euclidean space. We demonstrate that pure strategy equivalence can still be applied when this condition is not met by approximating the continuous game. We illustrate this using the Bertrand price game characterised by Judd et al. [13].

Bertrand price games have been extensively studied as an economic model for determining prices in competitive settings. In this example, we consider two firms, x and y , which respectively produce a different good for price p_x and p_y . There are three types of customers, which have a distinct demand for both goods. The first type of customer has linear demand curves $d_{x,1}$ and $d_{y,1}$ and only wants the good from firm x ,

$$d_{x,1}(p_x, p_y) = a - p_x \quad d_{y,1}(p_x, p_y) = 0, \quad (20)$$

with a signifying all factors, other than price, which influence the demand. The demand function for the third type of customer is defined analogously for the good of firm y ,

$$d_{x,3}(p_x, p_y) = 0 \quad d_{y,3}(p_x, p_y) = a - p_y. \quad (21)$$

Finally, the second type of customer has a demand for both goods,

$$d_{x,2}(p_x, p_y) = n \cdot p_x^{-\sigma} \left(p_x^{1-\sigma} + p_y^{1-\sigma} \right)^{(\gamma-\sigma)/(-1+\sigma)} \quad (22)$$

$$d_{y,2}(p_x, p_y) = n \cdot p_y^{-\sigma} \left(p_y^{1-\sigma} + p_x^{1-\sigma} \right)^{(\gamma-\sigma)/(-1+\sigma)}. \quad (23)$$


with n the number of type two customers, σ the elasticity of substitution between x and y and γ the elasticity of demand for the composite good. The total demand for each good, respectively d_x and d_y , is given by summing the individual demands for each type. Finally, let m be the unit cost of production for each firm, then the profit for both firms is defined as,

$$r_x(p_x, p_y) = (p_x - m) \cdot d_x(p_x, p_y) \quad (24)$$

$$r_y(p_x, p_y) = (p_y - m) \cdot d_y(p_x, p_y). \quad (25)$$

p_x	p_y	r_x	r_y
2.168	25.157	724.337	608.981
25.157	2.168	608.981	724.337

Table 1: The Nash equilibrium strategies and their profits. The highest profit for firm x and y are highlighted.



(a) Learning curves for the equilibrium at $(2.168, 25.157)$. (b) Learning curves for the equilibrium at $(25.157, 2.168)$.

Figure 4: Learning curves for the Bertrand price game in the interval $[1, 30]$.

The range of possible prices considered in the game is the open interval $(0, +\infty)$. As such, strategy spaces in the continuous game are non-compact, thus violating a necessary condition for pure strategy equivalence. We can resolve this, however, by making compact convex approximations of the strategy spaces and using these instead. We do this by constraining prices to be in a closed interval $[p_{\min}, p_{\max}]$, which ensures that strategy sets are 1-simplices as in the previous example. Because of this approximation, we may reuse the same strategy bijection as defined in Eqs. (18) and (19). Note that approximating the continuous game by altering strategy sets may remove existing equilibria from reach or introduce new ones. In this particular example, as we are both lower and upper bounding the strategy sets, it is possible that an equilibrium falls outside of the bounds and a new equilibrium is created in the MONFG which is not an equilibrium in the original game.

For the following experiments, we define $\sigma = 3, \gamma = 2, n = 2700, m = 1$ and $a = 50$. With these parameters, the price game has two distinct Nash equilibria, shown in Table 1.

4.3.1 Suitable Approximation. To give both equilibria a chance of being selected, we set $p_{\min} = 1$ and $p_{\max} = 30$. Every execution of the fictitious play algorithm is run for 200 iterations and results are averaged over 1000 trials as in the previous section.

In Fig. 4, we show the trajectories leading to the two equilibria. In earlier episodes, the trajectories are non-smooth and show high standard deviation, as the best response computation from a limited history of play leads agents to change strategies rapidly. However, once beliefs converge after approximately 150 episodes, a Nash equilibrium is consistently played. We also find that the learning trajectories for both equilibria are similar, showing that the individual trajectory that is followed is mostly determined by randomisation early on in the learning process. These results demonstrate that multi-objective algorithms can be applied even to approximations of continuous games, given that these approximations sufficiently capture the original game.




Figure 5: Learning curves for the Bertrand price game in the interval [4, 30].

4.3.2 Unsuitable Approximation. Next, we consider what happens when an unsuitable approximation of the strategy sets is used. Specifically, we raise the minimum price to 4, which renders both equilibria from Table 1 impossible. Intuitively, these equilibria had one firm that opted for a mass-market strategy with lower prices and another that opted for a niche strategy with higher prices. By raising the minimum price, we render this mass-market strategy impossible. We show the resulting learning trajectories for this experiment in Fig. 5.

We find that throughout all trials the firms rapidly converge to the joint strategy (22.987, 22.987), which leads to a profit of 673.38 for both. In fact, this joint strategy is a Nash equilibrium in the MONFG but not a Nash equilibrium of the original continuous game, thus demonstrating the mentioned limitations of the approximation technique. Interestingly, the Nash equilibrium in the approximate MONFG leads to a joint strategy with higher social welfare when considering both the total sum of profits as well as the maximum lowest profit. Specifically, it returns a total profit of 1346.754, while both equilibria from Table 1 lead to a total profit of 1333.318 and has a higher lowest profit. As such, even when using insufficient approximations for the original game, our contributions may result in interesting solutions from, e.g., a mechanism design perspective.

We note that designing appropriate approximations for arbitrary continuous games is a nontrivial task. For most interesting applications, suitable convex compact approximations will not be given and thus require leveraging domain knowledge or post-processing to confirm the retrieved solution in the original game.

5 RELATED WORK

Multi-objective games were first introduced by Blackwell [2] and have since been studied broadly. One solution concept that is often considered are Pareto-Nash equilibria, which take a utility agnostic approach and are thus defined over the vectorial payoffs directly [12]. When taking a utility-based approach [21], it has been shown that Nash equilibria need not exist [19]. Follow up work showed that assuming only quasiconcave utility functions [22] is a sufficient condition to guarantee existence again. From a computational perspective, reinforcement learning algorithms and additional techniques such as opponent modelling [20, 27] and communication [23] have recently been successfully explored.

The second game model we considered in this work are continuous games. General existence guarantees are known in these games, with for example work on Nash equilibria [7] and correlated equilibria [8]. From a computational perspective, both fictitious play [6] and no-regret learning have been explored [11], with the latter also obtaining strong convergence guarantees. We note that while we adhere to the definition of a continuous game by Stein et al. [25], there exist other definitions for continuous games which place different assumptions on the strategy sets or utility functions [1, 6, 11]. Finally, polynomial games, a subset of continuous games, have been covered in detail with theoretical and algorithmic results for their Nash equilibria and correlated equilibria [25, 26].

Our work is related to other equivalence notions in game theory. The first notable example of such an equivalence notion is strategic equivalence [16]. An advantage of strategic equivalence is that Nash equilibria are preserved, thus being a useful construct for computing Nash equilibria in games. For example, the Nash equilibria of an otherwise intractable game might be computed by constructing a strategically equivalent zero-sum game for which efficient solving methods do exist [10].

Pure strategy equivalence, as defined in Section 3.2, is most closely related to the concept of a game isomorphism which defines two games to be isomorphic when there exists a mapping from one to the other [5]. Two variants of a game isomorphism are defined, namely a strong and a weak isomorphism, with a strong isomorphism preserving all Nash equilibria and a weak isomorphism preserving only the pure strategy Nash equilibria.

6 CONCLUSION

We contribute a novel equivalence class, called pure strategy equivalence, between continuous games and multi-objective normal-form games. We show that for every continuous game whose strategy spaces are convex subsets of an Euclidean space, a pure strategy equivalent MONFG can be constructed and vice versa. Moreover, this equivalence entails the persistence of Nash equilibria. We demonstrate the applicability of pure strategy equivalence by learning Nash equilibria in two continuous games utilizing a multi-objective fictitious play algorithm.

The equivalence notion introduces a range of new theoretical contributions and computational approaches for both models. For the theoretical aspect, it is known that other game models which allow for more complex interactions, such as Bayesian games and extensive form games, can be reduced to normal-form games [16]. Formulating a bridge between continuous games and multi-objective games thus opens the possibility for additional equivalence results between these games with an infinite number of pure strategies and the related multi-objective variant. From an algorithmic perspective, we expect that this will allow continuous games to be solved more efficiently, as the tabular structure of MONFGs can be more appropriate for computational approaches. For future work, we aim to evaluate this on larger continuous games with more complex structures. Finally, as the equivalence notion is not unique, an interesting question remains how to find the best pure strategy equivalent game for any given game.

ACKNOWLEDGMENTS

We are grateful to Raphael Avalos for helpful discussions during the early stages of this work. WR and RR are supported by the Research Foundation – Flanders (FWO), grant numbers 1197622N and 1286223N and CG is supported by the Marie-Curie grant 101063180. This research was supported by funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” program.

REFERENCES

- [1] Lukáš Adam, Rostislav Horčík, Tomáš Kasl, and Tomáš Kroupa. 2021. Double Oracle Algorithm for Computing Equilibria in Continuous Games. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (May 2021), 5070–5077. <https://doi.org/10.1609/aaai.v35i6.16641>
- [2] David Blackwell. 1954. An Analog of the Minimax Theorem for Vector Payoffs. *Pacific J. Math.* 6, 1 (1954), 1–8. <https://doi.org/10.2140/pjm.1956.6.1>
- [3] Glen E. Bredon. 1993. General Topology. In *Topology and Geometry*, Glen E. Bredon (Ed.). Springer New York, New York, NY, 1–62. https://doi.org/10.1007/978-1-4757-6848-0_1
- [4] Haifeng Fu. 2021. On the Existence of Pareto Undominated Mixed-Strategy Nash Equilibrium in Normal-Form Games with Infinite Actions. *Economics Letters* 201 (April 2021), 109771. <https://doi.org/10.1016/j.econlet.2021.109771>
- [5] Joaquim Gabarró, Alina García, and Maria Serna. 2011. The complexity of game isomorphism. *Theoretical Computer Science* 412 (2011), 6675–6695. Issue 48. <https://doi.org/10.1016/j.tcs.2011.07.022>
- [6] Sam Ganzfried. 2021. Algorithm for Computing Approximate Nash Equilibrium in Continuous Games with Application to Continuous Blotto. *Games* 12, 2 (June 2021), 47. <https://doi.org/10.3390/g12020047>
- [7] I. L. Glicksberg. 1952. A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points. *Proc. Amer. Math. Soc.* 3, 1 (1952), 170–170. <https://doi.org/10.2307/2032478>
- [8] Sergiu Hart and David Schmeidler. 1989. Existence of Correlated Equilibria. *Mathematics of Operations Research* 14, 1 (Dec. 1989), 18–25.
- [9] Conor F. Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Raymond, Timothy Verstraeten, Luisa M. Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai A. Irissappane, Patrick Mannion, Ann Nowé, Gabriel Ramos, Marcello Restelli, Peter Vamplew, and Diederik M. Roijers. 2022. A Practical Guide to Multi-Objective Reinforcement Learning and Planning. *Autonomous Agents and Multi-Agent Systems* 36, 1 (April 2022), 26. <https://doi.org/10.1007/s10458-022-09552-y>
- [10] Joseph Lee Heyman. 2019. *On the Computation of Strategically Equivalent Games*. Ph.D. Dissertation. The Ohio State University.
- [11] Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. 2021. Adaptive Learning in Continuous Games: Optimal Regret Bounds and Convergence to Nash Equilibrium. In *Proceedings of Thirty Fourth Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 134)*, Mikhail Belkin and Samory Kpotufe (Eds.). PMLR, 2388–2422.
- [12] Anisse Ismaili. 2018. On Existence, Mixtures, Computation and Efficiency in Multi-Objective Games. In *PRIMA 2018: Principles and Practice of Multi-Agent Systems*, Tim Miller, Nir Oren, Yuko Sakurai, Itsuki Noda, Bastin Tony Roy Savarimuthu, and Tran Cao Son (Eds.). Springer International Publishing, Cham, 210–225.
- [13] Kenneth L. Judd, Philipp Renner, and Karl Schmedders. 2012. Finding All Pure-Strategy Equilibria in Games with Continuous Strategies. *Quantitative Economics* 3, 2 (July 2012), 289–331. <https://doi.org/10.3982/QE165>
- [14] Alistair Letcher, David Balduzzi, Sébastien Racanière, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. 2019. Differentiable Game Mechanics. *Journal of Machine Learning Research* 20, 84 (2019), 1–40.
- [15] Junlin Lu, Patrick Mannion, and Karl Mason. 2022. A Multi-Objective Multi-Agent Deep Reinforcement Learning Approach to Residential Appliance Scheduling. *IET Smart Grid* n/a, n/a (May 2022). <https://doi.org/10.1049/stg2.12068>
- [16] Michael Maschler, Eilon Solan, and Shmuel Zamir. 2013. *Game Theory*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9780511794216>
- [17] Pablo A. Parrilo. 2006. Polynomial Games and Sum of Squares Optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*. IEEE, San Diego, CA, USA, 2855–2860. <https://doi.org/10.1109/CDC.2006.377261>
- [18] Roxana Rădulescu, Patrick Mannion, Diederik M. Roijers, and Ann Nowé. 2020. Multi-Objective Multi-Agent Decision Making: A Utility-Based Analysis and Survey. *Autonomous Agents and Multi-Agent Systems* 34, 1 (April 2020), 10–10. <https://doi.org/10.1007/s10458-019-09433-x>
- [19] Roxana Rădulescu, Patrick Mannion, Yijie Zhang, Diederik M. Roijers, and Ann Nowé. 2020. A Utility-Based Analysis of Equilibria in Multi-Objective Normal-Form Games. *The Knowledge Engineering Review* 35 (2020), e32–e32. <https://doi.org/10.1017/S0269888920000351>
- [20] Roxana Rădulescu, Timothy Verstraeten, Yijie Zhang, Patrick Mannion, Diederik M. Roijers, and Ann Nowé. 2022. Opponent Learning Awareness and Modelling in Multi-Objective Normal Form Games. *Neural Computing and Applications* 34, 3 (Feb. 2022), 1759–1781. <https://doi.org/10.1007/s00521-021-06184-3>
- [21] Diederik M. Roijers and Shimon Whiteson. 2017. Multi-Objective Decision Making. In *Synthesis Lectures on Artificial Intelligence and Machine Learning*, Vol. 34. Morgan and Claypool, 129–129. <https://doi.org/10.2200/S00765ED1V01Y201704AIM034>
- [22] Willem Röpke, Diederik M. Roijers, Ann Nowé, and Roxana Rădulescu. 2022. On Nash Equilibria in Normal-Form Games with Vectorial Payoffs. *Autonomous Agents and Multi-Agent Systems* 36, 2 (Oct. 2022), 53. <https://doi.org/10.1007/s10458-022-09582-6>
- [23] Willem Röpke, Diederik M. Roijers, Ann Nowé, and Roxana Rădulescu. 2022. Preference Communication in Multi-Objective Normal-Form Games. *Neural Computing and Applications* (July 2022). <https://doi.org/10.1007/s00521-022-07533-6>
- [24] Willem Röpke, Carla Groenland, Roxana Rădulescu, Ann Nowé, and Diederik M. Roijers. 2023. Bridging the Gap Between Single and Multi Objective Games. <https://doi.org/10.48550/ARXIV.2301.05755>
- [25] Noah D. Stein, Asuman Ozdaglar, and Pablo A. Parrilo. 2008. Separable and Low-Rank Continuous Games. *International Journal of Game Theory* 37, 4 (Dec. 2008), 475–504. <https://doi.org/10.1007/s00182-008-0129-2>
- [26] Noah D. Stein, Pablo A. Parrilo, and Asuman Ozdaglar. 2011. Correlated Equilibria in Continuous Games: Characterization and Computation. *Games and Economic Behavior* 71, 2 (March 2011), 436–455. <https://doi.org/10.1016/j.geb.2010.04.004>
- [27] Yijie Zhang, Roxana Rădulescu, Patrick Mannion, Diederik M. Roijers, and Ann Nowé. 2020. Opponent Modelling for Reinforcement Learning in Multi-Objective Normal Form Games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Auckland, New Zealand, 2080–2082–2080–2082.