The Danger Of Arrogance: Welfare Equilibra As A Solution To Stackelberg Self-Play In Non-Coincidental Games

Jake Levi¹ Chris Lu² Timon Willi² Christian Schroeder de Witt² Jakob Foerster²

Abstract

The increasing prevalence of multi-agent learning systems in society necessitates understanding how to learn effective and safe policies in general-sum multi-agent environments against a variety of opponents, including self-play. General-sum learning is difficult because of nonstationary opponents and misaligned incentives. Our first main contribution is to show that many recent approaches to general-sum learning can be derived as approximations to Stackelberg strategies, which suggests a framework for developing new multi-agent learning algorithms. We then define non-coincidental games as games in which the Stackelberg strategy profile is not a Nash Equilibrium. This notably includes several canonical matrix games and provides a normative theory for why existing algorithms fail in selfplay in such games. We address this problem by introducing Welfare Equilibria (WE) as a generalisation of Stackelberg Strategies, which can recover desirable Nash Equilibria even in noncoincidental games. Finally, we introduce Welfare Function Search (WelFuSe) as a practical approach to finding desirable WE against unknown opponents, which finds more mutually desirable solutions in self-play, while preserving performance against naive learning opponents.

1. Introduction

The use of machine learning in multi-agent systems is becoming increasingly prevalent in society (Zhang et al., 2021), but learning in multi-agent systems presents several challenges. Firstly, different agents might have misaligned or conflicting incentives (Subramanian et al., 2023), in which case agents won't necessarily act in the interests of other agents. Secondly, the learning environment is nonstationary, because other agents are simultaneously learning and updating their policies. Non-stationarity violates standard theoretical assumptions made in reinforcement learning (Sutton & Barto, 2018), and empirically makes learning significantly more challenging (Wang et al., 2022).

Real-world case-studies of multi-agent systems have shown how nefarious human agents can derail the behaviour of artificial learning agents, causing emotional harm to innocent observers and reputational damage to the learning agent's creators (Wolf et al., 2017). Such examples highlight the possible negative real-world consequences of deploying naively designed learning agents in multi-agent systems, and hence the necessity of adapting to the possible incentives and behaviour of other agents. These considerations are especially important in safetycritical environments (Kiran et al., 2021).

Opponent shaping (OS) is an approach to general-sum learning in which agents explicitly consider the opponent's incentives and behaviour, and adapt their own behaviour in order to *shape* the opponent's future learning process. For example, under self-play LOLA (Foerster et al., 2018) converges to prosocial solutions which incentivise opponent cooperation in canonical games such as the Iterated Prisoners' Dilemma (IPD). However, it has been shown (Willi et al., 2022) that OS algorithms can behave arrogantly in games such as the chicken game, leading to catastrophic outcomes in self-play. Avoiding catastrophe in self-play is fundamentally important for any learning algorithm that is to be deployed outside of simulation, because otherwise a malicious opponent could derail performance simply by choosing to use the same learning algorithm.

Our first contribution, in Section 5.1, is to show that Stackelberg strategies (Simaan & Cruz Jr, 1973) when chosen by *both* players represent a sensible solution concept in many two-player games. In Section 5.2 we show that many OS algorithms can be derived as approximations to Stackelberg strategies, and in Section 5.3 we use this framework to derive new example algorithms which have qualitative advantages over existing approaches in small games. In Section 6.1 we introduce *non-coincidental games* as games in which the Stackelberg strategy profile is not a

¹Department of Computer Science, University of Oxford, Oxford, United Kingdom ²Department of Engineering, University of Oxford, Oxford, United Kingdom. Correspondence to: Jake Levi <jake.levi@stcatz.ox.ac.uk>.

Preprint. Under review.

Nash Equilibrium, which notably includes several canonical matrix games (such as the chicken game), and helps explain why several OS algorithms which approximate Stackelberg strategies also fail in similar cases. We address this problem in Section 6.2 by introducing Welfare Equilibria (WE) as an abstract generalisation of Stackelberg strategies, which can find *desirable* solutions in self-play in a broader class of games, including non-coincidental games. Lastly, in Section 6.3 we introduce Welfare Function Search (WelFuSe) as a practical approach to finding desirable WE against unknown opponents, which considers the problem of choosing a welfare function as a bandit problem, which is solved using posterior sampling. We demonstrate how WelFuSe finds more desirable solutions in OS self-play while preserving performance against naive learning opponents. The payoff tables for all matrix games we describe are included in Table 2 in the Appendix, and code is publicly available online¹.

2. Background

Agents in multi-agent systems may have misaligned or conflicting incentives, which can be analysed using game theory. A game consists of a fixed number of players, each of which chooses a strategy, and the collection of strategies chosen by each player is known as the strategy profile. Each player has a reward function, which returns to them a scalar reward as a function of any given strategy profile. A given player's strategy is a best response (BR) to a given strategy profile when there is no other strategy available to that player which increases the value of their reward function (assuming the rest of the strategy profile remains fixed). A strategy profile is a Nash Equilibrium (NE) when every player's strategy is a BR to all other players, in which case no agent is incentivised to deviate from their choice of strategy (Nash, 1950).

A Stackelberg strategy is an abstract choice of strategy that either player can choose to play in almost any twoplayer game (Simaan & Cruz Jr, 1973), modelled on the concept of Stackelberg games. Suppose a two-player game G consists of strategies x and y (one for each player) and reward functions $R^{x}(x, y)$ and $R^{y}(x, y)$, and x is chosen to be a Stackelberg strategy. The player that chooses x assumes the opponent will be able to first observe and then play a BR to any given choice of x, therefore the opponent's strategy can be described by the "Opponent BR Function", denoted by $y^*(x)$ and shown below in Equation 2. Once the opponent BR function is known, the reward for any given choice of x is a function only of x (and not of y), which can be maximised using standard optimisation approaches, leading to the Stackelberg strategy x^* , shown in Equation 3. The opponent BR functions and Stackelberg strategies for either player in the game G are summarised below:

$$x^*(y) = \operatorname*{argmax}_{x} \Big[R^x(x, y) \Big] \tag{1}$$

$$y^*(x) = \underset{y}{\operatorname{argmax}} \left[R^y(x, y) \right]$$
(2)

$$x^* = \operatorname*{argmax}_{x} \left[R^x \left(x, y^*(x) \right) \right]$$
(3)

$$y^* = \underset{y}{\operatorname{argmax}} \left[R^y \Big(x^*(y), y \Big) \right] \tag{4}$$

A simple approach to *learning* in multi-agent systems is naive learning (NL). A NL agent takes steps of gradient ascent on their own reward function, given the current strategies of all other players at each time step, assuming those strategies are constant. There are many simple scenarios in which NL fails to converge in self-play (Singh et al., 2000). OS approaches address this problem by anticipating and shaping the update step of the opponent. The first OS method is LOLA (Foerster et al., 2018), which includes a Taylor series expansion of its reward function after simulating an opponent NL update in its optimisation objective. LOLA achieved impressive results such as converging to the tit-for-tat (TFT) NE strategy in IPD self-play. However, later work (Letcher et al., 2018) showed that LOLA can fail to preserve NE, and introduced SOS as an alternative, which interpolates between LOLA and an earlier approach known as LookAhead (Zhang & Lesser, 2010). This addresses the problem of "arrogance" and encourages SOS to converge to stable-fix points. COLA (Willi et al., 2022) addresses the inconsistency in which LOLA implicitly assumes that the opponent is a NL agent by explicitly learning the opponent's update function, and using the learned update function to train a policy. The same work also introduced Exact-LOLA (ELOLA), which is similar to LOLA except it does not approximate the perturbed reward function using Taylor series. M-FOS (Lu et al., 2022b; Fung et al., 2023) achieves particularly good performance against NL agents by learning which strategy to play next as a function of both agents' most recently played strategies. This function does not use an explicit model of the opponent, and is learned over multiple episodes against the opponent. SHAPER (Khan et al., 2023) is parameterised with an RNN which captures memory over multiple time scales, allowing it to scale up to high-dimensional and nplayer games (Souly et al., 2023). The Good Shepherd (Balaguer et al., 2022) anticipates the opponent's long-term behaviour by simulating many naive updates of the opponent's parameters, and using these hypothetical future opponent parameters when optimising the reward function.

OS approaches have found desirable solutions in a variety of classic matrix games, as well as extensive-form games such as the Coin Game (Foerster et al., 2018). However,

¹https://github.com/jakelevi1996/welfare_equilibria_public

there are differentiable games which do not contain any NE, in which many approaches (including OS approaches) fail to converge. Even defining a solution concept in such games, which do not contain any NE, is non-trivial. In Section 5.1 we focus on one such example (provided by (Letcher, 2020) in the proof of Theorem 1, page 5) which we refer to as the "Impossible Market", which consists of two players choosing strategies $x \in \mathbb{R}$ and $y \in \mathbb{R}$ with the following reward functions:

$$R^{x}(x,y) = -\frac{x^{6}}{6} + \frac{x^{2}}{2} - xy - \frac{1}{4} \left(\frac{y^{4}}{1+x^{2}} - \frac{x^{4}}{1+y^{2}} \right)$$
(5)
$$R^{y}(x,y) = -\frac{y^{6}}{6} + \frac{y^{2}}{2} + xy + \frac{1}{4} \left(\frac{y^{4}}{1+x^{2}} - \frac{x^{4}}{1+y^{2}} \right)$$
(6)

3. Related Work

Much previous work has explored Stackelberg strategies in various applications (Pita et al., 2009; Yin et al., 2010; 2012), and has also argued that Stackelberg equilibria (in which a leader plays a Stackelberg strategy and a follower plays a BR) should be considered as solution concepts, even in games which are not Stackelberg games (Conitzer, 2016). Using Query Oracles as an abstraction for the opponent BR has facilitated the extension of Stackelberg strategies to deep reinforcement learning (Gerstgrasser & Parkes, 2023). Using commitment schedules has allowed Stackelberg strategies to be learned without access to the opponent reward function (Loftin et al., 2023). However, a common theme in previous work on Stackelberg strategies is the assumption that the opponent always plays a BR to the Stackelberg strategy, introducing possibly arbitrary asymmetry into the game, rather than considering the implications of both players choosing Stackelberg strategies. In contrast, we will consider Stackelberg strategies in selfplay, and generalisations of Stackelberg strategies which display more desirable behaviour in self-play in a broader range of environments.

Beyond Stackelberg strategies, previous work has investigated alternative solution concepts and approaches to learning and equilibrium selection. An Active Markov Game accounts for agents in a game updating their strategies over time, and an Active Equilibrium is a solution concept in which no agent could improve their long term average reward by changing their update function (Kim et al., 2022a;b). FURTHER is an approach for learning in an Active Markov Game using variational inference to approximate opponent update functions (Kim et al., 2022a). The Active Equilibrium solution concept is closely related to a NE in a meta-game, in which the strategy in the meta-game is the update function, and the reward in the meta-game is long-term average reward over time. This interpretation illuminates the connection with MFOS (Lu et al., 2022b), which uses reinforcement learning to shape the opponent learning in the meta-game. Modelling Opponent Learning (MOL) uses a two-phase approach, which first learns the game structure and BR of the opponent, and then guides the opponent's learning in the second phase (Hu et al., 2023).

4. Problem Settings

We consider two different problem settings. The first, in Section 5, is a traditional general-sum learning setting similar to that used by LOLA (Foerster et al., 2018), in which both agents have full access to the environment and the opponent's most recent strategy, including gradients. This could apply to an offline learning setting against a simulated opponent, after which the parameters are frozen and deployed online. The second problem setting in Section 6 assumes a *meta-learning* approach, which uses batches and multiple episodes of learning. Agent strategies are reset between each episode, which has more in common with general-sum meta-learning approaches such as MFOS (Lu et al., 2022b). The resulting shaping strategy is then deployed in the real world. This assumes that at test time the shaper interacts with the shapee over the course of the shapee's training horizon.

5. The Stackelberg Framework For General-Sum Learning

5.1. The Stackelberg Strategy Profile As A Solution Concept

Stackelberg strategies are usually motivated by assuming asymmetry between players, such as one player learning more quickly than the other (Simaan & Cruz Jr, 1973) or having extra information (Chen & Cruz, 1972). However, there are situations in which it is sensible for both players to bilaterally choose Stackelberg strategies, such as in the Impossible Market which was introduced in Section 2, which does not contain any NE. It is straightforward to approximate the opponent BR functions and Stackelberg strategies for each player in the Impossible Market using grid search. The results of doing so are shown in Figure 5 in the Appendix, leading to the Stackelberg strategies $x^* = 0$ and $y^* = 0$. If x and y both play Stackelberg strategies in the Impossible Market, they both receive a reward of 0. Even if either player switches to a BR, the Stackelberg agent receives a reward of -0.654. This outcome is still better for the Stackelberg agent than the worst-case reward that both agents would periodically receive if they both used learning algorithms which converge to a limit cycle (Letcher, 2020), which is typically less than -1. Furthermore, both players cannot play BR strategies simultaneously because the game does not contain any NE. The Stackelberg strategy therefore represents a robust choice of strategy for each agent in the Impossible Market.

In general it is possible for both agents to choose Stackelberg strategies in any two-player game, and therefore the Stackelberg strategy profile (the strategy profile in which both players choose Stackelberg strategies) represents a solution concept for two-player games. Notably, the Stackelberg strategy profile may be distinct from a Stackelberg equilibrium, in which one player plays a Stackelberg strategy and the opponent plays a BR (Conitzer, 2016). Unlike NE, the Stackelberg strategy profile provides a unique solution concept in almost every two-player game (whereas games may have multiple NE), and the Stackelberg strategy profile also generalises to two-player games which do not contain any NE (such as the Impossible Market). Included in the Appendix are depictions of the Stackelberg strategy profile for the games of Matching Pennies (Figure 6), Stag Hunt (Figure 7), Prisoners' Dilemma (Figure 8), Awkward Game (an asymmetric, general-sum, two-player 2×2 matrix game, containing only one NE which is mixed, shown in Figure 9), and IpdTftAlldMix (a version of IPD in which both players must choose a 1D strategy, referring to a parameter which interpolates between the alldefect and TFT strategies, shown in Figure 10), showing that the Stackelberg strategy profile locates the most desirable NE (in terms of maximising both players' rewards) in all five cases, regardless of whether that NE is pure or mixed.

5.2. Opponent-Shaping Derived Via The Stackelberg Framework

It is straightforward to approximate the BR functions and Stackelberg strategies in the Impossible Market with grid search because the strategies for both players are 1D. In games with high-dimensional action spaces it will be impractical to use grid search, and other types of approximation are needed. Here, we demonstrate that approximating argmax in several different ways recovers different OS algorithms, as well as LookAhead. LookAhead is an "opponent-aware" algorithm which, while aware of the opponent's update step, does not shape it.

Suppose that at time t our strategy is x_t and the opponent's strategy is y_t . If we approximate $\hat{y}_t \approx y^*(x_t)$ in Equation 3 with a single step of gradient ascent with learning rate α , and we approximate $x_{t+1} \approx x^*$ with a single step of gradient ascent with learning rate η (without gradient flow from \hat{y}_t to x_{t+1}), we recover the LookAhead learning update:

$$x_{t+1} = x_t + \eta \frac{\partial}{\partial x_t} \left[R^x \left(x_t, \hat{y}_t \right) \right]$$

where $\hat{y}_t = y_t + \alpha \frac{\partial}{\partial y_t} \left[R^y \left(x_t, y_t \right) \right]$ (7)

If instead we do include gradient flow from $\hat{y}_t(x_t)$ (now written as a function of x_t) to x_{t+1} , we recover the ELOLA learning update:

$$x_{t+1} = x_t + \eta \frac{\partial}{\partial x_t} \left[R^x \left(x_t, \hat{y}_t(x_t) \right) \right]$$
where $\hat{y}_t(x_t) = y_t + \alpha \frac{\partial}{\partial y_t} \left[R^y \left(x_t, y_t \right) \right]$
(8)

If we also approximate the reward function R^x with a first-order Taylor series with respect to $\hat{y}_t(x_t)$ then we recover the original LOLA update. If we interpolate between LOLA and LookAhead (which have both been shown to be approximations of Stackelberg strategies) then we recover SOS (which is therefore also an approximation of a Stackelberg strategy). If we approximate $y^*(x)$ in Equation 3 with many steps of gradient ascent (while still approximating x^* with a single step of gradient ascent) then we recover the Good Shepherd learning update.

The OS algorithms we considered so far use explicit update rules for choosing strategies on successive steps of a game. MFOS is distinct from such algorithms in the sense that MFOS learns an update rule which is fixed between time steps but varied between episodes, in order to maximise the expected discounted return in each episode. The MFOS update rule is learned from experience against its opponent, and therefore the extent to which MFOS fits in to the Stackelberg framework depends on the nature of the opponent. Against an opponent which always plays an approximate BR (which is a reasonable approximation for a rational opponent), MFOS learns to play the strategy which maximises reward against the opponent stategy, which approximates the Stackelberg strategy to the extent that the opponent plays an approximate BR. This interpretation is consistent with the observation that MFOS learns to play a ZD extortion strategy against a "Look-Ahead Best Response" agent in IPD (Lu et al., 2022b), which is the best strategy to play against an opponent that plays an approximate BR, and is therefore the Stackelberg strategy in IPD.

5.3. Deriving New Approximate Stackelberg Learning Algorithms

In Section 5.2 we showed that several OS algorithms approximate Stackelberg strategies. In this subsection we demonstrate that different approximations of the Stackelberg strategy can be used to devise new learning algorithms

The Danger Of Arrogance: Welfare Equilibra As A Solution To Stackelberg Self-Play In Non-Coincidental Games



Figure 1: Self-play phase portraits for different algorithms in the Impossible Market. Gradient arrow-lengths are normalised for visual clarity and do not represent gradient magnitude.

with qualitative advantages over existing approaches. Our intention here is simply to demonstrate the possibilities available within this framework. Developing these specific algorithms and further exploring the possibilities within this framework will be explored in future work.

In Section 5.1 we demonstrated that Stackelberg strategies offer a sensible solution concept in the Impossible Market. However previous work (Letcher, 2020) has shown that many algorithms (including several OS algorithms) actually fail to converge in the Impossible Market. Suppose we approximate the inner argmax used to calculate $y^*(x_t)$ in Equation 3 by randomly sampling N different opponent strategies, calculating the opponent's reward for each one, and assuming the opponent chooses the strategy which maximises their reward. We then update our own strategy using a single step of gradient ascent on our own reward function, evaluated using our current strategy and this approximate opponent BR. This can be summarised in the following learning update:

$$x_{t+1} = x_t + \eta \frac{\partial}{\partial x_t} \Big[R^x \left(x_t, \hat{y}_t \right) \Big]$$

where
$$\begin{cases} \hat{y}_t = \max_{n \in \{1, \dots, N\}} \Big[R^y (x_t, y_t + \varepsilon_n) \Big] \\ \varepsilon_n \sim \mathcal{N}(0, \sigma^2) \end{cases}$$
(9)

This update uses a SAmpling-based approximation to the inner argmax and a Gradient-Ascent-based approximation to the outer argmax, so we refer to this learning algorithm as "SaGa". The behaviour of SaGa self-play in the Impossible Market is shown in a phase portrait in Figure 1, alongside equivalent phase portraits for NL and ELOLA. These results demonstrate how NL and ELOLA self-play both consistently converge to qualitatively similar limit-cycles, whereas SaGa self-play consistently converges to the Stackelberg strategy profile, simply by using a sampling-based approximation to the opponent BR.

Like any algorithm which uses gradient ascent, existing OS algorithms are vulnerable to getting stuck in nonglobal local maxima. This problem is particularly prevalent for games such as Stag Hunt and IpdTftAlldMix, for which the reward function against a perfect BR opponent is not concave (see Figures 7 and 10 in the Appendix). It is therefore natural to consider whether we can achieve more robust performance by using sampling-based approximations for the outer argmax as well as the inner argmax, which can be achieved using the following learning update:

$$x_{t+1} = (1-\eta)x_t + \eta \max_{m \in \{1,...,M\}} \left[R^x \left(x_t + \varepsilon_m^x, \hat{y}_t \left(\varepsilon_m^x \right) \right) \right]$$

where
$$\begin{cases} \hat{y}_t \left(\varepsilon_m^x \right) &= \max_{n \in \{1,...,N\}} \left[R^y \left(x_t + \varepsilon_m^x, y_t + \varepsilon_n^y \right) \right] \\ \varepsilon_n^y, \varepsilon_m^x &\sim \mathcal{N}(0, \sigma^2) \end{cases}$$
(10)

This update uses a SAmpling-based approximation to the inner argmax *and* a SAmpling-based approximation to the outer argmax, so we refer to this learning algorithm as "SaSa". The behaviour of SaSa self-play in the Stag Hunt game is shown in a phase portrait in Figure 2, along-side equivalent phase portraits for NL and ELOLA. These results demonstrate how NL and ELOLA self-play both converge to sub-optimal NE from relatively large areas of the state-space, whereas SaSa self-play consistently converges to the optimal NE, simply by using sampling-based argmax approximations.

The results for SaSa against NL in IPD are included in Figure 3. SaSa achieves a stable mean reward of -0.730, which is greater than the reward for mutual-TFT and suggests that SaSa on average learns a ZD extortion strategy (Press & Dyson, 2012), demonstrating that SaSa can scale up to games with moderately high-dimensional action spaces.



Figure 2: Self-play phase portraits for different algorithms in the Stag Hunt game. Axes refer to probability of hunting a stag for each player respectively, with the upper right corner corresponding to the optimal NE.

6. Welfare Equilibria

6.1. Non-Coincidental Games And The Chicken Catastrophe

As mentioned in Section 5.1, Stackelberg strategy profiles find optimal NE in a variety of zero-sum and general-sum two-player games. A common property of many of these games is that the most rewarding NE for one player coincides with the most rewarding NE for the other player, so we refer to such games as "Coincidental games". In general however, there are many games in which this property does not hold, such as the Chicken Game. The Stackelberg strategy for either player in the Chicken Game is to drive straight (to which the opponent BR is to chicken out), however if both players drive straight then they both experience the unique worst possible outcome, in which neither player's strategy is a BR (shown in Figure 11 in the Appendix). We formally define non-coincidental games as games in which the Stackelberg strategy profile is not a NE², and refer to the failure of the Stackelberg strategy profile to find a desirable outcome in the chicken game (a non-coincidental canonical matrix game) as "The Chicken Catastrophe".

In Section 5.2 we showed that many OS algorithms approximate Stackelberg strategies. Considering that the Stackelberg strategy profile fundamentally fails in the Chicken Game, this provides a normative theory for why OS algorithms which approximate Stackelberg strategies *should* *also* fail in self-play in the Chicken Game. This analysis is consistent with empirical evidence (Willi et al., 2022), and is problematic due to the importance of finding desirable solutions in self-play (which we highlighted in Section 1). In the following subsections we present a generalisation of Stackelberg strategies which is capable of finding NE in self-play in non-coincidental games. Our analysis excludes M-FOS meta-self-play and the SOS algorithm, which both address arrogance and are outside our framework.

6.2. Addressing The Chicken Catastrophe

A simple solution for the Chicken Catastrophe is for each player to maximise egalitarian social welfare (defined in Equation 15) instead of self-reward, while still assuming the opponent plays a BR. If both players take this approach, then they both concurrently select the unique mixed NE in the Chicken Game, and both receive a better reward than if they had both played Stackelberg strategies (shown in Figure 12 in the Appendix). However, this approach fails in self-play in the non-coincidental game Bach Or Stravinsky (BOS). Both pure NE in BOS have maximal and equal egalitarian social welfare, so there is no way for both players to consistently favour one NE over the other in a way which is fair to both players. A solution to BOS self-play is for each player to maximise fairness (defined in Equation 16) while assuming the opponent plays a BR. If both players take this approach, then they both concurrently select the unique mixed NE in BOS, which is maximally fair to both players, and also better than failing to coordinate (shown in Figure 16). However, this approach fails in self-play in a game we introduce and refer to as the Eagle Game. In the Eagle Game, maximising egalitarian social welfare leads to better outcomes for both players than maximising selfreward or fairness, assuming the opponent plays a BR in all cases (shown in Figure 23).

The previous paragraph demonstrates that choosing an ap-

²Interestingly, the Impossible Market is a non-coincidental game, although the Stackelberg strategy profile achieves desirable behaviour in this case. One response to this observation follows from the concept of arrogance penalties, which are introduced in Section A.1 of the Appendix. Specifically, the Impossible Market has negative arrogance penalties, so both players are rewarded for "arrogantly" assuming an opponent BR and choosing a Stackelberg strategy. This is in contrast to the other non-coincidental games we consider, wherein mutual arrogance has a cost.



Figure 3: Results for 5000 steps of learning with SaSa ($\eta = 0.1, M = 20, N = 20, \sigma = 1$) against NL ($\eta = 0.1$) in IPD (using discounted returns with $\gamma = 0.96$), averaged over 100 trials. Final mean rewards are -0.72972 and -1.97547 for SaSa and NL respectively.

propriate welfare function and then maximising that welfare function (while assuming that the opponent will play a *greedy* BR) allows each player to effectively choose between desirable NE strategies in a wide variety of games. However, to our knowledge there is no single welfare function that consistently leads to desirable NE strategies in every possible game. This motivates us to define the "Welfare Equilibrium" (WE) strategy for *any* given choice of welfare function as shown below, followed by examples of empircally useful welfare functions (the opponent BR functions are defined as in Equations 1 and 2):

$$x_{\rm WE}^* = \operatorname*{argmax}_{x} \left[w^x \left(x, y^*(x) \right) \right] \tag{11}$$

$$y_{\rm WE}^* = \underset{y}{\operatorname{argmax}} \left[w^y \left(x^*(y), y \right) \right]$$
(12)

$$w_{\text{greedy}}^x(x,y) = R^x(x,y) \tag{13}$$

$$w_{\text{greedy}}^{y}(x,y) = R^{y}(x,y) \tag{14}$$

$$w_{\text{egalitarian}}(x,y) = \min\left(R^x(x,y), R^y(x,y)\right)$$
(15)

$$w_{\text{fairness}}(x,y) = -\left|R^x(x,y) - R^y(x,y)\right|$$
(16)

The Stackelberg strategy for either player is a special case of a WE strategy in which that player chooses to maximise a greedy welfare function (self-reward). The WE strategy is therefore a generalisation of the Stackelberg strategy. Included in the Appendix are depictions of WE strategy profiles for the Chicken Game (Figure 12), Baby Chicken Game (a variant of the Chicken Game with a less severe punishment for mutually driving straight, which leads to a more intelligible depiction of the WE profile, shown in Figure 14), BOS (Figure 16), Tandem Game (Figure 18), Ultimatum Game (Figure 20), and Eagle Game (Figure 23), showing that WE strategy profiles with appropriate welfare functions can locate desirable NE and provide better rewards for both players than the equivalent Stackelberg strategy profile in all six cases.

Table 1: Rewards in Tandem Game for WE strategy profiles with different welfare functions. Strategies are restricted to the range [-2, 3] for numerical tractability, as in Figure 18.

	Greedy (G)	Egalitarian (E)	Fairness (F)
G	(-30.00, -30.00)	(-6.24, -11.24)	(-6.24, -11.24)
Ε	(-11.24, -6.24)	(0.00 , 0.00)	(0.00, 0.00)
F	(-11.24, -6.24)	(0.00 , 0.00)	(0.00, 0.00)

A natural question to ask is when (if ever) it is actually in a player's interest to maximise any welfare function other than a greedy welfare function. The performance of any learning algorithm in a multi-agent system depends on the nature and dynamics of the opponent. Against an opponent that will always play a BR, the Stackelberg strategy (equivalent to a greedy WE strategy) is always the best strategy to play (by definition). However, as outlined in Section 1, it is also important to consider how an algorithm performs in self-play, and for two WE agents in self-play in noncoincidental games, the greedy welfare function is often not the best welfare function to choose. This is demonstrated in Table 1, which shows the possible rewards for WE agents in self-play choosing between greedy, egalitarian and fairness welfare functions and then playing WE strategies in the Tandem Game (Letcher et al., 2018). In this example, not only is it sometimes in a player's interest to maximise a non-greedy welfare function, but in fact choosing the greedy welfare function is a *strictly dominated strategy*. This means that for either player, regardless of whichever welfare function is chosen by the opponent, a better reward would be achieved by choosing egalitarian or fairness welfare functions instead of a greedy welfare function. Therefore in this example, both players have a greedy incentive to maximise a non-greedy welfare function.

The welfare functions considered so far are not invariant to affine transformations of either player's reward function. This can be addressed by introducing "arrogance penalties", which are discussed in Section A.1 of the Appendix.



Figure 4: Comparing ELOLA ($\eta = 0.1, \alpha = 25$, averaged over 100 trials) with WelFuSeElola (same ELOLA hyperparameters, e = 3, s = 1000, b = 100, averaged over five random seeds) against different opponents in Chicken Game. Top row: against NL. Bottom row: self-play. Left column: rewards for ELOLA. Centre column: rewards for WelFuSeElola. Right column: welfare functions chosen by WelFuSeElola in each episode.

6.3. Welfare Function Search

The previous subsection demonstrated that the best welfare function to choose in a WE strategy depends on the game and the nature of the opponent, but did not explain how to choose the best welfare function using a practical algorithm, which we introduce in this subsection. The welfare function is effectively an optimisation objective function (given a suitable approximation of the opponent BR function), but rather than searching over the entire continuous space of all possible welfare functions (Lu et al., 2022a), we instead take a simpler approach and assume access to a finite set of pre-defined welfare functions, such as the three that were considered in Section 6.2, which have been shown to find mutually desirable NE in a variety of canonical non-coincidental games. We also assume access to an inner OS algorithm (such as ELOLA), which is able to optimise a given choice of welfare function, rather than the usual self-reward. In the OS update step, the opponent is always assumed to maximise self-reward, consistent with the definition of WE in Equations 11 and 12. This leads to the algorithm "Welfare Function Search" (WelFuSe), a practical algorithm for adaptively choosing a welfare function from experience, which preserves performance against NL while avoiding catastrophe in self-play. The principle behind WelFuSe is to treat the problem of choosing a welfare function as a discrete bandit problem, which is solved over multiple episodes using a batched variant of posterior sampling (Sutton & Barto, 2018), in order to maximise final *self-reward* from each episode. After each episode, all agents' strategies are reset, and a new batch of welfare functions are sampled. The full WelFuSe algorithm is described in Algorithm 1 in the Appendix.

The results from using WelFuSe with ELOLA as the inner OS algorithm (which we jointly refer to as "Wel-FuSeElola") in the Chicken Game are shown in Figure 4 in the Appendix. These results demonstrate that against NL, WelFuSeElola quickly learns to reject egalitarian and fairness welfare functions in order to optimise a greedy welfare function (Figure 4c). This leads to the optimal reward against NL (Figure 4b), which is equal to the performance of ELOLA (Figure 4a). However in self-play (when both agents are WelFuSeElola agents, playing against each other with separate welfare function distributions, while each assumes their opponent takes an approximate NL update step), WelFuSeElola learns to reject the greedy welfare function and instead optimise a mixture of egalitarian and fairness welfare functions (Figure 4f), leading to maximally fair and egalitarian outcomes for both WelFuSeElola self-play agents (Figure 4e). This is much more desirable for both agents than the catastrophe experienced by ELOLA agents in self-play (Figure 4d), and we therefore conclude that WelFuSeElola has solved the chicken catastrophe.

We end this subsection by emphasising that WelFuSe can be applied given any reasonable choice of inner OS algorithm, and therefore WelFuSe should be considered as an *extension* to existing OS algorithms, rather than a mutually exclusive alternative. As we have shown, WelFuSe learns to select a non-greedy welfare function when it is in the agent's interest to do so. In other cases, WelFuSe learns to reject all other welfare functions and learns to only optimise a greedy welfare function, in which case WelFuSe simply reduces to its inner OS algorithm.

7. Conclusions

In Section 5.1 we showed that Stackelberg strategy profiles provide a sensible solution concept in a variety of games, which we refer to as coincidental games, as well as in the Impossible Market, which contains no NE. In Section 5.2 we showed that Stackelberg strategies represent a unifying framework from which many existing OS algorithms can be derived as approximations. In Section 5.3 we demonstrated the value of this framework by using it to derive new algorithms which have qualitative advantages over previous approaches. For example, SaGa was able to consistently converge in the Impossible Market (unlike many learning algorithms), and SaSa was able to consistently converge to the global optimum in Stag Hunt while avoiding the nonglobal local optimum. SaSa was also able to strongly dominate NL in IPD in a single episode.

In Section 6.1 we defined non-coincidental games as games in which the Stackelberg strategy profile is not a NE, which notably includes several canonical matrix games, and illustrates why OS algorithms can fail in self-play in such games. Avoiding catastrophe in self-play is important for any learning algorithm to be safe to deploy in multi-agent systems in the real world. To this end, in Section 6.2 we introduced WE as a generalisation of the Stackelberg framework, which can recover desirable NE solutions in noncoincidental games. We showed that against another WE agent in the Tandem game, choosing a greedy welfare function is a strictly dominated strategy, and therefore in this example both players have a greedy incentive to maximise a non-greedy welfare function. In Section 6.3 we introduced WelFuSe as a practical approach to choosing between welfare functions, which was able to preserve the performance of ELOLA against NL, while also avoiding catastrophe in self-play.

Despite these contributions, this work offers alluring directions for future work, such as developing new OS algorithms using more sophisticated approximations to Stackelberg strategies, and more sophisticated approaches for choosing effective welfare functions. Overall, we hope that this work takes us closer towards understanding how to design multi-agent learning algorithms that are both safe and effective in the real world.

Acknowledgements

This work was supported by the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems (grant number EP/S024050/1). CSDW received generous support by the Cooperative AI Foundation. This work was supported by Armasuisse Science+Technology.

Impact Statement

This paper presents contributions which are intended to improve understanding and learning outcomes in generalsum multi-agent systems. We hope that this will contribute to improved outcomes in multi-agent systems in the real world, including for models which learn while interacting with humans. Such models have caused negative realworld outcomes in the past (Wolf et al., 2017) due to naive design, and we hope that our contributions may help to design better models which avoid such outcomes in the future.

References

- Balaguer, J., Koster, R., Summerfield, C., and Tacchetti, A. The good shepherd: An oracle agent for mechanism design. arXiv preprint arXiv:2202.10135, 2022. 2
- Chen, C. and Cruz, J. Stackelburg solution for two-person games with biased information patterns. *IEEE Transactions on Automatic Control*, 17(6):791–798, 1972. 3
- Conitzer, V. On stackelberg mixed strategies. *Synthese*, 193(3):689–703, 2016. 3, 4
- Foerster, J., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I. Learning with opponentlearning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 122–130, 2018. 1, 2, 3
- Fung, K., Zhang, Q., Lu, C., Willi, T., and Foerster, J. N. Analyzing the sample complexity of model-free opponent shaping. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023. 2
- Gerstgrasser, M. and Parkes, D. C. Oracles & followers: Stackelberg equilibria in deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 11213–11236. PMLR, 2023. 3
- Hu, Y., Han, C., Li, H., and Guo, T. Modeling opponent learning in multiagent repeated games. *Applied Intelligence*, 53(13):17194–17210, 2023. 3
- Khan, A., Willi, T., Kwan, N., Tacchetti, A., Lu, C., Grefenstette, E., Rocktäschel, T., and Foerster, J. Scaling opponent shaping to high dimensional games. *arXiv preprint arXiv:2312.12568*, 2023. 2

- Kim, D.-K., Riemer, M., Liu, M., Foerster, J., Everett, M., Sun, C., Tesauro, G., and How, J. P. Influencing longterm behavior in multiagent reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 18808–18821, 2022a. 3
- Kim, D.-K., Riemer, M., Liu, M., Foerster, J. N., Tesauro, G., and How, J. P. Game-theoretical perspectives on active equilibria: A preferred solution concept over nash equilibria. arXiv preprint arXiv:2210.16175, 2022b. 3
- Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Al Sallab, A. A., Yogamani, S., and Pérez, P. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23 (6):4909–4926, 2021. 1
- Letcher, A. On the impossibility of global convergence in multi-loss optimization. *arXiv preprint arXiv:2005.12649*, 2020. 3, 5
- Letcher, A., Foerster, J., Balduzzi, D., Rocktäschel, T., and Whiteson, S. Stable opponent shaping in differentiable games. arXiv preprint arXiv:1811.08469, 2018. 2, 7
- Loftin, R., Çelikok, M. M., van Hoof, H., Kaski, S., and Oliehoek, F. A. Uncoupled learning of differential stackelberg equilibria with commitments. *arXiv preprint arXiv:2302.03438*, 2023. 3
- Lu, C., Kuba, J., Letcher, A., Metz, L., Schroeder de Witt, C., and Foerster, J. Discovered policy optimisation. *Advances in Neural Information Processing Systems*, 35: 16455–16468, 2022a. 8
- Lu, C., Willi, T., De Witt, C. A. S., and Foerster, J. Modelfree opponent shaping. In *International Conference on Machine Learning*, pp. 14398–14411. PMLR, 2022b. 2, 3, 4
- Nash, J. F. Equilibrium points in n-person games. Proceedings of the national academy of sciences, 36(1):48–49, 1950. 2
- Pita, J., Jain, M., Ordónez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., and Kraus, S. Using game theory for los angeles airport security. *AI magazine*, 30 (1):43–43, 2009. 3
- Press, W. H. and Dyson, F. J. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences*, 109(26):10409–10413, 2012. 5
- Simaan, M. and Cruz Jr, J. B. On the stackelberg strategy in nonzero-sum games. *Journal of Optimization Theory* and Applications, 11(5):533–555, 1973. 1, 2, 3

- Singh, S., Kearns, M. J., and Mansour, Y. Nash convergence of gradient dynamics in general-sum games. In UAI, pp. 541–548, 2000. 2
- Souly, A., Willi, T., Khan, A., Kirk, R., Lu, C., Grefenstette, E., and Rocktäschel, T. Leading the pack: N-player opponent shaping. arXiv preprint arXiv:2312.12564, 2023. 2
- Subramanian, J., Sinha, A., and Mahajan, A. Robustness and sample complexity of model-based marl for generalsum markov games. *Dynamic Games and Applications*, pp. 1–33, 2023. 1
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018. 1, 8
- Wang, Y., Damani, M., Wang, P., Cao, Y., and Sartoretti, G. Distributed reinforcement learning for robot teams: a review. *Current Robotics Reports*, 3(4):239–257, 2022. 1
- Willi, T., Letcher, A. H., Treutlein, J., and Foerster, J. Cola: consistent learning with opponent-learning awareness. In *International Conference on Machine Learning*, pp. 23804–23831. PMLR, 2022. 1, 2, 6
- Wolf, M. J., Miller, K., and Grodzinsky, F. S. Why we should have seen that coming: comments on microsoft's tay "experiment," and wider implications. *Acm Sigcas Computers and Society*, 47(3):54–64, 2017. 1, 9
- Yin, Z., Korzhyk, D., Kiekintveld, C., Conitzer, V., and Tambe, M. Stackelberg vs. nash in security games: interchangeability, equivalence, and uniqueness. In AAMAS, volume 10, pp. 6, 2010. 3
- Yin, Z., Jiang, A. X., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., and Sullivan, J. P. Trusts: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI magazine*, 33(4): 59–59, 2012. 3
- Zhang, C. and Lesser, V. Multi-agent learning with policy prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, pp. 927–934, 2010. 2
- Zhang, K., Yang, Z., and Başar, T. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control*, pp. 321–384, 2021. 1

A. Appendix

A.1. Arrogance Penalties And Invariant Welfare Functions

The egalitarian and fairness welfare functions provide mutually desirable NE in a variety of non-coincidental games, however they share a common disadvantage, which is that unlike BR functions and Stackelberg strategies, they are not invariant to shifting and positive scaling of either player's reward function. We can address this issue by introducing "arrogance penalties". When a player decides to play the Stackelberg strategy (drive straight) in the chicken game, they effectively assume that the opponent will play a BR (chicken out) to the Stackelberg strategy, in which case the former player expects a reward of +1. We refer to the reward for either player when they play a Stackelberg strategy and the opponent plays a BR as the "Stackelberg baseline". When both players choose Stackelberg strategies in the chicken game and mutually drive straight, they both receive a reward of -100, which is 101 less than the Stackelberg baseline. This discrepency between expected and received rewards for each player is effectively a penalty for arrogantly assuming that they could "call the shots" while the opponent would "fall in line". We therefore refer to the difference between the Stackelberg baseline and the reward from the Stackelberg strategy profile as the "arrogance penalty" for each player respectively. Following this intuition, we can define the shift-normalised reward functions $\pi^x(x, y)$ and $\pi^y(x, y)$ for any strategy profile as the reward that each player receives relative to their Stackelberg baseline:

$$\pi^{x}(x,y) = R^{x}(x,y) - R^{x}\left(x^{*},y^{*}\left(x^{*}\right)\right)$$
(17)

$$\pi^{y}(x,y) = R^{y}(x,y) - R^{y}\left(x^{*}\left(y^{*}\right), y^{*}\right)$$
(18)

We note that for either player, if their opponent will always play a BR, then the shift-normalised reward function for the former player is negative everywhere except when they play a Stackelberg strategy, at which point the shiftnormalised reward function is zero (this follows from the definition of the Stackelberg strategies in equations 3 and 4). We also note that, because both players' BR functions and Stackelberg strategies are invariant to shifting of either player's reward function, so too are $\pi^x(x, y)$ and $\pi^y(x, y)$, because any shift in $R^x(x, y)$ or $R^y(x, y)$ would be cancelled out between the two terms in equations 17 and 18 respectively. This allows us to define the shift-normalised egalitarian welfare function as follows:

$$\pi_{\text{egalitarian}}(x,y) = \min\left(\pi^x(x,y), \pi^y(x,y)\right)$$
(19)

We note that in the special case of coincidental games, when both players use the shift-normalised egalitarian welfare function, the resulting WE strategy profile is always the Stackelberg strategy profile. The shift-normalised egalitarian welfare function is invariant to shifts in either player's reward function, but it is not invariant to scaling. To address this issue, for non-coincidental games, we can introduce the affinely-normalised reward functions (equal to the shift-normalised reward functions normalised by the absolute arrogance penalties) and affinely-normalised egalitarian welfare function as follows:

$$\bar{\pi}^{x}(x,y) = \frac{R^{x}(x,y) - R^{x}\left(x^{*},y^{*}\left(x^{*}\right)\right)}{\left|R^{x}\left(x^{*},y^{*}\left(x^{*}\right)\right) - R^{x}\left(x^{*},y^{*}\right)\right|}$$
(20)
$$\bar{\pi}^{y}(x,y) = \frac{R^{y}(x,y) - R^{y}\left(x^{*}\left(y^{*}\right),y^{*}\right)}{\left|R^{y}\left(x^{*}\left(y^{*}\right),y^{*}\right) - R^{y}\left(x^{*},y^{*}\right)\right|}$$
(21)
$$\bar{\pi}_{\text{egalitarian}}(x,y) = \min\left(\bar{\pi}^{x}(x,y),\bar{\pi}^{y}(x,y)\right)$$
(22)

The affinely-normalised fairness welfare function can be defined analogously. The affinely-normalised egalitarian welfare function (Equation 22) is invariant to shifting and positive scaling of either player's reward function. Therefore, if it is used by both players as the welfare function in a WE strategy profile, then the strategies chosen by both players are also invariant to shifting and positive scaling of either player's reward function. Furthermore, in symmetric games, the Stackelberg baseline and arrogance penalty are necessarily symmetric for both players, which implies that the affinely-normalised egalitarian WE strategy profile is simply equivalent to the egalitarian WE strategy profile in symmetric games.

What is offered by the affinely-normalised egalitarian welfare function in terms of mathematical appeal is compensated by its greater computational burden in asymmetric games, because it depends on knowing the Stackelberg strategies for both players and their respective BRs before it can be used to calculate a WE strategy. This would limit the applicability to online learning settings, although it could be used if there was the opportunity to approximate Stackelberg strategies and BRs offline before the start of online learning. This approach would also be suitable to offline learning settings, in which the Stackelberg strategies, BRs, and WE strategies could be learned sequentially. In general, we leave the development of invariant welfare functions which are more conducive to efficient practical online implementation as an interesting direction for future work.

(a) Payoff table for Prisoners' Dilemma

	Cooperate	Defect
Cooperate	(-1.0, -1.0)	(-3.0, 0.0)
Defect	(0.0, -3.0)	(-2.0 , - 2.0)

(b) Payoff table for Matching Pennies

	Heads	Tails
Heads	(1.0, -1.0)	(-1.0 , 1.0)
Tails	(-1.0 , 1.0)	(1.0 , -1.0)

(c) Payoff table for Stag Hunt

	Stag	Hare
Stag	(10.0, 10.0)	(1.0, 8.0)
Hare	(8.0, 1.0)	(5.0, 5.0)

(d) Payoff table for Chicken Game

	Chicken out	Drive straight
Chicken out	(0.0, 0.0)	(-1.0, 1.0)
Drive straight	(1.0 , -1.0)	(-100.0, -100.0)

(e) Payoff table for Bach Or Stravinsky (BOS)

	Bach	Stravinsky
Bach	(2.0, 1.0)	(0.0, 0.0)
Stravinsky	(0.0, 0.0)	(1.0, 2.0)

(f) Payoff table for Baby Chicken Game

	Chicken out	Drive straight
Chicken out	(0.0, 0.0)	(-1.0, 1.0)
Drive straight	(1.0, -1.0)	(-3.0 , - 3.0)

(g) Payoff table for Awkward Game

	Cooperate	Defect
Cooperate	(3.0, 1.0)	(1.0, 3.0)
Defect	(2.0, 5.0)	(4.0, 2.0)

(h) Payoff table for Eagle Game

	Cooperate	Defect
Cooperate	(4.0, 1.0)	(-4.0, -1.0)
Defect	(-2.0, -3.0)	(2.0, 3.0)

Table 2: Payoff tables for matrix games

Algorithm 1 Welfare Function Search (WelFuSe)

Require: Set of welfare functions $\mathcal{W} = \{w_1, \dots, w_{|\mathcal{W}|}\}$, number of episodes *e*, steps per episode *s*, batch size *b*, inner OS algorithm $os(w_k, s)$ 1: for $j \in (1, ..., b)$ do $k \sim \mathcal{U}(\{1, \dots, |\mathcal{W}|\})$ // Sample initial welfare functions uniformly 2: $w^{(1,j)} \leftarrow w_k$ 3: 4: end for 5: for $i \in (1, ..., e)$ do for $j \in (1, \ldots, b)$ do 6: $os(w^{(i,j)},s)$ // Optimise welfare function using OS 7: Observe final reward $r^{(i,j)}$ 8: 9: end for 10: Reset all agents' parameters // Consider each batch-index 11: for $j \in (1, ..., b)$ do 12: // Consider each welfare function w_k 13: 14: for $k \in (1, \ldots, |\mathcal{W}|)$ do $m_k \sim \mathcal{U}(\{n \in \{1, \dots, b\} : w^{(i,n)} = w_k\})$ // Sample batch-index which used w_k 15: end for 16: $\hat{k} \leftarrow \operatorname{argmax} \left[r^{(i,m_k)} \right]$ // Select welfare function with best sampled reward 17: $w^{(i+1,j)} \leftarrow w_{\hat{k}}$ // Set welfare function for next episode 18: 19: end for 20: end for







Figure 6: Stackelberg strategy profile (Greedy WE) for MatchingPennies $x^*=0.500, y^*=0.500, R^x=0.000, R^y=0.000$







Figure 8: Stackelberg strategy profile (Greedy WE) for PrisonersDilemma $x^* = 0.000, y^* = 0.000, R^x = -2.000, R^y = -2.000$







Figure 10: Stackelberg strategy profile (Greedy WE) for IpdTftAlldMix $x^* = 1.000, y^* = 1.000, R^x = -1.000, R^y = -1.000$



Figure 11: Stackelberg strategy profile (Greedy WE) for ChickenGame $x^* = 0.000, y^* = 0.000, R^x = -100.000, R^y = -100.000$



Figure 12: Egalitarian WE for ChickenGame $x^* = 0.989, y^* = 0.989, R^x = -0.011, R^y = -0.011$



Figure 13: Stackelberg strategy profile (Greedy WE) for BabyChickenGame $x^* = 0.000, y^* = 0.000, R^x = -3.000, R^y = -3.000$



Figure 14: Egalitarian WE for BabyChickenGame $x^* = 0.666, y^* = 0.666, R^x = -0.334, R^y = -0.334$



Figure 15: Stackelberg strategy profile (Greedy WE) for Bach Or Stravinsky $x^* = 1.000, y^* = 0.000, R^x = 0.000, R^y = 0.000$



Figure 16: Fairness WE for Bach Or Stravinsky $x^* = 0.666, y^* = 0.334, R^x = 0.667, R^y = 0.667$







Figure 18: Egalitarian WE for Tandem $x^* = 0.500, y^* = 0.500, R^x = 0.000, R^y = 0.000$















Figure 22: Fairness WE for EagleGame $x^* = 0.000, y^* = 0.500, R^x = 0.000, R^y = 0.000$



Figure 23: Egalitarian WE for EagleGame $x^* = 0.000, y^* = 0.000, R^x = 2.000, R^y = 3.000$