Advanced Equilibrium Analysis in Game-Theoretic Models: Extending Evolutionary Algorithms to Complex Strategic Scenarios

Natalia Trankova ,
MSc - Skolkovo Institute of Science and Technology, nattrankova@gmail.com

Olga Chumakova,
BSc – National Research University Higher School of Economics, olchumakova@edu.hse.ru

Andrei Zhuk,
MSc – University of Pennsylvania, andrei.zhuk@alumni.upenn.edu

Dmitrii Rykunov,
BSc – National Research University Higher School of Economics, dprykunov@gmail.com

Ivan Giganov,
MSc - Northwestern University, igiganov@nes.ru

Yaroslav Starukhin - QuantumBlack,
yaroslav_starukhin@quantumblack.com

Ivan Serov - QuantumBlack,
ivan.serov.a@gmail.com

DOI: 10.31364/SCIRJ/v12.i10.2024.P10241000
http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P10241000

Abstract: This article presents a comprehensive empirical and theoretical evaluation of an evolutionary optimization algorithm designed to locate multiple equilibria in game-theoretic models. Building upon a general-purpose framework established in [1], the algorithm is extended to complex strategic scenarios, including mixed strategies, dynamic environments, and high-dimensional strategy spaces. Its performance is rigorously analyzed across various models, examining the impact of parameters such as population size and mutation rate on convergence dynamics.

The results demonstrate the algorithm's robustness in identifying Evolutionarily Stable Strategies (ESS) and navigating multiple equilibria, even for challenging scenarios. Theoretical and empirical evidence support its effectiveness. These findings have significant implications for modeling strategic behavior in economics and social sciences.

Keywords: **Game theory, evolutionary algorithms, Nash equilibrium, prisoner's dilemma, computational modeling, strategic interactions, multiple equilibria, genetic algorithm, evolutionary game theory, evolutionary stable strategies (ESS), algorithmic game theory, agent-based modeling, convergence dynamics.**

1 Introduction

Accurately predicting outcomes in strategic interactions is a core challenge in game theory, mainly when dealing with complex social and economic systems. While traditional frameworks such as Nash Equilibrium [2] have provided valuable insights, they often prove insufficient in scenarios characterized by multiple equilibria. This shortcoming limits game-theoretic models' explanatory and predictive power, especially when agents make decisions based on uncertainty and incomplete information.

Recent developments in evolutionary game theory have introduced computational methodologies that systematically navigate the strategy space to identify multiple equilibria [3], [4]. Evolutionary algorithms offer a versatile framework for equilibrium analysis by simulating dynamic interactions among agents. Nevertheless, these algorithms frequently encounter difficulties with convergence in high-dimensional strategy spaces and often require meticulous parameter tuning, which constrains their practical application [1] [5] [6]. For instance, while coevolutionary algorithms have demonstrated efficacy in identifying Nash equilibria in specific models like symmetric Cournot games, [7] [8], their application has largely been restricted to limited scenarios such as two-player or symmetrical setups. Furthermore, single-population evolutionary algorithms have shown reduced effectiveness in handling cases involving mixed strategies [9].

To address these limitations, previous research [1] introduced a general-purpose evolutionary optimization algorithm to systematically locate multiple equilibria in non-cooperative game-theoretic models. This algorithm, incorporating elements of genetic algorithms and coevolutionary dynamics, was validated on classical benchmarks such as the Prisoner's Dilemma, demonstrating its capability to identify evolutionary stable strategies (ESS) and evolutionary stable sets. Despite these promising results, questions remained about the algorithm's robustness and efficiency in more complex, high-dimensional scenarios, as well as its sensitivity to different parameter configurations.

The present study builds on this foundation by extending the algorithm's empirical and theoretical analysis to explore its performance under a broader range of conditions. The algorithm is evaluated across diverse game-theoretic models, focusing on the impact of critical parameters such as population size, mutation rates, and interaction frequencies on convergence behavior. This assessment aims to deepen an understanding of the algorithm's operational dynamics and its efficacy in navigating intricate strategy spaces.

Additionally, this research delves into the theoretical aspects of the algorithm's behavior, particularly its ability to converge to mixed strategy equilibria and evolutionary stable sets in environments characterized by uncorrelated asymmetry. This paper provides formal and empirical evidence to elucidate the conditions under which the algorithm successfully identifies and stabilizes multiple equilibria, thereby contributing to a more nuanced understanding of equilibrium selection in complex systems.

The results of this work are applicable in game-theoretic models across various disciplines, including economics, political science, and behavioral ecology. By enhancing the computational tools available for analyzing strategic interactions, this study offers a robust framework for researchers and practitioners seeking to model and predict the behavior of agents operating in complex and uncertain environments.

2 Model and Algorithm Specification

2.1 Model of Social Interaction

This study builds upon the general-purpose game-theoretic model introduced in the previous work [1], where agents interact strategically within a defined social framework. The core model conceptualizes each agent's decision-making process as a series of choices from a finite set of actions guided by maximizing individual payoffs. An evolutionary optimization algorithm is employed to explore the strategy space to address the challenge of locating multiple equilibria in such interactions.

2.1.1 Model Extensions and Modifications

To accommodate the model for more complex strategic scenarios, several enhancements to the original model are introduced:

1. Dynamic Strategy Adaptation:
   - Incorporating a dynamic adaptation mechanism allows agents to adjust their strategies based on the evolving behavior of other agents. This modification more accurately captures the iterative nature of real-world decision-making processes, especially in environments with changing conditions or incomplete information.
2. Handling Mixed Strategies:
   - Including mixed strategies expands the model beyond the original focus on pure strategy. Agents can now select probabilistic combinations of actions, providing a more realistic representation of strategic interactions. This is especially useful in scenarios where pure strategies are insufficient to capture the complexity of agent behavior.
3. Extended Parameter Space:
   - Expanding the model's parameter range introduces stochastic elements, adding randomness to agent decision-making and payoff outcomes. This change enables a more comprehensive exploration of strategy spaces and offers insights into equilibrium stability under varied conditions.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

2.2 Notation

The notation used in this article:

- $\Gamma = (A,\ \Delta,\ \Psi)$: game
- A: The set of agents participating in the game, $A = \{a_1, a_2, \dots, a_n\}$, $|A| = n$
- Δ_i: The action set of agent i, representing all possible actions they can take. $\Delta = \{D_1, \dots, D_n\}$,
  - where D_i – decision-space, $D_i = \{d_{i_1}, \dots, d_{i_m}\}$,
    - where d_{i_j} – decision-set. $\forall \mathrm{d}_{i_j}$ is a discrete set or compact space.
    - $D_i \neq \emptyset$
  - $\Delta \cong \mathrm{A}; |\Delta| = |\mathrm{A}| = \mathrm{n}$
- Ψ: payoff function
  - $\prod_{d \in \Delta} d \to P$
  - $\mathrm{P} \cong \mathrm{A}; |\mathrm{P}| = |\mathrm{A}| = \mathrm{n}$
  - $P = \{p_1, \dots, p_n\}, p_i \in \mathbb{R}$
- $p_i(\phi_i, \phi_{-i})$: The payoff function of agent i, representing their expected payoff given their strategy ϕ_i, and the strategies of all other agents ϕ_{-i}
- ϕ_i: The strategy of agent i, which can be either a pure strategy (selecting one action with certainty) or a mixed strategy (a probability distribution over Δ_i).
- ϕ_{-i}: The strategy profile of all agents except agent i.
- Φ: The strategy profile of all agents, $\Phi = (\phi_1, \phi_2, \dots, \phi_n)$
- $E[p_i]$: The expected payoff for agent i, calculated over the set of possible outcomes based on their strategy and the strategies of other agents.

Social interaction is defined as a non-cooperative game characterized by the following elements:

- Agents and Action Sets: $A = \{a_1, a_2, \dots, a_n\}$, denote the set of agents. Each agent has an associated action set Δ_i , representing the possible actions available to the agent.
- Strategy Profile: A strategy profile $\Phi = (\phi_1, \phi_2, \dots, \phi_n)$ consists of individual strategies ϕ_i, where each strategy ϕ_i is a probability distribution over the action set Δ_i. This allows for the representation of both pure and mixed strategies.
- Payoff Function: Each agent i aims to maximize its expected payoff, represented by the function $p_i : \Phi \times \Delta_{-\mathrm{i}} \to \mathbb{R}$ where Δ_{-i} denotes the strategy space of all agents except i. The expected payoff is calculated based on the chosen strategy profile and the outcomes of strategic interactions.

2.3 Equilibrium Definition

In game-theoretic models, an equilibrium represents a state where no agent can improve their payoff by unilaterally changing their strategy, given the strategies of other agents. This section introduces new equilibrium concepts tailored to the current study's focus on mixed strategies and complex agent interactions, building on the foundational definitions provided in the previous work [1].

2.3.1 Nash Equilibrium (NE): A Nash equilibrium is a strategy profile $\Phi = (\phi_1, \phi_2, \dots, \phi_n)$ such that no agent can improve their expected payoff by unilaterally changing their strategy. Formally, for each agent i and any alternative strategy $\phi_i' \neq \phi_i$:

$$E[p_i(\phi_i, \phi_{-i})] \ \geq \ E[p_i(\phi_i', \phi_{-i})]$$

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

2.3.2 Evolutionary Stable Strategy (ESS): An ESS is a refinement of Nash Equilibrium used in evolutionary game theory. It represents a strategy ϕ^* that, if adopted by most of the population, cannot be invaded by any alternative strategy ϕ_i'. For ϕ^* to be evolutionarily stable, it must satisfy:

$$E[p_i(\phi^*, \phi^*)] \geq E[p_i(\phi', \phi^*)]$$

for all $\phi' \neq \phi^*$.

Mixed Strategy Equilibrium (MSE): Mixed strategies allow agents to choose probabilistic combinations of actions, representing complex behavioral patterns more accurately than pure strategies. A Mixed Strategy Equilibrium occurs when agents assign probabilities to multiple actions. For a strategy profile $\Phi = (\phi_1, \phi_2, \ldots, \phi_n)$, each agent's strategy ϕ_i maximizes their expected payoff given the strategies of others. For each agent i:

$$E[p_i(\phi_i, \phi_{-i})] \geq E[p_i(\phi_i', \phi_{-i})]$$

2.3.3 where ϕ_i is a probability distribution over Δ_i.

In the model, agents are assumed to be able to choose their actions according to any rules. These rules are represented for each agent by a set of probability distributions over the agent's decision sets, each of which is called a strategy. This incorporates mixed strategies into the model.

$$\forall d \in \Delta, \exists \phi: \phi \text{ is a random variable with probability distribution over } d: \forall \phi: \phi \in \Phi \Phi_i \coloneqq \text{set of } \phi \text{ corresponding to } d \in D_i, |\Phi_i| = |D_i|$$

2.3.4 Evolutionary Stable Set (ESS Set): An ESS Set is a collection of strategy profiles that remains stable under evolutionary dynamics. It generalizes the ESS concept to polymorphic populations, where multiple strategies coexist in stable proportions. For an ESS Set S, any minor deviation by a subset of the population from $\Phi \in S$ does not result in a shift outside of S.

2.3.5 Convergence Dynamics and Stability Criteria: A strategy profile Φ is stable if small perturbations in the strategy of any agent do not lead to an increase in expected payoff. Formally, for any small change ϵ in strategy ϕ_i:

$$E[p_i(\phi_i, \phi_{-i})] \geq E[p_i(\phi_i + \epsilon, \phi_{-i})]$$

where ϵ represents a small deviation in strategy ϕ_i.

2.3.6 Agents' Optimization Problem

Agents aim to maximize their expected payoffs by tweaking their strategies. Therefore, each agent solves the following optimization problem:

$$\forall a_i \in A: \Phi_i = arg\ arg\ E[p_i \in \Psi(\Phi_i \times \Delta_{-i})] \text{ , given } \Delta_i \text{ is constant.}$$

2.3.7 Stable State

A stable state (equilibrium) in the model is defined as a state where no agent could increase their expected payoff, given that other strategies are stable. At least one such point always exists because, in any game-theoretic model, a Nash equilibrium exists, which this condition holds by definition [2].

$$\forall a \in A: \exists \Phi_a: E[p_a \in \Psi(\Phi_a \times \Delta_{-a})] \geq E[p_a \in \Psi((\Phi_a)' \times \Delta_{-a})]$$

2.3.8 Optimization Dynamic

Agents are assumed to change their strategies by small perturbations with a drift towards strategies yielding higher expected payoff until there would be no possible perturbations yielding higher expected payoff than current strategy yields. In this model, agents always use totally mixed strategies due to the non-zero probability of perturbing their strategy by selecting a previously unused action.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

2.4 Algorithm Outline

Genetic Algorithms (GAs) are a subclass of Evolutionary Algorithms (EAs), which are optimization techniques inspired by Darwinian evolution. In GAs, solutions are represented as individuals within a population. Each individual (or species) encodes a potential solution in the search space and is evaluated using a fitness function—for example, the value of an objective function in maximization tasks. The algorithm evolves the population over successive generations, favoring individuals with higher fitness while maintaining diversity. The key steps of a GA are:

1. Initialization: Generate an initial population uniformly distributed across the search space to ensure broad exploration.

2. Evaluation: Assign each individual a fitness value based on their performance on the target function, differentiating species within the population.

3. Selection: Use methods like the roulette wheel or tournament selection to select a subset of high-fitness individuals to serve as parents for the next generation.

4. Crossover: Produce offspring by combining traits of selected parents. For instance, averaging parent components can create new individuals in real-valued vector representations.

5. Mutation: Introduce random changes to some traits of individuals (including offspring) to explore new areas of the search space and maintain genetic diversity, preventing premature convergence.

6. Replacement: Form the new generation by combining parents and offspring, often maintaining a constant population size. Strategies like elitism may retain the best individuals.

GAs evolve the population toward optimal or near-optimal solutions by iteratively applying these steps. The balance between exploiting high-fitness individuals and exploring new possibilities allows GAs to effectively navigate complex search spaces.

As input, the algorithm expects a game model as outlined in the Model of Social Interaction section. It is important to highlight that for the algorithm to work, it is sufficient to know only the structure of the action set Δ (agents and their decision sets) and an interface to compute the corresponding payoffs; there is no need to know how these payoffs are calculated. As output, the algorithm provides a found equilibrium or a statement indicating failure to find one. In either case, the output should contain a log of convergence dynamics, which includes individuals' encodings in each generation and their fitness values. This algorithm can be initiated multiple times to find multiple equilibria.

5 Results

5.1 Multiplicity of Evolutionary Stable Strategies Definitions

As has been shown by Rykunov D.P. [1], the algorithm converges to ESS. However, the model employs mixed-strategy approaches, which are not widely used in the evolutionary game theory literature and introduce unique properties of system dynamics.

First, it should be noted that there are two definitions of Evolutionarily Stable Strategy (ESS): the classic one by [10] and the stricter one by [9]. These definitions differ slightly and do not always coincide. Due to the lack of prior work on this topic, there is ambiguity regarding which equilibrium concept an algorithm like the one proposed here converges to. An empirical experiment involving a game called "Harm thy Neighbor" is arranged to determine this. This game has two Nash Equilibria in pure strategies: [A; A] and [B; B]. According to the ESS definition by [10], both Nash Equilibria are ESS. In contrast, under the ESS definition by [9], neither of the Nash Equilibria is ESS—there is no ESS in this game.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

		Player 2	
		A	B
Player 1	A	2; 2	1; 2
	B	2; 1	2; 2

Figure 1: HarmThyNeighbour game's payoff matrix

To determine to which definition of an Evolutionarily Stable Strategy (ESS) the algorithm converges, the game is run without termination criteria. The expectation is that if the algorithm converges to the stricter definition of ESS, it will fail to reach a stable state in this setup. Even if it occasionally attains a generation where the best-fitted solution represents one of the Nash Equilibria (NE), it will diverge from it, resulting in the oscillating behavior of the population. Conversely, consistent convergence behavior will be observed if the algorithm converges to the classic definition of ESS. The population would adhere to one of the equilibria, and the entire population—not just the best-fitted solution—would converge to it.

According to the base model defined in [1], the algorithm depends on five parameters: POPSIZE, NPAIRS, NGAMES, DROPOUT_RATE, and MUTATION_MAGNITUDE. For this experiment, these parameters are set based on the convergence results from the prior work [1]:

Algorithm run parameters: popsize: 54, npairs: 20, ngames: 1, dropout_rate: 37, mutation_magnitude: 8

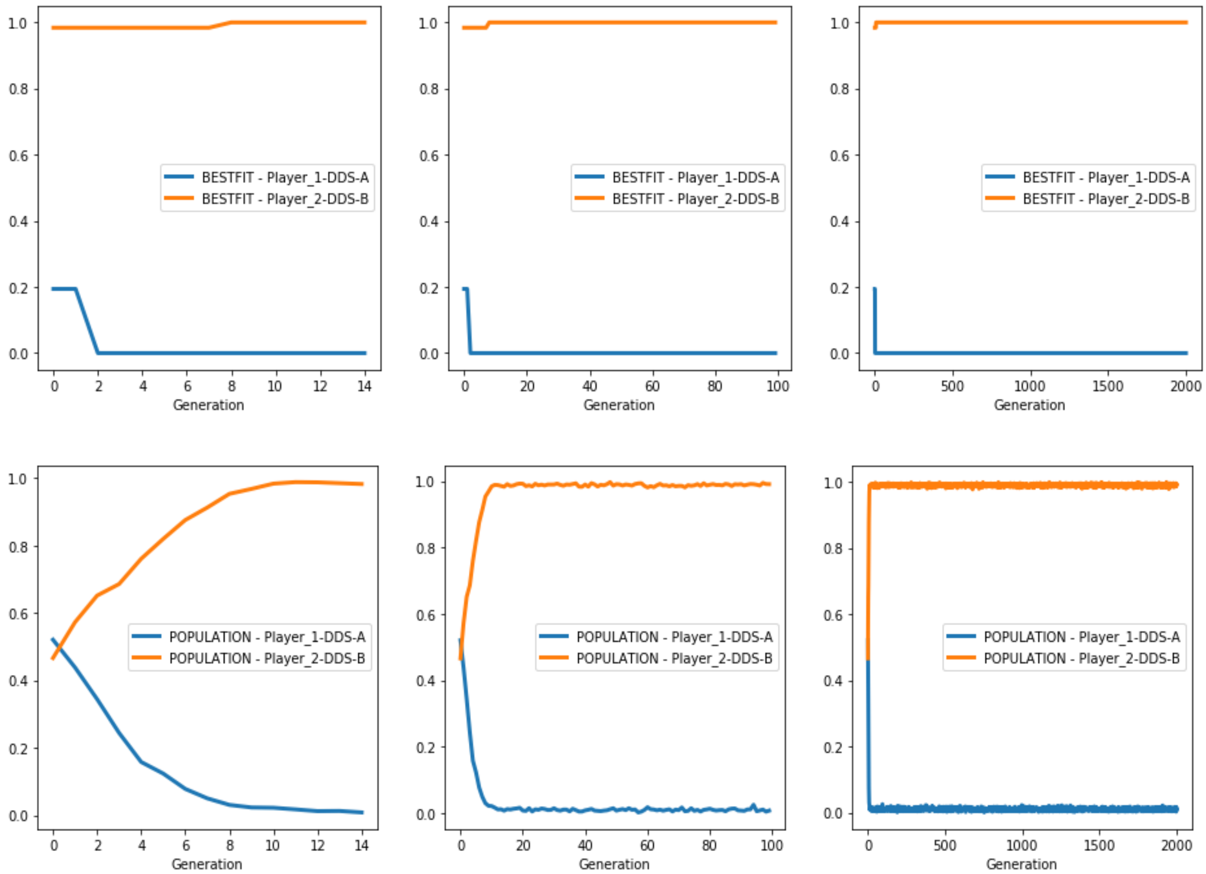

Figure 2: Convergence to an ESS of classic definition

Data obtained through the algorithm's run is presented in Figure 2. Each graph shows the dynamic evolution of equilibrium strategies set for the model by demonstrating two lines: mixed strategy of the first player as a frequency of playing $[A]$ move and

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

mixed strategy of the second player as a frequency of playing [B] move. In the first row, strategies of best-fitted individuals in each generation are depicted, while in the second row, the average strategy of the subpopulations is shown. Graphs differ by column only in the number of generations plotted.

From the graph, it could be readily seen that both best-fitted individuals' strategies and average subpopulations' strategies converge to one of the expected equilibriums $[B; B]$ – best-fitted individuals at 7^{th} generation (never diverge in further generations) and average subpopulations' strategies at approximately 15^{th} generation (fluctuate slightly around equilibrium point for all generations, never diverging significantly). It shows that the algorithm can converge to an ESS of the classic definition of [10].

5.2 Note on Convergence of Mixed Strategies Equilibria

5.2.1 Mixed Evolutionary Stable Strategy

ESS was defined originally for pure-strategist symmetrical game-theoretic models of coevolution in terms of replicator dynamic in differential equations [10] and stood for stable proportions of pure-strategist individuals with specific strategies in population; therefore, it could be treated as a stable state of population. Later in the work of [11], the ESS concept was revisited; this work introduced a stricter definition and presented an analysis of actual individuals' strategies in the context of ESS; also, the distinction between pure-strategist and mixed-strategist models was emphasized. The fundamental definition of an ESS suggests that it represents "an individual phenotype which can be established and maintained predominantly by natural selection in evolutionary competition with alternative types" [11] [12], which is presumed to be the case for both pure-strategist and mixed-strategist models. Pure ESS thus stands for monomorphic populations (populations in which only one pure strategy is presented), and mixed ESS stands for polymorphic populations (populations in which more than one pure strategy – regardless of being included in a mixed strategy or standing by itself – are presented). This means that, as with the case of pure-strategist models, mixed ESS in a mixed-strategist model suggests that a stable set of individual phenotypes must be presented in the population.

One of the main findings on the topic is that mixed ESS could only exist in mixed-strategist models if mixed strategies payoff is a linear combination of payoffs of its components if mixed strategies are sufficiently few in number [13]. Sufficiently few in this context means the space of mixed strategies available to individuals is of larger dimensionality than the space of possible pure strategies. This result (absence of mixed ESS) holds for virtually all mixed-strategist game-theoretic models. It is important because the algorithm uses mixed strategies by design, hence treating all classic game-theoretic problems formulated as Normal Form Games and Games in Extensive Form as mixed-strategist models. On the other hand, this result is a feature of ESS and is relevant only to a particular class of game-theoretic models – linear models. Indeed, the algorithm proposed in this work can also work with non-linear models.

5.2.2 Mixed Evolutionary Stable Set as an Alternative to ES Strategy

As an extension of the ESS concept, the Evolutionary Stable Set was introduced in [9]. ES set concept defines stability for a population as a whole, characterized by a population strategy that is a weighted average of the phenotypical content of the population, e.g., the average of individuals' strategies. In the context of mixed-strategist models, an arbitrarily large number of population states could capture each population strategy. Population state is defined as a particular phenotypical profile of a population, e.g., a set of individuals with specific strategies.

The stability of ES sets appears in the presence of a compact space of population states, from which, once converged, the population would not diverge. In this case, the state of a population converges to an ES set and is expected to drift around the ES set, preserving a stable population strategy.

In the context of mixed-strategist models, it is helpful to introduce a degenerate model, which is the same model treated as a pure-strategist one. It is shown that for a mixed-strategist model, every ESS of the model's degenerate corresponds to an ES set of the mixed-strategist model [13]. Because of the definition of ESS, particularly its uniqueness in its neighborhood, the corresponding ES set would represent just a single population strategy. Therefore, *the suggested algorithm could capture a mixed ESS of a degenerate model through population dynamics*.

5.2.3 Uncorrelated Asymmetry

An uncorrelated asymmetry is a concept introduced in the context of evolutionary game theory for symmetrical games representing coevolutionary dynamics [10]. It could be interpreted as follows: presence or absence of uncorrelated asymmetry determines whether agents in a game-theoretic model know which role they play in it; alternatively, whether agents know which role they have been assigned; or whether their strategies could be dependent on roles they have been assigned. It is essential in symmetrical games. Due to [14] in the presence of information asymmetry (uncorrelated asymmetry), every ESS (if any exists) is a strict Nash Equilibrium [15].

The proposed model introduces uncorrelated asymmetry by design, establishing separate subpopulations for each decision-set in a given model, thereby creating distinct subpopulations for each player. Species evolving in separate subpopulations evolve

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

ISSN 2201-2796

independently from each other. In each fitness evaluation, these species play the same roles. As stated in [14], "A sufficient condition for information asymmetry is satisfied if the two opponents in a contest always have roles which are different from each other."

To tackle game-theoretic models in a setup without uncorrelated asymmetry, a workaround extension for the algorithm was developed. It should be noted that a scenario of an absence of uncorrelated asymmetry is established only for agents with identical decision sets, which could naturally represent a case of agents being unaware of their role in fitness evaluation because there is no difference between them by design. The extension's Central idea is to ensure that subpopulations corresponding to decision sets united by information symmetry are identical before fitness evaluations. In each generation, one of such subpopulations is selected, and other subpopulations of the same information symmetry union are substituted with duplicates of the selected one. All decision sets with information symmetry are represented in this setup by the same subpopulation. So, because of the random nature of the matching algorithm [1] and identical structure of the decision sets of this subpopulation, it is evaluated as with random roles assigned unless there is an asymmetry in payoffs between decision sets (this case is not considered in this paper and suggested for further research).

5.2.4 Empirical Results on Convergence to Mixed ESS in the Absence of Uncorrelated Asymmetry

The algorithm's behavior on the task of locating Mixed ESS is investigated using a classic game from evolutionary game theory studies known as "Hawk Dove." It has two pure strategies Nash Equilibria $[Hawk;\ Dove]$ and $[Dove;\ Hawk]$ and one mixed Nash Equilibria $\left[\left(\frac{V}{C}\right) * Hawk + \left(1-\frac{V}{C}\right) * Dove;\ \left(\frac{V}{C}\right) * Hawk + \left(1-\frac{V}{C}\right) * Dove\right]$, when $V < C$, for a payoff matrix of the game, please refer to Figure 3. In uncorrelated asymmetry, the game has two ESS corresponding to pure NE. Without uncorrelated asymmetry, the game has only one ESS corresponding to mixed NE.

		Player 2	
		Hawk	Dove
Player 1	Hawk	(V-C)/2; (V-C)/2	V; 0
	Dove	0; V	V/2; V/2

Figure 3: HawkDove game's payoff matrix

Two models are constructed based on Hawk Dove game with different payoff matrices, resulting in different proportions of $V \div C$ Thus, there is a different mixed ESS to capture the difference in the convergence of population strategy to Mixed ESS, if any. HawkDoveSkewedHawk model in Figure 4 favors the Hawk strategy in Mixed ESS, and HawkDoveSkewedDove favors the Dove strategy in Mixed ESS.

		Player 2	
		Hawk	Dove
Player 1	Hawk	-2.5; -2.5	15; 0
	Dove	0; 15	7.5; 7.5

Figure 4: HawkDoveSkewedHawk model's payoff matrix

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

		Player 2	
		Hawk	Dove
Player 1	Hawk	-7.5; -7.5	5; 0
	Dove	0; 5	2.5; 2.5

Figure 5: HawkDoveSkewedDove model's payoff matrix

Two algorithm runs are conducted on each model, with and without uncorrelated asymmetry, for 1000 generations. For each run, a graph is constructed to display the best-fitted strategy in the population and the population strategy for each generation.

Algorithm run parameters: popsize: 54, npairs: 20, ngames: 1, dropout_rate: 18, mutation_magnitude: 8

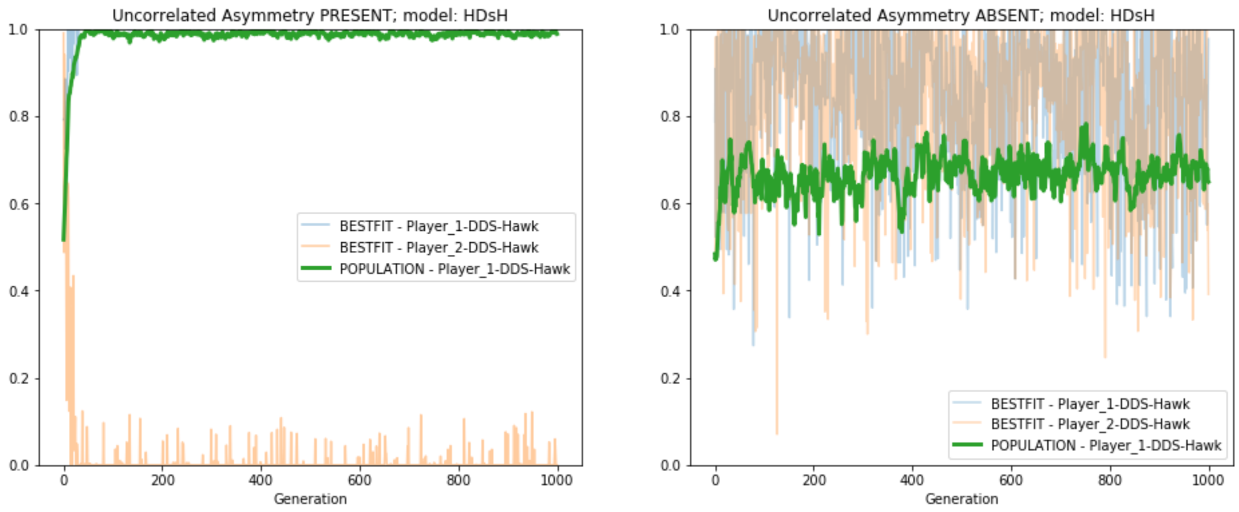

Figure 6: Mixed ESS in HawkDoveSkewedHawk

Algorithm run parameters: popsize: 54, npairs: 20, ngames: 1, dropout_rate: 18, mutation_magnitude: 8

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

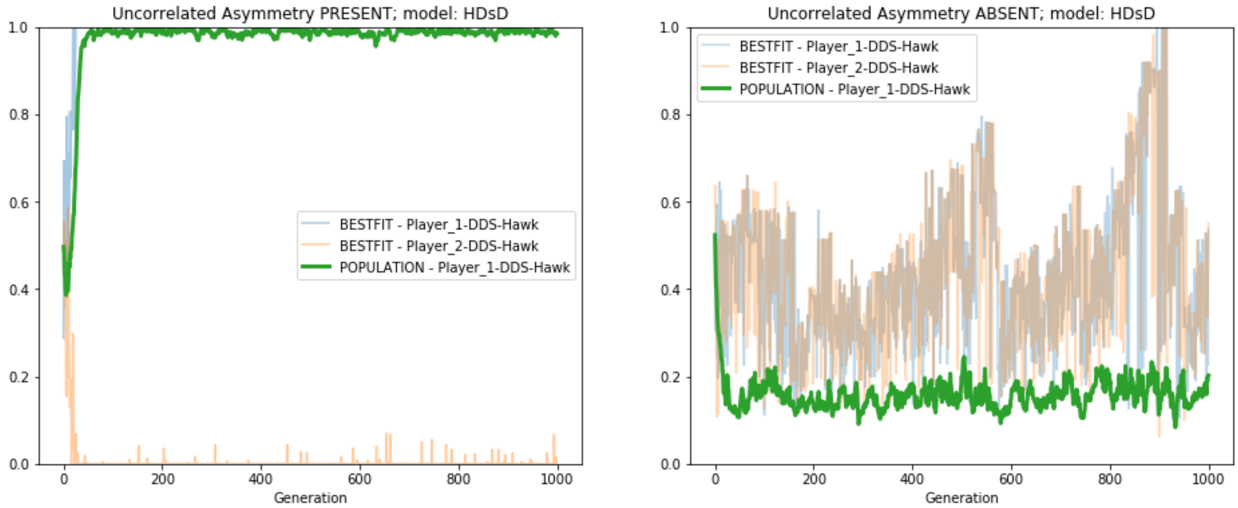

Figure 7: Mixed ESS in HawkDoveSkewedDove

The data obtained from running the algorithm is presented in four graphs: two graphs in Figure 6 (for the model favoring the Hawk strategy in mixed ESS) and two graphs in Figure 7 (for the model favoring the Dove strategy in mixed ESS). Graph sets for both models have an identical structure. Graphs on the right-hand side show the algorithm's output running with modification simulating an absence of uncorrelated asymmetry. Graphs on the left-hand side show the output of the algorithm running with the default setup, thus introducing uncorrelated asymmetry to the model.

Each graph depicts the dynamic evolution of an equilibrium strategy set for the model by three lines. The first one (light blue line) – mixed strategy of the best-fitted specie representing the first player as a frequency of playing $[Hawk]$ move. The second (light orange line) – mixed strategy of the best-fitted specie representing the second player as a frequency of playing $[Hawk]$ move. The third (bold green line) – average mixed population strategy as a frequency of playing $[Hawk]$ move.

As expected, under uncorrelated asymmetry both models converged in terms of individuals' strategies to one of the ESS represented by pure NE, in this case – $[Hawk;\ Dove]$. The same models evaluated in coevolutionary style with a setup of absence of uncorrelated asymmetry have shown non-convergence of individual strategies, which could be seen on graphs Figure 6 and Figure 7 by the highly volatile random walk process of best-fitted strategy. At the same time, on the same graphs, it could be seen clearly, that population strategy represented on the graphs by frequency of $[Hawk]$ pure strategy quickly converges to low volatile random walk process around values approximately equal to theoretical Mixed ESS – 0.75 for HawkDoveSkewedHawk model and 0.25 for HawkDoveSkewedDove. It demonstrates the possibility of exploring mixed ESS by the proposed algorithm and, in general, shows how ES sets could be analyzed by the same means.

5.3 Multiple Equilibria Statistics

The proposed algorithm can locate multiple equilibria in game-theoretic models. Employing an Iterated Local Search (ILS) approach with random initial seeds effectively explores the solution space. The theoretical foundation of this method aligns with propositions from evolutionary game theory, which postulate that player populations converge to Evolutionarily Stable (ES) sets when started in their neighborhood [9]. This justifies treating the ILS with random seeds as asymptotically capable of locating all equilibria of the model. One should develop convergence metrics for the algorithm to utilize the iterated local search method. Used as a termination criterion for the search process, a convergence metric could significantly impact the total amount of computational resources consumed for the search process.

Such convergence metrics might differ from one application of the algorithm to another. The following factors should be considered while developing termination criteria for the algorithm: demanded level of accuracy in the location of equilibria in terms of proportions of false positive and false negative returns and numerical precision of found equilibria position. It is advised to change several termination criteria between meta-heuristic ILS iterations, i.e., roughly locating equilibria states, and then reinitializing the algorithm in the located spots with different parameters and termination criteria to enhance numerical approximation of equilibria. It is also advisable to conduct a supervised preliminary exploration of the problem and the algorithm's performance to accumulate prior knowledge of the problem before initialization of the automated search procedure because it might give insights regarding the problem's structure and guide parameters and termination criteria fine-tuning.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

6 Conclusions and Further Research

This study offers a detailed empirical evaluation of an evolutionary optimization algorithm for locating multiple equilibria in game-theoretic models. Systematic analysis of its performance across various game scenarios has demonstrated the algorithm's capability to converge to Evolutionarily Stable Strategies (ESS) and handle complex mixed-strategy equilibria. The results indicate that the algorithm is robust across various parameter settings, with convergence dynamics significantly influenced by population size, mutation rates, and the number of pairwise interactions. The empirical findings provide valuable insights into the algorithm's operational characteristics, particularly in navigating high-dimensional strategy spaces and identifying stable outcomes in the presence of multiple equilibria. These insights have practical implications for applying game-theoretic models in social sciences and economics, where accurately predicting the behavior of interacting agents is crucial. By leveraging the strengths of evolutionary game theory, this work contributes to a deeper understanding of equilibrium selection and strategic stability in complex systems.

Works Cited

[1] D. P. Rykunov, Evolving Strategies: A Numerical Approach to Game-Theoretic Equilibrium Analysis, 2024.

[2] J. Nash, "Non-Cooperative Games," *The Annals of Mathematics*, pp. 286-296, 1951.

[3] R. D. M. A. M. McKelvey, Computation of Equilibria in Finite Games, Handbook of Computational Economics, vol. 1, 1st ed., H. M. Amman, D. A. Kendrick, and J. Rust, Eds. Amsterdam, Netherlands: Elsevier, pp. 87-142., 1996.

[4] F. B. E. B. K. Mattheos K. Protopapas, "Coevolutionary Genetic Algorithms for Establishing Nash Equilibrium in Symmetric Cournot Games," *Advances in Decision Sciences*, 2010.

[5] T.-L. Y. Wei-Kai Lin, "The Co-Evolvability of Games in Coevolutionary Genetic Algorithms," Taiwan Evolutionary Intelligence Laboratory (TEIL), Taiwan, 2009.

[6] B. v. Stengel, "Computation of Nash equilibria in finite games: introduction to the symposium," *Economic Theory*, 2010.

[7] K. A. D. Jong, Evolutionary Computation, Cambridge, Massachusetts; London, England: The MIT Press, 2006.

[8] M. Preuss, Multimodal Optimization by Means of Evolutionary Algorithms, Münster, Germany: Springer International Publishing Switzerland, 2015.

[9] B. Thomas, "On evolutionarily stable sets," *Journal of Mathematical Biology*, pp. 105-115, 1985.

[10] J. M. Smith, Evolution and the Theory of Games, Cambridge University Press, 1982.

[11] B. Thomas, "Evolutionary stability: States and strategies," *Theoretical Population Biology*, pp. 49-67, 1984.

[12] J. Maynard Smith and G. Price, "The logic of animal conflict," *Nature*, pp. 15-18, 1973.

[13] B. Thomas, "Evolutionarily Stable Sets in Mixed-Strategist Models," *Theoretical Population Biology*, pp. 332-341, 1985.

[14] R. Selten, "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Conflicts," in *Models of Strategic Rationality*, Springer Netherlands, 1988, pp. 67-75.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX

[15] L. Samuelson, Evolutionary Games and Equilibrium Selection, Cambridge, MA: The MIT Press, 1997.

http://dx.doi.org/10.31364/SCIRJ/v12.i10.2024.P1024XX