

Article

Genetic Algorithms for Pareto Optimization in Bayesian Cournot Games Under Incomplete Cost Information

David Carfi ^{1,*} , Alessia Donato ²  and Emanuele Perrone ³¹ Department of Physics (MIFT), University of Messina, 98166 Messina, Italy² Department of Engineering, University of Messina, 98166 Messina, Italy; adonato@unime.it³ Independent Researcher, 00142 Rome, Italy; eperrone.ww@gmail.com

* Correspondence: dcarfi@unime.it

Abstract

This paper develops a practical computational framework for the Bayesian Cournot model with bilateral incomplete cost information, where each player is uncertain about the opponent's marginal cost, drawn from a continuous compact interval $[c_*, c^*]$ with $0 < c_* < c^* < \infty$. The infinite dimensionality of the functional strategy spaces (mappings from types to production quantities) renders analytical closed-form solutions infeasible in this continuous-type setting. To overcome this challenge, we restrict the strategy spaces to finite-dimensional differentiable sub-manifolds—specifically, one-parameter families of oscillatory functions (cosine, sine, and mixed forms). After suitable affine Q -rescaling to map the oscillatory range into the production interval $[0, Q]$, and with parameter ranges satisfying $\alpha, \beta > (\pi/2)/c_*$, these curves ensure near-exhaustivity: the joint production map $(\alpha, \beta) \mapsto (x_\alpha(s), y_\beta(t))$ covers $[0, Q]^2$ densely for every fixed cost pair (s, t) , thereby recovering (up to density and closure) the full ex-post payoff space. We introduce the ex-post payoff mapping $\Phi(s, t, x, y) = (e_s(x, y)(t), f_t(x, y)(s))$, which collects every realizable payoff pair once nature draws the types and players select their strategies. The image of Φ defines the general payoff space of the game, and its non-dominated points constitute the general ex-post Pareto frontier—all efficient realized outcomes across type-strategy realizations, without dependence on private probability measures over types. Using multi-objective genetic algorithms, we numerically approximate this frontier (and selected collusive compromises) within the restricted but representative sub-manifolds. The resulting frontiers are computationally accessible, robust to parameter variations, and validated through hypervolume convergence, sensitivity analysis, and comparisons with NSGA-II, PSO and scalarization methods. The findings are significant because they provide decision-makers in oligopolistic markets (e.g., electric vehicles) with viable, implementable production policies that explore efficient trade-offs under genuine cost uncertainty, without requiring explicit forecasts of the opponent's type distribution—a limitation of traditional expected-utility approaches. By focusing on ex-post efficiency, the method reveals belief-independent compromise solutions that may guide tacit coordination or collusive outcomes in real-world strategic settings.



Academic Editor: Manuel Alberto M. Ferreira

Received: 18 December 2025

Revised: 16 February 2026

Accepted: 20 February 2026

Published: 25 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: duopoly; normal form games; differentiable games; Bayesian models; multi-objective optimization; genetic algorithms; ex-post Pareto frontier; oscillatory strategies; incomplete information; Cournot competition

MSC: 91A10; 91B55; 90C29; 90C59; 91A44; 46E30; 28A33; 42A38

1. Introduction

1.1. Overview

This Introduction sets the stage for this paper by progressively narrowing from the broad context of Bayesian games under incomplete information to the specific challenges of Pareto frontiers in continuous-type Cournot duopolies, and finally to our proposed solution and contributions. It is structured in four logical parts: (1) general definition and motivation of Bayesian games, (2) core research questions and gaps in existing approaches, (3) overview of multi-objective optimization and genetic algorithms as the computational framework, and (4) summary of the paper's dual contributions and structure. While comprehensive to ensure self-containment, the section is focused on motivating the shift to an ex-post, belief-independent Pareto frontier.

1.2. General Definition and Motivation of Bayesian Games

Bayesian games represent a class of strategic interactions characterized by incomplete information, where players must make decisions without full knowledge of certain key parameters affecting the game's outcomes. In these models, rational players attempt to anticipate and optimize their strategies based on expectations about the game's results, but their ability to do so is significantly impaired by uncertainty regarding aspects such as the other players' cost structures, preferences, or capabilities. This uncertainty is typically formalized through the concept of "types," where each player is assigned a type drawn from a predefined set, and this type determines specific attributes like a varying cost function. For instance, in economic settings, a player's type might reflect their marginal production costs, which are private information not observable by opponents [1]. As a result, players form beliefs—often represented as probability distributions—over the possible types of their rivals, and they select strategies that maximize their expected payoffs conditional on these beliefs.

1.3. Core Research Questions and Gaps in Existing Approaches

In the context of Bayesian games, identifying the Pareto frontier poses substantial challenges. The Pareto frontier, or Pareto front, consists of all outcomes where no player can improve their payoff without simultaneously worsening at least one other player's payoff. Unlike in complete-information games, where the payoff space can often be mapped explicitly [2], the incomplete information in Bayesian models introduces complexity due to the probabilistic nature of types and the resulting expected utilities. Computing the Pareto frontier requires evaluating trade-offs across a potentially vast set of strategy profiles, accounting for all possible type realizations and their associated probabilities. Moreover, selecting optimal solutions from this frontier—those that represent fair or efficient compromises—becomes even more arduous, as it demands balancing multiple objectives under uncertainty. This difficulty is exacerbated in infinite-type spaces, such as continuous distributions of costs, where analytical closed-form solutions are rarely feasible. To address these issues, we pose two fundamental research questions that guide our inquiry:

- In what precise manner should Pareto boundaries be conceptualized and analyzed within Bayesian games, particularly when dealing with incomplete information and multi-dimensional strategy spaces?
- How can we identify and evaluate potential compromise solutions on these boundaries without relying on explicit forecasts or assumptions about the probability measures governing the type spaces?

Our proposed solution to these questions leverages multi-objective optimization techniques, specifically employing genetic algorithms (GAs) as a computational framework.

1.4. Multi-Objective Optimization and Genetic Algorithms as the Computational Framework

Multi-objective optimization arises in scenarios where a decision-maker must simultaneously optimize two or more conflicting objective functions. This paradigm, often synonymous with Pareto optimization, acknowledges that real-world problems frequently involve inherent trade-offs: improving one objective may degrade another. For example, in the domain of freight transportation logistics, a company might seek to minimize delivery times and operational costs while maximizing the volume of goods transported [3]. Similarly, in environmental policy applications, objectives could include reducing pollutant emissions from vehicles while enhancing overall fleet performance and efficiency. In game-theoretic contexts, each player acts as an agent striving to maximize their individual profit, utility, or payoff, leading to a natural multi-objective structure where collective optimality requires navigating these conflicts.

It is crucial to recognize that, unlike single-objective optimization, multi-objective problems do not typically yield a unique optimal solution that excels in all dimensions. Instead, they produce a set of Pareto-optimal solutions—points where no further improvement in one objective is possible without a detriment to at least one other. This set, known as the Pareto front, can be finite or infinite in cardinality, and all points on it are deemed non-dominated, meaning they represent equally viable compromises depending on the decision-maker's priorities. In practice, the challenge lies in approximating this front when exact computation is intractable, as is often the case in complex systems like Bayesian games.

In Bayesian games specifically, deriving the Pareto-optimal solutions in closed form is notoriously difficult due to the interplay of uncertainty, high-dimensional strategy spaces, and non-linear payoff functions. Consequently, problem-solving in this domain is frequently reframed as the task of computing or approximating a representative subset of these Pareto-optimal solutions, providing decision-makers with a practical array of choices [4,5].

Evolutionary algorithms have emerged as a robust class of methods for tackling such multi-objective optimization problems by mimicking natural evolutionary processes to explore and refine solution spaces. Among these, popular variants include particle swarm optimization (which simulates social behaviors like bird flocking), simulated annealing (inspired by metallurgical cooling processes), and genetic algorithms (drawing from principles of natural selection and genetics) [6,7]. Genetic algorithms, in particular, operate by maintaining a population of candidate solutions, iteratively applying operators such as selection (favoring fitter individuals), crossover (combining traits from parents to create offspring), and mutation (introducing random variations) to evolve toward better approximations of the optimum.

The primary strength of evolutionary algorithms lies in their ability to generate diverse sets of solutions in a single run, facilitating the approximation of the entire Pareto front rather than isolated points. This population-based approach allows for broad exploration of the search space, making it well-suited to non-convex, multimodal problems where traditional gradient-based methods might falter. However, a notable limitation is that these algorithms do not guarantee true Pareto optimality; the solutions produced are only assured to be non-dominated within the generated set, meaning one cannot definitively claim they are globally optimal without exhaustive verification. To mitigate this, practitioners often employ large population sizes, which increase computational diversity and the likelihood of capturing high-quality approximations, though at the expense of greater resource demands.

By integrating these concepts, our work advances the application of genetic algorithms to Bayesian Cournot games, enabling the practical determination of Pareto boundaries even in challenging infinite-dimensional settings.

1.5. Our Contribution

After proposing a method to select on the Pareto frontier possible strategies in the finite case (see Appendix B), we propose, in this paper, a practical implementation of the Cournot Bayesian model in which both players do not know the marginal cost of the other. This study brings considerable difficulties because of the infinite dimensionality of the strategy spaces. Therefore, we propose a simplified approach for the decision-maker by restricting the strategy spaces to finite-dimensional differentiable sub-manifolds of the initial strategy spaces. In particular, in this paper, we propose curves of Bayesian strategies. The proposed method, even if it cannot be considered exhaustive or complete, allows the decision-maker to take optimal decisions via genetic algorithms while choosing the sub-manifolds as viable strategies. Specifically, we concentrate our attention on the determination of Pareto boundaries and some collusive agreements. Moreover, we emphasize that the necessity to consider such restriction of the strategy space is due to the assumption that the player cannot or does not desire to forecast the probability measures on type spaces.

2. Literature Review

In the literature, Bayesian games have recently been used in industrial organization contexts; for instance, to secure critical infrastructure when the adversary's characteristics are uncertain in order to identify optimal defense strategies [8] or to find the optimal pricing that maximizes the utilities in problems of electricity transaction [9].

Bayesian games and evolutionary algorithms are used together in the literature in different contexts. For instance, ref. [10] applies Bayesian optimization and a Genetic Algorithm in a problem of chaotic circuits in order to identify various chaogate designs. Ref. [11] faced the problem of solving complex Bayesian games by applying natural evolution strategies.

Generally, studies of Bayesian games focus on the existence of Nash equilibria; for instance, in the case of infinitely many players [12], or with the existence of approximate equilibria for the class of multicriteria Bayesian games [13]. M. Binois considered the uncertainty quantification on Pareto fronts and high-dimensional strategies in Bayesian optimization [14]. Ref. [15] calculated the Pareto front approximation of Bayesian optimization problems.

In related fields, multi-objective Bayesian optimization (MBO) has advanced methods for approximating Pareto frontiers under uncertainty. For instance, Suzuki et al. [16] introduce Pareto-frontier entropy search (PFES), extending entropy-based BO to incorporate the entropy of the Pareto frontier, capturing trade-offs among objectives with exact analytical computations. This addresses limitations in prior entropy-based MBO approaches, which often ignore inter-objective dependencies or rely on approximations, and it extends to decoupled settings where objectives are observed separately. Empirical results on benchmarks demonstrate PFES's superiority in exploring non-dominated solutions. However, PFES and similar MBO frameworks (e.g., [17,18]) primarily assume finite-dimensional input spaces and focus on non-strategic optimization without considering game-theoretic interactions or incomplete information over player types. In Bayesian games, the strategy spaces are often infinite-dimensional (functional mappings from types to actions), and optimization must account for ex-post realized payoffs across all type realizations, independent of private beliefs. Prior work in game-theoretic BO [19] explores model-based approximations but does not address continuous-type uncertainty or belief-independent frontiers. Our approach fills this gap by restricting to differentiable sub-manifolds (ensuring exhaustivity via joint surjectivity) and computing the general ex-post Pareto frontier, providing a computational bridge between MBO techniques and strategic Bayesian models.

Many authors use Bayesian games in the classic models of oligopoly with finitely many possible marginal costs (see [20–22]). The case of continuous uncertainty on marginal costs is studied only in [23] (to the best of our knowledge), for what concerns the Nash equilibria. We do not know of any authors who have studied the Pareto frontiers of Bayesian Cournot games in the continuous case. Here, we propose a computational method capable of determining Pareto frontiers using evolutionary algorithms when the decision-makers opt to move inside viable strategy sub-manifolds instead of the entire infinite-dimensional Lebesgue spaces $\mathcal{L}_1$.

3. Methods

In this paper, we analyze the Bayesian game associated with the deterministic game studied in Appendix B and introduced in [23]. A Bayesian game is a normal-form game in which it is assumed that players have incomplete information about strategies and payoffs [24–26]. In our case, the strategies are functional strategies, defined on the respective type-spaces of the players. Therefore, we go from finite-dimensional strategy spaces to infinite-dimensional ones.

Our aim is to find the whole Pareto frontier. However, the complexity of the proposed model with payoff functions of our type does not allow us to find and draw the points of the Pareto frontier in a closed form. For this aim, genetic algorithms are useful because they find and draw the optimal Pareto points. Furthermore, at each iteration, thanks to the ability of adaptation, genetic algorithms keep only the best solutions and transform them with genetic operators (mutation, reproduction and crossover) in order to improve them further [27].

In order to solve the multi-objective optimization problem defined in our Bayesian game theory model, we apply Genetic algorithms (GAs). GAs are very useful when the problem is not easy to solve in a closed form [28] and the strategy space of the players is defined by non-linear functions. In our case, the strategies themselves are infinite-dimensional vectors (that is, non-linear smooth functions), and we need to consider infinitely many strategies and their transformations in order to construct the Pareto boundaries.

We show in Appendix B how the model can be solved in closed form in the deterministic case, and, in the main text, how it can be solved in the case of parametric oscillatory strategies. Therefore, the novelty lies in the description of the Bayesian game model and in suggesting the solution to the decision-maker with an a posteriori method.

4. Theoretical Background

The Cournot–Bayesian Model: Uncertainty regarding the Costs of Both Players

We refer to [23] for the determination of the players' payoff functions in the case of uncertainty regarding the costs of both players.

The type set of the first player and of the second player is respectively the marginal cost interval

$$S := [c_*, c^*]$$

and

$$T := [c_*, c^*],$$

with $0 < c_* < c^* < \infty$.

The payoff strategies are as follows.

The first player's s^{th} type adopts the strategy set:

$$E_s := [0, Q],$$

so that the first player's total strategy profile space is the Cartesian product

$$\times_{s \in S} E_s = [0, Q]^S,$$

the set of all functions from S into $[0, Q]$. In other terms, any strategy of the first player is an infinite profile strategy

$$x : S \rightarrow [0, Q],$$

specifying, for any type $s \in S$, a certain production $x(s)$ in $[0, Q]$.

Assumption 1. *As a reasonable strategy set of the first player, we restrict our attention to the convex strategy sub-space:*

$$\mathbb{E} := \{x \in \mathcal{L}_p^1(S) : x(s) \in [0, Q]\}, \quad (1)$$

of all real bounded p -integrable functions on S with values in $[0, Q]$, where

$$p : C^0(S) \rightarrow \mathbb{R}$$

is a probability measure (private information of the first player, determined by the nature) of the compact type space S of the first player.

The second player's t^{th} type adopts the strategy set:

$$F_t := [0, Q],$$

so that the second player's total strategy profile space is the Cartesian product

$$\times_{t \in T} F_t = [0, Q]^T,$$

the set of all functions from T into $[0, Q]$. In other terms, any strategy of the second player is an infinite profile strategy

$$y : T \rightarrow [0, Q],$$

specifying, for any type $t \in T$, a certain production $y(t)$ in $[0, Q]$.

Assumption 2. *As a reasonable strategy set of the second player, we restrict our attention to the convex strategy sub-space:*

$$\mathbb{F} := \{y \in \mathcal{L}_q^1(T) : y(t) \in [0, Q]\}, \quad (2)$$

of all real bounded q -integrable functions on T with values in $[0, Q]$, where

$$q : C^0(T) \rightarrow \mathbb{R}$$

is a probability measure (private information of the second player, determined by the nature) of the compact type space T of the second player.

Note that the space $\mathcal{L}_q^1(T)$ is a space of functions and not a space of equivalence classes of functions, so two functions can be different even if they are equal almost everywhere with respect to the measure q .

We indicate the first player's type s by the index s and its stochastic payoff function by e_s . The payoff function

$$e_s : \mathbb{E} \times \mathbb{F} \rightarrow \mathcal{L}_q^1(T),$$

is defined, for every $(x, y) \in \mathbb{E} \times \mathbb{F}$, by:

$$e_s(x, y) = x(s)(Q - x(s) - y) - sx(s). \quad (3)$$

That payoff is the random variable $e_s(x, y) : T \rightarrow \mathbb{R}$ defined by

$$e_s(x, y)(t) = x(s)(Q - x(s) - y(t)) - sx(s), \quad (4)$$

for every type $t \in T$ of the second player.

The payoff function of the second player's t^{th} type

$$f_t : \mathbb{E} \times \mathbb{F} \rightarrow \mathcal{L}_p^1(S)$$

is defined, for every $(x, y) \in \mathbb{E} \times \mathbb{F}$, by:

$$f_t(x, y) = y(t)(Q - x - y(t)) - ty(t). \quad (5)$$

5. Results

This section presents the computational core of this paper. It progresses logically from the genetic algorithm methodology and strategy restrictions (Section 5.1), through concrete numerical examples and frontier visualizations (Section 5.1.3), to the theoretical justification of why the restrictions work (ex-post mapping and exhaustivity in Section 5.2), and finally to rigorous validation and robustness checks (Section 5.3). This structure allows readers to follow the path from implementation to results to reliability, with figures and tables providing visual and quantitative support at each step.

5.1. Study of the Bayesian Cournot's Normal Form Game by Genetic Algorithms

The payoff function that we shall examine is defined by Formulas (3) and (5):

$$(e_s, f_t) : \mathbb{E} \times \mathbb{F} \rightarrow \mathcal{L}_q^1(T) \times \mathcal{L}_p^1(S)$$

for every pair of types $(s, t) \in S \times T$, where

$$(e_s, f_t)(x, y) = (x(s)(Q - x(s) - y) - sx(s), y(t)(Q - x - y(t)) - ty(t)) \quad (6)$$

for every $(x, y) \in \mathbb{E} \times \mathbb{F}$.

5.1.1. Study Case Functions

Our aim is to construct the Pareto boundaries when restricting the strategy spaces to some family of functions. In other terms, we do not consider the entire space of integrable functions with values in $[0, Q]$, but we consider only some curves of possible strategy functions. Specifically, we shall examine some curves of oscillatory strategy functions and other types, those indicated in Table 1.

Table 1. Some exemplary curves.

	$x_\alpha(s)$	$y_\beta(t)$
(1)	$\cos(\alpha s)$	$\cos(\beta t)$
(2)	$\sin(\alpha s)$	$\sin(\beta t)$
(3)	$\cos(\alpha s)$	$\sin(\beta t)$
(4)	$\exp(-\alpha^2 s^2)$	$\exp(-\beta^2 t^2)$
(5)	$\alpha^2 s^2$	$\beta^2 t^2$

In other terms, we consider two curves of strategy functions, those defined by

$$x : A \rightarrow \mathcal{L}_p^1(S) : \alpha \mapsto x_\alpha$$

and

$$y : B \rightarrow \mathcal{L}_q^1(T) : \beta \mapsto y_\beta,$$

where A and B are two convenient compact intervals of the real line.

We desire to underline that a decision-maker very hardly could use the entire space $\mathcal{L}^1$ in order to create a production strategy, in practice. Therefore, when the decision-maker does not know its probability measure and he prefers to consider all possible scenarios, he should somehow restrict its strategy space to a more handable strategy subset. We suggest to use curves of continuous functions that are universally measurable; that is, they are measurable with respect to any probability measure.

Moreover, we are suggesting the first possible step in restricting the entire strategy space to sub-manifolds of finite dimension.

5.1.2. The Algorithm Method

The steps of the proposed genetic algorithm procedure are the following.

- There are assigned ranges of variation for the strategies s and t and for the two parameters α and β .
- The algorithm creates an initial population. Each individual of the population represents a possible solution to the problem; therefore, each individual is associated with a set of four variables (s , t , α and β) with values randomly selected within the variation ranges previously assigned.
- It is a problem of multi-objective optimization, so the payoff of the two players (the two objective functions that has to be minimized) is calculated for all the individuals in the population.
- The Pareto front plots the function values for all non-inferior solutions. The graph shows, in the payoff space, the points of intersection of the two objective functions (the payoff functions) for all the individuals, eliminating points with better earnings from time to time. At the end of the first generation, the points that belongs to the minimal Pareto boundary are shown; that is, all points that cannot be worsened.
- The algorithm stops if certain stopping criteria are reached (for instance, the multi-objective function value is less than a specified tolerance); otherwise, the algorithm passes to the next generation by manipulating the genes of the individuals using genetic transformation operators (reproduction, cross-over, and mutation).
- The multi-objective function value of the individuals of the new generation is then evaluated, and the algorithm stops if a stop criterion is reached, returning the values of s , t , α and β ; otherwise, it continues with a new generation, and so on.

5.1.3. Numerical Example

Figures 1 and 2 represent the Pareto frontier, in case of strategy functions of cosine type, when the parameters and the types vary in certain fixed intervals. Specifically, the types s and t belong to the interval $[0, 0.5]$, while the parameters α and β belong to two intervals, $[1, 2]$ and $[1, 5]$. The parameters of the GA are chosen in this example as follows: population size, 200; number of generations, 400; probability of crossover, 0.8; and probability of mutation, 0.1.

From the optimization process, we obtain some points constituting the minimal Pareto frontier: a set of points with coordinates $(e_s(x, y)(t), f_t(x, y)(s))$ which cannot be improved with respect to the inverse usual order of the plane. Moreover, we also obtain (via retro-

imaging) the strategic quadruples (s, t, α, β) determining the Pareto minimal payoffs. We show in Table 2 the results in the case $1 < (\alpha, \beta) < 5$.




Figure 1. $\alpha, \beta \in [1, 2]$.




Figure 2. $\alpha, \beta \in [1, 5]$.


Table 2. Results of Analysis (1).

s	t	α	β	$e_s(t)$	$f_t(s)$	s	t	α	β	$e_s(t)$	$f_t(s)$
0.01	0.39	3.33	4.70	−1.24	0.22	0.46	0.01	4.80	2.52	0.70	−1.60
0.28	0.21	4.51	3.86	−0.21	−0.57	0.05	0.31	3.16	4.46	−0.77	−0.10
0.01	0.44	3.08	4.38	−1.35	0.33	0.03	0.38	3.50	4.50	−1.11	0.10
0.01	0.43	2.88	4.45	−1.31	0.29	0.02	0.47	1.91	4.71	−1.58	0.67
0.00	0.46	3.23	4.52	−1.47	0.48	0.22	0.29	4.38	4.52	−0.54	−0.24
0.50	0.01	4.99	2.45	1.04	−1.79	0.07	0.33	3.49	4.49	−0.83	−0.06
0.04	0.38	3.05	4.56	−1.13	0.14	0.02	0.49	2.00	4.67	−1.62	0.74
0.02	0.49	1.99	4.49	−1.56	0.63	0.49	0.01	4.94	2.74	0.93	−1.73
0.44	0.01	4.95	3.32	0.65	−1.56	0.39	0.01	4.83	3.33	0.31	−1.32
0.50	0.00	4.99	2.46	1.04	−1.79	0.27	0.15	4.80	4.58	−0.18	−0.63

Table 2. Cont.

s	t	α	β	$e_s(t)$	$f_t(s)$	s	t	α	β	$e_s(t)$	$f_t(s)$
0.01	0.48	3.06	4.48	−1.53	0.58	0.41	0.01	4.73	3.76	0.35	−1.35
0.33	0.02	4.95	2.94	0.05	−1.05	0.33	0.01	4.66	3.34	−0.03	−0.95
0.48	0.01	4.70	2.52	0.76	−1.64	0.01	0.49	1.97	4.97	−1.75	0.95
0.48	0.01	4.97	2.52	0.88	−1.71	0.01	0.43	2.92	4.39	−1.29	0.26
0.34	0.01	4.89	2.66	0.08	−1.09	0.01	0.50	1.27	5.00	−1.78	1.02
0.24	0.24	4.25	4.51	−0.40	−0.36	0.03	0.44	1.64	4.59	−1.41	0.43
0.49	0.01	4.94	2.50	0.97	−1.76	0.02	0.39	3.16	4.54	−1.16	0.15
0.28	0.23	3.71	4.61	−0.37	−0.37	0.38	0.01	4.66	3.45	0.16	−1.19
0.02	0.40	3.01	4.67	−1.25	0.25	0.02	0.49	2.61	4.79	−1.69	0.86
0.40	0.01	4.69	2.71	0.26	−1.28	0.23	0.30	3.82	4.64	−0.59	−0.17
0.23	0.29	4.49	3.89	−0.43	−0.33	0.41	0.01	4.94	2.94	0.43	−1.41
0.04	0.32	3.24	4.65	−0.89	−0.05	0.01	0.48	2.98	4.80	−1.65	0.79
0.02	0.33	3.24	4.63	−0.94	−0.02	0.45	0.03	4.63	2.83	0.52	−1.47
0.27	0.21	4.80	4.58	−0.24	−0.54	0.03	0.45	3.40	4.54	−1.44	0.48
0.01	0.50	2.23	4.85	−1.74	0.93	0.33	0.02	4.84	3.25	0.01	−1.00
0.01	0.50	2.89	4.80	−1.72	0.90	0.02	0.34	3.63	4.70	−1.00	0.02
0.26	0.22	4.27	4.32	−0.31	−0.44	0.09	0.30	3.27	4.46	−0.69	−0.12
0.22	0.28	4.73	4.38	−0.48	−0.29	0.29	0.11	4.71	4.30	−0.12	−0.71
0.09	0.32	3.14	4.49	−0.79	−0.07	0.38	0.04	4.61	2.85	0.16	−1.15
0.01	0.50	1.24	5.00	−1.79	1.04	0.38	0.01	4.81	2.84	0.21	−1.23
0.01	0.45	2.08	4.42	−1.40	0.38	0.01	0.49	1.87	4.75	−1.67	0.82
0.27	0.25	4.43	3.75	−0.28	−0.47	0.31	0.18	4.53	4.28	−0.14	−0.66
0.04	0.40	2.76	4.48	−1.20	0.20	0.45	0.01	4.66	3.34	0.55	−1.51
0.30	0.10	4.68	3.46	−0.09	−0.75	0.46	0.01	4.67	2.57	0.62	−1.55
0.27	0.22	4.72	4.59	−0.28	−0.49	0.37	0.01	4.94	2.50	0.21	−1.24

In Figures 3 and 4, we present the determination of the Pareto frontier using the Non-Sorting Genetic Algorithm (NSGA-II) method [6] and the Multi-Objective Particle Swarm Optimization (MOPSO) method [29]. All the circles in figures 3 and 4 represent payoff points approximating the minimal solutions, therefore they are more concentrated about the Pareto boundary and the central part of the Pareto boundary. Moreover, those darker than the others are minimal points. The Pareto frontier seems similar to that shown in Figure 2 obtained using the multi-objective genetic algorithm optimization method.

Figure 3. $\alpha, \beta \in [1, 5], s, t \in [0, 0.5]$.




Figure 4. $\alpha, \beta \in [1, 5], s, t \in [0, 0.5]$.

We propose, as possible solution, the point $K' = (0.21, -1.24)$ (*minimum collective loss solution*), which is one point characterized by the minimum collective loss (the minimum of the function $e + f$) upon the Pareto boundary (see Figure 5). The corresponding K in the strategy space is

$$K = (0.37, 0.01, 4.94, 2.50).$$




Figure 5. $\alpha, \beta \in [1, 5], s, t \in [0, 0.5]$.

5.1.4. Remarks

From the analysis of the above example, we can easily see that the range of the parameters α and β has a significant influence on the shape of the Pareto frontier. Moreover, we should also take into account the incidence of the variation of other algorithm parameters (e.g., mutation or crossover). Indeed, the genetic algorithm (GA) implementation has been here presented in an applied manner, focusing on its practical utility for approximating the Pareto frontier in the restricted sub-manifolds. However, it is necessary a deeper analysis of convergence, sensitivity, and/or comparisons with other optimization methods. We addresses these points through a more solid theoretical foundation (Section 5.2) and enhanced validation (Section 5.3).

5.2. Ex-Post Payoff Mapping and Exhaustivity of Restricted Oscillatory Strategies

To precisely identify the object approximated by our multi-objective genetic algorithm, we introduce the *ex-post payoff mapping*:

$$\Phi : S \times T \times \mathbb{E} \times \mathbb{F} \rightarrow \mathbb{R}^2, \quad (7)$$

defined by

$$\Phi(s, t, x, y) = (e_s(x, y)(t), f_t(x, y)(s)), \quad (8)$$

where the components are the realized payoffs to player 1 and player 2, respectively, when player 1's marginal cost (type) is s , player 2's is t , and the players select production functions $x \in \mathbb{E}$ and $y \in \mathbb{F}$.

This mapping aggregates *every possible realized payoff pair* occurring after nature draws the types and the players choose their strategies. The *payoff space* of the Bayesian Cournot game is its image:

$$\text{Image}(\Phi) = \{ \Phi(s, t, x, y) \in \mathbb{R}^2 \mid (s, t) \in S \times T, (x, y) \in \mathbb{E} \times \mathbb{F} \}. \quad (9)$$

The *general Pareto frontier* consists of the non-dominated points in $\text{Image}(\Phi)$: a payoff pair (u_1, u_2) belongs to the frontier if no other realizable pair (u'_1, u'_2) satisfies $u'_1 \geq u_1$, $u'_2 \geq u_2$ with at least one strict inequality.

To facilitate analysis and avoid degenerate behaviors near zero cost, we adopt bounded type intervals away from zero:

$$S = T = [c_*, c^*], \quad 0 < c_* < c^* < \infty. \quad (10)$$

This realistic choice reflects strictly positive marginal costs and keeps the compact sets well-behaved for integration and oscillation coverage.

The payoff space admits the decomposition

$$\text{Image}(\Phi) = \bigcup_{(s,t) \in S \times T} P(s, t), \quad (11)$$

where $P(s, t)$ denotes the deterministic Cournot payoff set for fixed costs (s, t) : the image of the profit functions over all production pairs $(q_1, q_2) \in [0, Q] \times [0, Q]$.

For any fixed (x, y) , the sectional map $(s, t) \mapsto \Phi(s, t, x, y)$ covers only the profits arising from the joint production rectangle:

$$R(x, y) = \{ (x(s), y(t)) \mid s \in S, t \in T \} \subseteq [0, Q] \times [0, Q]. \quad (12)$$

Full recovery of $\text{Image}(\Phi)$ therefore requires that the union over chosen strategies densely (or exactly) covers $[0, Q]^2$ in productions for every (s, t) .

Theorem 1 (Joint surjectivity per type pair). *A sufficient condition for the restricted parametric families $\{x_\alpha \mid \alpha \in A\} \subseteq \mathbb{E}$ and $\{y_\beta \mid \beta \in B\} \subseteq \mathbb{F}$ to generate the entire payoff space (up to closure/density) is the following: for every fixed $(s, t) \in S \times T$, the map*

$$A \times B \rightarrow [0, Q] \times [0, Q], \quad (\alpha, \beta) \mapsto (x_\alpha(s), y_\beta(t)) \quad (13)$$

is surjective.

Proof. Given any (s, t) and any target productions $q_1, q_2 \in [0, Q]$, surjectivity ensures there exist $\alpha \in A$, $\beta \in B$ such that $x_\alpha(s) = q_1$ and $y_\beta(t) = q_2$. Thus, the sectional payoffs over

(α, β) cover the full deterministic set $P(s, t)$ for that (s, t) . Union over all (s, t) then yields $\text{Image}(\Phi)$. $\square$

Our oscillatory families (Table 1)—cosine, sine, and mixed—satisfy this condition (or achieve near-surjectivity) under Q -rescaled normalization and appropriate parameter ranges. Consider standard affine transformations that map the oscillatory range $[-1, 1]$ to $[0, Q]$; for example,

$$x_\alpha(s) = \frac{Q}{2}(1 + \cos(\alpha s)) \quad \text{or} \quad x_\alpha(s) = \frac{Q}{2}(1 + \sin(\alpha s)), \quad (14)$$

and analogously for $y_\beta(t)$. Both $\cos$ and $\sin$ restricted to arguments in $[0, \pi/2]$ are *bijective* onto $[0, 1]$:

- $\cos : [0, \pi/2] \rightarrow [0, 1]$ is continuous, strictly decreasing (from 1 to 0), hence bijective (inverse: $\arccos$).
- $\sin : [0, \pi/2] \rightarrow [0, 1]$ is continuous, strictly increasing (from 0 to 1), hence bijective (inverse: $\arcsin$).

For fixed $s \geq c_*$, the argument range is $\alpha s \in [\alpha c_*, \alpha c^*]$. To ensure the effective argument covers at least one full monotonic interval $[0, \pi/2]$ (guaranteeing the normalized function reaches from 0 to Q or Q to 0), it suffices that

$$\alpha c_* \geq \pi/2 \quad \Leftrightarrow \quad \alpha \geq \frac{\pi/2}{c_*}. \quad (15)$$

The same holds for $\beta \geq (\pi/2)/c_*$. In our numerical cases ($\alpha, \beta \in [1, 2]$ or $[1, 5]$), choosing c_* small but positive (e.g., $c_* \approx 0.1$ – 0.2) yields thresholds around 7.85–15.7, which are exceeded in the upper parts of the intervals $[1, 5]$. Larger α, β allow multiple oscillations or wider coverage, ensuring dense filling of $[0, Q]$ for each s (and symmetrically for t). The joint map $(\alpha, \beta) \mapsto (x_\alpha(s), y_\beta(t))$ thus covers $[0, Q]^2$ densely—or exactly in the limit—for every (s, t) .

This exhaustivity property justifies restricting to low-dimensional (one-parameter) curves: they remain viable and implementable while generating a representative (often nearly complete) payoff space and Pareto frontier under continuous bilateral cost uncertainty. Moreover, by focusing on ex-post realized outcomes across all type realizations, the approach avoids explicit dependence on the private probability measures p and q , aligning with the decision-maker's preference for robust strategies without forecasting beliefs over the opponent's cost distribution.

5.3. Validation and Robustness Analysis

To address potential concerns regarding the genetic algorithm's (GA) implementation, we provide convergence analysis, parameter sensitivity tests, comparisons with standard multi-objective methods, and validation against reference solutions like Bayesian Nash equilibria (BNE). All tests use the oscillatory restrictions (cosine family, unless specified) with $Q = 1$, $[c_*, c^*] = [0.1, 0.5]$ and the following GA parameters: population, 100; generations, 100; crossover, 0.8; mutation, 0.05.

Convergence is monitored via the hypervolume indicator (HV) [30], measuring the volume dominated by the approximated Pareto front relative to a reference point (e.g., nadir $(0, 0)$). HV stabilizes after ≈ 40 generations (plateau at $HV \approx 0.85$), indicating reliable approximation without early stopping bias.

Parameter sensitivity is assessed by varying mutation rate (0.01–0.1) and parameter ranges $\alpha, \beta \in [1, 5]$ vs. $[1, 10]$. Table 3 reports HV deviations $< 5\%$ for mutation and $< 3\%$ for

ranges, confirming robustness under exhaustivity (joint surjectivity ensures dense coverage for $\alpha, \beta > (\pi/2)/c_* \approx 15.7$, exceeded in upper intervals).

Comparisons with standards include Non-Sorting Genetic Algorithm [6] (NSGA-II, implemented via pygmo library), Particle Swarm Optimization (PSO, via pygmo library) and scalarization (weighted sums, $\lambda \in [0, 1]$). On a discretized benchmark (10 types), our GA achieves HV 0.88 vs. NSGA-II 0.86 (similar, as our GA uses analogous sorting), PSO 0.80 (converges faster initially but plateaus lower due to less diversity on non-convex frontiers) and scalarization 0.75 (poorer on non-convex parts). Scalarization fails to fill concavities, while population-based methods like ours excel.

For BNE validation, we discretize to finite types (uniform grid on $[c_*, c^*]$) and compute BNE analytically (affine strategies [31]). BNE points lie inside our ex-post frontier (e.g., dominated by 10–20% in payoffs), as expected: BNE optimizes expectations under beliefs, while our frontier captures all efficient ex-post realizations, independent of p/q .

These analyses validate the GA's effectiveness for our ex-post objective, distinguishing it from belief-dependent references like BNE.

Table 3. Parameter sensitivity (average over 10 runs).

Parameter	Variation	HV Deviation (%)
Mutation rate	0.01 vs. 0.1	4.2
α, β range	[1, 5] vs. [1, 10]	2.8

5.4. Pareto Boundary with Different Oscillatory Functions: A Comparison

We analyze the Pareto frontiers of the Bayesian Cournot game using all the different types of functions (1)–(5) proposed in Table 1 (see Figures 6–20).

Moreover, we distinguish the case of costs s and t fixed equal; s and t varying in the interval 0.15–0.35; s and t varying in the interval 0.2–0.4. The parameters α and β vary in $[0, \pi/(2c_*)]$, as previously determined, in the cosine, sine and mixed cases. In the exponential case, we use the interval $[0, 100]$ for α and β to approximate the interval $[0, \infty]$. In the polynomial case, α and β vary in $[0, 1/c_*]$.




Figure 6. Analysis (1), $s, t \in [0.2, 0.4]$.




Figure 7. Analysis (1), $s, t \in [0.15, 0.35]$.




Figure 8. Analysis (1), $s, t = 0.25$.




Figure 9. Analysis (2), $s, t \in [0.2, 0.4]$.




Figure 10. Analysis (2), $s, t \in [0.15, 0.35]$.




Figure 11. Analysis (2), $s, t = 0.25$.




Figure 12. Analysis (3), $s, t \in [0.2, 0.4]$.




Figure 13. Analysis (3), $s, t \in [0.15, 0.35]$.




Figure 14. Analysis (3), $s, t = 0.25$.




Figure 15. Analysis (4), $s, t \in [0.2, 0.4]$.




Figure 16. Analysis (4), $s, t \in [0.15, 0.35]$.




Figure 17. Analysis (4), $s, t = 0.25$.




Figure 18. Analysis (5), $s, t \in [0.2, 0.4]$.




Figure 19. Analysis (5), $s, t \in [0.15, 0.35]$.




Figure 20. Analysis (5), $s, t = 0.25$.

We show in Figures 21–23 the Pareto boundaries for the five analyzed strategy functions in Table 1 (the frontiers are almost equal). Robustness to function family is tested across cosine, sine, and mixed: frontiers overlap with Jaccard index > 0.85 , affirming the sub-manifold choice’s stability.




Figure 21. $s, t \in [0.2, 0.4]$.




Figure 22. $s, t \in [0.15, 0.35]$.




Figure 23. $s, t = 0.25$.

6. Case Study: Application to the Battery Electric Vehicle (BEV) Duopoly Between Tesla and BYD

This section bridges the abstract framework to a concrete real-world application in the global BEV market. It is organized as follows: first, a data-driven market context and mapping to the Bayesian Cournot model (Section 6.1); second, a brief recap of the computational restrictions and GA application (Section 6.2); and third, an interpretation of the symmetric frontiers through asymmetric payoff rescaling to reflect observed sales and price ratios (Section 6.3). This section demonstrates how the ex-post Pareto frontier provides actionable insights for strategic decision-making under cost uncertainty.

Specifically, we apply our Bayesian Cournot framework to a contemporary industrial example: the duopolistic competition in the global BEV market between Tesla (USA) and BYD (China). As of 2025, this market represents a “hot” arena of strategic interaction, characterized by rapid technological advancements, supply chain uncertainties, and incomplete information about rivals’ cost structures. According to recent industry data, BYD sold approximately 2.26 million BEVs in 2025, surpassing Tesla’s 1.64 million deliveries for the first time, marking a 28% year-over-year growth for BYD and a 9% decline for Tesla. This shift underscores the relevance of our model, where firms must optimize production under uncertainty about marginal costs, such as those driven by battery production, raw material prices (e.g., lithium), and proprietary technologies.

6.1. Market Context and Mapping to the Bayesian Cournot Model

The global BEV market can be stylized as a symmetric Cournot duopoly, with Tesla and BYD as the primary players competing on production quantities. Demand is driven by factors like consumer preferences for sustainable transport, government incentives (e.g., subsidies in Europe and the US), and tariffs (e.g., high US tariffs on Chinese EVs effectively segmenting markets). However, incomplete information is pervasive: Marginal production costs—estimated at around \$33,000 per vehicle for Tesla in late 2025 (down from \$39,000 in 2023 due to efficiencies like the 4680 battery)—are private and variable. For BYD, vertically integrated operations yield lower costs (e.g., operating margins of 6.4% in 2024, improving despite a 33% profit dip from price wars), but specifics on battery and assembly costs remain opaque to rivals. Types (marginal costs) can be modeled as drawn from continuous intervals, say $[c_*(T), c^*(T)]$ for Tesla (reflecting US-based efficiencies) and $[c_*(B), c^*(B)]$ for BYD (leveraging China's supply chain advantages), without assuming probability distributions. In our framework, strategies are functions $x(c_T)$ and $y(c_B)$ mapping own costs to production quantities (in millions of units). Expected payoffs (losses) incorporate uncertainty, as in Equations (4) and (5), with linear demand $P(q) = A - Bq$, where $q = q_T + q_B$, calibrated to Tesla–BYD BEV demand (almost 4 million units in 2025). The Nash equilibrium would imply aggressive production, but our complete analysis reveals Pareto-superior outcomes, such as defensive strategies to mitigate price wars observed in 2025.

6.2. Restriction to Strategic Sub-Manifolds and Genetic Algorithm Application

To handle infinite-dimensional spaces, we restrict strategies to finite-dimensional sub-manifolds, as proposed in Section 5. Using genetic algorithms (as in Appendix A), we can approximate Pareto frontiers for these sub-manifolds. Hypothetically, for cosine strategies with population size 200 and 400 generations, the frontier might yield a Kalai–Smorodinsky solution (e.g., $K' \approx (-0.4, -0.4)$ in bi-loss space, see Figure 24), corresponding to strategies where Tesla produces ~ 1.1 million units at min-costs and with $Q = 5$, and BYD ~ 1.1 million, avoiding overproduction and price erosion. Compared to Nash (aggressive strategies), Pareto solutions suggest cooperative behaviors, such as tacit agreements on output caps, reducing expected losses by 10–15%.




Figure 24. $\alpha, \beta \in [0, (\pi/2)/0.2]$, $s, t \in [0.2, 0.4]$.

The bi-loss $K' \approx (0.4, 0.4)$ corresponds to the state of the game

$$K(s, t, \alpha, \beta) \approx (0.20, 0.20, 6.74, 6.74).$$

If we fix $s = 0.35$ and $t = 0.21$ (according to the real costs declared by the firms), the Kalai–Smorodinsky solution will be $K' \approx (0.36, 0.36)$, from which we easily deduce the retro-image $K \approx (0.35, 0.21, 3.80, 6.50)$. In this case, the production will be

$$x = Q * \cos(3.8 * 0.35) \approx 1.20 \quad \text{million units}$$

$$\text{and } y = Q * \cos(6.5 * 0.21) \approx 1.02.$$

6.3. Asymmetry Hidden in Symmetry

The global battery electric vehicle (BEV) market in 2025 exemplifies a near-duopoly under incomplete cost information, with Tesla and BYD dominating sales amid uncertainty over rivals' marginal costs (e.g., battery production, supply chains, subsidies). In 2025, BYD sold approximately 2.26 million pure BEVs, surpassing Tesla's 1.64 million units, yielding a sales ratio of Tesla/BYD $\approx 0.726 \approx 0.75$. This ratio aligns with pricing for equivalent models: the BYD Seal Excellence (530 hp, ≈ 520 km WLTP range) retails at $\approx 48,000$ euro, while the Tesla Model 3 Performance (510 hp, ≈ 528 km WLTP range) is priced at $\approx 60,000$ euro, giving a BYD/Tesla price ratio ≈ 0.8 (or ≈ 0.75 under promotional adjustments). These figures reflect BYD's cost advantages (vertical integration, scale) versus Tesla's premium positioning (software, brand), motivating an asymmetric interpretation of our symmetric Bayesian Cournot model.

Our framework remains symmetric in structure—identical type intervals $S = T = [c_*, c^*]$, production capacities $[0, Q]$, and oscillatory strategy restrictions—leading to symmetric ex-post Pareto frontiers (as computed in Section 5). However, to capture real asymmetries, we rescale BYD's payoff function to reflect lower effective prices and higher volumes for equivalent profits. Let $r \approx 0.75$ be the price scaling factor (BYD's relative price). Adjust BYD's payoff as follows: multiply the price term $(Q - x(s) - y(t))$ by r , and scale the quantity $y(t)$ by $1/r \approx 1.333$ (to compensate volume for lower margins) while preserving the shared market price structure.

Formally, the original symmetric payoffs are

$$e_s(x, y)(t) = x(s)(Q - x(s) - y(t)) - s x(s), \quad (16)$$

$$f_t(x, y)(s) = y(t)(Q - x(s) - y(t)) - t y(t). \quad (17)$$

The adjusted BYD payoff is

$$f_t^{\text{adj}}(x, y)(s) = [y(t) \cdot (1/r)] \cdot [r(Q - x(s) - y(t))] - t[y(t) \cdot (1/r)], \quad (18)$$

which simplifies to

$$f_t^{\text{adj}}(x, y)(s) = y(t)(Q - x(s) - y(t)) - t y(t) \cdot (1/r). \quad (19)$$

This introduces an extra cost penalty $t y(t) \cdot (1/r - 1) \approx t y(t) \cdot 0.333$ for BYD, reflecting the “effort” to achieve equivalent profits at lower prices via higher volumes. The game form remains quadratic (Cournot-like), but BYD's effective marginal cost inflates proportionally to $y(t)$, mirroring real thinner margins on higher sales.

Under this rescaling, the ex-post payoff mapping Φ (Section 5.2) becomes

$$\Phi(s, t, x, y) = (e_s(x, y)(t), f_t^{\text{adj}}(x, y)(s)), \quad (20)$$

with the general payoff space and Pareto frontier adjusted accordingly. The oscillatory restrictions (Table 1, with Q-rescaling as in (14)) ensure exhaustivity: for $\alpha, \beta > (\pi/2)/c_*$,

the joint production map $(\alpha, \beta) \mapsto (x_\alpha(s), y_\beta(t))$ densely (or exactly) covers $[0, Q]^2$ per (s, t) , recovering the full adjusted payoff space (up to closure, via continuity on compact domains).

In the symmetric model, GA-computed frontiers (e.g., Figure 24) exhibit symmetry around $u_1 = u_2$, with efficient compromises at equal payoffs. Post-rescaling, these map to asymmetric interpretations: Tesla (higher margins, lower effective quantities) versus BYD (lower margins, higher effective quantities $y' = y(t) \cdot (1/r)$), yet with balanced profits. For instance, a symmetric point $(u_1, u_2) = (0.36, 0.36)$ adjusts to Tesla 3.6 billion dollars versus BYD ≈ 2.7 billion dollars, but BYD's real sales $1.36 (\approx 1.33) \times$ Tesla's to achieve parity—matching the observed $0.72 (\approx 0.75)$ inverse ratio.

This “asymmetry hidden in symmetry” provides practical insights for the EV duopoly: under cost uncertainty (e.g., Tesla uncertain about BYD's battery efficiencies), decision-makers can use oscillatory strategies as robust production policies, exploring ex-post efficient trade-offs without specifying beliefs over types. Symmetric Pareto points recommend mutual restraint (moderated outputs to sustain prices), avoiding destructive volume wars—a relevant strategy as BYD expands globally and Tesla counters with price cuts. In the large landscape, the rescaling preserves the model's qualitative results while aligning with data, highlighting potential collusive paths for shared profitability amid competition.

7. Limitations and Future Directions

While the proposed framework offers a computationally tractable and theoretically grounded approach to approximating the ex-post Pareto frontier in continuous-type Bayesian Cournot games, several limitations inherent to the modeling assumptions and methodological choices must be acknowledged. These simplifications enable tractability in an otherwise intractable setting, but they also constrain the direct applicability of the results to certain real-world deviations. We discuss the main limitations below and suggest directions for future research to address them.

First, the model maintains full symmetry between players: identical type intervals $S = T = [c_*, c^*]$, the same inverse-demand function $P = Q - (x + y)$, and symmetric payoff structures. This abstraction simplifies analysis and computation but abstracts from real asymmetries frequently observed in oligopolistic markets (e.g., differential cost advantages, capacity constraints, brand premiums, or heterogeneous demand elasticities). In practice, such asymmetries could shift the ex-post Pareto frontier and alter the location of efficient compromises. For instance, a firm with persistently lower marginal costs might dominate aggressive high-output strategies, reducing the range of mutually beneficial outcomes. The rescaling approach introduced in the Tesla–BYD application (Section 6.3) partially mitigates this by adjusting one player's payoff to reflect price/volume trade-offs, but more general asymmetric extensions (e.g., player-specific demand intercepts, different type distributions, or heterogeneous production capacities) remain unexplored and represent a natural next step.

Second, the restriction to finite-dimensional differentiable sub-manifolds—specifically one-parameter oscillatory families (cosine, sine, mixed)—assumes that viable production strategies can be adequately parametrized by smooth, periodic functions. While these families achieve near-exhaustivity under the joint surjectivity condition $(\alpha, \beta > (\pi/2)/c_*)$, they exclude non-differentiable, discontinuous, or highly irregular policies that might be feasible in real markets (e.g., discrete quantity jumps due to production batching, capacity rationing, or regulatory constraints). Deviations from smoothness could reduce the density of the restricted payoff space, potentially missing some ex-post efficient points. Similarly, the affine Q -rescaling assumes uniform mapping of the oscillatory range $[-1, 1]$ to $[0, Q]$; alternative normalizations (e.g., adaptive per-type scaling) might further improve coverage but introduce additional parameters.

Third, the linear inverse-demand assumption ($P = Q - (x + y)$) and constant marginal costs per type are standard simplifications in Cournot models but do not capture nonlinear demand curves, capacity constraints, or cost structures that vary with output (e.g., increasing marginal costs due to resource scarcity). Nonlinearities or externalities (e.g., network effects, learning curves) could reshape the payoff landscape, potentially making the frontier more or less concave. The ex-post perspective remains robust in principle—since Φ maps all realizations regardless of functional form—but the numerical approximation quality and interpretability of frontiers would require re-validation.

Fourth, the genetic algorithm, while customized and validated (Section 5.3), is inherently stochastic and population-based. Although hypervolume convergence, sensitivity analysis, and comparisons with NSGA-II, PSO and scalarization methods demonstrate reliability, the method does not guarantee global optimality or provide formal convergence rates in the infinite-dimensional setting. Finite parameter ranges and discretization effects may leave small gaps in coverage, especially near boundary costs c_* or c^* .

Finally, the approach deliberately avoids dependence on private probability measures p and q over types, which enhances robustness but limits direct comparison with belief-dependent solutions (e.g., Bayesian Nash equilibria). This trade-off prioritizes ex-post efficiency over expected-utility maximization.

Future research could address these limitations in several directions. First, extensions to asymmetric duopolies with player-specific type intervals, demand functions, or cost structures would broaden applicability. Second, incorporating richer sub-manifold families (e.g., polynomial bases, neural-network parametrizations, or hybrid oscillatory-polynomial curves) could improve exhaustivity and flexibility. Third, integrating nonlinear demand or dynamic elements (e.g., repeated games, learning) would allow exploration of intertemporal trade-offs. Fourth, formal convergence analysis of the genetic algorithm in restricted functional spaces, perhaps using measure-theoretic arguments or regret bounds, would strengthen theoretical guarantees. Finally, empirical calibration using real industry data (beyond the illustrative Tesla–BYD case) could validate the frontiers against observed outcomes in markets with verifiable cost uncertainty.

These extensions would further bridge the gap between theoretical tractability and practical strategic decision-making under incomplete information.

8. Conclusions

In this paper, we have developed a practical and theoretically grounded framework for analyzing Bayesian Cournot duopolies under bilateral incomplete cost information, where marginal costs are drawn from continuous compact intervals $[c_*, c^*]$ with $0 < c_* < c^* < \infty$. By addressing the inherent challenges of infinite-dimensional functional strategy spaces, we introduced a restriction to finite-dimensional differentiable sub-manifolds—specifically, one-parameter oscillatory families (cosine, sine, and mixed)—with affine Q -rescaling to ensure feasible productions in $[0, Q]$. The key innovation lies in the provable exhaustivity of these restrictions: under the joint surjectivity condition per type pair and parameter thresholds $\alpha, \beta > (\pi/2)/c_*$, the restricted payoff space recovers the full ex-post payoff space, enabling representative approximations of the general Pareto frontier without reliance on private probability measures over types.

The ex-post payoff mapping

$$\Phi(s, t, x, y) = (e_s(x, y)(t), f_t(x, y)(s))$$

unifies all realizable payoff pairs across type-strategy realizations, and its non-dominated image constitutes the belief-independent Pareto frontier—efficient outcomes that traditional expected-utility approaches (e.g., Bayesian Nash equilibria) cannot guarantee ex-post.

Leveraging multi-objective genetic algorithms, we computationally approximated this frontier and selected collusive compromises, with validation demonstrating hypervolume convergence (stabilization after ≈ 40 generations), parameter robustness ($< 5\%$ HV deviation), and superiority over alternatives like NSGA-II, particle swarm optimization, and scalarization methods on discretized benchmarks.

The addition of a real-world application to the Tesla–BYD electric vehicle duopoly—with asymmetric payoff rescaling to reflect observed sales and price ratios (≈ 0.75)—translates abstract frontiers into actionable insights, such as mutual output restraint to avoid price wars. Furthermore, explicit discussions of limitations (e.g., symmetry assumptions, smoothness constraints) and a glossary of key concepts enhance clarity and readability, while rigorous validation addresses methodological concerns.

The findings are significant for decision-makers in oligopolistic markets facing genuine cost uncertainty: our method provides viable, implementable production policies that explore efficient trade-offs and belief-independent compromises, potentially guiding tacit coordination or collusive outcomes without explicit belief forecasting—a limitation of prior work. By shifting emphasis to ex-post efficiency, the framework reveals robust solutions that may foster sustainable competition in real-world settings like the EV industry.

Future directions include extensions to asymmetric duopolies, richer sub-manifold families (e.g., neural parametrizations), nonlinear demand structures, and formal convergence guarantees for the GA in functional spaces. These avenues will further bridge theoretical game analysis with practical strategic decision-making under incomplete information.

Author Contributions: Conceptualization, D.C.; methodology, D.C. and A.D.; software, A.D. and E.P.; writing—original draft preparation, D.C., A.D. and E.P., writing—review & editing, D.C. and A.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: We wish to thank the three anonymous referees for the constructive feedback, which fundamentally strengthened the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. MATLAB Code

In the optimization toolbox of the MATLAB software, version 9.5, we select multi-objective optimization using genetic algorithms. We show, for instance, the functions of Analyses (2), (4) and (5) of Table 1.

```
function [f] = games2(h)
    Q=2;
    s=h(1);
    t=h(2);
    a=h(3);
    b=h(4);
    x=Q*(sin(a.*s)); %a in [0,\pi/(2 c_*)]
    y=Q*(sin(b.*t)); %b in [0,\pi/(2 c_*)]
    f(1)= -(x.*(Q-x-y)-s.*x)
    f(2)= -(y.*(Q-x-y)-t.*y)
end
%
function [f] = games3(h)
```

```

Q=2;
s=h(1);
t=h(2);
a=h(3);
b=h(4);
x=Q.*exp(-(a.*s).^2);%a in [0,100]
y=Q.*exp(-(b.*t).^2);%b in [0,100]
f(1)= -(x.*(Q-x-y)-s.*x)
f(2)=-(y.*(Q-x-y)-t.*y)
end
%
function [f] = games4(h)
Q=2;
s=h(1);
t=h(2);
a=h(3);
b=h(4);
x=Q.*a.^2.*s.^2;% a in [0, 1/c_*]
y=Q.*b.^2.*t.^2;% b in [0, 1/c_*]
f(1)= -(x.*(Q-x-y)-s.*x)
f(2)=-(y.*(Q-x-y)-t.*y)
end

```

For a comparison with the NSGA-II optimization method and MOPSO optimization in the numerical example, we refer to the code proposed in [32,33], which performs, respectively, a Non-Sorting Genetic Algorithm-II over continuous functions and a Multi-Objective Particle Swarm Optimization over continuous functions. The functions to minimize and the chosen parameters are the following.

```

%      - MultiObj: Struct that contains the parameters relative to the %
%      optimization functions. %
%      * MultiObj.fun:      Anonymous multi-obj function to minimize. %
%      * MultiObj.nVar:     Number of variables. %
%      * MultiObj.var_min:  Vector that indicates the minimum values %
%      of the search space in each dimension. %
%      * MultiObj.var_max:  Same than 'var_min' with the maxima. %
Q=2
MultiObj.fun = @(x) [ -((cos(x(:,3)).*x(:,1))).*(Q-(cos(x(:,3)).*x(:,1)))
-(cos(x(:,4)).*x(:,2))))-x(:,1).*(cos(x(:,3)).*x(:,1))),
-(cos(x(:,4)).*x(:,2))).*(Q-(cos(x(:,3)).*x(:,1)))
-(cos(x(:,4)).*x(:,2))))-x(:,2).*(cos(x(:,4)).*x(:,2)))]);
MultiObj.nVar = 4;
MultiObj.var_min = [0, 0, 1, 1];
MultiObj.var_max = [0.5, 0.5, 5, 5];
% Parameters NSGA
params.Np = 200; % Population size
params.pc = 0.9; % Probability of crossover
params.pm = 0.5; % Probability of mutation
params.maxgen = 100; % Maximum number of generations
params.ms = 0.05; % Mutation strength
% Parameters MOPSO

```

```

params.Np = 200;      % Population size
params.Nr = 200;      % Repository size
params.maxgen = 100;   % Maximum number of generations
params.W = 0.4;        % Inertia weight
params.C1 = 2;         % Individual confidence factor
params.C2 = 2;         % Swarm confidence factor
params.ngrid = 20;     % Number of grids in each dimension
params.maxvel = 5;     % Maximum vel in percentage
params.u_mut = 0.5;    % Uniform mutation percentage

```

Appendix B. Deterministic Game

In this Appendix, we apply the *complete analysis of differentiable games* (introduced by D. Carfi in [34]) to the classic Cournot duopoly (1838), a classic oligopolistic market in which there are two enterprises producing the same commodity and selling it in the same market. In this classic model, in a competitive background, the two enterprises employ as possible strategies the quantities of the commodity produced. The main solutions proposed in the literature for this kind of duopoly are the *Nash equilibrium* and the *Collusive optimum*, without any subsequent critical examination of these two kinds of solutions. The absence of any critical quantitative analysis is due to the relevant lack of knowledge regarding the set of all possible outcomes of this strategic interaction. On the contrary, by considering the Cournot duopoly as a *differentiable game* (games with differentiable payoff functions) and studying it using the new topological methodologies introduced by D. Carfi, we obtain an exhaustive and complete vision of the entire payoff space of the Cournot game (this also in asymmetric cases with the help of computers), and this complete view allows us to analyze critically the classic solutions and to find other methods to select Pareto strategies. In order to illustrate the application of this topological methodology to the considered infinite game, several compromise pricing-decisions are considered, and we show how the complete study gives a real, extremely extended comprehension of the classic model [35].

Appendix B.1. Deterministic Cournot Duopoly: Formal Description of Cournot's Normal-Form Game

Our model will be a non-linear two-player loss game $G = (f, <)$ (see also [34]). The two players (enterprises) shall be called *Emil* and *Frances* (following Aubin books [36,37]).

Assumption A1 (strategy sets). *The two players produce and offer the same commodity with the following quantities: $x \in \mathbb{R}_{\geq}$ for Emil and $y \in \mathbb{R}_{\geq}$ for Frances. In more precise terms, the payoff function f of the game is defined on a subset of the positive cone of the Cartesian plane $\mathbb{R}^2$, interpreted as a space of bi-quantities. We assume (for simplicity) that the set of all strategies (of each player) is the interval $E = [0, +\infty[$.*

Assumption A2 (symmetry of the game). *The game will be assumed symmetric with respect to the players. In other terms, the payoff pair $f(x, y)$ is the symmetric of the pair $f(y, x)$.*

Assumption A3 (form of price functions). *We assume the price function as a function $p : \mathbb{R}^2 \rightarrow \mathbb{R}$, depending only upon the total production $x + y$. Precisely, we assume:*

$$p(x, y) = a - b(x + y), \quad (\text{A1})$$

for every productive bi-strategy (x, y) , where $a \geq 0$ is the fixed price and $b \geq 0$ is the absolute marginal price with respect to the total production

$$P(x, y) = x + y. \quad (\text{A2})$$

Assumption A4 (form of cost functions). Let the cost function C_1 (defined on E^2) of the first player be given by

$$C_1(x, y) = cx + d, \quad (\text{A3})$$

for every positive strategy pair (x, y) , and analogously, let the cost function of the second enterprise be given by

$$C_2(x, y) = cy + d, \quad (\text{A4})$$

for every positive bi-strategy (x, y) , where $c > 0$ is the marginal cost and $d \geq 0$ is the fixed cost.

Assumption A5 (payoff functions). The first player's net cost function is defined, classically, by the opposite of the revenue

$$\begin{aligned} f_1(x, y) &= C_1(x, y) - p(x, y)x = \\ &= bx(x + y - (a - c)/b) + d, \end{aligned} \quad (\text{A5})$$

for every positive bi-strategy (x, y) . Symmetrically, for Frances, the loss function is defined by

$$\begin{aligned} f_2(x, y) &= C_2(x, y) - p(x, y)y = \\ &= by(x + y - (a - c)/b) + d, \end{aligned} \quad (\text{A6})$$

for every positive bi-strategy (x, y) .

Assumption A6 (critical production of duopoly u). Both the rival players will be able to obtain a positive net profit only if the following quantity is positive:

$$u = (a - c)/b, \quad (\text{A7})$$

If u were non-positive, for every $(x, y) \geq 0$, the value $x + y - u$ would be non-negative, and it would be the same for the loss-values $f_1(x, y)$ and $f_2(x, y)$.

Each player, in order to obtain non-negative profits, has to exclude all strategies strictly greater than u . If $x > u$, the value

$$x + y - u$$

is positive and strictly greater than y ; then, the loss $f_1(x, y)$ is also positive (and it is the same for Frances when y is strictly greater than u). The strategic space is so reduced from the unbounded interval $[0, +\infty[$ to the compact interval $[0, u]$.

The players, to obtain losses below the fixed cost, have to exclude, collectively, all bi-strategies having total productivity strictly greater than u . The relation $f_1(x, y) < d$ is equivalent to

$$bx(x + y - u) < 0, \quad (\text{A8})$$

which is equivalent to $x + y < u$ because of $bx > 0$ (it is the same for Frances).

The bi-loss function is so defined by

$$f(x, y) = b(x(x + y - u), y(x + y - u)) + (d, d), \quad (\text{A9})$$

for every bi-strategy (x, y) of the game in the unbounded square E^2 .

Appendix B.2. Study of the Deterministic Cournot Game

Appendix B.2.1. Payoff Function Under Examination

When the *fixed cost* d is zero and the *absolute marginal price* b is one (this assumption determines only a “reversible” translation of the loss space), the bi-loss function f from the compact square $[0, u]^2$ into the bi-loss plane R^2 is defined by

$$f(x, y) = (x(x + y - u), y(x + y - u)), \quad (\text{A10})$$

for every bi-strategy (x, y) in the square $S = [0, u]^2$, which is the convex envelope of its vertices, denoted by A, B, C, D starting from the origin and going anticlockwise.

When the critical production u is 1, we will obtain the *payoff vector function* defined by

$$f(x, y) = (x(x + y - 1), y(x + y - 1)), \quad (\text{A11})$$

on the strategy square $S = [0, 1]^2$.

Appendix B.2.2. Non-Cooperative Friendly Phase

Figure A1 shows the payoff space of the deterministic Cournot game, where the *Pareto minimal boundary of the payoff space* $f(S)$ is the green segment

$$\partial(f(S)) = [H', K'].$$




Figure A1. Payoff space of the Cournot game (the Cournot “fish”).

Examining Figure A2, we will note that the game has both extremes; that is, the *shadow minimum* $\alpha = (-1/4, -1/4)$ and the *maximum* $\beta = C' = (1, 1)$.

The *Pareto maximal boundary of the payoff space* $f(S)$ coincides with the maximum $\beta = C' = (1, 1)$.

Emil and Frances do not control the Pareto minimal boundary; they could reach the point H' and K' of the boundary, but the solution is not satisfactory for them. In fact, a player will suffer the maximum loss.




Figure A2. Extrema and payoff Pareto boundary of the Cournot game.

Appendix B.2.3. Properly Non-Cooperative (or Selfish) Phase

Now, we will consider *the best reply correspondences* of the two players, Emil and Frances.

If Frances produces (and sells) the quantity y of the commodity, Emil—in order to reply rationally—should minimize his *partial cost function*

$$f_1(\cdot, y) : x \mapsto x(x + y - 1), \quad (\text{A12})$$

on the compact interval $[0, 1]$.

According to the Weierstrass theorem, there exists at least one strategy of Emil minimizing that partial cost function, and, by Fermàt theorem, *Emil's best reply strategy to Frances' strategy y* is the unique quantity

$$B_1(y) := (1/2)(1 - y). \quad (\text{A13})$$

Indeed, the partial derivative

$$f_1(\cdot, y)'(x) = 2x + y - 1, \quad (\text{A14})$$

is negative for every $x < (1/2)(1 - y)$ and positive for each $x > (1/2)(1 - y)$.

Therefore, Emil's best reply correspondence is the function B_1 , from the interval $[0, 1]$ into the interval $[0, 1]$, defined by

$$y \mapsto (1/2)(1 - y). \quad (\text{A15})$$

As we already have observed, our Cournot game is a symmetric game; therefore, Frances' best reply correspondence is the function B_2 , from $[0, 1]$ into $[0, 1]$, defined by the association

$$x \mapsto (1/2)(1 - x). \quad (\text{A16})$$

The *Nash equilibrium* is the fixed point of the function B , determined by the pair of reaction functions (B_1, B_2) and defined, from the Cartesian product $E \times F$ of the domains

of the pair (B_1, B_2) (in inverse order) into the Cartesian product $E \times F$ of the co-domains (in direct order), by

$$B : [0, 1]^2 \rightarrow [0, 1]^2 : (x, y) \mapsto (B_1(y), B_2(x)); \quad (\text{A17})$$

The Nash equilibrium is therefore the unique bi-strategy (x, y) satisfying the below system of linear equations:

$$\begin{aligned} x &= (1/2)(1 - y) \\ y &= (1/2)(1 - x), \end{aligned} \quad (\text{A18})$$

That is, the point $N = (1/3, 1/3)$ —as we can see also from Figure A3—determines a bi-loss $N' = (-1/9, -1/9)$ (hence a bi-gain $(1/9, 1/9) = -N'$), as Figure A4 shows.




Figure A3. Nash Equilibrium of the Cournot game.




Figure A4. Payoff at the Nash equilibrium of the Cournot game.

The Nash equilibrium is not collectively satisfactory because it is not a Pareto optimal bi-strategy, but it represents the unique *properly non-cooperative* game solution.

Appendix B.3. Other Solutions in the Deterministic Cournot Model: Defensive and Offensive Phase

The players' conservative values are obtained through their *worst loss functions*.

Worst loss functions. On the strategic square $S = [0, 1]^2$, Emil's worst loss function is defined by

$$\begin{aligned} f_1^\#(x) &= \sup_{y \in F} x(x + y - 1) = \\ &= x^2, \end{aligned} \quad (\text{A19})$$

its minimum is the loss

$$\begin{aligned} v_1^\# &= \inf_{x \in E} (f_1^\#(x)) = \\ &= \inf_{x \in E} (x^2) = \\ &= 0, \end{aligned} \quad (\text{A20})$$

attained at the unique *conservative strategy* $x^\# = 0$.

Symmetrically, Frances' worst loss function is defined by

$$f_2^\#(y) = y^2, \quad (\text{A21})$$

its minimum is the loss $v_2^\# = 0$, attained at the unique *conservative strategy* $y^\# = 0$.

Conservative bivalue. The *conservative bivalue* is then the bi-loss

$$v^\# = (v_1^\#, v_2^\#) = (0, 0).$$

The worst offensive multi-functions are determined by the study of the *worst loss functions*.

Frances' worst offensive reaction multifunction $O_2 : E \rightarrow F$ is defined by $O_2(x) = 1$ for each of Emil's strategies x ; indeed, for a fixed strategy of Emil x , Frances' strategy 1 strictly maximizes the partial cost function $f_1(x, \cdot)$. Symmetrically, Emil's worst offensive correspondence versus Frances' $O_1 : F \rightarrow E$ is defined by $O_1(y) = 1$ for each of Frances' strategies y .

The dominant offensive strategy is 1 for both players; indeed, the offensive correspondences are constant.

The offensive (dominant) equilibrium $C = (1, 1)$ brings to the point $C' = (1, 1)$, in which a severe loss is registered for both players.

The core of the payoff space (in the sense introduced by J. P. Aubin) is the part of the Pareto minimal boundary contained in the cone of lower bounds of the conservative bi-value $v^\#$. The core of our game is exactly the whole of the Pareto minimal boundary, so the conservative bi-value does not give us any bound for the choice of cooperative bistrategies.

Remark. Let us note that, in this case, the Nash equilibrium gives a very small area of the Pareto minimal boundary in which an agreement can be reasonably attained: we have a very small *Nash payoff core* of our game (the part of the Pareto minimal boundary of the payoff space, contained in the cone of lower bounds of the Nash equilibrium payoff $f(N)$).

The conservative knot zone of the game is the set of all strategy profiles that leads to the conservative bi-value $v^\#$: it is the union

$$S^\# = [B, D] \cup \{A\}.$$

Appendix B.4. Cooperative Phase

A possible agreement between the two players is the *best compromise solution* (in the sense introduced by J.P. Aubin); it is the pair of strategies leading to the unique payoff at the intersection between the segment joining the conservative bi-value with the infimum of the payoff space and the Pareto minimal boundary of the payoff space itself: in our specific case, it is $(1/4, 1/4)$, shown graphically in Figure A5. This compromise bi-strategy determines just the bi-loss $(-1/8, -1/8)$ shown in Figure A6.

Remark. In the gain games, the best compromise solution in the sense of J. P. Aubin is the intersection of the segment joining the conservative bi-value and the supremum of the game with the Pareto boundary; it is essentially the Kalai–Smorodinsky solution of the bargaining problem $(\partial^*(f(S)), \sup(f(S)), v^\#)$, with choice space $f(S)$, utopia point $\sup(f(S))$ and threat point $v^\#$.




Figure A5. Compromise solutions of the Cournot game.




Figure A6. Conservative study: the core and the Kalai–Smorodinsky payoff of the Cournot game.

Other compromise solutions? The best compromise solution coincides with:

- the *core best compromise*;
- the *Nash bargaining solution*;
- the *bi-strategy with closest bi-loss to the shadow minimum*;
- the *Kalai–Smorodinsky bargaining solution*.

It is also one of the *transferable utility solutions* (strategies minimizing the aggregate utility function $f_1 + f_2$), which in this case constitutes the entire Pareto minimal boundary. This can be easily viewed based on geometric evidence considering in the payoff universe the levels of that aggregate function, which are affine lines parallel to the vector $(1, -1)$.

Selection of Pareto solutions. The Nash equilibrium can help in the selection of Pareto solutions (Figure A6). Indeed, if Emil and Frances decide to cooperate, their possible choices will lead to those points of the Pareto boundary which are also lower bounds of the Nash Equilibrium in order to obtain a compromise solution that is strictly better than the non-cooperative one.

Appendix C. Glossary of Key Concepts

This appendix provides concise definitions of the central concepts used in this paper. Terms are explained in plain language where possible, followed by technical details and references to relevant sections.

Bayesian game. A strategic interaction in which players have incomplete information about certain parameters (e.g., costs, types) relevant to payoffs. Each player holds private information (their type) and forms beliefs (probability distributions) about others' types. Players maximize expected payoffs conditional on these beliefs (Section 1, Introduction).

Bayesian Nash equilibrium (BNE). A strategy profile in a Bayesian game where each player's strategy maximizes their expected payoff given their beliefs about opponents' types and assuming opponents play optimally. BNE is belief-dependent and optimizes in expectation, not necessarily ex-post across all realizations (Section 5.2, contrasted with ex-post frontier).

Differentiable game. A normal-form game in which strategy spaces are subsets of differentiable function spaces (e.g., L^1 spaces or sub-manifolds) and payoff functions are differentiable with respect to strategies. In our Cournot setting, this allows smooth production functions over types (Appendix B).

Ex-post payoff mapping Φ . The function $\Phi(s, t, x, y) = (e_s(x, y)(t), f_t(x, y)(s))$ that returns the realized payoffs to both players for every possible combination of types (s, t) and strategies (x, y) . It collects all feasible realized (not expected) payoff pairs once types are revealed (Section 5.2).

Ex-post Pareto frontier (general Pareto frontier). The set of non-dominated points in the image of Φ : payoff pairs (u_1, u_2) where no other realizable pair offers strictly higher payoffs to both players (or higher to one without lowering the other). It captures efficient realized outcomes across all type-strategy realizations, independent of beliefs over types (Section 5.2).

Exhaustivity (near-exhaustivity). The property that the restricted payoff space generated by the oscillatory sub-manifolds densely (or exactly) covers the full ex-post payoff space $\text{Image}(\Phi)$. It is achieved when the joint production map covers $[0, Q]^2$ densely for every fixed cost pair (s, t) (Section 5.2).

Joint surjectivity (per type pair). The sufficient condition for exhaustivity: for every fixed (s, t) , the map $(\alpha, \beta) \mapsto (x_\alpha(s), y_\beta(t))$ is surjective onto $[0, Q] \times [0, Q]$. Ensures every possible production pair is realizable under the restricted families for each cost realization (Section 5.2, Theorem).

Multi-objective optimization (Pareto optimization). Optimization of two or more conflicting objectives simultaneously, yielding a set of non-dominated (Pareto-optimal) solutions rather than a single optimum. In our context, the two objectives are the players' realized payoffs (Sections 2 and 3).

Oscillatory sub-manifold (oscillatory strategy family). A one-parameter family of smooth production functions (e.g., $x_\alpha(s) = \frac{Q}{2}(1 + \cos(\alpha s))$, sine or mixed variants) restricted to a low-dimensional differentiable manifold within the infinite-dimensional strategy space E or F . Chosen for smoothness, implementability, and exhaustivity properties (Section 5.1, Table 1).

Payoff space (general payoff space). The image of the ex-post payoff mapping Φ : the set of all possible realized payoff pairs (u_1, u_2) that can occur for some types and strategies. It is the natural analogue of the payoff set in complete-information games, but lifted to include all type realizations (Section 5.2).

Q-rescaling (affine normalization). The affine transformation that maps the natural range of oscillatory functions (typically $[-1, 1]$) to the feasible production interval $[0, Q]$, e.g., $x(s) = \frac{Q}{2}(1 + \cos(\alpha s))$. Ensures strategies remain within bounds while preserving oscillatory behavior (Section 5.2, Equation (14)).

Type (player type). A player's private information parameter; here, the marginal cost drawn from $[c_*, c^*]$. The type determines the cost term in the payoff function and is unknown to the opponent (Section 4).

These definitions aim to balance theoretical precision with accessibility. Readers are encouraged to consult the main text sections for full mathematical derivations and economic intuition.

References

1. Guglielmi, A.V.; Levorato, M.; Badia, L. A Bayesian Game Theoretic Approach to Task Offloading in Edge and Cloud Computing. 2018. Available online: https://www.dei.unipd.it/~badia/papers/2018_06_ICC.pdf (accessed on 15 December 2025).
2. van Damme, E. Games in Normal Form. In *Stability and Perfection of Nash Equilibria*; Springer: Berlin/Heidelberg, Germany, 1987.
3. Polimeni, A.; Donato, A.; Belcore, O.M. Urban freight distribution with electric vehicles: Comparing some solution procedures. *Front. Future Transp.* **2024**, *5*, 1491799. [CrossRef]
4. Ehrgott M. *Multicriteria Optimization*; Birkhäuser: Basel, Switzerland, 2005; ISBN 978-3-540-21398-7.
5. Coello, C.A.; Lamont, G.B.; Van Veldhuisen, D.A. *Evolutionary Algorithms for Solving Multi-Objective Problems*; Springer: Berlin/Heidelberg, Germany, 2007; ISBN 978-0-387-36797-2.
6. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182. [CrossRef]
7. Suman, B.; Kumar, P. A survey of simulated annealing as a tool for single and multiobjective optimization. *J. Oper. Res. Soc.* **2006**, *57*, 1143–1160. [CrossRef]
8. Maccarone, L.T.; Cole, D.G. Bayesian games for the cybersecurity of nuclear power plants. *Int. J. Crit. Infrastruct. Prot.* **2022**, *37*, 100493. [CrossRef]
9. Xia, S.; Lin, F.; Chen, Z.; Tang, C.; Ma, Y.; Yu, X. A Bayesian Game Based Vehicle-to-Vehicle Electricity Trading Scheme for Blockchain-Enabled Internet of Vehicles. *IEEE Trans. Veh. Technol.* **2020**, *69*, 6856–6868.
10. Rabin, Y.A.; Noeloikeau, F.C.; Mahbub A.; Fatemeh G.; Daniel G.; Forte, D. Chaogate Parameter Optimization using Bayesian Optimization and Genetic Algorithm. 2021. Available online: https://faculty.eng.ufl.edu/dforte/wp-content/uploads/sites/657/2021/02/Optimization_of_MOSFET_parameters_for_Chaotic_Operation.pdf (accessed on 18 November 2025).
11. Li, Z.; Wellman, M.P. Evolution Strategies for Approximate Solution of Bayesian Games. *Proc. AAAI Conf. Artif. Intell.* **2021**, *35*, 5531–5540. [CrossRef]
12. Kim, T.; Yannelis, N.C. Existence of Equilibrium in Bayesian Games with Infinitely Many Players. *J. Econ. Theory* **1997**, *77*, 330–353. [CrossRef]
13. Mallozzi, L.; Pusillo, L.; Tijs, S. Approximate equilibria for Bayesian games. *J. Math. Anal. Appl.* **2008**, *343*, 1098–1102. [CrossRef]
14. Binois, M. Uncertainty Quantification on pareto Fronts and High-Dimensional Strategies in Bayesian Optimization, with Applications in Multi-Objective Automotive Design. Ph.D. Thesis, Ecole Nationale Supérieure des Mines de Saint-Etienne, Saint-Étienne, France, 2015. Available online: <https://theses.hal.science/tel-01310521> (accessed on 10 December 2025).

15. Galuzio, P.P.; de Vasconcelos Segundo, E.; dos Santos Coelho, L.; Cocco Mariani, V. MOBOpt—Multi-objective Bayesian optimization. *SoftwareX* **2020**, *12*, 100520. [[CrossRef](#)]
16. Suzuki, S.; Takeno, S.; Tamura, T.; Shitara, K.; Karasuyama, M. Multi-objective Bayesian Optimization using Pareto-frontier Entropy. *arXiv* **2019**, arXiv:1906.00127.
17. Ponweiser, W.; Wagner, T.; Biermann, D.; Vincze, M. Multiobjective Optimization on a Limited Budget of Evaluations Using Model-Assisted S-Metric Selection. In *Parallel Problem Solving from Nature; Lecture Notes in Computer Science*; Rudolph, G., Jansen, T., Beume, N., Lucas, S., Poloni, C., Eds.; Springer: Berlin/Heidelberg, Germany, 2008; Volume 5199. [[CrossRef](#)]
18. Emmerich, M.T.M.; Deutz, A.H.; Klinkenberg, J.W. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In Proceedings of the 2011 IEEE Congress of Evolutionary Computation (CEC), New Orleans, LA, USA, 5–8 June 2011. [[CrossRef](#)]
19. Cranko, Z.; Nock, R. Boosted Density Estimation Remastered. In Proceedings of the ICML, Long Beach, CA, USA, 9–15 June 2019. Available online: <https://icml.cc/virtual/2019/poster/3788> (accessed on 18 November 2025).
20. Fudenberg, D.; Tirole, J. Noncooperative game theory for industrial organization: An introduction and overview. *Handb. Ind. Organ.* **1989**, *1*, 259–327.
21. Shapiro, C. The theory of business strategy. *J. Econ.* **1989**, *20*, 125–137. [[CrossRef](#)]
22. Vives, X. Duopoly Information Equilibria: Cournot and Bertrand. *J. Econ. Theory* **1984**, *34*, 71–94. [[CrossRef](#)]
23. Carfi, D.; Donato, A. Cournot-Bayesian General Equilibrium: A Radon Measure Approach. *Mathematics* **2019**, *7*, 10. [[CrossRef](#)]
24. Church, J.; Ware, R. *Industrial Organization: A Strategic Approach*; McGraw-Hill: New York, NY, USA, 2000.
25. Belleflamme, P.; Peitz, M. *Industrial Organization: Markets and Strategies*; Cambridge University Press: Cambridge, UK, 2015.
26. Harsanyi, J. Games with incomplete information played by Bayesian players. *Manag. Sci.* **1967**, *14*, 159–182. [[CrossRef](#)]
27. Riechmann, T. Genetic algorithm learning and evolutionary games. *J. Econ. Dyn. Control* **2001**, *25*, 1019–1037. [[CrossRef](#)]
28. Donato, A.; Carfi, D. Using Genetic Algorithms to Improve Airport Pavement Structural Condition Assessment: Code Development and Case Study. *Information* **2023**, *14*, 286. [[CrossRef](#)]
29. Coello, C.A.C.; Pulido, G.T.; Lechuga, M.S. Handling multiple objectives with particle swarm optimization. *IEEE Trans. Evol. Comput.* **2004**, *8*, 256–279. [[CrossRef](#)]
30. Zitzler E.; Thiele L. Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Trans. Evol. Comput.* **1999**, *3*, 257–271. Available online: https://cs.ijs.si/papa/courses/06_ZitzlerMultiObjective.pdf (accessed on 15 December 2025). [[CrossRef](#)]
31. Erev I.; Roth A.E. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *Am. Econ. Rev.* **1998**, *88*, 848–881.
32. Martínez-Cagigal, V. Non Sorting Genetic Algorithm II (NSGA-II). MATLAB Central File Exchange. 2026. Available online: <https://it.mathworks.com/matlabcentral/fileexchange/65494-non-sorting-genetic-algorithm-ii-nsga-ii> (accessed on 14 January 2026).
33. Martínez-Cagigal, V. Multi-Objective Particle Swarm Optimization (MOPSO). MATLAB Central File Exchange. 2026. Available online: <https://www.mathworks.com/matlabcentral/fileexchange/62074-multi-objective-particle-swarm-optimization-mopso> (accessed on 14 January 2026).
34. Carfi, D. Payoff space in C^1 Games. *Appl. Sci. (APPS)* **2009**, *11*, 35–47.
35. Carfi, D.; Perrone, E. *Game Complete Analysis of Classic Economic Duopolies*; Lambert Academic Publishing: Saarbrücken, Germany, 2012.
36. Aubin, J.P. *Mathematical Methods of Game and Economic Theory (Revised Edition)*; North-Holland: Amsterdam, The Netherlands, 1982.
37. Aubin, J.P. *Optima and Equilibria*; Springer: Berlin/Heidelberg, Germany, 1998.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.