Journal of Machine Learning **ISSN**: 2790-2048(e), 2790-203X(p)

Stochastic Delay Differential Games: Financial Modeling and Machine Learning Algorithms

Robert Balkin * [1], Hector D. Ceniceros † [1], and Ruimeng Hu ‡ [1,2]

[1]Department of Mathematics, University of California, Santa Barbara, CA 93106-3080, USA.
[2]Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106-3080, USA.

Abstract. In this paper, we propose a numerical methodology for finding the closed-loop Nash equilibrium of stochastic delay differential games through deep learning. These games are prevalent in finance and economics where multi-agent interaction and delayed effects are often desired features in a model, but are introduced at the expense of increased dimensionality of the problem. This increased dimensionality is especially significant as that arising from the number of players is coupled with the potential infinite dimensionality caused by the delay. Our approach involves parameterizing the controls of each player using distinct recurrent neural networks. These recurrent neural network-based controls are then trained using a modified version of Brown's fictitious play, incorporating deep learning techniques. To evaluate the effectiveness of our methodology, we test it on finance-related problems with known solutions. Furthermore, we also develop new problems and derive their analytical Nash equilibrium solutions, which serve as additional benchmarks for assessing the performance of our proposed deep learning approach.

Keywords:
Stochastic delay differential games,
Deep fictitious play,
Recurrent neural networks,
Portfolio games.

Article Info.:
Volume: 3
Number: 1
Pages: 23 - 63
Date: March/2024
doi.org/10.4208/jml.230713

Article History:
Received: 13/07/2023
Accepted: 15/02/2024

Communicated by:
Jiequn Han

1 Introduction

Stochastic delay differential games combine game theory and stochastic control problems with delay. These control problems encompass various models applicable to economics, advertising, and finance. For instance, in determining a firm's optimal advertising policy, Gozzi and Marinelli [10] consider a model which incorporates the delayed impact of advertising expenditures on the firm's goodwill. Similarly, in finance, optimal investment and consumption decisions could also take into account delayed market features as is done by Pang and Hussain [23]. Furthermore, these delayed stochastic control problems can often be extended to incorporate interaction with competitors, who can influence both the underlying system dynamics and the objectives of individual actors. In the context of such scenarios, the stochastic control problem with delay can be further extended to a stochastic delay differential game. This framework captures the interaction among all

*Corresponding author. rbalkin@ucsb.edu
†ceniceros@ucsb.edu
‡rhu@ucsb.edu

participants (or players) who select their controls to optimize their objectives. The controls of each player affect the system dynamics, which are modeled as a system of stochastic delay differential equations (SDDEs). The outcome of the game is represented by the concept of Nash equilibrium, which is a collection of all players' choices, ensuring that no player has the incentive to deviate unilaterally.

Despite introducing mathematical and computational challenges, incorporating delay is crucial for developing more complex and realistic models that capture real-world phenomena. For instance, in the analysis of systemic risk, Carmona *et al.* [6] model bank lending and borrowing as a stochastic differential game without delay, assuming a specific form of bank repayments at time t. To enhance the model's realism, the same authors in collaboration with Mousavi [5] consider banks that must repay their borrowings at time t by time $t + \tau$, introducing a delayed factor into the governing dynamics. While this model effectively captures the nature of delayed repayments, it also increases the mathematical and computational complexity of the underlying problem. However, given the widespread occurrence and realistic nature of delayed problems, it is essential to address the computational challenges they present.

The primary reason for these difficulties lies in the inherent dimensionality of the problem. Stochastic differential games already face the curse of dimensionality when the number of players, denoted as N, is large. Adding to this inherent complexity, stochastic delay differential games introduce a possibly infinite-dimensional component, as the drift and volatility of the associated SDDE depend on the entire path. To formalize this, we note that one can employ the approach of dynamic programming to characterize the value functions associated with the closed-loop Nash equilibrium through a system of Hamilton-Jacobi-Bellman (HJB) equations, enabling the determination of Nash equilibrium controls. However, the resulting HJB equations in the delayed case involve derivatives with respect to variables in an infinite-dimensional Hilbert space as detailed in the book by Fabbri *et al.* [8, Section 2.6.8]. Numerically solving this HJB system would require an additional high-dimensional approximation to handle the infinite dimensionality arising from the delay. However, deep learning methodologies are natural choices for solving problems with high dimensionality and have been used in similar instances. For example, Fouque and Zhang [9] parameterize the optimal control with neural networks to solve a mean field control problem arising from an inter-bank lending model with delayed repayments, and Han and Hu [14] solve the stochastic control problems with delay using neural networks.

To address the challenge of high dimensionality in these problems, we propose a deep learning-based method that effectively handles the delay. Inspired by the approach presented in [14], which utilizes recurrent neural networks (RNNs) to solve stochastic control problems with delay, we introduce an algorithm for finding the Nash equilibrium of stochastic delay differential games. Specifically, we parameterize players' controls using RNNs and approximate their objective functions by sampling the game dynamics under these RNN-based controls. The parameters of the RNNs are then optimized using the concept of deep fictitious play, as introduced in [12,16]. The utilization of neural network-based control functions enables us to reformulate the problem in a finite-dimensional setting. Now, the optimization for a given player revolves around selecting the neural network parameters. Moreover, employing RNNs, in particular, allows us to effectively