REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
12-02-2020	Final Report	3-Apr-2017 - 2-Feb-2020

4. TITLE AND SUBTITLE
Final Report: Large Deviations in Multi-Agent Systems

5a. CONTRACT NUMBER
W911NF-17-1-0134

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

6. AUTHORS

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
University of Wisconsin - Madison
Suite 6401
21 N Park Street
Madison, WI 53715 -1218

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
70671-NS.8

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON / 19b. TELEPHONE NUMBER
UU	UU	UU	UU		William Sandholm / 608-263-3858

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

RPPR Final Report
as of 15-Apr-2020

Agency Code:

Proposal Number: 70671NS **Agreement Number: W911NF-17-1-0134**

INVESTIGATOR(S):

Name: William Sandholm
Email: whs@ssc.wisc.edu
Phone Number: 6082633858
Principal: Y

Organization: **University of Wisconsin - Madison**
Address: Suite 6401, Madison, WI 537151218
Country: USA
DUNS Number: 161202122 EIN: 396006492

Report Date: 02-May-2020 Date Received: 12-Feb-2020
Final Report for Period Beginning 03-Apr-2017 and Ending 02-Feb-2020

Title: Large Deviations in Multi-Agent Systems

Begin Performance Period: 03-Apr-2017 **End Performance Period:** 02-Feb-2020
Report Term: 0-Other

Submitted By: William Sandholm Email: whs@ssc.wisc.edu
Phone: (608) 263-3858

STEM Degrees: 0 **STEM Participants:** 1

Major Goals: Large multi-agent systems are basic models in wide variety of disciplines, ranging from the social sciences to engineering and the physical sciences. In many applications, these systems are best understood from a game-theoretic perspective, with each agent being endowed with a payoff function that it aims to maximize. To obtain a dynamic model of agents' behavior, one also assigns each agent a revision protocol, which describes how the agent uses the information it possesses to decide when to switch actions, and which action to choose.

Most work in evolutionary game theory has focused on two questions. One, equilibrium convergence, considers which multi-agent systems will achieve equilibrium configurations over moderate time spans. The other, long-run equilibrium selection, considers models in which agents sometimes choose suboptimal actions, and describes which equilibrium will be played in a large proportion of periods over long enough time spans.

For a full understanding of large multi-agent systems, one must address a third question: that of equilibrium breakdown. Here the aim is to understand how and when equilibrium is likely to unravel, and which new equilibrium, if any, is likely to arise in its place. While such questions are of basic importance, they have attracted limited attention in the literature, in part because of the technical demands they impose.

In our first major goal, we use methods from large deviations theory to study escape from and transitions among equilibria in large multi-agent systems. The analysis of large deviations in games takes the theory of equilibrium convergence as its prerequisite, and in turn, this analysis provides new, general, tractable methods for the study of equilibrium selection.

In the basic model considered here, the behavior of revising agents is described by a noisy best response protocol, under which a revising agent typically chooses an optimal action, but places positive probability on all actions. To account for agents' incentives to avoid costly mistakes, we assume that the probability of any given suboptimal choice depends on its payoff consequences.

We first consider the question of large deviations in the large population limit, holding the noise level in agents' decisions fixed. We establish a large deviations principle, which describes the rate of decay of the probability of observing sample paths from any given set as the population size grows large. This description is couched in terms of solutions to optimal control problems. Combining this result with other techniques should provide exact characterizations of escape from, transitions among, and long-run selection of equilibria. We then study the large population double limit, in which the former limit is followed by taking the noise level to zero. Taking this second limit makes the control problems noted above much more amenable to analysis.

In order to solve the latter optimal control problems, we propose new results on solutions of HJB equations with state constraints. These results should allow the control problems to be solved in far more environments than is possible using the existing theory.

The second major goal of the project is to develop computational tools to study evolutionary game dynamics. Doing so will allow us to understand how these dynamics behave away from the limiting regimes in which theorems are usually proved. The software also allows for rapid explorations and simulations of evolutionary game dynamics, and is presented in a simple graphical user interface.

The third major goal of the project is to investigate new evolutionary game dynamics, especially ones uncovered by computational work. These dynamics can extend those classes for which equilibrium convergence and stability results are known to hold, or, in the opposite direction, can exhibit novel stability and instability properties that address longstanding questions in game theory.

Accomplishments: This project lead to six publications and one additional submitted paper. The two papers listed next solve various problems related to large deviations under evolutionary game dynamics:

[1] "Sample Path Large Deviations for Stochastic Evolutionary Game Dynamics" (with Mathias Staudigl). Mathematics of Operations Research 43 (2018), 1348-1377.

Abstract: We study a model of stochastic evolutionary game dynamics in which the probabilities that agents choose suboptimal actions are dependent on payoff consequences. We prove a sample path large deviation principle, characterizing the rate of decay of the probability that the sample path of the evolutionary process lies in a prespecified set as the population size approaches infinity. We use these results to describe excursion rates and stationary distribution asymptotics in settings where the mean dynamic admits a globally attracting state, and we compute these rates explicitly for the case of logit choice in potential games.

[2] "Hamilton-Jacobi Equations with Semilinear Costs and State Constraints, with Applications to Large Deviations in Games" (with Hung Tran and Srinivas Arigapudi). Submitted to Mathematics of Operations Research.

Abstract: We characterize solutions of a class of time-homogeneous optimal control problems with semilinear running costs and state constraints as maximal viscosity subsolutions to Hamilton-Jacobi equations, and show that optimal solutions to these problems can be constructed explicitly. We present applications to large deviations problems arising in evolutionary game theory.

The following two papers report on software developed for computation and simulations of evolutionary game dynamics.

[3] "EvoDyn-3s: A Mathematica Computable Document to Analyze Evolutionary Dynamics in 3-Strategy Games (with Luis Izquierdo and Segis Izquierdo). SoftwareX 7 (2018), 226-233.

Abstract: EvoDyn-3s generates phase portraits of evolutionary dynamics, as well as data for the analysis of their equilibria. The considered evolutionary dynamics are ordinary differential equations based on adaptive processes taking place in a population of players who are randomly and repeatedly matched in couples to play a 2-player symmetric normal-form game with three strategies. EvoDyn-3s calculates the rest points of the dynamics using exact arithmetic, and represents them. It also provides the eigenvalues of the Jacobian of the dynamics at the isolated rest points, which are useful to evaluate their local stability. The user only needs to specify the 3×3 payoff matrix of the game and choose the dynamics.

[4] "An Introduction to ABED: Agent-Based Simulation of Evolutionary Game Dynamics" (with Luis Izquierdo and Segis Izquierdo). Games and Economic Behavior 118 (2019), 434-462.

Abstract: ABED is free and open-source software for simulating evolutionary game dynamics in finite populations. We explain how ABED can be used to simulate a wide range of dynamics considered in the literature and many novel dynamics. In doing so, we introduce a general model of revisions for dynamic evolutionary models, one that decomposes strategy updates into selection of candidate strategies, payoff determination, and choice among candidates. Using examples, we explore ways in which simulations can complement theory in increasing our understanding of strategic interactions in finite populations.

The final three papers study a variety of novel evolutionary game dynamics. Articles [6] and [7] were directly inspired by work using the software developed in [4].

[5] “Riemannian Game Dynamics” (with Panayotis Mertikopoulos). Journal of Economic Theory 177 (2018), 315-364.

Abstract: We study a class of evolutionary game dynamics defined by balancing a gain determined by the game’s payoffs against a cost of motion that captures the difficulty with which the population moves between states. Costs of motion are represented by a Riemannian metric, i.e., a state-dependent inner product on the set of population states. The replicator dynamics and the (Euclidean) projection dynamics are the archetypal examples of the class we study. Like these representative dynamics, all Riemannian game dynamics satisfy certain basic desiderata, including positive correlation, local stability of interior ESSs, and global convergence in potential games. When the underlying Riemannian metric satisfies a Hessian integrability condition, the resulting dynamics preserve many further properties of the replicator and projection dynamics. We examine the close connections between Hessian game dynamics and reinforcement learning in normal form games, extending and elucidating a well-known link between the replicator dynamics and exponential reinforcement learning.

[6] “Best Experienced Payoff Dynamics and Cooperation in the Centipede Game” (with Segis Izquierdo and Luis Izquierdo). Theoretical Economics 14 (2019), 1347-1386.

Abstract: We study population game dynamics under which each revising agent tests each of his strategies a fixed number of times, with each play of each strategy being against a newly drawn opponent, and chooses the strategy whose total payoff was highest. In the centipede game, these best experienced payoff dynamics lead to cooperative play. When strategies are tested once, play at the almost globally stable state is concentrated on the last few nodes of the game, with the proportions of agents playing each strategy being largely independent of the length of the game. Testing strategies many times leads to cyclical play.

[7] “Stability for Best Experienced Payoff Dynamics” (with Luis R. Izquierdo and Segis S. Izquierdo). Journal of Economic Theory (2020), 104957.

Abstract: We study a family of population game dynamics under which each revising agent randomly selects a set of strategies according to a given test-set rule; tests each strategy in this set a fixed number of times, with each play of each strategy being against a newly drawn opponent; and chooses the strategy whose total payoff was highest, breaking ties according to a given tie-breaking rule. These dynamics need not respect dominance and related properties except as the number of trials become large. Strict Nash equilibria are rest points but need not be stable. We provide a variety of sufficient conditions for stability and for instability, and illustrate their use through a range of applications from the literature.

Training Opportunities: This grant provided funding for one research assistant. I hired graduate student Srinivas Arigapudi for the duration of the grant. I worked closely with Arigapudi during this time and our interactions lead to a number of working papers related to the topics in this proposal.

Results Dissemination: Work was disseminated through publications as described above and through a number of public seminars.

Honors and Awards: In 2019 I was named a Fellow of the Game Theory Society.

Protocol Activity Status:

Technology Transfer: Nothing to Report

PARTICIPANTS:

Participant Type: Graduate Student (research assistant)
Participant: Srinivas Arigapudi
Person Months Worked: 9.00 **Funding Support:**
Project Contribution:

International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

Nothing to report in the uploaded pdf (see accomplishments).