# The Resilience Game: A New Formalization of Resilience for Groups of Goal-Oriented Autonomous Agents

Extended Abstract

Michael A. Goodrich
Brigham Young University
Computer Science Department
Provo, Utah, USA
mike@cs.byu.edu

Jennifer Leaf
Julie A. Adams
CoRIS Institute
Oregon State University
Corvallis, OR, USA
leafj,julie.a.adams@oregonstate.edu

Matthias Scheutz
Department of Computer Science
Tufts University, Medford, MA, USA
Matthias.Scheutz@tufts.edu

## ABSTRACT

Groups of autonomous robots should be resilient. They should have the ability to cope with unknown events, long-lasting alterations to the environment, degradation of capacities, robot losses, and changes to communication networks. This paper presents a multiagent resilience formulation for goal-based agents. The formulation applies to mixed motive groups where agent goals have commonalities but are not perfectly aligned. Resilient groups must not only be resilient to chance exogenous perturbations but also intentional endogenous perturbations among the agents. Defining resilience using expected utilities leads to a new way of looking at multiagent resilience, namely the *resilience game.* The resilience game makes it possible to use the notion of equilibrium from game theory to evaluate how the intentional stances of agents determine when multiagent algorithms are resilient. A guided diffusion of innovations problem is used to demonstrate how the resilience game provides insight into the effectiveness of various joint algorithms.

## KEYWORDS

Autonomous agents; resilience; multiple agents.

**ACM Reference Format:**
Michael A. Goodrich, Jennifer Leaf, Julie A. Adams, and Matthias Scheutz. 2023. The Resilience Game: A New Formalization of Resilience for Groups of Goal-Oriented Autonomous Agents: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

## 1 INTRODUCTION

Robots should exhibit resilience, which is the ability to continue successful operations even when the robots are subjected to unanticipated conditions. Defining resilience for a group of autonomous agents is challenging when the goals of the agents do not perfectly align. When agents have mixed motives, such as when they have different achievement goals *achievement goals* [24], an exogenous perturbation can benefit some agents but not others. An agent that is negatively affected by a perturbation might have an incentive to change its behavior, driving the system from equilibrium.

There are limitations in previous definitions of resilience that make them difficult to apply to multiagent groups. *Engineering resilience* defines resilience as the ability to maintain a desired property (e.g., stability, performance threshold) or return to a desired equilibrium state after a perturbation [2, 8, 9, 13, 14, 19, 20]. *Switching resilience* is the ability to (re)establish stability, perhaps around a new equilibrium state, when a different set of operating parameters or a new set of goals arise [10]. A *viability kernel* defines resilience as a system's ability to stay within bounded region of the state space [1]. The viability framework assumes only that the system's goal is reachable and does not impose the existence of a baseline or equilibrium condition. *Ecological resilience* defines resilience as the ability to retain a system's *identity* and usually applies to natural multiagent phenomena such as predator-prey systems that do not have fixed stable points in their dynamics [4, 5, 11, 12, 15].

## 2 THE RESILIENCE GAME

There is a tension among the assessments made by goal-driven agents when their goals do not align; a joint algorithm that is good for one agent might not be good for another. This tension leads to a set of mathematical games collectively called the *resilience game.* The resilience game accounts for how *exogenous perturbations* can induce a cascade of *intentional perturbations* when agents have incentives to change behaviors. Given a model of world uncertainty, the expected utility of any joint algorithm can be obtained. Joint algorithms include both cooperative algorithms and algorithms in which an agent unilaterally deviates from the joint algoirthm.

A mathematical game can be used to model the incentive structure associated with whether an agent has an incentive to unilaterally deviate from a joint algorithm. Exogenous perturbations alter the payoffs for the joint algorithms. The *resilience game* is a meta-game that contains one game for each exogenous perturbation. An algorithm is an equilibrium of the resilience game if it is an equilibrium for each exogenous perturbation.

## 3 CASE STUDY

***World Model.*** The world in the *guided diffusion of innovations* problem is a network of nodes where each node represents an individual who might adopt a particular technology[3, 6, 16, 17]. Each edge in the network represents a relationship between individuals. This paper refers to the individuals in the network as *adopters* to differentiate them from *agents* who are coordinating their joint algorithm. Experiments used an assortative network with 80 adopters constructed using the assortativity mixing algorithm [18].

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.* 

Adopters choose to adopt a new "technology" from the abstract set $\{b, m, c\}$, indicating blue, magenta, and cyan. The adopters are influenced in their decisions by two forces: how many of their neighbors have adopted a technology at time $t$ and whether or not one of the agents is "making a sales pitch" at time $t$. The state of the world at time $t$ is the set of adopters, the colors of the adopters $\{b, m, c\}$, and the edges connecting the adopters. The initial state consists of all agents uncommitted and three *early adopters* having chosen unique initial technologies.

Adopters make decisions based both on the actions of the agents and on prior decisions made by their neighbors. Two decision rules from the literature are applied [3, 6]: In the *absolute threshold* rule, an agent adopts a new technology if $\kappa_{\text{fixed}} \in \{1, 2, 3\}$ or more of its neighbors have adopted. In the *fractional threshold* rule, an agent adopts a new technology if a portion $\kappa_{\text{fraction}} \in \{0.15, 0.25, 0.35, 0.45\}$ of its neighbors have adopted. A sales agent is considered one of the adopter's neighbors if the agent is contacting the adopter, but adopters can make decisions based only on their neighbors when no agent is in contact.

***Agent Model.*** A "sales manager" assigns each agent both a technology color and one of the four unique communities detected by the betweenness algorithm [7]. Each *early adopter* resides in the community assigned to the corresponding sales agent. Coordinating sales agents can *Cooperate* by "selling" to the neighbor of an adopter with largest degree within their assigned area or *TeamUp* by rotating through each others areas coordinating their "sales calls" to convince more agents to adopt more rapidly. Each agent also has a *Defect* algorithm in which the agent can sell to any adopter.

***Experiment Design and Results.*** A series of simulation experiments were performed under various conditions. The independent variables are the set of intentional and exogenous perturbations. The joint algorithms are *Cooperate*, *TeamUp*, all agents *Defect*, and *CooperateDefect* and *TeamUpDefect*, in which the cyan agent defects and the other agents play Cooperate or TeamUp, respectively. The exogenous perturbation occurs when adopter agents change their decision rules varying across the seven conditions for $\kappa_{\text{fixed}}$ and $\kappa_{\text{fraction}}$.

The *baseline world conditions* used the FixedThreshold adopter strategy with $\kappa_{\text{fixed}} = 2$. Figs. 1(a)–(b) show the expected utility of the cyan agent on the $x$-axis and the sum of the utility of the magenta and blue agent on the $y$-axis. The diagonal line represents *efficiency*, which occurs when all adopters select a technology. The vertical line is a representative satisficing threshold. The *fair allocation* is for each agent to receive an expected utility of 33%.

Agents seeking a Bayes-Nash [21] equilibrium always have a unilateral incentive to defect from *Cooperate* and *TeamUp* for all exogenous perturbations. The payoff vector is inefficient when all agents defect. For difficult perturbations, the incentive to defect for Bayes-Nash agents yields very inefficient results.

By contrast, agents seeking a satisficing equilibrium [22, 23] have no unilateral incentive to defect from the Cooperate or TeamUp joint solutions because their individual payoffs exceed the aspiration level. Even when an exogenous perturbation occurs and payoffs drop for the TeamUp strategy, satisficing agents have no incentive to change because their payoffs are still satisficing. When

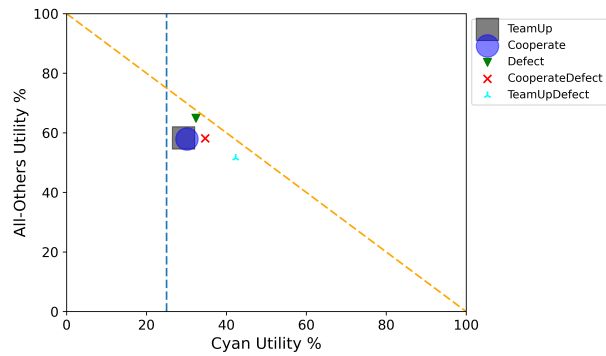


**(a) Intentional deviation from baseline under condition StrongAssortative. The WeakAssortative pattern is similar.**

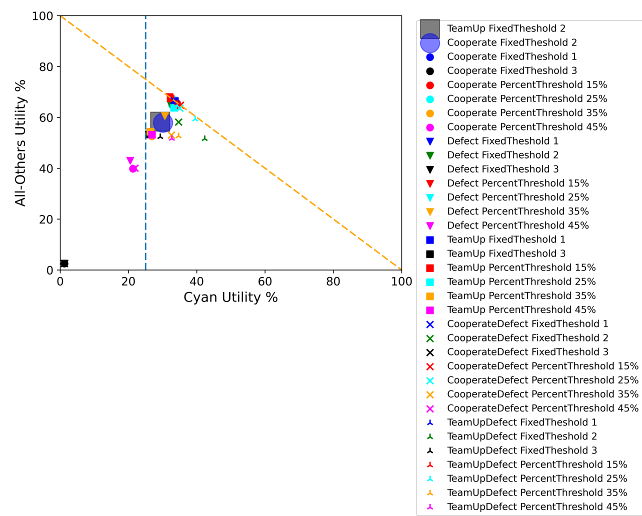


**(b) Perturbations to the adopter strategies for the StrongAssortative condition. The WeakAssortative pattern is similar.**

**Figure 1: Resilience game example.**

satisficing agents are using the Cooperate joint strategy, the difficult worlds ($\kappa_{\text{fixed}} = 3$ and $\kappa_{\text{fraction}} = 0.45$) induce an incentive to change their strategies, but not by defecting since defecting does cause the agents to obtain payoffs about the aspiration level.

## 4 SUMMARY

The resilience game helps make it clear that whether or not a joint strategy is resilient to exogenous and intentional perturbations depends on the intentional stance of the agents. The game shows that cooperating Bayes-Nash agents are not resilient for the guided-diffusion of innovations game because individual defections lead to poor payoffs when exogenous perturbations are challenging. Cooperating satisficing agents are more likely to be resilient to both exogenous and intentional perturbations if the aspiration level is not too high. This points to an important and interesting tension between seeking optimal solutions versus seeking resilient solutions.

## ACKNOWLEDGMENTS

This work was supported by the U.S. Office of Naval Research (N00014-18-1-2831). The authors thank two undergraduate research assistants, Elijah Pettitt and Nathan Riches, for running the simulation experiments.

## REFERENCES

[1] Christophe Béné and Luc Doyen. 2018. From Resistance to Transformation: A Generic Metric of Resilience Through Viability. *Earth's Future* 6, 7 (2018), 979–996.

[2] Fridolin Simon Brand and Kurt Jax. 2007. Focusing the Meaning(s) of Resilience: Resilience as a Descriptive Concept and a Boundary Object. *Ecology and Society* 12, 1 (2007). https://doi.org/10.5751/ES-02029-120123

[3] Damon Centola, Víctor M. Eguíluz, and Michael W. Macy. 2007. Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications* 374, 1 (2007), 449–456. https://doi.org/10.1016/j.physa.2006.06.018 arXiv:0504165 [physics]

[4] Graeme S Cumming and John Collier. 2005. Change and identity in complex systems. *Ecology and Society* 10, 1 (2005). https://doi.org/10.5751/ES-01252-100129

[5] Graeme S Cumming and Garry D Peterson. 2017. Unifying Research on Social-Ecological Resilience and Collapse. *Trends in Ecology & Evolution* 32, 9 (2017). https://doi.org/10.1016/j.tree.2017.06.014

[6] David Easley and Jon Kleinberg. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world.* Cambridge university press.

[7] Michelle Girvan and Mark EJ Newman. 2002. Community structure in social and biological networks. *Proceedings of the national academy of sciences* 99, 12 (2002), 7821–7826.

[8] V. Grimm and Christian Wissel. 1997. Babel, or the ecological stability discussions: An inventory and analysis of terminology and a guide for avoiding confusion. *Oecologia* 109, 3 (1997), 323–334. https://doi.org/10.1007/s004420050090

[9] Dave Hodgson, Jenni L. McDonald, and David J. Hosken. 2015. What do you mean, 'resilient'? *Trends in Ecology & Evolution* 30, 9 (2015), 503–506. https://doi.org/10.1016/J.TREE.2015.06.010

[10] Robert R. Hoffman and P. A. Hancock. 2017. Measuring Resilience. *Human Factors* 59, 4 (2017), 564–581. https://doi.org/10.1177/0018720816686248

[11] C. S. Holling. 1973. Resilience and Stability of Ecological Systems. *Annual review of ecology and systematics* 4 (1973), 1–23. https://doi.org/10.2307/2096802

[12] C S Holling. 1996. Engineering Resilience versus Ecological Resilience. In *Engineering Within Ecological Constraints.* The National Academies Press, Washington, D.C., 31–43. https://doi.org/10.17226/4919

[13] Seyedmohsen Hosseini, Kash Barker, and Jose E. Ramirez-Marquez. 2016. A review of definitions and measures of system resilience. *Reliability Engineering and System Safety* 145 (2016), 47–61. https://doi.org/10.1016/j.ress.2015.08.006 arXiv:0809.5001

[14] Johannes Ingrisch and Michael Bahn. 2018. Towards a Comparable Quantification of Resilience. *Trends in Ecology & Evolution* 33, 4 (2018), 251–259. https://doi.org/10.1016/J.TREE.2018.01.013

[15] Richard C Lewontin. 1969. The meaning of stability. In *Brookhaven symposia in biology*, Vol. 22. 13–24.

[16] Nigel Meade and Towhidul Islam. 2006. Modelling and forecasting the diffusion of innovation–A 25-year review. *International Journal of forecasting* 22, 3 (2006), 519–545.

[17] Mark Newman. 2018. *Networks.* Oxford university press.

[18] Mark EJ Newman and Michelle Girvan. 2003. Mixing patterns and community structure in networks. In *Statistical Mechanics of Complex Networks.* Springer, 66–87.

[19] Stuart L. Pimm. 1984. The complexity and stability of ecosystems. *Nature* 307, 5949 (1984), 321–326. https://doi.org/10.1038/307321a0

[20] Allyson E. Quinlan, Marta Berbés-Blázquez, L. Jamila Haider, and Garry D. Peterson. 2016. Measuring and assessing resilience: Broadening understanding through multiple disciplinary perspectives. *Journal of Applied Ecology* 53, 3 (2016), 677–687. https://doi.org/10.1111/1365-2664.12550

[21] Y. Shoham and K. Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations.* Cambridge University Press.

[22] Herbert A Simon. 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics* 69, 1 (1955), 99–118.

[23] Jeff L Stimpson and Michael A Goodrich. 2003. Learning to cooperate in a social dilemma: A satisficing approach to bargaining. In *Proceedings of the 20th international conference on machine learning (ICML-03).* 728–735.

[24] M. Birna Van Riemsdijk, Mehdi Dastani, and Michael Winikoff. 2008. Goals in agent systems: A unifying framework. In *Proc. 7th Int. Joint Conf. Autonomous Agents and Multiagent Systems*, Vol. 2. 702–709.