Hertie School Fall Semester 2025

Course Syllabus

GRAD-E1489: Algorithmic Game Theory and Governance

1. General information

Course Format	Onsite
Instructor(s)	Asya Magazinnik
Instructor's E-mail	a.magazinnik@hertie-school.org
Assistant	Amanda Slater; slater@hertie-school.org
Instructor's Office Hours	By appointment, Thursdays, 10-12. Please email Amanda Slater to schedule an appointment.

Link to Study, Examination and Admission Rules and MIA, MDS and MPP Module Handbooks

For information on **course room, times and session dates,** please consult the Course Schedule on *MyStudies*.

Instructor Information:

Asya Magazinnik is the Professor of Social Data Science at the Hertie School. Her research interests include electoral geography, federalism, local politics, and law enforcement. She also works on political methodology, in particular at the intersection of causal inference and formal theory. Her work has appeared in the American Journal of Political Science, the Journal of Politics, and other outlets. Previously, she was an Assistant Professor of Political Science at MIT. She earned a PhD in Political Science from Princeton University in 2020 and holds an MPP from the Harris School at the University of Chicago.

2. Course Contents and Learning Objectives

Course contents:

This course is about how the emergent field of algorithmic game theory is transforming society and governance. Advances at the nexus of computer science and game theory have been applied to learning in AI systems; to the design and governance of online platforms and cryptocurrencies; to improving the efficiency of energy markets and transportation networks; to the fair allocation of shared resources; and to matching organ donors to receivers. The course begins by teaching students the basic tools of game theory, including Nash equilibrium, backward induction, expected utility, static and dynamic games of complete and incomplete information, repeated games, mechanism design, and social choice theory. With every topic, students are introduced to algorithmic approaches as well as cutting-edge applications to society and governance.

Teaching style:

Each session will be a mix of some lecture on the game theoretic concepts and some in-depth discussion of the applied readings. Reading in advance and active participation in class discussion are essential to success.

Prerequisites:

Mathematics for Data Science is a prerequisite for this course. Students are expected to have a working knowledge of optimization, multivariate calculus, and probability theory.

I also expect a willingness to work hard on unfamiliar materials. While the Math for Data Science course will give students the minimal baseline for understanding the course materials, we will be reading some papers at the research frontier – which is a long way from the fundamentals. Students are expected to read closely and re-read texts, to be actively engaged in class and bring questions to work through together, and to consult the instructor for additional assistance.

The course will occasionally rely on some methods of mathematical proof (e.g., proof by contradiction, proof by induction), but prior exposure to them is not assumed.

Diversity Statement:

In this classroom and at the Hertie School more broadly, we respect and value differences in race, gender, gender identity, ethnicity, age, physical and language abilities, culture, religion, sexual orientation, and other identities. Such differences contribute to the strength of our community, and upholding an atmosphere of inclusion and mutual respect is of the utmost importance.

3. Grading and Assignments

Composition of Final Grade:

Assignment 1: 3 problem sets	Deadline: 8-10 days after distribution	Submit via Moodle	60% (20% each)
Assignment 2: Final Project	Deadline: December 16th	Submit via Moodle	30%
Assignment 3: Presentation	Deadline: In class (12th session)	In class	10%

Assignment Details

Assignment 1: Problem Sets

This class will have three problem sets. These assignments reinforce the basic game theoretic foundations we learn in class and may consist of solving problems by hand or in code. They will be distributed after sessions 3, 5, and 7.

Start assignments early. Please note that the problem sets are designed to take time and thought, and cannot be (successfully) completed the night before they are due. Please read and begin them as soon as they have been assigned in order to give you ample time to struggle with the problems on your own, seek help from each other and the instructor, and write up clear solutions.

I ask that all solutions be typed using your preferred LaTeX editor. Not only does this make the instructor's job of grading much easier, but the exercise of writing up your work clearly and understandably is an important skill in its own right for any practicing data scientist.

Please show all work. I should be able to trace every step of your process, whether it's a mathematical proof, a solution to a problem, or code. **I should be able to successfully run any of your code on my machine; replication is key!**

You may work alone or in groups or 2-3 people of your choosing. If you work in a group, you may submit one set of solutions for the group.

Assignment 2: Final Project

The final project for the course will be a 6-8 page paper on a topic of your choosing (approved in advance by the instructor):

1) **Discussing the policy implications of a topic.**
   Taking a topic from the syllabus (or beyond) – e.g., no-regret learning or a stable matching algorithm – discuss how the tool is currently being used in public policy, the potential it has to improve governance and society, and the potential risks of this technology.

2) **Understanding and implementing an approach.**
   Here, you may select a topic from the syllabus (or beyond) to carefully learn and implement on your own in code. You can draw from any chapter of the *Algorithmic Game Theory* text, Roughgarden's lecture notes, the *Machine Learning* special issue, or the arXiv Computer Science and Game Theory repository (all listed in the "General Readings" section below), among other sources.

Students are welcome to work in groups of 2-3 or alone – as they wish. Groups should submit one submission for the whole group.

Assignment 3: Presentation

In Session 12, each student (or group) will make a short presentation of their progress on the final project to the class.

Late submission of assignments: For each day the assignment is turned in late, the grade will be reduced by 10% (e.g. submission two days after the deadline would result in 20% grade deduction).

Attendance: Students are expected to be present and prepared for each class session. Active participation during lectures and seminar discussions is essential. Please note that students can miss up to two sessions (out of twelve) if no course assignments are affected. **Students do not need to**

inform the course instructor(s) in advance if they are unable to attend a session. For further information please consult the Examination Rules §10.

Academic Integrity: The Hertie School is committed to the standards of good academic and ethical conduct. Any violation of these standards shall be subject to disciplinary action. Plagiarism, misuse of AI, free riding in group work, and other deceitful actions are not tolerated. See Examination Rules §16, the Hertie Plagiarism Policy, and the Hertie Guidelines for Artificial Intelligence Tools.

Compensation for Disadvantages: If a student furnishes evidence that he or she is not able to take an examination as required in whole or in part due to disability or permanent illness, the Examination Committee may upon written request approve learning accommodation(s). In this respect, the submission of adequate certificates may be required. See Examination Rules §14.

Extenuating circumstances: An extension can be granted due to extenuating circumstances (i.e., for reasons like illness, personal loss or hardship, or caring duties). In such cases, please contact the course instructor and Examination Office *in advance* of the assignment deadline.

General Readings:

All readings are either available online, or the relevant portions will be uploaded to the Moodle site.

1. Gibbons, Robert. 1992. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press.
2. MIT OpenCourseWare (OCW) Game Theory Course: https://ocw.mit.edu/courses/17-810-game-theory-spring-2021/pages/lecture-slides/.
3. Nisan, Noam, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (eds.) 2007. *Algorithmic Game Theory*. New York: Cambridge University Press.
4. Roughgarden, Tim. 2016. *Twenty Lectures on Algorithmic Game Theory.* New York: Cambridge University Press.
5. Roth, Alvin E. 2016. *Who Gets What – and Why: The New Economics of Matchmaking and Market Design.* Mariner Books.
6. Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values.* W.W. Norton & Company.

Other References:

- Computer Science and Game Theory papers on arXiv (https://arxiv.org/list/cs.GT/recent)
- Machine Learning, Special Issue on Learning & Computational Game Theory (https://link.springer.com/journal/10994/volumes-and-issues/67-1)
- Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations (https://www.masfoundations.org/mas.pdf)
- Artificial Intelligence, Special Issue on Foundations of Multiagent Learning (https://www.sciencedirect.com/journal/artificial-intelligence/vol/171/issue/7)
- Ioannis Panageas' Intro to Algorithmic Game Theory Lecture Notes: (https://panageas.github.io/agt2022/)

Course Sessions and Readings

All course readings can be accessed on the course Moodle page.

Session 1: An Introduction to Rational Choice Theory	
Required Readings	Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* Oxford University Press. Chapter 1. Stanford Encyclopedia of Philosophy, "Normative Theories of Rational Choice: Expected Utility"
Application Readings	Shoham, Yoav. 2008. "Computer Science and Game Theory." *Communications of the ACM* 51(8).

Session 2: Games in Strategic Form and Nash Equilibrium	
Required Readings	Pure Strategies: Gibbons 1.1 Mixed Strategies: Gibbons 1.3.A Magazinnik OCW, Lecture 2-3
Application Readings	Daskalakis, Constantinos, Paul W. Goldberg, and Christos H. Papadimitriou. 2008. "The Complexity of Computing a Nash Equilibrium." Porter, Ryan, Eugene Nudelman, and Yoav Shoham. 2008. "Simple search methods for finding a Nash equilibrium." *Games and Economic Behavior* 63: 642-662.

Session 3: Dynamic Games of Complete Information	
Required Readings	Perfect Information: Gibbons 2.1A, 2.1B Imperfect Information: Gibbons 2.2A, 2.2B, 2.4 Magazinnik OCW, Lecture 4
Application Readings	Panageas Lecture Notes on No-regret learning Jacob, Athul Paul, Yikang Shen, Gabriele Farina, and Jacob Andreas. 2024. "The Consensus Game: Language Model Generation via Equilibrium Search." ICLR. "Game Theory Can Make AI More Correct and Efficient." 2025. Quanta Magazine.

Session 4: Repeated Games I: Foundations	
Required Readings	Finitely Repeated Games: Gibbons 2.3A Infinitely Repeated Games: Gibbons 2.3B Magazinnik OCW, Lecture 5
Application Readings	Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values.* W.W. Norton & Company. Chapter titled "Agency."

Session 5: Repeated Games II: Applications to Learning and Artificial Intelligence

Application Readings	Shoham, Yoav, Rob Powers, and Trond Grenager. 2007. "If multi-agent learning is the answer, what is the question?" *Artificial Intelligence* 171: 365-377. Foerster, Jakob N., Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2017. "Learning with Opponent-Learning Awareness." ArXiv. - See also: OpenAI, "Learning to Model Other Minds" Akata, Elif, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2025. "Playing Repeated Games with Large Language Models." *Nature Human Behaviour.*

Session 6: Static Games of Incomplete Information	
Required Readings	Gibbons 3 Magazinnik OCW, Lecture 6
Application Readings	Vorobeychik, Yevgeniy, Michael P. Wellman, and Satinder Singh. 2007. "Learning Payoff Functions in Infinite Games." *Machine Learning* 67: 145-168.

Session 7: Dynamic Games of Incomplete Information	
Required Readings	Gibbons 4.1, 4.2, 4.3A Magazinnik OCW, Lecture 7

Session 8: Algorithmic Mechanism Design	
Required Readings	Nisan, Noam. 2007. "Introduction to Mechanism Design (for Computer Scientists)." In *Algorithmic Game Theory*. Parkes, David C. 2007. "Online Mechanisms." In *Algorithmic Game Theory.* Roughgarden, Tim. 2016. *Twenty Lectures on Algorithmic Game Theory.* Lecture 2 (Mechanism Design Basics) and Lecture 4 (Algorithmic Mechanism Design).

Session 9: The Price of Anarchy	
Required Readings	Johari, Ramesh. 2007. "The Price of Anarchy and the Design of Scalable Resource Allocation Mechanisms." In *Algorithmic Game Theory.* Roughgarden, Tim. 2016. *Twenty Lectures on Algorithmic Game Theory.* Lecture 11 (Selfish Routing and the Price of Anarchy).

Session 10: Computational Social Choice

Required Readings	Magazinnik OCW, Lecture 8 Roughgarden, Tim. 2016. *Twenty Lectures on Algorithmic Game Theory.* Lecture 10 (Kidney Exchange and Stable Matching).
Application Readings	McElfresh, Duncan C., et al. 2023. "Matching Algorithms for Blood Donation." *Nature Machine Intelligence* 5: 1108-1118. Roth, Alvin C. 2015. *Who Gets What – And Why.* Boston: Houghton Mifflin Harcourt. Part III: "Design Interventions to Make Markets Smarter, Thicker, and Faster."

Session 11: Guest Lecture (TBD)	
Required Readings	TBD

Session 12: Student Presentations	
Required Readings	None