

# DS 574: Algorithmic Game Theory — Fall 2026

**Instructor:** Prof. Kira Goldner

**Email:** goldner@bu.edu

**Office Hours:** Tuesday 3:15-4:15PM and by appointment

**Office Location:** CCDS 1339, 665 Commonwealth Ave

**Lectures:** Tuesday and Thursday 2:00-3:15PM, Location TBD

**Course Description:** This course is an introduction to the interdisciplinary area of Algorithmic Game Theory: where computational perspectives are applied to economic problems, and economic techniques are brought to problems from computer science. We will explore a broad range of topics at the frontier of new research, starting with some of the fundamentals, such as welfare-maximizing auctions and types of Nash Equilibria. Throughout the semester, the class will also learn about prevalent topics such as (1) Data Science & Incentives, (2) Mechanism Design for Social Good, and (3) optimization and robustness in mechanism design. As part of this course, students will engage in a (guided) research project, experiencing the various parts of conducting original research.

The course is aimed at graduate students but will be accessible to motivated advanced undergraduate or masters students with the appropriate background.

**Learning Objectives:** The objectives of this course are to:

- **understand the core content:** understand the basic concepts underlying algorithmic game theory and mechanism design and how they may be applied to create interesting research directions.
- **learn about cutting edge research directions:** examples are at the intersection of algorithmic mechanism design and: economics, data science, policy, and more.
- **get hands-on problem-solving and critical-thinking experience:** this includes thinking critically about questions in mechanism design and synthesizing research in the area.

**Prerequisites:** **Strong proficiency** in DS 120, DS 121, and DS122 and DS 320 (or equivalent), or instructor's consent. Students should have:

- Mathematical maturity and comfort with formal proofs
- A solid understanding of probability (discrete and continuous random variables, moments, and conditional probability)
- Familiarity with algorithms and computational efficiency.

**Course website:** <https://www.kiragoldner.com/teaching/DS574/>. There will also be a Piazza website for the course: <https://piazza.com/bu/fall12026/ds574> (access code “AMD”), and Gradescope: <https://www.gradescope.com/courses/532743> (entry code ZZV4DV).

**BU CDS PhD Core:** This course satisfies Mathematical Foundations or Optimization Algorithms in the Methodology Core, as well as Theoretical Foundations or DS for Social & Behavioral Sciences in the Subject Core.

**BU CS PhD Algorithms & Theory Depth Requirement:** This course not yet approved but slated to be approved for the Computer Science Algorithms & Theory Depth Requirement.

**BU Hub:** This course satisfies **Social Inquiry I**.

**Social Inquiry I:** This course uses tools and concepts from economics such as game theory, behavioral economics, welfare, equilibrium, and much more in order to reason about how to allocate resources to individuals and how they will behave in response, both in interacting with the algorithm and in interacting with each other. These models in economics and particularly behavioral economics are also built on top of foundations in psychology.

**Feedback on Learning Outcomes:** Students will complete problem sets that require them to solve mechanism design problems, reasoning about the designer, the individuals, and how they all interact. The problem sets will also incorporate different economic theories and the ethical implications behind them. Students must be creative about their solutions. They will be encouraged to work collaboratively in small teams to learn from one another. They will receive prompt grading with feedback on their solutions, and thus on their creativity, and how to iterate from there. During lectures, students will complete in-class exercises on mechanism design problems. In-class exercises will include opportunities to observe ideas from other students as well as to hear feedback from other students and from the instructor. The mechanism design for social good problem formulation and the final project allow opportunities for students to experiment with creativity and iterate on it, as well as to make ethical choices based on what they have learned in the course.

**Books and Other Course Materials:** There are no required books. There will be suggested readings from various textbooks and lecture notes suggested along with the course.

**Assignments and Grading:**

- Problem sets: 30% in total.
- Mechanism Design for Social Good problem formulation: 15%.
- Final project: 45%.
- Class participation—in class and via piazza: 10%.

### Homework (30%):

- Expect to spend at least 10 hours per assignment on homework.
- **Late policy:** You may use up to 4 late days throughout the semester, but not more than 2 days on a given assignment. For each instance, you may only use an integer number of late days. Outside of this policy, no late submissions will be accepted.
- Your written assignments must be prepared with LaTeX, not handwritten.
- You must hand in your homework via Gradescope, which will be due at 11:59pm on the day assigned.
- **Regrade requests:** Regrade requests must be submitted within 7 days of receiving the graded assignment and only via Gradescope. You must also submit an explanation detailing which problems were graded incorrectly and an argument that the submitted solution is indeed correct. Regrades may only be requested if it is believed that a correct answer was marked as incorrect, not because insufficient partial credit was given to an incorrect or partially correct solution. If you request a regrade, you accept that the entire assignment/exam will be regraded, not just the problem(s) believed to be graded incorrectly.

### Homework Collaboration Policy:

- For many people, algorithmic problem-solving is a collaborative endeavor. As such, you may work with up to two other classmates on the bi-monthly homeworks for the course. However, the assignments you hand in must be written up by yourself and represent your own thoughts and work. In particular, you may discuss ideas with your classmates in person, but as a rough rule, nobody should leave the room with anything written down. If you really understand the discussion, you should be able to reconstruct it on your own. **As a hard rule, you must write up your arguments and problem sets individually. You may not use the internet or other references other than the course materials, unless told otherwise. Please see the course Generative AI policy.**
- You must write your collaborators' names on the top of your assignment. Crediting one's peers is an important habit. If you do not work with collaborators, list "Collaborators: None."
- Finally, make sure you adhere to BU's academic conduct policy, which I take very seriously: <https://www.bu.edu/academics/policies/academic-conduct-code/>. In particular, homework making up part of your grade is a gift: if you are caught cheating on homework, you will be reported to the Academic Conduct Committee, and the homework will be down-weighted to zero. In this case, in-person quizzes will take the place of homework for you.

**Mechanism Design for Social Good Problem Formulation (15%):** Students will choose a social good domain and will formulate a mechanism design question in this area. You must defend why this is an important problem within this domain using domain-related sources, and will defend why this is an important problem for mechanism design to solve using what we have learned in the course. You will also identify an ideal domain expert with whom you would hope to collaborate if you were to pursue this project.

**Final Project (45%):** You will choose a research question that we have not covered in class to investigate. You will either learn to what extent it has been solved (by reading a research paper or lecture notes) or will attempt to do original research on it. You will then present what you have learned to the class. To aid in the experience of conducting research for those who may not have engaged in research before, the course will contain several discussions and check points on how to complete various key steps:

- choosing a high-level topic,
- identifying literature to review,
- completing a literature review,
- writing up motivation and the literature review,
- identifying a concrete question to investigate,
- making progress on the concrete question,
- putting together a high-quality intuition-focused presentation and giving it, and
- putting together a high-quality report.

See the schedule of lectures for the list of discussions and the project description and rubric on the course website for what is expected.

**Participation (10%):** In-class participation will be graded via participation cards every two weeks. Each student will receive a participation card with their name on it, and when they participate in good faith during class, the instructor will collect the card (counting their participation for this period) and return it at the end of the two-week period. Class attendance will not otherwise be taken. Piazza statistics will be used to supplement in-class participation.

**Curves:** Each individual assignment is graded with its statistics in mind, with some type of curve potentially being applied, although it may not be visible to students. As a result, overall grades follow the exact formula laid out here. No overall curve will be applied at the end of the course.

**Reasonable Accommodations:** If you are a student with a disability or believe you might have a disability that requires accommodation, please contact the Office for Disability Services (ODS) at (617) 353-3658 or [access@bu.edu](mailto:access@bu.edu) to coordinate any reasonable accommodation requests. ODS is located at 25 Buick Street on the 3rd floor. If you have an accommodation, please report it to me within the first two weeks of the semester.

**Generative AI Policy:** I do not use LLMs in the design or preparation of course materials. You are not permitted to use LLMs to solve homework or generate writing for class assignments. You may use them like a search engine: to point you toward other resources (papers, LaTeX commands, etc.). If you are unsure whether LLM use is permitted in your use case, ask, otherwise, the assumption is that it is not permitted. Violation of this policy is a violation of BU's academic integrity policy.

## Tentative Outline of Class Meetings

- Lecture 1: Overview and Policies, Intro to AGT
- Lecture 2: Projects, Incentive Compatibility, & The Revelation Principle
- Lecture 3: Bayesian Settings and Revenue Equivalence
- Lecture 4: Myersonian Virtual Welfare and Quantile Space
- Lecture 5: Ironing Virtual Values
- Lecture 6: Multidimensional Settings and VCG, **Discussion:** Project Ideas
- Lecture 7: Ascending Auctions I
- Lecture 8: Ascending Auctions II & Walrasian Equilibria
- Lecture 9: Prophet Inequalities and Balanced Prices
- Lecture 10: KVV, Prior Independence: Bulow-Klemperer & Single Sample
- Lecture 11: Gains from Trade in Two-Sided Markets
- Lecture 12: **Discussion:** Starting research and making progress
- Lecture 13: Mechanism Design for Social Good I: Health Insurance Markets
- Lecture 14: Mechanism Design for Social Good II: Kidney Exchange
- Lecture 15: Mechanism Design for Social Good III: Democracy, Summary of MD4SG Directions
- Lecture 16: Interdependent Values I
- Lecture 17: Interdependent Valuations II
- Lecture 18: Behavioral Economics and Mechanism Design I
- Lecture 19: Behavioral Economics II
- Lecture 20: **Discussion:** How to speak and write
- Lecture 21: Fair Division
- Lecture 22: Machine Learning and Incentives
- Lecture 23: Collage of Topics (Blockchain, Prediction Markets, Cap and Trade, Tournament Design)
- Lecture 22: Project Presentations
- Lecture 23: Project Presentations
- Lecture 24: Project Presentations