

# Seminar on Incentives and Learning

## (Winter 2022-23)

Hi and welcome to the seminar! We meet every Monday 12:30-14:20. Every week, three students will have 90 minutes to present a paper to the class. In the remaining 10 minutes we will have a discussion led by student “reviewers” of the paper.

Course requirements:

- Main task: Your main task is to carefully read, understand, and present a paper to the class with your partners. Grading is individual.
  - a. After your presentation please upload your slides (e.g. to Dropbox) and share a link in the table below (see the first row for reference). The slides may be used for reference by the other students, and will be used for grading.
- Secondary task (mandatory): Read one additional paper and serve as a “reviewer” (reviewing the paper, not its presentation) - write a 1-page review of the paper, and lead the 10-20 minute discussion of it in class. Submit the review via [Google Form](#) before the meeting in which the paper is discussed. The format will be similar to peer reviews in top conferences.
- Attendance: You may miss up to one meeting on weeks #2 to #13 (in case of special circumstances like “miluim”/sickness/etc. please email me). It is possible (but not advised) to participate by [Zoom link](#); Meeting ID: 955 2126 3682; Passcode: bandit
- Bonus points for helpful participation in class!

Presentation requirements (also used as grading criteria):

- Situate the paper you’re presenting within the past and present literature - don’t forget to survey follow-up work.
  - a. Look for the journal version of your paper and for online talks given by the authors.
- Please make sure to pick one or more main technical results or proof techniques and explain the technical ideas behind them. You may also pick a technical topic that comes up in the paper (for example, VC dimension, Gittins index, etc.) and teach it to the class.
  - a. There is no need to cover in your talk all the results in the paper.
- Critical thinking: Come up with directions for future research and/or identify weaknesses in the paper, both as a speaker and as a reviewer.
- Please practice your presentation to make sure it is streamlined, clear and engaging.
  - a. Use figures and examples wherever possible.
  - b. Please number your slides.

Materials and books: See below.

Paper allocations:

- ~~This Google document will be open to editing on Sunday Oct 30 at 17:00 PM, please sign yourself up twice - as a presenter and a reviewer.~~
- ~~In extreme cases (like the existence of a slot with a single presenter) I may have to re-assign you - so please wait for confirmation before you start working.~~
- ~~Needless to say, please respect your fellow students and do not modify names. Unethical behavior will not be accepted.~~

| Date   | Paper  | Topic   | Presenter(s)  | Discussant(s)  |
|--|--|---|---|--|
| #1<br>Oct 24<br><a href="#">link</a>                           | -  | Introduction to the seminar + guest lecture on "Auctions in the Wild" | 1. ענבל טלגם-כהן<br>2. אופיר פרידלר,<br>חברת Outbrain | -  |
| #2<br>Oct 31<br><a href="#">link</a>                           | P. Dhangwatnotai, T. Roughgarden and Q. Yan. <b>Revenue Maximization with a Single Sample.</b> Games and Economic Behavior 2015. | Introduction to mechanism design and sample complexity                | 1. שגיא לוי<br>2. יונתן מדן                           | -  |
| #3<br>Nov 7<br><a href="#">link</a>                            | Braverman et al. <b>Selling to a no-regret buyer.</b> EC'18.   | Mechanism design with learning agent                                  | 1. עידו קולמן<br>2. דורין שטיימן<br>3. שני גורן       | 1. גיא הורוביץ<br>2. יונתן סומר<br>3. תום אזולאי                       |
| #4<br>Nov 14<br><a href="#">link</a>                           | Kolumbus and Nisan. <b>Auctions between regret-minimizing agents.</b> WWW'22.  | Mechanism design with learning agents                                 | 1. אור מרקובצקי<br>2. ליאור מוטולה<br>3. דניאל יחזקאל | 1. איתן בלור<br>2. יואב יעבץ<br>3. גלעד שמרלר                          |
| #5<br>Nov 21<br><a href="#">link</a>                           | Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C. Parkes. <b>Optimal Auctions through Deep Learning.</b> ICML'19.         | Mechanism design via deep learning                                    | 1. חן דוידוב<br>2. תומר מלניק<br>3. שליו ריסין        | 1. shani goren<br>2. ido kolman<br>3. dorin shteyman<br>4. Liran Cohen |
| #6<br>Nov 28<br><a href="#">link</a><br><a href="#">(pptx)</a> | M. Hardt, N. Megiddo, C. Papadimitriou and M. Wootters. <b>Strategic Classification.</b> ITCS'16.                                | Strategic classification - gaming                                     | 1. איתן בלור<br>2. יואב יעבץ<br>3. גלעד שמרלר         | 1. אור מרקובצקי<br>2. ליאור מוטולה<br>3. דניאל יחזקאל                  |
| #7<br>Dec 5<br><a href="#">link</a><br><a href="#">link</a>    | J. Kleinberg, S. Mullainathan and M. Raghavan. <b>How Do Classifiers Induce Agents To Invest Effort Strategically?</b> EC'19.    | Strategic classification - incentivizing effort                       | 1. דניאל בלינקי<br>2. נעם משה<br>3. אלון זיו          | 1. אורן רוזנברג<br>2. תומר מלניק<br>3. שליו ריסין                      |

|   |   |  |   |   |
|---|---|--|---|---|
| #8<br>Dec 12<br><a href="#">link</a>                        | Sundaram et al.<br><b>PAC-Learning for Strategic Classification.</b><br>ICML'21.  | Strategic classification -<br>PAC learning                                     | 1. נטע שור<br>2. סהר אלבז<br>3. רחלי צפובצקי          | 1. יואב האריס<br>2. נסים יחיא<br>3. אלון זיו                        |
| Dec 19  | No class  |  |   |   |
| #9<br>Dec 26<br><a href="#">link</a>                        | Bacon et al. <b>Predicting Your Own Effort.</b><br>AAMAS'12.<br>Waggoner. <b>Proper Scoring Rules.</b> Blog post from "The Tiger's Stripes" 2018. | Scoring rules: How can we incentivize an agent to make an accurate prediction? | 1. ניב גלעדי<br>2. רועי מחפוד<br>3. אורן כץ           | 1. רחלי צפובצקי<br>2. תמיר שור                                      |
| #10<br>Jan 2<br><a href="#">link</a>                        | Bates et al.<br><b>Principal-agent hypothesis testing.</b><br>Working paper.  | Hypothesis testing   | 1. תום אזולאי<br>2. דוד ולנסי<br>3. לירן כהן          | 1. יונתן מדן<br>2. נעם משה<br>3. דניאל בלינקי<br>4. סהר אלבז        |
| #11<br>Jan 9<br><a href="#">link</a><br><a href="#">lab</a> | Arora et al.<br><b>Generalization and Equilibrium in Generative Adversarial Nets (GANs).</b> ICML'17.   | GANs   | 1. אורן רוזנברג<br>2. שניר גרין<br>3. תמיר שור        | 1. ברק גחטן<br>2. שגיא לוי<br>3. נטע שור                            |
| #12<br>Jan 16<br><a href="#">link</a>                       | Blum et al. <b>One for One, or All for All: Equilibria and Optimality of Collaboration in Federated Learning.</b><br>ICML'21.                     | Federated learning   | 1. Andrew Elashkin<br>2. גיא הורוביץ<br>3. יונתן סומר | 1. אורן כץ<br>2. חן דוידוב<br>3. ניב גלעדי                          |
| #13<br>Jan 23<br><a href="#">link</a>                       | J. Kleinberg, S. Mullainathan and M. Raghavan. <b>Inherent Trade-Offs in the Fair Determination of Risk Scores.</b> ITCS'17.                      | Fairness   | 1. ברק גחטן<br>2. יואב האריס<br>3. נסים יחיא          | 1. Andrew Elashkin<br>2. שניר גרין<br>3. רועי מחפוד<br>4. דוד ולנסי |

# Materials on AGT and learning

Courses, tutorials, talks etc.:

- U. Chicago Fall 2022 course: <https://www.haifeng-xu.com/cmssc35401fa22/index.htm>
- MIT Spring 2019 course: <http://vsyrgkanis.com/6853sp19/>
- TAU Fall 2018 course: <http://advanced-topics-ml-agt-tau-2018.wikidot.com/>
- NeurIPS'21 workshops: <https://www.strategic-ml.com/>  
<https://sites.google.com/view/strategicml/about>
- EC'19 workshop: [Learning in Presence of Strategic Behavior](#)
- [EC'17 tutorial](#) on incentivizing and coordinating exploration
- [EC'16 tutorial](#) on algorithmic game theory and data science
- Dagstuhl 2017 seminar:  
<https://www.dagstuhl.de/de/programm/kalender/semhp/?seminr=17251>
- Papadimitriou's 2015 talk at the Simons Institute:  
<https://simons.berkeley.edu/talks/christos-papadimitriou-2015-11-20>
- 

Books:

1. [Twenty Lectures on Algorithmic Game Theory](#), by Tim Roughgarden, Cambridge University Press, 2016.
2. [Mechanism Design and Approximation](#), by Jason Hartline.
3. [Algorithmic Game Theory](#), by Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay V. Vazirani (eds.), Cambridge University Press, 2007.
4. [Online Learning and Online Convex Optimization](#), by Shai Shalev-Shwartz, Foundations and Trends in Machine Learning, 2011.
5. [Understanding Machine Learning: From Theory to Algorithms](#), by Shai Shalev-Shwartz and Shai Ben-David, Cambridge University Press, 2014.
6. [Prediction, Learning, and Games](#), by N. Cesa-Bianchi and G. Lugosi, Cambridge University Press, 2006.
7. [Introduction to Multi-Armed Bandits](#), by A. Slivkins, Foundations and Trends in Machine Learning, 2019.

See also on [my homepage](#) under "Recommended links for students" the links "How to read a paper" and "Tim Roughgarden's AGT courses".