ORIE 4154/5154: Pricing and Market Design

Optimization, learning, and incentives in modern markets

Fall 2026 Syllabus

> *Markets are human artifacts—they are tools that we build ourselves to help us [...] a market that can operate freely is like a wheel that can turn freely: it needs an axle and well-oiled bearings. How to provide that axle and keep those bearings well oiled is what market design is about.*
>
> Alvin Roth, "Remarks on Market Design" (2023)

TL;DR

Every market must decide *who gets what, and on what terms*. The central question of this course is:

When do simple prices work, when do they fail, and what replaces them?

Across the course, prices will recur as tools to **extract surplus, ration scarcity, decentralize allocations, correct externalities, and experiment and learn**. We will combine ideas from **operations research, economics, and computer science** to study demand, scarce capacity, customer choice, strategic behavior, private information, and institutions such as auctions and matching. The questions we will study are increasingly important in digital platforms, where algorithms – and now increasingly autonomous agents – both learn from markets and change the data the markets generate.

This is intended to be a mathematically substantive, model-driven course for senior OR/CS undergraduates and master's students. That said, I will aim to keep our organizing principle as **economic question first, mathematical machinery second**: we will introduce an optimization, probability, learning, or game-theoretic tool when a market-design problem demands it. Most lectures will begin with concrete market questions, develop a model, and try to extract reusable principles rather than presenting the math in isolation.

Lectures	Tuesday/Thursday 1:25pm–2:40pm, CIS 142
Canvas site	`https://canvas.cornell.edu/courses/88431`
Ed Discussion	`https://edstem.org/us/courses/103820/`
Instructor	Sid Banerjee
Email	sb2373@cornell.edu
Office	Rhodes Hall 229
Office hours	**[TBA: office hours]**
Prerequisites	ORIE 3300 and ORIE 3500 (or equivalent); comfort with linear optimization, basic probability, and mathematical modeling
Deliverables	5-6 homeworks, final exam (or project), and class participation requirements

Course Details

Course Description

Sometimes a market allocates through a posted price. Sometimes it uses an auction, an admission rule, a matching algorithm, or a reputation system. Increasingly, these decisions are made by algorithms that simultaneously learn from data and alter the data they observe.

The course asks how firms and platforms should set prices and allocate resources when customers have heterogeneous values, resources are scarce, demand must be learned, and participants may behave strategically. Questions include:

- Why might charging every customer exactly what they are willing to pay be a terrible long-run pricing policy?
- How can LP dual variables become actual market prices?
- Why should a seller sometimes reject a profitable customer?
- How can a firm learn a good price without sacrificing too much revenue while experimenting?
- Why can a model become *more confidently wrong* as it collects more data?
- Why can adding one more bidder outperform a more sophisticated auction?
- When should we use prices, auctions, matching, reputation, or another market rule?

We begin with the simplest benchmark—known buyer values—and progressively add unknown demand, sequential learning, scarce capacity, substitution, strategic behavior, and asymmetric information. Several mathematical connections recur:

$$\begin{array}{rcl} \text{LP duality} & \longleftrightarrow & \text{market prices} \\ \text{confidence bounds} & \longleftrightarrow & \text{exploration/exploitation} \\ \text{concentration} & \longleftrightarrow & \text{large-market predictability} \\ \text{Nash equilibrium} & \longleftrightarrow & \text{strategic interaction} \\ \text{incentive compatibility} & \longleftrightarrow & \text{truthful market design} \end{array}$$

A recurring question: what changes when the market gets large? A small market may require exact dynamic or strategic analysis. In a large market, aggregate behavior can become predictable, competition and matching opportunities can thicken, and individual participants may become less pivotal. Yet scale also makes exact optimization computationally impossible. This tension—scale breaks exact methods while creating structure that makes approximations powerful—is a central Operations Research perspective on markets. We will move repeatedly between exact small-market models and large-market approximations.

Applications include airlines and hotels, online retail, advertising, cloud and platform markets, marketplaces, matching systems, and other allocation problems with scarce resources and private information.

Course Arc

We will build up a realistic market, with which each new source of difficulty forcing a new modeling and/or mathematical idea.

1. **Pricing and demand.** We begin with known buyer values and ask what prices can accomplish: surplus division, market clearing, decentralized assignment, and congestion control. We then derive demand from value distributions and study elasticity, hazard rates, quantile revenue, virtual values, and optimal posted prices.

2. **Learning while earning.** When demand is unknown, experimentation affects current revenue. We introduce stochastic multi-armed bandits, concentration bounds, regret, and optimism/UCB. A recurring lesson is that actions determine which data are observed.

3. **Scarcity, scale, and online allocation.** Finite capacity introduces opportunity cost. We use Littlewood's rule and a small dynamic program as exact benchmarks, then formulate network revenue management and use a fluid LP and its bid prices as tractable approximations. After exposing the gap between fluid predictions and stochastic reality, we develop confidence-aware decisions and the Bayes Selector viewpoint: predict the actions of an offline clairvoyant and control the cost of disagreements.

4. **Choice and substitution.** When customers choose among available products, product-level demand is not independent of the offer set. We study the spiral-down effect, random-utility and multinomial-logit models, and assortment optimization.

5. **Auctions, competition, and mechanisms.** We compare posted prices and auctions, introduce strategic interaction and equilibrium, and study auction formats, single-parameter mechanism design, Myerson's framework, and the role of market thickness. Results such as Bulow–Klemperer illustrate how competition can sometimes substitute for more sophisticated optimization.

6. **Information structure and richer pricing.** We compare observable segmentation with hidden customer types and self-selecting menus. We then move to nonlinear pricing, versioning, bundling, and multidimensional values, emphasizing participation constraints, incentive constraints, information rents, and cannibalization.

7. **Market design beyond posted prices.** The final part develops selected models of multi-item allocation and VCG, pricing under competition, adverse selection, reputation, matching without money, and platforms. These topics show why market design must account not only for optimization but also for information, incentives, trust, and institutional rules.

The precise balance may shift with pace. Possible extensions include overbooking, finite-inventory dynamic pricing, proper scoring rules, censored-demand estimation, and deeper multi-parameter or auction models.

Learning Goals

By the end of the course, students should be able to:

- formulate and solve basic pricing and allocation models from buyer values, demand, choice, and resource constraints;
- interpret LP dual variables and dynamic marginal values as economic prices or opportunity costs;
- analyze demand learning using concentration bounds, regret, optimism, and the feedback between decisions and observations;
- distinguish exact, fluid/first-order, and clairvoyant benchmarks, and reason about how scale changes both performance and tractability;
- model substitution using random utility and MNL and optimize simple assortments;
- analyze basic strategic and informational problems and compare prices, auctions, mechanisms, reputation, and matching as market-design interventions.

Prerequisites and Course Level

Students should be comfortable with **linear optimization**, **basic probability**, calculus, and mathematical modeling, at approximately the level of ORIE 3300 and ORIE 3500 (or equivalent). In particular, we will use LP duality and complementary slackness; random variables, expectation, conditional probability, and common distributions; and elementary calculus.

Prior exposure to economics, game theory, stochastic processes, or algorithms is useful but *not required*; the necessary concepts will be introduced as they arise. Some assignments may involve computation or simulation, so familiarity with Python (or a comparable language) will be helpful.

Communication and Class Meetings

We have a public course website, which will be the main source of information for the course. Notes, assignments, reading material and other links will be posted here regularly, and we will tell you in class about updates.

We will use Ed Discussion as a place to ask and to answer questions about the course content and the homework. Please post your questions to this site and feel free to answer other students' questions there. Participation on Ed will be taken into account when determining your level of class participation for this course. To contact course staff, **please only send us private messages via Ed** (no email).

If you are not enrolled on Canvas, or Ed, please email the instructor asap.

Class meetings combine mathematical models, worked examples, and discussion of how assumptions change market outcomes. Attendance will not be recorded. Some in-class activities or quizzes may contribute to participation credit; the policy and treatment of missed activities will be specified once we reach steady state in class.

Resources

There is no single required textbook. Lecture notes and selected readings will be used throughout the semester; different references support different parts of the course.

- **Course notes.** Notes will be posted on the course website. Notes may be available for later lectures, but treat these as stubs which are not updated; we will try to update them based on what we discuss in class within a week. Students are encouraged to bring a printed copy for annotation, or take notes that they can compare later with the handouts.
- **Pricing and economic perspective.** R. Vohra and L. Krishnamurthi, *Principles of Pricing: An Analytical Approach*, provides a recurring narrative for demand, monopoly pricing, discrimination, bundling, auctions, and competition.
- **Revenue management and choice.** K. Talluri and G. van Ryzin, *The Theory and Practice of Revenue Management*, is the principal technical reference for capacity control, fluid and network RM, bid prices, and choice-based models.
- **Learning and online decisions.** A. Slivkins, *Introduction to Multi-Armed Bandits*, together with selected papers and course notes on online allocation and the Bayes Selector framework.
- **Game theory, auctions, and mechanisms.** T. Roughgarden, *Twenty Lectures on Algorithmic Game Theory*; A. Karlin and Y. Peres, *Game Theory, Alive*; and P. Milgrom, *Putting Auction Theory to Work*.
- **Additional readings.** Original papers and instructor notes will be used for selected topics including spiral-down, matching, reputation, and richer pricing or mechanism-design models.

Course Expectations

Academic Standards and Use of External Tools

Every student is expected to abide by the Cornell University Code of Academic Integrity: https://theuniversityfaculty.cornell.edu/dean/academic-integrity/code-of-academic-integrity/.

Submitted work must follow the collaboration rules stated for each assignment. Students must acknowledge collaborators and any permitted external resources. Buying, selling, reposting, or otherwise redistributing course materials without permission is not allowed.

You may use AI systems—including generative tools, coding assistants, and agents—for all work in this course. You may decide when and how to use them without requesting permission, and course activities and assessments are designed with that availability in mind. AI use is optional unless a particular activity requires it. You remain responsible for understanding the work you submit and for its correctness and provenance. You should review and test AI-generated contributions and be prepared to explain, evaluate, and modify any part of the submission. Ordinary expectations for citation and licensing continue to apply. Cite identifiable human-authored code, text, data, and ideas. Comply with the license terms governing any external materials you incorporate, including requirements for attribution, copyright notices, modification, or redistribution. These obligations also apply to sources discovered through AI. Unless the assignment states otherwise, *you do not need to cite AI merely because it assisted you, nor do you need to submit a record of your interactions with it.* This policy reflects the learning goals and assessment design of this course rather than a judgment that AI should always be used; other courses may appropriately prohibit the use of AI.

Personal Conduct and Inclusive Learning

Students are expected to be respectful and professional in class and in online communication. Cornell supports an inclusive learning environment in which differences are respected and recognized as a source of strength. Everyone in the course should make a good-faith effort to understand perspectives, behaviors, and worldviews that may differ from their own.

If you have concerns about your learning, grades, progress, or other difficulties, please contact the instructor or course staff. Additional Cornell resources are available at https://caringcommunity.cornell.edu.

Accommodations for Students with Disabilities

Cornell University is committed to full inclusion in its educational programs and services. Student Disability Services (SDS) determines eligibility and works with students, faculty, and staff to recommend reasonable accommodations. Students are encouraged to contact SDS as early as possible: https://sds.cornell.edu.