Decision Making Under Uncertainty

Module Leader

Name: Dr Ryan Cory-Wright
E-mail: r.cory-wright@imperial.ac.uk
Office hours: By appointment
Homepage: ryancorywright.github.io

Course Summary and Objectives

Since the fields of Operations Research and Management Science were founded in the Second World War and its aftermath, one of the core questions addressed by both fields has been how to make high-quality decisions under uncertainty. George Dantzig, one of the founders of both fields, described the problem that originated his research in the foreword to Linear Programming I (1997) as:

"the problem of planning or scheduling dynamically over time, particularly when there is uncertainty about the values of coefficients in the equations."

In this course, we study three modelling paradigms for making high-quality decisions under uncertainty. These paradigms are: stochastic optimization (for approximately four weeks), robust optimization (for approximately four weeks), and dynamic optimization (for approximately two weeks). Motivated by Dantzig's powerful opening line in his monograph:

"the final test of a theory is its capacity to solve the problems which originated it."

we also illustrate these modelling paradigms by discussing their applications across various areas of operations management and analytics, including supply chain management, revenue management, healthcare, and energy among others.

Learning Outcomes

On successful completion of the module, students should be able to:

- Describe the core theory for three of the main paradigms for decision making under uncertainty: stochastic optimization, robust optimization, and dynamic optimization.
- Summarize the main application areas of these paradigms in operations management and analytics.
- Interpret operational problems and model them as mathematical problems using these paradigms.
- Analyze and criticize research papers in the field.

Assessment

The overall goals of this class are to:

- Introduce the main research questions in decision-making under uncertainty.
- Summarize some of the most promising research directions of the past 20 years.
- Prepare you to perform research in decision making under uncertainty.

With these goals in mind, we will grade each student according to performance on:

- **30% from two individual assignments (15% each)**: We plan to assign two homework assignments, which are to be completed individually and submitted at the beginning of the lecture on the due date. Most problems in these assignments will be taken from textbooks in the reference list at the end of this syllabus, some are short and are meant to reinforce important lecture material. Some problems will require programming;

assistance will be available, and students are welcome to help each other with coding, but each student must be the author of the code they submit.

- **10% research paper critique:** Each student will present on a topic related to a course theme, where possible giving high-level intuition, a motivating example/potential applications, and connections to other themes and topics. The presentation should be 15 minutes, with 2-3 additional minutes of discussion time for the class. We suggest 10 main slides, which should be submitted to the course staff before the start of the class where the student presents. Students will need to select a topic to present on (by the end of week 2), and the lecture which they think it would be most appropriate for them to present in, in coordination with the course staff. They should not present before week 4, in week 7 (midterm) or week 10 (final project presentations). Students are encouraged, although not required, to select a topic/paper from the main reading list at the end of this syllabus for their research paper critique.

- **30% midterm quiz:** There will be a 180-minute midterm quiz which takes place in lecture 7, on the content in the first six weeks of the class. We will provide a practice quiz of a similar difficulty a few days before the quiz. The quiz will be designed so that most students should finish before the 180-minute mark.

- **30% final project**: This project should demonstrate the practical application of some techniques covered in the course through a short write-up (and probably some original code) and a five-minute presentation. Each student should discuss their final project ideas with the course staff well in advance of the course dates. For particularly ambitious projects, students can make small teams if approved by the course staff. Each team member will be expected to contribute as much as a student working on an individual project, so group write-ups and presentations will be longer. Projects with synergy with ongoing research are allowed, although you will need to explain what you did during the time the course was running.

A lecture-by-lecture preview

Part of this class consists of ten three-hour lectures (each lecture includes one or two breaks). Students are expected to attend all sessions in-person (health permitting). To get as much out of this class as possible, we suggest that you spend at least as much time on reading the papers and textbooks referenced in the lectures/reviewing the lectures as you spend in class. Lecture slides/handouts will be posted online in advance, so students need not take copious notes.

We now present a brief summary of the intended content for each lecture (time permitting):

1. **Introduction to Decision Making Under Uncertainty**
   - Why decision-making under uncertainty?
   - Different approaches to decision-making under uncertainty
   - Review of probability theory: Law of Large Numbers, Central Limit Theorem
   - Review of optimization theory: Convexity, Duality, KKT Optimality Conditions

2. **Two-Stage Stochastic Optimization and Sample-Average Approximation**

- Homework #1 released at end of class
- Students should pick the week they will do their critical paper review in by the end of this week, and email Ryan to tell him the week
- Motivation: OLS regression and almost sure convergence
- Sample average approximation: Theory
- A closed-form solution to The Newsvendor Problem
- Sample average approximation: Algorithmics
    - Benders decomposition 101
- Can we do better than sample average approximation?

3. Personalized+ (Contextual) Sample-Average Approximation

- Can we improve on SAA?
- Predictive to prescriptive analytics
- Smart "predict, then optimize"
- Illustrating predictive to prescriptive analytics computationally
    - Using Julia/JuMP/Gurobi

4. Robust Optimization

- Homework #1 due at start of class
- The robust "trick"
- Ell_1, ell_infinity and budget uncertainty sets
- Constraint-wise reformulations
- Application: Lasso regression and the partial equivalence between regularization robustness
- Homework #2 released at end of class

5. Robust Optimization II

- A reminder of strong duality and support functions
- Tractable reformulations of general convex uncertainty sets

6. Robust Optimization III

- Adaptive and adjustable robust optimization
- Linear decision rules

7. Midterm

- 180 min midterm quiz (30%)

8. Distributionally Robust Optimization

- The conservatism and intractability of robust and stochastic optimization
- Distributionally robust optimization to the rescue

- Data-driven DRO

9. Dynamic Optimization (a.k.a. Dynamic Programming)

- Homework #2 due at start of class
- State Space, Actions, Recursion, Boundary Conditions
- Application: The Knapsack Problem
- The Curse of Dimensionality
- Guidelines for when Dynamic Optimization is a good/bad idea

10. Dynamic Optimization and SAA

- SDDP
- Guidelines for when Dynamic Optimization is a good/bad idea
- Course summary: which modeling paradigm works best when?
- Last 90 minutes: Final project presentations.

Reference List

Note: except where explicitly indicated otherwise with a *, it should be possible to access all references listed here using the Imperial College Library website (or your university library if you are not an Imperial student). Please contact the library if it is not possible to access a reference, or a member of the teaching team if the library is unable to assist. In some cases, it can also be quicker to find a preprint of the article listed here on arXiv or Optimization Online. Some of the textbooks listed here are also freely available on the author's websites.

You may also ask to present a research critique on a paper not on this list, provided it fits within the overall scope of the class. In this case, you should discuss this with the teaching team ahead of time.

Stochastic Optimization:

Recommended Review Articles and Textbooks:

- Birge, J., Louveaux, F. (2011) Introduction to Stochastic Programming*. Springer*.
- Kall, P., Wallace, S. (1994) Stochastic Programming. (*Wiley*).
- Shapiro, A., Dentcheva, D., Ruszczynski, A. (2021*)* Lectures on Stochastic Programming: Modeling and Theory. *Society for Industrial and Applied Mathematics*.
- Shapiro, A., Philpott, A. (2007) A Tutorial on Stochastic Programming. Technical Report. Available at: https://stoprog.org/sites/default/files/SPTutorial/TutorialSP.pdf

Two-Stage Stochastic Optimization and Sample-Average Approximation:

- Kleywegt, A. J., Shapiro, A., Homem-de-Mello, T. (2002). The Sample Average Approximation Method for Stochastic Discrete Optimization. *SIAM Journal on Optimization*, *12*(2), 479-502.
- Mak, W.-K., Morton, D.P., Wood, R.K. (1999) Monte Carlo Bounding Techniques for Determining Solution Quality in Stochastic Programs. *Operations Research Letters* 24:47-56.

Personalized Sample-Average Approximation:

- Ban, G-Y., & Rudin, C. (2019). The Big-Data Newsvendor: Practical Insights From Machine Learning. *Operations Research,* 67(1), 90-108.
- Bertsimas, D., & Kallus, N. (2020). From Predictive to Prescriptive Analytics. *Management Science*, *66*(3), 1025-1044.
- Elmachtoub, A. N., & Grigas, P. (2022). Smart "Predict, Then Optimize". *Management Science*, *68*(1), 9-26.
- Ho-Nguyen, N., & Kilinc-Karzan, F., Risk Guarantees for End-to-End Prediction and Optimization Processes. *Management Science,* 68(12), 8680-8698.
- Kallus, N., Mao, X. (2022). Stochastic Optimization Forests. *Management Science*.
- Loke, G. G., Tang, Q., & Xiao, Y. (2022). Decision-Driven Regularization: A Blended Model for Predict-Then-Optimize. *Available at SSRN 3623006*.

Multi-Stage Stochastic Optimization:
- Dantzig, G.B., Infanger, G. (1993) Multi-Stage Linear Programs for Portfolio Optimization. *Annals of Operations Research,* 45(1):57—76.
- Pereira, M.V.F., Pinto, L.M.V.G. (1991) Multi-Stage Stochastic Optimization Applied to Energy Planning. *Mathematical Programming* 52:359—375.
- Shapiro, A. (2011) Analysis of Stochastic Dual Dynamic Programming Method. *European Journal of Operational Research,* 209(1):63—72.

Robust Optimization

Recommended Review Articles and Textbooks:
- Bertsimas, D., Brown, D., Caramanis, C. (2011) Theory and Applications of Robust optimization. *SIAM Review* 53(3): 464—501.
- Bertsimas, D., Den Hertog, D. (2022) Robust and Adaptive Optimization. *Dynamic Ideas Press*. (*)
- Rahimian, H., Mehrotra, S. (2019) Distributionally Robust Optimization: A Review. *arXiv preprint arXiv 1908.05659.*

Static Robust Optimization:
- Chap. 2 of Bertsimas, D., Den Hertog, D. (2022) Robust and Adaptive Optimization. *Dynamic Ideas Press*. (*)
- Ben-Tal, A., Nemirovski, A. (1998) Robust Convex Optimization. *Mathematics of Operations Research* 23(4):769-805.
- Ben-Tal, A., Nemirovski, A. (1999) Robust Solutions of Uncertain Linear Programs. *Operations Research Letters* 25(1):1-13.
- Bertsimas, D., Pachamanova, D., Sim, M. Robust linear optimization under general norms. *Operations Research Letters* 32(6):510-516.
- Bertsimas, D., Sim, M. (2004) The Price of Robustness. *Operations Research* 52(1): 35—53.
- Bertsimas, D., Sim, M. (2006) Tractable Approximations to Robust Conic Optimization Problems. *Mathematical Programming* 107(1-2): 5—36.
- Fan, X., and Hanasusanto, G. A. (2024). A Decision Rule Approach for Two-Stage Data-Driven Distributionally Robust Optimization Problems with Random Recourse. *INFORMS Journal on Computing*, 36(2), 526–542.
- Soyster, A. (1973) Convex Programming With Set-Inclusive Constraints and Applications to Inexact Linear Programming. *Operations Research* 21(5):1019—1175.
- Soyster, A. (1974) A Duality Theory for Convex Programming with Set-Inclusive Constraints. *Operations Research* 22(4):683—915.

Applications of Static Robust Optimization:
- Chap. 21-22, 25-27 of Bertsimas, D., Den Hertog, D. (2022) Robust and Adaptive Optimization. *Dynamic Ideas Press*. (*)
- Bertsimas, D., Copenhaver, MS. (2018) Characterization of the Equivalence of Robustification and Regularization in Linear and Matrix Regression*. European Journal of Operational Research* 270(3):931—942.
- Bertsimas, D., Dunn, J., Pawlowski, X., Zhou, Y. (2019) Robust classification*. INFORMS Journal on Optimization* 1(1):2—34.
- Bertsimas, D., Thiele, A. (2006) A robust optimization approach to inventory theory. *Operations Research* 54:150—168.

- Cheng, T., Dai, M., Zhang, X. (2024) Robust Drone Delivery with Weather Information. *Manufacturing & Service Operations Management,* 26(4):1441–1460.
- El-Ghaoui, L., Lebret, H. (1997) Robust solutions to least-square problems with uncertain data matrices. *SIAM Journal on Matrix Analysis and Applications* 18(4):1035—1064.
- Goldfarb, D., Iyengar, G. (2003) Robust Portfolio Selection Problems. *Mathematics of Operations Research* 28(1):1-38.
- Perakis, G., Sim, M., Tang, Q., and Xiong, P. (2023). Robust Pricing and Production with Information Partitioning and Adaptation. *Management Science*, 69(3), 1398–1419.
- Xu, H., Caramanis, C., Mannor, M. (2009) Robustness and Regularization of Support Vector Machines. *Journal of Machine Learning Research* 10: 1485—1510.
- Xu, H., Caramanis, C., Mannor, M. (2010) Robust regression and Lasso. *Advances in Neural Information Processing Systems* 56(7): 1801—1808.

Distributionally Robust Optimization:

- Chap. 19-20 of Bertsimas, D., Den Hertog, D. (2022) Robust and Adaptive Optimization. *Dynamic Ideas Press*. (*)
- Delage, E., Ye, Y. (2010) Distributionally Robust Optimization Under Moment Uncertainty With Application to Data-Driven Problems. *Operations Research* 55(3):98—112.
- Gao, R. (2023) Finite-Sample Guarantees for Wasserstein Distributionally Robust Optimization: Breaking the Curse of Dimensionality. *Operations Research* 71(6):2291-2306.
- Gao, R., Kleywegt, A. (2023) Distributionally Robust Stochastic Optimization With Wasserstein Distance. *Mathematics of Operations Research* 48(2):603-655.
- Goh, J., Sim, M. Distributionally Robust Optimization and its Tractable Approximations. *Operations Research* 58(4):902-917.
- Kuhn, D., Shafiee, S., and Wiesemann, W. *(2025).* Distributionally Robust Optimization*. Acta Numerica, 34, 579–804.*
- Van Parys, B. P., Esfahani, P. M., Kuhn, D. From Data to Decisions: Distributionally Robust is Optimal. *Management Science* 67(6):3387—3402.
- Wang, I., Becker, C., Van Parys, B., and Stellato, B. (2025). Mean robust optimization*. Mathematical Programming,* 213(1–2), 1235–1277.
- Wiesemann, W., Kuhn, D., Sim, M. (2014) Distributionally Robust Convex Optimization. *Operations Research* 62(6):1358—1376.
- Zhen, J., Kuhn, D., Wiesemann, W. (2025) A Unified Theory of Robust and Distributionally Robust Optimization via the Primal-Worst-Equals-Dual-Best Principle. *Operations Research*, 73(2):862–878

Applications of Distributionally Robust Optimization:

- Selvi, A., Belbasi, M.R., Haugh, M. B., Wiesemann, W. (2022) Wasserstein Logistic Regression with Mixed Features. *Advances in Neural Information Processing Systems* 35:16691—16704.
- Selvi, A., Liu, H., Wiesemann, W. (2025) Differential Privacy via Distributionally Robust Optimization. *Operations Research*, articles in advance.
- Wang, S., and Delage, E. (2024). A Column Generation Scheme for Distributionally Robust Multi-Item Newsvendor Problems. *INFORMS Journal on Computing*, 36(3), 849–867
- Xu, H., Mannor, S. (2010) Distributionally Robust Markov Decision Processes. *Advances in Neural Information Processing Systems* 23.

Optimization Bias and Out-of-Sample Disappointment:

- Gupta, V., Huang, M., Rusmevichientong, P. (2022). Debiasing In-Sample Policy Performance for Small-Data, Large-Scale Optimization. *Operations Research*.
- Gupta, V., Rusmevichientong, P. (2021). Small-Data, Large-Scale Linear Optimization With Uncertain Objectives. *Management Science*, *67*(1), 220-241.
- Rice, L., Wong, E., Kolter, J. Z. (2020) Overfitting in Adversarially Robust Deep Learning. *International Conference on Machine Learning* 8093—8104.
- Smith, J.E., Winkler, R. L. (2006) The Optimizer's Curse: Skepticism and Post-Decision Surprise in Decision Analysis. *Management Science* 52(3):311-322.
- Tsybakov, A. B. (2004). *Introduction to Nonparametric Estimation*. Chapter 3.4.

Adaptive Optimization:

- Bertsimas, D., Iancu, D., Parrilo, P. (2010) Optimality of Affine Policies in Multi-Stage Robust Optimization. *Mathematics of Operations Research* 35(2):363-394.
- Bertsimas, D., Litvinov, E., Sun, X.A., Zhao, J., Zheng, T. Adaptive Robust Optimization for the Security Constrained Unit Commitment Problem. *IEEE Transactions on Power Systems* 28(1):52—63.

Dynamic Optimization

Recommended Review Articles and Textbooks:

- Chapter 6 of: Kleinberg, J., Tardos, E. (2006) *Algorithm Design.* Pearson Education
- Bertsekas, D. (2012) *Dynamic Programming and Optimal Control: Volumes I-II.* Athena Scientific.

Applications of Stochastic Dynamic Optimization:

- D'Aeth, J. C. et al. (2023) Optimal Health Care Scheduling During the SARS-CoV-2 Pandemic, *Management Science*.
- Philpott, A. B., Mason, A. (2001) Optimal Yacht Routes Under Uncertainty *SNAME Chesapeake Sailing Yacht Symposium.*