Fall, 2024, Mini 2: 94-833:
Decision Analysis and Multi-Criteria Decision Making

Lecture: Tues/Thurs 2:00 – 3:20 PM in HbH 1007
Recitation: Friday 9:30 – 10:50 AM in HbH 1007

Instructor:
Jon Caulkins
HbH 2219
(412) 268-9590
caulkins@andrew.cmu.edu

Course Secretary:
Adrienne McCorkle
Hamburg Hall 2250
(412) 268-9576

Office Hours: Tuesday & Thursday 1:00 – 1:50 PM and Friday 11 AM – 12 Noon.[1]

Course Description:

This course complements and extends 95-760 Decision Making under Uncertainty and both two-mini sequences in Management Science (90-722 & 90-760 and 90-755 & 90-775) by addressing three topics in managerial decision making:

- MCDM (Multi-Criteria Decision Making)
- Decision Analysis
- Decision processes

MCDM is a collection of methods for trading-off different alternatives' performance on multiple conflicting attributes or objectives. A classic example from personal life is deciding which car to buy (attributes: cost, appearance, MPG, cargo capacity, etc.) or which apartment to rent (cost, size, distance to work, neighborhood amenities, etc.). Classic professional examples include which vendor to choose in response to an RFP or which public policy to follow (cost to taxpayers, scale of one or more intended benefits, scale of various unintended benefits, etc.). MCDM methods include weighted sum scoring models, swing weights, TOPSIS, DEA, AHP, and rank-based methods.

Decision Analysis is the prescriptive model for rationally maximizing subjective expected utility in the face of uncertainty; it is particularly powerful for dealing with sequential decisions, quantifying the value of information, assessing & incorporating subjective probabilities, and doing Bayesian updates of probabilities as new information becomes available.

[1] A *Chronicle of Higher Ed* article said "about 11 percent of instructors [surveyed] described one benefit of office hours as discussing with students their career goals and opportunities that could support them, like internships and conducting research. None of the students provided that reason. And while about 20 percent of professors indicated one-on-one time or getting to know one another as a purpose of office hours, only about six percent of students did". FWIW, in a smaller 2nd year course, I think good reasons to come to office hours include career discussions and getting to know each other, and I'd add talking about your capstone projects and just shooting the breeze about current events, new ideas, science, etc.

Decision process considerations go beyond the paradigm of a single well-defined decision maker and mathematical method. We will only have time to introduce this area, but topics we cover include industrial analytics, composite indicators ("US News & World Report" style ratings), balanced scorecards, "dashboards" of key performance indicators, group processes, and matching a decision method to the circumstances at hand.

Course Objectives:

The first two topics are skills oriented. Their objective is to empower students to apply the methods in professional practice (and in personal life decisions). Students should leave the class able to apply the methods to routine problems, and with a foundation for further self-study when the decision's stakes warrant use of more advanced versions of these methods.

The last component is more conceptual, providing exposure to issues one confronts when bridging between textbook methods and real-world implementation. The twin learning goals are knowledge about the issues and developing critical thinking skills concerning the "meta decision" of when to employ which decision method, and the criteria relevant to making those judgments.

Prerequisites:

The only formal prerequisite is fluency with algebra & Excel and basic knowledge of probability theory and distributions (random variables, conditional probability, Bayes Rule, etc.). However, this course is usually taken after either 95-760 Decision Making Under Uncertainty or one of the two Management Science sequences (90-722 & 90-760 or 90-755 & 90-775), and the course is pitched at a second-year level. First-year students who are comfortable with quantitative methods are welcome, but first-year students without a quantitative bent are encouraged to wait until their second-year.

Canvas:

Course materials will be posted to Canvas. Monitor it for announcements (e.g., changes to assignments).

Readings should be read ***before*** the class in which they will be discussed.

Recitations:

Recitation are optional and informal. During the first two weeks, we'll be vetting ideas for what decision to analyze in HW #1 and HW #2. If you're not in recitation, we can do that one on one. Originally the course didn't have recitations, but there had been demand for "collective office hours". Some weeks I'll bring some exercises we can do, and I can work textbook-like problems if asked, but often the most fun comes from spending time in open-ended discussion of "structuring" homework problems, systems synthesis projects, problems from your internships & professional work, etc.

Grading:

Course grade is based on homework (38%), class participation (15%), project presentation/memo/technical report (27%), and the final exam (20%).

Class participation means *participation*, not just attendance. You will score poorly on class participation if all you do is show up on time for every class.

Project presentations will be Tuesday November 26th, with other graded items related to the project due on other dates (see project description on Canvas for details). I intentionally make the project due a little earlier, and make the final week of class a little easier, to compensate for many other courses having the greatest workload in the final week.

There will four homework assignments, plus two instances of required peer feedback. Type the parts of HW that are mostly words, but you may (neatly!) handwrite parts involving algebra, decision trees, derivations, etc. that would be time consuming to type. HW is due at the beginning of class. Late assignments will receive a 0.

HW #1 & #2 must be done individually. HW #3 & #4 can be done individually but I encourage you to work in groups of two, three, or four. If you work in a group: (1) submit one HW for the group and everyone in the group will receive the same grade and (2) more will be expected of group HW in terms of clarity of presentation and exposition, absence of arithmetic errors, quality of writing, etc. Within a group you may collaborate in any way you choose. Across groups feel free to discuss concepts and approaches, but there should never be sharing of anything electronic or stating of answers or result obtained.

You may talk with others in the class about HW #1 & #2 because (as you will see) you will all be doing your own unique analyses, so there is no risk of copying answers from each other. Ask me for clarification if necessary.

Because you can benefit from feedback and from seeing what other people do on HW #1 & #2 (there is no single "answer key"), I will require you to read and write a one-page review/critique of another person's HW, due nine days after the HW gets submitted. You'll send that to your classmate and cc me. I'll "grade" your critique out of 1 point based mostly on effort & professionalism. I would like you to find your own pairs for this purpose, rather than assigning papers randomly, because some people will write about things that they might prefer to share only with people they know/trust.

Taping or Recording Classroom Activities

No student may record or tape any classroom activity without my express written consent. If a student believes that he/she is disabled and needs to record or tape classroom activities, he/she should contact the Office of Disability Resources to request an appropriate accommodation.

Academic Integrity:

I am not kind to cheaters. If you cheat, I will fail in you in the course, report you to the administration and support expulsion from the University. If you have any questions about what might or might not constitute cheating, ask me.

Accommodations for Students with Disabilities:

If you have an accommodations letter from the Disability Resources office, please discuss your needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate.

Readings & Detailed Course Outline

The required readings are mostly from chapters I have written plus Chapter 14 (Decision Analysis) from Cliff Ragsdale's *Spreadsheet Modeling and Decision Analysis: A Practical Introduction to Management Science* (any edition from the 6th edition on will do). The book is available in various formats from the publisher, Cengage.

Part I: Multi-Criteria Decision Making

It is always good to do the reading before the class in which the material is discussed, but that is imperative for the classes on AHP and rank-based methods. Also: There is quite a bit to read before the 2nd class, so please budget time for that.

The readings, including my chapters are available on Canvas.

Class #1 Intro and weighted sum models: Read before class Sections 1.1 – 1.5 (8 pages) from my book chapters on MCDM. Ragsdale also offers a 4-page intro to MCDM (Sections 14.16-14.17), but I don't think you need to read it.

Class #2 Finish weighted sum models & TOPSIS: Read 1.6-1.7 (5 pages), Section 2.3 on swing weights (2 pages), and watch a 22-minute video *before* class. After class you should skim Sections 2.1-2.2, 2.4, and 2.8whch review things we'll cover in class.

Before doing HW #1, read about (1) Identifying and organizing objectives/criteria in pages 45-55 of Keeney, R.L. (2013) Identifying, prioritizing, and using multiple objectives. *EURO Journal on Decision Processes*, 1:45-67 and (2) Creating single attribute value functions in pages 119 – 134 of Belton and Stewart (2001) *Multi Criteria Decision Analysis: An Integrated Approach*.

If you are curious you can read what they say about swing weighting (Keeney pp.56-64 and Belton & Stewart pp.135-138), and Belton & Stewart pp.139-141 are also useful for distinguishing what they call relative and cumulative weight when objectives are nested in a tree structure.

Optional: Ragsdale explains how to implement DEA with Excel's Solver, but that is not part of this course (Section 3.15).
Optional/Note: In class I will contrast MCDM & Goal Programming (GP), but I know not all of you learned GP in your optimization class, so I will not test you on GP ideas.

Class #3: DEA & AHP: Read my Chapter 2.5 (on DEA) and 2.6 (on AHP), but you can save the last 3 pages of 2.6 (on details of AHP calculations and AHP express) until after class, and then you can just skim them. Ragsdale's 6-page section 14.18 on AHP is very good on the computations, but briefer on background. I'd call that an optional but useful 2nd source.

Class #4: Rank-based methods: Read my Chapter 2.7 (11 pages on rank-based methods) and 2.9 (2-page review of MCDM).

I suggest but do not require you to read the rest of my Chapter 1 at some point and

Spradlin, T., "A Lexicon of Decision Making", DSSResources.COM, 03/05/2004.
(http://dssresources.com/papers/features/spradlin/spradlin03052004.html)
If you read this before the first day of class, that'd great, but there is a lot of other material to read in the first week so deferring it until the 2nd or 3rd week is fine.

Other completely optional resources on MCDM you might find useful in the future:
Belton, Valerie and Theodor J Stewart (2001). *Multi Criteria Decision Analysis: An Integrated Approach*. Springer.

If you want to use MCDM professionally and wish you understood more than we could cover in this class, read the Belton and Stewart book. It is excellent on scoring models, creation of value function, and the relationship between MCDM and DEA.

Ralph Keeney is famous for observing that failure to make the right choice is often rooted in failing to correctly identify what you really care about. He espouses Value-Focused Decision making, and you can tap into his wisdom at various levels of intensity.

- Keeney, R. *Value-Focused Thinking: A Path to Creative Decision Making*. Harvard University Press, 1992. The definitive treatment.
- Keeney, R.L. (1994). Creativity in Decision Making with Value-Focused Decision Making," *Sloan Management Review*, Summer, 33-41. Quick read, with practical tips on how to identify decision criteria.

There are many books and literally thousands of articles on DEA. I do not know which is best. You can browse the web resources (e.g., http://deazone.com/en/) as well as I can, but I think the original papers that established the method are still useful because they were written to explain it to people who had never heard of it before,

Charnes, A., W. Cooper, & E., Rhodes (1978) "Measuring the efficiency of decision-making units," *European Journal of Operational Research* vol. 2, pp. 429–444.

Banker, R.D., R.F. Charnes, & W.W. Cooper (1984) "Some Models for Estimating Technical and Scale Inefficiencies in Data Envelopment Analysis, *Management Science* vol. 30, pp. 1078–1092.

For the practicalities of doing DEA in a spreadsheet without specialized software, see Section 3.14 of Ragsdale, Cliff T. *Spreadsheet Modeling and Decision Analysis: A Practical Introduction to Management Science*.

For more on TOPSIS, see:
Hwang C.L. and K. Yoon, *Multiple Attribute Decision Making: Methods and Applications*, Springer-Verlag, New York, NY, 1981.

For more on rank-based methods, see:
Börgers, Christoph (2010). *Mathematics of Social Choice: Voting Compensation and Division*, Philadelphia: SIAM.

Part II: Decision Analysis

The main text is chapters (Chapters 3, 4, and 5) that I have written and posted on Canvas, plus portions of the Decision Analysis chapter (Chapter 14) of Cliff Ragsdale's *Spreadsheet Modeling and Decision Analysis*.

First Pair of Classes: The learning objectives concern payoff matrices, the value of perfect and imperfect information, decision trees, and "flipping" trees. My strategy is to cover all (or almost all) of these ideas in the first class, and devote the 2nd class to examples/cases. That might seem to imply you should read everything before the 1st class, but some ideas are easier to grasp if you see them first (in the 1st class) before reading about them, so those sections are listed as readings for the 2nd class. Also, you can choose between relying primarily on chapters I wrote or on Ragsdale's Chapter, but note that Ragsdale treatment is not as deep.

First Class: Read Sections 3.1 - 3.3 (6 pages on payoff matrices & EVPI) and Sections 4.1-4.2 (2 pages on decision trees).

Or, if you want Ragsdale's take on these topics, that'd be Sections 14.0-14.4 (5-page intro & payoff matrices), 14.8 (2 pages on EVPI), & 14.9 (3 pages on decision trees). Skip 14.10 & 14.12 (good stuff, but it is specific to one software platform).

Second Class: Read the cases we'll work (I'll decide which based on how the first class goes). Also, read or skim the readings on the more advanced ideas from the first class. If they made perfect sense in class, then you can skim. If the ideas were not clear, allow more yourself more time to read more thoroughly.

From my chapters, that'd be Sections 3.4-3.5 (imperfect info in payoff matrices), 4.3 (detailed worked example parallel to the market research example from class), Section 4.4 (Value of imperfect info in a big tree), and 4.5 ("flipping trees", 3 pages). The rest of Section 4 covers optional ideas that I usually defer until Cass #4 to allow enough time to work examples.

The corresponding sections in Ragsdale's book are 14.11, & 14.13 (8 pages on decision trees and flipping trees) and 14.14 (3 pages on conditional probability and Bayes' Rule).

Second Pair of Classes: The primary learning objectives concern (1) risk-return frontiers, (2) stochastic dominance, and (3) utility theory. The last is the most important, and it includes modeling DM's subjective utility functions and the associated risk premium (RP) and certainty equivalent (CE). Usually there is time to cover some other ideas that could logically belong to week #1, but I think are better deferred so the 2nd day of week #1 can focus on examples/cases.

Third Class

Risk-return frontiers and stochastic dominance are covered in detail in my Chapter 5, which is an adapted version of what I usually assign in 90-760 Management Science II (so it's pitched at a more introductory level than are Chapters 3 & 4). It's long (19 pages), but you can skim/skip 5.1-5.2 (which are on non-probabilistic decision rules, which Ragsdale

covers in 14.5-14.6) and you are not required to learn the newsvendor problem or yield management applications (Sections 5.6-5.7). Ragsdale does not cover those ideas at all.

Ragsdale covers utility theory in Section 14.15 (6 pages), and I introduce it in Section 5.4 (2 pages).

Fourth Class Read my summary of key points in week #2 concerning utility theory.

The advanced topics I cover in Chapter 4 are
4.6: Bayesian updating of a continuous variable (4 pages)
4.7: Sequential decision analysis (1-page brief intro)
4.8: Modeling continuous uncertainties (2 pages)
4.9: Eliciting subjective probabilities (2 pages)

Time permitting, I will discuss challenges to the MAUT view of the world, such as Allais' Paradox and Framing effects. Those are super cool ideas that are covered in detail elsewhere on campus. There is no assigned reading for those. Mostly I just want to explore the (interesting) boundary between MAUT as a prescriptive model and humans as (flawed?) decision makers.

If you want to learn more about Decision Analysis, an easy-to-read and excellent older textbook is:
Clemen, Robert T. and Terrence Reilly (2004). *Making Hard Decisions with Decision Tools*, Duxbury, Pacific Grove.

The definitive advanced treatment of Decision Analysis methods is still:
Keeney, Ralph L. and Howard Raiffa (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs,* Wiley and Sons.

Part III: Decision Processes

See Canvas pages for details, but the readings are:

Tuesday: Ranking Colleges and Other Things
Read my chapter on US News & World Report Type Rankings and then
Read Gladwell (2011) "The Order of Things: What College Rankings Tell Us"

Thursday: Engineering Excellent Decision Processes
Skim Wikipedia entry on "The Wisdom of Crowds"
Skim Adamatsky and Ilachinski, "The Wisdom of Slime"
Read Davenport & Harris, *Analytics at Work*, on Craft vs. Industrial Analytics
Read Davenport & Harris, *Analytics at Work*, on Managing Decisions as Processes
Read prep memo on "Classroom exercise on decision process design"