MBA 635 Topics: Strategic Decision Intelligence

Spring 2025 (revised 2025-03-08)

Bill Foote

Contents

Course Description

Strategic Decision Intelligence develops both the process and result of a seemingly heartless component of intelligence, namely the power of human reasoning. A decision is a compound of the organizational and individual experience, understanding, judgment, in a word knowledge, coupled with action, namely the decision. Typical strategic decisions include investment, financing, partnering, pricing, market selection, production and distribution location, insurance, innovation, joint ventures and any number of portfolio alternatives such as product and sales mix, talent hierarchies, market mix, partner mix, incentives.

By a strategic decision we mean the following:

- A decision which substantially affects the achievement of the vision, mission, and purpose of the organization. Such a decision would often be deferred to the Board of Directors or Trustees of the organization with regard for the delegation of authority from the Board to Management.
- A materially substantive decision such as the acquisition of a competitor in a new market, representing a material change in potential revenues, numerous legal, tax, environmental, political, and even cultural changes in the organization, reflecting societal changes outside of the corporation.
- A decision for which there is anticipated a strong and persistent reaction, possibly retaliatory, possible and new entrant or even an exit by a competitor, including the competitor's supply chain.
- A decision, whose prosecution requires a major restructuring of the organization, including culture, commitments, resourcing, legal and compliance requirements, hierarchies, networks, and markets.

Strategic Decision Intelligence operates within the management paradigm of being-knowing-doing with the emphasis on the knowing component. SDI generates simulations of decision maker interactions using system dynamics principles and practices and is grounded in group participation behavior. The generative model is the statistical causal model fed to a Bayesian inferential information-theoretic engine. The generative model enabling Bayesian inference embeds four jointly operative decision paradigms (analytical, directive, conceptual, behavioral). These decision paradigms reach into and inform the being-doing aspects of management through configurations of power (concentrated-diffuse) and loyalty (institutional-positional) in people, and their supporting social and organizational structures, and the behavioral formation of expectations. In this way SDI accommodates and integrates collaborative, rule-based, executive, and negotiated decision styles. It implies that artificial intelligence (AI) and machine learning (ML) do not have the realistic potential to replace human intelligence, rather they are indeed tools we might deploy to enhance human intelligence and well-being. SDI prioritizes persons at the interface of persons who make decisions consistent with data and the technology persons make to inform, govern, and execute those decisions.

Throughout the course we apply Bayesian data analytics to calibrate, fit, simulate, and infer conclusions consistent with the causal model of the decisions we are studying. The causal model will be built in VensimPLE, read into R, where we can mash the model with data and use Stan to make probabilistic inferences and R's ggplot2 to visualize and animate results. During the first 5 of 7 weeks we will build basic system dynamics causal models of various business decision, estimate probability models of decisions based on the causal models, and infer probabilistic results using Machine Learning and information criteria. During weeks 6 and 7 we will apply our decision technology to the problem of describing, explaining, and inferring decision outcomes for Public Benefit Corporations culminating in a final course project to be presented at the end of the term.

For this course we will cross three computing platforms. **VensimPLE** will help us build and simulate generative causal models, visualize results, and develop scenarios for decision makers. The **R programming language (with R Studio - Posit)**, the tidyverse of data, optimization, numerical integration, and visualization packages will provide a platform for analysis, inference, and visualization of results. The **Stan (for Stanislaus Ulam) probabilistic programming library** with its ability to estimate systems of differential equations (the underlying mathematics of system dynamics) using Hamiltonian Monte Carlo simulation will allow us to estimate both volatility and uncertainty within the causal models we have built for decision makers to consume.

Is is possible to use just paper and pencil, perhaps spreadsheets? Yes, but talk to the instructor about your approach. Spreadsheets also offer backstop tech-stacks to extend and automate spreadsheet routines. **VensimPLE** only requires some baseline training we provide during the course and only a smattering of elementary algebra. In the end, as in the beginning, we will paper and pencil our models. Here is a Google Sheets example of a fairly complex workforce dynamics model based on a *Medium* article by Daniel Sepulveda Estay, A System Dynamics model using spreadsheets — a case in Healthcare Operations.

Course topics

Here is a summary of the weekly topics.

Week	Topic	Notes
Week 1	Building a generative model of decision maker interaction: collaboration and competition	Simulating a decision anthropology of gift with Lotka, Volterra, and Schindler
Week 2	Endogenizing external factors and limiting growth and innovation	Constraints on decision maker innovation in markets with Meadows, Guardini, and Hirschfeld
Week 3	Counting with Bayes to infer (learn)	Updating knowledge to infer decision consistencies with McElreath, Ockham, and Jaynes
Week 4	Acting on Bayes to build decision alternatives consistent with data	The classic two state - two decision model with Zellner, Laplace and Lagrange
Week 5	Inferring the impact, uncertainty and predictability of decision tradeoffs	An inference machine for language and diversity with Watanabe, Nettle and Searle
Week 6	Expectations formation, learning, priooritization, and managerial decisions	Decision inference in a behavioral model with Sterman, Turing, and Aquinas
Week 7	Compendium: final case preparation and presentation	A trolley ride with (Philipa) Foot

Note Well: We will consult with several collaborators (e.g., Aquinas and Meadows) in our discussion boards and lecture notes throughout the course.

Resources

Readings

The main resource for the course is online at systemdynamics101.com including these books and articles (more of each in the following weeks):

1. Jim Duggan, *System Dynamics in R*, 2016. A rentable and for purchase online edition is accessible here.
2. William G. Foote. 2022. *Stats Two: Bayesian Data Analysis using R, Stan, Rethinking, and the Tidyverse*. Access here. This book compiles much of the material needed for the course from other sources and on its own. It is inspired by E.T. Jaynes *Theory of Probability: The Logic of Science*. Access a pre-print copy is in this directory. and owes most of its impetus, and not a few of the models, from the next resource by Richard McElreath.
3. Richard McElreath. 2020. *Statistical Rethinking: A Bayesian Course with examples in R and Stan* details through this site We will be working through elements of the first 5 chapters, but will move quickly to the next to the last chapter on Lynx and Hare competition.

4. System Dynamics Review - 2021 - Andrade and Duggan - A Bayesian approach to calibrate system dynamics models using Hamiltonian Monte Carlo.This, along with Jair Andrade's calibration tutorial will provide much of the inference workflow for our 7 week journey through this topic.

5. Gelman and Shalizi's Bayesian Hypothetico-Deductive Methodology (2013). Do we agree with this fairly scathing approach? There are excellent examples of the inferential process implied by Bayesian analytics of which we well need be aware. Do we count instances of an event (frequentist) or do we count the ways in which events are consistent with our hypotheses about the events (Bayesian and Aristotle / Aguinas)?

6. Tom Fiddaman's Vensim Calibration notes. We will be using some of the workflow here. However, most of that tasking will be inside of R with Stan.

7. R for Graduate Students. Simply, and very cohesively, arranged and developed introduction to R using the **tidyverse**.

8. Contemporary statistical inference for infectious disease models using Stan. Epidemics 29 (2019). Anastasia Chatzilenaa et her colleagues have compiled a well-worked explanation of the basic Bayesian model structures for inferring conclusions from system dynamics models. Abstract: This paper is concerned with the application of recent statistical advances to inference of infectious disease dynamics. We describe the fitting of a class of epidemic models using Hamiltonian Monte Carlo and variational inference as implemented in the freely available Stan software. We apply the two methods to real data from outbreaks as well as routinely collected observations. Our results suggest that both inference methods are computationally feasible in this context, and show a trade-off between statistical efficiency versus computational speed. The latter appears particularly relevant for real-time applications.

9. Foote, W. G. (2024). Neither a Beast Nor a God: A Philosophical Anthropology of Humanistic Management. Humanistic Management Journal, 1–45.. I wrote this article (gratefully published!) as my ongoing notes for understanding the reasons for the reasons we have when we manage, decide, organize, essentially giving ourselves to one another. The mindset (and heartset too) here is the underpinning of Strategic Decision Intelligence. Two poignant examples (the lidmen who cooking coal for steel production, and the families enslaved by warlords to mine rare earths) goaded me into writing this really long, involved, and highly textured article. It continues to put my preconceptions in dire peril. This is in service to C.S. Lewis's characterization of the educational goal to water the deserts, not cut the jungles.

Additional Resources

1. Hadley Wickham, Danielle Navarro, Thomas Lin Pedersen, *ggplot2: Elegant Graphics for Data Analysis (3e)*, online with books, slides, datasets, R scripts, numerous examples here, *grammar of graphics* package resources for numerous ways to visualize data of all types. We will frequently need this resource to glean important insights from the simplest to the most complex of our models.

2. Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*, Springer Science & Business Media, 2009. The authors have a tutorial site with a downloadable edition of this book, R scripts and other materials accessible here. These authors are at the tip of the machine learning spear in the frequentist tradition.

3. Yanchang Zhao, *Data Mining with R*. 2020, online with books and other resources all accessible here all from the frequentist tradition of model formulation, inference and interpretation.

4. The maths: Harry Hochstadt's *Differential equations : a modern approach (1963-4)* details in the first 84 pages the math behind system dynamics, namely solving systems of simultaneous differential equations.

5. The numerics: Joel Ferziger's *Numerical Methods for Engineering Applications, 1978.* Yes, FORTRAN. I moved most of these routines to MATLAB and APL2 in the 80's.

Weekly live sessions will expand on key aspects of chapters from Duggan, Foote and McElreath, and others as needed, and prepare us for weekly assignments. Vensim mdl models, R scripts, Stan code, along with `RMarkdown` source files, and data sets accompany each week.

Access

VensimPLE

1. With your Manhattan College sign-on credentials use Manhattan College remote desktop (DLS 210, 309 or 314) to access VensimPLE.
2. Download (for free) VensiMPLE. Choose your Windows or Mac OSX platform, your email, and you will receive instructions to follow.

R, Python and Stan

Access R in several possible ways:

1. With your Manhattan College sign-on credentials use Manhattan College remote desktop (DLS 210, 309 or 314 to use R Studio, the integrated development environment (IDE) we will use.
2. Use the RStudio cloud platform. This option does not support the Stan library and the cmdstanr package, yet. It does support the slimmed down rethinking package, which will suffice for our first 3 weeks.
3. Long and Teetor (2019, chapter 1) provide a tutorial to access R and Rstudio.
4. Access and install Stan through the CmdStanR site. Those of the Python persuasion may access and install Stan through the CmdStanPy Github site (scroll down to Read.me). The sites will direct you to install the **CmdStan** C++ library which power the R and Python interfaces.

Spreadsheets

It is possible to code everything we do in this course in a spreadsheet, if we use macros and subroutines as well as strict adherence to the principles and practice of spreadsheet engineering. Two brands are available: Google Sheets and Microsoft 365. Google Sheets requires us to use Javascript, while Microsoft 365 uses Visual Basic for Applications to build automation tools such as macros, user-defined functions and subroutines. Here is a Google Sheets example of a fairly complex workforce dynamics model

Data

Most of the data for the course comes from R packages such as **rethinking@slim**, a non-Stan version of **rethinking** both by Richard McElreath.

1. Set up a working directory on your computer by creating an R project in RStudio. Typically this is located in the user's documents directory. With this working directory go to RStudio and create a new project to you will save your files.

2. Within the working directory, you can set up a data directory called `data`. This is a sub-directory of your working directory.

When you set up an `R Project` in a directory, forever more that directory will be your working directory, at least for the work being done in that directory. Choose other directories to set up other projects.

Statistics tutorials

The practice of business analytics benefits greatly from advances in statistics and operations research. Here is a statistics primer that we can use to refresh our understanding and use of basic concepts and models often deployed in business analytics. You can access the Statistical Rethinking site here

Requirements

The course is officially online. However optional live sessions each week (they will be video'd!) will help keep all of us awake to one another's needs throughout the course including advice on solving some of the thornier problems we will face.

There are six (6) assignments to be submitted to the course Learning Management System (LMS) site. In addition there is a participation grade.

- Assignments 1 through 5 are each worth 13% of the final grade, that is, for a total of 65% of the final grade. These will be completed during weeks 1 through 5 of the course. They will typically include an upload of a model and a short interpretation of results.
- Assignment 6 is worth 25% of the final grade and will be completed during weeks 6 and 7 of the course. This last assignment is designed to pull course concepts together into a comprehensive system dynamics model replete with interpretation and ready for policy analysis. The final project artifacts will be deposited into an R Package hosted on the student's Github account.
- Final assessment prerequisites: the student presents a 5th week summary of course work to date and a proposal for the final project. Written hand-outs (3-5 pages or slides) and a video will constitute the artifacts for this assessment prerequisite. This is a binary (all or nothing) requirement and will serve to advise the student of gaps in knowledge and skill, all to position the student for a successful completion of the final project.
- Participation is worth 10% of the final grade. This grade will be primarily evidenced through posted comments on the course blog site regarding modeling vignettes and issues.

We will work in teams and submit individual assignments and posts. Individual final grades are A-F (integer ranges): A (>95), A- (90-94), B+ (85-89), B (80-84), C+ (75-79), C (70-74), D (65-69), F (<65).

Rubrics

Assessment scale for all assignments:

5 - The student simply commands the subject.

4 - There are smaller calculation, informational, inferential errors, minor misunderstandings or gaps in the student's knowledge

3 - Major misunderstandings or gaps in the student's knowledge, leaps of logic, or the student is sometimes unable to apply basic concepts, or there are some major calculation, informational, inferential errors: nevertheless the average performance is satisfying and remediation is possible.

2 - There are major gaps in the knowledge and skills of the student.

1 - The student can list concepts and results covered but cannot reproduce their essence and correctly use them in practical situations presented during the course.

0 - The student does not really know the subject at all!

Indicative integer thresholds (the instructor reserves the authority to judge the assignment of scores across the breadth and depth of observed coursework): 5 –95 pct, 4 – 85 pct, 3 – 75 pct, 2 – 60 pct, 1 – 55 pct of the maximum assessment score. These will be mapped to letter grades at the end of the course.

This general rubric guides assessment of coursework.

- **Words:** The text is laid out cleanly, with clear divisions and transitions between sections and subsections. The writing itself is well-organized, free of grammatical and other mechanical errors, divided into complete sentences, logically grouped into paragraphs and sections, and easy to follow from the presumed level of knowledge.
- **Numbers:** All numerical results or summaries are reported to suitable precision, and with appropriate measures of volatility and uncertainty attached when applicable.
- **Pictures:** All figures and tables shown are relevant to the argument for ultimate conclusions. Figures and tables are easy to read, with informative captions, titles, axis labels and legends, and are placed near the relevant pieces of text.
- **Code:** The code is formatted and organized so that it is easy for others to read and understand. It is indented, commented, and uses meaningful names. It only includes computations which are actually needed to answer the analytical questions, and avoids redundancy. Code borrowed from the notes, from books, or from resources found online is explicitly acknowledged and sourced in the comments. Functions or procedures not directly taken from the notes have accompanying tests which check whether the code does what it is supposed to. All code runs, and the `R Markdown` file `knits` to `html_document` output, or other output agreed with the instructor.
- **Modeling:** Model specifications are described clearly and in appropriate detail. There are clear explanations of how estimating the model helps to answer the analytical questions, and rationales for all modeling choices. If multiple models are compared, they are all clearly described, along with the rationale for considering multiple models, and the reasons for selecting one model over another, or for using multiple models simultaneously.
- **Inference:** The actual estimation and simulation of model parameters or estimated functions is technically correct. All calculations based on estimates are clearly explained, and also technically correct. All estimates or derived quantities are accompanied with appropriate measures of uncertainty. Information criteria are described and used to explain outcomes in terms of predictive power including volatility and uncertainty (e.g., over- and underfitting). Exceptions are noted, assumptions are challenged, recommendations for ongoing work are stated.
- **Conclusions:** The substantive, analytical questions are all answered as precisely as the data and the model allow. The chain of reasoning from estimation results about the model, or derived quantities, to substantive conclusions is both clear and convincing. Contingent answers (for example, "if X, then Y, but if A, then B, else C") are likewise described as warranted by the model and data. If uncertainties in the data and model mean the answers to some questions must be imprecise, this too is reflected in the conclusions.
- **Sources:** All sources used, whether in conversation, print, online, or otherwise, are listed and acknowledged where they used in code, words, pictures, and any other components of the analysis.

Other matters

Assignment Formatting

All assignments must be turned in electronically, through the learning management system, by each student. All assignments will involve writing a combination of code and actual prose. You must submit your assignment in a format which allows for the combination of the two, and the automatic execution of all your code. The easiest way to do this is to use `R Markdown`. `R Markdown` also allows the use of interactive modeling through `Shiny` applications.

Work submitted as Word files, unformatted plain text, etc., are not acceptable at any time during the course. Each assignment will require the submission of at least one `R Markdown` script file and the `html` file that the `R Markdown` script generates. When using data sets, this course will only use `csv` (comma separated variable files generated by Excel or in text files or calls to data using APIs. If the submission uses a `csv` file, that file must also be submitted with the `R Markdown` script and generated `html` output files. The student may also submit a supplemental `R` script file, suitably commented, that represents the `R` code chunks in the `R Markdown` script.

Managing the data base of submitted assignments throughout the course will be aided by standards including file name construction for assignment submission. To this end, every file submitted must have a file name which includes the student's name, course identifier, and clearly indicates the type of assignment (project) and its number (week). Here is the format we will use: `yourName_courseidentifier_Assignment#.ext`, where # is the week number and ext is the file name extension.

For example W.G. Foote would submit **RMarkdown** file with this filename:

- `wgfoote_MBAC611_Assignment1.Rmd`, where the file extension `Rmd` is the extension that `RStudio` uses for `R Markdown` documents.
- File extensions `R`, `html`, and `csv` are the other three admissible file types.

Course Specific Policies

Students are expected to behave in a professional and courteous manner at all times when interacting with all members of the course learning community. Respect for others is demonstrated through attendance, meaningful participation, and punctuality. Every effort should be made to be present for each session, if not feasible, view the recording of each session, especially since weekly assignments will be made conditional on content in live sessions.

All projects must be completed and submitted by the due dates and times set out. This will allow the entire class to review and revise submissions in a timely fashion. Submissions to `lms.manhattan.edu` are based on Eastern Time (UTC +5).

- Late submissions will result in student inability to accumulate the knowledge needed to advance to the next week's coverage of course topics.
- Late submission will also delay necessary instructor feedback to the student in a timely fashion.
- As the course continues to layer on more skills and capabilities, a late submission with inaccurate or incorrect implementations of financial applications will only deprecate the student's ability to successfully complete future assignments, let alone the final project.

So, don't be late. Due dates are posted in the LMS (Learning Management System) to help you pace your progress through the very quickly rolling out of the 7 weeks of the course term. Due dates are not deadlines. However, there is one deadline from the Registrar: all grades (including Incompletes) must be posted within 48 hours of the end of the course term.

Academic integrity

The Manhattan College Academic Integrity Policy holds students accountable for the integrity of the work they submit. Students should be familiar with the Policy and know that it is their responsibility to learn about instructor and general academic expectations with regard to proper collection, usage, and citation of sources in written work. The policy also governs the integrity of work submitted in exams and assignments as well as the veracity of signatures on attendance sheets and other verifications of participation in class activities. For more information and the complete policy, see the Manhattan College Catalog.

Students with disabilities

If you need academic accommodations due to a disability, then you should immediately register with the Director of the Specialized Resource Center (SRC). The SRC at Manhattan College authorizes special accommodations for students with disabilities. If you have a documented disability and you wish to discuss academic accommodations, please contact me within the first week of class.