

Decision Lab

Foundations of Decision-Making Tools for Competitive Strategy

Dr. Jose Mendoza
Competitive Strategy
Spring 2026

Table of Contents

1. Introduction: The Purpose of the Decision Lab
2. The Decision-Making Tools
3. The Analytic Hierarchy Process (AHP)
4. Scenario Planning
5. Decision Trees
6. Negotiation Theory and Simulation
7. Integrating the Tools: From Analysis to Action
8. References

1. Introduction: The Purpose of the Decision Lab

The Decision Lab is a suite of interactive tools designed to develop and sharpen the decision-making capabilities of students in competitive strategy courses. In the practice of strategy, managers and leaders are routinely confronted with decisions that are complex, uncertain, multi-dimensional, and often contested by parties with competing interests. Traditional classroom instruction provides the theoretical foundations for understanding these challenges, but theory alone is insufficient. Students must practice making decisions under realistic conditions to develop the judgment and analytical discipline that effective strategists require.

The Decision Lab addresses this gap by providing hands-on access to four distinct categories of decision-making tools, each grounded in a well-established body of academic theory. These tools are not abstract exercises; they simulate the kinds of analytical challenges that arise in real-world strategic decision-making, from evaluating competing alternatives against multiple criteria to planning for uncertain futures, quantifying risk, and negotiating agreements with counterparts who have their own objectives and constraints.

The four decision-making domains covered by the Decision Lab are:

- The Analytic Hierarchy Process (AHP) — for structured multi-criteria decision analysis
- Scenario Planning — for strategic thinking under deep uncertainty
- Decision Trees — for sequential decision-making under risk
- Negotiation Simulation — for multi-party bargaining and agreement design

Each of these domains has a rich intellectual history, and the tools in the Decision Lab are designed to bring that theory to life through direct student engagement. This document provides the theoretical background for each tool, offering students a foundation in the principles and methods they will apply during the course.

2. The Decision-Making Tools

The Decision Lab includes the following interactive tools, each corresponding to one of the four decision-making domains:

AHP Studio implements the Analytic Hierarchy Process developed by Thomas L. Saaty. It enables students to structure complex decisions as hierarchies of criteria and alternatives, perform pairwise comparisons to establish priority weights, check the consistency of their judgments, conduct sensitivity analysis, and generate comprehensive decision reports. The tool supports group decision-making through participant participation links, consensus measurement using Kendall's coefficient of concordance, and multi-round Delphi iteration for convergence.

Scenario Sim is a scenario planning simulation that guides students through the process of identifying driving forces, critical uncertainties, and predetermined elements to construct and compare multiple plausible future scenarios. It allows students to practice the structured imagination that scenario planning demands, developing strategic options that are robust across different possible futures.

Decision Tree Sandbox provides an interactive environment for constructing and analyzing decision trees. Students define decision nodes, chance nodes, and outcomes, assign probabilities and payoffs, and use the tool to calculate expected values and identify optimal decision paths. Sensitivity analysis features allow students to explore how changes in probabilities or payoffs affect the recommended course of action.

Negotiation Sim places students in structured negotiation exercises where they negotiate multi-issue deals against either an AI counterpart or fellow students. The simulation scores outcomes against an objective rubric that weighs economic value, strategic alignment, relationship preservation, and information management, providing quantitative feedback on negotiation performance across multiple rounds.

The sections that follow provide the theoretical foundations for each of these four domains.

3. The Analytic Hierarchy Process (AHP)

3.1 Origins and Purpose

The Analytic Hierarchy Process was developed by Thomas L. Saaty in the 1970s and formally presented in his 1980 book of the same name. AHP was conceived as a practical framework for making decisions that involve multiple criteria, particularly in situations where the decision-maker must weigh qualitative and quantitative factors simultaneously. Saaty's insight was that human beings are more reliable at making relative judgments between two items at a time than at producing absolute measurements of complex, intangible factors. AHP exploits this cognitive strength through a systematic process of pairwise comparisons.

AHP has since been applied to an extraordinarily wide range of problems, from corporate resource allocation and supplier selection to public policy evaluation, healthcare prioritization, and environmental management. Its strength lies in its ability to decompose an overwhelmingly complex decision into a hierarchy of manageable sub-problems, then mathematically synthesize the results into a coherent overall ranking.

3.2 The Hierarchical Structure

Every AHP analysis begins with structuring the decision as a hierarchy. At the top sits the goal: the overarching objective the decision-maker is trying to achieve. Below the goal are the criteria, which are the dimensions along which alternatives will be evaluated. Criteria may themselves be decomposed into sub-criteria if greater specificity is needed. At the bottom of the hierarchy are the alternatives: the concrete options under consideration.

For example, a company choosing a new market to enter might define the goal as "Select the most attractive market for expansion." The criteria might include market size, competitive intensity, regulatory environment, and cultural alignment. Sub-criteria under regulatory environment might include intellectual property protection, tariff levels, and tax incentives. The alternatives might be three candidate countries.

3.3 Pairwise Comparisons and the Fundamental Scale

Once the hierarchy is established, the decision-maker performs pairwise comparisons at each level. For each pair of criteria, the decision-maker answers the question: “With respect to the goal, how much more important is Criterion A than Criterion B?” Judgments are recorded using Saaty’s fundamental scale, which ranges from 1 (equal importance) to 9 (extreme importance of one element over another). The reciprocal property ensures that if A is judged five times as important as B, then B is automatically one-fifth as important as A.

These comparisons produce a reciprocal matrix. For a set of n criteria, the matrix is $n \times n$, with ones along the diagonal and reciprocal values in symmetric positions. The principal eigenvector of this matrix yields the priority weights, which represent the relative importance of each criterion on a ratio scale. The eigenvalue method has been shown to produce meaningful priority vectors that preserve the rank order and proportional relationships implied by the original judgments (Saaty, 1980).

3.4 Consistency and the Consistency Ratio

A critical feature of AHP is its built-in mechanism for detecting inconsistency in judgments. Perfect consistency would mean that if A is twice as important as B, and B is three times as important as C, then A should be six times as important as C. In practice, human judgments are rarely perfectly consistent, but excessive inconsistency signals that the decision-maker’s judgments are unreliable and should be revised.

AHP measures consistency through the Consistency Ratio (CR), which compares the consistency of the actual judgments to the consistency that would be expected from a matrix of random judgments of the same size. Saaty proposed a threshold of $CR \leq 0.10$ (10%) as acceptable; higher values suggest the decision-maker should revisit their comparisons. This self-correcting feature is one of AHP’s most important contributions to decision methodology.

3.5 Synthesis and Sensitivity Analysis

After priority weights are computed for criteria and for alternatives with respect to each criterion, global synthesis combines them into an overall ranking. Each alternative’s global priority is the weighted sum of its local priorities under each criterion, where the weights are the criteria

priorities. The result is a single ranking of alternatives that reflects all the criteria and their relative importance.

Sensitivity analysis is the final step, and it addresses the question: “How robust is this ranking?” By systematically varying the weight assigned to each criterion, the analyst can observe whether the top-ranked alternative changes. If the ranking is stable across a wide range of weight variations, the decision-maker can proceed with confidence. If small changes in weights produce rank reversals, the decision is more sensitive and warrants additional deliberation.

3.6 Group Decision-Making with AHP

AHP can be extended to group decision-making by aggregating individual pairwise comparison matrices. The standard approach uses the geometric mean of corresponding matrix entries to produce a consensus matrix (Saaty and Vargas, 2012). Kendall’s coefficient of concordance (W) can be used to measure the degree of agreement among group members, and the Delphi method can be employed to iteratively converge on shared priorities through multiple rounds of anonymous comparison and feedback.

3.7 Example: Selecting a Cloud Infrastructure Provider

A mid-sized software company must choose among three cloud infrastructure providers—AWS, Azure, and Google Cloud—for hosting a new product line. The decision involves trade-offs across multiple dimensions that are difficult to compare directly. The CTO assembles a team and structures the problem using AHP.

Step 1: Define the hierarchy. The goal is “Select the best cloud provider.” The team identifies four criteria: Cost (monthly and annual pricing tiers), Reliability (uptime guarantees and track record), Scalability (ability to handle traffic spikes and geographic expansion), and Developer Experience (documentation quality, tooling, and existing team expertise). The three alternatives sit at the bottom of the hierarchy.

Step 2: Pairwise comparisons of criteria. The team compares the four criteria in pairs with respect to the goal. They judge Reliability to be strongly more important than Cost (5 on Saaty’s scale), Scalability moderately more important than Cost (3), and Developer Experience equally to moderately more important than Cost (2). They judge Reliability slightly more important than

Scalability (2) and moderately more important than Developer Experience (3). Scalability and Developer Experience are judged roughly equal (1). These six judgments populate a 4×4 reciprocal matrix.

Table 1. Pairwise Comparison Matrix (Criteria vs. Goal)

	Cost	Reliability	Scalability	Dev. Exp.
Cost	1	1/5	1/3	1/2
Reliability	5	1	2	3
Scalability	3	1/2	1	1
Dev. Exp.	2	1/3	1	1

Values above the diagonal are direct judgments; values below are reciprocals.

Step 3: Compute priorities and check consistency. The eigenvector method yields the following priority weights. The Consistency Ratio is 0.03, well below the 0.10 threshold, confirming that the judgments are logically coherent.

Table 2. Criteria Priority Weights

Criterion	Priority Weight
Reliability	0.47
Scalability	0.23
Developer Experience	0.18
Cost	0.12

Consistency Ratio = 0.03 (acceptable threshold: ≤ 0.10)

Step 4: Compare alternatives under each criterion. The team then performs pairwise comparisons of the three providers with respect to each criterion. For example, under Reliability, AWS is judged moderately better than Azure (3) and strongly better than Google Cloud (5); Azure is judged moderately better than Google Cloud (3). Similar comparison matrices are constructed for Cost, Scalability, and Developer Experience.

Step 5: Synthesize and analyze. Global synthesis combines the criteria weights with the local alternative priorities:

Table 3. Global Synthesis — Final Ranking

Alternative	Cost (0.12)	Reliab. (0.47)	Scalab. (0.23)	Dev. Exp. (0.18)	Global Priority
AWS	0.26	0.54	0.43	0.40	0.45

Azure	0.33	0.30	0.35	0.35	0.33
Google Cloud	0.41	0.16	0.22	0.25	0.22

AWS ranks first. Sensitivity analysis shows the ranking is stable unless Cost weight exceeds 0.35.

The team concludes that AWS is the robust choice under their current priorities, and the decision report documents the rationale for stakeholder review.

This example illustrates how AHP converts a subjective, multi-dimensional decision into a transparent, defensible ranking with built-in consistency checking and sensitivity testing.

4. Scenario Planning

4.1 Origins and Purpose

Scenario planning emerged as a strategic planning discipline at the RAND Corporation in the 1950s and was developed into a practical business methodology by Pierre Wack and his colleagues at Royal Dutch/Shell in the 1970s. Peter Schwartz, who worked with Wack at Shell and later founded the Global Business Network, popularized the approach in his 1991 book, *The Art of the Long View*. The methodology has since been adopted widely in corporate strategy, public policy, military planning, and organizational foresight.

The central premise of scenario planning is that the future is inherently uncertain, and that traditional forecasting methods, which attempt to predict a single most likely future, are inadequate for strategic decision-making in volatile, complex environments. Instead of asking “What will happen?”, scenario planning asks “What could happen?” and “Are we prepared for it?”

4.2 The Logic of Scenarios

A scenario is not a prediction; it is a plausible, internally consistent narrative about how the future might unfold. The purpose of developing multiple scenarios (typically two to four) is not to identify the “right” future but to stretch the organization’s thinking and prepare it for a range of possibilities. Scenarios challenge the “official future”—the implicit, often unexamined assumptions about the trajectory of the business environment—and reveal strategic vulnerabilities and opportunities that would otherwise remain hidden.

The value of scenario planning lies not only in the scenarios themselves but in the process of creating them. The structured conversation required to identify driving forces, isolate critical uncertainties, and construct coherent narratives forces decision-makers to surface and examine their assumptions, consider perspectives they might otherwise ignore, and develop a shared strategic vocabulary (van der Heijden, 2005).

4.3 The Scenario Planning Process

While variations exist, most scenario planning processes follow a common sequence of steps:

1. Identify the focal issue or decision. What strategic question does the organization need to answer? This anchors the entire exercise.
2. Identify the key forces in the local environment. What factors in the immediate business context (customers, competitors, suppliers, regulators) will most influence the outcome?
3. Identify the driving forces. What macro-level trends and forces (technological, economic, political, social, environmental) are shaping the broader landscape?
4. Rank by importance and uncertainty. Not all forces are equally important or equally uncertain. The most useful scenario axes are built from forces that are both highly important and highly uncertain.
5. Select scenario logics. Choose two or three critical uncertainties to serve as the axes of variation. These define the structural differences between scenarios.
6. Flesh out the scenarios. Develop each scenario into a coherent narrative, integrating the driving forces and predetermined elements into a story of how the future might unfold.
7. Assess implications. For each scenario, ask: What would this mean for our strategy? What opportunities and threats would emerge? What capabilities would we need?
8. Identify signposts. Determine early indicators that would signal which scenario (or combination of scenarios) is beginning to materialize, enabling early strategic adaptation.

4.4 Predetermined Elements and Critical Uncertainties

A key analytical distinction in scenario planning is between predetermined elements and critical uncertainties. Predetermined elements are trends or forces that are highly likely to continue regardless of which scenario unfolds. Demographic shifts, the retirement of a generation of workers, or the expiration of a major patent are examples of predetermined elements. These factors appear in every scenario; they are the fixed backdrop against which the uncertain dynamics play out.

Critical uncertainties, by contrast, are the forces whose future direction is genuinely unknown and whose resolution would have significant strategic consequences. The outcome of a pending regulation, the pace of technological adoption, or the response of a key competitor are examples.

These uncertainties form the axes of the scenario framework and define the structural differences between alternative futures.

4.5 From Scenarios to Strategy

The ultimate objective of scenario planning is not to produce scenarios for their own sake but to improve the quality of strategic decisions. A robust strategy is one that performs acceptably across multiple scenarios, rather than one that is optimized for a single predicted future.

Schwartz (1991) emphasizes that scenarios serve as “wind tunnels” for testing strategic options: if a proposed strategy fails catastrophically under one plausible scenario, it should be modified or complemented with contingency plans.

Scenario planning also builds organizational resilience by expanding the mental models of decision-makers. Leaders who have rehearsed multiple futures are less likely to be caught off guard by unexpected developments and more likely to recognize emerging patterns early enough to respond effectively.

4.6 Example: An Automaker’s Electric Vehicle Strategy

A traditional automaker is deciding how aggressively to invest in electric vehicle (EV) manufacturing over the next decade. The strategic planning team uses scenario planning to stress-test their investment options.

Focal issue: Should the company commit \$8 billion to building dedicated EV production lines, or pursue a more flexible hybrid approach that can produce both internal combustion and electric vehicles?

Driving forces and critical uncertainties: The team identifies dozens of forces shaping the industry. After ranking them by importance and uncertainty, two critical uncertainties emerge as the primary axes of variation. The first is the pace of government regulation: governments may impose aggressive EV mandates with near-term deadlines, or they may adopt a slower, market-driven transition with weaker regulatory pressure. The second is the trajectory of battery technology costs: a breakthrough in solid-state batteries could cut costs dramatically, or costs could decline only incrementally along current lithium-ion improvement curves.

These two uncertainties define a 2×2 matrix of four scenarios:

Table 4. Scenario Matrix — Electric Vehicle Strategy

	Battery Breakthrough	Incremental Battery Improvement
Aggressive Regulation	A — Electric Surge <i>Fast, mandate-driven transition; early movers dominate</i>	B — Policy Push <i>Mandates outpace economics; margin pressure on all OEMs</i>
Weak Regulation	C — Tech Leapfrog <i>Market-pull transition; cost advantage to tech leaders</i>	D — Slow Lane <i>Gradual shift; combustion dominates; EV overcapacity punished</i>

Scenario A — “Electric Surge” (aggressive regulation + battery breakthrough): Governments ban new combustion vehicle sales by 2032. Solid-state batteries cut pack costs below \$50/kWh, making EVs cheaper than combustion vehicles without subsidies. Consumer adoption accelerates rapidly. Companies that invested early in dedicated EV lines capture dominant market share. Late movers face stranded combustion assets and a scramble for battery supply.

Scenario B — “Policy Push” (aggressive regulation + incremental battery improvement): Governments impose strict EV mandates, but battery costs decline slowly. EVs remain more expensive than combustion vehicles, and consumers comply reluctantly. Manufacturers face margin pressure as they must sell EVs at a loss to meet quotas. Companies with flexible production lines that can adjust output ratios between EV and combustion models manage the transition more profitably.

Scenario C — “Tech Leapfrog” (weak regulation + battery breakthrough): Governments take a hands-off approach, but cheap, high-performance batteries make EVs irresistibly attractive to consumers on pure economics. Adoption is driven by market pull rather than regulatory push. The transition is fast but uneven across regions. First movers in EV technology gain a cost advantage, but the absence of mandates means combustion vehicles remain legal and continue to sell in some segments.

Scenario D — “Slow Lane” (weak regulation + incremental battery improvement): Neither regulation nor economics forces a rapid transition. EVs grow their market share gradually but remain a premium segment. Combustion and hybrid vehicles dominate for another decade.

Companies that over-invested in EV capacity face underutilized factories. The flexible hybrid approach delivers the best returns.

Strategic implications: The team tests both investment options against all four scenarios:

Table 5. Strategy Performance Across Scenarios

Strategy	A: Electric Surge	B: Policy Push	C: Tech Leapfrog	D: Slow Lane
\$8B Dedicated EV	Excellent	Adequate	Good	Poor
Flexible Hybrid	Adequate	Good	Adequate	Good
Staged Commitment	Good	Good	Good	Adequate

Shaded cells indicate the best-performing strategy for that scenario.

The team recommends a staged commitment: invest \$3 billion in a first dedicated EV line immediately (capturing upside in the favorable scenarios), while designing the remaining capacity for flexible conversion and monitoring signposts—battery cost milestones, regulatory drafts, and consumer adoption rates—that would trigger the full commitment.

The example demonstrates how scenario planning does not seek to predict which future will arrive, but instead produces strategies that are resilient across multiple plausible futures.

5. Decision Trees

5.1 Origins and Purpose

Decision tree analysis is one of the foundational tools of decision science. The methodology was formalized in the 1960s, most notably by Howard Raiffa in his 1968 work *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Decision trees provide a structured, visual framework for analyzing sequential decisions under risk, where the outcomes of current decisions depend on uncertain future events and on subsequent decisions that will be made in response to those events.

The power of decision trees lies in their ability to make the structure of a decision problem explicit. By mapping out the full sequence of decisions, chance events, and outcomes, the analyst can evaluate complex problems that would be difficult or impossible to reason through intuitively. Decision trees are widely used in business strategy, capital investment analysis, pharmaceutical development, litigation, and resource management.

5.2 Structure of a Decision Tree

A decision tree consists of three types of nodes connected by branches:

- Decision nodes (represented as squares) are points where the decision-maker must choose among alternatives. Each branch from a decision node represents one available option.
- Chance nodes (represented as circles) are points where the outcome is determined by uncertainty. Each branch from a chance node represents a possible state of nature, and each branch is assigned a probability reflecting the likelihood of that state.
- Terminal nodes (represented as endpoints) indicate the final outcomes of a particular sequence of decisions and chance events. Each terminal node is assigned a payoff value reflecting the desirability of that outcome.

The tree is read from left to right, following the chronological sequence of decisions and events. Decision nodes come first when the decision-maker must act before knowing the outcome of an uncertain event; chance nodes come first when the uncertainty is resolved before the decision is made.

5.3 Expected Value and the Rollback Method

The standard method for solving a decision tree is backward induction, also called the rollback or folding-back method. Starting from the terminal nodes on the right side of the tree and working backward:

1. At each chance node, compute the expected monetary value (EMV) as the probability-weighted average of the values of the branches emanating from that node.
2. At each decision node, select the branch with the highest expected value (assuming the decision-maker is an expected value maximizer).
3. Continue rolling back until the root of the tree is reached, at which point the optimal initial decision is identified.

The expected value framework assumes that the decision-maker is risk-neutral—that is, indifferent between a certain payoff and an uncertain gamble with the same expected value. For decisions involving large stakes relative to the decision-maker's wealth, this assumption may not hold, and utility theory can be incorporated to reflect risk aversion or risk-seeking preferences (von Neumann and Morgenstern, 1944).

5.4 The Value of Information

One of the most powerful applications of decision trees is calculating the value of information. Before committing to a decision, the decision-maker often has the option to gather additional information—through market research, prototype testing, expert consultation, or other means. Decision trees allow the analyst to compute the Expected Value of Perfect Information (EVPI), which is the maximum amount the decision-maker should be willing to pay for information that would completely resolve the uncertainty before the decision is made.

The EVPI is the difference between the expected value of the optimal decision made with perfect information and the expected value of the optimal decision made without that information. If the EVPI is small relative to the cost of the information, the decision-maker should proceed without it; if the EVPI is large, further investigation is warranted. This framework transforms the often vague question of “should we do more research?” into a precise economic calculation.

5.5 Sensitivity Analysis in Decision Trees

Real-world probability and payoff estimates are inherently uncertain. Sensitivity analysis addresses this by systematically varying key assumptions and observing how the optimal decision changes. One common approach is to vary a single probability or payoff value across its plausible range while holding all other values fixed, producing a sensitivity chart that shows the threshold values at which the optimal decision switches.

When multiple uncertain parameters interact, scenario analysis or Monte Carlo simulation can be used to explore the joint effects. The result is a richer understanding of which uncertainties matter most for the decision and where additional information would be most valuable.

5.6 Example: A Pharmaceutical Company's Drug Development Decision

A pharmaceutical company has a promising compound for treating a chronic respiratory condition. The company must decide whether to invest \$40 million in Phase III clinical trials. If the trials succeed, the drug could be submitted to regulators for approval. The decision tree maps out the sequence of decisions and uncertainties.


Decision node (root): Invest \$40M in Phase III trials, or abandon the compound and license it to a larger firm for a guaranteed \$12M.

Chance node 1 (trial outcome): Based on Phase II results and therapeutic area benchmarks, the team estimates a 60% probability that trials succeed and a 40% probability they fail. If trials fail, the \$40M investment is lost entirely (net payoff: $-\$40\text{M}$).

Decision node 2 (post-success): If trials succeed, the company can submit for regulatory approval (additional \$5M cost) or sell the successful trial data to a partner for \$30M (net payoff: $\$30\text{M} - \$40\text{M} = -\$10\text{M}$, but with certainty).

Chance node 2 (regulatory outcome): If the company submits, there is an 80% chance of approval and a 20% chance of rejection. If approved, projected net present value of future revenue (after all costs) is \$200M. If rejected, the net payoff is $-\$45\text{M}$ (the \$40M trial investment plus \$5M submission cost).

The following diagram illustrates the complete decision tree for this problem. Squares represent decision nodes, circles represent chance nodes, and diamonds represent terminal payoffs. The optimal path is highlighted.



Solving by backward induction:

Starting from the terminal nodes and working backward, the EMV at each node is computed:

Table 6. Decision Tree Rollback Calculation

Node	Type	Calculation	Value
Regulatory outcome	Chance	$(0.80 \times \$200M) + (0.20 \times -\$45M)$	\$151.0M
Post-success choice	Decision	$\max(\text{Submit: } \$151M, \text{ Sell data: } -\$10M)$	\$151.0M
Trial outcome	Chance	$(0.60 \times \$151M) + (0.40 \times -\$40M)$	\$74.6M
Root decision	Decision	$\max(\text{Invest: } \$74.6M, \text{ License: } \$12M)$	\$74.6M

Optimal strategy: Invest in Phase III trials, then submit for regulatory approval if trials succeed.

Value of information: Suppose the company could commission a \$2M biomarker study that would perfectly predict the trial outcome before committing the \$40M:

Table 7. Expected Value of Perfect Information (EVPI)

Measure	Value
EMV without information	\$74.6M
EMV with perfect information	\$95.4M

EVPI (difference)	\$20.8M
Cost of biomarker study	\$2.0M

Since EVPI (\$20.8M) far exceeds the study cost (\$2M), the study is clearly worth commissioning.

Sensitivity analysis: The team explores how the decision changes if the probability of trial success varies. At a trial success probability of approximately 17%, the EMV of investing drops to \$12M (equal to the licensing alternative), and below that threshold, licensing becomes the preferred choice. Since the team's 60% estimate is well above this threshold, the decision to invest is robust to substantial downward revision of the success probability.

This example illustrates how decision trees transform a complex, multi-stage decision under uncertainty into a structured calculation, and how EVPI analysis provides a rigorous answer to the question of whether to invest in additional information before committing.

6. Negotiation Theory and Simulation

6.1 Origins and Purpose

Negotiation is among the most fundamental skills in business and strategy. Whether negotiating contracts, joint ventures, labor agreements, or mergers and acquisitions, the ability to reach agreements that create and capture value is essential to competitive success. The modern academic study of negotiation draws on multiple intellectual traditions, including game theory, behavioral economics, psychology, and law.

Two landmark works shaped the field as it is practiced today. Howard Raiffa's *The Art and Science of Negotiation* (1982) provided a rigorous analytical framework grounded in decision analysis and game theory, while Roger Fisher and William Ury's *Getting to Yes* (1981) introduced the concept of principled negotiation and made negotiation theory accessible to a broad audience. Together, these works established the dual emphasis on analytical rigor and interpersonal skill that characterizes modern negotiation education.

6.2 Distributive and Integrative Bargaining

Negotiation theory distinguishes between two fundamental paradigms. Distributive bargaining (also called zero-sum or win-lose negotiation) involves dividing a fixed pool of value between the parties. In a purely distributive negotiation, any gain for one party comes at the expense of the other. Classic examples include haggling over price, dividing a budget, or splitting a territory. The key concept in distributive bargaining is the Zone of Possible Agreement (ZOPA): the range between the parties' reservation prices (the worst deal each party would accept). If the ZOPA exists, agreement is possible; if not, the parties will walk away.

Integrative bargaining (also called win-win or value-creating negotiation) recognizes that most real-world negotiations involve multiple issues on which the parties have different priorities. By trading concessions on issues of lower priority for gains on issues of higher priority, both parties can achieve outcomes that are superior to simple compromise. Fisher and Ury (1981) captured this insight with their principle of "focusing on interests, not positions," arguing that negotiators should look beyond their stated demands to understand the underlying needs and concerns that drive those demands.

6.3 BATNA and Reservation Value

One of the most important concepts in negotiation theory is the Best Alternative to a Negotiated Agreement (BATNA), introduced by Fisher and Ury (1981). A party's BATNA is the best course of action available if the current negotiation fails. The value of the BATNA sets the reservation value: the minimum outcome the party should accept in the negotiation. Any proposed agreement that is worse than the BATNA should be rejected.

BATNA is powerful because it shifts the focus from the psychology of the bargaining table to the objective reality of outside options. A negotiator with a strong BATNA can afford to take greater risks, make fewer concessions, and walk away from unfavorable deals. Conversely, a negotiator with a weak BATNA is in a vulnerable position. Improving one's BATNA before and during a negotiation is therefore one of the most effective strategies available.

6.4 Multi-Issue Negotiation and Scoring

In practice, most strategic negotiations involve multiple issues that must be resolved simultaneously. A brand partnership, for example, might involve negotiations over licensing fees, royalty rates, campaign duration, exclusivity provisions, performance bonuses, termination clauses, and usage rights. Each party assigns different weights to these issues based on their strategic priorities and constraints.

Raiffa (1982) demonstrated that multi-issue negotiations can be analyzed using scoring systems that assign points to outcomes on each issue, weighted by their importance. This transforms a complex, seemingly intractable negotiation into a structured optimization problem. The scoring framework also reveals the potential for value creation: when the parties' priorities differ, trades across issues can generate joint gains that make both parties better off than any equal-split compromise.

6.5 The Role of Information

Information asymmetry is a defining feature of most negotiations. Each party typically knows more about their own priorities, constraints, and alternatives than the other side does. The strategic management of information—what to reveal, what to conceal, and when to share—is a central challenge in negotiation.

Revealing information about one's priorities can facilitate value-creating trades, but it also creates vulnerability to exploitation. Experienced negotiators develop skills in strategic disclosure: sharing enough information to enable integrative agreements while protecting sensitive details about reservation values and alternatives. The Negotiation Sim scores students on information management alongside economic value and relationship preservation, reinforcing the multi-dimensional nature of negotiation skill.

6.6 Behavioral Dimensions

While analytical frameworks provide the structural foundation for negotiation, the practice of negotiation is deeply influenced by psychological and behavioral factors. Anchoring effects, framing biases, overconfidence, reactive devaluation (dismissing proposals simply because they come from the other side), and emotional escalation can all distort negotiation outcomes.

Research by Bazerman and Neale (1992) demonstrated that even experienced negotiators frequently fail to achieve efficient outcomes due to systematic cognitive biases. Awareness of these biases is a necessary complement to analytical skill: the most effective negotiators combine rigorous preparation with the ability to recognize and manage psychological dynamics at the table.

6.7 Example: Negotiating a Brand Partnership Agreement

A consumer electronics startup (NovaTech) wants to co-brand a new line of smart home devices with a major retailer (HomeFirst). The negotiation involves eight issues, and each side has different priorities. This example illustrates how multi-issue negotiation creates opportunities for value creation.

The issues and starting positions:

Table 8. Negotiation Issues and Opening Positions

Issue	NovaTech (Asks)	HomeFirst (Offers)
Base licensing fee	\$500,000/year	\$200,000/year
Royalty rate	8%	3%
Campaign duration	36 months	12 months
Exclusivity	Non-exclusive	Exclusive (smart home)
Performance bonus	2% above \$5M sales	None

Termination notice	12 months	3 months
Usage rights	Limited co-branding	Full sublicensing
Min. marketing spend	\$1,000,000/year	\$250,000/year

Hidden priorities and the potential for trades: NovaTech’s primary concerns are the royalty rate and campaign duration—the startup needs revenue and time to build brand recognition. The base fee matters less because NovaTech values the marketing exposure that HomeFirst’s retail network provides. HomeFirst, conversely, cares most about exclusivity and usage rights—the retailer wants to prevent NovaTech from cutting similar deals with competitors and needs broad rights to feature the brand across its marketing channels. HomeFirst is relatively flexible on the royalty rate because its margins on smart home products are strong.

Identifying the ZOPA: NovaTech’s BATNA is a smaller deal with an online-only retailer worth approximately \$180,000/year in guaranteed revenue and moderate brand exposure. HomeFirst’s BATNA is to launch its own private-label smart home line, which would cost approximately \$2M in development and carry execution risk. Both parties are better off reaching agreement, so a ZOPA exists.

Value-creating trades: By focusing on interests rather than positions, the negotiators discover several high-value trades:

Table 9. Agreed Terms — Final Deal vs. Opening Positions

Issue	NovaTech Ask	HomeFirst Ask	Agreed Term	Who Conceded?
Base fee	\$500K	\$200K	\$300K	NovaTech
Royalty rate	8%	3%	6%	Both
Campaign duration	36 mo.	12 mo.	24 mo.	Both
Exclusivity	Non-excl.	Exclusive	Exclusive	NovaTech
Perf. bonus	2% > \$5M	None	1.5% > \$3M	Both
Termination	12 mo.	3 mo.	6 mo.	Both
Usage rights	Limited	Full	Co-brand w/ approval	Both
Min. mktg. spend	\$1M	\$250K	\$750K	Both

The key integrative trade: NovaTech conceded exclusivity (high value to HomeFirst) in exchange for duration, royalty, and marketing spend (high value to NovaTech).

Scoring the outcome: Using a weighted scoring rubric, NovaTech evaluates the negotiated result:

Table 10. NovaTech Negotiation Scorecard

Dimension	Weight	Assessment	Score (1–100)
Economic Value	40%	Projected \$680K/yr vs. \$180K BATNA	82
Strategic Alignment	30%	24-mo. duration + \$750K mktg. commitment	78
Relationship	20%	Both sides' top priorities met; goodwill	85
Info. Management	10%	BATNA protected; learned rival's costs	75
Weighted Total	100%		80.6

$$\text{Score} = (82 \times 0.40) + (78 \times 0.30) + (85 \times 0.20) + (75 \times 0.10) = 80.6 \text{ out of } 100.$$

This example demonstrates how multi-issue negotiations, when approached with an integrative mindset, can generate outcomes that are superior to simple split-the-difference compromises. The key insight is that differences in priorities create the opportunity for mutually beneficial trades.

7. Integrating the Tools: From Analysis to Action

The four decision-making frameworks presented in this document are not isolated techniques; they are complementary perspectives on the common challenge of making good decisions under uncertainty and complexity. In the practice of competitive strategy, these tools are often used in combination.

Scenario planning establishes the range of possible futures a firm might face, identifying the critical uncertainties that will shape the competitive landscape. Decision trees provide a rigorous framework for evaluating strategic options under those uncertainties, quantifying the expected value of different courses of action and the value of gathering additional information before committing. AHP offers a structured approach for evaluating alternatives when the decision involves multiple qualitative and quantitative criteria that must be weighed against one another. And negotiation theory provides the frameworks needed to reach and structure agreements with partners, competitors, regulators, and other stakeholders whose cooperation is essential to executing a chosen strategy.

The Decision Lab brings these tools together in a single environment, allowing students to develop fluency in each methodology and, more importantly, to develop the strategic judgment needed to know which tool is appropriate for which type of decision problem. The goal is not to produce technicians who can mechanically apply a formula, but strategists who can diagnose a decision, select the right analytical approach, execute it with rigor, and act on the results with confidence.

8. References

- Bazerman, M. H., & Moore, D. A. (2013). *Judgment in managerial decision making* (8th ed.). Wiley.
- Bazerman, M. H., & Neale, M. A. (1992). *Negotiating rationally*. Free Press.
- Clemen, R. T., & Reilly, T. (2013). *Making hard decisions with DecisionTools* (3rd ed.). Cengage Learning.
- De Felice, F., & Petrillo, A. (Eds.). (2025). *The art of decision making: Applying AHP in practice*. IntechOpen.
- Fisher, R., Ury, W., & Patton, B. (2011). *Getting to yes: Negotiating agreement without giving in* (3rd ed.). Penguin. (Original work published 1981)
- Hammond, J. S., Keeney, R. L., & Raiffa, H. (1999). *Smart choices: A practical guide to making better decisions*. Harvard Business School Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Lax, D. A., & Sebenius, J. K. (1986). *The manager as negotiator: Bargaining for cooperation and competitive gain*. Free Press.
- Malhotra, D., & Bazerman, M. H. (2007). *Negotiation genius: How to overcome obstacles and achieve brilliant results at the bargaining table and beyond*. Bantam Books.
- Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Addison-Wesley.
- Raiffa, H. (1982). *The art and science of negotiation*. Harvard University Press.
- Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. McGraw-Hill.
- Saaty, T. L., & Vargas, L. G. (2012). *Models, methods, concepts & applications of the analytic hierarchy process* (2nd ed.). Springer.
- Schoemaker, P. J. H. (1995). Scenario planning: A tool for strategic thinking. *Sloan Management Review*, 36(2), 25–40.

- Schoemaker, P. J. H. (2022). *Advanced introduction to scenario planning*. Edward Elgar.
- Schwartz, P. (1991). *The art of the long view: Planning for the future in an uncertain world*. Doubleday/Currency.
- Spetzler, C., Winter, H., & Meyer, J. (2016). *Decision quality: Value creation from better business decisions*. Wiley.
- Thompson, L. L. (2020). *The mind and heart of the negotiator* (7th ed.). Pearson.
- van der Heijden, K. (2005). *Scenarios: The art of strategic conversation* (2nd ed.). Wiley.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Wack, P. (1985). Scenarios: Uncharted waters ahead. *Harvard Business Review*, 63(5), 72–89.

End of Reading Materials