



Department of Economics and Finance
Corso di Laurea in Economics and Business

Chair: Management

Managerial Decision-Making Under Uncertainty: A Strategic and Organizational Perspective

Prof. Francesca Vicentini

Mehmet Poyraz Ercins

SUPERVISOR

CANDIDATE

Academic Year 2025/2026

Table of Contents

Introduction

Chapter 1: Decision-Making Under True Uncertainty: Theoretical Foundations

- 1.1 Risk and True Uncertainty
- 1.2 Bounded Rationality and Satisficing
- 1.3 Heuristics and Cognitive Biases
- 1.4 Behavioral Distortions in Strategic Contexts

Chapter 2: Strategic Decision Processes in Uncertain Environments

- 2.1 Process Rationality and Decision Effectiveness
- 2.2 Speed and Comprehensiveness in High-Velocity Contexts
- 2.3 Structured Mechanisms for Reducing Bias
- 2.4 The Structure–Flexibility Trade-off

Chapter 3: Toward a Framework for Decision Quality Under Uncertainty

- 3.1 Integrating Behavioral and Process Perspectives
- 3.2 Development of the Conceptual Framework
- 3.3 Theoretical Implications
- 3.4 Limitations and Future Research

Conclusion

References

Introduction

Decision-making is, in some form, what management is. Any resource allocation, any strategic commitment, runs through a decision made by someone who doesn't know what's going to happen next. The gap between how management theory has traditionally described that process and how it actually works in organizations is, to put it generously, substantial.

The classical model assumed rationality: a decision-maker with clear preferences, access to relevant information, and the capacity to identify the value-maximizing option. It was mathematically elegant. It was also a poor description of what actually happens. Seventy years of evidence from economics, psychology, and management research makes that hard to dispute. Decision-makers don't optimize. They satisfice¹. They rely on heuristics that work well under some conditions and badly under others. They're subject to cognitive biases they can't reliably detect or correct². And they operate inside organizations with structural biases, political dynamics, and attention bottlenecks that compound individual errors in ways nobody fully anticipated.

What follows from that evidence is the subject of this thesis. Decision quality under genuine uncertainty can't be understood by looking at behavioral factors and process design separately. The two are interdependent in ways that matter practically. Cognitive constraints shape what any process can actually deliver. Process design determines which cognitive constraints cause the most damage. And the organizational context in which both operate shapes which configurations produce good decisions and which don't.

The argument develops across three chapters. Chapter 1 covers the theoretical foundations: Knight's distinction between calculable risk and true uncertainty³, Simon's account of bounded rationality and satisficing⁴, the heuristics and biases program of Kahneman and Tversky⁵, and the organizational amplification of cognitive limits documented by Cyert and March⁶ and Hambrick and Mason⁷. Chapter 2 turns to what the empirical literature says about strategic decision processes: how process design affects decision outcomes, how speed

¹ Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118.

² Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.

³ Knight, F. H. (1921). *Risk, uncertainty, and profit*. University of Chicago Press.

⁴ Simon, H. A. (1955). *A behavioral model of rational choice*, 1955.

⁵ Kahneman & Tversky, *Judgment under uncertainty: Heuristics and biases*, 1974.

⁶ Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Prentice-Hall.

⁷ Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 193–206.

and comprehensiveness trade off across different environments, and what structured mechanisms can do to limit the damage cognitive biases cause. Chapter 3 proposes an integrative framework, draws out the implications for behavioral strategy research⁸, and identifies where the field still needs empirical work.

Two research questions organize the analysis. The first: how does genuine uncertainty affect strategic decision-making, given the limits of bounded rationality and the distortions introduced by cognitive bias? The second: how can decision processes be designed to improve decision quality under true uncertainty, particularly in balancing analytical structure against adaptive flexibility? The questions are related but distinct. The first is mainly descriptive, asking where and why decision-making fails. The second is mainly prescriptive, asking what conditions make it possible to do better.

A note on scope. This thesis is a theoretical contribution, not an empirical one. It draws on a defined body of papers in strategic management, organizational behavior, and decision theory to build an integrated conceptual framework. No new data is presented and no hypotheses are tested in the formal sense. The contribution is in the synthesis: bringing behavioral and process perspectives together that have largely developed in parallel, and identifying where future empirical work should focus. The final chapter argues that this kind of groundwork is both legitimate and necessary as a precondition for the more specific empirical work the field needs.

⁸ Powell, T. C., Lovullo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1369–1386.

Chapter 1: Decision-Making Under True Uncertainty: Theoretical Foundations

1.1 Risk and True Uncertainty

Decision theory for most of the twentieth century had a solution to the problem of an unknowable future: converting it into something calculable. Uncertainty became risk, outcomes got assigned probabilities, and the rest was left to expected utility theory. Assign numbers. Weight by likelihood. Optimize. Clean. Logical. And, as it turns out, only partially right.

Frank Knight, writing in 1921, saw the problem earlier than most. His distinction between risk and uncertainty is where this thesis starts.⁹ Risk, in Knight's terms, describes situations where a probability distribution over outcomes can be constructed. It might be difficult. Manufacturer pricing a new product line, or a bank structuring a mortgage portfolio, but the math is doable, and the uncertainty is, in principle, insurable. True uncertainty is categorically different. No reliable probability distribution exists. The outcomes themselves may be partially unknown; the mechanisms connecting decisions to results are opaque; history offers no template close enough to be useful.¹⁰ Knight argued this kind of uncertainty characterizes most of what actually matters in business. And he was right. The line isn't just about having imprecise data or unreliable forecasts. Risk has tools for that: confidence intervals, sensitivity analysis, scenario planning. What makes something genuinely uncertain in Knight's sense is that the situation is too novel, or too singular, for any probability distribution to be meaningfully constructed in the first place. You're not estimating badly. You're trying to estimate something that doesn't have a stable base rate or a known causal structure to draw from.

People tend to cite Knight's distinction as if it were a neat taxonomy, two boxes for sorting decisions into. The deeper implication usually gets skipped over. For Knight, true uncertainty isn't just a harder version of risk. It's the economic source of profit. In a world of pure calculable risk, markets price everything correctly and competition erodes returns to zero: if everyone can compute the same expected values, no one earns above what the risk warrants. Profit, in Knight's account, is the reward for bearing uncertainty that others find too opaque to

⁹Knight, F. H. (1921). Risk, Uncertainty, and Profit. University of Chicago Press, 19–20.

¹⁰Knight, F. H. (1921). Risk, Uncertainty, and Profit. University of Chicago Press, 46–48.

act on.¹¹ The firm that moves decisively when everyone else is waiting for clarity doesn't just survive. It earns returns no actuarial table could have predicted. The distinction isn't philosophical. It's operational, and it's where competitive advantage actually comes from.

What follows from this is consequential. Classical decision theory, including the entire toolset of expected utility theory, works only when you have probabilities to plug in. When you don't, the optimization logic doesn't just become imprecise. It largely breaks down. There's nothing reliable to maximize over: no expected values to compare, no probability weights to attach. That's not a flaw in the model. It's a category error.

Cyert and March (1963) came at this from a different direction, arguing that firms don't optimize even when the math might theoretically permit it.¹² Cognitive limits, informational noise, and organizational complexity make full computation impossible in practice. March (1994) extended this further, observing that organizations facing uncertainty don't compute; they reach for rules, precedents, and identity.¹³ The operative question shifts from 'what is the expected value of each option?' to something much simpler: 'what do organizations like ours do in situations like this?' And here's the thing: under true uncertainty, that substitution isn't a failure of rationality. It may be the only reasonable response available.

These three arguments aren't running in parallel by coincidence. They're making the same point at different levels: Knight shows classical rationality breaking at the market, Cyert and March show it breaking inside the firm, and March's later work shows it breaking at the level of how individuals actually decide under pressure. The rational actor model doesn't just underperform under genuine uncertainty. It applies to a world that mostly doesn't exist.

True uncertainty also changes what strategic planning can and can't do — a point that tends to get lost in discussions that treat planning as generically valuable. Plans are, by their nature, probability statements: we expect these conditions, so we'll take these actions, and these outcomes will follow. Under true uncertainty, the expected conditions may not materialize, the causal links may not hold, and the outcomes may look nothing like the projection. This doesn't make planning useless. But it does shift what planning is for. The value of a plan under genuine uncertainty lies less in its specific forecasts and more in what the planning process builds:

¹¹Knight, F. H. (1921). Risk, Uncertainty, and Profit. University of Chicago Press, 232–238.

¹²Cyert, R. M., & March, J. G. (1963). A Behavioral Theory of the Firm. Prentice-Hall, 114–127.

¹³March, J. G. (1994). A Primer on Decision Making: How Decisions Happen. Free Press, 1–8.

shared understanding, identified assumptions, agreed decision triggers. A plan, under these conditions, is a framework for adaptation rather than a commitment to a course.

Strip away the rational actor model and the picture that replaces it is more complicated, and more honest. Under genuine uncertainty, the challenge isn't choosing optimally among alternatives with known probability-weighted outcomes. It's acting coherently when the structure of the problem itself is unclear. That's the starting point for everything that follows in this chapter.

1.2 Bounded Rationality and Satisficing

Simon's concept of bounded rationality is sometimes read as a concession to human weakness, a polite acknowledgment that real people fall short of perfect rationality. That reading undersells the claim. Simon was making a structural argument, not a psychological one.

The 1955 paper laid out two constraints that, operating together, make optimization not merely difficult but genuinely impossible.¹⁴ Cognitive limits come first: decision-makers can't hold complex multi-variable problems in their heads, can't run real-time calculations across many interdependent factors, and tend to narrow a problem's frame rather than expand it when under pressure. Informational limits follow: the full option set usually isn't known at the start, consequences are partially opaque, and the criteria for evaluation often shift during the decision process itself. These aren't weaknesses that better training could eliminate; they're structural features of the situation, and they make the classical model not harder but wrong.

Simon's alternative was satisficing.¹⁵ Rather than searching for the best option, decision-makers search for one that clears an acceptable threshold, what Simon called the aspiration level. Search stops when something qualifies; if nothing does, the threshold adjusts downward, and if options come easily it moves up. (There's something almost quietly radical about this framing: it says the criterion itself is a variable, not a fixed target, which is a significant departure from how classical theory treats preferences.) That adaptability is a significant part of what makes the concept robust across environments where rigid optimization simply can't function.

¹⁴Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–100.

¹⁵Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 101–104.

Satisficing is sometimes misread as choosing mediocrity by design. That misses what Simon described.¹⁶ The search process is sequential and structured: options are evaluated one by one against the aspiration level; search continues until something qualifies, and the threshold drifts if the search runs long. The result is neither random nor optimal in any classical sense. It's rational given the actual limits in play, and it stays responsive when conditions change. Optimization, by contrast, requires conditions that are difficult to satisfy outside of carefully bounded problems.

Cyert and March (1963) took satisficing and applied it at the level of the firm.¹⁷ Their behavioral theory of the firm changed how scholars thought about organizations, and the core insight still holds: firms aren't unitary rational actors optimizing a well-defined objective. They're coalitions of subgroups with partially conflicting interests, making decisions through bargaining and sequential attention rather than coordinated calculation. One useful concept they developed is organizational slack — the excess resources that accumulate when firm performance exceeds aspiration levels.¹⁸ Slack absorbs shocks, gives the organization room to adapt, and it tends to emerge from satisficing behavior as much as from any deliberate policy to build buffers — often it is both at once. In other words, the buffer can be a byproduct of the decision process as much as a conscious design choice.

Four mechanisms, in Cyert and March's account, allow organizations to manage uncertainty without resolving it.¹⁹ Quasi-resolution of conflict means subgroup goals are never truly reconciled; decisions are sequenced so different goals receive attention at different times, producing the appearance of coherence without genuine coordination. Uncertainty avoidance means firms orient toward short-run feedback rather than long-run forecasting. Problemistic search means the firm looks locally for solutions, in the vicinity of whatever problem just surfaced, stopping when something workable appears. Organizational learning means decision rules change through experience, but slowly and conservatively, never far from precedent.²⁰ What these mechanisms produce, in combination, are organizations that muddle through rather than plan through. Under genuine uncertainty, there may not be a better option.

¹⁶Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 104–110.

¹⁷Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, 36–38.

¹⁸Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, 42.

¹⁹Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, 117–120.

²⁰Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, 101–108.

One point that tends to get lost in summaries of this tradition comes from March (1994): decision-makers don't use a single mode of reasoning consistently.²¹ At times they reason consequentially, comparing expected outcomes. Under uncertainty, when outcomes are genuinely unclear, they tend to shift into rule-based reasoning instead — asking what someone in their role, in their organization, is supposed to do in a situation like this. That's less irrationality than a pragmatic adaptation to the limits of what calculation can actually deliver, one that draws on accumulated institutional knowledge that no freshly built analysis could fully replicate.

When a crisis arrives, the first organizational impulse is usually to demand more analysis: gather data, model scenarios, run the numbers.²² But when the crisis involves true uncertainty, that analysis can produce the feeling of rigor without its substance — a distinction that matters quite a lot in practice, even if it's hard to see from inside the process. Rules and routines embedded in organizational practice often encode better judgment than any freshly constructed model, precisely because they've been refined through repeated exposure to roughly similar situations. The hard part, knowing when to trust the embedded rule and when to override it, is a human call that no framework fully captures.

All of this connects, eventually, to the executive level.²³ Managers don't perceive their environments neutrally. Professional histories, educational backgrounds, and personal experiences function as cognitive filters: what gets noticed, what gets interpreted as a threat, what looks like an opportunity. Two executives with genuinely different career histories can look at the same competitive situation and come away with different pictures of it. Gavetti and Levinthal (2000) sharpened this by distinguishing forward-looking cognition, reasoning from mental models about genuinely novel configurations, from backward-looking experiential search that relies on trial-and-error learning.²⁴ When true uncertainty is the norm, novelty is frequent and historical analogies can mislead as often as they guide, particularly when the current situation departs significantly from anything the manager has encountered before — which makes the quality of a manager's forward-looking mental models matter more than is usually acknowledged.

²¹March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press, 57–68.

²²March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press, 8–17.

²³Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 193–194.

²⁴Gavetti, G., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113–115.

1.3 Heuristics and Cognitive Biases

If bounded rationality showed that optimal decision-making is structurally impossible under realistic conditions, Kahneman and Tversky's research program added something more troubling: the ways in which decision-makers fall short are consistent and predictable. The errors aren't random. They follow recognizable patterns, repeat across different types of problems, and can be documented with enough regularity to be studied systematically.

Their 1974 paper identified three heuristics that appear widely across judgment tasks.²⁵ Availability: people estimate probabilities based on how easily examples come to mind, which means recent, vivid, or emotionally salient events get overweighted relative to their actual frequency. Representativeness: people judge likelihood by resemblance to a prototype, routinely ignoring base rates and sample size. Anchoring: people start from a reference point and adjust toward a final estimate, but the adjustment is almost always insufficient; the starting number pulls the answer toward it even when that number was arbitrary or clearly uninformative. (One classic experimental demonstration had subjects spin a wheel of fortune before estimating the percentage of African countries in the UN; the random number from the wheel still influenced their answers, which is striking enough that it bears repeating: a number the subjects knew was random still pulled their estimates.)²⁶

Each of these shows up in recognizable ways in organizational contexts. Availability bias means firms often overreact to competitive threats that have recently materialized in dramatic form, and underreact to threats that have been building quietly. A rival who's been inactive for two years feels less dangerous, even if their underlying position hasn't weakened at all. Representativeness leads managers to categorize new situations by surface resemblance to past ones, sometimes missing the feature that makes the new situation genuinely different. A new entrant who looks like previous also-rans gets dismissed on those grounds, right up until it doesn't look so similar anymore. Anchoring shows up in budgeting, in acquisition pricing, in

²⁵Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1125.

²⁶Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1127–1128.

market size estimates: the first number to appear on the table pulls subsequent deliberation in its direction, regardless of how well-founded it was.²⁷

Tversky and Kahneman (1981) extended this line with their work on framing effects.²⁸ The same outcome described as a gain and as a loss produces systematically different choices, even when the options are objectively identical. This finding has considerable organizational relevance. The way a strategic choice is presented, not only its content, can determine how it's received and what gets approved. A workforce reduction pitched as an efficiency initiative gets a different response than the same proposal framed as a crisis response — the numbers are identical, but the framing is not, and the difference in framing produces real differences in outcomes that aren't fully correctable just by pointing out the equivalence.

The organizational implication worth making explicit is that whoever controls how a decision problem is defined has disproportionate influence on how it gets decided. The executive who establishes the reference point, who determines what the choice is being compared against, exercises a form of power that isn't captured in formal authority. A CFO who frames an acquisition price against the firm's cost of capital produces different deliberation than one who frames it against a competitor's recent deal. Neither frame is factually wrong. But they're not equivalent, and the choice between them is rarely examined as carefully as the choice between the options themselves.

The picture gets complicated at this point, and the complication matters. Gigerenzer and Gaissmaier (2011) pushed back against the conclusion that heuristics are simply errors waiting to be corrected by more rigorous analysis.²⁹ Their argument is that many heuristics are ecologically rational: well-suited to the informational environments in which they typically operate. In high-noise, data-poor conditions, a simple decision rule can outperform a sophisticated model, because the model's apparent precision is illusory when its inputs are unreliable. More data doesn't always improve decisions, and more analysis sometimes makes them worse.

²⁷Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1129–1131.

²⁸Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–455.

²⁹Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 454–457.

This complicates a tempting but too-easy conclusion: that firms should systematically replace heuristic judgment with formal analysis wherever feasible. Gigerenzer's evidence says that isn't right as a general principle.³⁰ A senior executive with thirty years in a sector deciding whether to enter a new market has something an analytical model doesn't have: accumulated experience with how projections in that sector tend to go wrong, and what the warning signs usually look like. The judgment call, 'this market looks too crowded and the regulatory trajectory is wrong,' may draw on genuinely superior information even if it can't be stated as a formal argument. The behavioral literature maps when heuristics produce errors. The ecological rationality literature maps when they don't. Both bodies of evidence are real. The reconciliation between them is still incomplete.

This isn't something that resolves neatly.³¹ Under high stakes and genuine uncertainty, figuring out in advance whether a given situation will benefit from structured analysis or from experienced judgment is itself a difficult judgment. That question will recur in Chapter 2, in a more applied form.

1.4 Behavioral Distortions in Strategic Contexts

Individual cognitive limits don't disappear when people join organizations. They tend to compound. The move from individual to organizational decision-making introduces additional layers of distortion: multiple decision-makers whose biases can reinforce each other, formal structures and incentive systems that introduce their own pressures, and cultures that can make certain errors functionally expected.

Work by Schwenk (1984) was among the first to map how cognitive simplification shows up specifically in the context of strategic decisions.³² Prior hypotheses persist against contradictory evidence — confirmation bias, in the language of psychology. Analogical reasoning is common and often genuinely useful, but it can badly mislead when the current situation is genuinely novel, which is, somewhat inconveniently, precisely the condition this thesis focuses on. And then there's the illusion of control: managers routinely overestimate how

³⁰Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 467–472.

³¹Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 462–465.

³²Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 111–113.

much influence they have over outcomes, leading to overconfidence in strategic plans and a systematic failure to appreciate how much the environment can surprise even experienced decision-makers.³³ These aren't occasional lapses; they're stable features of how strategic cognition tends to operate when the stakes are high and the time is short.

A second pattern Schwenk traced is harder to see from inside the organization: escalating commitment to a strategic direction as contrary evidence accumulates.³⁴ Once a direction has been publicly endorsed, particularly by senior executives whose reputations are attached to it, the costs of changing course grow steadily. Resources keep flowing toward the original choice, data that cuts against it gets reinterpreted, and dissenting views get pushed to the margins. (Kodak's reluctance to abandon film-based photography through the 1990s, despite having invented the digital camera in 1975, is one well-documented case; the pattern appears in sector after sector.) The sunk cost fallacy is operating here, but it's compounded by social dynamics and incentive structures that make reversal genuinely costly, not just psychologically uncomfortable. Under true uncertainty, where feedback is slow and ambiguous, this can run for a very long time before anyone outside the organization would recognize it as a problem.

Upper echelons theory, as Hambrick and Mason (1984) developed it, showed that these individual-level distortions aggregate upward into observable organizational patterns.³⁵ A CEO whose mental model of competitive dynamics was formed over three decades in one industry will amplify certain signals, notice certain threats, and pursue certain opportunities regardless of whether that model actually fits the current context. A top management team drawn from similar backgrounds will tend to converge on similar interpretations of ambiguous data. The implication, and it's a significant one, is that strategic blind spots tend to be stable over time — not because individual managers fail to learn, but because the people generating those blind spots share the same cognitive history and keep arriving at the same conclusions.

What this makes explicit, something process-focused frameworks sometimes understate, is that decision-makers aren't interchangeable.³⁶ Swapping one executive for

³³Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 116–119.

³⁴Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 113–116.

³⁵Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 196–198.

³⁶Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 198–201.

another changes the decision-making system in ways that go beyond differences in technical inputs or analytical skill. Different cognitive lenses, different professional histories, different thresholds for risk: these aren't just noise around some common rational process. They're causes of different strategic outcomes, and Hambrick and Mason's empirical predictions follow directly from this: organizations led by executives with particular background profiles should produce systematically different strategies, visible in the firm's actual choices. Those predictions have found broad empirical support across a range of organizational and industry contexts.

There's a related issue that surfaces when individual cognition is connected to how organizations learn over time. Gavetti and Levinthal (2000) traced this through a distinction between forward-looking cognition and backward-looking experiential search that is, in retrospect, fairly obvious, though it took some time for the field to formalize it.³⁷ Learning from experience works well when past situations provide reliable guidance for present ones. But under true uncertainty, that same experience can quietly become a trap: the organization gets very good at handling situations it has already encountered, while the situations that will actually shape its future look quite different. Forward-looking cognition, meaning the capacity to reason from mental models about genuinely new configurations, ones that haven't been rehearsed, becomes the scarce resource, and it tends to be considerably scarcer than most organizations appreciate until it's tested.

Powell, Lovallo, and Fox (2011) made the case for converting this body of work into a systematic research program.³⁸ Their argument was direct: strategic management has relied too long on an overly rationalistic conception of the strategic actor, one that produces tractable models but poor descriptions of how actual decisions get made. Real decisions are made by cognitively limited people, under genuine uncertainty, inside organizations that have their own structural tendencies to distort information and suppress dissent. Building behavioral realism into strategic theory isn't a concession to messiness — it's applying the same standard of rigor to the actual problem, rather than to a cleaner proxy for it.³⁹

³⁷Gavetti, G., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 119–124.

³⁸Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1370–1372.

³⁹Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1375–1378.

What behavioral strategy ultimately calls for is a shift in where the theoretical action is located.⁴⁰ Traditional strategic management placed it in competitive position, resource profiles, and industry structure — the external environment and the firm's position within it. Behavioral strategy places it, at least in part, inside the cognition of decision-makers: the mental models they carry, the biases they exhibit, the attention they allocate, and the processes through which individual judgments get aggregated into organizational choices. Calling this a rejection of earlier frameworks would be too strong. It's more accurate to say it's a reorientation, one that stops treating the human element as a residual source of noise to be modeled away and starts treating it as a primary determinant of strategic outcomes.

None of this is peripheral to a theory of strategic decision-making — it's close to the center of it, and arguably has been since at least the 1960s.⁴¹ Under true uncertainty, decision-makers don't simply fall short of some optimal standard and then muddle through as best they can. They reason through a different kind of logic in many respects, one that's heuristic, rule-based, shaped by organizational context and the personal histories of the people involved. Chapter 2 takes up the practical question that follows: not how to design away these limits, but what it looks like to build decision processes that actually account for them.

⁴⁰Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1372–1374.

⁴¹March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press, 175–180.

Chapter 2: Strategic Decision Processes in Uncertain Environments

2.1 Process Rationality and Decision Effectiveness

Chapter 1 ended on a genuinely uncomfortable note. Not conceptually uncomfortable, though there was plenty of that, but practically so. Managers can't see most of the biases that shape their thinking. Organizations don't cancel those biases out; they compound them. And awareness, by itself, doesn't help much. You can understand the planning fallacy perfectly and still produce wildly optimistic project timelines. So if better individual cognition isn't the answer, what is? That's what this chapter is about: not the psychology of how decisions go wrong, but whether the structure of the decision process itself can compensate for what individual judgment can't fix.

Dean and Sharfman (1996) produced one of the few empirical studies that actually tests whether the decision process, as opposed to decision content, affects outcomes.⁴² Their answer is both encouraging and sobering. When managers gathered relevant information and used it analytically, decision quality improved. But the effect was substantially weakened in organizations where self-interest and coalition dynamics were dominant.⁴³ Analytical rigor still happened. The information just got distorted before it could influence anything. The process looked rational from the outside. Inside, it wasn't.

Process rationality doesn't mean a formal sequence of steps. Dean and Sharfman's definition is functional rather than structural: it's about whether relevant information is actually gathered and genuinely applied to the choice at hand. An informal discussion where people honestly surface disagreements and work through them can be more rational in this sense than an elaborate review cycle that produces polished documentation while the real decisions happen elsewhere. The formal framework and the functional reality diverge constantly in most organizations, and it's the functional reality that determines what happens to decision quality.

Nutt (1999) spent years tracking organizational decisions and found something striking: roughly half failed outright, abandoned or delivering far less than expected.⁴⁴ The main culprits weren't bad information or genuinely hard problems. They were process problems. Managers typically considered only one or two options before committing, which means the analysis

⁴² Dean, J. W., & Sharfman, M. P. (1996). Does decision process matter? A study of strategic decision-making effectiveness. *Academy of Management Journal*, 39(2), 368–396.

⁴³ Dean & Sharfman, Does decision process matter?, 1996.

⁴⁴ Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13(4), 75–90.

came after the decision, not before it. Stakeholders got handled rather than consulted. Political opposition derailed decisions that were otherwise sound. And the biggest issue was premature commitment: choosing a direction early and treating the subsequent evaluation as a formality.⁴⁵

What makes these failures so persistent is that they're invisible from inside. A single recommended option isn't called premature commitment. It's called having done the analysis. Handling stakeholders isn't labeled political maneuvering; it's managing relationships. The vocabulary of organizational life wraps these failure modes in language that makes them sound like good practice. You don't recognize them as failures until the outcome forces the point, and by then you're already explaining them away as implementation problems or bad luck.

Organizations are political. That's not cynical, it's structural. Different groups have different interests, coalitions form around options that serve them, and information gets disclosed selectively. None of this is inherently pathological; political processes often surface genuine value disagreements that purely analytical frameworks can't resolve. The problem is what happens when politics dominates the deliberation entirely. The analytical work becomes decoration. Arguments get made, information gets cited, and the conclusion was always going to be whatever the dominant coalition wanted. Dean and Sharfman's finding that process rationality lost its predictive power under high political behavior is almost obvious in retrospect: the mechanism only works when the information it generates actually reaches the decision.

Put all of this together and a clear picture emerges, one that's different from what most organizations actually act on. Decision quality doesn't come mainly from having smart people or good data. It comes from whether the process is structured to surface real alternatives, whether the analysis genuinely shapes the choice rather than justifying it after the fact, and whether the political environment allows honest evaluation. None of those conditions are given. They have to be deliberately built. And the gap between a process that looks rational and one that actually functions that way can be enormous.

What this means practically is that process rationality has to be protected, not just designed. An organization can build a formally sound decision process, run it for years, and gradually watch it hollow out as informal norms develop around it. The review exists on the calendar. Attendance is good. The documentation gets filed. But the people in the room have

⁴⁵ Nutt, Surprising but true, 1999.

learned which conclusions the process is supposed to produce, and the analytical work adjusts accordingly. This isn't cynicism or bad faith; it happens incrementally, and at each step it looks like reasonable accommodation to organizational reality. The end state is a process that performs legitimacy rather than producing it. Recognizing that dynamic while it's happening, rather than after it has fully set in, is one of the harder governance problems any organization faces.

2.2 Speed and Comprehensiveness in High-Velocity Contexts

The intuition that speed and decision quality trade off is nearly universal among managers. Take more time, get a better answer. Move faster, get a rougher one. It sounds right. The empirical evidence doesn't really support it, and the reasons why are worth understanding carefully.

Eisenhardt (1989) studied executive teams in the semiconductor industry, an environment where technology shifted rapidly, competitive moves came fast, and being late on a major decision was often fatal.⁴⁶ The expectation going in might have been that fast-moving firms were cutting corners and getting lucky. What she actually found was different. The fastest teams weren't skipping analysis. They were doing it differently. Real-time data instead of historical projections. Advisors consulted simultaneously rather than in sequence. And multiple options developed in parallel rather than committing to a single path and revising when it failed.

The multiple-alternatives finding is worth dwelling on.⁴⁷ Running several options simultaneously sounds like it would take longer. More work, more analysis. It's actually faster, because it avoids what sequential commitment produces: pick a path, discover the flaws under real conditions, reverse course, start over. The comparison isn't between one thorough analysis and several shallow ones. It's between one long revision cycle and several shorter ones that converge. Speed and quality weren't in tension in Eisenhardt's high-performing firms. They came from the same source: a process architecture that was genuinely different from the sequential default most organizations use.

⁴⁶ Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32(3), 543–576.

⁴⁷ Eisenhardt, Making fast strategic decisions in high-velocity environments, 1989.

There's also something important in Eisenhardt's finding about real-time versus historical data.⁴⁸ In a stable environment, last year's numbers are a reasonable guide to this year. In a fast-moving one, they're not just less accurate; they're actively misleading, because the rigor they appear to provide generates confidence that isn't warranted. The market you modeled three months ago may already be a different market. A technology platform that looked dominant in your analysis may already be facing a viable substitute. Real-time metrics are messier. But at least they reflect what's actually happening.

There's a structural consequence to that shift that rarely gets discussed. Historical financial models can be built by analysts who weren't present for the events they're modeling. Real-time operational information tends to live closest to the people actually running the operations, who are often junior to the executives making the decisions. Getting that knowledge into a strategic deliberation in a form that is timely and usable isn't automatic. Organizations that haven't deliberately built channels for it keep reaching for the historical models, even when everyone in the room knows those models are stale.

Fredrickson and Mitchell (1984) add an important qualification.⁴⁹ They found that comprehensive decision processes, those that try to integrate all relevant information before committing, were associated with better performance in stable industries. In unstable industries, the relationship reversed. Comprehensiveness was negatively associated with performance. The logic is fairly straightforward: in a stable environment, the information you gathered at the start of the process is still valid when the decision gets executed. In an unstable one, it may not be. Thoroughness stops being a reliable indicator of quality when the underlying situation is itself changing.⁵⁰

The failure mode this creates is underappreciated. It's not just that the analysis takes too long. It's that a comprehensive analysis of a changing situation generates false confidence: a highly elaborated picture of what was true some weeks ago, presented with the authority of rigor. The organizations most at risk from this are the ones with the strongest analytical cultures, because those cultures make it hardest to recognize when thoroughness has stopped being a virtue. The analysis looks exactly like it always has. The problem is that the world it's describing has moved on.

⁴⁸ Eisenhardt, Making fast strategic decisions in high-velocity environments, 1989.

⁴⁹ Fredrickson, J. W., & Mitchell, T. R. (1984). Strategic decision processes: Comprehensiveness and performance in an industry with an unstable environment. *Academy of Management Journal*, 27(2), 399–423.

⁵⁰ Fredrickson & Mitchell, Strategic decision processes, 1984.

Eisenhardt and Fredrickson and Mitchell are making the same point from different angles. The process needs to match the environment it's operating in. In a stable setting, more comprehensive analysis is better. In a genuinely uncertain one, it can make things worse, not because rigor is bad, but because rigor applied to an outdated picture produces well-supported wrong answers. The key question isn't how much information is available. It's whether what has been gathered will still be valid when the decision gets executed.

The trickier problem is knowing which environment you're actually in. Stable and unstable aren't fixed categories. They describe conditions that change, sometimes slowly and sometimes suddenly. An organization that spent a decade building sophisticated analytical capabilities in a relatively stable market can be genuinely slow to recognize that those capabilities have become a liability. The evidence arrives as outcomes: decisions that used to work are now failing. By the time enough failures have accumulated to shift the culture, the cost has been paid many times over. Naming that dynamic at least creates the possibility of monitoring for it.

2.3 Structured Mechanisms for Reducing Bias

Knowing about cognitive bias doesn't help you fix it. That's the frustrating finding running through the debiasing literature. Awareness rarely eliminates a distortion. Study the anchoring effect in detail and you'll still anchor on the first number you see. The biases that distort strategic decisions operate largely below the level of conscious deliberation, which means the most intuitive response, telling people to be more careful, is largely ineffective. The question that actually matters is whether process-level interventions can accomplish what individual-level ones can't.

Lovallo and Sibony (2010) made the most sustained attempt to answer that question practically.⁵¹ Their starting premise is that debiasing has to happen at the process level, because individual debiasing is both difficult to produce and unreliable when produced. They identify three mechanisms worth examining. Devil's advocacy gives someone the structural role of making the strongest possible case against the preferred option, not as a courtesy but as a genuine feature of how the deliberation runs. Pre-mortem analysis inverts the framing: instead of evaluating whether a decision is sound, participants assume it has already been made and

⁵¹ Lovallo, D., & Sibony, O. (2010). The case for behavioral strategy. *McKinsey Quarterly*, March 2010, 30–43.

failed, then work backward to identify the most plausible causes.⁵² This matters because commitment effects normally make honest critique difficult after a direction is chosen. Pre-mortem sidesteps that by treating failure as the premise, not something to be argued for.⁵³

Reference class forecasting works through a different mechanism. Rather than trying to reason more carefully about the specific decision, it forces attention to base rates: how have comparable decisions, made by comparable organizations, actually turned out?⁵⁴ The purpose is to counter inside-view thinking, which focuses on the particulars of the current case and tends to be optimistic in ways that aggregate data does not. It's hard to sustain overconfidence about a new product launch when you can see that most comparable launches underperformed their projections. The mechanism doesn't ask people to reason differently. It changes what they're required to look at before committing.

Getting the reference class right, though, is harder than it sounds. Define it too broadly and the comparison loses meaning: technology launches don't all fail for the same reasons, and grouping them indiscriminately washes out exactly the variation that matters. Define it too narrowly and you don't have enough cases to generate meaningful base rates. There's also a discipline problem. The class needs to be defined before examining the specifics of the current case, because once you've seen the details, motivated reasoning will push toward whichever reference class produces the most comfortable prior. Easy to specify as a rule, very hard to sustain under decision pressure.

These mechanisms are well established. They also fail a lot, and the reasons are instructive. Devil's advocacy ritualizes easily. Someone is assigned the critic's role, delivers it, and the deliberation continues more or less as it would have without them. The mechanism works when the designated advocate is genuinely empowered to press the case and when senior leaders signal through their conduct, not just their words, that the criticism is being taken seriously. Pre-mortem has a similar failure mode: a list of risks gets generated, filed alongside the project plan, and rarely influences the actual decision. Whether these interventions do what they're supposed to depends almost entirely on the organizational environment they're embedded in.

⁵² Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.

⁵³ Lovallo & Sibony, The case for behavioral strategy, 2010.

⁵⁴ Lovallo, D., & Kahneman, D. (2003). Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81(7), 56–63.

Dane and Pratt (2007) introduce a complication that deserves more than a footnote.⁵⁵ Their argument is that framing intuition as simple irrationality misses something important. Intuition is pattern recognition, accumulated experience stored in non-declarative memory, operating quickly and largely below conscious access. In domains where the expert has genuine experience with real feedback, where patterns recur and the current situation resembles the environment where the intuition formed, it can be highly accurate. Sometimes more accurate than a formal model built from the same information by someone without the relevant domain background.

The question, then, isn't whether to trust intuition in general. That's the wrong frame.⁵⁶ The right question is domain-specific: where did this intuition form, and how similar is the situation now to the environment that produced it? A judgment developed over decades in one market context may be highly reliable within that setting and nearly useless outside it, even when the same executive makes both calls with the same apparent confidence. The gap between where an intuition formed and where it's being applied is exactly where structured deliberation adds the most value.

Powell, Lovallo, and Fox (2011)⁵⁷ push this further. Their argument is that the most consequential strategic errors aren't analytical failures: they're failures of the questions asked before analysis begins. Which options are considered? Which comparisons are made? What counts as relevant evidence? These framing choices determine the decision space within which subsequent reasoning operates, and they're almost entirely invisible once the analysis begins. Challenging the frame is harder than challenging the conclusion, because frames feel self-evident in a way that conclusions don't. Structured deliberation is most valuable at the framing stage: surfacing the assumptions that define what the decision is actually about, before those assumptions harden into the background conditions that analysis takes for granted.

2.4 The Structure–Flexibility Trade-off

Everything in the preceding sections rests on a tension that hasn't been named directly. Structure helps. The evidence on that is solid: systematic information-gathering, developing

⁵⁵ Dane, E., & Pratt, M. G. (2007). Exploring intuition and its role in managerial decision making. *Academy of Management Review*, 32(1), 33–54.

⁵⁶ Dane & Pratt, Exploring intuition and its role in managerial decision making, 2007.

⁵⁷ Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1369–1386.

multiple alternatives, explicit bias-checking, deliberate documentation, all of these improve decision quality in measurable ways. But structure costs things too. It takes time. It creates rigidity. It can produce decisions that are procedurally correct and substantively mediocre, optimized for compliance rather than genuine problem-solving. In situations requiring fast improvisation, or that fall outside the parameters the structure was designed for, it can fail completely. This isn't a peripheral qualification. It's a central feature of the problem.

Weick (1995) comes at this from a direction most decision process literature ignores.⁵⁸ His concern isn't how to structure deliberation. It's how people construct meaning when a situation is genuinely ambiguous and existing frameworks don't cleanly apply. His central finding is counterintuitive: sensemaking is retrospective. People don't understand a situation and then act on it. They act, and then construct an account of what the situation must have been given what they did. Under genuine ambiguity, understanding often can't precede action. You have to do something first to get a foothold on what's actually happening.

His illustration of this is the Mann Gulch disaster, where a crew of experienced smokejumpers died after encountering fire behavior they had never seen and couldn't interpret quickly enough.⁵⁹ The coordination structures that normally kept them functioning as a unit collapsed under conditions their training hadn't anticipated. When the foreman tried to communicate an escape route that violated everything they knew, the structure that had served them became an obstacle rather than a resource. The lesson isn't that training is counterproductive. It's that structures built for anticipated conditions fail when conditions exceed them, and that exceeding the anticipated conditions is, under genuine uncertainty, normal, not exceptional.

What Weick's analysis contributes to the decision process literature is a different account of what makes a process good. Most of that literature defines quality analytically: did the process surface the right information, consider the right alternatives, apply the right checks? Weick shifts the question toward adaptive capacity: can the structure flex when conditions exceed its design parameters, and can people within it maintain enough situational awareness to recognize when that's happening? These are genuinely different questions. A process optimized for analytical quality in stable conditions can score well on the first and catastrophically on the second. An organization that treats good process design as a one-time

⁵⁸ Weick, K. E. (1995). *Sensemaking in Organizations*. Sage Publications.

⁵⁹ Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.

engineering problem, rather than an ongoing management challenge, is essentially building for the conditions it already understands while leaving itself exposed to the ones it doesn't.

The same pattern appears in organizational decision-making, usually in less catastrophic form. A planning process designed for three-year horizons becomes a liability when the industry is moving quarterly. A resource allocation framework built around stable business unit boundaries blocks the most significant opportunities, which keep cutting across those boundaries. A risk assessment calibrated for supply chain disruptions provides no useful guidance when the discontinuity is regulatory. None of these structures are wrong in themselves. They're just optimized for conditions that may no longer exist. And the organizations most vulnerable to this failure mode are often the most successful ones, because success deepens the investment in whatever produced it.

Teece, Pisano, and Shuen (1997) address this at the organizational level.⁶⁰ Competitive advantage under rapidly changing conditions, they argue, doesn't come from holding particular resources or executing particular strategies. It comes from the firm's ability to integrate, build, and reconfigure internal and external competences to address rapidly changing environments. Dynamic capabilities are, in effect, meta-processes: not routines for making specific decisions, but organizational capacities for revising and replacing those routines when they stop fitting the environment. The practical challenge this creates is considerable: revising existing decision routines requires the people who run them to evaluate them objectively, setting aside cognitive attachments and institutional investments accumulated over years. The people most capable of executing a process are often the least positioned to recognize when it has become a constraint.

The resolution to the structure-flexibility tension isn't to choose between them.⁶¹ Effective organizations need both, operating at different levels. Structure at the level of individual decisions: surface alternatives, check for obvious biases, ensure the relevant knowledge is in the room, document the reasoning in ways that support honest review later. Flexibility at the meta-level: the organizational capacity to recognize when existing decision structures have stopped fitting the environment, and to revise them before the cost becomes undeniable. That second requirement is considerably harder than the first.

⁶⁰ Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.

⁶¹ Teece, D. J. (2016). Dynamic capabilities and entrepreneurial management in large organizations: Toward a theory of the entrepreneurial firm. *European Economic Review*, 86, 202–216.

Part of what makes it hard is that the signals are usually ambiguous. A decision process isn't obviously broken when it starts to fail. The outputs look similar. The deliberations take roughly the same amount of time. The documentation still gets produced. What changes is what the process is actually doing relative to what it appears to be doing, and that gap is invisible to anyone inside the organization who has an interest in not seeing it. The organizations that manage this well tend to share one feature: they've invested in deliberate, arm's-length review of their decision processes, not just their decision outcomes. They have people with standing to ask whether the process produced the decision or merely documented it, and those people are protected from the organizational pressure that would normally silence that question.

Building that meta-level flexibility deliberately is genuinely difficult. The people best positioned to evaluate a decision process are the ones who run it, and familiarity is exactly what makes objective assessment hard. Post-mortem reviews help when they're structured to separate process quality from outcome quality. A decision can be well-made and still produce a bad result; it can be poorly-made and still get lucky. Organizations that sustain genuine process learning tend to be ones where senior leadership actively maintains that distinction, where a fortunate outcome gets examined as carefully as a costly failure, and where the question isn't just what went wrong but how the decision was actually made.

After working through four sections of process findings, the natural instinct is to synthesize them into a general framework, a set of recommendations that work in most situations with adjustments at the margins. The literature doesn't support that move. The most consistent finding across the studies surveyed in this chapter is that the right process depends heavily on context. Not as a qualification appended to the main argument. As the main argument.

Ocasio (1997) provides the clearest theoretical account of why this is so.⁶² His argument is that what decision-makers attend to shapes what they decide, and that what they attend to is structured by the organizational rules, resources, and communication channels their context provides. Two firms facing identical environmental changes may respond completely differently, not because their managers differ in capability, but because their organizations route attention to different signals, frame problems through different lenses, and bring different interpretive frameworks to bear. The decision process is embedded in an attention structure,

⁶² Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1), 187–206.

and that structure is a property of the organizational context, not of the formal process design sitting on top of it.

For process design, this creates an uncomfortable implication. The same mechanism can work very differently in different organizational contexts. A formal strategic review that functions as a genuine deliberative forum in one organization becomes a documentation ritual in another. The structure is identical, but one produces information that actually influences choices while the other produces records of choices already made through other channels. Changing the structure without changing the context that determines what the structure does rarely produces the intended results. This is one reason process improvement programs disappoint so often: they treat the formal framework as the operative variable when the operative variable is what that framework actually does in a specific organizational setting.

These contextual factors don't operate in isolation, and the combinations matter. High velocity combined with poor knowledge distribution is a particularly difficult configuration: the situation is changing fast enough that you can't slow down to gather more information, and the information that exists is held by people who aren't in the room when the decision gets made. Strong hierarchy under high political pressure creates a different problem: the formal mechanisms get deployed, but the organizational environment makes honest evaluation impossible. Diagnosing which configuration you're operating in is a necessary prior step to thinking about process design, and most process frameworks skip it entirely.

Velocity matters, as Eisenhardt, and Fredrickson and Mitchell, establish clearly, but it's one dimension among several.⁶³ How reversible is the commitment? A decision revisable within three months warrants a much lighter process than one locking in a five-year trajectory. How asymmetric are the costs of error versus delay? Most processes don't calibrate to this. And where is the relevant knowledge? Some decisions concentrate expertise at the top of the hierarchy; others distribute it across levels and functions in ways the formal decision structure doesn't reflect. When the people with the most relevant knowledge don't have decision authority, the design challenge isn't optimizing the deliberation among those who do. It's figuring out how to move that knowledge upward in usable form.

⁶³ Eisenhardt, Making fast strategic decisions in high-velocity environments, 1989; Fredrickson & Mitchell, Strategic decision processes, 1984.

Culture may be the most consequential contextual variable and the hardest to design around.⁶⁴ It determines what can actually be said in a room without professional cost: whether a challenge lands as valuable critique or as unwelcome friction, whether a scenario showing project failure gets genuine consideration or polite acknowledgment followed by dismissal. Devil's advocacy operates completely differently in a culture that honestly values critical input versus one where hierarchy is strong and bad news travels slowly. The same goes for pre-mortem analysis and reference class forecasting. None of these mechanisms are culture-independent. They all require a social environment where honest evaluation is possible and where surfacing uncomfortable conclusions doesn't carry a penalty. Process designs that ignore culture tend to produce outputs that look correct on specification and fail in the room.⁶⁵

Chapter 3 examines these contextual factors from a more integrative angle, drawing on the behavioral and process material developed across Chapters 1 and 2 to ask what conditions determine whether any decision process actually works as intended. The research surveyed in this chapter establishes what good process design looks like in principle. What it establishes less clearly, and what subsequent analysis needs to address, is why organizations that understand these findings so often fail to act on them. Understanding that gap requires moving from process design as a technical problem to process design as an organizational and political one. The mechanisms for reducing bias, the tools for managing velocity, the structures for capturing distributed knowledge, all of them involve doing something that someone in the organization has an interest in not doing. That is the territory Chapter 3 enters.

Chapter 2 has traced the conditions under which structured decision processes improve outcomes, the specific mechanisms through which they do so, and the contextual factors that determine whether those mechanisms function as intended. What hasn't been attempted yet is an integration: a framework that brings together the behavioral foundations of Chapter 1 with the process analysis developed here into a coherent account of what decision quality under genuine uncertainty actually requires. That synthesis is the task of Chapter 3.

⁶⁴ Lovallo & Sibony, *The case for behavioral strategy*, 2010.

⁶⁵ Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Perspectives*, 13(4), 75–90.

Chapter 3: Toward a Framework for Decision Quality Under Uncertainty

3.1 Integrating Behavioral and Process Perspectives

Chapter 1 and Chapter 2 were running parallel tracks. That was deliberate — the behavioral material needed its own space before the process design question made much sense. But they're not separate topics. They bear on the same underlying problem, and this chapter tries to bring them together.

Why they've stayed separate is worth noting. The behavioral literature developed mainly within psychology and organizational behavior — its own journals, its own methods, its own ideas about what counts as a good finding. The process literature developed within strategic management, focused on executive teams and competitive outcomes. When behavioral work did reach strategic management, mainly through Kahneman and Tversky, it arrived as a list of cognitive problems to be aware of rather than as material to integrate with process design. The two bodies of work kept running in parallel. What gets lost when they stay separate is the interaction — the fact that which biases matter most depends heavily on the process context, and which process designs work depends heavily on the cognitive features of the people inside them.

The problem is: what makes a strategic decision good when the decision-maker is cognitively limited, the environment is genuinely uncertain, and the organization that produces the decision introduces its own distortions? The answer can't be 'getting to the correct choice' — under genuine Knightian uncertainty, there's no correct choice available in advance to benchmark against. The future can't be calculated. The probability distribution over outcomes can't be reliably specified. What a framework for decision quality can do is identify the conditions under which processes are more likely to produce choices that are well-adapted to the environment, internally coherent, and actually informed by available evidence rather than by cognitive distortion or by whoever happened to dominate the room.

The asymmetry is worth sitting with. When a decision goes well, organizations tend to attribute it to the quality of their analysis, the experience of the decision-makers, the soundness of the process. When it goes badly, the same factors get re-examined — and then adjusted after the fact. This makes learning about decision quality genuinely hard. You can't tell from a single outcome whether the process that produced it was good or lucky. You can't tell from a pattern of outcomes either, unless you're tracking the process in enough detail to separate its effects

from the environment's. Most organizations don't do that. The gap between what decision quality actually is and what gets used as a proxy for it is where a lot of the failure accumulates.

A framework that takes this seriously has to be explicit about what it means by quality. The definition used here is procedural rather than substantive: a decision process produces quality outcomes to the extent that it surfaces real alternatives, ensures that relevant information actually reaches and shapes the choice, and builds in mechanisms that reduce the distorting effects of cognitive bias and political dynamics. This doesn't guarantee a good outcome — under genuine uncertainty, nothing does. But it specifies conditions that make good outcomes more likely and bad ones more recoverable. That's the most useful thing a framework can offer when the future can't be known in advance.

This also means the framework has to be honest about what it can't do. It can't tell you whether a specific decision was good. By the time you know the outcome, you're looking backward, and the temptation to read quality off the result is almost irresistible. What the framework can do is describe what a process looked like while it was happening — whether real alternatives were on the table, whether the information that shaped the choice was actually the most relevant information available, whether the people with standing to raise problems felt safe doing so. Those are the observable features. They're also the ones that can be improved.

That framing requires treating bounded rationality not as a bug to be fixed but as a permanent feature of the situation. You can't give a manager a different cognitive architecture. You can't eliminate the planning fallacy by explaining it. The question is whether you can design processes and organizational contexts that make cognitive limits less damaging — and under what conditions that design actually works.⁶⁶

Knight was clear that the practical question isn't eliminating uncertainty, which is impossible, but acting constructively within it. March pushed further — the hardest decisions involve genuine ambiguity not just about outcomes, but about what the right criteria for evaluation even are. Gigerenzer makes essentially the same point from a different angle: any decision strategy has to be judged relative to the environment it's operating in, not against some

⁶⁶Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99.

universal standard of rationality. None of these are fringe methodological positions. They're what you're left with once you take uncertainty seriously.⁶⁷⁶⁸⁶⁹

3.2 Development of the Conceptual Framework

The framework has three components. The main claim is about how they interact, not about any one of them in isolation.

The cognitive component covers the individual and organizational-level factors from Chapter 1 — bounded rationality, satisficing, heuristic reasoning, cognitive bias, and the organizational dynamics that amplify individual-level distortions into collective ones. These can't be redesigned from the outside. What you can do is build processes and contexts that create conditions in which cognitive limits are less likely to cause serious damage. That connection — from cognitive features to process and context design — is what the framework is built around.⁷⁰

The organizational amplification point deserves more than a passing reference. Individual biases are well-documented and reasonably predictable. What's less predictable is what happens to them under hierarchy, time pressure, and when careers are on the line. Groupthink is the famous version, but it's actually less common than something subtler: information gets filtered as it moves up, problems get framed in terms of how they'll land rather than what's actually true, and people advocate for positions rather than explore them. The cognitive features that produce manageable errors in individual judgment produce systematic institutional distortions in collective ones. Process design that doesn't account for this will consistently underperform.

The process component covers the design choices about how the decision-making process is organized: how information gets gathered, how alternatives get generated and evaluated, who is involved in deliberation, what mechanisms exist for checking bias, and how quickly the process moves relative to how fast the environment is shifting. The finding from Chapter 2 is that these choices have measurable effects — but the optimal configuration varies.

⁶⁷Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. University of Chicago Press, 311–313.

⁶⁸March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press, 175–191.

⁶⁹Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 453.

⁷⁰Dean, J. W., & Sharfman, M. P. (1996). Does decision process matter? A study of strategic decision-making effectiveness. *Academy of Management Journal*, 39(2), 382–385.

A high-velocity environment needs a different process architecture than a stable one. A decision that depends on tacit expert judgment needs different deliberative structures than one where explicit data analysis is feasible. These aren't exceptions to a general rule. They are the general rule.⁷¹

One implication of this that gets missed: the process needs to be calibrated not just to the decision but to where in the organization that decision is happening. A decision made at the business unit level in a stable environment by a team with deep domain experience probably needs a lighter process than most frameworks prescribe. A novel strategic commitment made at board level in a fast-changing environment where the most relevant knowledge sits two layers down the hierarchy probably needs a heavier one. The error isn't using the wrong process in general. It's applying the same process regardless of those differences.

The contextual component is where the other two either work or don't. Context here means three things: the degree of genuine environmental uncertainty; the organizational features — culture, incentives, communication channels, political dynamics — that shape what any process actually does regardless of how it's designed on paper; and the distribution of knowledge across the organization. Where does the most relevant information actually sit? Is the process designed to reach it?⁷²

The knowledge distribution point is particularly underappreciated in the process literature. Organizations routinely make decisions at levels of the hierarchy where the most relevant expertise doesn't sit. The people who know what's actually happening in the market, in operations, with customers, are often not the people who make the final call. This isn't always a design failure — there are real reasons why strategic authority sits where it does. But it means the design question isn't just 'how do we deliberate well among the people in the room?' It's 'how do we get the knowledge that isn't in the room into the conversation in a form that's timely and usable?' That's a harder problem, and most process frameworks skip it.

These three interact in ways that have consequences for how you think about decision quality. Cognitive constraints are universal, but whether they actually damage outcomes depends on the process context — the same bias that produces a catastrophic error in an unstructured deliberation may barely register in a process with appropriate checks. Process

⁷¹Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32(3), 566–571.

⁷²Weick, K. E. (1995). *Sensemaking in Organizations*. SAGE Publications, 88–100.

designs that reduce the consequences of cognitive bias in one configuration may be ineffective or actively harmful in another: comprehensive information-gathering improves outcomes when information is stable but degrades them when conditions change faster than the process can complete. And the organizational context shapes the functional reality of the process in ways that are often invisible to whoever designed it.⁷³

The appropriate unit of analysis, then, isn't the individual decision-maker or the process design — it's the configuration of all three taken together. Change any one dimension and you shift what the others imply. That's what makes this integrative rather than just a list of factors that matter.⁷⁴

One way to see this is to think about what goes wrong when only one dimension gets addressed. Organizations that focus exclusively on the cognitive dimension — running bias-awareness training, bringing in behavioral economists to sit in on decisions — often find that the interventions don't stick. The individual-level changes dissipate under organizational pressure. Organizations that focus exclusively on process — implementing formal review stages, creating devil's advocate roles, requiring pre-mortem analyses — find that the mechanisms ritualize. They go through the motions without doing what they were designed to do. The contextual dimension is almost always what's missing: the cultural and political environment that determines whether the cognitive and process interventions have any grip. Addressing all three together is harder. It's also the only version that works.

3.3 Theoretical Implications

The framework bears on several strands of theory in strategic management. Three are worth being direct about.

Powell, Lovallo, and Fox's (2011) call for behavioral realism in strategic management pointed in the right direction. But the practical implication of behavioral strategy isn't 'eliminate biases' — that's largely impossible. It's 'design processes and organizational contexts that reduce the damage biases are likely to do in the specific conditions you're operating in.' That's a narrower claim, but a more actionable one. It requires knowing which cognitive

⁷³Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 528–532.

⁷⁴Fredrickson, J. W., & Mitchell, T. R. (1984). Strategic decision processes: Comprehensiveness and performance in an unstable environment. *Academy of Management Journal*, 27(2), 417–419.

characteristics are most consequential under which contextual configurations — which is what behavioral strategy as a research program has been slower to specify.⁷⁵

There's a practical implication here for how organizations typically use behavioral strategy findings. The default application is training — executives learn what the planning fallacy is, what anchoring looks like, how overconfidence manifests. The evidence that awareness changes behavior is weak. The evidence that process-level changes can reduce the damage these biases do is considerably stronger. That's a significant shift in where to focus the intervention: not on making individuals less biased, but on making the conditions under which their judgment operates more protective of the decision quality. The two approaches aren't mutually exclusive, but organizations that stop at awareness training are getting the less effective half of the program.

Hambrick and Mason's (1984) upper echelons framework showed that the cognitive profiles of top managers shape strategic choices in systematic ways. What the current framework adds is that these effects are mediated by process and context. The same cognitive profile at the top will produce different strategic outcomes depending on what processes exist for challenging executive judgment, what channels exist for surfacing information that contradicts executive priors, and what the cultural context does to the dynamics of deliberation. Upper echelons effects are real. They're not fixed.⁷⁶

There's also a research design implication worth naming. Most upper echelons studies treat executive characteristics as predictors and strategic outcomes as outcomes, with process largely as a black box in between. If the framework here is right, that approach will systematically underestimate upper echelons effects in some contexts — where the process genuinely mediates — and overestimate them in others — where the process buffers against them. Measuring process directly and treating it as a mediator rather than a control would produce more interpretable findings. The upper echelons literature has generated a lot of empirical work. The mediating role of process design is one of the questions that work has consistently left open.

Ocasio's (1997) attention-based view established that what firms notice depends on the rules and structures that channel attention. What the framework here adds is a question that

⁷⁵Powell, T. C., Lovullo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1381–1383.

⁷⁶Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 193–206.

view leaves mostly open: whether the attention being allocated is well-calibrated to actual environmental conditions. An organization whose attention structures were built for a more predictable environment may systematically misallocate attention when conditions become genuinely uncertain — focusing on signals that used to matter and missing the ones that now do. Designing attention structures with genuine uncertainty in mind is a more demanding requirement than current models typically acknowledge.⁷⁷

This connects to something organizations deal with constantly and rarely solve: what information actually gets into a strategic deliberation. Most organizations answer this by bringing in what's easy to bring in — what finance can model, what strategy teams can report, what's already in the management information system. The harder-to-access information — what the people closest to customers are actually seeing, what the organization's own processes are producing in terms of unintended consequences — tends to get filtered out before it reaches the room. An attention structure designed with genuine uncertainty in mind would actively compensate for that. Most don't, and the consequences show up as decisions that were well-analyzed but poorly informed.

All three theoretical implications point toward a common gap in the existing literature: the conditions under which the mechanisms it describes actually work. The behavioral strategy literature knows biases exist. The upper echelons literature knows executive cognition matters. The attention-based view knows that organizational structure shapes what gets noticed. What's less developed in each is the boundary conditions — when these effects are large versus small, when they can be managed versus when they compound into something that no process design can correct for. The framework proposed here is partly an attempt to map those boundary conditions by specifying the interactions between cognitive, process, and contextual factors.

3.4 Limitations and Future Research

Any honest accounting of a conceptual framework starts with what it doesn't do. This one is built on a synthesis of existing theory, not on primary empirical research. It connects arguments already in the literature and specifies relationships between variables, but it doesn't produce new data on how decisions actually unfold in real organizations. A framework like

⁷⁷Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1), 200–202.

this can clarify the relevant questions and specify what relationships to test. It can't substitute for testing them.

The synthesis is also partial in its theoretical coverage. The framework draws on bounded rationality, behavioral strategy, dynamic capabilities, and the attention-based view. Other traditions that bear on decision quality under uncertainty aren't fully represented. Institutional theory has a great deal to say about how organizational context shapes decision behavior through mechanisms not fully captured here. Sensemaking theory, present through Weick's work, could be developed further as a foundation for understanding how cognitive and contextual factors interact in genuinely ambiguous situations. These are gaps in scope, not fatal problems, but they affect how much weight to place on any specific proposition.

The framework is also necessarily general — it identifies dimensions and proposes relationships without specifying precise boundary conditions. How uncertain does an environment have to be before comprehensive analysis starts hurting rather than helping? How politically charged does an organization have to be before process rationality stops working the way Dean and Sharfman observed? The framework raises these questions without answering them. The answers are likely to be contingent on factors it doesn't fully capture.⁷⁸

What this means practically is that the framework can't be applied like a checklist. It doesn't say 'if you have high environmental uncertainty and strong political dynamics, use process type X.' It says that under those conditions, process designs that assume clean information and honest deliberation will underperform relative to what they look like on paper. That's a diagnostic claim, not a prescriptive one. The prescription has to be worked out in the specific context by people who know it. Which is itself something the framework would predict — context-independent prescriptions are exactly what the framework argues against.

There is also a practitioner gap that the framework opens without fully addressing. Most of what the decision-making literature produces is aimed at researchers. The framework here has practical implications — it suggests where organizations should focus attention, what kinds of process changes are likely to matter, which contextual factors are most consequential — but it doesn't translate those implications into specific interventions. That translation is difficult, because interventions that work in one organizational context often don't transfer. This is part of why the consulting literature on decision quality tends to underperform: it packages general

⁷⁸Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13(4), 86–88.

principles into context-independent tools, and then the context defeats the tool. A research agenda that takes the framework seriously would include work that gets closer to the organizational level — not just measuring the variables but understanding how they interact in the specific configurations that real organizations inhabit.

Time is a third gap. The framework is essentially cross-sectional — it describes relationships between cognitive, process, and contextual factors at a given moment. But organizations operate over time, and the dynamic capabilities Teece, Pisano, and Shuen (1997) describe involve precisely the capacity to revise these dimensions in response to environmental shifts. How quickly organizations can do this, what enables and blocks it, and how leadership cognition interacts with organizational change dynamics — those are questions the framework opens without answering.⁷⁹

The empirical gap is the most obvious priority. Research that tracks the full triad of cognitive characteristics, process design, and organizational context across a range of decision episodes would be valuable — most existing work examines one or two of these while holding the others constant or treating them as controls. Case-based comparative research examining how the same process produces different outcomes in different organizational contexts could test the interaction hypotheses directly.⁸⁰

Less obvious but probably as important: the mechanisms through which organizations actually develop the capacity to revise their decision processes. Dynamic capabilities theory describes what this looks like at the organizational level. The micro-level processes through which decision routines get noticed, evaluated, and changed are less well understood. Closing that gap would connect the conceptual framework here to the practical question of how organizations under genuine uncertainty can be designed to improve over time.⁸¹

The existing literature on decision-making under uncertainty runs its behavioral and process strands in parallel more than in genuine conversation with each other. What the framework proposed here does is bring them together and specify the conditions under which each factor matters most. That's groundwork for empirical work. It's not a substitute for it.

⁷⁹Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 532–534.

⁸⁰Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 120–124.

⁸¹Gavetti, G., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 128–134.

What would constitute progress on that empirical agenda is worth being specific about. The most valuable studies would follow decisions longitudinally — tracking the process as it unfolds rather than reconstructing it after the outcome is known. Retrospective accounts of decision processes are systematically distorted by outcome knowledge: people remember the process differently depending on how things turned out. Prospective tracking is harder and more expensive, but it's the only way to get clean measures of the process variables the framework identifies as consequential. That kind of research is rare in the strategy literature. It's rare because it's slow and organizationally demanding. It's also what the field needs most.

Conclusion

When I started working on this thesis, I expected the problem to be cleaner than it turned out to be. The question of how managers decide under genuine uncertainty⁸² sounds like it should have a reasonably tidy answer. It doesn't. The more carefully you look at the separate literatures, the more obvious it becomes that none of them is quite sufficient on its own.

Chapter 1 establishes something uncomfortable. Individual reasoning is reliably unreliable under uncertainty. Overconfidence, anchoring, the planning fallacy: Kahneman and Tversky⁸³ didn't discover these as curiosities. They documented them as predictable features of how people think when stakes are high and information is thin.

Chapters 2 and 3 ask what you can do about it. Chapter 2 covers what the empirical literature says about strategic decision processes: when structured processes actually improve decision quality, how speed and comprehensiveness trade off across different environments, what mechanisms like devil's advocacy and pre-mortem can do to limit the damage cognitive bias causes, and how the dynamic capabilities framework argues that firms need the capacity to sense environmental shifts and reconfigure their decision routines in response. What I found less convincing was how rarely these frameworks ask who is actually doing the sensing and reconfiguring. Chapter 3 adds that question. It turns out the organizational conditions matter at least as much as the tools: whether people can surface problems, whether dissent reaches decisions before they're locked in. A firm with rigid hierarchy and weak psychological safety will struggle under uncertainty regardless of how sophisticated its strategy looks on paper.

Putting all three levels together was the part of this thesis I found most difficult to get right. The failure modes differ at each level, and so do the interventions, but they interact in ways that aren't always obvious. Sophisticated strategy tools don't help if the leaders using them are subject to the biases Chapter 1 describes. And behavioral training won't fix that if the organizational structure punishes the dissent those leaders need to hear. You can't address one dimension and leave the rest alone. The framework in Chapter 3 tries to hold all three in the same frame. The aim is diagnosis: where is the breakdown actually happening, and what would have any grip on it.

⁸² Knight, F. H. (1921). Risk, uncertainty, and profit. University of Chicago Press.

⁸³ Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.

The limits here are worth being direct about, and I've tried to be honest with myself about them throughout. No organizations were studied. No data was collected. The argument is theoretical, which means it inherits the gaps in the literatures it draws on. The dynamic capabilities tradition is still contested on what sensing and reconfiguring look like in practice. The behavioral debiasing research⁸⁴ is considerably messier in field conditions than in labs. These are real constraints. They limit how confidently the argument can be stated.

The natural next step is empirical work, and it's probably the most important gap this thesis leaves open. The core question is whether organizations that take the cognitive and organizational dimensions seriously actually make better strategic decisions under genuine uncertainty, and whether the interactions the framework describes show up in the data. Longitudinal case studies would be the right method. A cross-sectional snapshot can't tell you whether decision quality shifts as conditions change. The framework generates the hypotheses. Testing them is someone else's job.

None of this makes the theoretical work pointless. Frameworks that organize existing knowledge have value before anyone tests them. They focus attention and give practitioners a vocabulary for problems that otherwise get handled entirely on instinct. What struck me most, working through this material, was how consistently the different literatures point to the same gap: the absence of an integrated account that treats cognitive, strategic, and organizational factors as parts of the same problem. The argument here is modest, even if the terrain isn't: uncertainty in management isn't one problem. It's several, operating simultaneously at different levels. Most of the literature, read in isolation, offers a tidier picture. That's probably why it underperforms.

⁸⁴ Lovallo, D., & Sibony, O. (2010). The case for behavioral strategy. McKinsey Quarterly.

References

- Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Prentice-Hall.
- Dane, E., & Pratt, M. G. (2007). Exploring intuition and its role in managerial decision making. *Academy of Management Review*, 32(1), 33–54.
- Dean, J. W., & Sharfman, M. P. (1996). Does decision process matter? A study of strategic decision-making effectiveness. *Academy of Management Journal*, 39(2), 368–396.
- Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32(3), 543–576.
- Fredrickson, J. W., & Mitchell, T. R. (1984). Strategic decision processes: Comprehensiveness and performance in an unstable environment. *Academy of Management Journal*, 27(2), 399–423.
- Gavetti, G., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113–137.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482.
- Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 193–206.
- Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.
- Knight, F. H. (1921). *Risk, uncertainty, and profit*. University of Chicago Press.
- Lovallo, D., & Kahneman, D. (2003). Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81(7), 56–63.
- Lovallo, D., & Sibony, O. (2010). The case for behavioral strategy. *McKinsey Quarterly*.
- March, J. G. (1994). *A primer on decision making: How decisions happen*. Free Press.
- Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13(4), 75–90.
- Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1), 187–206.
- Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1369–1386.

- Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 111–128.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.
- Teece, D. J. (2016). Dynamic capabilities and entrepreneurial management in large organizations: Toward a theory of the entrepreneurial firm. *European Economic Review*, 86, 202–216.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458.
- Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
- Weick, K. E. (1995). *Sensemaking in organizations*. SAGE Publications.