Decision-Dominant Logic:
A Leadership Playbook for the Algorithmic Age

Prof. Dr. Roger Moser
Chairman SatSure & Faculty, University of St.Gallen / Macquarie University
Stars Symposium Alumnus

Leading in the Algorithmic Age

Over the last decades, leaders have been taught two big stories about how advantage is created. The first came from *Service-Dominant Logic*: focus on superior service, co-create value with customers, and build deep relationships that rivals struggle to copy. The second came from *Network Effect-Dominant Logic*: design platforms where every new user makes the system more valuable, then scale faster than anyone else. Both lenses still matter. Many successful firms are built on exactly these ideas.

But the competitive frontier has shifted. **In the algorithmic age**, advantage is no longer defined only by better services or larger network effects. It is defined by how quickly and reliably you can turn signals into better decisions at scale. Every function in your company now sits inside an invisible learning loop: data is captured, models are updated, policies are changed, and the next set of decisions is shaped in real time. The quality of those loops determines whether you move faster than the market or are slowly outpaced by it.

Service-Dominant Logic tells you who you create value with. Network Effect-Dominant Logic tells you where scale amplifies that value. **Decision-Dominant Logic** tells you how you actually convert both into a compounding advantage. **It treats decisions as the primary design unit of the organization:** which decisions matter most, who owns them, what information they use, how they are tested, and how they are improved over time. In an environment saturated with data, AI, and algorithms, this is the level at which leadership can still make the decisive difference.

The algorithmic age does not replace earlier logics. It reorders them. Services, platforms, and their inherent network effects remain important, but without a deliberate architecture for the decisions that run through them, they underperform. Decision-Dominant Logic provides that architecture. It asks leaders to step beyond buying tools or hiring more data scientists and instead build the algorithmic flywheels that compound learning from every significant choice the organization makes.

What follows is a practical leadership approach for doing exactly that.

Decisions as an Asset You Can Manage

Today, the most valuable asset in your company is not data, talent, capital, or code—it is the ongoing stream of decisions made by you and your people every hour, week, and quarter. Decisions set direction, allocate attention, pull resources, and quietly reshape the future in increments. When decisions are treated as the primary unit of value,

unexpected clarity emerges: waste becomes visible, learning becomes measurable, and competitive advantage becomes compounding.

Decision-Dominant Logic starts with a simple mindset shift for leaders: instead of asking "What should we build?" or "Which tool should we buy?", ***ask "Which decisions must we get right—again and again—to win?"*** Name those decisions, connect them to signals, record intent, study outcomes, and feed the learnings back into the next round. That is how you build an algorithmic flywheel: data sharpens models, models inform choices, choices produce outcomes, and outcomes generate new data. Each loop turns faster and truer, refocusing with every use.

Practical Example: Operations in Action

This is not an ivory-tower concept. Imagine a frontline team managing order flows through a distribution center. Their recurring decision is when to expedite a shipment. Traditionally, each expedite is a one-off judgment call. In a decision-dominant environment, every expedite is a small hypothesis: "If backlog exceeds X and carrier reliability drops below Y, expediting increases customer retention by Z." Triggers are explicit, guardrails are clear, outcomes measured. At the end of the week, managers gather for twenty minutes - not to blame or celebrate, but to ask: "What did we expect? What happened? What will we change?" This tone, curious rather than judgmental, turns work into an ongoing, humane experiment.

Applying the Logic to Strategy

While frontline teams make hundreds of decisions daily, senior teams make a few big calls whose impact lasts for years. Decision-Dominant Logic treats each strategic choice as a portfolio of tests: "What would have to be true for this bet to succeed? Which early signals would force us to adapt?" Here, leaders focus on learning quickly whether an idea deserves to scale. Quarterly reviews shift from theatre to disciplined conversations about evidence, thresholds, and the courage to change course when reality dictates.

Two Levels, One Habit

On the operations side, the habit is short cycles of explicit intent followed by rapid reflection. Strategically, the habit becomes clear tests with pre-committed trigger points for adaptation. Together, these create an organization that keeps its attention where truth is likely to reveal itself: in the outcomes of its own choices.

The Effect on Leadership and Organizational Culture

Executives may worry that this approach slows them down; in reality, precision enables speed. When triggers, options, and guardrails are explicit, teams stop debating shadows and start moving. Predictable cadence in reviewing outcomes reduces drama and increases momentum. When reviews focus on learning, not judgment, people surface weak signals early, releasing energy back into the work. The human dimension matters: decisions are made under pressure, which narrows perspective and raises errors of pride and fear. The antidote is not toughness slogans, but leaders who stabilize the room, separate the person from the problem, and model constructive self-talk when

outcomes disappoint. A steady tone becomes competitive advantage. Leadership is easier when decisions are grounded in shared facts and assumptions rather than defensive debate.

Building the Pattern: Outcomes, Decisions, Questions

Clarifying outcomes comes first. Whether walking the floor or joining a meeting, call out the result that matters now: the customer experience you want to protect, the service level you will not compromise, the margin or risk you are willing to accept this week. When a shared outcome is visible, preferences fade and alignment grows.

Next, define the decision. Make it specific, owned, and bounded. What exact call is being made? Who owns it? What options are really available, and what is not? In high-stakes moments, this might be a choice to release a product, change a price, or reroute supply. On the frontline, it could be the decision to expedite, override, or wait. Defined decisions turn vague action points into clear commitments for future review and improvement.

Formulate sharp questions. Three often suffice: What do we expect will happen if we take this path? What early signal would compel a pivot or fast adjustment? What constraint could we relax to learn faster without serious risk? These questions turn judgment from a leap of faith into a designed act, honouring agency and shrinking uncertainty.

Leadership Rituals for Continuous Improvement

Incorporate this approach as a weekly leadership ritual. Pick five material decisions from the past few days or weeks and replay each in about three minutes. For each, the owner states the outcome intended, the choice made, the trigger that justified it, and the actual outcome observed. The group asks: Where did expectations differ from reality? What will we decide differently next time? Which rule or model needs updating? The purpose is not finding heroes or culprits, but better outcomes, sharper decisions, and smarter questions.

This scales up for executive forums. Each portfolio bet begins with a crisp outcome statement: what success looks like in numbers or behaviors. Define the few decisive choices that drive the bet. Then address the questions: What must be true for success? Which signals could prove us wrong early? What will our posture be if those signals appear? This creates social acceptance of retiring logic that no longer works, without humiliation, preserving the ability to place bold bets again.

Mapping Decision Patterns Across the Organization

Over time, make the pattern visible through a living map of your firm's recurring outcomes that create or capture most value. For each critical outcome, list key decisions, their owners, triggers, menus of action, guardrails, and at least one clear outcome metric (but not more than three). Alongside each decision, record the hypotheses, disconfirming signals, and cadence for review. This map is not a decorative poster, it is a practical guide for how the organization connects outcomes, decisions, and questions to compete in different markets.

Signs Your Flywheel Is Working

As leaders consistently clarify outcomes, define decisions, and ask sharper questions, a few telltale signs will emerge: people start speaking in hypotheses, cycle times reduce because conversations no longer start from zero, and reversals become less emotional and more affordable because conditions were agreed in advance. Leaders find they can be both demanding and kind, precise and patient. The culture grows quieter and more candid, not out of apathy, but from focus on the next outcome, the next decision, and the next question to advance the system.

Your Role as a Leader

If you lead a large enterprise, your role is to set the cadence and protect the tone. Set cadence by ensuring every important decision has explicit outcomes and progress is reviewed promptly. Protect tone by rewarding learning, especially when outcomes disappoint. When a bet fails and the team brings evidence, treat it as performance, and you will see more of it. When a bet succeeds and the team refines future practice, amplify it for continuous improvement.

Start small. Focus on one customer journey, plant, or product line. List ten recurring, high-impact decisions in terms of value creation or capture. Instrument one with a visible trigger, outcome metric, and lightweight log, then run a fifteen-minute replay weekly. The following month, apply the model to a strategic initiative with decisive tests and disconfirming signals - put those signals on the calendar. Keep going.

Paradoxically, the more you formalize decision-making practices, the more human the work becomes. People feel trusted because expectations are clear. They feel safe because reviews are anchored in truth, not ego. They feel proud because improvement is visible and shared. This sustains pace without burning out the organization.

Decision-Dominant Logic is not about worshiping data or algorithms - it is about honouring the craft of good choices and relentless learning. Algorithmic flywheels compound this craft. Build it once in one part of the business, then again in another, and soon posture shifts: the company stands taller, argues less, adapts more, wastes fewer cycles, and earns more trust. That posture is what advantage feels like from the inside.

There will always be market noise, volatility in costs, and surprises in technology. You cannot control that. You can control the quality, cadence, and character of your decisions. Make those your focus, protect the tone, and let outcomes be your teacher. If you do, your organization will not just perform - it will learn at a pace competitors cannot match, and that compounding effect will become your signature towards employees, customers and competitors.

You can learn more about Decision-Dominant Logic through
Dr. Moser's LinkedIn Newsletter

****

Appendix:
How AI/ML and GenAI Operationalize the Decision-Dominant Logic through Algorithmic Flywheels

The promise of Decision-Dominant Logic in the form of algorithmic flywheels becomes real when any human/artificial intelligence is blended into the flow of work and into the cadence of strategic choice. Think of AI/ML as a torque amplifier. It sharpens the signals that trigger action, proposes stronger options at the moment of choice, and turns outcomes into cleaner learning. GenAI adds a conversational layer that drafts the artifacts leaders need, so discipline does not become bureaucracy. What follows is a narrative look at how this plays out in practice at two levels.

An Operational Example: Same-Day Fulfillment in a National Retailer

Picture a retail operations director who owns same-day fulfillment across 14 micro-fulfillment centers. The recurring decision is when to expedite a customer order from another node. In the old world, each expedite is a harried judgment call. In the new world, the team instruments the decision.

A small forecasting model watches demand and carrier reliability. An anomaly detector sits on top of picker throughput to spot stalls. When the backlog at a site rises above an adaptive threshold and the carrier on that route has slipped below a reliability floor, the system proposes an expedite. A short GenAI brief appears in the team's console with three elements the director trusts: what signal just fired, the ranked options with projected impact on on-time-in-full and contribution margin, and the guardrails that must hold if the action is taken. The brief also shows a costed counterfactual that answers the question people actually ask under pressure, which is what happens if we do nothing for one more hour.

The director still decides. The human check is fast because the object in front of them is standardized. After the shift, the flywheel continues. Logs capture the intent behind the choice, the cost of the action, and the lagged customer outcome. GenAI drafts the weekly two-page replay for the stand-up. It contrasts expectation with reality, calls out where the model's ranking did not match the observed result, and proposes one tweak to the trigger or the guardrail for the next cycle. People do not waste energy building slides. They spend their attention on the few places where judgment must evolve. After four weeks the expedite rate drops, on-time performance rises, and the room is calmer because everyone can see how learning is compounding.

A Strategic Example: The Annual Capital Allocation Across Growth Bets

Now zoom out to a once-a-year decision. The executive team must allocate next year's growth capital across three bets: a new market entry, a product line extension, and a pricing architecture change. In the old world, this meeting is theater. Long decks, thin evidence, and hardened positions. In the new world, the flywheel supplies an evidence engine.

Three months before the decision window, a small strategy cell uses GenAI to draft a set of explicit assumptions for each bet and the few disconfirming signals that would kill

the logic early. Market and customer telemetry feeds a set of lightweight models that watch demand elasticity, switching costs, and competitor moves. A simulation sandbox lets the team practise choices safely. The CFO can ask to see what happens if conversion moves by half a point while acquisition costs rise by ten percent, and the room can see the ripple through margin, capacity, and cash.

In prep, GenAI produces a three-page narrative per bet. It opens with the plain-language thesis and the two numbers that would most rapidly prove it wrong, then lays out the investment ramp, the pre-committed pivot points, and the first quarter's learning plan. It also generates a red team critique for each option so the room hears an intelligent counter-story before anchoring. On the day, the discussion is shorter and cleaner. The group is not voting on whose deck is more persuasive. It is agreeing on which assumptions deserve capital, what evidence must arrive by specific dates, and how the organization will adapt if the world votes no.

The flywheel does not end with the vote. Throughout the year, the same models watch the agreed signals. At each quarterly forum the GenAI brief presents a one-page update in the same structure. If a disconfirming signal crosses the line, the team already knows the posture. It redeploys capital without drama and publishes a short learning memo that improves the next generation of bets. The firm earns the right to be bold because it has made reversals cheap in advance.

The Pattern That Links Both Stories

At the operational level, AI tightens triggers and GenAI standardizes decision craft, so teams move faster with less noise. At the strategic level, AI surfaces weak signals and GenAI turns assumptions into living tests, so big choices learn in public. In both cases humans keep ownership of intent, ethics, and tone. The technology makes the work more disciplined without making it more brittle. The cadence makes learning visible. Over time, people speak the same language from the warehouse floor to the boardroom table, and the organization acquires a posture that competitors find hard to copy: calm under pressure, precise in choice, and relentless in learning.