The Innovator's Solution

By Clayton M. Christensen and Michael E. Raynor — 20-Minute Read Summary

Daily Book, Episode 47 — Sequel to The Innovator's Dilemma (Episode 46)

The Book in 60 Seconds

The Innovator's Dilemma (1997) explained why great companies fail: by listening to their best customers, investing in better products, and chasing higher margins, they get disrupted from below. The Innovator's Solution (2003) answers the question executives kept asking afterward: what do I actually do about it on Monday morning?

Christensen and Michael Raynor turn the theory of disruption into a management playbook. The book gives you tools to spot a genuine disruptive opportunity, to test whether your own organization can pursue it, and to structure the effort so it survives inside a big company. Its core ideas: disruption comes in two forms (low-end and new-market); customers “hire” products to get jobs done; an organization's capabilities — its resources, processes, and values — determine what it can and cannot do; disruptive businesses need patient money; and disruption can become a repeatable capability, not a one-off stroke of luck.

If The Innovator's Dilemma is the diagnosis, The Innovator's Solution is the treatment plan.

The Central Argument

The book opens with a problem of method. Most management advice is based on stories: a successful company did X, so you should do X. Christensen argues that stories are not enough, because the same action can succeed in one circumstance and fail in another. What managers need is a theory — a statement of what causes what, and under which circumstances it holds.

The theory of disruption is such a theory. It says: when a technology or business model improves faster than customers can use it, the basis of competition changes; simpler, cheaper, more convenient products enter at the bottom or in entirely new markets; incumbents, behaving rationally, flee upmarket; and the entrant climbs until it owns the market.

The Innovator's Solution takes that theory and asks how a manager can use it predictively — to see a disruption before it happens, to build one deliberately, and to defend against one. The book is organized around a series of questions a manager should ask:

1. Is this opportunity genuinely disruptive? Or is it a sustaining innovation in disguise?
2. What job is the customer trying to get done? And who is not being served at all?
3. Can my organization do this? Do we have the processes and values it requires?
4. Is the money patient enough? Will the funding let the business grow on its own timeline?
5. Where will the profits in this industry go? And how do I position for that?

Each question gets a framework, and each framework gets tested against real cases.

Two Kinds of Disruption

The first refinement of the theory is the most important. There are two different kinds of disruptive footholds, and confusing them is dangerous because they behave differently.

Low-end disruption targets the least demanding, least profitable customers at the bottom of an existing market. The classic case is the steel minimill. Nucor and its peers began by making rebar — the cheapest, lowest-quality steel product, the kind integrated mills barely bothered with. The minimills' cost structure let them profit at prices the integrated mills couldn't match. Then they got better, moved into larger steel products, and eventually took the whole market. The incumbents didn't fight for rebar because it wasn't worth fighting for — and by the time the fight mattered, it was too late.

New-market disruption is different and, Christensen argues, even more important. It doesn't steal existing customers at all. It creates customers who didn't exist — people who couldn't use the product because it was too expensive, too complicated, or too inconvenient. These are non-consumers, and they represent the largest untapped opportunity in business.

The book's great example is Sony and the transistor radio. In the 1950s, the major American electronics companies made radios that sat in the living room: large wooden consoles, expensive, beautiful pieces of furniture. Sony introduced a small pocket radio that sounded worse. The incumbents dismissed it — who would want a worse radio? But Sony wasn't selling to living-room radio buyers. It was selling to teenagers who had never owned a radio, because a radio had always belonged to the family, in the living room, controlled by the parents. A pocket radio was a personal radio. It created an entirely new market.

Honda's entry into the American motorcycle market follows the same pattern. The big American and British manufacturers sold large, powerful, expensive bikes to serious enthusiasts. Honda arrived with tiny 50cc bikes — the Super Cub — and was laughed at. But Honda wasn't selling to bikers. It was selling to people who had never ridden a motorcycle, who wanted cheap, easy transport. Once those new customers were in, Honda moved upmarket into bigger bikes and eventually dominated.

The test: an idea is genuinely disruptive if (a) it is aimed at customers the incumbents are happy to ignore — either the least profitable customers or people who aren't customers at all — and (b) the business can make money at a low price, with a cost structure designed for that low price from the start. Both conditions matter. A cheaper product built on an expensive cost structure will eventually drift upmarket and become just another incumbent.

Asymmetric motivation

Why don't incumbents fight back? Because both sides are motivated in opposite directions. The entrant wants to move upmarket, where the money is. The incumbent wants to flee upmarket, where the margins are. So the entrant climbs and the incumbent retreats, and they never actually fight until it's too late. Disruption is not a battle. It is a slow, polite handover of a market — and both sides are behaving rationally given their incentives.

Jobs to Be Done

The most famous idea to come out of this book is jobs to be done: customers don't buy products; they hire products to get a job done in their lives. The idea later became the subject of Christensen's book Competing Against Luck (2016), but it is introduced and developed here.

The book's example is the milkshake. A fast-food chain wanted to sell more milkshakes, so it did what every company does: it asked milkshake drinkers what they wanted. Better flavor, they said. Thicker. More chocolate. The company improved the milkshake, and sales didn't move.

Then Christensen's team watched when and why people actually bought milkshakes. Almost half of all milkshakes were sold before 8:30 in the morning, to one person, alone, in a car. Those morning buyers weren't really buying a milkshake. They were hiring it to do a job: survive a long, boring drive to work, with one hand on the wheel, staying full until mid-morning. A milkshake was perfect — it took twenty minutes to finish through a thin straw, it was filling, and it didn't spill easily.

The milkshake's real morning competitors weren't other milkshakes. They were bagels, bananas, doughnuts, and coffee. A bagel is dry, boring, and finished in two minutes. A banana is gone in thirty seconds. A doughnut leaves crumbs on your suit. The milkshake won because it did the job better.

Contrast that with the afternoon milkshake buyers: mostly parents buying treats for their kids after school. The job there was completely different — be a good parent, give your child a treat, show love, and do it quickly before dinner. Those parents didn't care about thickness or the straw. Same product, same restaurant, two completely different jobs.

The implications are far-reaching:

- Segment by job and circumstance, not by customer attributes. Age, income, and demographics tell you almost nothing about why someone buys. The morning commuter and the afternoon parent look the same on paper and hire the same product for totally different reasons.
- Improving a product for one job can make it worse for another. Thicker is better for the parent (the child takes longer to finish it) and worse for the commuter (now they can't get it through the straw).
- Traditional market research fails with disruptive products. Ask people what they want and they describe a slightly better version of what they already have. Ask what job they're trying to get done, and watch how they do it today, and you find the opening.
- Non-consumption is a signal. People who have never bought your kind of product aren't necessarily uninterested. They may be waiting for a product that actually fits the job.

Jobs to be done connects directly back to the two kinds of disruption. A low-end disruption serves a job that existing customers are overserved on — they pay for performance they don't need. A new-market disruption serves a job that nobody is serving at all.

More cases

The book applies the jobs lens to several industries. Charles Schwab's online trading: the big brokerages sold advice, research, and personal service at high commissions. Schwab's online product offered none of that — just a cheap way to buy and sell shares. It served a new group of self-directed investors and people who had never invested because the old model was too expensive and intimidating.

Intuit's QuickBooks: accounting software was sold to trained bookkeepers. QuickBooks was built for the small-business owner doing the books at the kitchen table, late at night, with no training. The job wasn't “produce a perfect set of accounts.” It was “get the books done quickly, without learning accounting, and feel confident the tax return is right.”

Southwest Airlines: a low-end disruption of the major airlines — no assigned seats, no meals, no hub-and-spoke system, just point-to-point flights, fast turnarounds, and low prices. And Southwest didn't only steal passengers; it created them. People who would have driven four hundred miles now flew, because flying had become cheap enough and simple enough to beat driving. New consumption is the signature of a real disruption.

RPV: Resources, Processes, Values

The most practically useful framework in the book for managers is RPV — resources, processes, and values. It is a diagnostic for what an organization can and cannot do.

- Resources are people, money, technology, brand, patents — the things a company has.
- Processes are the way work actually gets done — routines, checklists, how a decision travels through the organization, the habits that let a company do the same thing reliably, thousands of times.
- Values are the criteria by which the company decides what is worth doing — what counts as a big enough opportunity, what margins are acceptable, what kind of customer matters.

When a company is young, its capabilities live mostly in its resources, in its people. As it grows and succeeds, its capabilities migrate into its processes and values. That is what makes a big company reliable. But those same processes and values become disabilities when the world changes. Capabilities define disabilities. The processes and values that make a company great at sustaining innovation make it terrible at disruptive innovation, because disruption requires doing things the company's processes aren't built for and pursuing opportunities its values say are too small or too low-margin to matter.

Before launching a disruptive effort, ask three questions:

1. Do we have the resources — the people, money, and technology?
2. Do we have the processes — the way of working — that this new business needs?
3. Do our values allow it? Will the organization get excited about a market that is small today, with low margins, serving customers we have always ignored?

For most big companies, the honest answer to the second and third questions is no. Their processes are built to serve demanding customers with high-performance products. Their values say a new opportunity has to be big enough to move the needle for a billion-dollar company. A disruptive market starts small and low-margin by definition, so the company's own immune system rejects it.

The prescription is structural, not motivational. If the disruptive business doesn't fit your processes and values, don't try to force it inside the existing organization. Create a separate organization with its own processes, its own values, its own cost structure, and its own definition of success. Give it the resources it needs — and then leave it alone.

A separate disruptive unit needs its own profit-and-loss statement, so it can be judged on its own economics; a leader genuinely committed to the new business, not someone doing it as a side assignment; and protection from the main business, which will always want its resources back when times get tough.

The nuance matters: sometimes you don't need a separate organization. If your processes actually fit the opportunity, integrating it can work. RPV is a diagnostic tool. It tells you where the mismatch is, so you can decide whether to fix the organization, spin the new business out, or acquire the capability from outside. The problem is structural, and the fix is structural — you don't pep-talk your way out of a values mismatch.

Good Money vs. Bad Money

Christensen draws a distinction between good money and bad money — not clean versus dirty, but patient versus impatient.

Good money is capital that understands a disruptive business starts small, and accepts modest returns in the early years while the business finds its footing. Bad money demands big, fast returns. And bad money forces a company into sustaining innovation: if you take money that expects a huge return quickly, you can't afford to pursue a small new market. You have to go after the biggest, most demanding customers, because that's the only way to generate the returns your investors expect.

The amount of money you take, and the returns you promise, determine your strategy more than your intentions do. You can have the best disruptive idea in the world, and if you fund it with impatient capital, you will end up building a sustaining product for the mainstream market.

This explains a familiar startup pattern: a company takes venture money, the rounds get bigger, the growth targets get steeper, and suddenly it's chasing enterprise customers it never wanted, because that's where the revenue is.

It applies inside big companies, too. When a big company funds an internal venture, the funding usually comes with the expectation that it will grow fast enough to matter to the parent. A disruptive business doing ten million dollars in revenue is a triumph for a startup and a rounding error for a ten-billion-dollar company — so the internal venture gets killed, not because it failed, but because it succeeded too slowly for the parent's values.

The prescription: match the money to the market. A disruptive business needs a cost structure that can be profitable at a small scale, and investors — internal or external — who are comfortable with that. One of senior management's main jobs is to protect a young disruptive business from the parent company's own growth expectations until it is strong enough to stand on its own.

There is a beautiful inversion here: having less money can be a strategic advantage, because it forces you to build a business that works at a low price point, serving customers nobody else wants. By the time you've built that business, you have a cost structure and a customer base the incumbent can't match without destroying its own margins.

The Law of Conservation of Attractive Profits

The book's most abstract-sounding idea has enormous practical consequences. In any industry, there is always a place where the profits concentrate — and that place moves over time, in a predictable way.

When a product is new and “not good enough,” the companies that win are the ones that integrate everything — controlling the whole system, because integration is the only way to make a not-good-enough product work. Think of the early personal computer or the early smartphone. Profits concentrated at the integrated companies.

But once the product gets “good enough” — once the technology overshoots what customers can use — the architecture opens up. The product becomes modular: standard parts, standard interfaces, anyone can assemble it. And when that happens, the profits migrate away from the assemblers and toward the parts of the value chain that are still proprietary, still hard to copy. In personal computers, the profits moved from the computer makers to the makers of the microprocessor and the operating system. The assemblers were left competing on price.

The same pattern plays out everywhere. In smartphones, profits concentrated at Apple, which stayed integrated, and at the companies that owned the hard parts — the chips and the app store. In electric cars, the open question is where the profits will settle as the architecture opens up: the battery, the software, the charging network, or the car itself? The law says the profits will go wherever the bottleneck is — wherever the part is that is still not good enough and still proprietary.

For disruption, this matters because it tells you where to position. If you're entering a market where the architecture is becoming modular, don't try to be the assembler. Own the piece that is still hard. If you're an incumbent in an integrated position, expect modularization to come for your margins, and think about which proprietary piece you can own next.

The Disruptive Growth Engine

The book's final move is to assemble all of these tools into something it calls a disruptive growth engine: a deliberate, repeatable capability for launching disruptive businesses — not once, but over and over.

The engine has four parts:

1. A theory. Good management theory is not academic decoration; it is a practical tool. A theory tells you what will happen and why, so you can make decisions before the data exists.
2. A process for finding the jobs. A disciplined way to look for jobs to be done and non-consumers — the places where a new market is waiting.
3. Structural discipline. The willingness to spin out the new business when your processes and values don't fit, and to fund it with patient money.
4. Patience. Letting the business grow on its own timeline, protected from the parent company's growth expectations.

This reframes innovation from a lightning strike into a capability. Most companies treat innovation as something that happens when a genius has an idea. Christensen argues that disruption can be a managed process, like quality control or supply-chain management. You can build an organization that is systematically good at finding and growing new markets.

The practical checklist, distilled:

1. Find a job that isn't getting done, or a group of non-consumers.
2. Ask whether the job is in a new market or at the low end of an existing one.
3. Check whether the business can be profitable at a low price, with a cost structure built for that price.
4. Run the RPV test on your own organization, and be honest about whether your processes and values can support it.

5. Make sure the money is patient.
6. If the fit is wrong, spin the new business out and protect it.
7. Ask the hardest question: are you willing to pursue an opportunity that is small today, knowing it might cannibalize your existing business tomorrow?

That last question is almost a moral one for a manager. The disruptive business will always look like a bad idea on the numbers. The decision to pursue it is a decision to bet on the future over the present — and most organizations, quite rationally, choose the present.

Disruption in Your Own Work

You don't need to run a company to use this book. The frameworks scale down.

If you lead a team, run the RPV test on your own group before taking on a new kind of work. A team that is excellent at shipping polished features for existing customers may be genuinely bad at exploring a rough new product — not because the people are wrong, but because the team's processes (its review cycles, its quality bar, its definition of “done”) and its values (what counts as a win) are tuned for the old work. The fix is the same at small scale: carve out a protected space with different rules, different metrics, and a different definition of success, and defend it from the main team's standards until the new thing is strong enough to survive them.

If you build products, make “jobs to be done” a habit. Before adding a feature, ask what job the user is hiring the product for in that moment — and what they're hiring instead when they don't use you. The answer is often humbling: your competitor is not the other app in your category, but a spreadsheet, a phone call, or simply doing nothing.

And if you invest your own time, treat it like the book treats money. Attention is patient capital or impatient capital. A side project funded with impatient attention — one that must "pay off" within weeks — will always drift toward the safe, sustaining version of itself. The disruptive version needs a longer runway and a smaller definition of success at the start.

Memorable Lines

> *"Customers don't buy products. They hire them to get jobs done."*

> *"Capabilities define disabilities."*

> *"A good theory tells you what will happen and why, so you can make decisions before the data exists."*

> *"The amount of money you take, and the returns you promise, determine your strategy more than your intentions do."*

Fair Criticisms

The Lepore critique. The most famous criticism of the disruption theory came from historian Jill Lepore in The New Yorker (2014). She argued that Christensen cherry-picked his cases, that many of the companies he said were doomed actually survived, and that "disruption" had become a buzzword justifying reckless destruction. There is truth in parts of this: Seagate, one of the disk-drive companies used as a cautionary tale, survived and adapted; some specific predictions did not pan out; and the word "disruption" has been wildly overused.

Christensen's reply. In 2015, in the Harvard Business Review, Christensen clarified what the theory does and does not claim. Disruption is a specific process, not a synonym for any new technology. Uber, for example, is not a disruption in his sense: it did not start at the low end or in a new market, but in the mainstream taxi market, serving existing customers with a better service. It is a sustaining innovation that happened to be very successful. The clarification makes the theory more falsifiable — and therefore more useful.

The limits. The theory works better as a way of thinking than as a precise prediction machine. It helps you ask the right questions; it doesn't guarantee the right answers. The book is also rooted in technology and manufacturing, and some of its frameworks are harder to apply to services and software. And it assumes a company can see a disruption coming if it looks carefully — in practice, the hardest disruptions to spot come from outside your category entirely. A hotel company in 2008 was not watching a website for renting air mattresses.

Why It's on the List

This is the sequel to The Innovator's Dilemma (Episode 46), and it completes the picture: the first book explains why disruption happens; this one tells you what to do about it. The “jobs to be done” framework alone has reshaped how product teams think about customers, and the RPV framework is one of the clearest explanations ever written of why big organizations can't do what looks, from outside, like the obvious thing. For anyone who manages a team, builds a product, or runs a company, this is the more useful of the two books.

How to Approach It

Read it with a specific problem in mind — a product you're building, a market you're watching, a team you're running. The frameworks reward application. It is denser than The Dilemma, more of a manual, so don't read it as a story; read it as a toolkit, and test each tool against something real.

If you haven't read The Innovator's Dilemma, this book recaps the core theory before going further, so it stands on its own reasonably well — but reading them in order is better.

What to Read Next

- Competing Against Luck (2016) — Christensen's full deep dive on jobs to be done; the most directly useful of his books for a product person.
- Good Strategy Bad Strategy by Richard Rumelt — the strategy companion to the disruption playbook.
- Only the Paranoid Survive by Andy Grove — the other classic on what to do when your industry is being turned upside down.
- The Design of Everyday Things by Don Norman — next in this series (Episode 48), the design-side companion: why everyday objects are so often confusing, and what good design actually is.