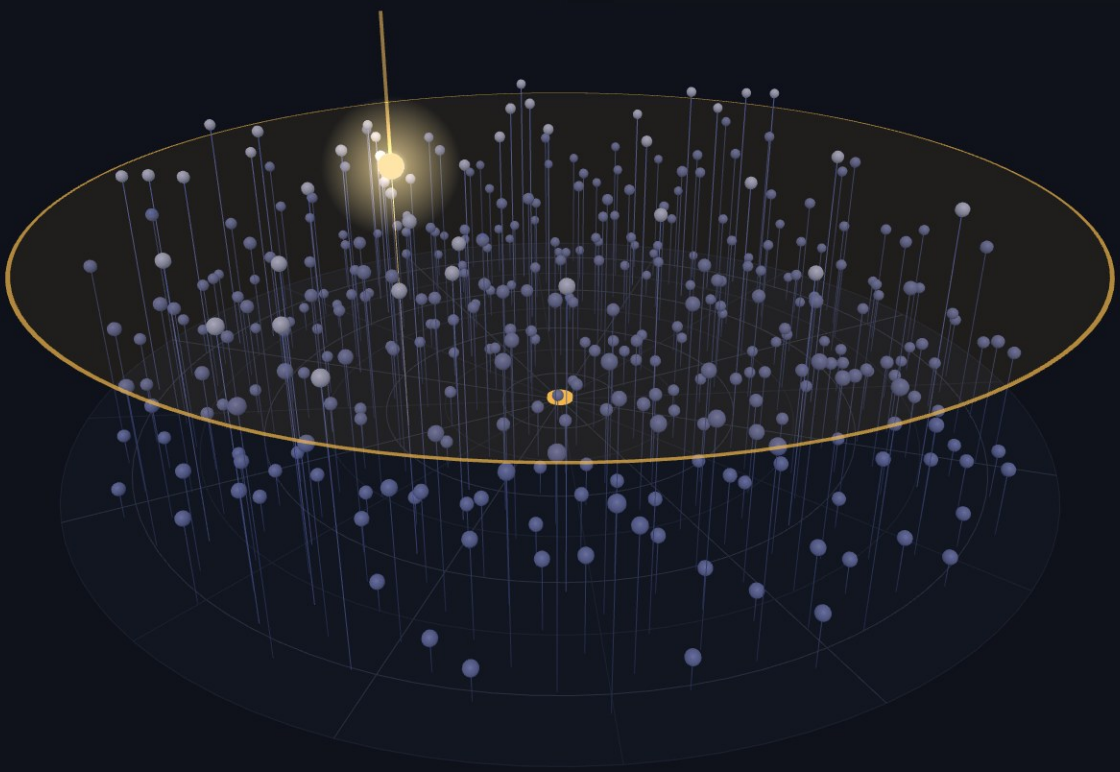
A BOOK

Good *Enough*

Herbert Simon's satisficing in the age of tireless search

JOHN RECTOR



A BOOK

Good *Enough*

Herbert Simon’s satisficing in the age of tireless search

JOHN RECTOR

johnrector.me · November 2026

CONTENTS

Part One · The Stopping Rule

1. Forty Minutes on a Friday
2. The Man Who Named Enough
3. The Bar Moves
4. Good Enough Was Often Better

Part Two · The Searcher That Never Tires

5. Search Was the Cost
6. A Delegate With No Reason to Stop
7. The Old Problem of When to Stop

Part Three · The Bar Is the Job

8. When Search Is Free, the Bar Is the Decision
9. Never Show Me the Runner-Up
10. The Treadmill and the Ratio

Part Four · Enough, Everywhere

11. The Rents on Stopping Early
12. The Firm That Never Stops Searching
13. Enough Is a Person

Part Five · The Coming Decade

14. What I Expect

Epilogue. Friday, Again

PART ONE

The Stopping Rule

Why every sensible person who ever lived has settled for good enough, and why that was never a failure.

CHAPTER 1

Forty Minutes on a Friday

> **FRIDAY, 7:12 P.M.**
>
> *You told yourself you would pick a place for dinner before you left the house. It is now seven forty, you have eleven tabs open, three of them are the same restaurant, and the one with the best reviews has nothing until nine.*

You know this evening. Everyone does. You are not short of options. You are drowning in them. Every place is fine. Two or three are probably wonderful. And somewhere in those tabs is the one that would be perfect, if only you could be sure which one it is. So you keep looking. The kitchen at the place you would have chosen in the first two minutes is filling up while you compare photographs of the same plate of pasta.

When you finally book, you feel a small, familiar flatness. Not because the choice is bad. It is a good choice. You feel it because you saw the others. Somewhere in the back of your mind the runner-up is still open in a tab, and it will sit at the table with you all night.

Here is the strange part. If you had been born a hundred years earlier, this evening could not have happened. You would have walked down the street, looked in two windows, smelled one kitchen and gone in. You would not have called that settling. You would have called it dinner.

The difference between those two evenings is not character. It is not discipline or modern impatience. It is the cost of looking. For almost all of human history, looking was expensive, so people looked a little and stopped. Then, in a single generation, looking became cheap for the eyes

but stayed expensive for the mind. We got the menus of every restaurant in the city and kept the same tired attention we had before. Friday night is what happens in the gap.

A third evening is now coming, and it is the subject of this book. In that evening, looking becomes free for the first time in history, not for you but for something that works for you. Your AI delegate can read every menu, every review and every reservation book in the city in the time it takes you to put on your shoes. It never gets tired. It never feels the flatness. It has no reason to stop.

That changes the oldest rule of choosing. And it moves the one part of the decision that cannot be handed off into the center of your life.

WHAT THIS BOOK ARGUES

In 1956 an economist and psychologist named Herbert Simon gave a name to the way real people choose. He called it satisficing: setting a bar for what is good enough and taking the first option that clears it. He showed that this is not laziness or a failure of reason. It is the right answer when looking costs something.

The argument of this book is simple. AI is about to make looking free for our delegates. When that happens, the reason for stopping early disappears, but the bar does not. The bar becomes the whole decision. The delegate can find the best of everything. Only you can say what is enough. The people, companies and products that understand this will feel calm and do well. The ones that turn free search into endless comparison will make everyone richer in options and poorer in peace.

THE CHAPTER IN ONE LINE

When looking is free, the only decision left is what counts as enough.

CHAPTER 2

The Man Who Named Enough

Herbert Simon spent most of his career at the Carnegie Institute of Technology in Pittsburgh, and he was hard to classify. He wrote about public administration, economics, psychology, management and computer science, and he did important work in all of them. In 1978 he received the Nobel Memorial Prize in Economics for his research on how decisions are actually made inside organizations.

His starting point was an observation so plain it was almost rude. The economics of his day assumed a decision maker who knew every option, could predict every outcome and always picked the best. Simon looked at real offices and real people and saw nothing of the kind. A purchasing manager did not survey every supplier on Earth. A family did not compare every house in the city. A chess player did not consider every possible move, because no one can. The number of possible games of chess is larger than the number of atoms in the observable universe.

So what do real people do? Simon's answer, first sketched in his 1947 book Administrative Behavior and sharpened in two papers in 1955 and 1956, was that they set an aspiration level, a bar, and they search until they find something that clears it. Then they stop. In the 1956 paper he needed a word for this, and he made one by joining satisfy and suffice. Satisficing.

TWO BLADES

Simon's most famous image for this came later. He said human rational behavior is shaped by a pair of scissors. One blade is the structure of the environment: how the options are laid out, how hard they are to find, how much each one reveals about the others. The other blade is the

capability of the person choosing: their time, their memory, their attention, their ability to compute.

FIGURE 1

Simon’s scissors

Blade one: The world	Blade two: The searcher
· How the options are laid out · How hard each one is to find · How much one option tells you about the others	· Time · Memory · Attention · The ability to compare

A decision is cut by both blades at once. Satisficing is the fit between a limited searcher and a costly world. AI changes the second blade for the delegate, and leaves it unchanged for you.

You cannot understand a decision by looking at only one blade. A good shopper in a well-organized market looks brilliant. The same shopper in a chaotic bazaar looks foolish. Nothing about the shopper changed. The other blade did.

This is the first principle the rest of the book stands on. Satisficing is not a flaw in the person. It is the fit between a limited mind and a costly world. When one blade changes, the right way to choose changes with it.

THE SAME WINTER

There is a detail in Simon's story that matters more now than it did then. In the same months that he was naming satisficing, he and his colleague Allen Newell, working with the programmer Cliff Shaw, built one of the first programs that could be called artificial intelligence. It was called the Logic Theorist, and it proved theorems from Whitehead and Russell's Principia Mathematica. Simon later recalled telling his students early in 1956 that over Christmas he and Newell had invented a thinking machine. They presented it that summer at the Dartmouth

workshop where the term artificial intelligence was coined.

The Logic Theorist did not search every possibility either. It used rules of thumb to look in promising places first and stopped when it found a proof. In other words, it satisficed. The man who explained why humans settle for good enough built the first machines that search, and he built them to settle too, because search was expensive for machines as well.

Seventy years later, the descendants of those machines are fast enough, cheap enough and tireless enough that the blade Simon held fixed has started to move.

THE CHAPTER IN ONE LINE

Satisficing is the fit between a limited searcher and a costly world. Change either blade and the right way to choose changes.

CHAPTER 3

The Bar Moves

The most important word in Simon's account is not satisficing. It is aspiration. The bar is where all the action is.

Where does the bar come from? Simon's answer was that it is learned. When you find good options easily, your aspiration rises. When you search and search and find nothing that clears it, your aspiration falls. The bar is not a fixed preference written somewhere inside you. It is a running estimate of what the world will give you, adjusted by what the world has given you lately.

You can see this in your own life. The first apartment you rent after college is a palace if it has a dishwasher. Five years later the same apartment is a compromise. Nothing about the apartment changed. Your bar rose, because the places you had lived in taught it to.

THE FIRM HAS A BAR TOO

In 1963 two of Simon's colleagues, Richard Cyert and James March, carried the idea from people to companies in A Behavioral Theory of the Firm. They argued that organizations do not maximize profit in any complete sense. They set targets based on their own past performance and on what similar firms achieve, and they keep doing what they are doing as long as results clear the target.

Only when results fall below the bar does the organization start to search. Cyert and March called this problemistic search. A company that is doing fine does not go looking for a better supplier, a better process or a better strategy. It goes looking when something hurts, and it looks first in the neighborhood of the problem.

Two more ideas came out of their work that will matter later in this book. The first is slack: the spare capacity, extra people and comfortable margins that build up when an organization is clearing its bar easily. Slack looks like waste. It is also what lets an organization absorb a shock without breaking. The second is that the bar rises and falls with results, so a long run of good years quietly raises what everyone inside considers normal.

> *The bar is not a preference written inside you. It is a running estimate of what the world will give you.*

WHY THIS MATTERS NOW

If the bar is learned from experience, then anything that changes your experience changes your bar. A tool that reliably finds you better options will, over time, raise what you consider good enough. This is not a side effect. It is the mechanism. We will come back to it in Part Three, because it is the reason a world that always delivers the best can start to feel like nothing special.

THE CHAPTER IN ONE LINE

Your bar is learned from what the world has handed you. Change what arrives and the bar moves with it.

CHAPTER 4

Good Enough Was Often Better

For a long time satisficing was treated as the poor cousin of optimizing. Optimizing was what you would do if you were smart enough and had the time. Satisficing was what you did because you were not and did not.

Two lines of research have turned that around.

FAST AND FRUGAL

The first came from the German psychologist Gerd Gigerenzer and his colleagues, starting in the 1990s. They studied simple rules of thumb, which they called fast and frugal heuristics, and found that in many real environments these rules did as well as elaborate calculations, and sometimes better. A rule that looks at one good cue and ignores the rest can beat a model that weighs everything, because the world is noisy and the elaborate model ends up fitting the noise.

Gigerenzer's phrase for this was ecological rationality. A decision method is not rational in the abstract. It is rational in an environment. That is Simon's scissors again. The rule of thumb is not a crippled optimizer. It is a tool shaped to a world where information is partial and time is short.

DOING BETTER, FEELING WORSE

The second line of research is the one that should worry anyone building AI. In the early 2000s the psychologist Barry Schwartz and his colleagues began to measure the difference between people who tend to maximize, always looking for the best, and people who tend to satisfice. In 2004 he gathered the work into a book called The Paradox of Choice.

The most striking study followed graduating college students through their job searches. The maximizers looked at more options and landed jobs that paid noticeably more. By any outside measure they did better. They also felt worse. They were less satisfied with the jobs they took, more anxious during the search and more likely to wonder whether they had chosen wrong. The researchers, Sheena Iyengar, Rachael Wells and Schwartz, titled the paper Doing Better but Feeling Worse.

The reason they found is the key to the rest of this book. The maximizers did not suffer from the search itself so much as from what the search left behind: the options they had seen and not taken. Every rejected job that was attractive in some way stayed with them. They fixated on what they had passed up. The satisficers had seen fewer alternatives, and so they had fewer ghosts.

> *The pain of maximizing does not come from the search. It comes from the runners-up you saw and did not take.*

TWO THINGS AT ONCE

So we arrive at the end of Part One holding two facts that seem to pull against each other. Maximizing finds better outcomes. Maximizing makes people unhappy with them. For the whole of human history, those two facts traveled together, because the person who did the searching was the same person who had to live with what the searching revealed.

That is the bond AI is about to break.

THE CHAPTER IN ONE LINE

Maximizing gets better results and leaves more regret, because the searcher sees the options it passed up.

PART TWO

The Searcher That Never Tires

What happens to a world built on the cost of looking when looking becomes free.

CHAPTER 5

Search Was the Cost

Look around and you will find that a surprising amount of the economy was built on one quiet fact: people stop looking early.

In 1961 the economist George Stigler wrote a paper called The Economics of Information. His question was why the same product sells for different prices in the same town. His answer was that finding the lowest price costs time and effort, so buyers search only a little, and sellers can charge more than the lowest price without losing everyone. The spread of prices in a market, he argued, is a measure of how ignorant buyers are, and how expensive it is for them to stop being ignorant.

Once you see it, you see it everywhere.

THE RENTS ON STOPPING EARLY

A brand is, among other things, a stopping rule. The old line that nobody ever got fired for buying IBM was not about computers. It was about the cost of searching and the risk of being blamed. Buying the known name let a manager stop looking with a clear conscience.

A default is a stopping rule that someone else wrote for you. The plan your employer picks, the setting the phone ships with, the box already checked on the form. Most people keep the default, not because they have studied it but because changing it means looking, and looking costs.

The top result in a search, the product at eye level on the shelf, the restaurant at the head of the list, the plumber who calls back first: each of these earns money because the person choosing will often take the

first acceptable thing they see. Whole industries exist to sell position, and position is valuable only because human search is short.

FIGURE 2

Five rents on stopping early

01 **The brand.** Buy the known name and stop looking with a clear conscience.

02 **The default.** A stopping rule someone else wrote for you.

03 **First position.** The top result, the eye-level shelf, the head of the list.

04 **First to answer.** The contractor who calls back today beats the better one who calls tomorrow.

05 **The renewal.** The policy that rolls over because reconsidering costs.

Each of these earns money because the person choosing takes the first acceptable thing they see. Each loses its value when the searcher is a delegate that never stops early.

None of this is a conspiracy. It is simply what happens when sellers meet buyers whose attention is limited. The sellers collect a rent on the buyer's stopping rule. The rent has been so steady for so long that most people mistake it for the natural shape of commerce.

THE SCALE OF THE QUIET ASSUMPTION

Think of how much of daily life rests on it. Why does a mediocre restaurant on a busy corner survive? Because people walking by satisfice. Why does the contractor who answers the phone get the job over the better contractor who calls back tomorrow? Because the homeowner has a bar, and the first one to clear it wins. Why do renewal notices work, and subscriptions persist, and insurance policies roll over year after year at prices no one would choose fresh? Because reconsidering costs, and so people do not.

Every one of these is a bet that the buyer will stop early. Hold that thought.

THE CHAPTER IN ONE LINE

Much of commerce collects a rent on the fact that people stop looking early. Brands, defaults and first position are all stopping rules.

CHAPTER 6

A Delegate With No Reason to Stop

Now change one blade of Simon's scissors.

An AI delegate, the AI that works on your behalf and answers to you, does not have your limits. It can read every menu in the city, every review, every price, every reservation book, every warranty and every complaint filed against the business in the last ten years. It can compare two hundred plumbers on the work they have actually done. It can read every clause of every insurance policy you might switch to. It does not get tired at seven forty on a Friday. It does not feel the flatness. For the delegate, the cost of one more look is close to nothing.

In Simon's terms, the capability blade of the searcher has just become enormous, and the environment blade has become transparent. The world is the same shape as before, but the searcher can now see all of it at once.

When that happens, the economic reason for satisficing disappears for the delegate. There is no point stopping at the first acceptable plumber when the delegate can examine all of them in a second. The delegate can, in principle, maximize.

THE BLADE THAT DID NOT MOVE

But notice what did not change. You did not change. Your attention is the same size it was before. Your evening is the same length. Your ability to hold eleven options in your head and compare them fairly is exactly what it was when there were only two windows on the street.

This is the fact that most early AI products got wrong. They gave the delegate infinite search and then handed the results back to the human. Here are twelve restaurants. Here are five flights. Here are three versions of your email. Here are the pros and cons of each.

That design takes the one person in the system with a bounded mind and floods it. It turns every user into Schwartz's maximizer, faced with more attractive rejected options than ever, at machine speed. It is Friday night with two hundred tabs.

> *The delegate's search became free. Your attention did not. Handing the delegate's results back to you only moves the flood.*

TWO SEARCHES, NOT ONE

The way out begins with seeing that there are now two searches where there used to be one. There is the search through the world, which the delegate can do completely. And there is the search through your own mind for what you actually want, which nobody can do for you in full, but which a good delegate can learn a great deal about.

The first search is about options. The second is about the bar. For all of history they were done by the same tired person at the same time. They can now be pulled apart.

THE CHAPTER IN ONE LINE

AI makes search through the world free for the delegate. It does nothing for your attention, so the results must not simply be handed back to you.

CHAPTER 7

The Old Problem of When to Stop

Mathematicians have a famous puzzle about stopping, and it shows exactly what AI changes and what it does not.

It is usually called the secretary problem, and around 1960 the writer Martin Gardner put it in front of a large audience. You are hiring for one position. Candidates arrive one at a time in random order. After each interview you must decide on the spot: hire this person, or send them away forever. You cannot go back. How do you give yourself the best chance of hiring the very best candidate?

The answer is beautiful. Look at roughly the first third of the candidates, a little over thirty-seven percent, and hire none of them. Use them only to learn what good looks like. Then hire the first person who is better than everyone you have seen so far. No strategy does better.

Look at what the puzzle is really about. It is not about intelligence. It is about two hard constraints: options arrive one at a time, and rejection is permanent. The whole art of the solution is managing the cost of saying no too early or yes too soon.

WHAT CHANGES

Many of life's real searches have the same shape. Apartments, jobs, contractors, even partners arrive one at a time, and the good ones do not wait. A great deal of human anxiety is the secretary problem felt from the inside: is this one good enough, or should I wait for something better that may never come?

An AI delegate does not dissolve that shape entirely. The apartment still rents to someone else. But it changes it drastically. A delegate can see many candidates in parallel instead of in sequence. It can keep watching a market continuously, so the good options are noticed the moment they appear. It can hold an offer open by acting quickly. In many searches, the one-at-a-time constraint simply goes away, and with it most of the mathematics of when to stop looking.

WHAT REMAINS

What does not go away is the other half of the question. Even when the delegate can see every candidate at once, someone still has to decide what counts as good. Is the apartment with the longer commute and the better light better or worse than the one with the short commute and the dark kitchen? The secretary problem quietly assumes you can rank every candidate against every other. Real life never gives you that. It gives you trade-offs, and only you know your exchange rates.

So the stopping problem changes its nature. It used to be: when should I stop looking? It becomes: when should I stop wanting better? The first was a problem of time. The second is a problem of the bar.

THE CHAPTER IN ONE LINE

AI removes much of the one-at-a-time constraint that made stopping hard. It leaves the question of what counts as good, which only you can answer.

PART THREE

The Bar Is the Job

When the delegate can find the best of everything, the one thing left for you is to say what is enough.

CHAPTER 8

When Search Is Free, the Bar Is the Decision

Here is the turn the whole book has been walking toward.

For all of human history, choosing had three parts, and the same person did all three. You searched the world for options. You compared the options you found. And you held a bar that told you when to stop. The three were tangled together so tightly that we never thought of them as separate. Searching was tiring, so the bar dropped. Comparing was painful, so you stopped searching. The bar was set by how much search you could stand.

AI pulls the three apart. The delegate can search the world completely. The delegate can compare everything it finds, against everything else, on every dimension it knows you care about, without fatigue and without regret. That leaves exactly one part of the decision that cannot be handed off whole: the bar. What counts as enough, for you, here, now.

> *The delegate can find the best of everything. Only you can say what is enough.*

When search is free, the bar is not one input among many. It is the decision. Everything else is execution.

FIGURE 3

Who holds each part of a choice

Part of the choice	Before	With a delegate
Search the world for options	You, until you are tired	The delegate, completely
Compare what you found	You, and you remember the runners-up	The delegate, on file, out of your evening
Hold the bar	You, set mostly by fatigue	You, and it is now the whole decision

Free search pulls apart three things one tired person used to do at once. Two of them move to the delegate. One of them cannot.

TWO BARS, NOT ONE

It helps to see that a bar has two edges.

The first is the floor: what is acceptable at all. The plumber must be licensed. The restaurant must have something you can eat. The flight must arrive before the meeting. Below the floor, nothing counts, however cheap or convenient.

The second is the stop: the point past which better is not worth more of your life. A restaurant that is slightly better but forty minutes further away. A flight that is twenty dollars cheaper with a connection. A contract clause that could be negotiated a little more in your favor at the cost of another week. The stop is where you tell the delegate that more is not better anymore.

People are usually clear about their floors and vague about their stops. That vagueness was hidden for most of history, because exhaustion set the stop for you. You stopped when you were tired. A delegate is never tired, so the stop has to come from somewhere else. If it does not come from you, it will come from whoever designed the delegate, or from

whoever is paying to be found by it.

HOW A DELEGATE LEARNS YOUR BAR

The good news is that you do not have to write your bar down in a form. A good delegate learns most of it the way your own bar was learned: from history. What you accepted. What you sent back. What you returned, complained about, reordered, recommended. Which restaurant you went back to without being asked. Which suggestions you ignored.

From that record a delegate can build a working model of your floors and stops in each part of your life, and it can keep the model current, because your bar moves as your experience moves. The model will be wrong at the edges, and the edges are where the delegate should ask. Not often. Not about everything. Only where the stakes are high and the bar is unknown.

That is a very different relationship from the one most people have with software today. The question is no longer which of these do you want. The question, asked rarely, is what would be enough.

MY OWN BAR

I live a version of this every day. I have told my AI, in plain words, that I trust it, and that it should bring to me only the matters worthy of my attention. Everything else, it decides and handles. I do not want a menu. I do not want three options and a table of pros and cons. When it brings me something, it should be a real decision that only I can make, or an action that cannot be undone.

That instruction is a bar. It is a bar on my attention rather than on restaurants, but it is the same kind of thing. It tells the delegate where good enough is, so that it can do all the searching and comparing on its side and spend almost none of mine. I have written about the habit behind it in My AI Delegate and You Are Not the User. This book is the theory under the habit.

THE CHAPTER IN ONE LINE

When search and comparison are free, the bar is the whole decision. Know your floors, and be honest about your stops.

CHAPTER 9

Never Show Me the Runner-Up

Go back to the graduating students in Chapter 4. The maximizers did better and felt worse, and the reason was the attractive options they had seen and rejected. Their regret was made of runners-up.

Now imagine a fourth kind of job seeker that did not exist in 2006. Call it the delegated maximizer. Their delegate looks at every opening, reads every review of every employer, models every commute and every salary path, and finds the best match against their bar. Then it brings them one offer, with a short note on why. It does not bring them the eleven near-misses. It keeps those on file.

This person gets the maximizer's result and the satisficer's peace. They have the better job, and they have no ghosts, because they never met the alternatives.

> *Regret needs a visible alternative. The delegate can do the comparing so that you never have to meet the runner-up.*

This is the single most important design principle for AI in daily life, and almost nobody building it says it out loud: show the choice, not the choices.

HIDING IS NOT LYING

There is an obvious objection. Is this not just a machine keeping things from you? It would be, if the delegate pretended there were no alternatives, or would not show them when asked. That would be a delegate you could not trust, and it would deserve to be fired.

The honest version is different. The delegate keeps every serious option it considered and the reason it passed each one over, on file. You can open the file whenever you want. You can ask why not the other one, and get a straight answer. I wrote about this record in The No File. What the delegate does not do is put the file in front of you uninvited. The runner-ups exist. They are not in your evening.

This is how we already treat the human professionals we trust most. A good doctor considers many diagnoses and tells you the one she believes. A good architect draws a dozen plans and shows you the one that works. You would not want either of them to hand you the whole pile and say you decide. Their value is that they carried the comparison so you did not have to.

WHEN TO SHOW TWO

There are times when the delegate should bring more than one option. When the decision cannot be undone. When the trade-off touches something only you can weigh, like a child, a health choice, or a person. When your bar for this kind of thing is genuinely unknown. In those cases the delegate should bring the smallest set that represents the real choice, usually two, framed as the decision it is.

What it should never do is bring you twelve because it found twelve. Finding is its job. Choosing what is enough is yours. Spreading the finding back across your attention is a failure of the first job dressed up as respect for the second.

THE CHAPTER IN ONE LINE

Show the choice, not the choices. Keep the runner-ups on file, open on request, and out of your evening.

CHAPTER 10

The Treadmill and the Ratio

There is a deeper trap in a world of tireless search, and to see it clearly I need the framework I have spent years developing, the Reality Equation.

In its simplest form it says that reality, as it is had by someone, is a ratio. R equals A over E. A is the Actual: what the world hands you, beyond anyone's power to change once it has arrived. E is Expectation: what your unconscious brings to meet it, already there before the Actual arrives. Surprise, the felt jolt of the moment, is the logarithm of that ratio. When the Actual exceeds the Expectation, you feel the lift I call the AAH. When it falls short, you feel the drop I call the OO. When they match, the ratio is one, and there is no jolt at all.

Expectation has two parts in my framework, Prediction and the Ideal. This chapter is about the first one, Prediction: the part that habit and history train.

NO ONE TOUCHES THE DENOMINATOR

The rule in the framework that matters most here is that no one can reach in and change either term directly. Not you, not your delegate, not any company. Your Expectation is not set by your conscious choices. It is inherited from your history. The only road to it runs through the Actual: something acts, the action leaves a trace in the world, the trace becomes part of what has happened to you, and what has happened to you slowly retrains what you predict.

Now put a maximizing delegate into that loop. Every week it finds you the best restaurant, the best price, the best plumber, the best flight. It does not touch your Expectation. It cannot. What it changes is which Actuals arrive. And Actuals are exactly what Prediction learns from. So,

week by week, the best becomes the predicted. The ratio drifts toward one. The lift goes quiet.

> *A delegate that always delivers the best does not make life feel better. It makes the best feel normal.*

This is the same thing Simon saw when he said the aspiration level rises when the bar is easily met, and what psychologists later called the hedonic treadmill. The Reality Equation says why it happens and why it cannot be talked around. You cannot decide to be delighted by the fourth perfect dinner in a row. The angle and the weight of your Expectation were set before the plate arrived. Choosing how to see it comes too late.

THE COMPARISON BORROWS THE SURPRISE

There is a second effect, and it explains the flatness on Friday night. When you open eleven tabs and study the photographs, those photographs are Actuals too. They arrive before the dinner does, and they train your Prediction for the dinner in advance. By the time the real plate is in front of you, much of its surprise has already been spent on the tabs.

In an earlier piece, The Conservation of Surprise, I argued that the total surprise between a first expectation and a final actual is fixed, and that hype borrows surprise from the event, at interest. Comparison shopping is hype you manufacture for yourself. Every runner-up you study pre-spends the evening. The delegate that keeps the comparison on its side of the line is not only sparing your attention. It is leaving the surprise where it belongs, in the moment itself.

HIGH IS NOT GOOD

I want to be careful here, because this is where people misread the equation. A high ratio is not the goal. A life tuned for constant lift is as

unwell as a life of constant drops. A ratio near one is stability, and stability is not suffering. A delegate that makes the best ordinary has not harmed you. It has made a part of your life reliable, and reliable is worth a great deal.

The point is narrower. A tireless searcher, left alone, will turn more and more of life into reliable sameness at the top, and you will not feel it as the top. So the Manager of a delegate has a real choice to make, not about how to feel, which no one controls, but about where to act. In which parts of life do you want the delegate to maximize? And in which do you want to keep some looking for yourself, some wandering, some walking down the street and going in where the kitchen smells good? Those are actions. They change what arrives. And what arrives is the only thing that ever changes what you expect.

THE CHAPTER IN ONE LINE

No one can set your Expectation, not even your delegate. It changes what arrives, and what arrives retrains what you predict. Decide where to let it maximize.

PART FOUR

Enough, Everywhere

What the end of costly search does to markets, to companies, and to the people in your life.

CHAPTER 11

The Rents on Stopping Early

In Chapter 5 we saw how much of commerce collects a rent on the fact that people stop looking early. The brand as a stopping rule. The default nobody changes. The top slot, the eye-level shelf, the contractor who calls back first, the policy that rolls over because reconsidering costs.

Now give every buyer a delegate that never stops early.

The rents do not shrink gradually. They collapse wherever delegates do the buying, because the thing they were charging for has disappeared. The delegate does not care which plumber answered first. It cares which plumber has the best record of work against your bar. It does not care what is at eye level. It reads every shelf. It does not keep the default because it is the default. It reads the default and every alternative, every year, at no cost.

WHAT REPLACES POSITION

If position stops paying, what starts? The answer follows directly from the bar. A seller now has to be the thing that clears a delegate's model of a particular buyer's floors and stops, and does so reliably enough that the delegate stops looking.

That rewards a few things that position used to let sellers skip. A record that a machine can read and believe. Promises that are kept and can be checked, which I have called the warranty layer. A real human face that answers for the work when something goes wrong, which is the subject of my book Face Value. And fit: being exactly right for a narrow set of people rather than acceptable to everyone walking past.

THE WAR FOR THE BAR

Sellers will not give up their rents quietly. When they can no longer capture the search, they will try to capture the bar.

You can already see the shape of it. Marketing used to aim at the moment of choice: the shelf, the ad beside the search result, the jingle in your head in the aisle. In a delegated world the moment of choice belongs to a machine that reads everything. So the target moves upstream, to the definition of enough. Persuade the buyer that no car without this feature is acceptable. Persuade them that good enough for a mattress means this certification. Shape the floor, and every delegate will do your selling for you.

> *When the delegate does the searching, advertising stops aiming at the shelf and starts aiming at the bar.*

This is why it matters so much whose delegate it is. A delegate built by the seller, or paid by the seller, will quietly carry the seller's bar. A delegate that is truly yours, centered on you and answerable to you, will carry yours, learned from your own history. It is the strongest argument I know for personal AI over enterprise AI, and for every business to let people bring their own. I made the case in Bring Your Own AI. The war for the bar is where that case gets decided.

THE CHAPTER IN ONE LINE

Rents on stopping early collapse when delegates do the buying. Sellers will fight for the bar instead, so the bar must live in a delegate that is yours.

CHAPTER 12

The Firm That Never Stops Searching

Cyert and March described companies that search only when something hurts. Results fall below the bar, and the firm goes looking for a fix, starting close to home. When results clear the bar, it carries on. Between searches, slack builds up: spare people, comfortable inventories, suppliers kept on out of habit, margins nobody bothered to squeeze.

AI gives companies the same thing it gives individuals: a searcher that never tires. And the temptation is obvious. Why search only when it hurts? Why not search all the time? Re-bid every supplier every week. Re-price every product every hour. Re-route every shipment, re-staff every shift, re-check every contract, continuously, forever.

Some of that will be pure gain. A great deal of what companies do out of habit is waste, and a tireless searcher will find it.

SLACK WAS A SHOCK ABSORBER

But slack was never only waste. It was the margin that let an organization absorb a surprise without breaking: the spare supplier who could take an urgent order, the extra hands on a bad day, the long relationship with a vendor who would answer the phone at midnight because you had been a good customer for ten years. A firm that searches continuously will strip that margin out, one sensible optimization at a time. It will look brilliant on an ordinary day and fragile on a bad one. I argued in Slack Is a Forecast that slack is a bet about the future. Continuous search quietly bets it all away.

There is a control problem too. In The Shower Knob I described how a loop starts to swing when the hand reacting to it becomes faster and stronger while the pipe stays the same length. A company that re-optimizes everything continuously, against suppliers and customers who are doing the same, is a building full of fast hands on long pipes.

MANAGEMENT BECOMES BAR-SETTING

So the job of running an organization changes. The searching can be delegated. The bar cannot. What counts as a good enough supplier, so that we stop re-bidding them and build a relationship instead? What margin is enough, so that we stop squeezing and leave room for the shock? Which processes are good enough to leave alone for a year?

In a world of free search, the best-run firms will be the ones that write down their stops on purpose. Not because they cannot find better, but because they have decided that better is not worth the brittleness. Satisficing, which began as a description of limited firms, becomes a strategy for unlimited ones. It is a close cousin of the argument in my book Work Expands: when help is free, the scarce thing is the judgment about where to stop.

THE CHAPTER IN ONE LINE

A firm that searches forever strips out the slack that kept it alive. In a world of free search, management is the deliberate setting of stops.

CHAPTER 13

Enough Is a Person

There is one domain where the logic of the tireless searcher becomes not only unwise but wrong, and it is the most important domain of all.

You do not maximize people.

Think about what it would mean. A delegate that keeps quietly scanning for a better friend than the friends you have. A better business partner than your partner. A better spouse than your spouse. A better employee than the one who has worked beside you for twelve years. Each scan could even be accurate. Somewhere among eight billion people there is, by some measure, a better match for almost anything. And to go looking would be a kind of betrayal.

We already have a preview. Dating apps took the search for a partner, which for most of history was short, local and strongly satisficing, and made it nearly endless. Many people found partners that way. Many others found the maximizer's misery: the next profile always one swipe away, every relationship shadowed by the runner-ups. The search got cheaper, and the commitment got harder, because commitment is what satisficing is called when the option is a person.

> *Commitment is what satisficing is called when the option is a person.*

RANK AND ORDER

I teach my students that delegation runs down, never up, and that rank and order matter. You delegate the small matters so you can keep what is worthy of your attention. And some things are worthy of your

attention by their nature: your customers, your friends, your family. Never have your AI deal with a person who deserves you. Handing that person down to a machine dishonors them.

The same rule governs the bar. When a person clears your bar, the search closes. Not because you could not find someone better on paper, but because the closing is the point. Loyalty is a bar that has been cleared and then deliberately retired. A friend is someone you have stopped comparing.

So the right instruction to a delegate here is not show me better people. It is the opposite. Protect the people I have chosen. Remember their birthdays, yes, but tell me, so that I am the one who calls. Keep their news in front of me. Never put a runner-up beside them. In the age of the tireless searcher, the most human act left may be to stop searching on purpose, and mean it.

THE CHAPTER IN ONE LINE

Never let the delegate maximize a person. When someone clears your bar, the search is over, and that ending is what loyalty is.

PART FIVE

The Coming Decade

Five predictions about life after the cost of looking, and how each of them could be wrong.

CHAPTER 14

What I Expect

Here is what I expect over roughly the next ten years, stated plainly enough that you will be able to tell whether I was right. Each prediction comes with what is already true, what has to happen, and where I am probably wrong.

PREDICTION 1

Assistants will stop handing you lists

The best AI products will answer most everyday requests with one choice and a short reason, not a list of options. Show alternatives will survive as a quiet link, not the main event. Products that keep answering with menus will feel exhausting next to the ones that simply choose, and people will leave them.

ALREADY TRUE	WHAT HAS TO HAPPEN	WHERE I AM PROBABLY WRONG
Many assistants already act on simple requests without asking, and the products people describe as calm are the ones that make a choice and own it. My own work with my AI already runs this way.	Delegates have to become trusted enough, and accountable enough, that people are comfortable letting them choose. That depends on the honest record of what was passed over, open on request.	People may turn out to want the menu, not for information but for the feeling of control. If choosing among options is itself a pleasure people refuse to give up, lists will survive and I will have misjudged what people want from their evenings.

PREDICTION 2

Every personal AI will carry your bar

Personal AI will keep a working model of your floors and stops for each part of your life, learned from your history, visible to you and editable in plain words. Enough will become a setting, like a privacy setting today, and the delegate will ask about it only at the edges where stakes are high and your bar is unknown.

ALREADY TRUE	WHAT HAS TO HAPPEN	WHERE I AM PROBABLY WRONG
Delegates already learn preferences from what people accept and reject, and people already write standing instructions for their AI. Mine says to bring me only what is worthy of my attention.	The model of the bar has to become portable, so it travels with the person when they move from one AI to another, rather than staying locked inside one company.	The bar may prove too fluid and too contradictory to model well. If people's sense of enough changes faster than any delegate can learn it, the delegate will keep guessing wrong, and people will go back to choosing for themselves.

PREDICTION 3

First position will stop paying

The value of default placement, the top search slot, the eye-level shelf and being first to call back will fall sharply wherever delegates do the buying. Brands will shift from being stopping rules to being promises: a record a machine can read, a warranty that is kept, and a face that answers for the work.

ALREADY TRUE	WHAT HAS TO HAPPEN	WHERE I AM PROBABLY WRONG
Search and shelf position have always been valuable because human searchers stop early. Delegates that read every option already exist, and agents are already being sent to compare prices and book services.	Delegates have to handle a large enough share of everyday buying that sellers notice the change in what wins.	Sellers may succeed in owning the delegates themselves. If most people use a delegate that a seller or platform controls, first position will simply move inside the delegate, and the rent will survive under a new name.

PREDICTION 4

The fight will move to the bar

As position loses value, marketing will shift its effort from the moment of choice to the definition of enough: persuading people what no acceptable product can lack. Regulators and consumer advocates will begin treating the shaping of a delegate's bar by a seller as a new kind of conflict of interest.

ALREADY TRUE

Marketing already shapes standards of acceptability through certifications, feature checklists and categories. The tools to target a delegate's preferences directly are arriving with the delegates.

WHAT HAS TO HAPPEN

Delegates have to be influential enough that sellers find it cheaper to shape the bar than to compete on the shelf.

WHERE I AM PROBABLY WRONG

People's bars may prove stubborn, rooted in their own history rather than in persuasion. If a delegate's bar truly comes from what the person has accepted and returned, sellers may find the bar much harder to move than the shelf ever was.

PREDICTION 5

Good enough will return as a virtue

Within a decade, well-run households and companies will talk about deliberate stops the way they now talk about budgets: what is enough, and where we refuse to keep searching. Keeping a supplier, a doctor or a friend without re-comparing will be understood as a strength. A few domains will be kept for searching by hand, for the pleasure of it.

ALREADY TRUE

The research on maximizers and regret is well known, and many people already sense that endless comparison is making them miserable. Companies already fear the brittleness that comes from optimizing every margin away.

WHAT HAS TO HAPPEN

Free search has to become ordinary enough that its costs are felt in daily life, not just predicted in books like this one.

WHERE I AM PROBABLY WRONG

The pull of better may simply win. If tireless search keeps delivering visible gains, people and firms may never choose to stop, and good enough will remain what it has been for decades, a phrase said with a shrug.

WHERE I AM PROBABLY WRONG OVERALL

The deepest way this book could fail is if delegates never become truly ours. Every prediction above assumes a delegate that answers to the person it works for and carries that person's bar. If the delegates people actually use are built to serve the people who sell to them, then free search will not free anyone. It will simply move the shelf inside the machine, and the bar will be written by someone else. That is not a technical question. It is the question of whose AI it is.

EPILOGUE

Friday, Again

> **FRIDAY, 6:58 P.M.**
>
> *You are putting on your shoes when the card appears. A table for two at seven thirty at the place around the corner you have never tried, and a short line: the kitchen is good, the noise is low, and they have the dish you have been ordering everywhere this month. You tap Try it.*

There are no tabs. You did not see the other forty restaurants your delegate considered, or the three it nearly chose, or the reasons it passed over each. They are on file if you ever want them. You will not.

The dinner is very good. You will remember it, partly because nothing pre-spent it. You walked in not knowing exactly what the room looked like, and the room got to surprise you.

Herbert Simon named satisficing in an age when looking was expensive, and for seventy years the word sounded like an apology. Good enough meant I ran out of time. It meant I could not afford better.

That meaning is ending. Looking is about to become free, not for us, but for the delegates that work for us. When it does, good enough stops being a confession about our limits and becomes the most important thing we say. It is the line between what the machine handles and what is ours. It is where the searching stops and the living starts.

The delegate can look at everything. It can find the best of anything you name. It cannot tell you when to stop wanting more. That was always yours. Now it is the whole decision.

Say it clearly. Enough.

BACKGROUND

Herbert A. Simon, Administrative Behavior (Macmillan, 1947).

Herbert A. Simon, "A Behavioral Model of Rational Choice," Quarterly Journal of Economics 69 (1955).

Herbert A. Simon, "Rational Choice and the Structure of the Environment," Psychological Review 63 (1956).

Herbert A. Simon, "Invariants of Human Behavior," Annual Review of Psychology 41 (1990).

Herbert A. Simon, Models of My Life (Basic Books, 1991).

Allen Newell and Herbert A. Simon, "The Logic Theory Machine," IRE Transactions on Information Theory (1956).

Richard M. Cyert and James G. March, A Behavioral Theory of the Firm (Prentice-Hall, 1963).

George J. Stigler, "The Economics of Information," Journal of Political Economy 69 (1961).

Gerd Gigerenzer, Peter M. Todd and the ABC Research Group, Simple Heuristics That Make Us Smart (Oxford University Press, 1999).

Barry Schwartz, The Paradox of Choice: Why More Is Less (Ecco, 2004).

Barry Schwartz et al., "Maximizing Versus Satisficing: Happiness Is a Matter of Choice," Journal of Personality and Social Psychology 83 (2002).

Sheena S. Iyengar, Rachael E. Wells and Barry Schwartz, "Doing Better but Feeling Worse: Looking for the 'Best' Job Undermines Satisfaction," Psychological Science 17 (2006). https://works.swarthmore.edu/fac-psychology/104/

Philip Brickman and Donald T. Campbell, "Hedonic Relativism and Planning the Good Society," in Adaptation-Level Theory (1971).

Martin Gardner, "Mathematical Games," Scientific American (February 1960).

Satisficing, overview https://en.wikipedia.org/wiki/Satisficing

John Rector, The Reality Equation.
https://johnrector.me/2025/08/18/the-reality-equation-6/

John Rector, The Conservation of Surprise.
https://johnrector.me/2026/09/04/conservation-of-surprise/

John Rector, My AI Delegate. https://johnrector.me/2026/10/30/ai-delegate/

John Rector, You Are Not the User. https://johnrector.me/2026/10/30/not-the-user/

John Rector, The No File. https://johnrector.me/2026/11/05/the-no-file-app-to-stack/

John Rector, The Warranty Layer.
https://johnrector.me/2026/11/06/the-warranty-layer-app-to-stack/

John Rector, Bring Your Own AI: The Policy Every Business Needs.
https://johnrector.me/2026/10/30/bring-your-own-ai-the-policy/

John Rector, Slack Is a Forecast.
https://johnrector.me/2026/10/21/slack-is-a-forecast/

John Rector, The Shower Knob. https://johnrector.me/2026/11/06/the-shower-knob/

John Rector, Try It Is the Business Model.
https://johnrector.me/2026/10/25/try-it-business-model/

John Rector, The Expensive Hour.
https://johnrector.me/2026/10/04/the-expensive-hour/

John Rector, The Last Line Is You. https://johnrector.me/2026/09/27/last-line-is-you/

Written and published November 8, 2026, at johnrector.me.