

ECSTATIC ECONOMICS · THE LUMINOUS BUSINESS METHODOLOGY

The Customer Already Told You

Finding the Gain in What Customers Already Say

Ammanuel Santa Anna

Santa Anna, A. (2026). The Customer Already Told You: Finding the Gain in What Customers Already Say. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989055
Luminous Library Licence 1.0 · business.hautelumiere.life/library/the-customer-already-told-you

THE LICENCE

The Customer Already Told You
Finding the Gain in What Customers Already Say

Ecstatic Economics · The Luminous Business Methodology. The Haute Lumière edition, 2026. © 2026 Ammanuel Santa Anna.

Released under the Luminous Library Licence 1.0 (hautelumiere.life/library-licence): free to read, share, quote, translate and teach, non-commercially, for ever. The Author keeps every other right, including every commercial right.

Published by Luminous Prosperity Inc. for Haute Lumière. Author ORCID 0009-0008-3244-0306.

DOI 10.5281/zenodo.23989055 · this record's page: business.hautelumiere.life/library/the-customer-already-told-you/

Please cite as: Santa Anna, A. (2026). The Customer Already Told You: Finding the Gain in What Customers Already Say. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989055

A NOTE BEFORE YOU READ

This work is offered for education and leadership development. It is not legal, financial, accounting, investment, medical or psychological advice; take the advice of a qualified professional who knows your circumstances. Companies and events are described from the public record as the author understands it; characterisations are the author's analysis. Examples are illustrative, and no particular outcome is promised to any reader.

ABSTRACT

Finding the gain in what customers already say: the second call, the closed ticket, the silent leaver, the repair and the measure. Much of what arrives at a firm's front door is its own earlier failure, coming back with the customer's account of where it happened. Finding the gain in what customers already say — the second call, the closed ticket, the silent leaver, the repair and the measure; seven chapters with a field manual.

CONTENTS

CHAPTER ONE The Second Call 4

CHAPTER TWO The Ticket Marked Resolved 16

CHAPTER THREE The Customer Who Left Quietly 27

CHAPTER FOUR The Apology and the Cause 38

CHAPTER FIVE The Calls That Stop Coming 50

CHAPTER SIX The First Ninety Days 61

CHAPTER SEVEN The Field Manual 75

NOTES References 90

CHAPTER ONE

The Second Call

Much of what arrives at a firm's front door is its own earlier failure, coming back with the customer's account of where it happened.



The Second Call

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER ONE

Ten past eight, the spare headset

It is ten past eight on a Monday in November, and the head of customer services at a housing association is sitting at the end of a row of desks with a spare headset on. She has not done this for two years. The screen above the row shows the figure the board likes best: calls answered within twenty seconds, ninety-something, green. Beside her, an adviser who has worked the repairs line since before the new phone system takes the fourth call of the morning, and the woman on the other end begins, as the second and third callers began, with the same four words. *I'm ringing again about.*

She is ringing again about her kitchen ceiling. Water came through it eleven days ago. A repair was booked, for a Thursday, between eight and one. Nobody came. She rang, and it was booked again, for the following Tuesday, between eight and one, and on Tuesday a plumber arrived who said the leak was from the flat above and that he could not get into the flat above. She is calling to ask whether anyone has told the flat above. The adviser is kind and quick. He books a third appointment and raises a note for the plumber to contact the neighbour. The call takes six minutes and is logged, correctly by every rule the system has, as one repair request handled.

The head of customer services is a composite, drawn from the way these mornings are described in the research below and in the house's own practice, and no firm is meant. What she feels at the end of the fourth call is not in any report. It is a slow tightening under the breastbone, the feeling of a person who has just learned that the number she presents every month is true and beside the point. The line is answering quickly. It is answering, quickly, the same woman three times.

Ammanuel Santa Anna's canon asks the question her body has just answered. In its chapter on walking a value stream, it suggests the walker notice what happens to her own body at each step of the work, and it puts one question first: "What happens in your chest when you enter the customer service department?"[1] The canon's own answer is a team under pressure, holding itself carefully. The answer this morning is more specific. The department is full of calls the organisation made for itself.

This chapter makes one claim. A large part of the work arriving at a firm's front door is the firm's own earlier failure, coming back, and every one of those calls carries the customer's exact account of where the failure happened. The organisation that counts them as demand, and answers them well, has paid twice for information it then throws away. The second call is the first place a firm stops hearing what its customer already told it.

Two kinds of demand

The distinction she has just heard was named by the British occupational psychologist and consultant John Seddon, and it is simple enough to sound obvious once it is said. In *Freedom from Command and Control* (first published in 2003) and then in *Systems Thinking in the Public Sector* (2008), he divided what arrives at a service organisation into two kinds.[2] Value demand is what the service exists to meet: a tenant reporting a leak for the first time, a customer opening an account, a resident asking how to apply for a parking permit. Failure demand is demand caused by the organisation failing to do something, or failing to do it right, for the customer: the second call about the leak, the call to ask why the letter said one thing and the website another, the call to chase a promise.

Seddon's account of where the idea came from is usually retold as a story about banks in the 1980s, when telephone work moved out of branches into call centres and the volume of calls rose faster than anyone had planned for. His explanation was that much of the new volume was customers ringing back because

their first contact had not resolved anything. The story is told here as it is told in secondary sources, and its details were not checked against his own text. What matters is the shape of the claim. The calls were not a fact about the customers. They were a fact about the system.

That shape is the whole of the idea's usefulness, and it is why it is worth defending against its own popularity. Failure demand is not a measure of staff performance. The adviser on the repairs line did nothing wrong; he did everything right, three times. The failure happened in the way the repair was scheduled, in the plumber's van, in the absence of any arrangement for reaching a neighbour, and it surfaced at the one desk in the building that had no power to fix any of it. A manager who reads failure demand as a figure about the people who answer the phone has read it upside down.

Santa Anna's canon describes the people at that desk with some care, in its chapter on the members of an organisation whose loss would most change it. It asks, as one step in finding them, "Who absorbs the system's emotional energy?" and its first answer is the person who deals with the frustrated customer. That is the repairs line. It absorbs, every day, the anger the system produces upstream, and in most organisations it is measured on how quickly it does so.

There is an honest limit to say here. Seddon's evidence comes largely from his own consulting interventions, written up as cases. He and others have reported very high proportions of failure demand in some public-sector call centres, and those figures circulate widely. They are reports from inside particular buildings, gathered by the people who were also doing the work of changing them, not independent measurements across many organisations. This book prints none of them, and for a plainer reason than doubt: a figure from somebody else's building says nothing reliable about yours. The only number that would tell the head of customer services anything is the one she could gather herself, with a spare headset, over a fortnight.



A target with a number on it

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER ONE

A target with a number on it

The idea travelled further than most management ideas do. In 2008 the British government included, among the one hundred and ninety-eight indicators in its new national framework for local authorities, one called National Indicator 14, *avoidable contact*: the proportion of a council's contact with the public that was of low or no value to the person making it. Councils began collecting it that autumn. In 2010 it was removed from the national set, and some councils kept it as a local measure because they had found it

useful.[3] The history is recorded mainly in the councils' own committee papers, which were read for this book; the central government guidance was not.

The episode is worth a paragraph because of what Seddon is reported to have thought of it. The concept was his, and he objected to its use as a national target. The objection, as it is usually summarised, was that a council told to report less avoidable contact would find ways to report less of it, by changing how calls were classified rather than by changing the work that caused them. The summary comes from secondary accounts and was not checked against his own words. The logic does not depend on who said it. A number that a council is judged on is a number the council will manage, and the easiest part of failure demand to manage is the label.

Santa Anna's canon draws the same line in its history of how programmes have been assessed. Counting outputs, it says, asks whether the organisation did what it said it would do, and that question matters. But as a whole approach to assessment, output counting "confuses activity with achievement." A repairs line that answers ninety-something per cent of calls inside twenty seconds has achieved something real. It has not achieved a mended ceiling.

The difference between using failure demand to study the work and using it as a target is the difference between a doctor's thermometer and a thermostat on the wall. The thermometer tells you something about the body that you then go and find out more about. The thermostat is wired to make the reading come out right. Santa Anna's book on doing business as a living system uses the same image about the way most firms listen to their customers. Surveys, focus groups and dashboards, it says, "are the auditory equivalent of a thermostat — they tell you if you're too hot or too cold, but they can't hear the music."

Where the call was made

The fourth call of the morning was made by a booking system that offers every tenant the same five-hour slot, and by a schedule that sends a plumber to a leak without anyone first asking where the water is coming from. It was made by a van that carries no key to the flat above and a job sheet that has no field for *neighbour to be contacted*. None of these is anyone's fault in the ordinary sense. Each was a reasonable decision by somebody, somewhere upstream, who never heard the call it caused.

That is the mechanism, and it is older than call centres. In the early 1950s, on the floor of a Toyota plant in Japan, the engineer Taiichi Ohno began building what became the Toyota Production System, written up in English in 1988.[4] Among its best-known contributions is a list of the forms waste takes in production, and one of them is defects: work that has to be done again because it was not done right. Ohno's practical insistence, as the canon tells it, was "follow the product". Trace the work from its beginning to its end and ask, at every step, whether that step creates value for the customer or creates waste. Seddon's contribution, in effect, was to notice that in a service organisation the defects are not stacked on a shelf at the end of the line. They walk back in through the front door and ask for the manager.

His own route into the work, and the one the head of customer services is now on, is to begin there, at the front door, and follow each call of failure demand back to the place it was made. Seddon insists that this is done by managers themselves, listening to real calls and reading real cases, rather than delegated to analysts who report on them. The reason is not that analysts are less able. It is that a manager who has heard the woman with the ceiling cannot afterwards read the green figure on the screen in the same way, and a manager who has only read a report about her can.

There is one careful evaluation of the approach in a peer-reviewed journal, and it is worth reading for its tone as well as its finding. In 2008 Michael Jackson, Nick Johnston and Seddon himself described pilots in three British social-housing organisations, sponsored by the government office then responsible for housing, which used the approach in different service areas.[5] The results, they reported, indicated gains in efficiency and in service performance, in tenants' satisfaction and in the morale of staff, "although securing and sustaining these results requires considerable organizational commitment." Seddon is an author of the paper, so it is not independent evidence; it is a careful account by people with an interest in the result. The qualification at the end of the sentence is the most useful part of it.

Santa Anna's canon has a case of its own that shows how the information is lost, and it is set in a hospital rather than a housing association. A regional healthcare network, the canon's case runs, could not understand why patients' satisfaction in its emergency departments kept falling. When the people closest to the patients were finally asked, the causes turned out to be relational and environmental, and the staff had known them for years. What had happened to their knowledge on its way up is told in one sentence:

> *These were problems that the frontline staff had been aware of for years — and had repeatedly raised through formal channels — but the intelligence had been filtered and reframed as it traveled up the hierarchy until it became unrecognizable: "Staff recommends additional customer service training," the executive summary read.*

The repairs line has its own version of that summary. It is the monthly report that says contact volumes are high and recommends more advisers. The advisers would recommend something else. They would recommend that somebody ring the flat above.

The weight on the line

It is worth staying a moment with the adviser, because the cost of failure demand is not only the six minutes of his time. He has taken a version of the fourth call perhaps forty times this month. He knows the five-hour slot, the plumber who cannot get into the flat above, the letter that tells tenants to expect a call that never comes. He apologises for all of them, sincerely, many times a day, for decisions he had no part in and cannot change. The canon describes what this does to a team over time, in its account of how an organisation and the people it serves shape each other: "A department that is structurally coupled with angry

customers will develop different capacities (and different pathologies) than one coupled with grateful ones."

The pathologies are recognisable to anyone who has run such a desk. People learn a soothing script and a flat voice. They learn which promises not to make. The best of them learn workarounds, a direct number for a particular supervisor, a way of flagging a job so it is seen, and the organisation comes to depend on those workarounds without knowing that they exist. The most experienced advisers often leave, because absorbing the same anger forty times a month for decisions they cannot affect is wearing in a way that has nothing to do with the number of calls.

Read the adviser's position as information, though, and it changes. He is the best-informed person in the building about where the repairs service fails. He hears every failure, in the customer's own words, within hours of its happening, and he hears them all in one place. Santa Anna's canon has a sentence about organisations that cannot use what their members know, and it fits the repairs line exactly. Where truths cannot travel, it says, "Over time, this produces organizations that are literally stupider than their members." The adviser knows why the calls come. The organisation, as an organisation, does not.



Listening at the front door

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER ONE

Listening at the front door

What the head of customer services does next is less dramatic than a restructuring and more demanding than a survey. For a fortnight she spends two mornings a week at the end of the row with the spare headset, and she asks two of her team leaders to do the same. They listen to real calls, and for each one they write down three things on a single sheet: what the person wanted, in the person's own words; whether it was value demand or failure demand; and, if it was failure, what had failed and where.

The sorting is done from the caller's point of view, not the organisation's. That is the decision that most changes the result. A call to chase a repair that has not happened is failure demand even if the repair is still within the organisation's own target time, because the caller would not have needed to ring if she had known when it would happen. A call to ask what a letter meant is failure demand, because a letter that needs a phone call to explain it has failed. The temptation, which every team that tries this meets in the first week, is to classify from inside: *that one was the customer's fault, she didn't read the letter.* Seddon's rule, as his followers teach it, is that the customer's reading of the letter is the only reading that counts, because it is the one that produced the call.

The advisers notice the change before the managers do. Santa Anna's canon records a sentence that people on the floor of large organisations say about the people above them, collected in a case about a group of facilities whose new chief operating officer went to listen: *"They don't understand what it's like on the floor."* For two mornings a week, the person who sets the targets for the repairs line is sitting on the floor, and the advisers start, cautiously, to tell her things that are not on the sheet.

By the end of the fortnight the three of them have sheets for some four hundred calls. They do not report a percentage to the board. They do something with the sheets that a percentage would hide. They tally the failures by where they were made, and the tally shows what such tallies usually show, which is that a handful of causes account for most of the second calls. In this composite the first is the five-hour slot that nobody can plan a day round. The second is the job that needs access to another flat. The third is a letter.

Chapter 7 carries the demand sheet, the sorting rules and the tally, with the mistakes that are made at each step, and Chapter 6 places them in the order of a first ninety days. What can be said here is how the fortnight changes the meeting that follows it. The head of customer services takes her three causes not to her own team, who already know them, but to the people who own the booking system, the schedule and the letter. She takes the sheets, with the callers' words on them. And she asks for one thing about each: that someone who owns the cause listens to ten of the calls it produces.

It is worth saying what she does not do. She does not set a target for failure demand. She does not rank teams by it, or put it on the screen above the row beside the twenty-second figure. She has read about National Indicator 14, or something like it, and she knows that the moment the share of failure demand becomes a number that people are judged by, the sheets will start to fill with value demand. The work of measuring what the change is worth, honestly enough that the finance director will sign it (Chapter 5), comes later, and it is done differently.

There are failure modes in the listening itself, and they are predictable. A manager who listens to calls once and then delegates the listening has learned something and taught nothing. A tally that is presented as a league table of departments turns a study into a blame exercise, and the departments will contest the classification rather than the cause. And a team that finds failure demand and responds by writing a better

script for the advisers has treated the symptom at the one desk that never caused it. The listening is where the work begins. The change happens upstream.

The fortnight also changes how she reads a different number. The repairs line marks most of its calls as resolved on the first contact, and that figure, too, is green. But the woman with the ceiling was resolved on her first contact, and on her second, and on her third. What it would take to count a case as finished only when the customer agrees it is finished is a question with its own history, and it begins with the ticket the screen says is closed (Chapter 2).

Listening in your own building

The article this chapter began as ended at the end of a row of desks. The book can go further, into the finer points a manager meets when she tries this in her own organisation, and the first of them is where to sit.

Listen where the demand arrives, and in every place it arrives. For most organisations that is more than the telephone. It is the email inbox, the web form, the letters that are scanned and queued, the counter or reception, and the conversations operatives and engineers have on doorsteps. Each channel carries its own mix. Letters often carry the cases that have already failed twice, because people write when ringing has not worked. Doorsteps carry the failures nobody has reported yet. A fortnight of listening that covers only the phone will understate failure demand in the places where it is most expensive.

Spread the listening across the week. Monday mornings carry the weekend's failures; Friday afternoons carry the chasing of promises made for that week. Two mornings and one afternoon a week, for two weeks, by three people, will usually produce enough calls to see the pattern. That is a planning guide from practice, not a finding; stop when the next twenty calls tell you nothing the last hundred did not.

Sort from the customer's point of view, every time, and keep a short list of the kinds of failure so that the tally can be read. Five kinds cover most of what a service organisation hears. *Chasing*: the customer is asking where something is. *Not done right*: it was done, and it was wrong. *Not understood*: a letter, a bill or a message needed explaining. *Sent elsewhere*: the customer was told to contact someone else, and has. And *repeat*: the same failure as before, again. Anything that does not fit goes in a sixth column, *not sure*, and the listeners discuss those together at the end of each session rather than forcing them into a box.

Tell the team what the listening is for before the first session, and tell them plainly: it is a study of the system, not an audit of the people answering, and no adviser's name goes on any sheet. Then keep the promise. The first time a sheet is used to discuss an individual adviser's handling of a call, the advisers will start choosing which calls the managers hear.

Take causes upstream with the evidence, not a summary. The tally shows where the failure was made; the sheets show what it sounded like. The most useful single request a manager can make of the person who owns an upstream system is the one the head of customer services made: listen to ten of the calls your sys-

tem produces. Ten is enough to hear the pattern and few enough to fit into a morning. Most owners who do it change something within the month, because it is hard to keep a rule after hearing ten people explain what it did to them.

And begin with the cheapest change. The first failures removed should be the ones that cost almost nothing to fix: a sentence in a letter, a field on a job sheet, a text message the evening before. They prove to the advisers that the listening leads somewhere, and to the owners upstream that the study is not a prelude to blame. Chapter 6 places this listening in the first month of a ninety-day plan, and Chapter 7 gives the demand sheet, the sorting rules and the cause tally as working pages.

Ten past eight, a month later

It is ten past eight on a Monday in December, and the head of customer services is at the end of the row again, with the same spare headset. The figure above the desks is still green. The booking system now offers two slots, morning and afternoon, and a text the evening before with an hour. Jobs that need access to another flat carry a flag, and the scheduler rings the neighbour before the plumber leaves the depot. The letter has been rewritten by two advisers and a tenant, and it is shorter.

None of this has emptied the line. People still ring about leaks, for the first time, and that is what the line is for. The adviser beside her takes a call from a man whose bathroom tap will not turn off, and books the job, and the man says thank you and goes. Two calls later a woman begins with four words, and the head of customer services writes them down on her sheet, and the place they came from beside them.

It is still the repairs line. It is quieter by one ceiling.

CHAPTER TWO

The Ticket Marked Resolved

A case counted by its contacts can be solved four times before it is solved once.



The Ticket Marked Resolved

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER TWO

Half past nine, the email

It is half past nine on a Wednesday evening, and a man is sitting at his kitchen table with the laptop open beside a plate he has pushed to one side. An email has just arrived from his energy supplier. *Your case has been resolved,* it says. *How did we do? Tell us in one click.* Below it are ten small numbers in a row, nought to ten.

It is the fourth such email in five weeks. The case is a bill for a meter reading taken before he moved in, from the previous tenant's last month. He rang in October and was told it would be corrected. He received

a corrected bill with the same figure on it. He used the chat on the website and was told the reading had been updated. He wrote a letter, because somebody at work said a letter goes to a different team. Each contact was closed. Each closure produced an email with ten numbers in it. The figure on the bill has not changed, and a reminder about it is due on Friday. His jaw has been set since the subject line.

He is a composite, drawn from the way these cases are described in the research below and in the house's own practice, and no firm is meant. Somewhere in the supplier's offices, on a screen he will never see, his case appears four times, as four contacts, each handled within its target time, each marked resolved at the first attempt. By the measures the contact centre reports, the supplier has served him well four times. By his own measure, it has not yet served him once.

Ammanuel Santa Anna's canon has a short list of signs that an organisation is hearing its own environment poorly, and one of them reads like a description of the man's last five weeks: "Customer feedback is collected but rarely changes organizational behavior in meaningful ways."[6] The email is feedback, collected. The bill is behaviour, unchanged.

This chapter makes one claim. A firm that measures its contacts instead of its customers' cases will report a problem solved many times before it is solved once, and the gap between the two counts is where the customer's information is lost a second time. The first time was the call itself, the second call that a repairs line answers well and learns nothing from (Chapter 1). The second is the ticket, closed, which turns the customer's unfinished account into a finished number.

A queue built for arrivals

The way most contact centres are measured comes from a respectable place, and it is worth understanding before it is criticised. The mathematics of telephone traffic is more than a century old, and the modern contact centre is managed with its descendants: forecasts of how many calls will arrive in each half-hour, how long each will take, and how many advisers are needed to answer a given share of them within a given number of seconds. In 2003 the operations researchers Noah Gans, Ger Koole and Avishai Mandelbaum published a long tutorial and review of the field, and their summary of it is careful in both directions.[7] Call centres, they wrote, are "fascinating sociotechnical systems in which the behavior of customers and employees is closely intertwined with physical performance measures", and in them traditional operational models "are of great value—and at the same time fundamentally limited".

The value is real. A contact centre that cannot forecast its calls will have too few people on a Monday morning and too many on a Thursday afternoon, and every customer will wait longer for it. The limit is in the unit. The models count arrivals. A call is an arrival, a chat is an arrival, a letter scanned into the system is an arrival, and each is given a reference, handled, and closed. The man at the kitchen table is not an arrival. He is one case that has arrived four times, by three routes, and no part of the arrival-counting apparatus is built to notice that the four are the same.

That is not a failure of the people who run the centre. It is what the instrument was designed to see. A queue built for arrivals will report on arrivals, and if the most important fact about a contact is that it is the fourth about the same thing, the queue will not report it. There is a cost hidden in that blindness, and it lands on the forecast itself. Each of the man's four contacts was an arrival the centre had predicted, from the arrivals of previous years, and staffed for. A contact centre that has always taken repeat contacts will forecast repeat contacts, hire for them, and meet them within its target time, and every one of those measures will be met. The firm is not only failing to see the repeats. It is paying, in advance and in good faith, for the people to answer them. The figure called first-contact resolution is meant to repair this, and in many organisations it is measured honestly. But it is often measured from the firm's side: an adviser ticks a box to say the matter was resolved, or the system assumes it was if the customer does not call back on the same channel within a few days. A letter, sent to a different team, is not a call back. The man's case was resolved at first contact four times.

Santa Anna's canon, writing about what organisations measure, defends the numbers before it questions them. Revenue, growth, market share, satisfaction scores, it says: "These metrics aren't bad — they carry real information." The contact centre's figures carry real information about the queue. They carry almost none about the man.



Six minutes, and the board on the wall

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER TWO

Six minutes, and the board on the wall

There is a second layer, and it concerns what the numbers do to the people who are measured by them. In many contact centres the figure each adviser is watched on is average handle time: how long, on average, each contact takes. It is a reasonable thing to know, for the forecast. It becomes something else when it is put on a board beside each adviser's name.

An adviser measured on handle time learns, quickly and without anyone telling her to, which kinds of call make her number worse. A billing dispute that needs a meter history pulled from another system, a read-

ing checked against the date of a tenancy, and a corrected bill raised by a team in another building takes twenty minutes to do properly. It takes six to log, promise, and close. She is not dishonest. She is doing what the board on the wall asks of her, and the board does not ask whether the bill was right. W. Edwards Deming, the statistician whose fourteen points for management were set out in *Out of the Crisis* (1986), argued against numerical quotas and work standards on the floor for exactly this reason: a quota measures the count, and the count can be met in ways that defeat the purpose.[8] The argument is summarised here from his book's familiar reading, not from its text, which was not reopened.

John Seddon, whose work on failure demand opened this book, made a related argument about service organisations in *Freedom from Command and Control*: that measures of activity, such as calls handled and time taken, drive the work towards the measure, and that the measures worth having describe the system's purpose from the customer's point of view, such as how long it takes, end to end, for the thing the customer wanted to be done.[9] That too is given in summary, from his book's argument as it is usually taught.

The Margin, another of the house's books, reads [what a number beside a name does to the person who reads it](https://hautelumiere.life/editions/rsd-master-rsd-in-acedemia-final-manuscriptt/page/), and the contact centre is a building full of such numbers. The adviser who closes the man's case in six minutes knows, as she closes it, that it will come back. She has learned not to mind, because minding does not change the board. Santa Anna's canon, in its history of how programmes have been assessed, describes the same confusion in a setting far from any contact centre:

> *But output counting, as a comprehensive approach to assessment, has a fundamental limitation: it confuses activity with achievement. A literacy program that distributes ten thousand books has produced an impressive output. But if those books sit unread on shelves — if the program failed to account for the cultural context, the reading levels of the recipients, or the absence of lighting in their homes — then the output tells us almost nothing about the program's actual contribution to human flourishing.*

The board on the wall counts books distributed. The man's bill is a book nobody read.

How did we do?

The email at the kitchen table asks a question, and the answer it collects is less informative than it looks. Satisfaction surveys sent after a contact are answered by the people who choose to answer them, and those people are not a random sample of the customers who were served.

The marketing scientists Abbie Griffin and John Hauser studied how firms gather the voice of the customer for product development, in a paper published in 1993, and among their findings was one they stated plainly in its abstract.[10] Their data, they wrote, "demonstrate a self-selection bias in satisfaction measures

that are used commonly for QFD and for corporate incentive programs", QFD being quality function deployment, the method many firms use to carry customers' needs into a product's design. Their study was of product development, not contact centres, and it should be read that way. But the mechanism travels easily. A customer whose problem was solved and who feels warmly towards the firm may answer. A customer whose problem was solved and who feels nothing may not. A customer whose problem was not solved, on his fourth email, may answer with a zero, or may close the laptop, which leaves no number at all.

Santa Anna's canon describes the same limit in a passage about the surveys organisations run on their own staff, and the sentence holds for customers too: "The survey captures what people are willing to report — not what they actually feel." The man is willing to report a great deal. He is not willing to report it into ten small numbers in a row, for the fourth time.

The house's Luminous 100, its reading of the frameworks every consultancy sells, makes a remark about the best-known one-question survey that applies to every survey of its kind. Of the Net Promoter Score it says: "The follow-up qualitative question ("why?") is where the real insight lives." The email at the kitchen table has no *why*. It has ten numbers, and a button.

There is a fairer word to say for the survey. It is cheap, it is regular, and a firm that watches its scores over years will see some things move that it would otherwise miss. The trouble is not the survey. It is the habit of treating the score as the customer's account, when the customer's account is in the contacts, in his own words, four times, and the score is what is left when the words are taken away.

Following one case home

What changes this is a different unit of attention, and it can be adopted in a week without new software. The director responsible for the contact centre, or the head of billing, or whoever will listen, takes ten customers who contacted the firm last month. Not ten contacts: ten customers. For each one, she gathers every contact the customer made about the same matter, by every route, calls and chats and letters and emails and the visit to the shop, using the customer's account rather than the ticket reference. And she reads them in order, as the customer lived them.

Santa Anna's canon gives the instruction in one line, as part of its method for reading an organisation through its processes rather than its reports: "Follow a customer complaint from receipt to resolution." The discipline is in the word *resolution*. It means the point at which the customer would say the thing was done, not the point at which a ticket was closed. For the man at the kitchen table, it is the day a bill arrives with the right figure on it, and nothing else.

For each case she writes down three numbers and one sentence. How many contacts did the customer make before the matter was done? How many days passed between the first contact and the last? How

many of the contacts were the customer chasing, or explaining again what he had already explained? And the sentence is the customer's own, taken from the first contact, of what he wanted.

The order of what she does matters more than the numbers. She reads the cases with an adviser beside her, one who handles such contacts every day, because the adviser will see at once where each case went wrong: the system that could not show a meter history, the team in another building with no telephone, the letter that went into a queue nobody owned. She does not ask the adviser why the case was closed. She asks what the adviser would have needed to finish it. The answers are usually specific, and they are usually the same three or four things.

Read the man at the kitchen table this way, and his case comes out on a single line. Five contacts by three routes, over forty-one days, four of them chasing or explaining again. His first sentence, from the call in October: *the reading on this bill is from before I moved in.* And the adviser's answer, when the director asks what she would have needed to finish it at the first contact, takes ten seconds. She would have needed to see the date the tenancy began beside the date of the reading, which the billing system holds and the contact screen does not show, and she would have needed the authority to raise a corrected bill herself, which is held by a team in another building. Neither is expensive. Neither has been asked for, because nobody before had read his five contacts in a row.

The other nine cases are different in their details and similar in their shape. A direct debit set up twice. A moving-house request that reached one system and not another. A complaint about a smart meter that was closed as a technical query. In most of them the information needed to finish the case was already somewhere in the firm, and in most of them the person who spoke to the customer could not see it or could not act on it.

Chapter 7 carries the case trace as a working page, with the rules for gathering contacts across channels and the count of clean cases, and the mistakes made at each step, and Chapter 6 places it in the order of a first ninety days. It is the customer's version of [walking one ordinary order backwards through a building, desk by desk]({{BOOK:orgdev-flow}}), and it shows the same thing a walk shows, which is that most of the time is spent waiting and most of the waiting is in the gaps between teams.

There are mistakes to watch for, and they are familiar. A director who chooses the ten worst cases, the ones that reached the chief executive's inbox, will design for a crisis rather than for the ordinary month. A trace that gathers only calls, because calls are easiest to find, will miss the letter, and the letter is often where the case went to die. And a director who reads the cases and then sets a target for contacts per case has repeated, one level up, the error the board on the wall made: the count will fall, and some of it will fall because customers have stopped trying. The reading is for finding causes. The measure that will later say whether the causes were removed has to be designed with more care than a target, and with the customers who give up (Chapter 3) counted as well as the ones who ring.

What the ten cases usually change first is a definition. A case is finished when the customer does not need to come back about it. That is a definition a firm can measure from its own records, by looking, after an agreed period, for any further contact from the same customer about the same matter, by any route. It is harder to measure than a ticked box. It is also the only version of first-contact resolution that the man at the kitchen table would recognise.



The second unit

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER TWO

The second unit

The deeper change is in what the director now sees when she looks at the contact centre's figures. Before, she had one unit, the contact, and a set of measures that described the queue very well. After, she has two, and the second, the case, describes the customer. The two disagree, often, and the disagreements are the most useful thing on her desk. A month in which contacts fall and cases take longer to finish is not a good month. A month in which contacts rise because a new letter invites customers to ring before a problem becomes a dispute may be an excellent one.

Santa Anna's canon describes the healthy version of this in a list of signs that information moves well through an organisation, and the second of them is the plainest: "Bad news travels as fast as good news." The case trace is how bad news about a bill travels to the person who can change the system that produced it, in the customer's own words, before the reminder goes out on Friday.

It also changes the adviser's position. She was measured on six minutes. She is now one of two people who read a case end to end, and the only one of the two who could have predicted, in the first minute of the man's first call, that the meter history would be the problem. A firm that asks her what she would have needed has found the cheapest source of improvement in the building. It is the same knowledge it has been paying her for, every day, and closing in six minutes.

Tracing cases in practice

The article this chapter began as ended at a kitchen table. The book can go further, into the finer points of tracing cases inside an organisation that was not built to show them.

Choose the ten at random from an ordinary month, from the list of customers who made any contact, and replace any case that is still open with the next one on the list. The point is to see what ordinary looks like. The cases that reached the chief executive's inbox will be traced anyway; they are not a sample of anything.

Gather by the customer, not the reference. In most organisations a customer's contacts live in several systems: the telephony platform, the case-management tool, the shared inbox, the scanned post, the notes field of the billing system, the shop's own log. Few of these share an identifier other than the customer's account or address. The convenor, or whoever does the gathering, should expect the first case to take a morning and the tenth to take forty minutes, as the routes through the systems become familiar. Write down, for each case, which systems had to be searched; the list is itself a finding, because it is the list an adviser would need to search to see the case whole.

Define *done* before you read. The definition that works is the customer's: the matter is finished when the customer did not need to come back about it within an agreed period. Agree the period at the start, and keep it the same for every case. A few weeks suits most service matters; a billing dispute may need a full

cycle. Write the customer's own first sentence at the top of each trace, because it is the only fair test of whether the case ended where it should have.

Read each case with an adviser, and ask one question at each contact: *what would the person handling this have needed to finish it here?* The answers fall, in most organisations, into four groups: information held elsewhere that the adviser could not see; authority held elsewhere that the adviser could not use; a hand-off to a team with no way back to the customer; and a rule that required the customer to do something the customer could not reasonably know to do. Count the answers by group across the ten cases. The largest group is usually where the first change belongs.

Count the clean cases. Of the ten, how many were done in one contact by the customer's definition? The figure is not a target and should not be reported as one. It is the plainest description an organisation can give itself of how often it serves a customer once, and it should be read beside the contact centre's own first-contact figure. Where the two disagree, the gap is the hidden repeat this chapter is about.

Make the first change at the point of contact. The cheapest improvements are usually to the screen and the mandate: show the adviser the one piece of information that would have finished most cases, and give the adviser authority to act on it within a stated limit. These cost little, and they move the finishing of a case to the person who already knows how. Then trace another ten in four weeks, and expect the largest group to have changed. Chapter 6 places the trace in the second month of the ninety days, and Chapter 7 gives the case trace and the clean-case count as working pages.

Half past nine, in February

It is half past nine on a Wednesday evening in February, and the man is at the kitchen table again, the laptop shut. In January, after his fifth contact, a different adviser rang him back, which nobody had done before, and said that she could see the reading came from the previous tenancy and that she had the authority to correct it herself, and she did, while he was on the phone. The bill that came a week later had the right figure on it. He has not contacted the supplier since.

No email arrived to say his case had been resolved. There was nothing left to resolve.

The plate is still beside the laptop, and this time he has finished his supper.

CHAPTER THREE

The Customer Who Left Quietly

The customer who leaves without a word has said something precise, by the one route a firm cannot read.



The Customer Who Left Quietly

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER THREE

Friday, the accounts list

It is a quarter to six on a Friday in March, and the owner of a builders' merchant is in the small office above the trade counter with the month's accounts list printed out in front of her. The yard below is closing. Forklifts are being parked, and somebody is sweeping sand off the concrete by the gate. She is reading the list as she reads it every month, account by account, and she stops at a name she has known for eleven years. A roofing contractor. He has bought from the yard every week since her father ran it: battens, felt, lead, the odd pallet of tiles. His last order was in the second week of January.

There is no complaint on his account. There is no note from the counter, no returned delivery, no unpaid invoice. She rings the counter and asks whether anyone has heard from him, and the man on the counter says no, now she mentions it, he hasn't been in. On the drive home, ten minutes later, she passes a merchant on the ring road and sees his van in their yard, tailgate down, being loaded.

The owner is a composite, drawn from the way these discoveries are described in the research below and in the house's own practice, and no firm is meant. What she feels on the ring road is not quite anger and not quite grief. It is the sinking a person feels when she learns that a relationship ended some time ago and she was the last to be told. Eleven years of Tuesdays, and he did not say a word.

Ammanuel Santa Anna's canon, in a book on doing business as a living system, puts the value of what she has missed more bluntly than most management texts would: "Complaints are the most valuable feedback you receive, and most businesses waste them."[11] The contractor did not even give her the chance to waste his. He took the most valuable feedback the yard could have received, whatever it was, and gave it to the merchant on the ring road instead, in the form of his custom.

This chapter makes one claim. The customer who leaves without a word has told the firm something precise, and the firm hears only that he has gone. A firm that wants its customers' information has to make it cheaper for them to say what is wrong than to leave, and has to go and ask the ones who leave anyway. Silence is not satisfaction. It is a message sent by the one route the firm cannot read.

Exit and voice

In 1970 the economist Albert Hirschman published a short book with a long subtitle, *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*.[12] Its argument begins with an observation that sounds simple and has kept economists and political scientists busy for half a century. When the quality of what a firm offers declines, its customers have two broad ways of responding. They can leave, which Hirschman called exit. Or they can stay and say something, which he called voice: complain, protest, write, argue, ask for it to be put right.

Economists, he observed, had thought almost entirely about exit, because exit is how a market is supposed to discipline a firm. A customer who leaves takes her money with her, the firm's revenue falls, and the firm, if it is paying attention, improves. Hirschman's point was that exit is a poor messenger. It tells the firm that something is wrong and almost nothing about what. Voice carries the information. A complaint says which delivery, which product, which promise, and why it mattered. A falling revenue line says only that it fell. The book's argument is summarised here as it is generally read; its text was not reopened for this book.

Hirschman added a third term, and it is the one the owner needs most. Loyalty, in his account, is what holds a customer in place long enough to use voice instead of exit: the attachment that makes a person willing to say *this isn't good enough* rather than quietly go elsewhere. A loyal customer is not one who never

complains. He is one who complains first. The contractor had eleven years of loyalty, and something in the last few months spent it without his ever raising his voice. Either the yard gave him no easy way to speak, or he had spoken, in some small way, and nothing happened.

Hirschman made one further observation that the owner would recognise, though she has never read him. The customers most sensitive to a fall in quality, he argued, tend to be the first to leave, because they notice it soonest and mind it most, and their going deprives the organisation of exactly the voices that could have told it what was wrong. The point is given here as it is usually summarised from his book. In a builders' merchant it has a face. The contractor who notices first that the felt arrived after his men had gone is the one who runs the tightest jobs, keeps the best clients and pays on time. The customer who does not notice, or does not mind, stays. A yard that loses its most exacting customers silently, one at a time, is left with the ones who ask least of it, and its accounts list will look steady for a long time before it does not.

Santa Anna's canon has a sentence that frames the whole of Hirschman's book from the side of the firm. Writing on the way an organisation and its customers shape each other, it says that the company shapes what customers expect through what it offers, and "customers shape the company's evolution through their responses." Exit is a response. So is voice. Only one of them can tell the yard what to change.



One complaint in three

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER THREE

One complaint in three

How often do dissatisfied customers speak? The question has been asked many times, and the oldest careful answer comes from a survey of American households carried out in 1975. Arthur Best and Alan Andreasen asked people about the problems they had had with common purchases, and what they had done about them, and published their findings in 1977.[13] Consumers, they reported, perceive problems with many of the products and services they buy, and "voice complaints concerning about one-third of those problems."

Where they did complain, satisfactory resolution came in somewhat more than half of the cases. Third parties, consumer agencies and the like, played a very small part.

That is one survey, of one country's households, nearly fifty years ago, and its figures should be read as a picture of a pattern rather than a constant of nature. But the pattern it describes has held up in the larger literature on complaint behaviour, and its shape is what matters for the owner. Most problems are never reported to the firm. Of those that are, a large share are not put right. And the customer's private responses, telling friends, quietly switching, are invisible to the firm by design. The marketing scholar Jagdip Singh, setting out a taxonomy of complaint behaviour in 1988, distinguished exactly these routes: voice to the seller, private responses such as word of mouth and exit, and complaints to third parties.[14] Only the first arrives at the counter.

A word is needed here about the figures that usually appear in this place. A series of studies carried out for the United States Office of Consumer Affairs in the late 1970s and 1980s, by a research firm known as TARP, is the source most often given for striking statistics about how few dissatisfied customers complain and how many others each one tells. Those statistics circulate in many versions, attributed to different years and different reports, and this book did not reach the primary reports to check any of them. It therefore prints none. Best and Andreasen's one-third is less dramatic and better sourced, and it is enough.

Santa Anna's canon describes the attitude that turns this pattern into a loss, in its passage on how ventures listen, and the practice it proposes instead:

> *Instead of resolving complaints and closing tickets, practice complaint archaeology — dig beneath the surface complaint to find the unmet need, the broken expectation, the relationship wound. A customer who says "Your software is slow" might really be saying "I feel unsupported when I'm under deadline pressure."*

A complaint that arrives at the yard, if one third of problems produce one, is not a single customer's grievance. It is a sample, and it is usually a sample of something larger than its words.

Why a firm should want more complaints

In 1987 the economists Claes Fornell and Birger Wernerfelt, starting from Hirschman's exit and voice, built a formal model of what they called defensive marketing: the effort a firm spends keeping the customers it has, as opposed to winning new ones.[15] Their conclusion, stated in the paper's abstract, contradicts the instinct of most managers. Many firms, they wrote, try to reduce the number of complaints they receive, and "this objective is found to be questionable." Their analysis suggested instead that complaints from dissatisfied customers "should be maximized subject to certain cost restrictions," because a customer who complains can be kept, and a customer who leaves silently must be replaced, usually at greater cost.

It is a theoretical model, and its conclusions depend on its assumptions, which this chapter does not reproduce. But its central inversion is worth holding in mind, because it changes what a falling complaint count means. A yard that receives fewer complaints this year than last has not necessarily become better. It may have become harder to complain to. The counter may be busier, the manager harder to find, the phone line answered by a recording. Or customers may have concluded, from the last time, that saying something changes nothing, which is the most expensive conclusion a customer can reach, because it is the one that precedes exit.

The Margin, another of the house's books, reads [why the people who most need a door are rarely the first to knock](https://hautelumiere.life/editions/rsd-master-rsd-in-acedemia-final-manuscriptt/page/), and the same logic runs through a trade counter. The customers most likely to complain are often the ones with the least at stake, and the ones with the most at stake, a contractor with a roof half on and a client waiting, are the ones with least time to stand at a counter and explain. They do not complain. They ring the merchant on the ring road.

Santa Anna's canon draws a distinction, in a quite different setting, that helps here. Writing about what happens when a team speaks about a loss, it separates grief from complaint, and it gives complaint a sentence of its own: "Complaint says: "Something is wrong, and someone should fix it."" The canon's interest there is in grief, and it warns against complaint becoming a habit. But for a firm and its customers the sentence is the definition of a gift. A complaint is a customer telling the firm that something is wrong and that the firm is the one who could fix it. It is, in Hirschman's terms, the customer choosing voice when exit was available. The silent contractor made the other choice.

The call to the customer who left

What the owner does on Monday is the oldest practice in service, and the one most firms delegate until it means nothing. She rings him.

She does it herself, not the counter, and not a survey. She does not ring to win him back, and she says so in her first sentence, because a call that opens with an offer is heard as a sales call and answered as one. She says that she noticed he had not ordered since January, that she was sorry not to have noticed sooner, and that she would be grateful to know what had gone wrong, if anything had. Then she waits.

The answer, when it comes, is ordinary, as these answers usually are. In January the yard moved its delivery slots for trade customers from first thing in the morning to a window between ten and two, to suit the new driver's rota. The contractor's men start on a roof at half past seven. Twice in January they stood on scaffolding for three hours with nothing to fix, and on the second occasion the felt arrived after they had gone. He mentioned it at the counter, once, to someone he did not know, who said he would pass it on. The merchant on the ring road delivers at seven. That was all.

The canon's book on doing business as a living system describes the attitude she needs for the call, in its account of what a listening firm does instead of surveying. Most firms listen through instruments, it says, but deep listening is "the practice of receiving another human being's experience with enough openness, patience, and presence that you perceive not just what they're saying but what they're needing, feeling, and becoming." The owner does not need the contractor's feelings. She needs what he needed, which was felt on a scaffold at half past seven.

Chapter 7 carries the call as a run sheet, with the opening sentence, the questions in the order they are asked, the mistakes at each step and how to read the answers, together with a page for reading each complaint that does arrive as a sample of the many that do not, and Chapter 6 places them in the order of a first ninety days. What can be said here is what she does with the answer. She does not argue with it, and she does not tell him about the driver's rota. She thanks him, says what she will look into, and says when she will ring back. Then she rings four more trade customers whose orders have fallen, and two of them mention the delivery window. The third mentions a price. The fourth says, after a pause, that he had not thought anyone at the yard would want to know, and then talks for twenty minutes. She writes down what each of them says in their own words, on one page, and at the bottom of the page she writes the only figure she needs for the decision, which is the sum of their four accounts last year, from her own ledger.

There are predictable ways to get this wrong. A call made by a junior member of staff, from a script, is a survey with a voice. A call that turns into an offer, a discount or a promise, teaches the customer that leaving is how to negotiate. A call made only to the largest accounts misses the small contractor who was about to become a large one. And a call that hears something and changes nothing is worse than no call at all, because it confirms the customer's conclusion that saying something is pointless.

She also changes the counter, which is where the contractor spoke once and was not heard. The man on the counter now writes down anything a trade customer says that sounds like *this didn't work for us*, in the customer's words, on a sheet that she reads every Friday with the accounts list. The sheet is not a complaints log, which would need a form and a category and a reference number. It is a page of sentences. The house's Luminous 100 says of mapping a customer's path through a firm that each point of contact is a meeting between two living systems, and "the quality of that meeting either nourishes both or depletes both." The counter is that meeting, fifty times a day. The page is filled on Monday and read on Friday, and by the following Monday something on it has changed, or the customer who said it has been told why not.

Not every customer who says something will be pleased by what happens next, and how a firm answers the ones it has let down is a subject of its own, with an evidence base more mixed than its folklore (Chapter 4). And not every customer who stops contacting a firm has left it. Some have stopped because their problem was finally solved, and some because they gave up, and a firm that counts the contacts that no longer arrive (Chapter 2) has to know which is which before it can believe the count.



What voice costs, and what it saves

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER THREE

What voice costs, and what it saves

It is worth being plain about the economics, without inventing a figure. The contractor bought every week for eleven years. The yard's own records say what his account was worth in a year; the owner has the number on the printed list in front of her, and it is the only number that matters for her decision. Against it stand the cost of the delivery window, which suited the rota, and the cost of the call, which was ten minutes of her Monday.

The deeper cost is the information that never arrives. For every contractor who leaves silently because of the delivery window, there may be others who are still ordering and grumbling on the scaffold, and others who have quietly moved their urgent orders elsewhere and kept only the routine ones. None of them appears as a complaint. All of them appear, eventually, in the accounts list, as a slow decline that nobody can explain. Fornell and Wernerfelt's model is a formal way of saying what the owner learned on the ring road: a complaint is cheaper than a departure, and a firm that makes complaining easy is buying its own information at the lowest price it will ever be offered.

Santa Anna's canon puts the larger stake in a sentence about where an organisation's future is decided. "The places where organizational systems meet — between departments, between the organization and its customers, between the organization and its community — are the zones of evolutionary potential." The trade counter is such a place. Each contractor who leaves knows why he moved his order, and the counter is the one place in the yard where he might have said so.

Hearing the ones who leave

The article this chapter began as ended above a trade counter. The book can go further, into the finer points of hearing customers who have stopped speaking, in organisations of every kind.

Find them in your own records first. A lapsed customer is one whose pattern has changed: a trade account that ordered weekly and has not ordered for a month, a subscriber who has stopped using the service before cancelling it, a patient who no longer books, a client whose renewals have shrunk. Each organisation has its own pattern and its own records. The list does not need to be clever. It needs to be read, every month, by someone senior enough to act on what it says.

Make the call senior and early. The call to a customer who has left works best when it comes from someone the customer would not expect to hear from, within weeks rather than months of the change, and without an offer. Weeks matter because the customer still remembers what happened. Seniority matters because it tells the customer the question is real. And the absence of an offer matters most of all: a call that opens with a discount is a negotiation, and it teaches the customer, and everyone the customer tells, that leaving is how to get one.

Open with what you noticed and ask one question. *I noticed you haven't ordered since January, I'm sorry not to have noticed sooner, and I'd be grateful to know what went wrong, if anything did.* Then wait, and write down what is said in the customer's own words. Ask a second question only to understand the first answer better. The book's run sheet in Chapter 7 gives the questions in order and the replies that close a call too early.

Read every complaint as a sample. If, as Best and Andreasen found, people voice complaints about something like one problem in three, every complaint that arrives stands for others that did not. The useful question about a complaint is therefore not only *how do we put this right?* but *who else has met this, and not*

said so? The answer is often in the records: the other deliveries in the same window, the other accounts on the same tariff, the other tenants on the same estate. A complaint read this way becomes a search, and the search usually finds customers who are still with the firm and still grumbling.

Make voice cheaper than exit. Most firms make complaining harder than leaving without meaning to: a phone line answered by a menu, a form with a reference field the customer does not have, a counter where the person who listens is not the person who can act. Each of these can be checked by trying it, as a customer would, on an ordinary afternoon. The counter page in Chapter 7 is the cheapest instrument there is, a sheet of sentences in the customer's words, read every week by the person who owns the business.

And do not celebrate a falling complaint count until you know why it fell. Fornell and Wernerfelt's model is a reminder that fewer complaints can mean a firm that has become harder to talk to, and Chapter 5 sets out how to tell the difference from records the firm already keeps. Chapter 6 places the lapsed-customer calls and the counter page in the ninety days.

Friday, the accounts list, in May

It is a quarter to six on a Friday in May, and the owner is in the office above the trade counter with the month's accounts list. The yard below is closing. Trade deliveries now go out from seven, and the driver's rota was changed to fit them, which took a fortnight and one difficult conversation. Beside the accounts list is the counter's page of sentences, three of them this week, one about a pallet of blocks that arrived with a broken strap, one about the price of lead, one about the coffee machine.

She reads down the list and stops at the roofer's name. Two orders in April, three in May. It is not every week, yet. He told her on the phone that the ring-road merchant was all right but that he missed being known.

She turns the page and reads the sentence about the strap.

CHAPTER FOUR

The Apology and the Cause

A failure well recovered is information twice, and most organisations keep only the first half.



The Apology and the Cause

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FOUR

Five to eleven, the desk

It is five to eleven on a Tuesday night, and the night manager of a city hotel is standing behind the front desk with the arrivals list in one hand. The lobby is quiet. The bar has closed, and the only sound is the revolving door turning slowly in the draught from the street. A woman comes through it with a wheeled case and a coat over her arm, off the last train, and gives her name. She booked the room in February, for a conference that starts at eight tomorrow. The night manager looks at the screen, and then at the list, and knows before he has finished reading that the hotel has given her room away.

It happens, on some nights in a busy month. The hotel sells a few more rooms than it has, on the reasonable expectation that some guests will not arrive, and on most nights the arithmetic works. Tonight it did not. He feels it as a heat at the back of the neck, the body's version of the sentence he is about to say, and he says it plainly. *I'm very sorry. We don't have your room tonight. That's our mistake, not yours, and I'm going to fix it now.*

He fixes it well. He has already rung the hotel across the square, which has a room. He walks her there himself, pulling her case, and pays for the room and for her breakfast, and tells her the hotel will cover a taxi to the conference and back. He gives her his name. In the morning she writes a review on her phone, between sessions, that says the staff were the most gracious she has met in years.

He is a composite, drawn from the way these nights are described in the research below and in the house's own practice, and no hotel is meant. So is she. What he does not see is what happens to her review on Thursday, when the revenue manager reads it at the weekly meeting and remarks that it shows the overbooking policy is working: the hotel sold the extra room, and the guest who was moved is happier than if nothing had happened.

Ammanuel Santa Anna's canon begins its account of how an organisation should treat a broken commitment with two flat sentences that the night manager would recognise: "Promises break. This is not a possibility to be guarded against. It is a certainty to be prepared for."[16] The night manager was prepared. The question for the hotel is what the revenue manager is prepared to learn.

This chapter makes one claim. A failure well recovered is information twice, once about how to put things right and once about what went wrong, and most organisations keep only the first. The evidence that a good recovery leaves a customer better off than no failure at all is weaker than its reputation, and a firm that relies on it has mistaken a well-handled cost for a source of loyalty. The recovery is half the work. The other half is the cause.

A paradox, and its folklore

The idea the revenue manager is reaching for has a name. In the service literature it is called the service recovery paradox: the claim that a customer who suffers a failure that is then put right very well can end up more satisfied, more loyal, or more likely to recommend the firm than a customer who suffered no failure at all. It was discussed in management writing from the late 1980s and named in the academic literature in the early 1990s, and it has been taught ever since, often as encouragement to front-line staff, sometimes as something closer to a strategy. The early history is given here in outline, from general knowledge of the field; the first papers were not reopened.

The idea is attractive because it is partly true and because it flatters the people who do the recovering. The night manager did something difficult well, at eleven at night, and the guest noticed. A good recovery is a

demonstration of exactly the qualities a customer most wants to know a firm has: that it will tell the truth when things go wrong, act quickly, and bear the cost. It would be strange if seeing that demonstration made no difference.

Santa Anna's canon, writing about the relationships inside a governing body rather than between a hotel and its guests, goes further than most management texts. Repair, it says, "practiced with skill and genuine care, can actually strengthen the relational fabric of the governance community beyond its pre-rupture state." And it adds, anticipating the objection: "This is not a naive claim." Inside a group of people who will work together for years, that may well be so, and the canon cites the research on relationships and restorative justice in support. The question for the hotel is whether it holds for a guest who stays two nights a year.

The first careful experimental test was published in 1998, by Amy Smith and Ruth Bolton, who asked real customers of a restaurant and a hotel to respond to a described failure and recovery at an organisation they had recently used.[17] Their conclusion was carefully balanced, and its second half is the one the revenue manager has not read. "Although excellent service recoveries can enhance customer satisfaction and increase repatronage intentions," they wrote, "viewing service failures as opportunities to impress customers with good service performance may involve substantial risks."



What the studies found together

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FOUR

What the studies found together

By the mid-2000s there were enough studies of the paradox to combine them, and in 2007 Celso Augusto de Matos, Jorge Luiz Henrique and Carlos Alberto Vargas Rossi published a meta-analysis.[18] Their abstract begins with an admission that is unusual in a paper about a famous idea: empirical tests of the paradox "produced mixed results, with only some studies supporting this paradox." When they combined the studies, they found a significant positive effect of a good recovery on satisfaction. They found no effect on repurchase intentions, on word of mouth, or on the firm's image. And they found that the effect on satisfac-

tion depended on how the studies were done: on whether they were cross-sectional or followed customers over time, whether the subjects were students, and what kind of service was being studied.

Read for the hotel, that is a precise and modest finding. A guest who is rescued well may say, soon afterwards, that she is satisfied, perhaps more satisfied than if nothing had gone wrong. The combined evidence does not show that she is more likely to come back, or to tell anyone, or to think better of the hotel. Her glowing review is real. It is one measure, taken the next morning.

The canon's own reasoning helps explain why the hotel's case differs from the governing body's. Its claim rests on time: "Relationships that have weathered rupture and repair develop a resilience and depth that relationships that have never been tested cannot match." A guest who stays two nights a year has not had a relationship with the hotel long enough for a single rupture and repair to become part of its history. She has had one bad arrival and one good walk across a square. The research suggests that, for her, the walk mostly cancels the arrival, and rarely more.

The next year Jens Michel and Matthew Meuter published a study designed to deal with the problem that had made earlier results so mixed.[19] If the paradox happens at all, they argued, it is a rare event, and most studies are too small to see it reliably. So they tested it in a bank, using more than eleven thousand interviews about real customer encounters. They found that the paradox existed, that it was rare, and that the differences, though statistically significant, were "not very large". Their conclusion for managers is the sentence the revenue manager most needs, and it is worth giving in full: "The implications of verifying a service recovery paradox do not suggest that ineffective service followed by an outstanding service recovery is a viable strategy."

None of this means the night manager's work was wasted. It means its value lies somewhere other than where the folklore puts it. A good recovery limits the damage of a failure, for the customer who suffered it, and that is worth doing well every time. It does not turn a failure into a gain, and a firm that budgets for failures on the expectation that recovery will pay for them is spending money to stand still.

Outcome, procedure, treatment

What makes a recovery good is better understood than whether it pays, and the best-known study of it is about complaints rather than hotels. In 1998 Stephen Tax, Stephen Brown and Murali Chandrashekaran studied customers' accounts of their own recent complaints, and the first thing they reported was sobering: "a majority of complaining customers were dissatisfied with recent complaint handling experiences."[20] The firms in question had been told something was wrong, and had mostly made it worse.

Using the idea of justice from social psychology, they showed that customers judge a complaint by three things. The outcome they receive: the room, the refund, the corrected bill. The procedure used to reach it: how quickly, how easily, how many times they had to explain. And the interpersonal treatment along the

way: whether they were believed, spoken to plainly, and treated as someone who had been wronged rather than someone being difficult. A recovery that gets the outcome right and the treatment wrong is remembered for the treatment. The study also found that a customer's earlier good experiences softened the effect of a bad recovery only "to a limited extent". Goodwill is not a reserve to be drawn down.

The night manager did all three. The outcome was a room, breakfast and a taxi. The procedure was immediate, at the desk, without a phone call to a supervisor or a form for the guest to fill in, because the hotel had given him the authority to act. And the treatment began with a sentence most organisations find hard to say. *The Cairn*, another of the house's books, reads [the sentences that are harder to hedge when they have to be said out loud](https://hautelumiere.life/editions/rsd-rsd-and-self-confidence/page/), and *that's our mistake, not yours* is one of them.

Santa Anna's canon, in its chapter on transforming conflict, draws the distinction that separates the night manager's sentence from the one most guests hear:

> *Acknowledgment is different from apology, though apology may be part of it. Apology says "I'm sorry." Acknowledgment says "I see what happened. I see what it cost. I understand — as fully as I am able — the impact of my actions on you, on us, and on the space between us."*

A guest who hears *we apologise for any inconvenience* has been given an apology. A guest who hears *we don't have your room, that's our mistake, and I'm going to fix it now* has been given an acknowledgment, and the difference is felt before it is understood. The canon names the form the first kind often takes, the apology that places the problem in the other person, and the hotel trade has its own version of it in the phrase *we're experiencing high demand tonight*.

The room that was always going to be there

There is a reason the paradox is weak, and a Japanese quality engineer described it long before anyone tested the paradox. In 1984 Noriaki Kano and his colleagues published a paper, in the journal of the Japanese Society for Quality Control, distinguishing kinds of quality by how their presence and absence feel to a customer.[21] Some qualities are what came to be called must-be: their absence causes strong dissatisfaction, and their presence is simply taken for granted. Some are one-dimensional: the more of them, the better. And some are attractive: their absence is not missed, but their presence delights. The original paper is in Japanese and was not read for this book; the categories are given as a 2013 review of the hundred and forty-seven studies that followed it describes the field, and as they are generally taught.[22]

A room, on the night it was booked, is a must-be quality. Nobody is delighted to find that a hotel has a room for them. Everybody is dismayed to find that it has not. The night manager's recovery added something attractive, a walk across the square, a taxi, a name, and for one guest on one night the attractive thing outweighed the missing one. But a must-be quality cannot be traded against attractive ones indefin-

itely. A guest who is moved across the square twice does not count two delightful walks. She counts two missing rooms.

That is why the revenue manager's reading of the review is the most expensive sentence of the week. It treats a must-be failure, caused by a rule the hotel chose, as if the recovery had settled it. Santa Anna's canon describes the step the hotel has skipped, in its account of what happens after a broken commitment is acknowledged. The impact assessment, it says, "also examines systemic factors", and its first question is whether the commitment was realistic from the start. The canon's principle is that the organisation answers for the conditions as well as the person answers for the act: "Sacred Accountability holds individuals responsible for their commitments and holds the system responsible for creating conditions in which commitments can be realistically fulfilled." The night manager kept his commitment. The overbooking rule broke the hotel's.



Two halves of every recovery

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FOUR

Two halves of every recovery

What changes this is not a better recovery. The night manager's was very good. It is the second half, which most organisations never build.

The first half stays as it is, and is protected. The people who face the customer have the authority to put it right on the spot, within limits that are written down and generous enough to use, because the procedure the customer experiences is the recovery. They say the sentence of acknowledgment, in their own words, without the phrase about inconvenience. They give their name. And someone rings the customer after-

wards, within a day or two, to ask whether what was done was enough. That call is the one most often skipped, and it is the one that tells the hotel whether its recovery worked from the guest's side rather than from the desk's.

It is worth saying what the first half costs the person who does it, because the cost is usually invisible. The night manager walked across the square at eleven, pulling a stranger's case, after a shift that began at three, and he did it well because he has done it before. Each time, he spends something that the hotel does not record: the heat at the back of the neck, the sentence of acknowledgment said with conviction for a decision he did not make, the half-hour afterwards in which the evening does not quite settle. Done once a month, it is part of the job. Done three times a week, because the rule that causes it has not changed, it is how a good night manager becomes a tired one, and then a former one. The recovery card protects him as well as the guest. It moves the weight of the failure, on paper, from the desk to the person who can remove its cause.

The second half begins with a card. Every recovery, however small, produces one: the date, what the customer was promised, what happened instead, what was done, what it cost, and, in the words of the person who did the recovering, why it happened. The night manager's card says *overbooked by two on a conference night; walked one guest to the hotel across the square*. Once a week the cards go not to the front-desk team, who already know, but to the person who owns the system that produced each failure. The overbooking cards go to the revenue manager. And the revenue manager reads them in the same meeting at which she reads the reviews. The two belong together. A review that praises a recovery and a card that records its cause are the same event seen from the guest's side and the hotel's, and a manager who reads only the first will always conclude that the system is working.

Santa Anna's canon, listing the kinds of truth an organisation finds hardest to say about itself, puts first the operational ones, and among them "The product feature that customers hate but the roadmap treats as a success." The overbooking rule is that feature, and the Thursday meeting was treating it as a success. The broken expectation on the night manager's card is a room. The cause is a rule, set for good reasons, that does not distinguish between a quiet Tuesday and the night before a conference that fills the city.

Chapter 7 carries the recovery card, the follow-up call and the weekly cause review as working pages, with the questions in order and the mistakes made at each step, and Chapter 6 places them in the order of a first ninety days. What can be said here is what the cause review usually finds. A small number of causes produce most of the cards, as they produce most of the second calls on a repairs line, and most of those causes are rules or systems that nobody at the desk can change. The overbooking rule is changed to hold back on conference nights. The cost of the walks falls. The night manager's Tuesdays become quieter, which he notices before anyone else does.

There are mistakes that are easy to make. A recovery budget that is treated as a fixed cost of doing business, rather than as a signal, will be spent every year on the same causes. A cause review that becomes a

hunt for whoever made the rule will teach the revenue manager to stop reading the cards. A recovery that pays customers generously and asks nothing, in effect buying their silence, removes the information from the system at the moment it arrives. And a firm that counts its recoveries and celebrates a rising number has measured its failures and called them service.

There is also the guest who is not recovered at all, because she did not say anything. Some customers who meet a failure do not mention it at the desk; they leave in the morning and do not come back, and the customer who leaves without a word (Chapter 3) is the subject of a chapter of its own.

A hotel that makes recovery visible and easy, a desk where it is plainly safe to say *this isn't right*, will hear about more of its failures. That will look, on the recovery count, like getting worse. It is the beginning of getting better, and whether it is better can only be known if the firm decides in advance how the fall in failures will be counted (Chapter 5).

Building the second half

The article this chapter began as ended at a hotel desk. The book can go further, into the finer points of building a recovery that teaches the organisation as well as rescuing the customer.

Write the authority down. The people who face customers need to know, before the night it is needed, what they may do without asking: what they may give, refund, move, rebook or waive, and up to what limit. The limit should be generous enough to use and written plainly enough to be remembered. An adviser who has to ring a supervisor before she can put something right has made the customer wait for the procedure, and the procedure is part of what the customer judges.

Keep the card short. Six lines are enough: the date; what the customer was promised; what happened instead; what was done to put it right; what it cost, in money or time; and, in the words of the person who did the recovering, why it happened. The last line is the one that matters, and it should be written in plain words without a category. *Overbooked by two on a conference night* is more useful than a code for *capacity*.

Make the follow-up call within two days. Someone other than the person who did the recovery rings the customer, asks whether what was done was enough, and writes down the answer. The call does two things. It tells the organisation whether the recovery worked from the customer's side, which is the only side that counts, and it tells the customer that the failure was taken seriously after the moment had passed. It should never contain an offer the customer did not ask for.

Hold the cause review weekly, and invite the owners. The front-line team does not need to read the cards; it wrote them. The people who need to read them are those who own the systems and rules that produced each failure: the revenue manager for overbooking, the head of housekeeping for rooms not ready, the head of IT for the booking that lost its surname. Group the week's cards by cause, give each group to its owner,

and ask for one thing within a week: either a change, or a plain reason why the cause will stay and what it will cost to keep.

Name the pattern, not the person. The canon's guidance on acknowledging difficult truths inside an organisation, that the effective version names systemic patterns rather than individuals, is the rule that keeps the cause review alive. A review that asks who set the rule teaches owners to stop attending. A review that asks what the rule produced, and what it would take to change it, keeps them in the room.

Read recovery spend as a signal. Most organisations record what recoveries cost as a line in the budget. Read it instead as a measure of the failures the recoveries were for, by cause, and watch which causes the spend comes from. A recovery budget that rises because the organisation has made it easier for customers to speak is a good sign; one that rises because the same cause is recurring is a cost that could be removed.

And check which failures are must-be. Kano's distinction gives a cause review its order of priority. A failure of something customers take for granted, a room, a correct bill, an appointment kept, cannot be recovered into delight more than once; it has to be prevented. Those causes go first. Chapter 6 places the card and the review in the ninety days, and Chapter 7 gives the recovery card, the follow-up call and the cause review as working pages.

Five to eleven, in April

It is five to eleven on a Tuesday night in April, and the night manager is behind the front desk with the arrivals list. The revolving door turns in the draught. A man comes through it with a rucksack, off the last train, and gives his name, and the night manager finds it on the screen, and the room is there. He hands over the key and wishes him goodnight.

On the shelf below the desk is the box of recovery cards. There are three in it this month. One is about a broken shower, one about a booking made under the wrong surname, and one about a late check-out that was promised and forgotten. None of them says *overbooked*.

He reads the arrivals list to the end, and every name on it has a room.

CHAPTER FIVE

The Calls That Stop Coming

A fall in what customers need to say is a gain only when the counting was agreed before the answer was known.



The Calls That Stop Coming

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FIVE

Quarter to six, the budget paper

It is a quarter to six on the last Thursday in March, and the finance director of a housing association is reading the paper the board will discuss on Monday. It is the budget for the coming year, and on page four, in a box, there is a line she did not write. Calls to the repairs line, it says, have fallen by a fifth since the autumn, following changes to how repairs are booked. The line proposes that next year's budget assume the fall continues, and that the contact centre's agency hours be reduced to match.

She believes something real has happened. She has sat at the end of the row with the spare headset herself, at the head of customer services' invitation, and heard fewer people begin with *I'm ringing again about*. But four other things sit in her stomach like a stone. The winter was mild, and in a mild winter fewer boilers fail. In January the association launched a way for tenants to report repairs online, and some of the calls that are no longer calls are now forms. The advisers have been sorting calls into value demand and failure demand since November, and a sorting that people know is being watched is a sorting that can drift. And she cannot rule out the hardest possibility of all: that some tenants have stopped calling because, after the third missed appointment, they concluded that calling did not help.

She is a composite, drawn from the way these evenings are described in the research below and in the house's own practice, and no firm is meant. If she lets the line stand and the fall is partly the weather, next winter's calls will arrive at a contact centre that has been cut to fit a mild one. If she strikes it out, the head of customer services, whose team did change the work, will learn that a gain her board can read is a gain her board will not believe.

Ammanuel Santa Anna's canon has a passage that sits over the box on page four. Writing on the history of how organisations assess what they do, it begins by defending the habit before it names the cost[23]:

> *The impulse to quantify is not, in itself, problematic — it is an expression of our longing to understand, to improve, to steward our resources wisely. But somewhere along the winding path from ancient tallying sticks to modern impact dashboards, something essential was lost. We became so good at counting that we forgot to ask what was worth counting in the first place.*

The calls are worth counting. Whether a fall in them is a gain depends on what else is counted beside them.

This chapter makes one claim, and it closes the argument the four chapters before it have made. The information customers give a firm, in second calls, in reopened cases, in silences and in recoveries, holds a gain, and the gain can be measured honestly enough for both sides of the firm to sign, but only if what counts as a gain is agreed before anyone knows the answer. A fall in what customers have to say is good news only when the firm can show they stopped needing to say it.

Exit dressed as improvement

The hardest of her four worries comes from the economist who wrote the book on it. In *Exit, Voice, and Loyalty* (1970), Albert Hirschman distinguished the customer who stays and complains from the customer who leaves, and argued that the second tells an organisation much less than the first.[24] A housing association's tenants cannot easily leave; that is part of what makes social housing different from a shop. But they can stop using voice. A tenant who has learned that reporting a repair produces a five-hour wait and a missed appointment may simply stop reporting, and live with the dripping tap, and the repairs line will record her as a call that did not come.

Read that way, a fall in calls is ambiguous in the most uncomfortable direction. Claes Fornell and Birger Wernerfelt, building a formal model on Hirschman's distinction in 1987, concluded that a firm's goal of reducing complaints is "questionable", and that complaints from dissatisfied customers should, within cost limits, be maximised rather than minimised.[25] Their model was about firms that can lose customers, not landlords who mostly cannot, and its conclusions follow from its assumptions. But the inversion it states is the finance director's worry exactly. A falling count of contacts can be a system that fails less, or a population that has stopped expecting anything.

The two have to be told apart by something other than the count itself, and the association already holds what it needs. Every year, by law, a gas engineer visits every rented home with a gas supply for a safety check.[26] When an engineer finds a repair the tenant never reported, a radiator that has not worked since October, a tap held shut with a cloth, it is written on the job sheet. That record is collected for a different purpose by a different team, and it does not depend on tenants calling. If calls fall and unreported repairs found at the safety check rise, the tenants have gone quiet. If both fall, the system is failing them less.

Santa Anna's canon puts the principle under that check in a line about how an organisation and the people it serves shape each other: "Recognize that you cannot change one side without changing both." The repairs service changed, and the tenants' behaviour changed in response. The measure has to read both.

A mild winter

The second worry is older and has been met before, in [the measurement of a gain on a factory floor]({{BOOK:orgdev-flow}}). A comparison between one period and another tells you little if the two periods differed for reasons nobody controlled. The autumn and winter just ended were mild. A cold winter produces more boiler failures, more frozen pipes and more condensation, and therefore more first calls and more second ones, whatever the repairs service does.

The answer is the one the long practice of measured savings gives. A baseline drawn from two or three years of the association's own records, long enough to show what an ordinary winter and a hard one look like. An adjustment for the weather, agreed before the results are known, calculated from a record the association does not control, such as the number of cold days in each month from the national weather service. And the comparison made not between this winter and last, but between this winter's calls and the number the history says a winter like this one would have produced.

There is a point here about which number to use, and it changes the argument on page four. Calls are the wrong unit for the gain, because calls rise and fall with the number of repairs, and the number of repairs rises and falls with the weather and the age of the housing. What the changes to the booking system were meant to reduce was not repairs but the calls each repair produces, and the visits each repair takes. The association records both. Its phone system can count calls about a job by the job's reference. Its repairs system counts visits per completed job, including those that ended with *no access* or *wrong part*. Contacts per

repair, and visits per repair, are much less sensitive to a mild winter than calls are, because the weather changes how many repairs there are more than it changes how many times each one has to be chased.

Santa Anna's canon, in its list of the principles by which a living system regulates itself, gives the warning that belongs over every comparison of this kind: "Map ALL loops before changing anything — unintended consequences live in unseen loops." The weather is one loop. The new website is another. The tenants' own conclusion about whether calling helps is a third, and it is the one most easily missed.



The label that moves

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FIVE

The label that moves

The third worry is about the sorting. Since November, the advisers have marked each call as value demand or failure demand, which is the practice that the head of customer services began with a spare headset (Chapter 1), and the share of failure demand has fallen. The finance director knows the history of what happens when such a share becomes a number people are judged by, because she has read the committee papers from the years when England's councils reported avoidable contact as a national indicator.[27] A label that matters to the people applying it is a label that will drift.

The social scientist Donald Campbell stated the general law in 1979, in a paper on evaluating public programmes: the more a quantitative indicator is used for social decision-making, the more it is subject to corruption pressures, and the more apt it is to distort the processes it was meant to monitor.[28] That is a paraphrase of his argument, not a quotation, and the paper was not reopened for this book. Nobody on the repairs line is dishonest. But a call about a missed appointment that could be heard as a new request, *can you book me in for Thursday?*, will be heard that way more often by a team that has been praised for the falling share.

The guard against drift is old and cheap. The sorting is used for study, and the measure that goes in the budget is not the share of failure demand at all; it is contacts per repair and visits per repair, which are counted by the systems rather than labelled by people. And once a month, someone with no stake in the result, from the internal audit team or a neighbouring landlord, listens to a random sample of recorded calls and sorts them blind. If the advisers' sorting and the blind sorting begin to diverge, the label has moved, and nobody needs to be blamed for it to be corrected.

The house's Luminous 100 says of lean thinking, the tradition from which the language of waste and failure demand descends, that "The philosophy is more important than the tools." The sorting sheet is a tool. The philosophy is that a call is information about the system, and a team that is rewarded for producing fewer of one label has been handed the tool and not the philosophy.

The channel that shifts

The fourth worry is the simplest to check and the one most often missed. In January the association launched online repair reporting. Every contact that moved from the phone to the website disappeared from the phone figures. Some of those contacts were first reports, which is what the website was for. Some were tenants chasing a repair that had not happened, by the route that no longer required them to wait on hold.

A measure that counts calls will report the website as a gain in the repairs service. It is nothing of the kind; it is the same demand arriving by another door. The case-level count, which follows one customer's matter across every route rather than one contact at a time, is the remedy, applied to a budget: every contact

about a repair, by any route, counted against that repair. If contacts per repair fall across all channels together, the system is failing less. If calls fall and forms rise by the same amount, the system has moved its queue.

Santa Anna's canon, in its method for walking a value stream, gives the question that catches this and every other blind spot in a measure, in one line: "Whose perspective is absent? Customers? Suppliers? The night shift? The temporary staff?" A budget line built from the phone system's figures is the contact centre's perspective. The tenant's is in the repairs system, the website, the gas engineer's job sheet and the complaint that went to the ombudsman. A measure that both sides can sign has to be built from enough of them that no single team's records can flatter it.

What a customer's word is worth

None of this would matter if the information customers give were worth little. The evidence that it is worth a good deal is real and should be stated at its proper strength. In 2004 Eugene Anderson, Claes Fornell and Sanal Mazvancheryl developed a framework for how customer satisfaction affects a firm's future cash flows, and tested it against firms' market values.[29] They found "a positive association between customer satisfaction and shareholder value", and "significant variation in the association across industries and firms." That is an association across many firms, not a promise about any one, and it says nothing about how much any particular change is worth. It is enough to say that what customers feel is not a soft quantity that finance can safely ignore. Santa Anna's book on doing business as a living system says of the stories customers tell about a firm that they are "the living data of your venture — richer, more dimensional, and more strategically valuable than any metric." The finance director does not need to believe all of that to believe the half that matters: the second call is data, and so is the call that never came.

For the housing association, the worth is more concrete and lives in its own records. Every second call is an adviser's time. Every second visit is an operative's morning, a van, and often a part ordered twice. Every repair a tenant stops reporting is a repair that will be found later, larger. The gain from the changes to the booking system is the difference between what those things cost in a year like this one and what the history says they would have cost, adjusted for the weather and the housing, read from records the association already keeps, and checked by someone with no stake in the answer.

What the association does with that gain is a separate question, and it is not this book's to answer. The board may decide to buy fewer agency hours, or to keep the hours and spend them ringing tenants back, or something else. Who decides, and how any part of a measured gain is shared with the people whose work produced it, are questions for each organisation and the people it covers, and they are the substance of an engagement rather than a book. What can be said is that every one of those decisions is easier when nobody in the room doubts the number.



Signing it

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER FIVE

Signing it

Picture the same finance director in the same office a year later, at the same hour.

In April, before the board took any decision, she sat down with the head of customer services, the head of repairs, an adviser from the line and a tenant from the association's scrutiny panel, and agreed what they would count. Contacts per completed repair, across every channel. Visits per completed repair, from the repairs system. A history of three years, so that the baseline showed hard winters and mild ones. An adjustment for cold days, from the weather service's published monthly figures, and one for homes added to or

sold from the stock. Two balancing measures beside the main ones: unreported repairs found at the annual gas safety check, and complaints to the ombudsman. A blind sample of recorded calls each month, sorted by the internal auditor. And the figures, every month, on the wall of the contact centre and in the tenants' newsletter, in the same form the board would see them.

The conversation was harder than she expected. The head of repairs wanted visits per repair left out, because his operatives' first-time fix rate was another department's target, and agreed to include it when the tenant on the panel said that a second visit was the thing tenants minded most. The head of customer services wanted the baseline taken from the autumn before the changes, which had been unusually bad, and the finance director explained, as gently as she could, that a baseline chosen at its worst point flatters the first year and punishes every year after it. The canon's governance chapter describes the rule that made the tenant's sentence count, in a practice where an organisation gathers what its members are noticing: "the frontline employee's observation carries equal weight with the CEO's." The scrutiny panel's tenant was the only person in the room who had waited in for a plumber.

In November the first cold snap arrived, and calls rose, and the board's chair rang her to ask whether the gain had vanished. She sent him the month's page. Calls were up. Contacts per repair, adjusted for the cold, were still below the baseline's range, and unreported repairs at the safety checks had not risen. The chair read it twice and said that he understood it, which was the first time anyone on the board had said that about a figure from the contact centre.

The page itself was plain. Each month's contacts per repair, plotted against the range the three years of history had shown for that month, with the adjustment for cold days written underneath in a sentence rather than a formula. Beside it, visits per repair, plotted the same way. Below them, the two balancing measures, and the result of the month's blind sample: how many of the forty recorded calls the auditor had sorted differently from the advisers, and which way. A point was treated as news only when it fell outside the history's range, or when several in a row moved the same way, which is the old discipline of telling the ordinary variation of a system from a real change in it.

In February the blind sample began to disagree. The auditor had sorted six calls as failure demand that the advisers had sorted as new requests, all of them tenants ringing to rebook appointments that had been cancelled by text. Nobody had been gaming anything. The new text messages had created a kind of call that the sorting rules did not mention, and the advisers had put it in the box that seemed closest. The rule was written into the sorting sheet that week, and the agreement the finance director had signed in April already said what to do when a system changed during the year: name the change, note the month, and read the figures on either side of it separately. It took one meeting.

Chapter 7 carries the measurement conversation for failure demand as a run sheet, with the questions in the order they are asked, the balancing measures and the one-page agreement, and Chapter 6 places it in the order of a first ninety days. It does not carry the design of any arrangement for sharing a gain, which

belongs to the organisation and the people it covers. Luminous Prosperity works with organisations to find the gains they already hold, in what their customers and their people already know, and to share in what is found; sharing is only honest when the measurement is, which is why the house has spent this much of a chapter on the part most consultancies leave to the end.

Every place a firm stops hearing its customer, the second call, the closed ticket, the silent departure and the recovery whose cause nobody reads (Chapter 4), holds a gain, and none of those gains can be counted until the counting is agreed. That is why the measure is settled first, and read last.

The measurement conversation for demand

The article this chapter began as ended with a signed page. The book can go further, into the finer points of a measurement conversation about what customers no longer need to say.

Put the right people in the room. The person responsible for the money. The person responsible for the operation that produces the contacts, and the person who owns the upstream system that produces the failures, because both their records will be read. Someone who speaks for the people whose work will change: an adviser, an operative. Someone who speaks for the customers, where the organisation has a way of finding one: a tenant panel, a customer council, a long-standing client. And the person who will do the calculation, and whoever will check it.

Choose units that follow the case, not the channel. Contacts per completed case, across every route by which customers reach the organisation. Visits, deliveries or attempts per completed job, where the work involves going somewhere. Days from first contact to done, by the customer's definition. These units are less sensitive than raw contact counts to the things nobody controls, such as weather or the size of the customer base, and they are counted by systems rather than labelled by people.

Draw the baseline long. Two or three years of the organisation's own records, where they exist, so that the ordinary range of each measure is visible, including its seasons. Copy the history and date it before anything changes, as Chapter 6 sets out, so that it cannot later be argued over.

Name the adjustments in advance, and take each from a record the organisation does not control. Weather from the national service's published figures. The size of the customer base from the organisation's own register. Changes in law or regulation by their dates. Each adjustment is written down with how it will be calculated before the first result is known. An adjustment proposed after the results are in should be refused, however reasonable it sounds, because the reason it sounds reasonable is that someone now knows which way it points.

Put exit beside every measure. The most important balancing measure for any fall in customer contact is a signal of customers who have stopped asking rather than stopped needing to ask. Every organisation has one, if it looks: problems found by its own staff that were never reported, complaints that went to an om-

budsman or regulator rather than the firm, cancellations, lapsed accounts, reviews written elsewhere. Choose one or two that the organisation already records, and agree that a fall in contacts with a rise in these is not a gain.

Check the labels blind. Where any part of the measure depends on people classifying contacts, have someone with no stake sort a random sample each month without seeing the original classification, and compare. Agreement is the evidence that the label has not moved. Disagreement is a finding about the rules, not an accusation against the people applying them.

Name what will change if the gain is real. A fall in failure demand frees capacity, and capacity is only money when something is done with it: hours not bought, overtime not worked, visits not made, a backlog cleared, customers rung back. The conversation should say, in plain words, which of those the organisation expects, because that is what finance will later look for in its own records.

What this conversation does not settle is how a measured gain is then shared with the people who found it: the proportions, the timing, the thresholds and the terms. Those depend on the organisation, its people and its history, and they are decided with the people who will be covered. For an organisation working with Luminous Prosperity they are the substance of the engagement, not of a book. Chapter 6 places the conversation in the last month of the ninety days, and Chapter 7 gives its questions in order and the one-page agreement.

Quarter to six, a year later

It is a quarter to six on the last Thursday in March, and the finance director is reading next year's budget paper. On page four there is a box, and this time she wrote it. Contacts per completed repair have fallen, it says, against a three-year baseline, after the adjustment for cold days and for the change in stock, and visits per repair have fallen with them. Unreported repairs found at safety checks are within their usual range. The figure is smaller than the one in last year's box. It is also, for the first time, a figure the head of customer services, the head of repairs and the tenant on the panel have each read and signed.

The stone that sat in her stomach a year ago is not there. She initials the page, and closes the paper.

A NOTE FROM THE HOUSE

A fall in what customers need to say is a gain only when both sides agreed how to count it before anyone knew the answer. That agreement is where a gainshare begins. Luminous Prosperity Inc. reads the front door with an organisation, and is paid as a share of the gain it helps find.

REQUEST A CONVERSATION

CHAPTER SIX

The First Ninety Days

Five practices set in order, from listening at the front door to measuring the calls that stop coming.



The First Ninety Days

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SIX

Monday, day one

It is half past seven on a Monday, and the chief executive of a service organisation of some five hundred people has the board's customer report open on the desk beside her coffee. It is a good report, as these reports go. Satisfaction is steady. Calls are answered quickly. Complaints are down a little on last year. She has read the first five chapters of a book like this one, or lived them, and she no longer believes that any of those three sentences means what it seems to. The temptation, which she feels as a restlessness in her hands, is to commission a programme by Friday: a name, a consultant, a customer platform, a launch.

She is a composite, drawn from the way these first weeks are described in the research the earlier chapters read and in the house's own practice, and no organisation is meant. Her organisation could be a housing association, an insurer, a utility, a large practice, a chain of shops: anything that serves people who can ring it, write to it, walk into it, or quietly stop using it. This chapter is written to her, and to anyone in her position: the person who runs something and wants to begin. It describes the first ninety days in order, with who is in the room, what each step needs, roughly how long it takes, the decisions that change the result, the mistakes made at each step and how to see them early, how to read what comes back, and what to do on the first Monday. The working pages it refers to are gathered in Chapter 7.

Two things should be said before the plan. The first is that nothing in it requires new software, a survey or a name. Every practice in this book reads what customers already say, in records the organisation already keeps, and a programme with a launch tends to produce a performance of listening rather than the thing itself. The second is that the order matters more than any single step. The chapters were written in the order a reader meets the problem. The ninety days run in the order the practices depend on each other.

Why this order

The five practices in this book are listening at the front door (Chapter 1), following cases home (Chapter 2), hearing the customers who leave (Chapter 3), the second half of every recovery (Chapter 4) and the measurement conversation (Chapter 5). In the first ninety days they run in that order, with one move brought forward.

The reasoning is this. Every later practice depends on knowing what customers actually say, in their own words, and on the people who lead the organisation having heard it themselves. So the first practice is the listening, because it gives the sponsor and her managers the ear that the case traces, the calls to leavers and the cause reviews all need. The case traces come next, because they show how the failures heard at the front door travel through the organisation, and because their first changes, at the point of contact, are cheap and visible. The calls to leavers come third, because a senior person who has already listened at the front door asks a lapsed customer better questions. The recovery cards come fourth, because by then there are owners upstream who have heard their own causes and are ready to read cards about them. And the measurement conversation comes last, because it measures what the other four have changed.

The one move brought forward is the history. Before day eleven, before any practice has altered anything, finance and the convenor copy and date two to three years of every record that bears on what customers say: contacts by channel, completed jobs and the visits each took, complaints, cancellations or lapsed accounts, and any record of problems found by staff that customers never reported. Chapter 5 explains why. A measure of what stopped can only be read against a history fixed before the stopping began.

Santa Anna's canon, describing the first rule of a practice in which an organisation gathers what its people are noticing, gives the discipline of the first month in one sentence:[30] "No solutions. The temptation to jump from observation to recommendation is suspended." The first ten days are spent finding out.

Who is in the room

Three roles carry the first ninety days, and they should be named before anything else is done.

The sponsor is the most senior person whose organisation's customers the practices will touch. In a small organisation it is the chief executive. In a large one it may be the head of a division or a region. The sponsor's job is not to manage the work. It is to listen herself, to make the calls to leavers that only a senior voice can make, to sit in the cause review often enough that owners attend it, and to protect the people who report what customers say. Expect it to take four to six hours a week in the first month and three to four after that.

The working group is small: five to seven people. It should include two people from the front line who hear customers every day, an adviser and, where the work involves visits, an operative or engineer; someone from finance who will later lead the measurement conversation; the owner of at least one upstream system that customers' failures come from, such as scheduling, billing or booking; and, where the organisation has a way to find one, someone who speaks for customers, such as a member of a tenant panel or a long-standing client. It helps to include one person who joined recently, because newcomers still hear what everyone else has stopped hearing. The group meets for an hour a week.

The convenor keeps the record: the demand sheets and tallies, the case traces, the notes of every call to a leaver, the recovery cards and what happened to each cause, and the dated copy of the history. Often this is someone from operations or from the contact centre's own planning team. The convenor needs about two hours a week, and three in the weeks of the case traces. The quality that matters most is that the front line trusts the convenor to write down what customers said rather than what the organisation would prefer them to have said.

One decision changes the result more than any other at this stage. The sponsor must decide that she will listen herself, at the front door, at least twice in the first month, and read at least three case traces end to end. A sponsor who receives the listening as a summary receives what the canon's healthcare case calls the executive summary, in which what the front line knew has been reframed on its way up until it recommends more training. Santa Anna's canon gives the consultant in that case the question every sponsor should ask herself before delegating the ear: "What is the cost of not trusting the intelligence that lives closest to the problem?"

Days one to ten: listening before changing

The first ten days change nothing visible. They gather four inputs, and the sponsor does part of the gathering herself.

The first is a list of every route by which customers reach the organisation, and who reads each. Telephone, email, web forms, chat, letters, the counter or reception, the doorstep, social media, the complaints that arrive through a regulator or an ombudsman. For each, the convenor writes down roughly how much arrives in a week, where it goes, and who, if anyone, reads it for patterns rather than for handling. Most organisations find at least one route that nobody reads.

The second is the set of reports about customers that reach the sponsor and the board, and what each counts. Answered within target, resolved at first contact, satisfaction score, complaints received, net promoter, whatever the organisation reports. For each the convenor notes, in a sentence, what unit it counts (a contact, a case, a survey response, a customer) and from whose side it is measured. This is not yet a criticism of the reports. It is a map of what the organisation can and cannot see.

The third is two sessions at the front door for the sponsor, of a morning each, with a spare headset or a chair beside the counter, and with nothing written down except what customers said. These are not yet the formal listening of days eleven to thirty. They are for the sponsor's own ear, so that when the sheets come back she has a sound to set them against.

The fourth is the inventory and dated copy of the history, described above, finished and filed unaltered before day eleven.

At the end of the ten days the working group meets for the first time and writes, on a single page, its first guess at where the organisation stops hearing its customers: two or three kinds of contact that seem to come round again, one or two routes that nobody reads, one or two groups of customers whose orders, bookings or contacts have fallen without explanation. It is a guess and should be labelled one. The page is kept and revisited on day ninety.

What goes wrong at this stage. The commonest mistake is to announce. A sponsor who tells the organisation that it is launching a customer programme will find that the front door becomes a performance of good service whenever she is near it. The early sign is a question in a corridor: *what is the new customer thing called?* If people are asking what it is called, it has been named, and the sponsor should stop naming it. The second mistake is to begin with a survey. A survey of customers in week one produces a score from the customers who answer surveys, which is the self-selected group Chapter 2 described. The third mistake is to buy something. A platform bought in the first ten days decides what the organisation will listen for before anyone has listened.



Days eleven to thirty: listening at the front door

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SIX

Days eleven to thirty: listening at the front door

In the third week the listening of Chapter 1 begins. The sponsor and two or three managers, ideally including the head of the contact centre and one manager from an upstream service, listen at every route customers use, on two mornings and one afternoon a week, for a fortnight. Each listener fills in the demand sheet in Chapter 7: what the customer wanted, in the customer's own words; value demand or failure demand, judged from the customer's side; and, for failure, its kind and where it was made.

The listening takes about three hours a session. At the end of each session the listeners spend twenty minutes together on the calls they marked *not sure*, and agree how to sort them. At the end of the fortnight the convenor tallies the sheets by where the failures were made, and the working group reads the tally for an hour.

The decision that changes the result is what is done with the tally. It is taken upstream, to the owners of the systems the failures came from, with the sheets, and each owner is asked to listen to ten of the calls their system produces. It is not turned into a target, a league table or a figure on the wall. The front line watches what happens to the first tally more closely than anything else in the ninety days, because it tells them whether the listening was a study or an audit.

What goes wrong. Listeners classify from inside: *the customer should have read the letter*. Only the phone is heard, and the letters and the doorstep are left out. An adviser's handling of a call is discussed by name, and the next session's calls become careful. The tally is presented to the board as a percentage, and the conversation becomes about whether the percentage is right. The early sign of each is the same: the advisers stop telling the listeners things that are not on the sheet. When that happens, the sponsor should ask one adviser she trusts, privately, what changed.

How to read what comes back. By day thirty the tally should show a handful of causes that account for most of the failure demand heard, and at least one owner should have listened to ten calls and changed something small: a sentence in a letter, a text the evening before, a field on a job sheet. Those first changes matter more than their size. They are how the front line learns that what it hears can travel.

Days twenty-one to forty-five: following cases home

Overlapping with the end of the listening, the convenor draws ten customers at random from an ordinary month and gathers every contact each made about the same matter, by every route, as Chapter 2 describes. The sponsor reads at least three of the ten end to end with an adviser beside her; managers from the working group read the rest in pairs, each pair with an adviser.

For each case the readers fill in the case trace in Chapter 7: the customer's first sentence, the contacts in order, the days elapsed, the contacts that were chasing or explaining again, and, at each contact, what the person handling it would have needed to finish it there. The convenor counts the answers by group, and the clean cases, the ones done in a single contact by the customer's definition.

Gathering the first case usually takes a morning, because the routes through the systems have to be found; the tenth takes under an hour. Reading a case with an adviser takes thirty to forty minutes. The working group spends ninety minutes on the ten together.

The decisions that change the result are about the point of contact. The first change should put in front of the person who speaks to the customer the one piece of information that would have finished most cases,

and give that person the authority to act on it within a written limit. These changes are cheap, they are visible to the front line within weeks, and they move the finishing of cases to the people who already know how. Larger changes, to systems or to the boundaries between teams, are named and given to owners, but they are not allowed to delay the first.

What goes wrong. The ten are chosen rather than drawn, and they are the worst ten. Only calls are gathered, because calls are easiest to find. The readers ask the adviser why a case was closed instead of what she would have needed. A target for contacts per case is set before anyone knows what moves it. The early signs are in the traces themselves: if every case reads as an exception, they were not drawn at random; if no trace contains a letter, the gathering is incomplete.

How to read what comes back. The traces produce three things worth keeping: the clean-case count, which should be read beside the organisation's own first-contact figure and the gap between them noted without drama; the groups of answers, which show where the first change belongs; and the list of systems that had to be searched to see each case whole, which is the list an adviser would need to see it too. Trace another ten in four weeks, and expect the largest group to have moved.

Days thirty-one to sixty: the ones who leave

In the fifth week the convenor produces, from the organisation's own records, a list of customers whose pattern has changed: accounts that ordered regularly and have stopped, subscribers who have stopped using the service, clients whose business has shrunk, whatever lapse looks like in this organisation. The sponsor, and one or two senior colleagues if the list is long, ring between ten and twenty of them over four weeks.

Each call follows the run sheet in Chapter 7. It opens with what was noticed and an apology for not noticing sooner, asks one question, and waits. It contains no offer. Santa Anna's book on doing business as a living system describes the spirit of such a call in its account of a weekly call to customers made for no purpose except to ask how they are: "Not a sales call. Not a feedback call." The call to a leaver has a purpose, but it should sound like that. The caller writes down what is said in the customer's own words and, at the end, says what she will look into and when she will ring back. Then she rings back.

At the same time the counter page of Chapter 3 is started at one or two places where customers speak to the organisation face to face or by phone: a sheet of sentences, in the customer's words, of anything that sounds like *this didn't work for us*. The sponsor or the owner of that part of the business reads it every Friday, and every line receives an answer within a week: changed, and how; not changed, and why; or passed to someone named, on a date.

Every complaint received in these weeks is also read as a sample, as Chapter 3 describes: for each, the convenor asks the records who else has met the same thing and not said so, and notes the answer beside the complaint.

The decision that changes the result is the second call. A customer who is rung, tells the truth, and hears nothing further has learned that speaking is pointless, and so has everyone she tells. The sponsor should not make more first calls than she can follow with a second.

What goes wrong. The calls are delegated to a junior member of staff with a script, and become a survey. An offer is made, and leaving becomes a negotiation. Only the largest accounts are rung. The counter page becomes a complaints form with categories, and the sentences stop. The early signs: lapsed customers who are rung decline to talk, which usually means the opening sounded like a sales call; the counter page has fewer sentences each week, which usually means the last week's were not answered.

How to read what comes back. By day sixty the notes of the calls and the counter page should show a few causes, some of which the listening at the front door has already heard and some of which it has not, because customers who leave are often the ones who never rang. Put the two lists side by side. A cause that appears on both is a cause worth a great deal.

Days forty-one to seventy: the second half of every recovery

Overlapping with the calls, the organisation builds the second half of its recoveries, as Chapter 4 describes. The first step is to write down the authority the front line has to put things right without asking, and the limit, and to make sure every person who faces customers knows it. The second is the recovery card: six lines, filled in for every recovery however small, with the cause in the words of the person who did the recovering. The third is the follow-up call, within two days, by someone other than the person who recovered, to ask whether what was done was enough.

The fourth is the weekly cause review. The convenor groups the week's cards by cause, and each group goes to the owner of the system or rule that produced it, who is asked for one thing within a week: a change, or a plain reason why the cause will stay and what it will cost to keep. The review takes thirty to forty-five minutes. The sponsor attends the first four and every fourth after that.

Santa Anna's canon, in its guidance on how an organisation names difficult truths about itself, gives the rule that keeps the review alive: "Effective acknowledgment names systemic patterns that individuals participate in but that no single individual controls". The cards describe rules and systems. The review asks what each produced, not who made it.

The decision that changes the result is who reads the cards. If they are read only by the front-line team, the organisation will become very good at recovering from the same failures. If they are read by the owners of the causes, it will start to have fewer of them.

What goes wrong. The authority is written so narrowly that nobody uses it. The cards acquire categories and drop the sentence about why. The cause review becomes a hunt for whoever made the rule, and owners stop attending. The recovery budget is cut because it has risen, when it has risen because customers can now speak. The early sign of each is on the cards themselves: if the *why* line grows shorter each week, people have stopped believing anyone reads it.

How to read what comes back. By day seventy the cause review should have changed at least one rule or system, and the convenor's record should show, for each cause, the number of cards before and after. Read those numbers with care: the change has been in place for weeks, not months. They are the first input to the last step.



Days sixty-one to ninety: measuring what stopped

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SIX

Days sixty-one to ninety: measuring what stopped

In the last month the working group chooses one or two of the changes made, the ones whose effect is most likely to show in records the organisation already keeps, and holds the measurement conversation of Chapter 5. Finance leads it. The operations members, the front-line members, the owner of the upstream system, the customer voice and whoever will calculate and check are in the room. The sponsor attends and says little.

The conversation follows the order in Chapter 7: what has been changed, in plain words; what unit follows the case rather than the channel; what the dated history shows about that unit's ordinary range and its seasons; what moves the unit for reasons nobody controls, and from which outside record each adjustment will be taken; which signal of customers who have stopped asking will sit beside the measure; how any labels will be checked blind; who calculates, and who checks; how the result will be shown to the people whose work produced it; and what will change, in the organisation's own records, if the gain is real. The output is one page, signed by the people in the room. The first conversation usually takes two meetings of ninety minutes.

By day ninety the first monthly figure should be on a wall where the front line can see it, plotted against the ordinary range of its history, with the balancing measure beside it, in the same form the board sees.

Santa Anna's canon describes what such a page is for, in its account of an organisation that hears well: "Every person in the organization becomes a sensor". The page on the wall is how the front line, which has been the organisation's best sensor all along, sees what its reports became.

The decision that changes the result is timing, and most of it was taken before day eleven, when the history was fixed. If a team had already begun to change its work before the history was copied, agree the measurement for that change from now. The earlier improvement is not counted. That is a real cost, and a small one beside a number nobody believes.

What goes wrong. Raw contact counts are used, and a mild season is mistaken for a gain. A web form is launched and the phone figures fall, and the shift is mistaken for a reduction. The balancing measure is left out because it is held by another team. An adjustment is proposed after the results are in. The early sign of each is the same: someone in the room, often the front-line member or the customer voice, goes quiet. The sponsor's job at this step is to notice who has stopped speaking.

What this step does not settle is how a measured gain is then shared with the people who found it: the proportions, the timing, the thresholds and the terms. Those depend on the organisation, its people and its history. They are decided with the people who will be covered, and for an organisation working with Luminous Prosperity they are the substance of the engagement, not of a book. What the first ninety days can do is make that later decision honest, by producing a gain that both sides already believe.

Day ninety: reading what you have

The working group meets on day ninety with six things on the table: the first-guess page from day ten, the demand tallies and the owners' first changes, the case traces with their clean-case count, the notes of the calls to leavers and the counter page, the recovery cards and the cause review's record, and the signed measurement page with its first figure.

It asks a short list of questions, in order. Where was the first guess right, and where was it wrong? Which failures heard at the front door have been removed upstream, and who removed them? How many of the ten cases were clean, and has the second set of ten moved? What did the customers who left say that nobody at the front door had heard? Which causes on the recovery cards have changed, and which owners have not yet answered? Is there a gain both sides believe, and what does it show so far? And then the question that matters most for the next ninety days: which customers, and which members of staff, have started saying things they did not say before?

What not to conclude is as important. Ninety days are too short to say that customers are better served in any lasting way; they are long enough to say whether the practices are being kept. A single month's fall in contacts is not yet a gain, for the reasons Chapter 5 gave. And an absence of findings is itself a finding. It usually means that one of the earlier practices is not yet real, and the first place to look is the front door: is the sponsor still listening herself?

Santa Anna's book on doing business as a living system gives, in its account of asking customers what they need, a four-word instruction that fits the whole of the ninety days: "Read every response yourself." The practices in this book can be delegated, and in a large organisation much of each must be. The reading cannot.

What to do on Monday

If the whole of this chapter is too much to begin with, three things can be done on the first Monday, and none of them needs anyone's permission.

Sit beside the people who answer your customers for one hour, wherever they answer them, and write down the first sentence of every contact in the customer's own words. Mark each one that begins *I'm ringing again*, or means it.

Take one customer who contacted you last month, and find every contact that customer made about the same matter, by every route. Count them, and count the days.

And find one customer who used to buy from you, or use you, and has stopped, and ring them yourself, with no offer, to ask what went wrong.

Monday, day ninety-one

It is half past seven on a Monday, and the chief executive's diary is full again. Two mornings this month are marked *front door*. On the wall of the contact centre, beside the figure for calls answered quickly, there is a second figure, contacts per completed case, plotted against three years of its own history, with a line beneath it for problems staff found that customers never reported. In the convenor's folder there are ten case

traces from August and ten from September, and the second ten have more clean cases than the first. On her desk is a note from a customer she rang in July, who rang back.

None of these has a name. Nobody has announced anything. The board's customer report is still on the desk, and she has written one question in its margin, the one she will ask on Thursday: *how do we know?*

CHAPTER SEVEN

The Field Manual

Twelve working pages, the demand sheet, the case trace, the call to the customer who left, the recovery card and the agreement, ready to carry into a room.



The Field Manual

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SEVEN

How to use these pages

The chapters of this book describe five practices as people doing them.[31] This chapter holds the working pages: the sheets, the run sheets, the question sets and the one-page templates that the earlier chapters point to and do not print. Each page is meant to be photocopied, filled in by hand, and kept in a folder that the convenor of Chapter 6 looks after. Nothing here needs software.

The pages follow the order of the first ninety days, which is also the order of the chapters. Each opens with a line saying where it comes from and when in the ninety days it is used. Where a page gives a script, the

words are a starting point; say them in your own voice. Where it gives a time, the time is a planning guide from practice, not a finding.

Two rules govern the whole manual. Write down what the customer said, in the customer's words, not what the organisation would have preferred. And answer everything you collect: a call to a leaver with no second call, a line on the counter page with no answer, a recovery card nobody upstream reads, teaches customers and staff the opposite of what the practice is for.

What these pages do not contain is the design of any arrangement for sharing a gain: proportions, thresholds, timing or terms. Those belong to each organisation and the people it covers, and they are decided separately.

1 · The demand sheet

From Chapter 1. Used in days eleven to thirty, at every route customers use.

Listener: Route: phone / email / web / letter / counter / doorstep / other Date and session:

Fill one row per contact. No adviser's name goes on this sheet.

#	WHAT THE CUSTOMER WANTED, IN THEIR OWN WORDS	VALUE OR FAILURE?	IF FAILURE: KIND (SEE PAGE 2)	WHERE WAS IT MADE?

At the end of the session (twenty minutes, all listeners together). Read aloud every row marked *not sure*, agree how to sort it, and write the agreed rule on page 2 so the next session sorts it the same way.

2 · Sorting rules: value or failure

From Chapter 1. Kept with the demand sheets, and added to after every session.

The one rule. Sort from the customer's side. If the customer would not have needed to make this contact had the organisation done something, or done it right, it is failure demand, whatever the organisation's own targets say.

The kinds of failure.

KIND	IT SOUNDS LIKE	EXAMPLE
Chasing	*Where is …? When will …? You said …*	A repair still within its target time, but the tenant was not told when
Not done right	*It came, and it's wrong*	The corrected bill shows the same figure
Not understood	*What does this mean?*	A letter, bill or message that needs a call to explain it
Sent elsewhere	*They told me to ring you*	A customer passed from one team to another, and back
Repeat	*It's happened again*	The same fault as last month
Not sure		Discussed at the end of the session

Rules agreed so far (add one line each time a *not sure* is settled):

1. _____
2. _____
3. _____
4. _____

Mistakes to watch for. Sorting from inside (*she should have read the letter*). Hearing only the phone. Discussing an adviser by name. Turning the share of failure demand into a target.

3 · The cause tally and the ten-call rule

From Chapter 1. Used at the end of the fortnight of listening, and repeated at day ninety.

The tally. From all the demand sheets, count the failure demand by where it was made, not by who answered it.

WHERE THE FAILURE WAS MADE (SYSTEM, RULE, TEAM, DOCUMENT)	COUNT	OWNER	OWNER HAS LISTENED TO TEN CALLS? (DATE)	FIRST CHANGE, AND DATE

The ten-call rule. For each of the largest causes, ask its owner to listen to ten of the contacts it produces, from recordings or beside an adviser, before any change is discussed.

First changes. Begin with the cheapest: a sentence, a field, a message the evening before. Record each change and its date on this page; the measurement of Chapter 5 will need the dates.

Reading it. A handful of causes usually account for most of the failure demand heard. Report the causes and their owners, not a percentage.



4 · The case trace

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SEVEN

4 · The case trace

From Chapter 2. Used in days twenty-one to forty-five, for ten customers drawn at random from an ordinary month; repeated four weeks later.

Customer (account, not ticket): Matter:

The customer's first sentence, from the first contact:

Done means: the customer did not need to come back about this within weeks (agree the period once, for all ten).

#	DATE	ROUTE	WHO HANDLED IT	WHAT THE CUSTOMER SAID OR ASKED	CHASING OR EXPLAINING AGAIN?	WHAT WOULD THE HANDLER HAVE NEEDED TO FINISH IT HERE?
1						
2						
3						
4						
5						
6						

Totals. Contacts before done: Days from first contact to done: Contacts that were chasing or explaining again:

Systems searched to see this case whole:

Read with: an adviser who handles such contacts every day. Ask what she would have needed, not why the case was closed.

5 · The clean-case count

From Chapter 2. Used once the ten traces are read.

Group the answers from the right-hand column of every trace.

GROUP	COUNT ACROSS THE TEN CASES	EXAMPLE
Information held elsewhere that the handler could not see		
Authority held elsewhere that the handler could not use		
A hand-off to a team with no way back to the customer		
A rule that asked the customer to do something they could not know to do		
Other		

The clean-case count. Of the ten, how many were done in a single contact, by the customer's definition? of 10.

Beside it: the organisation's own first-contact resolution figure for the same month: . The gap between the two is the hidden repeat.

First change at the point of contact. The one piece of information to put in front of the handler: . The authority to give the handler, and its written limit: .

Reading it. The count is a description, not a target. Trace another ten in four weeks; expect the largest group to have moved.

6 · The call to the customer who left: run sheet

From Chapter 3. Used in days thirty-one to sixty, by the sponsor or another senior person, for ten to twenty customers whose pattern has changed.

Before the call (five minutes). Read the customer's history: what they used to buy or use, when it changed, any contact or complaint on record. Decide nothing about why.

The opening (thirty seconds). *I noticed you haven't [ordered / booked / used the service] since [month]. I'm sorry not to have noticed sooner. I'd be grateful to know what went wrong, if anything did.* Then wait.

One question at a time. Ask a second question only to understand the first answer better: *When was that? What happened then? Did you mention it to anyone here?*

Write down what is said, in the customer's words.

WHAT THE CUSTOMER SAID	WHEN IT HAPPENED	HAD THEY TOLD ANYONE HERE? WHO?

The close. Thank them. Say what you will look into and when you will ring back. Make no offer. Santa Anna's book on doing business as a living system describes the question such a call should feel like: "No form. No survey. A genuine, human question."

The second call (within the time you promised). Say what was found and what changed, or why it did not.

Mistakes to watch for. A junior caller with a script. An offer in the first minute. Ringing only the largest accounts. A first call with no second.

Signs it is working. Customers talk for longer than you expected. Two calls in five mention the same cause. A customer you rang rings back.

7 · The counter page and the complaint read as a sample

From Chapter 3. Used from day thirty-one, at one or two places where customers speak to the organisation.

The counter page. Head it with one sentence: *Write down anything a customer says that sounds like "this didn't work for us", in their words.*

DATE	WHAT THE CUSTOMER SAID	CUSTOMER (IF THEY WISH)	ANSWER	ANSWERED ON

Read every Friday by the person who owns that part of the business. Every line receives one of three answers within a week: *changed, and how*; *not changed, and why*; *passed to (name), on (date)*.

The complaint read as a sample. For every complaint received, ask the records one question and write the answer beside it.

COMPLAINT (DATE, IN A SENTENCE)	WHO ELSE HAS MET THE SAME THING AND NOT SAID SO? (FROM THE RECORDS)	HOW MANY?	WHAT WAS DONE FOR THEM

Reading it. A falling number of lines or complaints is not yet good news; check page 10's balancing measure first.



8 · The recovery card and the follow-up call

HAUTE LUMIÈRE · THE CUSTOMER ALREADY TOLD YOU · CHAPTER SEVEN

8 · The recovery card and the follow-up call

From Chapter 4. Used from day forty-one, for every recovery, however small.

The authority, written down. What the people who face customers may do without asking, and up to what limit:

The card. Six lines.

Date
What the customer was promised
What happened instead
What was done to put it right
What it cost (money or time)
Why it happened, in your own words

The three things a customer judges (Tax, Brown and Chandrashekaran, 1998). Tick each.

- [ ] The outcome: what they received.
- [ ] The procedure: how fast, how easily, how many times they had to explain.
- [ ] The treatment: believed, spoken to plainly, told *that was our mistake*.

The follow-up call (within two days, by someone other than the person who recovered). *I'm ringing about [what happened] on [date]. Was what we did enough?* Write the answer on the back of the card. Make no offer the customer did not ask for.

9 · The weekly cause review

From Chapter 4. Thirty to forty-five minutes, every week from day forty-one.

Who attends. The convenor, with the week's cards. The owner of every system or rule that produced a card. The sponsor for the first four reviews and every fourth after that.

The order.

1. Group the week's cards by cause (five minutes).
2. Give each group to its owner, with the cards themselves.
3. For each group ask what the rule or system produced, not who made it.
4. Mark each cause as a must-be failure (something customers take for granted) or not. Must-be causes go first.

5. Agree what the owner will bring next week: a change, or a plain reason why the cause will stay and what it will cost to keep.

CAUSE	CARDS THIS WEEK	MUST-BE?	OWNER	CHANGE OR REASON, DUE	CARDS IN THE FOUR WEEKS SINCE

Mistakes to watch for. A hunt for whoever made the rule. Categories replacing the *why* line. The recovery budget cut because it rose. Owners who stop attending.

10 · The measurement conversation for demand

From Chapter 5. The history it reads is copied and dated by day ten (Chapter 6); the conversation is held in days sixty-one to ninety.

Who is in the room. The person responsible for the money. The person responsible for the operation, and the owner of the upstream system. Someone from the front line. Someone who speaks for customers, where possible. Whoever will calculate, and whoever will check.

The questions, in order. Do not skip ahead.

1. What, in plain words, has been changed?
2. Which unit follows the case rather than the channel: contacts per completed case, visits per completed job, days to done?
3. What does the dated history show about that unit over two to three years? Plot it. What does an ordinary month look like, in each season?
4. What could move the unit for reasons nobody here controls (weather, the size of the customer base, a change in law)? For each: which outside record, and how will it be allowed for?
5. Which signal of customers who have stopped asking will sit beside the measure (problems found by staff and never reported, complaints to a regulator or ombudsman, cancellations, lapsed accounts)?
6. Does any part of the measure depend on people labelling contacts? Who will sort a blind sample each month?
7. Are all routes counted together, so that a contact moving from phone to web is not read as a fall?
8. Who will calculate, and who with no stake will check?
9. How, and how often, will the result be shown to the people whose work produced it?
10. If the gain is real, what will change in the organisation's own records (hours not bought, visits not made, backlog cleared)?
11. What happens if a system or rule changes during the period?

The one-page agreement. Fill it in together and sign it before the change is read.

What has been changed, and when
Main unit, and the records it comes from
History used, and the ordinary range it shows
Adjustments, each with its outside record and method
Balancing measure: customers who have stopped asking
Blind sample of labels: who, how many, how often
Routes counted
Who calculates; who checks
How and how often the result is shown, and to whom
What will change in the records if the gain is real
What happens if a system or rule changes
Date agreed
Signed

Reading the monthly figure. Plot each month against the ordinary range of its history for that season. Treat a point as news only when it falls outside the range, or when several in a row move the same way. A fall in the main unit with a rise in the balancing measure is not a gain.

11 · The ninety days on one page

From Chapter 6.

DAYS	PRACTICE	WHO	TIME IT TAKES	WHAT IT PRODUCES	EARLY SIGN IT IS GOING WRONG
1–10	Listen: map the routes and the reports, two front-door sessions, fix a dated copy of the history	Sponsor, convenor	Sponsor about 5 hours a week	A first-guess page: where the organisation stops hearing	People asking what the customer programme is called
11–30	The front door: demand sheets, sorting, tally, ten-call rule (pages 1 to 3)	Sponsor and two or three managers	About 3 hours a session, three sessions a week, two weeks	Causes by owner; first cheap changes	Advisers stop saying things not on the sheet

DAYS	PRACTICE	WHO	TIME IT TAKES	WHAT IT PRODUCES	EARLY SIGN IT IS GOING WRONG
21–45	Following cases home (pages 4 and 5)	Convenor gathers; readers in pairs with an adviser	A morning for the first case, under an hour for the tenth; 30–40 minutes to read each	Clean-case count; groups; first change at the point of contact	Every case reads as an exception
31–60	The ones who leave: calls, counter page, complaints as samples (pages 6 and 7)	Sponsor and senior colleagues	15–20 minutes a call; 10 minutes a week for the page	Causes the front door never heard	Leavers decline to talk; fewer lines each week
41–70	The second half of recovery (pages 8 and 9)	Front line; owners; convenor	30–45 minutes a week	Causes changed upstream	The *why* line grows shorter
61–90	The measurement conversation (page 10)	Finance leads; operations; front line; customer voice; calculator and checker	Two meetings of 90 minutes	A signed page and a first monthly figure	Someone in the room goes quiet

12 · The day-ninety review and the first Monday

From Chapter 6. A meeting of the working group and the sponsor, about two hours.

Bring the first-guess page, the tallies and first changes, both sets of case traces, the notes of the calls and the counter page, the recovery cards and the cause review's record, and the signed measurement page.

Ask, in order:

1. Where was the first guess right, and where was it wrong?
2. Which failures heard at the front door have been removed upstream, and who removed them?
3. How many of the ten cases were clean, and has the second set of ten moved?
4. What did the customers who left say that nobody at the front door had heard?
5. Which causes on the recovery cards have changed, and which owners have not yet answered?
6. Is there a gain both sides believe, and what does it show so far, beside its balancing measure?
7. Which customers, and which members of staff, have started saying things they did not say before?
8. Which practice is not yet real, and what would make it real?

What not to conclude. That customers are better served in a lasting way (ninety days show whether the practices are kept). That one month's fall in contacts is a gain. That an absence of findings means an absence of failures; it usually means one practice is not yet real, and the first place to look is the front door.

The first Monday, for anyone who wants to begin with less.

- Sit beside the people who answer your customers for one hour, and write down the first sentence of every contact in the customer's words. Mark each that begins *I'm ringing again*, or means it.
- Take one customer who contacted you last month, find every contact they made about the same matter by every route, and count the contacts and the days.
- Ring one customer who has stopped buying from you, yourself, with no offer, and ask what went wrong.

What your customers have already said

Much of what arrives at a firm's front door is its own earlier failure, coming back with the customer's account of where it happened.

Luminous Prosperity Inc. takes this work into organisations as a gainshare. The flow is opened from the inside, with the people who already hold what the firm knows. The gain is counted from records the firm already keeps, against a measure both sides sign before anyone knows the answer, and the house is paid as a share of what is found: between thirty and ten per cent of it.

The book teaches the reading. What your own calls, tickets and quiet leavers are saying, and what it would be worth to need fewer of them, is read from your records. That begins with a conversation.

REQUEST A CONVERSATION

References

1. Sources for the author's words in this chapter: Santa Anna, A. *Integrated Value Stream Mapping*, "Somatic Stream Sensing" and the history of value stream mapping. The Luminous Developmental Canon; Santa Anna, A. *The Keystone Effect*, "Step 4: Map the Emotional Labor". The Luminous Developmental Canon; Santa Anna, A. *Progressive Impact Assessment*, "The First Wave: Counting Outputs". The Luminous Developmental Canon; Santa Anna, A. *Luminous Entrepreneurship: Doing Business In and As Living Systems*, Chapter 8, "Customers as Co-Creators". The Luminous Library; Santa Anna, A. *Conscious Governance Models*, "Principle 2", the case of the healthcare network. The Luminous Developmental Canon; Santa Anna, A. *Luminous Living Systems*, on structural coupling. The Luminous Developmental Canon; Santa Anna, A. *Organizational Soul Work*, "Consequence 4: The Erosion of Collective Intelligence". The Luminous Developmental Canon; Santa Anna, A. *Adaptive Resource Orchestration*, the case of the eleven facilities. The Luminous Developmental Canon.
2. Seddon, J. (2003). *Freedom from command and control: A better way to make the work work*. Vanguard Education; Productivity Press edition 2005 (verified in Open Library). Seddon, J. (2008). *Systems thinking in the public sector: The failure of the reform regime and a manifesto for a better way*. Triarchy Press (verified in Open Library). Neither book was reopened for this book; the definitions of value demand and failure demand are paraphrased, and the origin story in bank call centres is from secondary accounts.
3. National Indicator 14, "Reducing avoidable contact", in the national indicator set for English local authorities from 2008; removed from the national set for 2010/11. Read through council committee papers (for example Epping Forest District Council, "Measurement of Avoidable Contact"; Copeland Borough Council executive report, 29 June 2010), searched 30 Oct 2026. The central government guidance and Seddon's own statements on NI 14 were not reopened.
4. Ohno, T. (1988). *Toyota production system: Beyond large-scale production*. Productivity Press (verified in Open Library; not reopened). The list of wastes is given in its customary form.
5. Jackson, M. C., Johnston, N., & Seddon, J. (2008). Evaluating systems thinking in housing. *Journal of the Operational Research Society*, 59(2), 186–197. https://doi.org/10.1057/palgrave.jors.2602521 (The quoted sentence is from the abstract.)
6. Sources for the author's words in this chapter: Santa Anna, A. *Luminous Living Systems*, "Signs of compromised structural coupling", "Signs of healthy feedback" and "Process mapping". The Luminous Developmental Canon; Santa Anna, A. *Organizational Soul Work*, "The tyranny of the measurable". The Luminous Developmental Canon; Santa Anna, A. *Progressive Impact Assessment*, "The First Wave: Counting Outputs". The Luminous Developmental Canon; Santa Anna, A. *Systems Sensing Protocols*, "The Organizational Intelligence Gap". The Luminous Developmental Canon; Santa Anna, A. *The Luminous 100*, "Net Promoter Score (NPS)". Luminous Prosperity Inc.
7. Gans, N., Koole, G., & Mandelbaum, A. (2003). Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management*, 5(2), 79–141. https://doi.org/10.1287/msom.5.2.79.16071 (The quoted phrases are from the abstract; the history of telephone-traffic mathematics is summarised from general knowledge.)
8. Deming, W. E. (1986). *Out of the crisis*. MIT Center for Advanced Engineering Study (verified in Open Library; not reopened). The point on quotas and work standards is given in summary.
9. Seddon, J. (2003). *Freedom from command and control: A better way to make the work work*. Vanguard Education; Productivity Press edition 2005 (verified in Open Library; not reopened). The argument on activity measures and end-to-end measures is summarised as it is usually taught.
10. Griffin, A., & Hauser, J. R. (1993). The voice of the customer. *Marketing Science*, 12(1), 1–27. https://doi.org/10.1287/mksc.12.1.1 (The quoted sentence is from the abstract.)
11. Sources for the author's words in this chapter: Santa Anna, A. *Luminous Entrepreneurship: Doing Business In and As Living Systems*, Chapter 8, "Customers as Co-Creators" ("Deep Listening" and "The Complaint Archaeology"). The Luminous Library; Santa Anna, A. *Luminous Living Systems*, on structural coupling and its "Practical Implications". The Luminous Developmental Canon; Santa Anna, A. *Organizational Soul Work*, "Principle 5: Differentiate Grief from Complaint". The Luminous Developmental Canon; Santa Anna, A. *The Luminous 100*, "Customer Journey Mapping". Luminous Prosperity Inc.

12. Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*. Harvard University Press (verified in Open Library). The argument is summarised as it is generally read; the text was not reopened.

13. Best, A., & Andreasen, A. R. (1977). Consumer response to unsatisfactory purchases: A survey of perceiving defects, voicing complaints, and obtaining redress. *Law & Society Review*, 11(4), 701–742. https://doi.org/10.2307/3053179 (The quoted phrase and the findings are from the abstract.)

14. Singh, J. (1988). Consumer complaint intentions and behavior: Definitional and taxonomical issues. *Journal of Marketing*, 52(1), 93–107. https://doi.org/10.1177/002224298805200108 (The three routes are summarised from the paper's taxonomy as it is usually cited; the text was not reopened.)

15. Fornell, C., & Wernerfelt, B. (1987). Defensive marketing strategy by customer complaint management: A theoretical analysis. *Journal of Marketing Research*, 24(4), 337–346. https://doi.org/10.1177/002224378702400401 (The quoted phrases are from the abstract.)

16. Sources for the author's words in this chapter: Santa Anna, A. *Conscious Governance Models*, "Movement 4: Repair — When Promises Break", "Sacred Accountability". The Luminous Developmental Canon; Santa Anna, A. *Conflict Transformation Frameworks*, "Movement Two: Acknowledgment". The Luminous Developmental Canon; Santa Anna, A. *Organizational Soul Work*, "Type 1: Operational Truths". The Luminous Developmental Canon.

17. Smith, A. K., & Bolton, R. N. (1998). An experimental investigation of customer reactions to service failure and recovery encounters. *Journal of Service Research*, 1(1), 65–81. https://doi.org/10.1177/109467059800100106 (The quoted sentence is from the abstract. The early history of the paradox is given in outline from general knowledge; the first papers were not reopened.)

18. de Matos, C. A., Henrique, J. L., & Rossi, C. A. V. (2007). Service recovery paradox: A meta-analysis. *Journal of Service Research*, 10(1), 60–77. https://doi.org/10.1177/1094670507303012 (Findings and the quoted phrase from the abstract.)

19. Michel, S., & Meuter, M. L. (2008). The service recovery paradox: True but overrated? *International Journal of Service Industry Management*, 19(4), 441–457. https://doi.org/10.1108/09564230810891897 (Findings and quotations from the structured abstract.)

20. Tax, S. S., Brown, S. W., & Chandrashekaran, M. (1998). Customer evaluations of service complaint experiences: Implications for relationship marketing. *Journal of Marketing*, 62(2), 60–76. https://doi.org/10.1177/002224299806200205 (Findings and quotations from the abstract.)

21. Kano, N., Seraku, N., Takahashi, F., & Tsuji, S. (1984). Attractive quality and must-be quality. *Hinshitsu (Journal of the Japanese Society for Quality Control)*, 14(2), 147–156. (In Japanese; not read. No DOI was found in Crossref.)

22. Witell, L., Löfgren, M., & Dahlgaard, J. J. (2013). Theory of attractive quality and the Kano methodology: The past, the present, and the future. *Total Quality Management & Business Excellence*, 24(11–12), 1241–1252. https://doi.org/10.1080/14783363.2013.791117 (The count of 147 papers is from the abstract; the categories are given as generally taught.)

23. Sources for the author's words in this chapter: Santa Anna, A. *Progressive Impact Assessment*, "The Counting Impulse". The Luminous Developmental Canon; Santa Anna, A. *Luminous Living Systems*, structural coupling, "Practical Implications", and the living-system principles ("Feedback loops"). The Luminous Developmental Canon; Santa Anna, A. *Integrated Value Stream Mapping*, on whose perspective is absent from a mapping. The Luminous Developmental Canon; Santa Anna, A. *Conscious Governance Models*, "Principle 3: Distribute the Sensing" (the Sensing Circle). The Luminous Developmental Canon; Santa Anna, A. *The Luminous 100*, "Lean Thinking (Toyota Production System)". Luminous Prosperity Inc; Santa Anna, A. *Luminous Entrepreneurship: Doing Business In and As Living Systems*, Chapter 8, "The Story Harvest". The Luminous Library.

24. Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*. Harvard University Press (verified in Open Library; not reopened).

25. Fornell, C., & Wernerfelt, B. (1987). Defensive marketing strategy by customer complaint management: A theoretical analysis. *Journal of Marketing Research*, 24(4), 337–346. https://doi.org/10.1177/002224378702400401 (The quoted word is from the abstract.)

26. The Gas Safety (Installation and Use) Regulations 1998 (Great Britain), regulation 36, require landlords to arrange an annual gas safety check; stated from general knowledge, the regulation text not reopened.

27. National Indicator 14, "Reducing avoidable contact", English local authorities, 2008 to 2010, read through council committee papers searched 30 Oct 2026 (see the notes to Chapter 1).

28. Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67–90. https://doi.org/10.1016/0149-7189(79)90048-X (Paraphrased, not quoted; the text was not reopened.)

29. Anderson, E. W., Fornell, C., & Mazvancheryl, S. K. (2004). Customer satisfaction and shareholder value. *Journal of Marketing*, 68(4), 172–185. https://doi.org/10.1509/jmkg.68.4.172.42723 (The quoted phrases are from the abstract.)

30. Sources for this chapter: The practices and their sources are those of Chapters 1 to 5: Seddon on value demand and failure demand, Gans, Koole and Mandelbaum on contact-centre measurement, Griffin and Hauser on self-selection in satisfaction measures, Hirschman, Best and Andreasen, Singh, and Fornell and Wernerfelt on exit, voice and complaints, the service recovery literature and Kano's categories, and the measurement literature of Chapter 5. Full references are in those chapters' notes; Santa Anna, A. *Conscious Governance Models*, "Principle 3: Distribute the Sensing" (the Sensing Circle) and "Principle 2", the case of the healthcare network. The Luminous Developmental Canon; Santa Anna, A. *Luminous Entrepreneurship: Doing Business In and As Living Systems*, Chapter 8, "Customers as Co-Creators" ("The Customer Co-Creation Toolkit": "The Open Invitation" and "The Thank You Call"). The Luminous Library; Santa Anna, A. *Organizational Soul Work*, "Principle 4: Name the Pattern, Not the Person". The Luminous Developmental Canon; The time estimates in this chapter are planning guides from practice, not findings of a study; adjust them to the organisation.

31. Sources for this chapter: The sources behind these pages are given in full in the notes to Chapters 1 to 5: Seddon (2003, 2008), Ohno (1988) and Jackson, Johnston and Seddon (2008) for pages 1 to 3; Gans, Koole and Mandelbaum (2003), Deming (1986) and Griffin and Hauser (1993) for pages 4 and 5; Hirschman (1970), Best and Andreasen (1977), Singh (1988) and Fornell and Wernerfelt (1987) for pages 6 and 7; Tax, Brown and Chandrashekaran (1998), Smith and Bolton (1998), de Matos, Henrique and Rossi (2007), Michel and Meuter (2008) and Kano and colleagues (1984) for pages 8 and 9; Campbell (1979) and Anderson, Fornell and Mazvancheryl (2004) for page 10; Santa Anna, A. *Luminous Entrepreneurship: Doing Business In and As Living Systems*, Chapter 8, "The Customer Co-Creation Toolkit" ("The Open Invitation"). The Luminous Library; The times given are planning guides from practice, not findings of a study.

The Customer Already Told You

Ammanuel Santa Anna · Luminous Prosperity Inc. · Haute Lumière

Set in Cormorant Garamond and Manrope.

Read the Business Library at business.hautelumiere.life/library.

Luminous Library Licence 1.0 · hautelumiere.life/library-licence

REQUEST A CONVERSATION

business.hautelumiere.life/book/?kind=gainshare

The Business Library

business.hautelumiere.life/library/

The packages

hautelumiere.life/pricing/