

EARNED · CHAPTER TWO

The Signal That Has to Hurt

Chapter Two of Earned

Ammanuel Santa Anna

Santa Anna, A. (2026). The Signal That Has to Hurt. In Earned (ch. 2). Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989331
Luminous Library Licence 1.0 · business.hautelumiere.life/library/earned

THE LICENCE

The Signal That Has to Hurt
Chapter Two of Earned

Earned · Chapter Two. 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.23989331 · this record's page: business.hautelumiere.life/library/earned/03-the-signal-that-has-to-hurt/

Please cite as: Santa Anna, A. (2026). The Signal That Has to Hurt. In Earned (ch. 2). Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989331

Chapter Two of Earned. The whole book, with every chapter, is at business.hautelumiere.life/library/earned (DOI https://doi.org/10.5281/zenodo.23989328).

A NOTE BEFORE YOU READ

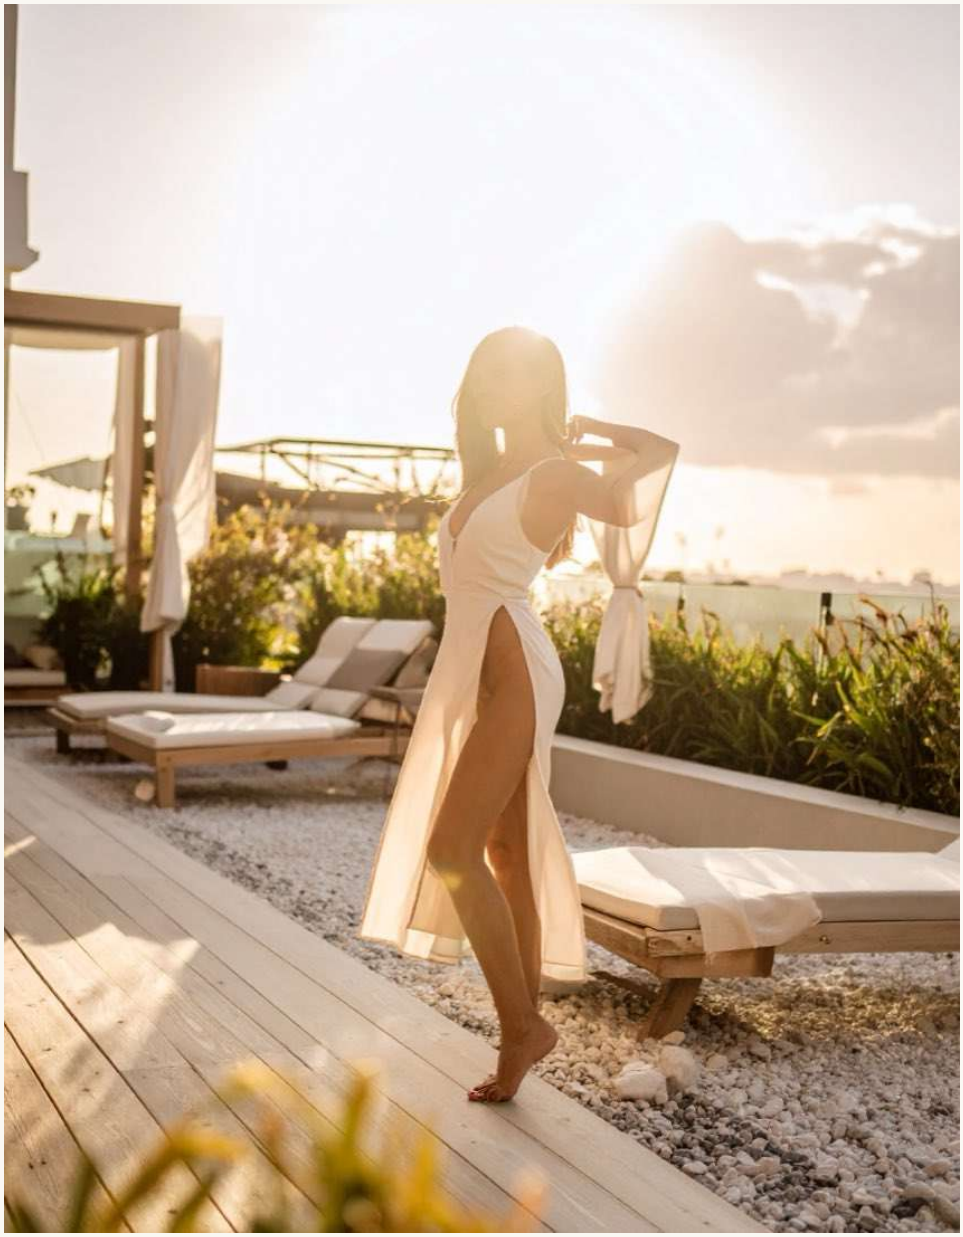
This work is offered for education and leadership development. It is not legal, financial, accounting, investment, medical or psychological advice; take the advice of a qualified professional who knows your circumstances. Companies and events are described from the public record as the author understands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

ABSTRACT

Take the standard trust apparatus a company assembles and run every piece of it through the test. Brand advertising. A national campaign about how much we care. Expensive — genuinely, painfully expensive. But could the unscrupulous competitor buy it? Yes, at exactly the same rate card. Nothing about the purchase order is harder for a company that intends to defraud you. There is a serious economic literature arguing that advertising spend can signal quality for repeat-purchase goods — burn money publicly, and you have credibly told the market you expect enough repeat business to earn it back — but notice how narrow that condition is. It requires that your revenue depends on people coming back, that quality is discoverable on first use, and that the customer reasons through the inference. Take away any one of those and the ad is just a purchase. We will return to what happens when the condition fails, because that failure has a name and a spectacular recent example.

CHAPTER TWO

The Signal That Has to Hurt



The Signal That Has to Hurt

HAUTE LUMIÈRE · EARNED · CHAPTER TWO

Running the test on the toolkit you already own

Take the standard trust apparatus a company assembles and run every piece of it through the test.

Brand advertising. A national campaign about how much we care. Expensive — genuinely, painfully expensive. But could the unscrupulous competitor buy it? Yes, at exactly the same rate card. Nothing about the purchase order is harder for a company that intends to defraud you. There is a serious economic literature arguing that advertising spend can signal quality for repeat-purchase goods — burn money publicly, and you have credibly told the market you expect enough repeat business to earn it back — but notice how narrow that condition is. It requires that your revenue depends on people coming back, that quality is discoverable on first use, and that the customer reasons through the inference. Take away any one of those and the ad is just a purchase. We will return to what happens when the condition fails, because that failure has a name and a spectacular recent example.

Values pages, manifestos, founder posts. Free. Structurally free. The most cynical company in your sector can publish a more moving one than yours, because it is unconstrained by having to mean it.

Security certifications. Here it gets interesting, because these do cost money — a SOC 2 Type II is real work, real auditor hours, real remediation. But apply the test properly. Can your least scrupulous competitor obtain one? In most enterprise software categories the honest answer is: yes, routinely, and they have. The auditee selects the auditor, pays the auditor, and scopes the systems under examination. The pass rate in the segment approaches unanimity. A certificate everyone in your category holds is not a signal; it is a licence. It tells the buyer you cleared the floor. That is a real and useful thing — it is not nothing to be above the floor — but it does not separate you from the company sitting next to you at the trade show, which is precisely what a signal is for.

A no-questions-asked refund. Now run it. Can the unscrupulous competitor offer it? Only if it is willing to eat the cost, and the cost lands asymmetrically: the company selling something that disappoints will have far more refunds claimed against it than the company selling something that satisfies. The promise is cheap for the good product and ruinous for the bad one. That is a separating signal — the structure is the same as the peacock's tail, and it works whether or not a single customer consciously reasons it through. It also has a bill attached, which is why so few companies write it without a clause. L.L. Bean carried an unconditional lifetime guarantee for roughly a century, and it was one of the most powerful trust assets in American retail. In February 2018 the company ended it, limiting returns to one year with proof of purchase, and said plainly why: too many people were returning items bought at yard sales, or worn for a decade, and the abuse had made the promise unaffordable. Note both halves of that. The signal worked *because* it was expensive, and it ended *because* it was expensive. There is no version where you get the credibility without the bill. If someone offers you one, they are selling you the counterfeit we come to shortly.

Published uptime and failure data. This passes, and it passes hardest when the numbers are bad. On 2 July 2019 Cloudflare pushed a single regular expression into a firewall rule that backtracked catastrophically

and consumed CPU across its entire global network; a large fraction of the internet returned errors for about half an hour. Cloudflare published the post-mortem the same day, in engineering detail, naming the regex, explaining that the change had gone out globally rather than in stages, and describing the deployment process that had permitted it. Two years earlier, GitLab had done something more extreme: during a late-night incident an engineer deleted the wrong database directory, and in the recovery GitLab discovered that several of its backup and replication mechanisms were not actually working. The company livestreamed the recovery, kept a public running document as it went, and afterwards published an account stating how many hours of customer data were unrecoverable.

Ask the test question of those. Could the negligent competitor publish that? It could publish *a* post-mortem. It could not publish that one, because that one contains the specific, checkable, damaging admissions that a negligent company's lawyers would never clear and whose contents its own engineers would rather bury. The cost of the disclosure falls on the disclosing company in proportion to how bad the underlying truth is — which is exactly the asymmetry the test is looking for. And a curious thing follows: publishing your failure rate makes the failure rate itself expensive to you, which is an incentive to reduce it. The signal is not merely evidence of the machinery. Over time it becomes part of the machinery.

Two companies that paid the bill

In December 2009 Domino's Pizza went on television and told the United States that its pizza was not good. Not in the euphemistic way brands say such things. The campaign opened with focus-group footage of ordinary customers saying the crust tasted like cardboard and the sauce tasted like ketchup, and internal footage of executives reading comment cards that said Domino's pizza was the worst they had ever eaten. Then the president of the company, Patrick Doyle, said on camera that they had heard it, that the recipe had been thrown out, and that the pizza had been rebuilt from the crust up.

Run the test. Could a company that had *not* reformulated have run that ad? Consider what happens if it does. The advertisement is an enormous, nationally broadcast invitation to re-test the product — it recruits millions of lapsed customers to place one more order specifically to see whether the claim is true. If the pizza is unchanged, the company has spent a fortune to remind the entire market of its worst attribute and then confirm it. The ad is not merely expensive; it is *self-detonating for the liar*. That is why it carried information, and it is why the ad was received as sincerity rather than as marketing. The market response was the largest domestic same-store sales jump the chain had recorded, on the order of fourteen percent in the following quarter — a move that essentially does not happen in mature restaurant chains.

The lesson is routinely mistold. It is repeated in brand circles as *honesty sells*, or *vulnerability is the new authenticity*, which is precisely the free sentence the unscrupulous competitor can also say. What Domino's actually did was construct a situation in which its own claim would be verified by the market within weeks, at scale, with the company's revenue riding on the verdict. It made the claim *checkable and the check immin-*

ent. Any company can be vulnerable in a video. Only a company that has genuinely changed the product can afford to hand the audience a test and a deadline.

Backblaze arrived at the same structure from the opposite direction — not repairing a reputation but building one from nothing, in a category where the buyer cannot inspect the thing being bought. Since 2013, the company has published quarterly statistics on every hard drive in its data centres: manufacturer, model number, drive-days observed, failures, annualised failure rate. Not a summary. The underlying data set, downloadable, model by model, including the years in which a particular Seagate three-terabyte model failed at several times the fleet average and Backblaze said so in public with the part number attached.

Look at who pays for that. Backblaze sells backup storage; it does not sell drives. Every quarter it hands its own suppliers' competitors a free, credible, independently-usable marketing asset, and hands its suppliers a public embarrassment, at a moment when it depends on those suppliers for procurement and pricing. It also exposes its own fleet composition and its own operational choices to second-guessing. The competitor who wants to match this cannot simply publish nicer numbers, because the data set is longitudinal and internally consistent — the drive-days accumulate quarter over quarter, models age in view, and a fabrication has to remain coherent against physics and against every other operator's experience for years. You cannot start this signal today. That is what makes it worth something: the cost is not the publication, it is the decade of not having stopped.

The counterfeit

Wherever a costly signal reliably produces a return, a cheap imitation of its surface appears and competes with it, and the imitation usually wins on distribution because it is cheaper to produce.

The clearest current specimen is the published cost breakdown. Everlane built a brand on what it called Radical Transparency: for each garment, a diagram showing the cost of materials, hardware, labour, duties, and transport, next to the company's price and a comparison to what a traditional retailer would charge. It is arresting the first time you see it, and it feels like the disclosure chapter of this book arriving early.

Now run the test. Could the least scrupulous competitor publish an identical breakdown tomorrow? Yes — and here is the sharp part — it could publish a *better* one. The numbers are self-reported, unaudited, and structurally unverifiable by the reader, who has no way to check a labour cost in a factory they cannot name and no standing to demand the ledger. Nothing in the disclosure creates exposure. Nothing in it can be falsified by a customer. And the content is chosen: a breakdown of unit costs is an argument that the price is fair, which is a marketing claim wearing a spreadsheet's clothes. It reveals the part of the operation the company was already comfortable revealing. Compare it to Backblaze's failure table, which reveals the part that costs the company something to reveal, and which its suppliers actively wish it would stop.

This is the moment transparency stopped being an expense and became a category. Once "transparent" is a market position, the incentive is to produce the *appearance* of disclosure at minimum cost, and the appearance is cheap: infographics, a supply chain map, a page called Our Standards. The tell is always the same and you can apply it without knowing anything about the industry. Real disclosure has a victim inside the company — a team whose numbers look worse, a product that sells less, a supplier who is angry, a lawyer who objected. If you read a transparency initiative and cannot identify who inside the building lost, you are reading marketing.



Why the certificate rots

HAUTE LUMIÈRE · EARNED · CHAPTER TWO

Why the certificate rots

There is a specific and predictable way that third-party verification decays, and it is worth understanding structurally because it is happening right now in several categories that currently feel solid.

The mechanism is issuer-pays. When the party being rated pays the rater, the rater's revenue depends on the rated party's satisfaction, and satisfaction correlates with favourable ratings. Nothing about this requires anyone to be corrupt. It requires only that the rater compete for business, that ratings be shoppable, and that time pass.

The canonical proof is the credit rating agencies before 2008. The agencies had moved decades earlier from investor-pays to issuer-pays, and by the mid-2000s the arrangers of structured mortgage products could take a proposed deal to multiple agencies, learn what each would rate it, and place the business accordingly. The agencies stamped triple-A — the same designation carried by sovereign debt — on enormous volumes of securities backed by loans that would default within eighteen months. Congressional investigators later surfaced internal messages from inside one of the agencies in which analysts joked that a deal could be structured by cows and they would rate it. Then most of those ratings were downgraded, many by ten notches or more, in a matter of months. The seal that the entire global financial system had used as a substitute for its own analysis had, under issuer-pays and competitive pressure, converged to worthlessness while retaining its full apparent form.

The consumer-web echo is smaller and instructive. TRUSTe sold a privacy seal that appeared on thousands of websites; in 2014 the Federal Trade Commission alleged it had failed to conduct the annual recertifications it promised in more than a thousand instances over seven years, and the company settled. The seal had been, for those sites, an image file.

The most telling case is not a rotted certificate but an abandoned one. Etsy certified as a B Corporation in 2012, when it was private, and the certification was central to its identity as a different kind of commerce company. It went public in 2015. In 2017, facing recertification, it declined — maintaining the status would have required restructuring as a public benefit corporation, and by then the company was under activist pressure, had changed chief executives, and had cut staff. Look closely at the timing, because it inverts the usual reading. The certification was cheap when Etsy was private and controlled by people who wanted it, and it became genuinely expensive at the exact moment it would have carried real information — the moment public shareholders would have had to be told that the company had bound itself. And at that moment, it was dropped. The signal's cost and its information content rose together, which is what the theory predicts, and the company paid the cost right up until the point where it was real.

The relocation

Here is the thing the test does to you once you have run it a few times, and it is not what most people expect.

You begin by using it as a filter on language — which claims to keep, which to delete. But every claim that survives has the same shape, and the shape is not a sentence about your character. It is a structure that will hurt you if you misbehave. The refund policy that costs you money if the product disappoints. The published failure rate that costs you enterprise deals if it worsens. The commitment made on a stage in front of the market with a verification deadline attached. In each case the company has not asserted its integrity. It has arranged to be punished for the absence of it, and then made that arrangement visible.

Which means the design problem is not character at all. Trust is not built by being trustworthy; it is built by placing yourself where being untrustworthy would be expensive. You are not trying to become the kind of company that keeps its word. You are trying to become the kind of company that would be *ruined* by breaking it — and then making that exposure legible to the people deciding whether to believe you.

This relocation is the most practically useful move in the book, because character does not survive contact with a bad quarter and exposure does. Intentions are unobservable to your customer, unenforceable by your board, and unreliable in your own successor, who has not read your founding memo and is under pressure you never faced. Exposure is observable, enforceable, and sits in the structure after you have left the chair. Chapter one said trust is a prediction customers make about what you will do when your interests and theirs come apart. Predictions are made about incentives, not about souls. Give them incentives to look at.

And it explains the otherwise baffling asymmetry that every operator has noticed and few can account for: why an enormous, beautiful, sincere brand campaign moves nothing, while one paragraph on a status page admitting exactly how much data was lost moves a great deal. The campaign is a report on your intentions from an interested party. The paragraph is a hostage.

The failure mode: costliness with no information in it

Now the edge past which this inverts, and it inverts hard, because a half-understood version of this chapter is more dangerous than not having read it.

The failure mode is theatrical costliness: spending real, verifiable, enormous money in a way that carries no information about the thing the customer actually needs to predict. Remember the word we planted earlier. The cost must be *differentially* higher for the company you don't want to be confused with. Money that both types can spend equally is not a signal; it is an expense with an audience.

FTX bought a Super Bowl advertisement in February 2022, starring Larry David, telling viewers not to miss out. It had already paid for the naming rights to the arena in Miami on a nineteen-year agreement worth well over a hundred million dollars. Both expenditures were real, verifiable, and vast — a viewer could confirm with their own eyes that this company had spent more than a struggling one could. And nine months later the company collapsed amid allegations that customer deposits had been used to fund an affiliated trading firm, and its founder was subsequently convicted of fraud.

The spend proved something. It proved FTX could spend. It proved nothing whatsoever about whether customer assets were segregated, which was the one question in the room, because a company misappropriating customer deposits can afford a Super Bowl ad *more easily* than an honest one — it is spending money that does not belong to it. The cost asymmetry ran backwards. That is theatrical costliness in its pure form, and the tell is available in advance without any inside knowledge: ask what specific claim the expenditure would be *falsified by*. If there is no answer — if the money would have been spent identically whether the underlying claim was true or false — the expense is decoration regardless of its size.

The same failure hides in quieter places. A sponsorship. A flagship office. A conference stage. A very large security team whose findings never reach a customer. Real money, competently spent, telling the market only that the money existed. If you are going to spend it, attach it to something falsifiable.

What to do on Monday

Open three pages: your homepage, your pricing page, your security or trust page. Take the first ten claims you find — not the headlines, the claims, the sentences that assert something about how you will behave.

Beside each one, write a single column, and give it a heading that is a question rather than a category: *what would this cost us if it were false?*

Most of your rows will come back blank, and the blank is the finding. A blank means the sentence is free — free to you, free to the company you least want to be compared to, and therefore invisible to the reader as a distinguishing fact. It has been occupying premium space on your most-read pages while doing no work at all. Delete it, or convert it.

Conversion is the more interesting half, and it is not a writing exercise. *We take your data seriously* converts into a published breach-notification commitment with a stated maximum number of hours and a named executive who owns the clock, which costs you the ability to manage the news cycle. *We stand behind our product* converts into a refund with no conditions and no window, which costs you money on your worst SKUs and will tell you, within one quarter, which those are. *We're transparent about performance* converts into last quarter's actual numbers, including the bad month, which costs you the deal you would have won on the average. In every case the conversion moves a sentence out of the marketing budget and into the

P&L, and that migration is the entire point — the claim only started carrying information at the moment it began appearing as a cost somewhere in your business.

Expect to be told that this is reckless, and expect the objection to arrive from your own general counsel, who will be doing their job correctly. The strongest version of their argument deserves a straight answer: specificity creates liability, and a published number can be used against you in a way a vague assurance cannot. That is true, and it is not a bug. That is the mechanism. The exposure your counsel is trying to eliminate is the identical exposure your customer is trying to detect. You cannot keep one without the other. What you can do — and what the rest of this book is about — is choose which exposures to accept deliberately, price them, and build the machinery that keeps you on the right side of them, rather than absorbing them by accident on your worst day.

Do the ten rows this week. Ten sentences, one column, one honest number or one blank. It will take an afternoon, it will embarrass you, and it will produce the first genuine inventory anyone at your company has ever made of what you are actually promising — which turns out to be a far larger and stranger surface than the three pages you just audited, and is where we go next.

YOUR PAGE

Chapter Two, carried forward. *Which decision in front of you this week is a deposit, and what will it cost you to make it?*

The Briefs

Chapter Two, in instruments



The Counterfeit Test: Four Words That Kill a Marketing Claim in a Meeting

HAUTE LUMIÈRE · EARNED · BRIEF 2.1

BRIEF 2.1

The Counterfeit Test: Four Words That Kill a Marketing Claim in a Meeting

Somebody has just put a line on the screen — *We put customers first* — and the room is nodding, because nobody has been given a reason not to. The line will cost about eleven thousand dollars to place and it will move nothing.

The move is to say four words out loud, every time, before the line is approved: "What's the liar's cost?"

Meaning: what would it cost a competitor who does *not* put customers first to print that identical sentence tomorrow morning? If the answer is zero — and for that sentence it is zero — the claim carries no information about you. It cannot, because it does not separate you from the company you are trying to be distinguished from. Anything a defector can copy for free is not a signal, it is decoration, and the customer's ear knows this even when their conscious mind does not. They have heard the sentence from the firm that shipped them the broken order.

The test works because information is carried by *differential cost*, not by content. A claim discriminates only when the bad actor's copy of it is expensive enough to be irrational. Hyundai's ten-year powertrain warranty in 1999 was not a statement about quality; it was a bill that only a company whose engines lasted could afford to pay. The sentence "our engines last" cost nothing and persuaded nobody. The warranty cost real money in every failure and persuaded a market. Same claim, two prices, two outcomes.

The condition is that the cost must be *conditional on lying* — visible, verifiable, and paid out precisely when the claim turns out false. A cost you would incur either way, like an expensive ad shoot, is just a cost. It signals a budget, not a truth.

The failure mode is using the test as a veto and stopping there. Teams that learn this question start killing copy without replacing it, and the site goes quiet while the competitor's cheap talk keeps running. The test is worthless unless it is paired with the follow-up: *what would we have to put at risk to make this claim true and expensive?* A team that only subtracts will conclude, wrongly, that trust is unbuildable.

Today: take the three claims on your homepage that the team is proudest of. Next to each, write the sentence a firm doing the opposite would have to write to copy it, and what that copy would cost them. If all three cost nothing, you have found the real state of your marketing, and it is the same as everyone's.



Signal Audit: Scoring Every Trust Claim on Your Homepage by What a Liar Would Pay to Copy It

HAUTE LUMIÈRE · EARNED · BRIEF 2.2

BRIEF 2.2

Signal Audit: Scoring Every Trust Claim on Your Homepage by What a Liar Would Pay to Copy It

Your homepage has somewhere between eight and forty trust-bearing elements on it: badges, testimonials, statistics, promises, logos, a security page link, a founder's note. Nobody has ever scored them, and the ones taking the most vertical space are usually the ones carrying the least.

The move: build a one-column spreadsheet of every such element and score each 0–3 on a single axis — *the cost, to a firm that does not deserve the claim, of displaying the identical element.*

0 — free. Anyone can write it. "Trusted by thousands." "Bank-grade security." "We care about your data." A values page.

1 — cheap. A fee, a form, a logo licence. Most badges. Most memberships. Most self-selected testimonials.

2 — costly but survivable. A real audit with real findings. A named customer who agreed to a recorded reference call. A published metric a rival could match by spending.

3 — unaffordable to a defector. A refund the customer triggers unilaterally. A published failure rate that will look bad in a bad quarter. A contractual penalty that pays out automatically. A disclosure that hands ammunition to a competitor.

The mechanism is separation. A claim informs only if the pooling equilibrium breaks — if the honest firm can afford it and the dishonest firm cannot. Your 0s and 1s are pooled: everyone displays them, including the frauds, so a rational customer correctly assigns them near-zero weight and you have paid for pixels. Your 2s and 3s are where the entire persuasive load of the page actually sits. Most sites score a total of four across the whole page and wonder why conversion is flat.

The audit only works if you score the *element as displayed*, not the underlying reality. You may genuinely have excellent security; if the page says "bank-grade security," that sentence is still a 0. The customer is scoring the artifact in front of them, not your intentions.

The failure mode is score inflation by the person who commissioned the element. Whoever bought the badge will argue it is a 2. Fix this structurally: have someone outside marketing score it, or better, score it as if you were the competitor — ask them what they'd have to pay, not what you paid.

Today: score the page, sum it, and then move your single highest-scoring element above the fold. That reordering alone is usually the largest free improvement available to a site, and it takes twenty minutes.

BRIEF 2.3

Warranties as Wagers: Pricing a Guarantee That Actually Says Something

Finance has just been handed a proposal for a money-back guarantee and has priced it the way they price everything: worst case, everyone claims, we lose the revenue line. The guarantee dies in that meeting, every time, because it was framed as an exposure rather than as a purchase.

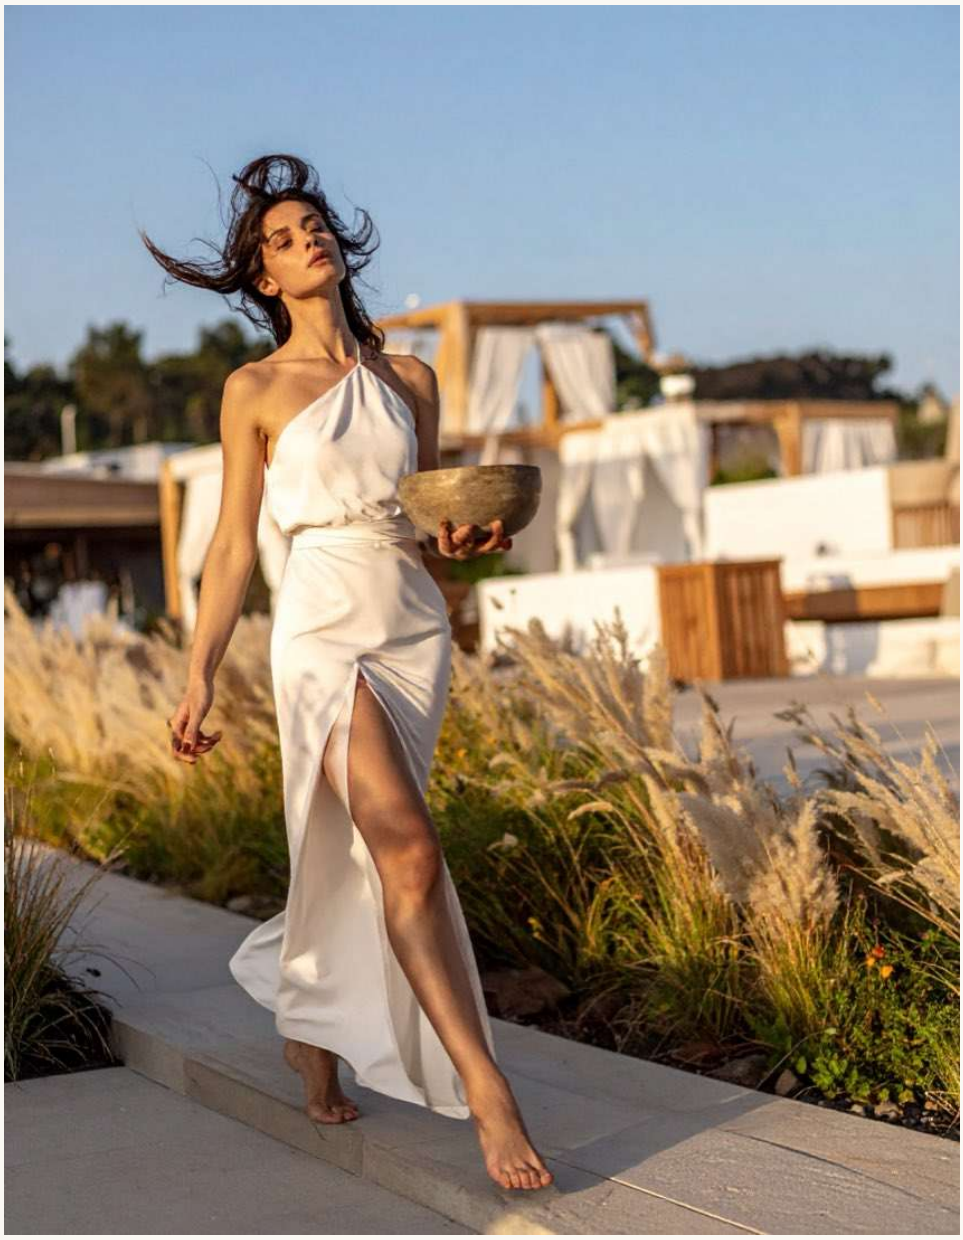
The move: price the guarantee as a *wager you expect to win*, and set its terms so that a worse version of your company would lose it.

The arithmetic is one line. Expected cost equals your true failure rate times the payout times the claim rate. If you sell a $400 product, genuinely fail 2% of the time, and 80% of failures claim, the guarantee costs you $6.40 per unit — 1.6% of price. Now ask the real question: what does the same guarantee cost a competitor failing at 9%? $28.80, or 7.2% of price, on margins that are probably thinner than yours. That gap *is* the signal. You are not buying insurance; you are buying an asymmetry that your competitor cannot afford to match and that the customer can infer without reading a word of your copy.

This tells you how to set the terms. Push the payout up and the friction down until the wager still clears your margin but breaks theirs. A guarantee that costs you 1.6% and them 7.2% is worth ten times a guarantee that costs you 0.3% and them 1.4%, even though the second looks safer on the spreadsheet. The conditions are that your failure rate must be genuinely known — a guarantee written on a guessed defect rate is a wager placed blind — and that claiming must be unilateral. Any approval step you insert transfers the cost back to the customer and collapses the asymmetry.

The failure mode is the ratchet. Guarantees get written in a confident year and quietly narrowed in a hard one, and the narrowing is read, correctly, as new information about the product. L.L. Bean's retreat from a lifetime return policy to one year in 2018 cost more in interpretation than it saved in returns, because the withdrawal of a signal is itself a signal, and a legible one. Do not offer terms you would rescind under pressure. Offer thinner terms you will hold through a bad quarter.

Today: pull your actual failure rate for one product line — the real number, from support tickets, not the QA target — and calculate what an unconditional 60-day guarantee would cost you and what it would cost the competitor you lose deals to. Bring both numbers to finance. The second one is the argument.



Why Your SOC 2 Persuades Procurement and Nobody Else

HAUTE LUMIÈRE · EARNED · BRIEF 2.4

BRIEF 2.4

Why Your SOC 2 Persuades Procurement and Nobody Else

The badge went on the site in March, the security page got a new section, and the marketing team is puzzled that self-serve conversion did not move at all — while enterprise deals started closing faster. Both effects are correct and they have the same explanation.

The move is to stop treating the SOC 2 as a trust signal and start treating it as what it is: a *liability transfer instrument for a specific buyer*, and to place it accordingly — deep in the sales motion, not on the homepage.

Here is the mechanism, and it is uncomfortable. You chose the auditor. You paid the auditor. You wrote the system description that defines the scope, and you selected which trust services criteria were in play. A firm with weak controls can obtain a clean Type II report by scoping narrowly and choosing an accommodating firm, and the market knows this, which is why the marginal cost to a defector is *low* — a fee and some paperwork — and therefore why the badge scores near the bottom on any honest signal audit. It is close to pooled: everyone selling to enterprise has one.

So why does it work on procurement? Because procurement is not asking "is this vendor safe." Procurement is asking "can I be blamed." The report is a defensible artifact in a file — it converts a personal career risk into an institutional process followed. That is a real job, worth real money, and the report does it well. It just is not the job of persuading a skeptical human that you will not lose their data.

The condition under which the report *does* carry information is when you publish what is normally hidden: the scope, the exceptions, the auditor's name, the period, and the management responses. Exceptions are costly to disclose. A defector will not print theirs. That is where the signal lives — in the part everyone redacts.

The failure mode is substitution. Teams get the report and stop building the expensive signals, because the compliance line item feels like it discharged the obligation. Two years later they have three certifications, an incident with no public postmortem, and a churn problem nobody can name. Compliance is a floor you must be standing on; it was never the thing you were standing *for*.

Today: find the exceptions section of your most recent report and ask your security lead one question — what would it cost us to publish this page, verbatim, with our responses? If the answer is "nothing much," publish it this week. If it is "quite a lot," you have just located the real work.



Publishing Your Failure Rates: The Backblaze Move, Adapted to Your Business

HAUTE LUMIÈRE · EARNED · BRIEF 2.5

BRIEF 2.5

Publishing Your Failure Rates: The Backblaze Move, Adapted to Your Business

Backblaze, a backup company with no obvious reason to do so, began publishing quarterly statistics on how often the hard drives in its data centers die — by manufacturer, by model, with the failure rates

named. Drive vendors were not pleased. The reports became the most-linked thing the company produced and are still cited by people who have never used the product.

The move: pick the one number your industry conventionally hides because it makes vendors look bad, measure it honestly, and publish it on a fixed cadence — including in the quarters where it is worse.

The mechanism has two parts and both matter. First, cost: a competitor with a bad number cannot copy this, because copying means printing their bad number. The disclosure is self-damaging in expectation, which is exactly what makes it credible. Second, and more powerful, the *cadence* is the commitment. A one-off report is a marketing decision. A standing quarterly obligation means you have pre-committed to publishing during the quarter it goes wrong, and everyone reading understands that you have. The schedule is the hostage.

Adaptation is a matter of finding your analogue. An agency publishes on-time delivery and budget variance across all engagements, not selected ones. A clinic publishes complication rates. A marketplace publishes dispute rates and how they resolved. A SaaS company publishes support first-response and resolution times, and every incident's postmortem, including the ones caused by its own bad deploy. The test for a good candidate: your sales team should flinch.

Conditions: the number must be defined publicly and stably, computed from a system you do not control by hand, and reported without survivorship filtering. Cherry-picked denominators are detected quickly and convert an asset into a scandal.

The failure mode is the quiet redefinition. Two bad quarters in and someone proposes excluding "planned maintenance" from downtime, or measuring response time from first human touch rather than ticket creation. This is worse than never publishing, because you have now taught an attentive audience that your numbers move when they are inconvenient. If you must change a definition, restate the full history under the new one and say why.

Today: name the number. One meeting, one whiteboard, the question being *what do we know about our own performance that we would be embarrassed to have printed?* Then measure it internally for one quarter before you commit publicly. Publish the definition first, the number second.



Issuer-Pays: Diagnosing Conflict in Any Certification You Are About to Buy

HAUTE LUMIÈRE · EARNED · BRIEF 2.6

BRIEF 2.6

Issuer-Pays: Diagnosing Conflict in Any Certification You Are About to Buy

A vendor is selling you a certification, a seal, a rating, or a "verified" badge, and the annual fee is quoted before the criteria are explained. That ordering is the diagnosis, but here is the full test.

The move: before buying any third-party attestation, answer five questions in writing. Who pays the certifier? Can the payer shop for a better verdict? Does the certifier lose money when it fails a client? Is the failure rate published? Can the badge be revoked, visibly, after issuance?

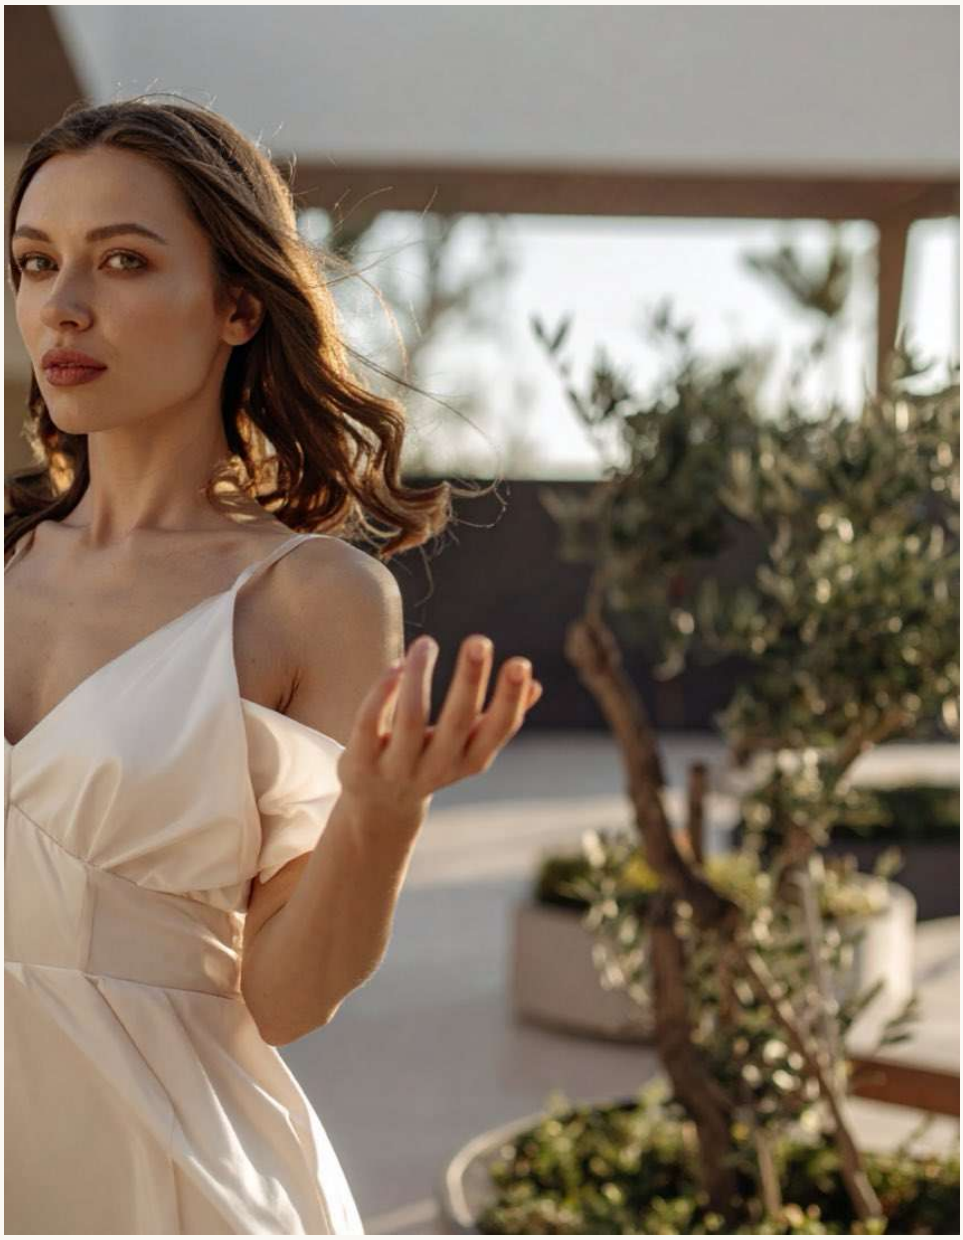
The mechanism is the one that broke the credit rating agencies. Moody's and Standard & Poor's moved from investor-pays to issuer-pays in the early 1970s, and by the 2000s the entity being rated selected and paid the rater and could take its structured-finance business elsewhere if the grade disappointed. No individual analyst had to be corrupt for the aggregate output to be worthless; the incentive gradient did the work. The lesson generalises exactly: *a certifier whose revenue depends on the approval of the certified will, over time, approve.* Not by fraud. By drift — softer criteria, generous scoping, a helpful pre-audit.

So the questions map to defences. Investor-pays or member-funded certifiers (Consumer Reports buys its test units at retail and takes no advertising) have the cleanest incentives and are rare. Where issuer-pays is unavoidable — most safety and security regimes — the mitigations are a published failure rate, visible revocation, and an inability to shop. A certifier that has never failed anyone has told you its price.

Apply the same test to certifications you *hold*. Your badge is being read by sophisticated buyers who run this diagnostic instinctively, which is why a seal with a 100% pass rate adds nothing to your page while a certification with real teeth carries weight far beyond its cost.

The failure mode is cynicism at the wrong altitude. This analysis makes it tempting to conclude all certification is theatre and skip it — but many issuer-pays regimes are load-bearing because the *certifier's own reputation* is a hostage worth more than any single client. Underwriters Laboratories is paid by manufacturers and is still worth having, because a UL failure that killed people would end UL. Ask whether the certifier has more to lose than you do. If yes, the conflict is contained.

Today: take the certification you already display most prominently and search for its published pass rate and its revocation list. If neither exists, you have learned what your badge is worth, and it is roughly the fee.



The Domino's Admission: When Saying the Product Is Bad Is the Cheapest Campaign You Will Ever Run

HAUTE LUMIÈRE · EARNED · BRIEF 2.7

BRIEF 2.7

The Domino's Admission: When Saying the Product Is Bad Is the Cheapest Campaign You Will Ever Run

In December 2009, Domino's ran national advertising in which its own customers said the crust tasted like cardboard and the sauce like ketchup, its chefs read the complaints on camera, and the company said, in

effect: they are right, we reformulated everything. Domestic same-store sales jumped sharply the following quarter — a swing of a kind that conventional campaigns of far greater spend do not produce.

The move: when your product's weakness is already common knowledge, stop defending it and state it more plainly than your critics do — then attach the fix.

The mechanism is cost again, but the cost here is reputational rather than financial. A competitor with an equally bad product cannot copy this ad, because copying it means broadcasting *their* defect to a market that had not fully noticed. The admission is only affordable to a firm that has actually fixed the thing, because the admission raises expectations to precisely the level where a continued failure is fatal. That is what makes it informative. And a second mechanism runs underneath: when the criticism is already circulating, denial costs you the assumption of shared reality with your customer. Saying it first buys that back at a discount, because you were going to pay for the complaint either way.

Three conditions, all necessary. The weakness must already be believed by the market — confessing a secret flaw is not brave, it is a disclosure event with different rules. The fix must be real, shipped, and immediately testable, because you are inviting the entire market to check. And the admission must be specific: "we weren't perfect" is cheap talk wearing a humble face, while "the crust tasted like cardboard" is a sentence a lawyer would have deleted, which is why it works.

The failure mode is confession without repair — the apology tour that ships nothing. Volkswagen in 2015 apologised comprehensively and at length, and the apology bought almost nothing, because the admission was extracted by regulators rather than volunteered, and because the structures that produced the cheating were the thing under question. An admission you were forced into carries no information; everyone knows what forced it. And an admission of a flaw you have not fixed does not lower expectations, it sharpens the audience's attention on the next failure.

Today: write down the one criticism of your product that your support team hears weekly and your marketing has never once acknowledged. Ask whether it is fixed. If it is, say it out loud in your own words before a reviewer does. If it isn't, you have your roadmap, not your campaign.



Hostages, Bonds, and Escrow: Borrowing Contract Design for Trust Signals

HAUTE LUMIÈRE · EARNED · BRIEF 2.8

BRIEF 2.8

Hostages, Bonds, and Escrow: Borrowing Contract Design for Trust Signals

Two firms who do not trust each other close deals every day, at scale, without either one becoming virtuous. Contract design solved this problem decades before marketing noticed it existed, and the instruments are sitting there unused by anyone building consumer trust.

The move: stop writing claims and start posting collateral. Three forms, each doing a different job.

The hostage. You place something you value in a position where your own bad behaviour destroys it. Oliver Williamson's insight was that a hostage need not be valuable to the *receiver* — it need only be costly to the giver. A published quarterly metric is a hostage. So is a founder's name on the guarantee, a public roadmap with dates, a no-NDA policy on customer references. The customer cannot seize it; your failure detonates it.

The bond. A sum posted in advance, forfeited on defined failure. Performance bonds in construction, surety in freight. The consumer analogue is the automatic service credit: not "contact us for compensation," which is a cost you control, but a credit issued by your own monitoring system without a human deciding. JetBlue's Customer Bill of Rights in February 2007 did this — self-imposed, schedule-based payouts for delays the airline itself caused.

The escrow. A third party holds the value until performance is verified, removing your discretion entirely. Source code escrow for enterprise software. Funds held until delivery. Deposits released on acceptance rather than on invoice.

The mechanism common to all three: they move the decision *out of your hands at the moment you would most want it back*. A promise is a claim about your future self. A bond is a constraint on your future self, imposed by your present self, and the customer can verify which one they are looking at by asking a single question — who decides whether this pays out? If the answer is "we do," it is a promise. If the answer is "the contract does," it is collateral.

The failure mode is collateral with a discretion clause hidden in it. Legal will want an exclusion for circumstances beyond reasonable control, and the moment that clause exists, sophisticated buyers reclassify the whole instrument as a promise, correctly. A small bond with no escape hatch outperforms a large one with three.

Today: take your service credit policy — you have one, buried in the SLA — and find out how a customer currently claims it. If it requires them to notice, ask, and be approved, rewrite it to fire automatically from your own uptime data. That single change converts a paragraph nobody reads into a bond.

BRIEF 2.9

Cheap Talk Inventory: Every Sentence on Your Site That Costs Nothing to Say

There is a page on your site — usually /about, sometimes /values — that took four weeks, three drafts, and a leadership offsite to produce, and no customer has ever changed a decision because of it. It is not badly written. It is free to write, and that is the whole problem.

The move: run a deletion audit. Export every customer-facing sentence that makes a claim about your character, competence, or care. Mark each one C if a company doing the opposite could publish it verbatim tomorrow at no cost. Then delete the C's — not rewrite, delete — and see what remains.

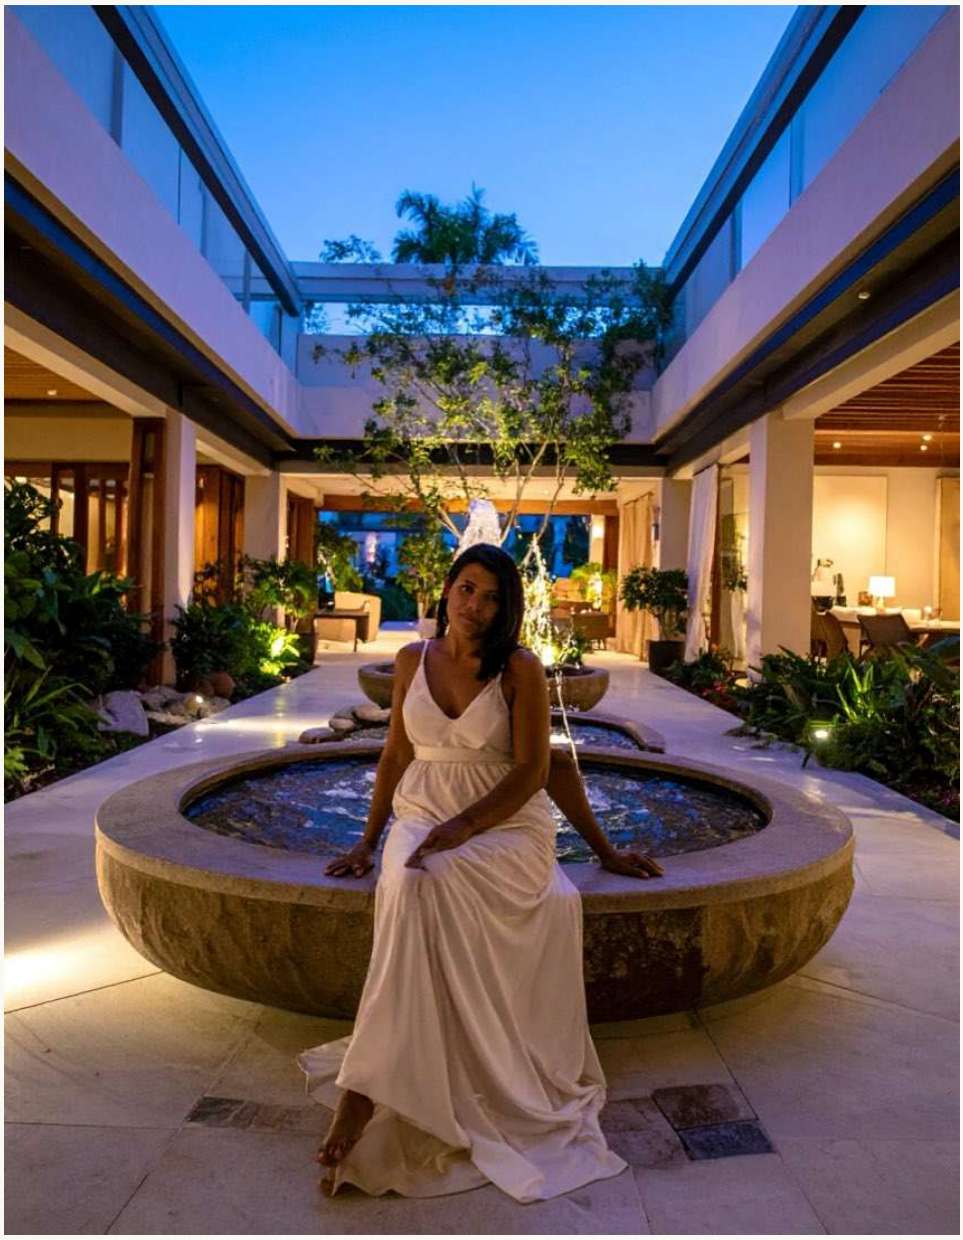
Economists call this cheap talk, and the formal result is bleak: a message that costs the sender nothing and cannot be verified conveys no information when interests diverge, so a rational receiver ignores it entirely. Your customer is doing this without knowing the term. They have been told *we're passionate about quality* by a company that shipped them a broken chair, and they have updated their priors accordingly. The sentence is not neutral now; it is faintly negative, because its presence marks you as the kind of firm that thinks the sentence works.

Cheap talk is not always useless — it coordinates when interests align, which is why "the checkout is on the next page" is fine and "we care about your experience" is not. The test is whether a defector's interests point the same way. If they'd say it too, it carries nothing.

What survives deletion is instructive. Usually: numbers, names, dates, mechanisms, and prices. *Median response time 14 minutes, published weekly. Every plan cancels in one click; here is the click. We do not sell data, and here is the contractual clause that makes that enforceable by you.* This is not a colder register — it reads warmer, because specificity is what respect for the reader sounds like.

The failure mode is deleting the C's and leaving the page empty, which teams then fill with more C's under deadline. Cheap talk is load-bearing *layout*: it holds space that finance never funded a real signal to occupy. So run the inventory with a builder in the room, not just an editor, and treat each deletion as a work item — what would have to be true, and what would it cost, for something with weight to stand here?

Today: open your values page. Count the sentences. Count how many survive the C test. Take that ratio to the person who owns the page — it is the most persuasive number about your marketing that you can produce in fifteen minutes.



The Apology Budget: Why 'Sorry' Reads as Cheap Talk, and What Would Not

HAUTE LUMIÈRE · EARNED · BRIEF 2.10

BRIEF 2.10

The Apology Budget: Why 'Sorry' Reads as Cheap Talk, and What Would Not

The outage lasted six hours, the postmortem is written, and the email going out tomorrow morning opens with *we're deeply sorry* and closes with *we take reliability extremely seriously*. It will cost nothing to send and it will recover nothing, and the second of those facts follows from the first.

The move: give every apology a budget, denominated before the incident, and pay it automatically. Not compensation on request. A pre-committed schedule — this class of failure costs us this much, paid to affected customers without them asking, within this many days.

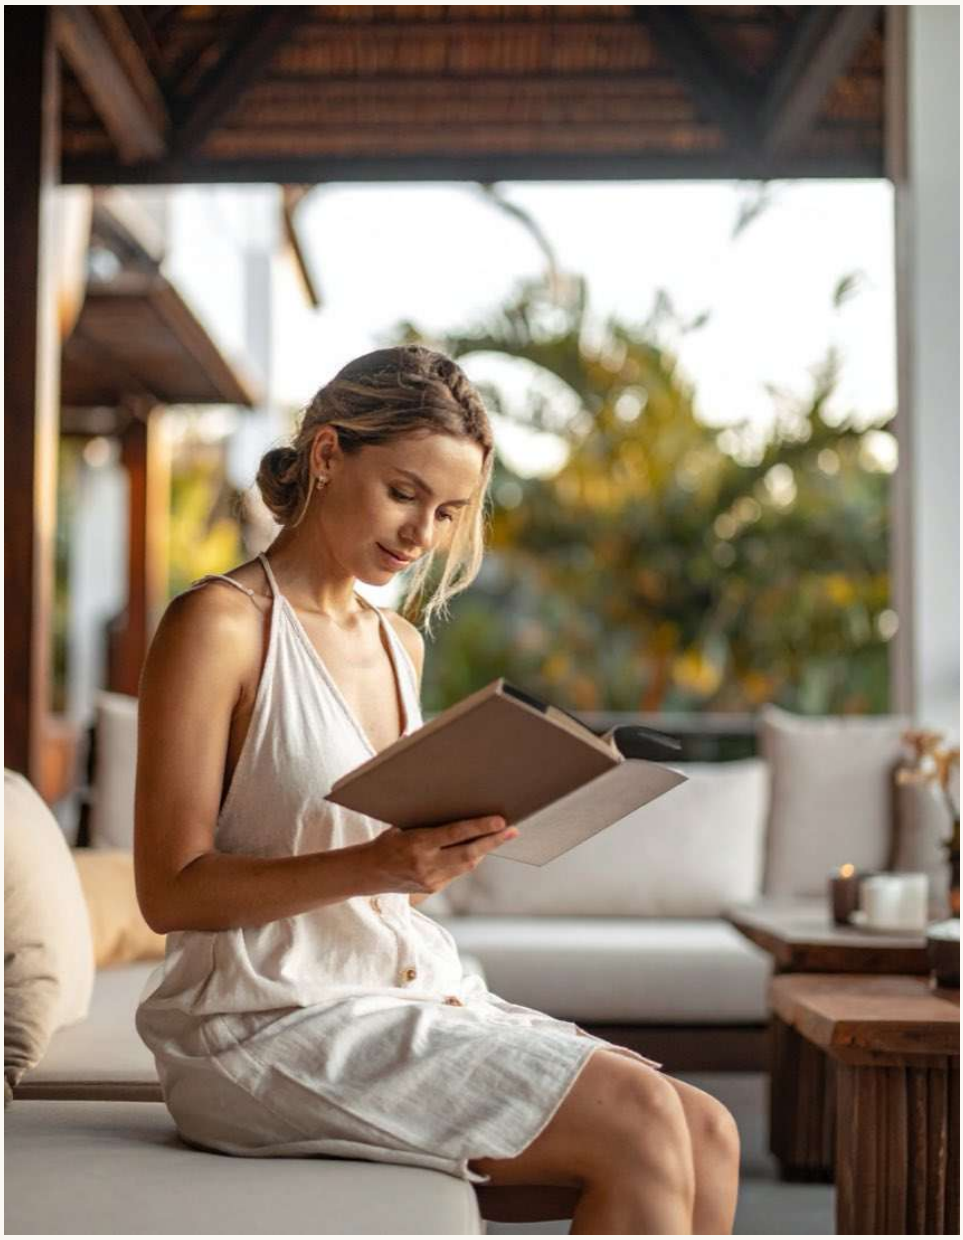
The mechanism is that remorse is unobservable and words about remorse are free, so the customer cannot distinguish a company that is sorry from a company that has a communications team. Both send the same email. What they *can* observe is money leaving your account for no reason other than your own failure, which the indifferent competitor will not do because it is expensive and they do not have to. That is the separation. The apology becomes informative at exactly the point it becomes costly.

Sequence matters as much as size. A payment made before the customer complains says something a payment made after negotiation cannot, because the negotiated one only proves you respond to pressure — which they already assumed. Pre-commitment matters more still: a budget set in advance means you have surrendered the option to be cheap on a bad day, and everyone can see the surrender.

Conditions: the budget must be public, the trigger mechanical, and the payment automatic. If a VP approves each one, you have built a discretionary fund, and discretionary funds shrink under margin pressure precisely when trust is being tested.

The failure mode is buying your way past a structural defect. Repeated payouts for the same recurring failure stop reading as integrity and start reading as a company that has priced its customers' suffering and finds the price acceptable — which is what it is. The budget is only a trust instrument if the failure rate is falling; otherwise it is a subscription to a problem. Track both numbers on the same chart and show them together.

Today: take your last significant incident. Calculate what a pre-committed, automatic credit would have cost you for that event. That number is small — it almost always is, and far smaller than the goodwill the apology email failed to buy. Bring it to the person who signs off on credits and ask for it as a standing line, not a case-by-case exception.



The essays · Chapter Two

HAUTE LUMIÈRE · EARNED · CHAPTER TWO

The Essays

Five assignments for the long evening

ESSAY 2.1

The prompt — If a signal only informs when it is costly, small enterprises that cannot afford the burn of capital to prove integrity appear structurally silenced by incumbents who can flood the zone with cheap

theater mimicking expenditure without risking solvency. The argument must demonstrate that this asymmetry does not condemn the small firm to permanent illegibility but instead forces a distinct separating equilibrium, one where the signal is not the waste of money but the imposition of structural constraint or temporal sacrifice that a low-quality incumbent cannot adopt without violating its own value extraction logic, and you must identify a concrete instrument that allows a firm of limited scale to credibly bind itself to high standards in a way that replicates the separation property of cost without requiring the capital reserves of a large organization.

What a serious answer has to do — You must establish that the mechanism of separation relies on the *impossibility of imitation* rather than the sheer magnitude of expenditure, showing how a small firm can embed a "handicap" into its operations that a malicious actor would find too expensive to maintain relative to the fraud's returns. The essay must identify a specific instrument—such as a governance structure, a revenue model, or a technical architecture—and demonstrate through a named case how it forces a separating equilibrium, while arguing past the cheap answer that the small firm should simply accept lower margins or raise capital to match the incumbent's signaling budget, which misunderstands that trust is built by preventing bad faith, not by outspending it.

Where to look — Examine cooperative enterprises that have adopted "one member, one vote" alongside veto rights for stakeholders, open-source software projects that enforce contribution licenses requiring derivative works to remain open, and specific small-scale service firms that have used "no-guarantee" contracts or personal asset pledges to break the trust deficit. Look also to the history of regulatory capture where small entrants used procedural complexity as a barrier to entry for opportunistic competitors, and to case studies of firms that grew by embedding compliance costs into their pricing model in ways that high-volume fraudsters could not sustain.

The length — 2,500 words minimum.

ESSAY 2.2

The prompt — Third-party certification attempts to substitute for costly signaling by pooling verification, yet the prevailing issuer-pays structure inverts the mechanism of trust: the firm selects and compensates the verifier, creating a dynamic where the verifier's revenue depends on pleasing the issuer, which reproduces the capture the regime claims to resolve. The argument must evaluate whether any existing certifica-

tion framework has durably resisted this inversion, and if the evidence points to chronic failure, you must design a regime in which the verifier's financial survival and legal exposure are tied to the long-term accuracy of the rating, such that the cost of a false positive—measured in downstream liability, lost renewal revenue, or reputation decay—structurally exceeds the immediate revenue from the contract, thereby forcing the verifier to internalize the risk of the signal.

What a serious answer has to do — You must establish that a durable certification regime requires the verifier to hold the "skin in the game" of the rating, demonstrating through a named case how a specific model shifts the cost of error from the public to the certifier, and you must argue past the naive proposal that issuers can simply be mandated to use a pool of verifiers without altering the payment architecture, showing instead that the payment flow must be decoupled from the issuer's discretion and linked to the outcome's verification by a third party or a mutualized risk pool.

Where to look — Analyze the insurance industry's underwriting model, where the insurer bears the loss and thus has a mechanism to accurately price risk, credit rating agencies and the regulatory shifts following the 2008 crisis, specific mutualized audit funds in the food safety sector where costs are shared and liability is pooled, and the history of standards bodies that collapsed when they reverted to issuer-pays dynamics. Look also for case studies of "follow-the-victim" payment structures in consumer protection, where the certifier is liable to the end-user for misrepresentation.

The length — 2,500 words minimum.

ESSAY 2.3

The prompt — The classic economic argument asserts that advertising spend signals confidence in product quality, because a firm selling a lemon cannot recover the sunk cost of advertising through repeat sales, whereas a firm selling quality can; yet this claim collapses when the market permits sufficient churn, when switching costs are negligible, or when the expected profit from a single fraudulent transaction exceeds the amortized cost of customer acquisition. The argument must evaluate this tension by showing precisely where the mechanism holds and where it inverts, naming the conditions under which advertising spend ceases to separate a high-intent firm from a fraudster, and demonstrating that high expenditure signals only the presence of a viable customer acquisition strategy, not the alignment of corporate intent with customer welfare, unless the firm's business model structurally ties revenue to long-term customer success.

What a serious answer has to do — You must establish the mechanism of the signal: advertising separates only when the Lifetime Value (LTV) of a customer is sensitive to quality and when the firm captures a sufficient share of that LTV to amortize the acquisition cost. The essay must identify a named case where high advertising coexisted with systemic fraud because the business model relied on volume and churn rather than retention, and argue past the shallow conclusion that "big ads mean big confidence" by showing that confidence in quality is a necessary but insufficient condition for the signal to hold, requiring the additional constraint that the firm's profit function is convex in reputation.

Where to look — Examine the direct-response marketing industry, particularly catalog businesses and subscription traps, where high ad spend accompanies low retention and high margins on the first sale. Look to the history of the pharmaceutical industry's shift to direct-to-consumer advertising and its correlation with prescribing patterns, the rise of ad-supported news media and its impact on editorial standards, and case studies of subscription services that use aggressive acquisition spending while deliberately obscuring cancellation mechanisms. Also review the economic literature on "churn and burn" models versus retention-based models to find the mathematical threshold where advertising ceases to be a signal of quality.

The length — 2,500 words minimum.

ESSAY 2.4

The prompt — Voluntary disclosure on one dimension reliably deflects scrutiny from another, creating a dynamic where transparency functions not as a signal of trust but as a shield that manages the observer's attention while protecting the mechanisms of value extraction. The argument must demonstrate where this deflection occurs, showing how high-intent firms are forced to disclose dimensions that are costly to falsify or costly to maintain, whereas low-intent firms can flood the zone with verifiable but irrelevant data to satisfy the appearance of openness. You must design a heuristic for a reader to distinguish between genuine disclosure, which exposes the firm to downside risk and invites adversarial verification, and strategic disclosure, which insulates the firm by satisfying regulatory or social demands through metrics that are cheap to produce and irrelevant to the core trust failure, and you must name the specific conditions under which a disclosure regime becomes a tool for capture rather than a mechanism for separation.

What a serious answer has to do — You must establish that the cost of disclosure is the key variable: genuine disclosure imposes a cost of compliance or risk of penalty that a fraudster cannot afford, while strategic

disclosure imposes a cost of curation that is low for a fraudster. The essay must identify a named case where a firm's disclosure of a benign metric correlated with the worsening of a hidden practice, and argue past the assumption that "more data equals more trust" by showing that trust requires the disclosure of *salient* risks, not just *available* data, and that a reader must look for the asymmetry between the ease of verifying the disclosed metric and the difficulty of auditing the omitted one.

Where to look — Investigate the history of ESG (Environmental, Social, and Governance) reporting, focusing on cases where high scores on disclosed metrics coincided with scandals in omitted categories. Look to the analysis of privacy policies and data handling practices, where firms disclose extensive lists of collected data while obscuring the secondary market for that data. Examine algorithmic accountability reports from tech firms, where firms disclose fairness metrics on specific subgroups while obscuring the overall bias of the system. Also review regulatory filings where companies used "safe harbor" disclosures to limit liability for forward-looking statements, and case studies of companies that used "transparency" to preempt stricter regulation by offering the appearance of self-policing.

The length — 2,500 words minimum.

ESSAY 2.5

The prompt — Regulation can be argued as the cheapest trust technology for an entire industry because it forces all participants to internalize the costs of trust failure, raising the floor and reducing the competitive advantage of bad actors who would otherwise undercut honest firms by externalizing those costs; in this view, regulation acts as a separating equilibrium where the marginal cost of compliance is lower for high-intent firms than for low-intent firms, effectively subsidizing trust by making adherence the path of least resistance. The argument must make this case by showing how regulation can level the playing field and reduce information asymmetry for the consumer, and then must construct the strongest counter-argument from the perspective of the single most honest firm in the industry, demonstrating how regulation can be captured to entrench incumbents, how compliance costs can stifle the very innovation that generates trust, and how the presence of a regulatory seal can create a moral hazard that allows low-intent firms to purchase legitimacy through compliance theater, thereby confusing the observer and eroding the distinction between regulated fraud and unregulated integrity.

What a serious answer has to do — You must establish that regulation functions as a trust technology only when the cost of compliance is convex in the degree of intent, meaning that a firm with malicious intent faces exponentially higher costs to comply than a firm with genuine intent. The essay must identify a named case where regulation successfully leveled the field by raising the cost of fraud, and a named case where regulation was captured to exclude small, honest entrants, arguing past the binary view that regulation is either good or bad by showing that the quality of the regulatory design—specifically its susceptibility to capture and its incentive structure for the regulated—is the determinant of whether regulation serves trust or serves the incumbent.

Where to look — Examine the history of healthcare regulation, including the impact of ERISA and HIPAA on small employer-sponsored plans and the trade-offs between security and access. Look to the banking sector and the Basel accords, analyzing how capital requirements affected small community banks versus large institutions. Investigate the fintech sector, where compliance costs have sometimes excluded small, mission-driven entrants while allowing large incumbents to absorb the burden. Review the literature on regulatory capture, the "co-option" of standards by industry groups, and case studies of firms that were destroyed by compliance requirements that were designed by their competitors. Also look to the distinction between "compliance" and "trust" in safety-critical industries, where regulation has saved lives but also created rigidities that delay innovation.

The length — 2,500 words minimum.

The Signal That Has to Hurt

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/