ESSAY · INSTITUTIONS, GOVERNANCE & SYSTEMS

The Invisible Curriculum of Organizations

How organizations really teach — and what it costs when the lesson changes

Lateral House · August 2026 · lateralhouse.com/essays/invisible-curriculum-of-organizations/

How this was made. This essay was AI-drafted. The author directed it, shaping the argument, structure, sourcing and edits, rather than writing each sentence personally.

In 1987, Paul O'Neill had recently taken over as CEO of Alcoa. A worker at one of its Arizona plants, a few weeks into the job, reached into a jammed six-foot mechanical arm and was killed. O'Neill gathered the plant's executives and, as he later recounted the moment, told them: "We killed this man. It's my failure of leadership. I caused his death. And it's the failure of all of you in the chain of command."

No training module tells an executive to own a death on the factory floor personally, out loud, in front of the people who could use it against him. But O'Neill didn't leave the lesson to chance, either. He also built a formal system to make sure it kept being taught: every injury, anywhere in the company, reported to his office within 24 hours with a corrective plan attached, and his home phone number handed to any employee who wanted to bypass a manager to report a hazard. Over the next thirteen years, by Duhigg's account, Alcoa's lost-workday injury rate fell by nearly 90 percent.

Look at what that system actually cost him, because the cost is the point. A 24-hour reporting deadline binds the CEO's own calendar, not just the plant's. A corrective plan attached to every report means someone senior reads every one. A direct line that bypasses the manager tells every manager that their filtering is not trusted on this one subject. Each component takes something away from someone with power, starting with O'Neill. That is what separates it from a safety poster.

The lesson generalizes further than it looks: the formal curriculum and the invisible curriculum aren't natural enemies. They drift apart when nobody bothers to align them — and they can be deliberately engineered to reinforce each other, at a cost most leaders aren't willing to pay. Most of the time, though, nobody pays it, and the two curricula pull apart on their own.

Ask experienced employees where they really learned how their organization works, and few say "in training." They remember a moment like O'Neill's — a reaction, a decision, a consequence — that told them more than any handbook did:

- A manager who took the blame publicly instead of quietly reassigning it.

- The promotion that went to someone everyone privately distrusted.
- The mistake nobody was punished for reporting, and the one that quietly ended someone's career.

These moments point to something simple: organizations teach constantly, whether or not anyone is officially teaching. This isn't a new discovery. Educators have called a version of it the "hidden curriculum" since Philip Jackson's 1968 book on classroom life — I use "invisible" rather than "hidden" only because nothing here is being concealed, which is what makes it hard to see — and organizational theorists from Chris Argyris to Edgar Schein have spent decades mapping how it operates inside companies. Steven Kerr named the reward-side version in 1975, in a paper about organizations that structurally reward one behavior while claiming to want another. Safety researchers have mapped the movement itself: Scott Snook's "practical drift" and Sidney Dekker's *Drift into Failure* both describe how local practice slides away from written procedure without anyone deciding that it should. What follows draws on that lineage and adds two working tools to it: **Curriculum Drift** and **Curriculum Debt** — ways of naming when an organization's two curricula pull apart, and what it costs when nobody corrects it. Drift is the rate; debt is the stock.

One caution before the cases. Almost all of them are visible because they ended badly, which is a selection problem, not a coincidence of research — the last section deals with what that does and doesn't prove. (A longer list of cases is in the appendix at the end.)

The Two Curricula That Shape Every Workplace

Every organization operates with two parallel curricula. The formal one is the visible, intentional system — training programs, process manuals, policy documents, leadership development — and it teaches what employees are supposed to know. The invisible one is everything employees actually experience: what gets rewarded or ignored, how leaders behave under pressure, which shortcuts are tolerated, who receives recognition, how meetings really run. It teaches what actually works here.

"Here" is doing more work in that sentence than it looks. A large organization rarely runs one invisible curriculum; it runs several, and the gap between a group's own and the company's is usually where the trouble starts. Danske Bank's Estonian branch is the sharpest example — a unit generating most of its own profit from clients the group's policy would never have approved, teaching its own staff something the group would not have recognized. Everything that follows applies to a unit as readily as to a company, and often more precisely.

> And when the two conflict, employees almost always believe the invisible one.

There's a simple reason for that, and it is the mechanism underneath everything that follows. What a policy says costs nothing to say. What a decision costs — a bonus withheld, a manager reassigned, a top performer actually let go — is expensive, and expensive signals are the ones people believe. Employees aren't being cynical when they trust the invisible curriculum over the formal one; they're pricing information correctly.

LATERAL HOUSE

This is costly signaling — the logic Michael Spence formalized in 1973, and which biologists know through Zahavi's handicap principle: a signal is believable in proportion to what it would cost to fake. A guarantee backed by a refund is more credible than one backed by an adjective, because only one of them can be called. Stating it plainly turns a complaint about culture into a test you can apply to a specific decision: what did this cost the person who made it, and who watched them pay?

The claim has some support beyond anecdote. Guiso, Sapienza and Zingales, surveying employees across more than 1,000 US firms, found that *advertised* corporate values — the ones on the wall — show no relationship to performance, while employees' perceptions of top management's integrity do. That is correlational survey data, and it establishes the first half of the argument rather than the whole of it: it shows the poster is noise and the experienced signal is the asset, not that the gap between them accumulates. The cases that follow do the rest of the work.

When Nobody Is Watching

Commonwealth Bank of Australia had values statements, an ethics framework, and a compliance program. In 2012, a software error silently switched off automated reporting on certain cash transactions — and nobody noticed for three years, until the bank itself identified the gap in 2015 and reported it. AUSTRAC, the anti-money-laundering regulator, brought civil proceedings in 2017 and settled for a A$700 million penalty. A related inquiry, run by banking regulator APRA into the bank's broader governance and culture, concluded that the institution had grown so comfortable with its own success that it had stopped looking closely at anything else: "continued financial success dulled the senses of the institution," in the panel's words.

This isn't a story about an incentive teaching people to look away — nobody has shown that. But where an organization spends attention is itself a costly signal, and spending none on a domain for three years tells everyone inside exactly what that domain is worth. It's a narrower, and in some ways more unsettling, one: an organization can be precise about the things it has decided matter and genuinely blind to the things it hasn't, and a compliance system nobody is watching will fail exactly like one nobody built.

The same principle scales down to the smallest management moments. A promotion, an appraisal, an escalation handled well or badly, a manager interrupting a junior colleague in a meeting — each is a teaching event, whether or not anyone intends it as one. Culture is set by hundreds of moments like these every week: every meeting teaches, every exception teaches, every silence teaches.

When the Curriculum Drifts

The two curricula are most dangerous when they begin to diverge quietly, at a rate nobody is measuring:

LATERAL HOUSE

- **Training says:** Follow the process → **Daily experience teaches:** Speed matters more than process.
- **Training says:** Speak up freely → **Daily experience teaches:** Dissent creates career problems.

This divergence is **Curriculum Drift**, and it runs in two directions.

Most often it looks like risk quietly getting relabeled as acceptable. Studying NASA's decision to launch the Challenger space shuttle in 1986 despite known problems with its O-ring seals, sociologist Diane Vaughan called this specific pattern **normalization of deviance** — a known risk reclassified as fine, simply because nothing went wrong the last several times. NASA made a close variant of the same mistake seventeen years later with Columbia; investigators noted afterward that the agency had fixed the specific technical cause but never actually corrected the cultural pattern the first inquiry had described. Drift like this is rarely one bad decision. It's a series of individually reasonable ones that add up to a risk nobody chose on purpose.

But drift can run the other way too, and the two need telling apart before anyone reaches for punishment. Sometimes the field is right and the manual is stale — a shortcut survives because it's genuinely faster or safer than a process written before conditions changed. That's not corruption; it's often the closest thing the organization has to real expertise, and punishing it teaches people to stop being right ahead of head office. Snook's account of practical drift makes the same observation from the other end: local practice adapts to local reality, and the written procedure is what falls behind. Toyota is the clearest case of an organization that decided not to leave this to chance. Under its production system, the written standard is not the authority; the people doing the work revise it, and a change that survives is written back into the standard, which is what kaizen means in practice. That is benign drift with a collection mechanism attached — the manual is expected to fall behind, so the organization goes looking for where it has. Most companies have the same drift and no such mechanism, which is why they experience it only as non-compliance.

Both kinds of drift look identical from a compliance dashboard. Telling them apart — output slipping versus the manual falling behind — is the actual skill, and it's a judgment call that normalization of deviance, on its own, won't make for you.

Boeing shows the pathological direction with the paper trail left intact. The 737 MAX program ran under sustained pressure to match a competing Airbus jet that was already taking orders Boeing couldn't win back, and that pressure shaped the design and certification of MCAS, the flight-control system whose behavior the two crashes traced back to. What congressional investigators described was not a single decision but years of accumulation, and the relabeling is visible in the design record: MCAS was allowed to depend on a single angle-of-attack sensor, and its authority to move the stabilizer was expanded late in development without the hazard classification being revisited. Each step was defensible on its own terms. Together they moved a system from one risk category into another without anyone recording that a category had changed — cost and schedule goals that, in the committee's words, "jeopardized the safety of the flying public." Two 737 MAX crashes killed 346 people. The company's accountability afterward was mostly financial and corporate rather than

LATERAL HOUSE

personal: four years of legal process, about $3.6 billion across two agreements with the Justice Department ($2.5 billion in 2021 and a further $1.14 billion in 2025), and no criminal conviction of the company itself.

Wells Fargo looks like the same story and isn't, which is why it belongs directly beside Boeing. Its Vision & Values statement, published in 2012, said: "We start with what the customer needs — not with what we want to sell them." Inside its branches, staff worked under a cross-sell target management itself branded "Gr-eight" — eight products per household — tracked daily and enforced with discipline up to termination for anyone who missed it. Nothing there drifted. The gap never widened, because it was built at full width, on purpose, and tracked daily by the people who built it. This is Kerr's problem exactly: rewarding A while hoping for B, from day one.

The consequences were the same either way. Employees who couldn't hit the number honestly, and couldn't afford to be fired for missing it, faced enormous pressure to cut corners — and thousands did, opening accounts customers hadn't asked for. Between January 2011 and March 2016, Wells Fargo fired approximately 5,300 employees over the resulting unauthorized-account practices. They went steadily, about 1 percent of the retail workforce a year, not in a single purge. This was the same company punishing people for behavior its own quota had pushed many of them toward. That double bind, not individual dishonesty alone, is what a quota this aggressive was training people to do at scale. The bank paid $185 million to settle with regulators in 2016 and $3 billion to resolve federal criminal and civil investigations in 2020. Nobody had to be secretly corrupt for this to happen. The quota was the curriculum. The values statement was the poster on the wall.

Wells Fargo is not drift; it is what drift is often mistaken for. The distinction matters practically. Drift is corrected by noticing movement — comparing what people do now against what they did last year. A contradiction designed in from the start shows no movement to notice, because it was never anywhere else. It is caught only by asking what the plan itself rewards, which is a different question and one almost nobody puts to their own targets.

There is a rival explanation for all of this, and it lands hardest right here. Perhaps these cultures didn't teach people at all — perhaps they selected them. Wells Fargo may not have taught anyone to open fake accounts so much as retained the people already willing to and pushed out the rest, a percent a year. If that's the mechanism, the remedy is hiring and exit interviews, not teaching events. Honestly, both are probably running at once, and they are hard to separate from inside. But it changes what the evidence would look like: if it's teaching, behavior shifts in people who stay; if it's selection, the population changes while individuals don't.

At first, gaps like these look survivable. But repeated contradictions change what people believe. They stop asking what the policy says and start asking what actually happens here. Once that shift happens, cynicism becomes rational.

Not every organization lets the metric win, and the mechanism runs identically in the other direction: the decision was costly, so it was believed. In 2009, as US retailers were freezing wages and cutting staff through the recession, Costco's CEO Jim Sinegal approved a $1.50-an-hour raise for hourly

LATERAL HOUSE

employees anyway, phased in over the next three years. His CFO, Richard Galanti, later said Sinegal's first reaction to being briefed on the downturn was: *"This economy is bad. We should be figuring out how to give them more, not less."* It fit a pattern: Wall Street analysts have complained for years, before and since, that Costco is too generous to its own people. A skeptic can reasonably answer that this is simply good labor economics for a high-volume, membership-fee business, where low turnover pays for itself — and that skeptic may be right about the arithmetic while missing what employees learned. The lesson wasn't the number on the pay stub. It was what the company chose to protect the moment protecting it stopped being free.

The Accumulation of Curriculum Debt

Organizations are familiar with technical debt: the shortcuts taken today that create larger engineering costs tomorrow. In the same way, organizations accumulate **Curriculum Debt**. If drift is the rate at which the two curricula separate, debt is the accumulated stock of separation nobody has corrected.

One qualification belongs here rather than at the end, because it changes how the metaphor should be read. This is a contingent liability, not a term loan. It has no fixed due date, and in many organizations it is never called at all. What it does is raise the probability and the size of a loss whose timing nobody can predict.

Every time leaders let a contradiction between stated values and observed behavior go unaddressed, they borrow against the future culture of the organization. Common examples: tolerating policy violations by high-performing rainmakers, ignoring toxic behavior because business results are strong, rewarding outcomes without examining the methods used to achieve them, allowing repeated process bypasses during high-pressure periods.

The reason this compounds rather than merely persisting is the whole justification for the metaphor, so it should be exact: **each unaddressed exception lowers the price of the next one.** The second waiver is cheaper to grant than the first, because the first established that waivers are grantable. The tenth is nearly free. What accumulates is a falling internal price on the rule itself.

Credit Suisse's 2021 loss on a single client, Archegos Capital Management, shows what unpaid debt looks like when it finally comes due. The bank's own investigators found the relationship had breached internal risk limits repeatedly and visibly for months, and that the risk function itself failed to escalate it or force the one control that would have capped the damage. The risk was conspicuous. People inside the bank saw it, at multiple points, and didn't act on it.

The investigators' own numbers make the failure starker, not more explicable: in 2020 the bank earned only around $16 million in actual revenue from the relationship, against exposure that would eventually cost it $5.5 billion when Archegos defaulted in March 2021. The bank was chasing a bigger number it never collected.

LATERAL HOUSE

No law was broken to produce this: the investigators found no fraud or illegal conduct by anyone at Credit Suisse itself — the limits were breached, repeatedly, but breaching an internal limit isn't a crime, just a decision nobody reversed. Nobody at the bank had to commit a crime to lose $5.5 billion. The organization simply kept choosing not to collect on a debt it had been accumulating in plain sight, for a fraction of what that debt would eventually cost.

A simple way to see how much debt has already accumulated, once: list every exception granted in the last twelve months — every waived approval, every violation left unaddressed because the numbers were good, every promotion that quietly contradicted a stated value. That list is the debt register. It's usually longer, and more revealing, than anyone expects. A minimal version has five columns:

EXCEPTION GRANTED	WHO BENEFITED	VALUE IT CONTRADICTED	DATE	REPAYMENT CONDITION
Missed a mandatory second sign-off on a large deal	Top-quartile sales rep, end of quarter	"No exceptions to the approval process"	Q1	Reviewed by someone the rep doesn't report to; any repeat within two quarters escalates automatically and the waiver stops renewing itself
Shipped a release without the required control check	The delivery team lead; the quarter's numbers	"The control check is never optional"	Q2	Not yet defined

The last column is the one that matters and the one everyone skips. An exception with a repayment condition attached — a review date, a reversal trigger, a cap on how many times it can recur — is a debt someone is managing. An exception with "not yet defined" in that column, six months after it was granted, is the debt compounding. That first sweep is deliberately wide, to show you the shape of the problem. The standing register that follows it has to be narrower or nobody will maintain it: only exceptions that needed discretionary sign-off above a defined level, reviewed quarterly by someone senior enough to close entries rather than only add them. The register isn't punitive; a leader who grants an exception, records it, and names the condition for closing it has done something different in kind from one who lets it pass unrecorded. Paying it down is rarely painless — it can mean withholding a bonus from a top performer who violated values, or dismantling a legacy system that quietly rewards the wrong behaviors.

The register has a real limit, and it isn't the one that first comes to mind. In 1985, after seeing damage on a previous flight's nozzle joint, NASA's Marshall Space Flight Center imposed a formal launch constraint on the shuttle's solid rocket boosters — then waived it for flight after flight in the run-up to Challenger, every waiver signed by the same project manager. The exception was recorded every time. What failed was where the record stopped: the constraint, its cause, and the waivers never reached the officials actually running the launch-readiness review. A register catches the exception. It doesn't catch an exception that never travels above the level that created it — that needs its own rule: a repayment condition renewed more than once has to be signed off by someone outside the chain that approved it the first time, not filed a level below where the real decision gets made.

LATERAL HOUSE

There's a second, more basic limit, and it applies to some of the worst cases here: a register only catches exceptions *to* policy. It has nothing to say about a policy that is itself the problem. Wells Fargo's quota was never logged as an exception — it was the plan, approved and tracked daily. The same is true of Purdue's bonus structure and Volkswagen's emissions targets. No amount of exception-tracking finds debt that was written into the org chart on purpose. That takes a different, harder question, and it isn't one a register can ask for you: not "what did we let slide," but "what are we currently measuring people on, and would we be comfortable if everyone hit the number exactly as designed?"

The Hidden Examination

Schools have formal examinations; organizations have them too, though nobody announces them. They arrive quietly, often at the end of a difficult day, when the easy answer and the honest answer point in different directions:

- Will you report a mistake?
- Will you escalate an emerging risk?
- Will you protect the customer when doing so is inconvenient?
- Will you follow procedure when nobody is watching?

In 2018, Wirecard's senior in-house counsel for Asia-Pacific, Pav Gill, escalated concerns about accounting irregularities in the company's Singapore office. When he reported what he found to Munich headquarters, the investigation was shut down and he was pushed out of the company. He kept a copy of his files, which eventually reached the Financial Times and became the basis of the reporting that forced the issue into the open.

By spring 2020, Wirecard's own supervisory board had commissioned KPMG as an independent special investigator to clear the company's name — and KPMG instead reported that it could not verify roughly half the company's revenue. Weeks later, a supposed €1.9 billion held in escrow turned out not to exist, and Wirecard collapsed: the first company ever to leave Germany's DAX stock index through insolvency. Gill's case is the hidden examination in its plainest form. The question was never whether Wirecard had a whistleblowing process. It did. The question was what happened to the person who used it.

Boeing sat the same examination and answered it the same way, on the factory floor rather than in a finance function. In June 2018, a senior manager on the 737 production line, Ed Pierson, wrote to the head of the program describing fatigue and chaos in production and asking that the line be temporarily shut down. He was refused. He met the program head again in July, repeated the concern, and retired on 1 August. His account is not the cause of the MCAS failures — those were designed and certified years earlier — but it is the same test, taken by a different person in a different part of the same company. Gill was expelled; Pierson left. The answers the company gave were the same. What every colleague of his learned that summer is not recorded in any report.

LATERAL HOUSE

> Employees learn not just from the questions, but from what happens to the people who answer them honestly.

If responsible escalation is punished, the invisible curriculum teaches silence. If integrity is protected, it teaches accountability.

A Quick Gut-Check for Leaders

This isn't a hypothetical exercise. In the mid-1970s, around the time he became CEO of Johnson & Johnson, James Burke ran a series of meetings challenging the company's senior managers to either recommit to the company's decades-old Credo clause by clause, rewrite it, or throw it out — because an untested statement of values was worse than none. The group voted to keep it, with updated language. A few years later, when cyanide-laced Tylenol capsules killed seven people in Chicago, Burke pulled 31 million bottles from shelves nationwide — against advice from the FBI and FDA that a full recall wasn't necessary — at a cost of roughly $100 million. He later said the Credo, tested years earlier by people who would have to live with it, gave him the standing to make that call. The point isn't the dollar figure. The test happened before the crisis, not during it — which is the only time a test of values is cheap enough to be honest. Johnson & Johnson's much later record, including billions in opioid- and talc-related liabilities, is a reminder that no Credo holds forever once the people who stress-tested it have moved on.

The entire point of the invisible curriculum is that leaders are often the last to see what it's teaching — so the most important instruction here isn't one of the questions below, it's who answers them. Ask people two levels below you, not yourself, and ask more than one of them separately rather than as a group.

After any significant decision, put a short set of questions to them:

1. What specific behavior just got rewarded here, and by whom?
2. Would you want every new employee to copy this behavior?
3. What story will people tell about this tomorrow?

These aren't an audit. They're a habit. Run them often enough, with people other than yourself doing the answering, and the invisible curriculum stops being invisible. One caveat, and it has a usable tell. If the answers keep coming back reassuring, that isn't proof everything is fine — in an organization where honesty is risky, reassurance is the safe answer. What separates a clean result from a frightened one is texture: people who actually feel safe volunteer small negatives unprompted, give concrete examples, and disagree with each other. Bland, uniform, positive answers from three people interviewed separately is the finding. Varied, specific, faintly critical answers are the clean bill.

Most readers of this essay cannot keep a company-wide register or poll two levels down, because most people in an organization have no authority over its exceptions. The instruments still scale down. A team lead's register is the list of times they let something slide for a good reason and never went back

LATERAL HOUSE

to it. Their gut-check is what they teach by how they handle the first mistake somebody reports to them — which is the moment their team learns whether reporting is safe, and is remembered long after any policy is forgotten. The layer that actually transmits most of the invisible curriculum is not the CEO. It's the manager translating pressure downward: the branch manager who enforced the daily Gr-eight target, not the executive who approved it.

The Real Question

The hardest case comes last, because it complicates everything above. Few companies have ever formalized their culture as thoroughly as Bridgewater Associates, whose 200-plus written Principles, recorded meetings, and real-time peer-rating systems were built explicitly to make values enforceable rather than aspirational. Journalist Rob Copeland's reporting on the firm — published as *The Fund* in 2023, and disputed by Bridgewater and its founder — found that the same apparatus could function as surveillance and score-settling, and that the founder's own standing inside the system was, in practice, unchallengeable, exactly inverting the transparency it was built to enforce.

If that reporting holds up, it is also the strongest objection to the register proposed above, and it deserves a direct answer rather than a graceful exit. Bridgewater did more documenting than this essay recommends, and on Copeland's account the documentation was captured by the people it was meant to constrain. Wells Fargo had a compliance function. Credit Suisse's risk function watched limits break for months. Why would a five-column table survive where those didn't?

It wouldn't, on its own. The register does no work as a record; it does its work only at the escalation rule — the requirement that a condition renewed more than once goes to someone outside the chain that granted it. That is the single component Challenger's waiver trail lacked, and it is the component that is always the first to be quietly dropped, because it is the only one that costs the grantor anything. A register without it is documentation, and documentation gets captured. This is the same test as everything else here: does the mechanism take something away from the person with power? If not, it will teach nothing.

Which brings up the fair objection to the whole framework. Every failure case here became visible because it eventually blew up — a settlement, a crash, an insolvency, an inquiry. That's how these stories get told at all. Plenty of organizations run quotas and cultures just as contradictory and never produce a headline, so a reader is entitled to ask what this framework would have said about Wells Fargo in 2013, and whether the answer is anything more than hindsight with vocabulary.

It is, but only because its testable parts are the boring ones, and they are all available before any bill arrives: how many entries in the register have "not yet defined" in the last column, and for how long; whether a repayment condition has ever actually been enforced against someone senior; whether three people two levels down, interviewed separately, give the same bland answer. Those are observable this

LATERAL HOUSE

quarter. None of them predicts a crisis. All of them tell you which direction you're moving. That is the most any of this claims to do. Most of the time you won't be able to tell from the inside which kind of organization you are — which is the argument for keeping the register, not against it.

None of this makes training or policy pointless. But the real culture of an organization lives in the space between what's written and what's repeatedly experienced — the space where trust is built or spent, where an organization's actual operating values get taught, tested, and passed on, whether or not anyone intended to teach them. Paul O'Neill's 24-hour reporting rule worked for a reason Bridgewater's Principles didn't: it applied to him first. He didn't just write it down. He made sure that when something went wrong, everyone in the room found out exactly what he thought that meant — and watched what happened to the first person who said it out loud, because that person was O'Neill himself. Writing the values down was never the hard part. What's hard is making sure the people who wrote them are graded by the same rules as everyone else. The question worth asking isn't "Did we communicate the policy clearly?" It's "What did our decisions teach today?" The answer to that question is usually a more accurate description of an organization than anything printed in its values statement.

Appendix: More Cases, For Reference

The essay ends above. What follows is a working reference, not more argument — a shortlist for anyone who wants a case that fits a specific audience better than the ones used in the text. Evidence quality varies, and each row says which kind it is: most rest on regulatory orders, court-appointed investigations, or admitted facts.

CASE	FORMAL CURRICULUM SAID	WHAT ACTUALLY GOT REWARDED	OUTCOME
Goldman Sachs / 1MDB (2011–2020) — *edge case*	Client-vetting controls repeatedly rejected financier Jho Low, unable to verify his wealth	Bond deals he helped arrange anyway went through Goldman's bond-underwriting committees, a different approval body from private-wealth client vetting — earning ~$600m in fees — and a banker was later found to have actively concealed Low's role from those committees	~$2.9bn coordinated global resolution (2020). Included deliberately as the harder case: this is less a story of incentives overriding culture than of individual fraud defeating it — a different failure mode worth telling apart from the others here
Danske Bank (2007–2015)	Group-wide anti-money-laundering policy and Code of Conduct	Its Estonian branch generated a large majority of its own profit from high-risk, non-resident clients, left outside group monitoring — a subculture running a curriculum the group did not share	Guilty plea; $2.06bn forfeited to US authorities (2022)
Volkswagen "Dieselgate" (2006–2015)	"Clean Diesel" certified compliant with US emissions law	An unreachable performance target under a CEO known to punish admissions of failure	$4.3bn in US federal penalties alone (total US cost closer to $25bn); CEO charged by the US but never extradited or tried there; two other executives served US prison terms

LATERAL HOUSE

Purdue Pharma (1996–2019)	A compliance program to detect and prevent opioid diversion	Sales bonuses that regularly exceeded base salary, tied directly to prescription volume	Guilty plea to federal fraud and kickback conspiracies; $8.3bn in penalties agreed (2020), largely unrecovered through the company's bankruptcy

Sources

Links point to the primary document where one is publicly available — regulator orders, court filings, and official inquiry reports. Books, paywalled articles, and documents whose original host is no longer online are cited without a link rather than pointed at a substitute.

- Intellectual lineage — Philip W. Jackson, *Life in Classrooms* (1968), on the "hidden curriculum"; Chris Argyris & Donald Schön on espoused theory vs. theory-in-use (1974 onward); Edgar Schein, *Organizational Culture and Leadership* (1985); Steven Kerr, "On the Folly of Rewarding A, While Hoping for B," (https://doi.org/10.5465/255378) *Academy of Management Journal* 18, no. 4 (1975): 769–783; Scott A. Snook, *Friendly Fire* (2000), on "practical drift"; Sidney Dekker, *Drift into Failure* (2011); on costly signaling, Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87, no. 3 (1973), and Amotz Zahavi's handicap principle (1975)
- Advertised vs. perceived values — Luigi Guiso, Paola Sapienza & Luigi Zingales, "The Value of Corporate Culture," *Journal of Financial Economics* 117, no. 1 (2015): 60–76
- Alcoa / Paul O'Neill — Charles Duhigg, *The Power of Habit* (2012), ch. 4. Both the dialogue and the ~90% injury-rate reduction rest on Duhigg's reporting: the dialogue is reconstructed by the author from later interviews rather than a contemporaneous transcript, and Alcoa published no injury rates in its SEC filings until 2002, so the figure cannot be checked against disclosure. Alcoa's FY2001 10-K reports a lost-workday rate of 0.16 — one lost-workday injury per 1.25 million hours worked
- Commonwealth Bank of Australia — APRA, *Prudential Inquiry into the Commonwealth Bank of Australia: Final Report* (Apr 30, 2018) (https://www.apra.gov.au/sites/default/files/CBA-Prudential-Inquiry_Final-Report_30042018.pdf); AUSTRAC, "AUSTRAC and CBA agree $700m penalty" (Jun 2018) (https://www.austrac.gov.au/news-and-media/media-release/austrac-and-cba-agree-700m-penalty) (see also CBA's own statement (https://www.commbank.com.au/guidance/newsroom/CBA-and-AUSTRAC-resolve-AMLCTF-proceedings-201806.html)); proceedings commenced 3 Aug 2017, penalty ordered by the Federal Court 20 Jun 2018
- Wells Fargo — CFPB Consent Order 2016-CFPB-0015 (Sept 8, 2016) (https://files.consumerfinance.gov/f/documents/092016_cfpb_WFBconsentorder.pdf), which imposes the $100m CFPB penalty; the $185m total comprises that plus $35m to the OCC and $50m to the City and County of Los Angeles. DOJ, "Wells Fargo Agrees to Pay $3 Billion" (Feb 21, 2020) (https://www.justice.gov/archives/opa/pr/wells-fargo-agrees-pay-3-billion-resolve-criminal-and-civil-investigations-sales-practices). The 5,300 terminations figure and the ~1%-a-year rate are from John Stumpf's testimony to the Senate Banking Committee (Sept 20, 2016) (https://www.govinfo.gov/content/pkg/CHRG-114shrg23001/html/CHRG-114shrg23001.htm), not from the CFPB order. Wells Fargo's 2012 Vision & Values statement

LATERAL HOUSE

- Boeing 737 MAX — House Transportation & Infrastructure Committee, *Final Committee Report: The Design, Development & Certification of the Boeing 737 MAX* (Sept 2020) (https://democrats-transportation.house.gov/imo/media/doc/2020.09.15%20FINAL%20737%20MAX%20Report%20for%20Public%20Release.pdf), from which the "jeopardized the safety of the flying public" finding is quoted; FAA, *Section 103 ODA for Transport Airplanes Expert Panel Review Report* (Feb 26, 2024) (https://www.faa.gov/newsroom/Sec103_ExpertPanelReview_Report_Final.pdf) — 27 findings, 53 recommendations; Ed Pierson, written testimony to the House Transportation Committee (Dec 11, 2019) (https://docs.house.gov/meetings/PW/PW00/20191211/110296/HHRG-116-PW00-Wstate-PiersonE-20191211.pdf); DOJ filings and resolutions, *United States v. The Boeing Company* (2021 deferred prosecution agreement; 2024 breach finding and abandoned guilty plea; 2025 non-prosecution agreement). The widely-quoted "designed by clowns… supervised by monkeys" message is from Boeing's own document release of Jan 9, 2020 (https://boeing.mediaroom.com/news-releases-statements?item=130600), and does not appear in the Committee report
- NASA Challenger / Columbia — Rogers Commission Report (1986) (https://www.nasa.gov/history/rogersrep/genindex.htm), esp. ch. VI "An Accident Rooted in History"; the Commission states that the launch constraint was waived "for each Shuttle flight after July 10, 1985" without giving a count, and its own flight table lists six such flights before Challenger. *Columbia Accident Investigation Board Report*, Vol. I (Aug 2003) (https://www.nasa.gov/wp-content/uploads/static/history/columbia/reports/CAIBreportv1.pdf), esp. ch. 8 "History as Cause," pp. 195–198. Diane Vaughan, *The Challenger Launch Decision* (1996)
- Costco — Bloomberg Businessweek, "Costco CEO Craig Jelinek Leads the Cheapest, Happiest Company in the World" (Jun 2013). Costco discloses a retention rate rather than a turnover rate, and only for employees with at least one year of service — around 90–94% in recent 10-K filings — which excludes first-year churn
- Credit Suisse / Archegos — Paul, Weiss, Rifkind, Wharton & Garrison, *Report of the Special Committee of the Board of Directors of Credit Suisse Group AG* (Jul 29, 2021) (https://www.sec.gov/Archives/edgar/data/1053092/000137036821000064/a210729-ex992.htm), which gives Archegos's 2020 revenue contribution as approximately $16 million; DOJ (SDNY), statement on the convictions of Bill Hwang and Patrick Halligan (Jul 10, 2024) (https://www.justice.gov/usao-sdny/pr/statement-us-attorney-damian-williams-convictions-bill-hwang-and-patrick-halligan) and the sentencing release (Dec 2024) (https://www.justice.gov/usao-sdny/pr/founder-and-head-archegos-capital-management-bill-hwang-sentenced-18-years-prison)
- Wirecard — KPMG, *Report on the Independent Special Investigation* (Apr 27, 2020); European Parliament briefing (2021); Dan McCrum / Financial Times investigative reporting (2015–2020)
- Johnson & Johnson / Tylenol — Harvard Business School case study, "Johnson & Johnson's Corporate Credo"; Knowledge@Wharton, "Tylenol and the Legacy of J&J's James Burke" (see also J&J's subsequent opioid and talc litigation record)
- Bridgewater Associates — Rob Copeland, *The Fund* (2023); contested by Bridgewater and Ray Dalio
- Appendix cases — DOJ, "Goldman Sachs Charged in Foreign Bribery Case and Agrees to Pay Over $2.9 Billion" (Oct 22, 2020) (https://www.justice.gov/archives/opa/pr/goldman-sachs-charged-foreign-bribery-case-and-agrees-pay-over-29-billion); Bruun & Hjejle report on Danske Bank (2018) and DOJ, "Danske Bank Pleads Guilty

LATERAL HOUSE

to Fraud on U.S. Banks" (Dec 13, 2022) (https://www.justice.gov/archives/opa/pr/danske-bank-pleads-guilty-fraud-us-banks-multi-billion-dollar-scheme-access-us-financial); DOJ, "Volkswagen AG Agrees to Plead Guilty and Pay $4.3 Billion" (Jan 11, 2017) (https://www.justice.gov/archives/opa/pr/volkswagen-ag-agrees-plead-guilty-and-pay-43-billion-criminal-and-civil-penalties-six) — which covers the penalties and the six indictments, but not the later sentencing of James Liang and Oliver Schmidt or the 2018 charging of Martin Winterkorn; DOJ, "Opioid Manufacturer Purdue Pharma Pleads Guilty to Fraud and Kickback Conspiracies" (Nov 2020) (https://www.justice.gov/archives/opa/pr/opioid-manufacturer-purdue-pharma-pleads-guilty-fraud-and-kickback-conspiracies) and the $8.3bn global resolution (Oct 2020) (https://www.justice.gov/archives/opa/pr/justice-department-announces-global-resolution-criminal-and-civil-investigations-opioid)

Cite this essay. Lateral House (2026). “The Invisible Curriculum of Organizations.” Lateral House, August 2026. https://lateralhouse.com/essays/invisible-curriculum-of-organizations/

© 2026 Lateral House. Published at lateralhouse.com. This PDF is for personal and educational reading; please link to the web version when sharing.

LATERAL HOUSE