O'Sullivan's First Law

Idea In Short

A policy that reliably produces the wrong result is not bad luck, it is bad design. O'Sullivan's First Law, originally coined to describe how institutions drift from their founding purpose, gives boards and executives a discipline: judge a policy by what it actually produces, not by what it was meant to produce. When a compliance rule, incentive scheme, or governance structure keeps generating the outcome leadership says it wants to avoid, the answer is not more enforcement or a fresh memo restating intent. The answer is to presume the design is wrong and rebuild it. This article applies that discipline to corporate governance, incentive design and policy review and sets out how boards can build in the kind of periodic, outcome-based audit that catches drift before it hardens into culture.

Executives spend a great deal of time debating whether a policy is well-intentioned. O'Sullivan's First Law asks a blunter question: does the policy work? The law began as an observation about political organizations, credited to journalist and columnist John O'Sullivan, who noted that institutions not explicitly built and defended around a conservative purpose gradually drift toward the opposite over time1. Strip out the partisan framing and a durable governance principle remains: organizations drift from their stated purpose unless something actively holds them in place and the direction of that drift is set by what the organization actually rewards, not by what its charter says. For boards, executives and policy designers, the law offers a diagnostic shortcut. When a policy produces the same unwanted outcome repeatedly, treat that outcome as evidence about the design, not as an accident to be explained away.

The origin of the observation

O'Sullivan made his observation in the 1980s while reflecting on why institutions such as universities, churches and media organizations tended to move in a consistent ideological direction over decades, regardless of who founded them or what their original charter said. Commentator Robert Conquest had earlier proposed a related idea, sometimes summarized as the notion that any organization not explicitly right-wing will become left-wing, which writers have traced back through Conquest's own reflections on institutional behavior2. Both observations share a mechanism rather than a political claim: institutions are staffed and re-staffed continuously and each new cohort inherits the formal rules but not necessarily the founders' vigilance about defending them. Absent active maintenance, the organization gradually optimizes for whatever is easiest to measure, reward, or avoid conflict over. That mechanism generalizes cleanly to corporate settings, where "explicitly right-wing" becomes "explicitly and continuously defended" and the drift shows up as scope creep, incentive gaming, or governance structures that no longer serve the purpose they were built for. Executives who dismiss the law as a political curiosity miss the structural insight underneath it.

Why organizations drift instead of holding steady

Drift happens because most organizational rules are written once and then interpreted repeatedly by people who were not in the room when the rule was made. A compliance policy drafted to prevent a specific abuse becomes, five years later, a box-ticking exercise administered by staff who know the letter of the rule but not the incident that prompted it. Sociologists call a related pattern goal displacement, where an organization's stated objective is gradually replaced by a proxy that is easier to measure and reward, such as a regulator substituting enforcement volume for actual compliance outcomes3. The same substitution happens inside companies whenever a metric becomes the target instead of remaining a proxy for the target. Sales teams optimize for the number that triggers commission, not necessarily for the customer relationship the commission was meant to reward. Risk committees optimize for the appearance of thorough review, not necessarily for catching the risk the review was designed to catch. None of this requires bad faith. It requires only the ordinary fact that people respond to what is measured and rewarded and that measurement systems, once built, rarely get revisited as carefully as they were designed.

The corporate governance case

Corporate boards exist substantially to prevent this kind of drift, which makes them a natural test case for the law. A board's formal charter typically states a broad purpose, protecting shareholder value, overseeing management, ensuring sound risk practices, but the charter alone does nothing to stop drift once directors rotate and the people who wrote the charter leave. Governance researchers have documented how boards accumulate blind spots over multi-year tenures and one widely cited response has been a push toward more active, continuous board engagement rather than periodic box-checking4. The practical version of O'Sullivan's First Law for a board is straightforward:

> if the same category of failure, a missed risk, an executive compensation structure that rewards short-term stock performance over durable value, a compliance breach that recurs despite prior remediation, keeps showing up on the board's agenda, the board should stop treating each instance as a one-off and start treating the pattern as proof that the underlying governance mechanism is broken

Decision rights and where drift hides

Drift often hides in decision rights rather than in the headline policy itself. A company can have an excellent stated risk appetite while its actual approval process routes every consequential decision through a single executive whose incentives differ from the board's stated priorities. Analysis of organizational decision-making consistently finds that unclear or informally shifted decision rights, not the absence of a written policy, are what produce outcomes at odds with stated strategy5. This is why a board reviewing a recurring bad outcome should ask not only "what does the policy say" but "who actually made this decision and what did they have to gain." The gap between the two answers is usually where the drift lives.

Reading outcomes as diagnostic evidence

The discipline O'Sullivan's First Law demands is to stop treating a repeated bad outcome as a communication failure and start treating it as data about the design. Economist Rob Norton's summary of the broader unintended consequences literature makes a similar point: outcomes that keep recurring under a stable set of rules are not really unintended, they are the predictable output of those rules interacting with the incentives of the people who operate under them6. A retail chain that keeps missing its inventory targets despite repeated staff retraining is not suffering from a training gap; it likely has an incentive or supply chain design that rewards the behavior producing the shortfall. A compliance function that keeps failing to catch the same category of violation is not short on vigilance; its review process likely samples the wrong transactions or reports to someone with a conflicting incentive.

The test is simple to state and hard to apply consistently:

> has this exact category of failure happened more than once under the same policy? If so, the policy, not the people executing it, is the more likely cause and it deserves redesign rather than another round of reinforcement training

Applying the law to policy design

Applying this discipline changes how a leadership team should run its policy review cycle. Instead of asking whether a policy is defensible in principle, the review should ask what the policy has actually produced since it was last examined, using the organization's own incident and performance data as the record. Research on hybrid organizations, including social enterprises balancing commercial and mission goals, shows that accountability structures which track outcomes against original purpose on a fixed schedule catch drift earlier than those relying on annual narrative reporting alone7. That finding translates directly to a corporate policy review: pair every governance policy with a small set of outcome metrics that would reveal drift and revisit those metrics on the same cycle as the policy itself, not on an ad hoc basis triggered only after a visible failure. This does not mean policies should chase short-term metric swings; a single bad quarter is not evidence of drift. Two or three recurrences of the same failure mode, under otherwise stable conditions, is.

Naming the mechanism before rewriting the rule

A common mistake when applying the law is to rewrite the rule without first naming the mechanism that produced the drift. A board that tightens a compliance policy after a violation, without asking why the previous version failed to prevent it, often ends up with a stricter version of the same flawed mechanism. The more durable fix names the specific incentive or reporting gap that let the old policy fail, states it in the revised policy's rationale and assigns a specific owner accountable for monitoring that exact mechanism going forward. This turns the policy review from a reactive patch into a structural correction, which is the difference between a rule that survives leadership turnover and one that drifts again as soon as the people who remember the original incident move on.

Guardrails against drift for boards and executives

Preventing drift is less about writing a better rule once and more about building a standing practice of checking rules against results. Boards benefit from rotating independent directors whose primary value is asking why a recurring outcome has not prompted a redesign, since incumbent directors and management are the people least likely to notice drift they have grown accustomed to. Executives benefit from pairing every incentive scheme with an explicit statement of the behavior it is meant to produce, reviewed against the behavior it actually produces at fixed intervals rather than only after a scandal. Policy designers, whether inside a regulatory agency or a corporate compliance function, benefit from distinguishing a true one-off surprise from a pattern, since only the pattern justifies redesign under this law. None of this requires cynicism about intentions. It requires the more modest and more useful habit of letting results, tracked consistently over time, carry more weight than the stated purpose of the policy that produced them.

- 1John O'Sullivan [Columnist]
- 2Conquest's Laws
- 3Exploring The Risk Of Goal Displacement In Regulatory Enforcement Agencies
- 4Corporate Governance 2.0
- 5Untangling Your Organization's Decision Making
- 6Unintended Consequences
- 7The Governance Of Social Enterprises

Summary

O'Sullivan's First Law started as an observation about institutional politics and travels well into corporate governance because both describe the same mechanism: structures drift from their stated purpose unless something actively holds them in place. For boards and executives, the practical discipline is to treat a repeated bad outcome as evidence about the policy that produced it, not as a reason to explain the outcome away. Incentive schemes, compliance functions and governance charters all drift toward whatever they actually reward, which is not always what their authors intended. The fix is neither cynicism nor rigid nostalgia for original intent. It is a habit of auditing policies against their real-world results on a fixed schedule, naming the mechanism that produced an unwanted outcome and rebuilding the design rather than patching the symptom. Organizations that make this a routine practice catch drift while it is still cheap to correct.