

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

THE FOOT OF THE VOLCANO

How Public Warning Fails
Against a Rare, Latent,
and Regime-Shifting Hazard

JAYSON COIL

COGPLEXITI · SYSTEMS & DECISION ANALYSIS
WHITE PAPER

The Foot of the Volcano

How Public Warning Fails Against a Rare, Latent, and Regime-Shifting Hazard

JAYSON COIL

July 2026

How to read this paper. This paper belongs to the Cogplexiti Systems & Decision Analysis line. It follows the line's five-part structure — the argument, the challenge to it, the credible alternative, and the recommendation — so a reader can weigh the case rather than receive it. Its subject is a single failure: a correctly warned member of the public, reasoning soundly, does not take the protective action the hazard requires. The paper does not treat message design, public-information craft, or the mechanics of alert and warning systems. Those are real, and they are covered elsewhere. This paper sits one layer beneath them. It asks why the failure persists even when the message is good, and it takes the failure apart to show that the persistence has a structure. The recommendation arrives last and stays short, because the contribution is the anatomy, not the fix.

A note on stance. This is analysis, not doctrine. Nothing here directs anyone to do anything. The paper applies the discipline used throughout this line: every actor is treated as locally rational, and the question is never who failed but what conditions made the outcome likely. That commitment is load-bearing here. The subject is a public that does not act, and the easy explanation — that the public is complacent, or under-informed, or in denial — is the explanation this paper is built to replace. Complacency is a verdict. This paper looks for a mechanism.

A note on citation. External constructs are attributed to their source literature at the level of their established claims, and fire-behavior claims are anchored to the fire-science literature at the same standard. Full sources appear in the References.

COPYRIGHT AND LICENSE

© 2026 Jayson Coil. All rights reserved. Published by Cogplexiti LLC.

ISBN: 979-8-9970824-6-8

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations used in reviews and other uses permitted by copyright law.

Use for Artificial Intelligence and Machine Learning. This work may not be used to train or develop artificial intelligence or machine learning systems, or be included in any dataset used for such purposes, without the prior written permission of the publisher.

Permissions. Direct requests to Cogplexiti LLC at permissions@cogplexiti.com.

EDITION INFORMATION

July 2026. Published by Cogplexiti LLC, Flagstaff, Arizona · jayson@cogplexiti.com · www.cogplexiti.com

Part of the Cogplexiti Analysis line (Systems & Decision Analysis), alongside The System Tier, The Incident Tier, Sustaining Decision Performance, and The Reveal.

DISCLAIMER

The views, frameworks, and recommendations expressed in this publication are those of the author and do not necessarily reflect the official policy or position of the Sedona Fire District, Southwest Incident Management Team 1, the National Wildfire Coordinating Group, the U.S. Forest Service, or any other agency or organization. This document is not doctrine. It is a professional analysis informed by operational experience, research, and study.

CONTENTS

1 · Executive Summary	4
2 · Support and Analysis.....	6
2.1 The substrate: outcome, not exposure	6
2.2 Divergence one: the baseline	6
2.3 Divergence two: the magnitude	10
2.4 The two divergences are not parallel.....	14
2.5 The coupling.....	15
3 · Red Team Assessment	17
3.1 The author's position.....	18
3.2 The single-hazard anchor.....	18
3.3 Does misclassification depend on discontinuity?	18
3.4 Does construal distance also touch experience?	19
3.5 Does motivated abstraction explain too much?	19
3.6 The two-route exhaustiveness claim.....	20
3.7 Is "regime" too strong a word?	20
3.8 The evidence base	21
4 · Dissenting Alternative.....	21
5 · Recommended Course of Action	23
5.1 Pre-commitment decided cold	23
5.2 Borrowed outcomes	23
5.3 Change the exposure, not the belief.....	24
5.4 What would show this wrong	24
5.5 Scope	25
References.....	26

1 • Executive Summary

A community lives at the foot of a dormant hazard. For years, nothing happens. The absence of events is read, correctly, as evidence — and the evidence points the wrong way. Then the event arrives, and it is not a larger version of anything the community has seen. It is a different thing entirely, and the model that served for years is useless in the window where it is finally needed.

This is the standing problem of public risk warning for rare, high-magnitude fire. The warning has to change a mental model that the receiver's own experience keeps confirming. That makes the task model replacement, not information transfer, and it is performed against the receiver's own data — data that validates the wrong model until the moment the right one is required and cannot be installed in time.

The failure has a single substrate. The public updates its risk model on outcome, not exposure — on what happened, not on what was risked. This is the same error the operational literature names in practitioners, relocated to the public: judging safety by whether the house burned rather than by whether exposure was reduced. Everything below is what that updating rule produces when it meets a hazard structured to mislead it.

The rule produces two failures, and they are different kinds of failure. The first is a baseline failure — a tracking failure over time, on the axis of probability. Over a run of quiet years, perceived risk falls while actual risk rises, because the events the receiver can observe understate the hazard the receiver cannot. The second is a magnitude failure — a representation failure in the moment, on the axis of consequence. When the event arrives, it exceeds anything in the receiver's experience, and it exceeds it not in size but in kind. The two failures are not variants of one problem. They are structurally different, and the paper takes them apart separately to show it.

The magnitude failure decomposes into an asymmetric structure: two ways the receiver's own experience fails to build the model, one way externally supplied knowledge fails to install it, and one factor — the cost of accepting the model — that makes every one of those failures more durable. The result is a model that can be neither grown from experience nor installed from outside. This asymmetry is not a defect in the account. It is a finding about how the phenomenon is built.


The load-bearing claim is that the two failures are coupled, and coupled through the warning itself. The signal loud enough to breach the magnitude problem is the same signal that, spent against a rare event, trains the discount that deepens the baseline problem. One channel is asked to do two jobs whose optimal settings run opposite. This is why the familiar prescription — warn only when warranted — does not resolve the bind. It trades one failure for the other.

The recommendation follows from the structure and is deliberately modest. Because the failure lives in perception, and perception is unreliable against this hazard, the durable protection routes around perception entirely, onto the built environment. The strongest version of the finding is that the communication problem is real, bounded, and partly unsolvable at the communication layer — and that the reliable remedy is structural, not rhetorical. That is a harder recommendation than a better warning, and a more defensible one.

FIGURE 1

The two divergences

One updating rule produces two structurally different failures — a tracking failure over time, on the axis of probability, and a representation failure in the moment, on the axis of consequence — coupled through the single instrument meant to correct them.



A community can hold an accurate picture of the fire and still discount it as rare; it can accept a rising base rate and still be unable to picture the event. The two failures are independent, and the warning cannot serve both at once.

2 · Support and Analysis

2.1 The substrate: outcome, not exposure

Begin with the rule, because both failures are the same rule meeting different features of the same hazard.

A person forms a belief about their own safety, and updates it as the world provides feedback. The available feedback is outcome. Did the fire come. Did the house burn. Did the evacuation prove necessary. These are the observable events, and a person who updates on them is doing what updating is for.

The defect is not in the person. It is that outcome is a poor signal for exposure when the two come apart, and for this hazard they come apart severely. Exposure is the risk that was run — the probability and magnitude of loss that the situation actually carried. Outcome is what happened to land. A person can run enormous exposure and draw a benign outcome many times in a row, and each benign outcome is read, correctly, as feedback. The reading is sound. The signal is misleading.

This distinction is not new to this line. The operational chapters draw it as the difference between decision quality and outcome: a decision is good if it was sound given what could be known, regardless of how it turned out, and judging a decision by its outcome is a named error (Coil, OLF Ch. 4). The practitioner version of the error is a commander who concludes a risky call was correct because the crew came home. The public version is a resident who concludes their exposure is acceptable because their neighborhood did not burn. It is the same error, one echelon removed from the fireline, and the line's own diagnostic tool is the key to it. The public is not failing to reason. It is reasoning correctly on a signal that the structure of the hazard has corrupted.

Everything that follows is a specification of how the corruption works. It works in two independent directions, and those directions are the two divergences.


2.2 Divergence one: the baseline

The first divergence is the drift of perceived risk away from actual risk over time. It runs on the probability axis — how likely — and it is a failure that accumulates across seasons rather than one that appears in a moment. The image is the community at the foot of the dormant volcano: the ground is quiet, the quiet is real, and the quiet is not safety.

FIGURE 2

The baseline divergence over time

Across a run of quiet years, perceived risk falls while actual risk climbs. The hazard accumulates as a stock — fuel load, interface density — behind a flat flow of small annual fires, and the event arrives where the gap is widest.



A perfectly calibrated observer of the annual fires would still misjudge the danger, because the quantity that matters was never among the quantities on display.

The baseline divergence is built from four components operating at three levels — the structure of the environment, the updating of the individual, and the composition of the population. The layering is part of the claim. The downward pull on perceived risk is overdetermined: it is produced independently at each level, so a fix that reaches one level leaves the others intact. This is the analytic reason exhortation does not work. A public-education campaign reaches, at best, the individual's updating. It does not touch the environment's structure or the population's composition, and those keep pulling.

Base rate — the environmental component

The annual probability of a catastrophic fire at any given location is low. Discounting a low-probability event is correct. A person who treated every fire season as an even chance of destruction would be miscalibrated in the other direction, and would exhaust

themselves and their neighbors on a risk that, in any given year, almost certainly will not arrive.

The defect is structural, not personal. A correct point estimate of a rare event still understates the risk-weighted exposure, because rarity and consequence pull in opposite directions and the mind holds the rarity more easily than the consequence. This is the well-documented tension in how people weigh rare, high-consequence events: the low probability dominates the judgment, and the severity is discounted along with the frequency (Slovic 1987; Tversky & Kahneman 1974). The signature of this component is actuarial — loss events per unit area per year — and the measurement confirms, rather than refutes, that the receiver is reasoning correctly on a genuinely low base rate. The problem is not that the base rate is misperceived. It is that a correctly perceived base rate is not, by itself, an adequate guide to a hazard whose entire weight sits in the tail.

Stock and flow — the environmental component

This is the deepest of the four, because it defeats even a perfectly calibrated observer.

The hazard does not live in the events the receiver can observe. What a resident sees, across years, is a flow: the annual fires, most of them small, most of them controlled. What actually determines the danger is a stock: accumulated fuel, the density of homes pushed into the wildland edge, the treatment deferred for another budget cycle. The stock is the hazard. The flow is the visible proxy, and it is a poor one. A resident who watched the flow attentively for a decade, and updated faithfully on it, would still misjudge the danger — because the quantity that matters was not among the quantities on display.

This line has already published the structure of this failure at a different scale. *The Reveal* documents an organizational insolvency that builds invisibly: a stock of unmet obligation accumulating behind a flow of apparently manageable annual activity, until the accumulated stock forces a collapse that the flow never signaled (Coil, *The Reveal*). The public version is the same insolvency pointed at a landscape instead of a staff. The fuel load and the interface density accumulate behind the reassuring flow of quiet seasons, and the accumulation is not an event, so it is not observed, so it is not updated on. The signature is a divergence plot: the modeled stock rising while the observed flow stays flat or falls. The widening gap between the two is the hidden hazard, and it is measurable even though it is invisible to the resident living inside it.

Renormalization — the individual component

Base rate is a fact about the world. Renormalization is what the individual's updating does with a run of that fact coming up empty.

Each uneventful season does not merely fail to raise the perceived risk. It lowers it. The internal reference point — the felt sense of what "normal" is — resets to the recent quiet rather than to the actual distribution of outcomes. A person who could, in principle, hold

"rare but catastrophic" as a stable belief instead drifts toward "quiet, as it has been," because the reference point tracks recent experience more tightly than it tracks a remembered statistic.

This is drift, described in the safety literature as the gradual normalization of a widening gap between assumed and actual risk, in which each uneventful step redefines the acceptable baseline downward (Vaughan 1996; Dekker 2011). The organizational versions are catastrophic precisely because no single step looks like a decision to accept more risk. The community version has the same quality. No resident decides, in any given year, that the neighborhood has become more dangerous and that this is acceptable. The reference point simply moves, one quiet season at a time, and the movement is not experienced as a choice. The signature is longitudinal: stated risk perception surveyed over time against constant or rising modeled exposure, showing perceived risk falling while actual risk holds — or mitigation spending decaying across a run of quiet years, as the felt need for it renormalizes away.

Survivorship — the population component

The first three components describe one receiver updating over time. The fourth describes who is left in the population doing the updating, and it operates at a level the other three do not reach.

A community exposed to a rare hazard is not a fixed population. It is a filtered one. Residents who are burned out and do not rebuild, or who are frightened off by a close call, exit. Those who remain are disproportionately those for whom the model "worked" — who stayed, and were not destroyed, and whose continued presence is itself evidence, to them, that staying was sound. The aggregate belief of the community is therefore biased toward safety not only because individuals renormalize, but because the sample from which the aggregate is drawn has been selected on the favorable outcome.

The Reveal names this mechanism at the organizational scale as selection on the favorable tail: a system that promotes or retains on the basis of surviving a risk comes to be populated by those who drew the benign outcome, and mistakes the draw for the skill (Coil, *The Reveal*). The community version is demographic rather than institutional, but the structure is identical. The long-tenure residents whose voices carry the most weight in local risk culture — who have "been here thirty years and seen it all" — are precisely the sample most biased by survivorship, because thirty years without the catastrophic event is exactly the record the rare hazard produces in most locations most of the time. The signature is a cohort comparison: the risk beliefs of long-tenure residents against those who left after events, and the shift in the community's belief composition following a loss.

What the layering establishes

Two of the four components are properties of the environment, one is a property of the individual's updating, and one is a property of the population's composition. Remove any single one, and the discount still stands on the remaining three. This is why the baseline divergence is resistant in a way that a simple information deficit is not. It is not one cause that a campaign could address. It is a stack of independent, reinforcing causes distributed across three levels, and a message can reach at most one of them.

2.3 Divergence two: the magnitude

The second divergence is different in kind. The baseline is a tracking failure that accumulates over time. The magnitude failure is a representation failure that appears in a moment: when the event arrives, the receiver cannot form an adequate picture of what it is. It runs on the consequence axis — how big — and its central feature is that the lethal event is not merely larger than the receiver's experience. It is categorically outside it.

These two divergences are orthogonal, and the orthogonality earns them separate treatment rather than one combined list. A community can hold an accurate picture of the firestorm regime — perhaps because it watched a distant catastrophe closely — and still discount it as rare and not-here. A community can accept a rising base rate and still be unable to represent the event when it arrives. The two fail independently. They meet at exactly one point: the accumulating stock of fuel and interface is both what the baseline fails to track and what makes the eventual event un-representable. But that shared cause produces two distinct failures — one a failure to track probability over time, one a failure to represent magnitude in the moment — and keeping the boundary sharp is what keeps the two sections from bleeding into each other.

The magnitude divergence decomposes into an asymmetric structure. There are two ways for a person to acquire a model of a magnitude they have not lived: grow it from their own experience, or receive it from outside. The decomposition follows those two routes and finds that each is blocked — the experience route by two distinct mechanisms, the transmission route by one — and that a fourth factor operates across both, making whichever block is active more durable. This is a three-plus-one structure, not a symmetric grid, and the asymmetry is developed in Section 2.4 as a finding rather than smoothed over here.

Misclassification — the first experience-route block

The receiver's model of "how bad a fire can be" is anchored on the worst fire the receiver has lived. That anchor does more than set a ceiling on imagined size. It sets a category. When a new event arrives, the receiver files it against the available categories, and the anchor governs which category the event is filed under. A regime-shifted fire, arriving in

a mind whose worst case is a severe structure fire, is filed as a severe structure fire — a bad one, an extreme one, but an instance of the known category.

This is an epistemic failure — a fact about the receiver's classification, not about the world. And it licenses a specific error: the receiver, having filed the event in the familiar category, runs the correct scaling operation for that category. Faced with what is coded as "a very bad fire," the receiver reasons as one reasons about very bad fires — worse than usual, but the same species, calling for a more urgent version of the usual response. The scaling is sound. The category is wrong. The output is a response calibrated to the wrong species of event.

The signature of misclassification is elicitable: ask a resident to describe their imagined worst case, in concrete terms — flame length, rate of spread, whether the house is defensible, whether there is time to leave — and compare the description to the actual regime. The gap between the two is the misclassification, and it is measurable before any fire arrives.

Discontinuity — the second experience-route block

Misclassification, alone, would suggest an obvious fix: show the receiver a bigger fire, and let them scale their anchor upward. Discontinuity is the reason that fix does not work.

The lethal event is not a larger instance of the familiar one. It is a different physical regime. The fire-behavior literature distinguishes fires whose spread is driven by ambient wind across surface fuels from fires whose behavior is dominated by their own convective plume — fires that generate their own weather, that loft embers in mass and cast them far beyond the flaming front, in which structures ignite from other structures largely independent of the wildland fuel around them (Rothermel 1991; Werth et al. 2011; Cohen 2000). The transition from the first regime to the second is not a change in degree. It is a change in the governing dynamics — the rules that determine how the fire moves and what it does to what it touches.

The consequence for the receiver is precise. A model built on the first regime cannot be extended to the second by scaling, because the target is not on the axis being scaled. Correct upward extrapolation from a lived anchor — "the worst I saw, but worse" — still lands wrong, because the worst-plus-worse is a bigger fire of the known kind, and the lethal event is a different kind. This is why misclassification and discontinuity are two blocks and not one. Misclassification says the receiver files the event in the wrong category. Discontinuity says that even correct scaling from the anchor would not reach the event, because the event is off the axis.

This block has a property the others lack: it is documented outside the receiver entirely. The discontinuity between fire regimes is a claim about fire, established in the fire-behavior literature, not a claim about public psychology. That makes "the public cannot extrapolate to the lethal regime" a statement about the structure of the world the public is

reasoning about, not a statement about a deficiency in the public. The receiver cannot get there from experience because experience of one regime does not contain the other, and no amount of reasoning within the first regime crosses into the second.

A caution belongs on the word regime, and it is developed in the Red Team: the discontinuity claimed here must be a real change in the fire's governing dynamics, established by fire behavior, not merely an event that was surprising or extreme. Extremity alone would not close the experience route, because an extreme instance of a known category is still, in principle, reachable by scaling. The claim requires the genuine discontinuity, and it stands or falls on the fire science, not on the severity of any particular loss.

THE REGIME IS REAL.

Plume-dominated fire generates its own wind, casts embers in mass well beyond the flaming front, and ignites structures from other structures rather than from the wildland fuel around them. This is documented fire behavior, not rhetorical emphasis. A resident who files the arriving event as a severe version of a familiar fire is not being careless. A severe fire is the worst event their experience contains, and the event belongs to a regime that experience does not.

Construal distance — the transmission-route block

If the model cannot be grown from experience, the remaining route is to install it from outside — to tell the receiver what the regime is, and have the telling take. Construal distance is why the telling does not take.

A threat held as distant — distant in time, in space, or in likelihood — is represented abstractly. This is the central finding of construal-level theory: psychological distance and abstraction move together, so that the more distant an object is held to be, the more it is represented in general, schematic, decontextualized terms rather than in concrete, perceptual, particular ones (Trope & Liberman 2010). A hazard filed as "unlikely, and not here, and not now" is filed as distant, and distance strips the representation of exactly the concrete detail that would make the regime legible. The receiver can be told, accurately, that the fire will generate its own weather and ignite the house from three directions at once — and the telling degrades, in the held representation, to a vague and abstract "very dangerous," because the threat is filed as distant and distant threats are held abstractly.

This is a transmission-route block specifically. It is not that the information is unavailable or wrong. It is that externally supplied information, however concrete at the point of delivery, thins toward abstraction in a mind that has filed the threat as distant — and the thinned representation is not adequate to drive the response the regime requires. The signature is the standard construal paradigm: manipulate the psychological distance at which a threat is held, and measure the concreteness of the representation that results.

Motivated abstraction — the modulator across both routes

The fourth factor is not a fifth block, and the difference is the crux of the asymmetry.

Accepting an adequate model of the lethal regime has a cost. The concrete picture — the house lost, the road closed, the night of it — entails an aversive and expensive action: leaving, spending on mitigation, disrupting a life deliberately built at the foot of the volcano. Because the concrete picture entails the costly action, the receiver has an interest in not forming the concrete picture. Keeping the threat abstract, ambiguous, or filed as continuous with the familiar is a way of not being obligated to act.

This factor does not sit on one route. It operates on every other block, and that is what makes it a modulator rather than a coordinate. It raises the persistence of misclassification: a receiver who could, with effort, recognize the event as out-of-category has an interest in not making the effort. It raises the persistence of the experience route generally: a receiver who could extrapolate from a near-miss has an interest in declining to. And it raises the persistence of construal distance: the abstraction that distance produces automatically is abstraction the receiver has an independent reason to preserve. Each block is made stickier by the receiver's interest in the block remaining in place.

The distinction between motivated abstraction and construal distance is worth holding, because they can look alike. Construal distance is cold: distance produces abstraction mechanically, regardless of what the receiver wants. Motivated abstraction is warm: the receiver has a stake in the abstraction, and works, non-consciously, to maintain it. The tell that they are two mechanisms and not one is an interaction. Hold psychological distance constant and raise the cost of the implied action, and the resistance to concretizing rises with the cost — which a purely distance-driven account cannot explain, because distance has not changed. The motivation responds to the cost of acting, not to the distance of the threat, and that independent responsiveness is the evidence that it is its own factor.

Motivated abstraction is also the component that most strains this paper's local-rationality commitment, and Section 3 addresses the strain directly rather than assuming it away.

Route closure is functional, not numerical

The asymmetry of this structure — two blocks on one route, one on the other, plus a modulator — invites the objection that the account is unbalanced and therefore incomplete. The objection mistakes the claim.

The load-bearing claim of the magnitude divergence is that both acquisition routes are closed: the model can be neither grown from experience nor installed from outside. That is a functional condition, not a numerical one. It requires at least one effective block on each route. It does not require an equal number of blocks on each route. The experience

route is closed by two mechanisms because it happens to fail for two independent reasons — the wrong category, and the unreachable regime. The transmission route is closed by one because construal distance is sufficient to close it. The modulator then raises the durability of whichever block is operating, on either route. Both routes are shut. The count on each side is a fact about how each route fails, not a requirement the account must satisfy for symmetry's sake.

A note on a route that is not a third route. A reader may ask about inference — whether a receiver could reason their way to the lethal regime without either experiencing it or being shown it. Inference is not a third source of model content. It is an operation performed on content supplied through experience, transmission, or both. Its premises here are physical and behavioral facts about fire, which reach the receiver only by experience or transmission — both already blocked — and its aversive conclusion is exactly what the modulator resists. Inference therefore does not open a route around the closure. It runs on inputs the closure has already shut, toward a conclusion the modulator already opposes.

2.4 The two divergences are not parallel

It would be tidy to present the baseline and magnitude divergences as two instances of one template. They are not, and the paper claims the difference as a finding rather than hiding it.

At the level above the mechanism, the two are the same. Both are what the outcome-updating rule produces when it meets a hazard structured to mislead it. The baseline is that rule failing to track probability over time; the magnitude is that rule failing to represent consequence in the moment. They share a generative logic. In that sense they are siblings.

Below that level, they are built differently. The baseline divergence is a stack of independent, reinforcing causes distributed across three levels — environmental, individual, population — each pulling perceived risk downward, no one of them sufficient to lift it if removed. The magnitude divergence is a route-closure structure — two acquisition routes, each blocked — with a cross-cutting persistence factor. These are not the same architecture. One is an accumulation; the other is a closure. Their kinship is at the level of process, not at the level of component geometry.


The general form of this is worth stating as a structural claim: two failures may be homologous at the level of the generative process while remaining heterogeneous at the level of mechanism. Forcing them into a shared component structure — insisting the magnitude divergence must also have three levels, or the baseline must also have a modulator — would impose a symmetry the phenomenon does not have, and would obscure the real result, which is that a single updating rule fails in two structurally

distinct ways against a single hazard. The decomposition is more faithful for admitting the asymmetry than it would be for manufacturing balance.

FIGURE 3

The magnitude failure is a 3+1 structure

Two acquisition routes feed the receiver's model of the regime. The experience route is closed by two independent mechanisms; the transmission route by one. A cross-cutting modulator — the cost of accepting the model — raises the durability of whichever block is active.



Closure is functional, not numerical. Both routes are shut; the count of blocks on each is a fact about how each route fails, not a symmetry the account must satisfy.

2.5 The coupling

The two divergences have been treated separately because they are separate failures. The central claim of the paper is that they are nonetheless coupled, and coupled through the one instrument meant to fix them: the warning.

Each divergence makes a demand on the warning, and the two demands run opposite. The magnitude divergence demands a signal that breaches the representation problem — loud, vivid, concrete, and rare enough to register as extraordinary, because a routine signal is filed in the routine category and scaled as routine. The baseline divergence demands the opposite: a signal held in reserve, because a signal spent often, against a rare event, trains the discount. Every warning that fires and is not followed by the catastrophic

outcome is a data point in the receiver's outcome-updating — and the data point says the warning overstated the risk. Spend the loud signal frequently, and the receiver learns, correctly on the outcomes observed, to discount it. The learning is sound. The signal has taught the discount.

This is one channel asked to do two jobs whose optimal frequencies are opposed. Breaching the magnitude ceiling wants the signal loud and used sparingly enough to stay extraordinary. Protecting the base rate wants the signal reserved to the point of near-silence, because every use that does not cash out in catastrophe erodes it. The same instrument cannot be simultaneously loud-and-rare-and-extraordinary and reserved-to-near-silence in the way each divergence separately requires.

This is the analytic content behind the familiar phrase "crying wolf." The craft literature already knows that over-warning erodes credibility, and prescribes the obvious correction: warn only when warranted. What the decomposition adds is that this is not a calibration problem with a correct setting waiting to be found. It is a structural bind. The prescription "warn only when warranted" resolves the baseline demand at the cost of the magnitude demand: reserve the signal, and it retains its meaning, but it is now used so rarely, and the standard for "warranted" is so high, that the receiver has little chance to build the model the magnitude divergence says they lack. Reserve protects credibility and starves representation. Spend builds representation and destroys credibility. The channel cannot serve both, and moving the dial does not find a setting that does. It only chooses which failure to accept.

There is a partial resolution, and its limit is instructive. The bind is a bind only if a single channel carries both jobs. Separate the channels, and the opposition eases. An ambient, chronic signal can do model maintenance — a steady, low-intensity message that the community lives in a fire environment, that the regime exists, that the baseline is what it is — without ever spending the maximal signal. A distinct, reserved, bright-line trigger can do action — one signal, fired rarely, meaning one thing: this line, now, go. The maximal signal is never spent on ambient risk, so it does not erode; the ambient signal never tries to carry the action, so its frequent use does not train a discount against the trigger. The two jobs are assigned to two instruments with two frequencies, and the opposition that made them incompatible on one channel is relaxed.

The limit is that fire does not always cash the trigger. The bright line is crossed, the order is given, and then the wind turns or the fire lays down and the neighborhood is spared. One uncashed trigger is a benign outcome following a maximal signal — and the receiver, updating on outcomes, registers exactly the data point the reserved channel was designed to avoid. The channel separation reduces the coupling. It does not sever it, because the trigger is still a signal, and a signal that fires without the outcome still teaches the discount, even a signal reserved for the bright line. This limit is not a flaw to be hidden. It belongs in the Red Team, because it bounds the one resolution the structure offers.

FIGURE 4

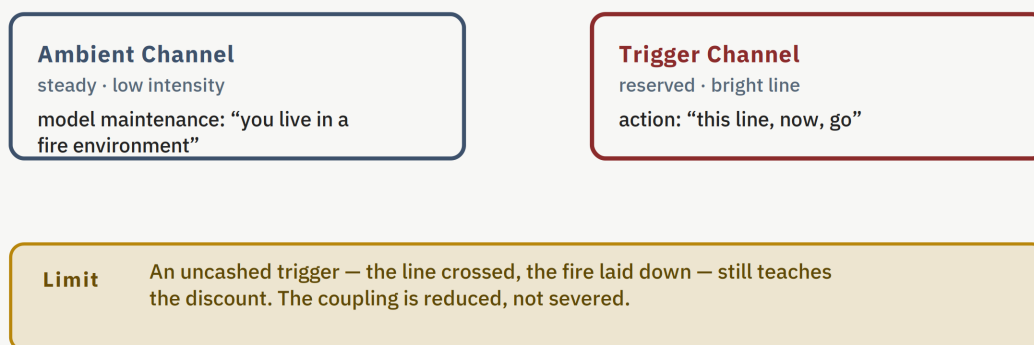
The coupling, and its partial resolution

One channel cannot serve both demands: spend the signal to build the model and it trains the discount; reserve it to protect credibility and it starves the model. Separating the ambient maintenance signal from the reserved trigger relaxes the bind — but does not sever it.

One Channel — The Bind



Two Channels — The Partial Resolution



The dial has no correct setting. Moving it toward “spend” buys the magnitude demand at the baseline’s expense, and toward “reserve” the reverse. The bind is structural, not a calibration error.

Warning is a mitigation of a problem it cannot solve. The reliable protection against a hazard the public cannot adequately perceive is a built environment that does not depend on the public perceiving it.

3 · Red Team Assessment

The analysis above is a structural account built largely from constructs imported into the fire domain and anchored on a single hazard type. That construction has specific vulnerabilities, and this section names them before a reader does.

3.1 The author's position

I command incidents, and issuing warnings to the public is part of that job. That places me on one side of the transaction this paper analyzes, and it introduces a specific bias vector: an account that locates the failure in the structure of the hazard and the receiver's updating is an account that does not locate it in the quality of the warnings issued by people like me. The reader should weigh the analysis with that interest visible. The internal check I can offer is that the paper's recommendation runs against the flattering reading — it concludes that better warning, the thing my side of the transaction controls, is not the reliable remedy, and routes the burden onto the built environment instead. An account written to protect the warner would not end there. But the vector is real, and it is declared rather than dissolved.

3.2 The single-hazard anchor

The entire structure is developed on wildland-urban interface fire. Every component is illustrated with fire, and the discontinuity block depends on a claim from fire-behavior science specifically. This is one hazard, and a structural account built on one case does not, by itself, establish that the structure generalizes. The volcano is a figure of speech in this paper; whether the same four-component baseline and the same three-plus-one magnitude structure describe flood, seismic, or pandemic risk is asserted nowhere and demonstrated nowhere. The honest scope is that this is the anatomy of one hazard's warning failure, and the generative logic — outcome-updating against a rare, latent, regime-shifting hazard — is a hypothesis about a class, not a proven property of it. Drafting should not write as though the fire case proves the class.

3.3 Does misclassification depend on discontinuity?

The sharpest internal objection is that misclassification and discontinuity are not two blocks but one, because misclassification only matters when a discontinuity is present to be missed. If the objection holds, the experience route has one block, not two, and the paper has manufactured an independence it does not have.

The defense is dissociation, and it requires showing that each occurs without the other. Misclassification without discontinuity: a very large fire, still within the wind-driven regime — no plume domination, no regime change — is filed by a receiver as an ordinary fire and scaled as one, and the receiver under-responds to a large-but-continuous event. The failure is real and the discontinuity is absent; the receiver simply mis-binned a big instance of the known kind. Discontinuity without misclassification: a receiver correctly senses that the arriving event is not normal — "this is not like anything I've seen" — and still has no adequate model of how the regime behaves or what it will do. The receiver has not mis-filed the event; they have correctly declined to file it in the familiar category, and are left without a category at all. The two dissociate, so they are two blocks. But the

objection identifies a real drafting hazard: the illustrations must include the pure cases, or the independence will read as asserted rather than shown.

3.4 Does construal distance also touch experience?

Construal distance is placed on the transmission route, but psychological distance plausibly shapes how a receiver interprets their own near-misses too — a close call held as "distant" in memory is construed abstractly, thinning the lesson it might have taught. If construal distance operates on the experience route as well, its placement as a transmission-specific block is unstable.

The defense narrows the mechanism rather than denying the overlap. On the transmission route, construal distance blocks the installation of externally supplied content: the concrete message thins to abstraction in a mind holding the threat as distant. That is a distinct operation from whatever distance does to the memory of a lived event, which is a modulation of experience-route content, not a block on transmission-route installation. The cleanest handling is to define the transmission-route block narrowly as the thinning of supplied content, and to treat any distance effect on near-miss memory as a separate and secondary matter — related, but not the same operation, and not what closes the transmission route. The objection is a warning against defining construal distance so broadly that it floats free of the route it is assigned to.

3.5 Does motivated abstraction explain too much?

A factor that raises the persistence of every other block, on both routes, is a factor at risk of explaining everything and therefore discriminating nothing. If motivated abstraction is invoked wherever a block persists, it becomes unfalsifiable — a name for persistence rather than an account of it.

The discipline is to specify where it should and should not be expected, and to make the specification testable. Motivated abstraction should be expected in proportion to the cost of the implied action: high where leaving is expensive, where the home is the primary asset, where the resident's life is deeply built into the place; low where the cost of the protective action is low. That predicts a specific, checkable pattern — resistance to concretizing the threat should scale with the personal cost of acting on it, holding the threat's distance and the receiver's information constant. Where the protective action is cheap and the receiver still resists the concrete picture, motivated abstraction is the wrong explanation and construal distance or misclassification is doing the work. The factor earns its place only if it is bounded this way; unbounded, it is the objection's target, and correctly so.

3.6 The two-route exhaustiveness claim

The magnitude analysis rests on there being two, and only two, ways to acquire a magnitude model — experience and transmission — with inference reduced to an operation on their outputs. If a reader can name a third genuine source of model content, the route-closure claim is incomplete.

The defense rests on a distinction that must be explicit in the account: the difference between a source of model content and an operation on content. Experience and transmission are sources — they introduce content the receiver did not have. Inference, analogy, and imagination are operations — they transform content already present, and cannot manufacture content from nothing. The claim of exhaustiveness is a claim about sources, and it is only as strong as that distinction is clear. If the account leaves "source" and "operation" blurred, the exhaustiveness claim is exposed; if it draws the line cleanly, the claim holds, because any candidate third route resolves into either a new source (which would need naming and does not appear to exist for magnitude models) or an operation on the two sources already treated.

3.7 Is "regime" too strong a word?

The discontinuity block leans on the claim that the lethal event is ontologically different — a real change in the governing dynamics of the fire, not merely a surprising or extreme instance. A reader may argue that "regime" and "ontological" overstate what is really a matter of degree, or that the discontinuity is model-relative — a discontinuity in our description, not in the world.

The defense is to define the term at the right strength and to rest it on the fire science. The claim is not that the event is surprising, and not that it is extreme. It is that the physical dynamics governing the fire's behavior change — that plume-dominated fire operates under different governing relationships than wind-driven surface fire, established as a matter of fire behavior and not of observer psychology. The fire-behavior literature supports that reading: whether a fire is wind-driven or plume-dominated is set by the balance between the power of the fire's own convection and the power of the ambient wind, and the two operate under different governing relationships rather than along a single scale of degree (Rothermel 1991; Werth et al. 2011). Were the fire science to support only a difference of degree, the discontinuity block would weaken toward the misclassification block and the experience route would close on one mechanism rather than two; because it supports the stronger reading, "regime" is warranted and the experience route is closed by a fact about fire. The block's strength rests on that fire-science distinction, not on the severity of any particular loss.

3.8 The evidence base

Most of the structure above is imported theory — construal-level theory, the risk-perception literature, the drift and normalization literature — applied to fire, plus fire cases used as illustration. Fire-specific empirical research at the level of these public-perception mechanisms is thin, and the catastrophic cases that motivate the paper are after-the-fact reconstructions, not designed studies. The account's empirical claims should be calibrated to that standard: the imported constructs are established in their home literatures, the fire application is argued rather than measured, and the signatures named for each component are proposals for how one would test the claim, not reports of tests already run. The paper's own evidence standard, applied to the paper, requires that the borrowed authority of the source literatures not be presented as fire-specific confirmation it does not yet have.

4 · Dissenting Alternative

The credible alternative is not a quarrel with any single component. It is a competing account of the whole phenomenon, and it is strong enough that the recommendation concedes part of it.

The dissent is the revealed-preference account. It holds that the paper has mistaken a preference problem for a cognition problem. On this reading, the resident who does not act is not running a broken model. They have weighed the costs and made a choice: the expense and disruption of leaving or hardening the home, set against a risk they judge — correctly, in most years — to be low, and they have decided the risk is worth bearing. That is not a failure of representation. It is a preference, and a legitimate one. People are entitled to accept risk to their own property in exchange for the value they get from living where and how they choose. The paper's elaborate cognitive machinery, on this view, is an overbuilt explanation for something simple: people know the risk and accept it, and calling that acceptance a "misperception" is the analyst substituting his risk tolerance for theirs.

This dissent is strong for a specific reason: the paper has already half-conceded it. Motivated abstraction is, on the paper's own account, a rational response to the cost of the protective action — the receiver keeps the threat abstract because concretizing it obligates an expensive act. The dissent presses exactly there. If the resistance to acting is a rational response to cost, then it is not a misperception at all; it is a preference operating as designed, and the paper has smuggled a preference into a cognitive frame by calling it a "modulator" of "blocks." The most honest version of the paper's own local-rationality commitment, the dissent argues, leads to the dissent's conclusion: everyone is behaving rationally, so there is no cognitive failure to explain, only a distribution of risk preferences the analyst happens to disagree with.

The dissent fails as the whole account, and the reason it fails is precise. A preference is only a preference if it is formed over an accurate representation of what is being chosen. A resident who declines to act because they have weighed the firestorm regime against the cost of leaving, and chosen to bear the risk, is exercising a genuine preference, and the paper has no quarrel with them. But the magnitude divergence is the claim that most residents cannot form an accurate representation of the firestorm regime — that the regime is unreachable from experience and thins to abstraction in transmission. A choice made over a representation that tops out at "a bad structure fire" is not a preference to bear the firestorm regime. It is a preference to bear a different, smaller thing that the resident has mistaken for the real one. The dissent assumes the accurate representation the paper spends Section 2.3 showing the receiver does not have. Where the representation is accurate, the dissent is right and the behavior is preference. Where the representation is the under-scaled one the magnitude divergence describes, the dissent is describing a preference over the wrong object.

The concession the dissent earns is real, and it is adopted directly into the recommendation. The dissent is correct that where representation is adequate, non-action is preference, and preference is not the analyst's to overrule. That means the goal of a remedy cannot be to force the protective action against a preference. It can only be to ensure the preference is formed over an accurate object — and, failing that, to protect the resident in a way that does not depend on the preference at all. This is precisely why the recommendation routes toward exposure-side, structural protection: hardening the built environment protects the resident whether their representation is accurate or not, and whether their preference is to act or not, which respects the autonomy the dissent correctly defends while covering the case the dissent cannot — the resident whose "preference" is an artifact of an under-scaled model. The dissent moves the recommendation away from persuasion and toward structure. It does not overturn the analysis; it disciplines the fix.

A second, weaker dissent deserves brief acknowledgment: the craft account, which holds that the failure is one of message quality and that well-designed, well-targeted, repeated risk communication does move behavior, so the paper overstates the structural closure to make the problem look harder than it is. This account is not wrong that message quality matters, and the paper does not claim warnings are useless. It claims they are insufficient against this hazard structure, and the coupling in Section 2.5 is the specific reason improving the message does not resolve the bind: a better message spent more often still trains the discount, and a better message held in reserve still starves the representation. The craft dissent improves the message; it does not escape the channel bind, which is a property of the structure and not of the message's quality.

5 · Recommended Course of Action

The recommendation is short, and its brevity is a claim. The contribution of this paper is the anatomy of the failure, not the remedy for it, and an anatomy that is honest about a hard structure should not close with a confident fix. What follows is not a set of prescriptions. It is the set of moves the structure permits, arranged in rising order of how much they concede — and the concession they build toward is that the reliable remedy is not a communication remedy at all.

5.1 Pre-commitment decided cold

The magnitude divergence says the receiver cannot summon an adequate representation in the window where the decision must be made. The first move takes the decision out of that window.

This line already carries the instrument, pointed at practitioners: decompose a master judgment into its load-bearing components, pair each with an observable indicator, and pre-decide the response to each indicator before the moment arrives (Coil, *Order It and It Comes*). Turned outward, the same discipline decomposes "am I safe" into observable, pre-specified conditions — a named fire on a named ridge, a specific evacuation level, a particular wind — and attaches a pre-decided action to each, settled on a cold day when the representation is not needed and imagination is not required. The receiver who has pre-committed does not have to picture the firestorm in the moment. They have to recognize a condition and execute a decision already made. This is the public version of a checklist against a known cognitive failure: it works precisely because it does not rely on the receiver's real-time judgment, which the paper has shown is not available.

The limit is that pre-commitment must be made when the threat is distant, and Section 2.3 shows that distance is when the threat is held most abstractly and the motivation to keep it abstract is strongest. The move fights the very conditions that make it necessary. It is a real move, and a bounded one.

5.2 Borrowed outcomes

If the receiver's own outcome library caps below the lethal regime, the only way to add an outcome at the right magnitude is to borrow one. A distant catastrophe — a Paradise, a Lahaina — supplies an outcome to a library that has none, and vicarious experience of a regime is the nearest available substitute for the experience the receiver cannot and should not acquire directly.

The move is real, and it is perishable. A borrowed outcome fades as it recedes in time, and it is subject to the same discount as any distant event — filed as "there, not here," construed abstractly, and weakened by the receiver's interest in keeping it that way. Borrowed outcomes can lift the ceiling temporarily; they do not install a durable model,

because the borrowing is itself a transmission, and transmission is a blocked route. The lift is partial and it decays, which is why it cannot be the load-bearing remedy.

5.3 Change the exposure, not the belief

The two moves above operate on the receiver's cognition, and the paper has spent its length showing that the receiver's cognition, against this hazard, is unreliable in specific, structural ways. The third move draws the conclusion the structure forces: route the protection off the unreliable variable and onto a reliable one.

Defensible construction and ignition-resistant materials protect a structure regardless of whether its occupant perceives the risk correctly, forms an accurate model, holds an adequate representation, or prefers to act. A hardened home does not require its owner to imagine the firestorm. It survives the ember cast whether the owner evacuated with a clear picture of the regime or fled at the last minute with no picture at all. This moves the burden of protection off perception — which the paper has shown is unreliable against this hazard — and onto the built environment, which is durable and inspectable and indifferent to the occupant's mental model.

This is the strong form of the recommendation, and it follows from the analysis rather than being appended to it. If the failure is structural in the receiver's perception, then a remedy that operates through the receiver's perception inherits the structural failure, and a remedy that operates on the physical exposure does not. The recommendation is therefore not "warn better." It is that warning is real, necessary, and bounded — a mitigation of a problem it cannot solve — and that the reliable protection against a hazard the public cannot adequately perceive is a built environment that does not depend on the public perceiving it. That is a harder recommendation than a better message, and it is the one the structure supports.

5.4 What would show this wrong

Per the line's discipline, the recommendation states in advance what evidence would count against it.

The account is wrong if fire-specific measurement of public risk perception fails to find the signatures the components predict — if perceived risk does not fall across quiet years against constant modeled exposure, if elicited worst-case representations do not gap systematically below the actual regime, if resistance to concretizing does not scale with the cost of the implied action. It is wrong if the fire-behavior literature does not support a genuine discontinuity in governing dynamics, in which case the discontinuity block collapses toward misclassification and the experience route closes on one mechanism. It is wrong about its remedy if hardened structures do not, in fact, outperform occupant perception as a predictor of survival — if the built-environment variable turns out to be a weaker protection than the perception variable it is meant to replace. And the

generalization beyond fire is wrong if the four-component baseline and the three-plus-one magnitude structure fail to describe the warning failures of other rare, latent, regime-shifting hazards, in which case the account is the anatomy of one hazard and not of a class. Each of these is checkable, and each would falsify a specific part of the structure rather than the whole, which is the property a decomposition is supposed to have.

5.5 Scope

This paper is the analysis of one clock. It treats the chronic problem — the standing misperception of a rare, latent, regime-shifting hazard by the public exposed to it — and it does not treat the acute problem of how a commander allocates a bounded defense and communicates that allocation in the hours of an incident. That is a different object, on a different timescale, and it belongs to a different paper. Nor is this a manual for warning craft, alert-system design, or public-information practice; those are staffed functions with their own literatures, and the account here sits beneath them, explaining why the failure they manage persists even when they are performed well. The paper's claim is bounded to what it has taken apart: a single updating rule, failing in two structurally distinct ways against a single hazard, coupled through the one instrument meant to correct it — and reliably countered not by that instrument but by the built environment that does not depend on it.

References

- Cohen, J. D. (2000). Preventing disaster: Home ignitability in the wildland-urban interface. *Journal of Forestry*, 98(3), 15–21.
- Coil, J. (2026). *The Operational Leader's Field Guide*. Flagstaff, AZ: Cogplexiti.
- Dekker, S. (2011). *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*. Farnham, UK: Ashgate.
- Rothermel, R. C. (1991). *Predicting Behavior and Size of Crown Fires in the Northern Rocky Mountains* (Research Paper INT-438). Ogden, UT: USDA Forest Service, Intermountain Research Station.
- Slovic, P. (1987). Perception of risk. *Science*, 236(4799), 280–285.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago, IL: University of Chicago Press.
- Werth, P. A., Potter, B. E., Clements, C. B., Finney, M. A., Goodrick, S. L., Alexander, M. E., Cruz, M. G., Forthofer, J. A., & McAllister, S. S. (2011). *Synthesis of Knowledge of Extreme Fire Behavior: Volume I for Fire Managers* (General Technical Report PNW-GTR-854). Portland, OR: USDA Forest Service, Pacific Northwest Research Station.

*The absence of events is read, correctly, as evidence —
and the evidence points the wrong way.*

A community lives at the foot of a dormant hazard. For years, nothing happens, and every quiet year teaches the same lesson: this place is safe. Then the event arrives, and it is not a larger version of anything the community has seen. It is a different thing entirely — and the model that served for years is useless in the window where it is finally needed.

The reliable remedy is structural, not rhetorical.

The Foot of the Volcano takes this failure apart. The public updates its risk model on outcome, not exposure — on what happened, not on what was risked — and that single rule produces two structurally different failures: perceived risk falling as actual risk climbs, and an event that exceeds the receiver's model in kind, not size. The two are coupled through the warning itself. The signal loud enough to breach one failure trains the discount that deepens the other. One channel is asked to do two jobs whose optimal settings run opposite.

The analysis runs the full discipline of the line — argument, red team, dissenting alternative, recommendation — and closes with a deliberately modest set of moves: pre-commitment decided cold, borrowed outcomes, and durable protection routed around perception entirely, onto the built environment. Written for emergency managers, public information officers, and leaders who must warn communities about a hazard their own experience keeps disproving.

Part of the Cogplexiti Systems & Decision Analysis line, alongside The System Tier, The Incident Tier, and Nobody's Job.

ABOUT THE AUTHOR

Jayson Coil is the founder of Cogplexiti LLC, an Assistant Chief with the Sedona Fire District, and a Type 1 Operations Section Chief on Southwest Incident Management Team 1. A U.S. Army veteran and red team instructor, he works at the intersection of operational command and the study of how organizations decide under uncertainty. He has helped develop the systems agencies use to manage major incidents and has presented for the National Fire Academy, the U.S. Forest Service, the Department of the Interior, FEMA, the Haas School of Business at UC Berkeley, and Los Alamos National Laboratory.



Cogplexiti

