

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

The Cost of Believing

Why a Parasite Moves the Public and a Firestorm Does Not, and What Comes Next

JAYSON COIL

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

WHITE PAPER

The Cost of Believing It

Why a Parasite Moves the Public and a Firestorm Does Not, and What Comes After Hardening

JAYSON COIL

Cogplexiti

Flagstaff, Arizona · cogplexiti.com

Copyright © 2026 Jayson Coil. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

This paper presents the author's independent analysis. It does not represent the official position, policy, or endorsement of any agency, department, or organization with which the author is or has been affiliated.

Cogplexiti · Systems & Decision Analysis — White Paper

ISBN: 979-8-9970824-7-5

First Edition · 2026

This paper supersedes the working draft circulated as The Foot of the Volcano. The earlier draft analyzed one hazard. This edition adds a second, and the comparison changes the finding.

Contents

How to Read This Paper	5
A note on stance.....	5
A note on evidence and citation	5
1. Executive Summary	7
2. Support and Analysis.....	9
2.1 The substrate: outcome, not exposure	9
2.2 The comparison case: a hazard the public acts on	10
2.3 Divergence one: the baseline	11
Base rate – the environmental component	12
Stock and flow – the environmental component	13
Renormalization – the individual component.....	13
Survivorship – the population component.....	14
What the layering establishes	15
2.4 Divergence two: the magnitude	15
Misclassification – the first experience-route block.....	16
Discontinuity – the second experience-route block	16
Construal distance – the transmission-route block	17
The cost of believing – the modulator across both routes	18
Route closure is functional, not numerical.....	20
2.5 What the comparison isolates.....	21
2.6 The coupling	22
2.7 The same failure under different institutions.....	24
3. Red Team Assessment	25
3.1 The author’s position.....	25
3.2 The two-hazard base	25
3.3 Is the comparison case fair?.....	26
3.4 Does misclassification depend on discontinuity?	26
3.5 Does construal distance also touch experience?	27
3.6 Does the cost of believing explain too much?	27
3.7 The two-route claim.....	28
3.8 Is <i>regime</i> too strong a word?	28
3.9 The evidence base	29

3.10 The levers in Section 5.4 are not costless	29
4. Dissenting Alternative.....	30
5. Recommended Course of Action	32
5.1 Pre-commitment decided cold.....	32
5.2 Borrowed outcomes	32
5.3 Change the exposure, not the belief.....	33
5.4 After hardening: four levers that are not mandates	33
Lever one: price the parcel, not the territory.....	34
Lever two: intervene at the transaction.....	35
Lever three: use the rebuild window.....	35
Lever four: make the hardened option the default.....	36
5.5 Why the timing matters	37
5.6 What would show this wrong	38
5.7 Scope.....	39
References.....	40

How to Read This Paper

This paper belongs to the Cogplexiti Systems and Decision Analysis line. It follows the line's five-part structure: the argument, the challenge to it, the credible alternative, and the recommendation. The structure lets a reader weigh the case rather than receive it.

The subject is one failure. A member of the public gets a correct warning, reasons soundly, and still does not take the protective action. This paper asks why that happens, and why it keeps happening even when the warning is good.

The paper does not cover message design, public information craft, or how alert systems work. Those are real subjects with their own literature and their own professionals. This paper sits one layer beneath them.

One thing is new in this edition. The failure is examined against a second hazard, where the public does act. The comparison is not decoration. It isolates the one variable that matters most, and that variable drives the recommendation.

A note on stance

This is analysis, not doctrine. Nothing here directs anyone to do anything.

Every actor is treated as locally rational. Local rationality means people act sensibly given what they know and what they face. The question is never who failed. It is what conditions made the outcome likely.

That commitment does real work here. The subject is a public that does not act. The easy explanation is that people are complacent, or uninformed, or in denial. That is the explanation this paper is built to replace. Complacency is a verdict. This paper looks for a mechanism.

Section 5 argues that the timing matters, and that some moves cannot wait. That is a claim about a rate of change, and it is supported by evidence. It is not an appeal, and the reader should hold it to the same standard as every other claim here.

A note on evidence and citation

Constructs borrowed from other fields are attributed to their source literature at the level of their established claims. Fire behavior claims are anchored to fire science at the same standard. Full sources appear in the References.

The foodborne case is treated as a class, not as one specific outbreak. What matters is the response pattern, and that pattern repeats. No case counts, product brands, or dates are

claimed here. The argument does not need them, and unverified numbers would weaken it.

European fire figures are provisional and are dated in the text at the point of use. Season totals climb, and national and satellite estimates differ in method and date. Where a national ministry figure and a satellite figure disagree, both are reported rather than reconciled.

1. Executive Summary

A community lives at the foot of a quiet volcano. For years, nothing happens. The quiet is real, and people read it as evidence. The evidence points the wrong way. Then the event arrives, and it is not a bigger version of anything the community has seen. It is a different thing. The model that worked for years is useless in the hour it is finally needed.

That is the standing problem of public warning for rare, high-consequence fire.

Now set a second hazard beside it. A foodborne parasite turns up in a packaged produce item. The warning goes out. People stop buying the product. Stores clear the shelves. The response is fast and broad, and it is often larger than the health risk by itself would justify.

Same public. Same week. Opposite errors.

One hazard is under-responded to, at enormous cost. The other is over-responded to, at almost none. Any account that explains only the fire case is incomplete, because the same people are making both calls.

The two errors have one source. People update their sense of safety on outcome, not exposure. Outcome is what happened. Exposure is the risk that was actually run. The two are not the same, and for some hazards they come apart badly.

This is the same error the operational literature names in commanders, moved one level down to the public. A commander who calls a risky decision good because the crew came home is making it. So is a resident who calls their exposure acceptable because the neighborhood did not burn. Neither is careless. Both are reasoning correctly on a signal the hazard has corrupted.

Against fire, that one rule produces two different failures. The first is a baseline failure. It runs on probability, and it builds over years. Across a run of quiet seasons, perceived risk falls while actual risk climbs, because the fires people can see understate the hazard they cannot.

The second is a magnitude failure. It runs on consequence, and it appears in a moment. When the event arrives, it exceeds anything in the person's experience. It exceeds it not in size but in kind.

The two failures are not versions of one problem. They are built differently, and the paper takes them apart separately to show it.

The magnitude failure has an uneven shape. There are two ways for a person to acquire a picture of something they have not lived: grow it from their own experience, or receive it from someone else. Two mechanisms block the first route. One blocks the second. A fourth factor makes whichever block is active harder to shift. The result is a picture that can be neither grown nor installed.

The load-bearing claim about fire 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 on a rare event, trains the discount that deepens the baseline problem. One channel is asked to do two jobs whose best settings run opposite. This is why the familiar advice, warn only when warranted, does not solve the bind. It trades one failure for the other.

The comparison case then explains why the same rule produces the opposite error for the outbreak.

The parasite arrives already concrete. It is in food a person ate this week, in a body they live in, from a company with a name. Nothing has to be imagined. And the action it asks for is nearly free: skip one product for two weeks. Believing it costs almost nothing.

Fire is the reverse on both counts. The lethal regime cannot be pictured from experience, and the action it asks for is expensive: leave, spend, or disrupt a life deliberately built at the foot of the volcano. That isolates the variable that carries the most weight in this paper, and gives it its title. It is not what people know. It is what it costs them to believe it.

The 2026 European season supplies a second test. The same failure is now visible under different building codes, different insurance markets, different fire services, and different property law. Populations with no fire history are receiving a regime their experience does not contain. That the failure appears under institutions this different is evidence that it is not an artifact of any one country's arrangements.

The recommendation follows from the structure and comes in two parts.

The first part is unchanged from the earlier draft. Because the failure lives in perception, and perception is unreliable against this hazard, the durable protection routes around perception and onto the built environment. A hardened home does not require its owner to imagine the firestorm.

The second part is new, and it is where the comparison case pays. If the binding constraint is the cost of believing, then hardening alone is too slow, because hardening is exactly the expensive action people avoid. Two further moves follow. Route the exposure signal into a channel people already attend to, and intervene at the moments when the implied action is cheapest. Section 5.4 develops four levers on that principle, none of them a government mandate and none of them fuels treatment.

FIGURE 1

One rule, two hazards, opposite errors

The same updating rule meets two hazards with opposite properties. Where the hazard is concrete and the implied action is cheap, the public over-responds. Where the hazard is un-picturable and the implied action is expensive, it under-responds.

ONE UPDATING RULE People update on outcome, not exposure		
	HAZARD A Foodborne outbreak	HAZARD B Interface fire
Hazard hidden in an unseen stock	No <i>the case count is the hazard</i>	Yes <i>fuel and interface accumulate unseen</i>
Lethal event differs in kind	No <i>case 2,000 is the same illness</i>	Yes <i>plume-dominated is another regime</i>
Arrives already concrete	Yes <i>in food eaten this week</i>	No <i>must be installed against distance</i>
Cost of acting on it	Near zero <i>skip one product</i>	Very high <i>leave, spend, or uproot a life</i>
Opposite errors, same public, same sound rule.	THE PUBLIC ACTS Product leaves the shelves. Response is fast and broad.	THE PUBLIC DOES NOT ACT Under-response, at catastrophic and irreversible cost.

The public is not inconsistent. It is applying one sound rule to two hazards that differ in exactly the ways that determine which direction the rule fails.

2. Support and Analysis

2.1 The substrate: outcome, not exposure

Start with the rule, because every failure below is that one rule meeting different features of a hazard.

A person forms a belief about their own safety and updates it as the world gives feedback. The feedback available is outcome. Did the fire come. Did the house burn. Did the evacuation turn out to be needed. These are the events a person can actually observe, and someone who updates on them is doing what updating is for.

The defect is not in the person. It is that outcome is a poor stand-in for exposure when the two come apart, and for some hazards they come apart severely.

Exposure is the risk that was run: the chance of loss and the size of it, as the situation actually carried them. Outcome is what happened to land. A person can run enormous exposure and draw a benign outcome many times in a row. Each benign outcome is read as feedback, and the reading is sound. The signal is what is wrong.

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, whatever it turned out to produce (Coil, OLFG Ch. 4). Judging a decision by its outcome is a named error.

The practitioner version is a commander who concludes a risky call was right 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 level off the fireline.

The public is not failing to reason. It is reasoning correctly on a signal the structure of the hazard has corrupted.

What follows specifies how the corruption works, and then why the same rule produces the opposite error against a different hazard.

2.2 The comparison case: a hazard the public acts on

The earlier draft of this analysis rested on fire alone. That was its largest weakness, and it was declared as such. A structure built on one case does not show that the structure means anything beyond that case.

The repair is not a second fire. It is a hazard where the same public makes the opposite error, because that lets the differences be named.

Take the foodborne outbreak as a class. A parasite or bacterium is traced to a packaged produce item. Public health issues a warning. Retailers pull the product. Consumers stop buying it, often well beyond the specific lot, the specific brand, and the specific region named. Coverage is loud and the language is vivid.

Four differences from fire are worth naming, and each one maps to a component developed later in this section.

The hazard is fully visible. Fire's danger lives in a stock a resident cannot see: accumulated fuel, homes pushed into the wildland edge, treatment deferred another budget cycle. An outbreak has no hidden stock. The case count is the hazard. What is observable is what is governing.

There is no change of kind. A severe fire and a firestorm are different physical events, not different sizes of one event. An outbreak is linear. The two-thousandth case is the same illness as the first. Nothing has to be imagined that experience cannot supply.

The threat arrives already concrete. Fire has to be described to a person who has never seen the lethal version. A foodborne parasite needs no description. It is in food eaten this week, in a body the person lives in, traced to a company with a name. It is close in time, close in space, and personal.

Acting on it is nearly free. The protective action is to skip one product for two weeks. That is the whole cost. Fire asks the person to leave a house, spend tens of thousands of dollars, or accept that the life they built in that place carries a risk they cannot buy their way out of.

A fifth difference sits alongside the four and is worth flagging because it does other work. An outbreak has a named party at fault. There is a grower, a processor, a retailer. Fire's causation is diffuse: decades of fuel accumulation, development decisions, deferred treatment, and a climate signal. Blame is available in the first case and not in the second, and available blame sustains attention.

WHAT THE COMPARISON IS AND IS NOT

The comparison does not claim the public over-reacts to outbreaks in a way that is foolish. That verdict would require behavioral measurement this paper does not have, and it would rest on the outcome that happened to land, which is the error this paper names in others.

The claim is narrower and observable: against the outbreak, the public acts. Product leaves shelves. Against the fire, it does not. Those two facts are enough to carry the argument.

Bottom line: the comparison is used to isolate variables, not to grade the public.


One more caution belongs here. Fire and foodborne illness are not equivalent in consequence, and nothing in this section suggests they are. The comparison holds the public constant and varies the hazard. That is its only job.

2.3 Divergence one: the baseline

The first divergence is the drift of perceived risk away from actual risk over time. It runs on probability, and it accumulates across seasons rather than appearing in a moment. The image is the community at the foot of the quiet volcano. The ground is still, the stillness is real, and the stillness 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 behind a flat flow of small annual fires, and the event arrives where the gap is widest.

**WHY A CALIBRATED OBSERVER STILL MISSES IT**

The stock is the hazard. The flow is the visible proxy, and it is a poor one. The quantity that matters was never among the quantities on display.

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 working at three levels: the structure of the environment, the updating of the individual, and the make-up of the population.

The layering is part of the claim. The downward pull on perceived risk is produced independently at each level. A fix that reaches one level leaves the others working. This is the analytical reason that telling people to worry more does not work. A public education campaign reaches, at best, the individual's updating. It does not touch the environment or the population, and those keep pulling.

Base rate — the environmental component

The annual chance of a catastrophic fire at any one address is low. Discounting a rare event is correct. A person who treated every season as a coin flip on losing the house

would be wrong in the other direction. They would exhaust themselves on a risk that, in any given year, almost certainly will not arrive.

The defect is structural, not personal. A correct estimate of a rare event still understates the risk-weighted exposure, because rarity and consequence pull opposite ways and the mind holds the rarity more easily. This is the documented tension in how people weigh rare, high-consequence events: the low probability dominates the judgment, and the severity gets discounted along with the frequency (Slovic 1987; Tversky & Kahneman 1974).

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.

Note what the outbreak case does here. Its base rate is also often misjudged, because seasonal outbreaks get covered as anomalies. But the error runs the other way, toward alarm. Base rate error has no built-in direction. It follows what is vivid, not what is large.

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 a resident can see. What a resident sees across years is a flow: the annual fires, most of them small, most of them controlled. What determines the danger is a stock: accumulated fuel, the density of homes at 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 closely for a decade, and updated faithfully on it, would still misjudge the danger. The quantity that mattered was never on display.

This line has already published this failure at another scale. *The Reveal* documents an organizational insolvency that builds invisibly: unmet obligation accumulating as a stock behind a flow of apparently manageable annual activity, until the stock forces a collapse the flow never signaled (Coil, *The Reveal*). The public version is the same insolvency pointed at a landscape instead of a staff.

The signature is a divergence plot: modeled stock rising while observed flow stays flat or falls. The gap between the two is the hidden hazard, and it is measurable even though it is invisible to the resident living inside it.

This is the component with no counterpart in the outbreak case, and it explains most of the difference in response. There is no invisible stock of parasite accumulating behind years of quiet lettuce. What can be seen is what is governing.

Renormalization — the individual component

Base rate is a fact about the world. Renormalization is what an individual's updating does with a run of that fact coming up empty. The word means the resetting of what counts as normal.

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 real distribution of outcomes. A person who could in principle hold "rare but catastrophic" as a stable belief drifts instead toward "quiet, as it has been."

This is drift, described in the safety literature as the gradual normalization of a widening gap between assumed and actual risk, where 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 measured over time against constant or rising modeled exposure, 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 person updating over time. The fourth describes who is left in the population doing the updating, and it works at a level the others 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, leave. Those who remain are disproportionately the ones for whom the model worked. They stayed, they were not destroyed, and their continued presence is itself evidence to them that staying was sound.

So the community's aggregate belief is biased toward safety twice over: because individuals renormalize, and because the sample producing the aggregate was selected on the favorable outcome.

The Reveal names this at the organizational scale as selection on the favorable tail. A system that promotes or retains people on the basis of having survived a risk comes to be staffed 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 the same. The long-tenure residents whose voices carry the most weight in local risk culture,

the ones who have been there thirty years and seen it all, are exactly the sample most biased by survivorship. Thirty years without the catastrophic event is precisely the record a 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. One is a property of the population's composition.

Remove any single one and the discount still stands on the remaining three. That is why the baseline divergence resists in a way a simple information gap does not. It is not one cause a campaign could address. It is a stack of independent, reinforcing causes spread across three levels, and a message can reach at most one of them.

2.4 Divergence two: the magnitude

The second divergence is a different kind of failure. The baseline is a tracking failure that builds over time. The magnitude failure is a representation failure that appears in a moment. When the event arrives, the person cannot form an adequate picture of what it is.

It runs on consequence, and its central feature is that the lethal event is not merely larger than the person's experience. It is outside it.

The two divergences are independent, and that independence earns them separate treatment. A community can hold an accurate picture of the firestorm, 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 picture the event when it arrives.

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-picturable. That shared cause produces two distinct failures, and keeping the boundary sharp is what keeps the two sections from bleeding together.

The magnitude divergence has an uneven shape. There are two ways to acquire a picture of a magnitude you have not lived: grow it from your own experience, or receive it from outside. The analysis follows those two routes and finds each blocked. The experience route is blocked by two separate mechanisms. The transmission route by one. A fourth factor works across both, making whichever block is active harder to move.

Misclassification — the first experience-route block

A person's sense of how bad a fire can get is anchored on the worst fire they have lived through. That anchor does more than set a ceiling on imagined size. It sets a category.

When a new event arrives, the person files it against the categories available, and the anchor governs which category it lands in. A regime-shifted fire, arriving in a mind whose worst case is a severe structure fire, gets filed as a severe structure fire. A bad one. An extreme one. But an instance of the known kind.

This is a failure of classification, not a fact about the world. And it licenses a specific error. Having filed the event in the familiar category, the person runs the correct scaling for that category. Faced with what is coded as a very bad fire, they reason as one reasons about very bad fires: worse than usual, 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 kind of event.

The signature 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. Compare the description to the actual regime. The gap 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 person a bigger fire and let them scale the anchor upward. Discontinuity is why that fix does not work.

The lethal event is not a larger instance of the familiar one. It is a different physical regime.

Regime here carries the sense it carries in government. A regime change is not more government, or worse government, or the same government turned up. It is a change in the rules by which everything operates. What was permitted becomes forbidden. What worked last week does not work today, and knowing the old rules well does not help.

A fire regime means the same thing. It is the set of relationships governing how the fire behaves, not a description of how bad the fire is.

The fire behavior literature separates two kinds of fire. In the first, ambient wind drives the spread across surface fuels. In the second, the fire's own convective plume dominates its behavior. The second kind generates its own weather. It lofts embers in mass and casts them far beyond the flaming front. Structures ignite from other structures, largely independent of the wildland fuel around them (Rothermel 1991; Werth et al. 2011; Cohen 2000).

The move from the first regime to the second is not a change in degree. It is a change in the rules that determine how the fire moves and what it does to what it touches.

The consequence for the person 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. The worst-plus-worse is a bigger fire of the known kind, and the lethal event is a different kind.

That is why misclassification and discontinuity are two blocks and not one. Misclassification says the person 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 person entirely. The discontinuity between fire regimes is a claim about fire, established in fire science, not a claim about public psychology. That makes "the public cannot extrapolate to the lethal regime" a statement about the world the public is reasoning about, not a statement about a deficiency in the public.

A caution belongs on the word regime, and Section 3 develops it. The discontinuity claimed here must be a real change in the fire's governing dynamics, established by fire behavior. Extremity alone would not close the experience route, because an extreme instance of a known category is still reachable by scaling in principle. The claim 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.

This is the component the outbreak case most clearly lacks. There is no second regime of foodborne illness that behaves under different rules. That absence is why the transmission of an outbreak warning succeeds where a fire warning fails.

Construal distance — the transmission-route block

If the picture cannot be grown from experience, the remaining route is to install it from outside. Tell the person what the regime is, and have the telling take. Construal distance is why the telling does not take.

Construal distance is the finding that a threat held as far away gets represented in vague, general terms rather than concrete ones. A threat held as distant in time, in space, or in likelihood is represented abstractly. This is the central result of construal-level theory: psychological distance and abstraction move together (Trope & Liberman 2010).

A hazard filed as "unlikely, and not here, and not now" is filed as distant, and distance strips out exactly the concrete detail that would make the regime legible. A person can be told, accurately, that the fire will make its own weather and light the house from three sides at once. In the held representation, the telling degrades to a vague "very dangerous," because the threat is filed as distant and distant threats are held abstractly.

This is a transmission-route block specifically. The information is neither unavailable nor 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. The thinned version is not enough to drive the response the regime requires.

The outbreak case is the clean contrast. Nothing has to survive the trip from abstract to concrete, because the hazard was never abstract. It is in a body the person lives in, from a package in their own refrigerator. Construal distance has nothing to work on.

The 2026 European season supplied an unusually clean version of this block. A fire that began near Los Gallardos, in Almería province in southern Spain, on 9 July 2026 killed thirteen people in what officials described as an expat community. Press reporting counted seven Britons among twelve foreign nationals killed. A visitor holds the hazard at maximum distance on every axis and carries no local experience base at all. Both routes to an adequate picture are closed at once.

What followed is the more useful detail. Andalusian emergency officials reported that most of the dead had not followed shelter-in-place instructions. Several died along a dry riverbed they had taken as an escape route. Seven died on foot after leaving their cars.

The instruction was issued and it was correct. Read through this paper's frame, the failure is not compliance. Sheltering in place is the counterintuitive act. It becomes intuitive only with a picture of the regime these residents had no route to acquire. Leaving by car is what a person does when the arriving event is filed as a large fire rather than as another kind of thing.

The cost of believing — the modulator across both routes

The fourth factor is not a fifth block, and the difference is the crux of the whole analysis. The earlier draft called this motivated abstraction. The plainer name is the cost of believing it, and the plainer name is the more accurate one.

Accepting an adequate picture of the lethal regime has a price. The concrete image, the house lost, the road closed, the night of it, obligates an expensive and unpleasant act:

leaving, spending on mitigation, disrupting a life deliberately built at the foot of the volcano.

Because the concrete picture obligates the costly action, a person has an interest in not forming the concrete picture. Keeping the threat vague, 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 works on every other block, and that is what makes it a modulator rather than a coordinate.

It raises the persistence of misclassification: a person 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 person who could extrapolate from a near miss has an interest in declining to. And it raises the persistence of construal distance: the vagueness that distance produces automatically is vagueness the person has an independent reason to keep.

The distinction between this factor and construal distance is worth holding, because they look alike. Construal distance is cold. Distance produces abstraction mechanically, whatever the person wants. The cost of believing is warm. The person has a stake in the abstraction and works, without deciding to, to maintain it.

The tell that they are two mechanisms is an interaction. Hold psychological distance constant and raise the cost of the implied action. Resistance to forming the concrete picture rises with the cost. A purely distance-driven account cannot explain that, because distance has not changed. The motivation responds to the cost of acting, not to the distance of the threat.

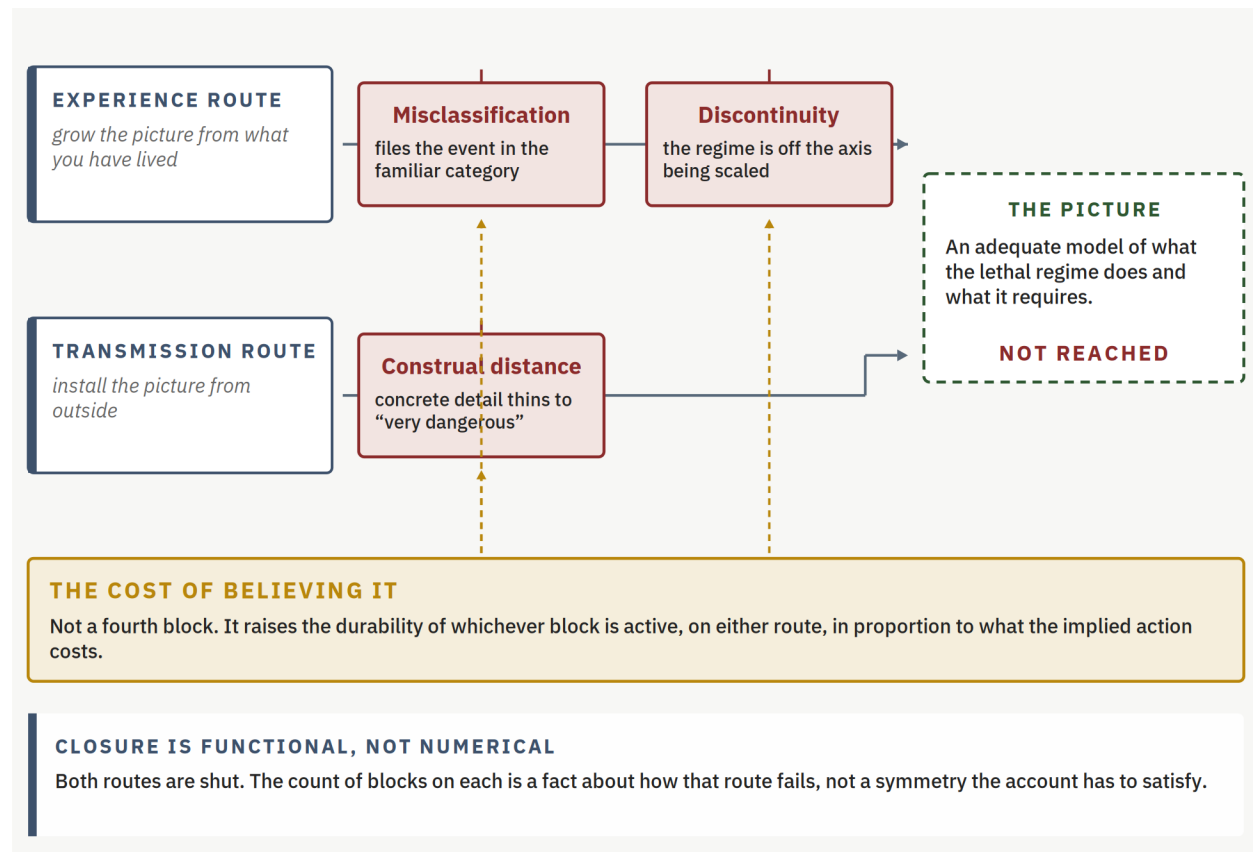
And here the comparison case does its most important work. The outbreak holds distance low and the cost of acting near zero. The public forms the concrete picture readily and acts on it. Fire holds distance high and the cost of acting very high. The public does neither.

That is the pattern the modulator predicts, observed across two hazards rather than argued within one.

FIGURE 3

The magnitude failure is a three-plus-one structure

Two routes feed a person's picture of the regime. The experience route is closed by two independent mechanisms; the transmission route by one. A cross-cutting factor, the cost of accepting the picture, 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 has to satisfy.

Route closure is functional, not numerical

The uneven shape of this structure invites an objection. Two blocks on one route, one on the other, plus a modifier, looks unbalanced and therefore incomplete. The objection mistakes the claim.

The load-bearing claim is that both acquisition routes are closed. The picture 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 on each.

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 enough to close it. The modulator then raises the durability of whichever block is working.

A note on a route that is not a third route. A reader may ask about inference, whether a person could reason their way to the lethal regime without either experiencing it or being shown it. Inference is not a third source of content. It is an operation performed on content supplied by experience or transmission. Its inputs here are physical facts about fire, which reach a person only by those two routes, both already blocked. And its unwelcome conclusion is exactly what the modulator resists.

2.5 What the comparison isolates

Set the two hazards against the components just developed, and the pattern resolves.

Of the four baseline components, only one has a strong counterpart in the outbreak case, and it runs in the opposite direction. Of the three magnitude blocks, none is present. And the modulator, the cost of believing, is near zero.

Two variables carry almost all of the difference.

Native concreteness. Whether the hazard arrives already in a form the person can picture, or has to be installed against distance. The outbreak arrives concrete. The fire regime cannot be installed at all.

The cost of the implied action. Whether believing it obligates something cheap or something expensive. Discard a product, or abandon a house.

This is a finding, and it has a direct consequence for what to do.

If the binding constraint were information, the remedy would be better information. The comparison says otherwise. The public that fails to act on fire is the same public that acts promptly on a hazard it can picture and afford to act on. Nothing about that public needs fixing.

So the useful moves are not moves that raise belief. They are moves that lower the cost of acting, and moves that route the exposure signal into a channel the person already attends to. Section 5.4 builds on exactly that.

THE PREDICTION THIS GENERATES

Public response to a hazard should scale with the cost of the implied action and with how concretely the hazard arrives, more than with modeled exposure.

This is checkable against existing recall compliance data and evacuation compliance data, and it does not require new instrumentation.

Diagnostic: where the protective action is cheap and the public still does not act, this account is the wrong one and something else is doing the work.

2.6 The coupling

The two divergences have been treated separately because they are separate failures. The claim of this section is that against fire 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 wants a signal that breaches the representation problem. Loud, vivid, concrete, and rare enough to register as extraordinary, because a routine signal gets filed in the routine category and scaled as routine.

The baseline divergence wants 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 catastrophe is a data point in the person's outcome updating, and the data point says the warning overstated the risk.

Spend the loud signal frequently and people learn, correctly on the outcomes they observed, to discount it. The learning is sound. The signal taught the discount.

One channel is being asked to do two jobs whose best frequencies are opposed. Breaching the magnitude ceiling wants the signal loud and used sparingly enough to stay extraordinary. Protecting the base rate wants it reserved nearly to silence, because every use that does not cash out erodes it.

This is the analytical content behind the 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. Reserve the signal and it keeps its meaning, but it is now used so rarely, and the standard for warranted is so high, that people get little chance to build the picture they lack. Reserve protects credibility and starves representation. Spend builds representation and destroys credibility.

Note that the outbreak channel does not carry this bind, and the reason is instructive. An outbreak warning usually cashes. The product is recalled, the illness stops, and the warning appears to have worked. The feedback loop is short and it validates. The fire warning that fires without an outcome is the exact data point that erodes it. The coupling is a property of hazards whose warnings usually do not cash.

There is a partial resolution for fire, 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. A distinct, reserved 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.

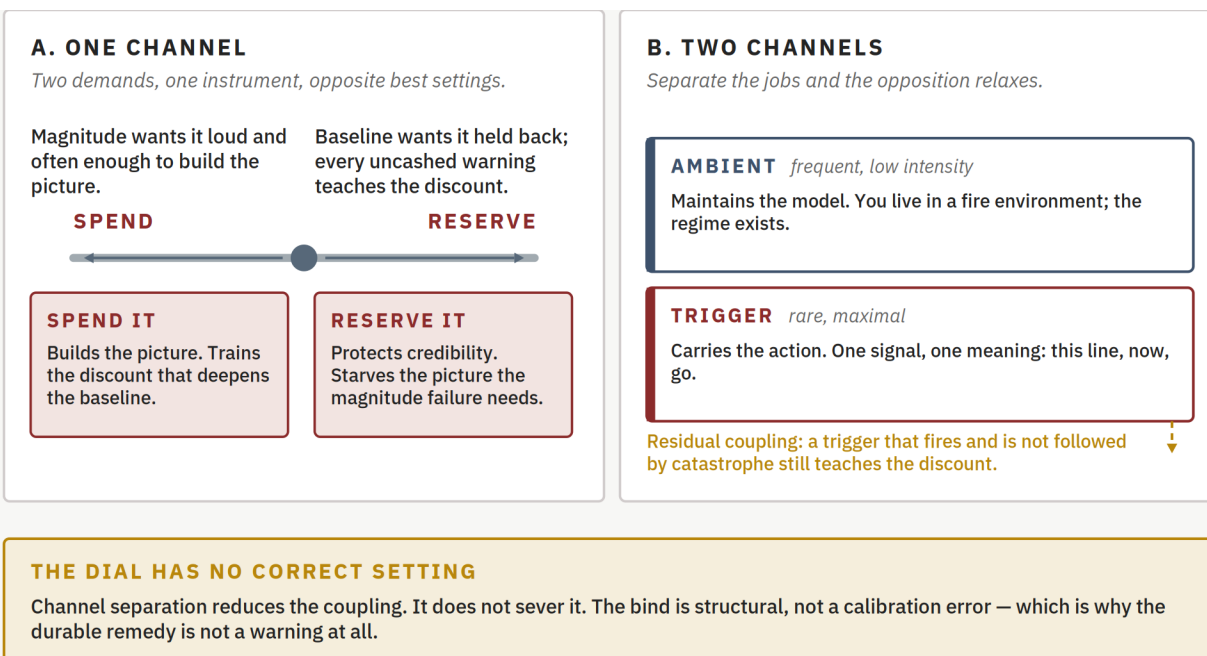
The limit is that fire does not always cash the trigger. The bright line is crossed, the order is given, and then the wind turns or the fire lays down and the neighborhood is spared. One uncashed trigger is a benign outcome following a maximal signal, and the person updating on outcomes registers exactly the data point the reserved channel was designed to avoid.

Channel separation reduces the coupling. It does not sever it. This limit is not a flaw to hide. 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 picture and it trains the discount; reserve it to protect credibility and it starves the picture. Separating the ambient maintenance signal from the reserved trigger relaxes the bind but does not sever it.



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.

2.7 The same failure under different institutions

A structure developed on American fire could be an artifact of American arrangements: American building codes, American insurance markets, American land use, American fire services. The 2026 European season provides a test of that objection, because the institutions differ on every one of those axes.

The European Forest Fire Information System forecast very extreme fire danger over the British Isles, France, and northern Iberia through 22 July 2026. France's most damaging fires that season burned on the Atlantic coast rather than the Mediterranean. The risk map now reaches well beyond southern Europe. These are provisional figures on a season still running.

Two features of that season bear directly on this analysis.

A live discontinuity in a population whose experience does not contain it. The Gironde fire in southwestern France burned intensely enough to generate its own thunderstorm, throwing lightning that started further fires. That behavior has been recorded in North American and Australian megafires. Reported in late July 2026, it had not previously been observed in France. The French Interior Ministry put national burned area near 115,000 hectares as of 26 July, a record for the country, with months of season remaining. Satellite trackers put the figure lower, which is the method divergence noted above.

The risk map extending into places with no fire history. Very extreme danger reaching the British Isles places the baseline divergence at its widest exactly where it is newest. Those populations have the least experience to update on, the least fire-adapted building stock, and the least developed public expectation that any of this applies to them.

The point is not that Europe is burning. The point is that the same failure appears under different codes, different markets, different services, and different property law. That is evidence the structure is not institutional. It is a property of how people update against a hazard with these features.

One caution. This is a second geography, not a second hazard. It strengthens the claim that the fire structure generalizes across institutions. It does not, by itself, establish that the structure describes flood, seismic, or pandemic risk. The comparison case in Section 2.2 does that work, and it does it by dissociation rather than by confirmation.

3. Red Team Assessment

The analysis above is a structural account built largely from constructs imported into the fire domain, anchored on one hazard closely and a second loosely. That construction has specific weaknesses. 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 puts me on one side of the transaction this paper analyzes, and it introduces a specific bias. An account that locates the failure in the structure of the hazard 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 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 elsewhere. An account written to protect the warner would not end there.

A second bias is new to this edition. Section 5.5 argues that some moves cannot wait. Arguments for urgency from someone in the response profession are arguments that tend to route resources toward the response profession. The check I can offer is that none of the four levers in Section 5.4 sends money or authority to fire agencies. They run through insurance, real estate transactions, rebuild permitting, and building products supply. If the urgency argument were self-serving, the levers would look different.

3.2 The two-hazard base

The earlier draft rested on one hazard and said so. This edition rests on one hazard developed closely and a second developed as a class, at low resolution and without case-specific evidence.

That is an improvement, and it is not a general theory. Two hazards do not establish a class. What the comparison establishes is narrower and worth stating precisely: it shows that the same updating rule can produce opposite errors, and it identifies which hazard properties correspond to which direction. That is a dissociation, which is a stronger form of evidence than a second confirming case, but it is still two points.

Whether the four-component baseline and the three-plus-one magnitude structure describe flood, seismic, or pandemic risk is asserted nowhere and demonstrated nowhere here.

3.3 Is the comparison case fair?

The sharpest objection to this edition is that the two cases are not measured the same way.

The fire case rests on a behavioral failure that can be observed: people warned, people did not evacuate, people died. The outbreak case, as usually presented, rests on how coverage was framed rather than on what the public actually did. Media volume is not public response, and treating the two as equivalent evidence would be an error.

There is a second problem with the usual framing. Calling the outbreak response disproportionate typically rests on the outcome that landed: few deaths, moderate hospitalization. That is judging a response by its outcome rather than by the exposure it faced, which is the error this paper names in Section 2.1. An analysis cannot name that error in the public and then commit it in its own comparison.

The defense is to narrow the claim to what is observable, and Section 2.2 does that. The paper does not need the public to have over-reacted. It needs only that the public acted: product left shelves, purchases fell, the protective action was taken. That is documented in recall compliance data and does not depend on any verdict about proportionality.

The objection stands as a drafting discipline. Any language implying that the public over-reacts to outbreaks should be struck, because the paper does not need it and cannot support it at this evidence level.

3.4 Does misclassification depend on discontinuity?

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

The defense is dissociation, and it requires showing each occurring without the other.

Misclassification without discontinuity: a very large fire, still wind-driven, no plume domination, no regime change, gets filed by a resident as an ordinary fire and scaled as one. The resident under-responds to a large but continuous event. The failure is real and the discontinuity is absent.

Discontinuity without misclassification: a resident correctly senses that the arriving event is not normal, that this is not like anything they have seen, and still has no picture of how the regime behaves or what it will do. They have not mis-filed the event. They have correctly declined to file it in the familiar category, and are left with no 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.

3.5 Does construal distance also touch experience?

Construal distance is placed on the transmission route. But psychological distance plausibly shapes how a person 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 the mechanism operates on both routes, its placement as transmission-specific is unstable.

The defense narrows the mechanism rather than denying the overlap. On the transmission route, construal distance blocks the installation of supplied content: the concrete message thins in a mind holding the threat as distant. That is a different operation from whatever distance does to the memory of a lived event, which modulates experience-route content rather than blocking transmission-route installation.

The cleanest handling is to define the transmission block narrowly as the thinning of supplied content, and to treat any distance effect on near-miss memory as secondary. The objection is a warning against defining construal distance so broadly that it floats free of the route it is assigned to.

3.6 Does the cost of believing explain too much?

A factor that raises the persistence of every other block, on both routes, and now also carries the comparison between two hazards, is a factor at risk of explaining everything and discriminating nothing. If it is invoked wherever a block persists, it becomes unfalsifiable: a name for persistence rather than an account of it.

This risk is higher in this edition than in the last, because the factor has been promoted to the title.

The discipline is to specify where it should and should not be expected, and to make the specification testable. The factor should be strong where the protective action is expensive, where the home is the primary asset, where a life is deeply built into a place. It should be weak or absent where the protective action is cheap.

That predicts a checkable pattern: resistance to forming the concrete picture should scale with the personal cost of acting on it, holding distance and information constant. Where the protective action is cheap and the person still resists the concrete picture, this factor is the wrong explanation and something else is doing the work.

The outbreak comparison is the first test of that boundary and it passes, because the cheap-action case shows the predicted absence of resistance. One passing test is not many.

3.7 The two-route claim

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

The defense rests on a distinction that has to stay explicit: the difference between a source of content and an operation on content. Experience and transmission are sources. They introduce content the person 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 blurs source and operation, the claim is exposed.

3.8 Is *regime* too strong a word?

The discontinuity block leans on the claim that the lethal event is genuinely different in kind, not merely a surprising or extreme instance. A reader may argue that regime overstates what is really a matter of degree, or that the discontinuity is a feature of our description rather than of the world.

The defense is to define the term at the right strength and 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: 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 literature supports that reading. Whether a fire is wind-driven or plume-dominated is set by one balance: the power of the fire's own convection against the power of the ambient wind. The two states run under different governing relationships. They are not two points on one scale of degree (Rothermel 1991; Werth et al. 2011).

Were the fire science to support only a difference of degree, the discontinuity block would collapse toward misclassification, and the experience route would close on one mechanism rather than two.

3.9 The evidence base

Most of the structure above is imported theory applied to fire, plus 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 rather than designed studies.

The paper's claims should be calibrated to that standard. The imported constructs are established in their home literatures. The fire application is argued rather than measured. The signatures named for each component are proposals for how one would test the claim, not reports of tests already run.

European figures carry an additional caution. They are season-to-date and provisional, national and satellite estimates differ in method, and totals will move. They are used here to show that the failure appears under different institutions, which is a qualitative claim and does not turn on any particular number.

3.10 The levers in Section 5.4 are not costless

The forward-looking section is the most exposed part of this paper, because it moves from anatomy to prescription and because each lever transfers a cost onto someone. Naming who is the minimum discipline.

Pricing exposure at the parcel produces uninsurability. A premium that tells the truth about exposure will, in some places, tell people their exposure is not insurable at any price they can pay. That is information working correctly and a household in crisis at the same time. The cost lands hardest on people with the least ability to move or retrofit.

The price signal does not finance the fix. Liao and Walls put a combined retrofit near forty thousand dollars, which annualizes to roughly ten times the average mitigation discount they measured. Their conclusion is that retrofit costs run orders of magnitude above the insurance savings. A signal that identifies the problem without funding the remedy converts a perception failure into a liquidity failure. That is progress only if the financing follows (Liao & Walls 2025).

Part of the gap is not correctable. The same study attributes small discounts partly to risk externalities: structure-to-structure spread and community fuel loads weaken the link between what a household spends and what its insurer expects to pay. Section 5.4 treats this as a ceiling on the pricing lever rather than a flaw in its execution.

Transaction disclosure suppresses value. Telling buyers the truth at the point of sale reduces prices in exposed communities. The people holding those assets are not the people who created the exposure.

Defaults raise the cost floor. Making hardened assemblies standard raises the cost of building anything, in a country with a housing shortage. That trade is real and it should be argued on its merits, not assumed away.

None of these costs argues against the levers. They argue against presenting the levers as free. A recommendation that hides its distributional effects is a recommendation that has not been thought through.

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 this 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, set against a risk they judge, correctly in most years, to be low. They decided the risk was 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 elaborate machinery above, on this view, is an overbuilt explanation for something simple. People know the risk and accept it, and calling that acceptance a misperception is an analyst substituting his risk tolerance for theirs.

This dissent is strong for a specific reason: this paper has already half-conceded it, and the new title concedes more.

The cost of believing is, on this paper's own account, a rational response to the cost of the protective action. The person keeps the threat vague because concretizing it obligates an expensive act. The dissent presses exactly there. If 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 this 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. There is only a distribution of risk preferences the analyst happens to disagree with.

The dissent fails as the whole account, and the reason is precise.

A preference is only a preference if it is formed over an accurate picture 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. This paper has no quarrel with them.

But the magnitude divergence is the claim that most residents cannot form an accurate picture of the firestorm regime. A choice made over a picture 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 picture that Section 2.4 shows the person does not have. Where the picture is accurate, the dissent is right and the behavior is preference. Where the picture is the under-scaled one, the dissent is describing a preference over the wrong object.

The comparison case sharpens this rather than softening it. Against the outbreak, the picture is accurate and the action is cheap, and people act. That is preference and cognition agreeing. Against fire, the picture is inaccurate and the action is expensive, and people do not act. The dissent cannot tell those two situations apart, because it has only one variable. This paper has two, and the two-hazard comparison is what shows both are working.

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. So the goal of a remedy cannot be to force the protective action against a preference. It can only be to make sure 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 the built environment and toward lowering the cost of acting rather than toward persuasion. Both respect the autonomy the dissent correctly defends. Both cover the case the dissent cannot: the resident whose preference is an artifact of an under-scaled picture.

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. It holds that the failure is one of message quality, that well-designed and repeated risk communication does move behavior, and that this paper overstates the closure to make the problem look harder than it is.

That account is not wrong that message quality matters, and this paper does not claim warnings are useless. It claims they are insufficient against this hazard structure. Section 2.6 gives 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 picture. The craft dissent improves the message. It does not escape the channel bind, which is a property of the structure rather than of the message.

5. Recommended Course of Action

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. The concession they build toward is that the reliable remedy is not a communication remedy at all.

Sections 5.1 through 5.3 are unchanged in substance from the earlier draft. Section 5.4 is new, and it exists because the comparison case identified a variable the earlier draft treated as background.

5.1 Pre-commitment decided cold

The magnitude divergence says a person cannot summon an adequate picture in the window where the decision has to be made. The first move takes the decision out of that window.

This line already carries the instrument, pointed at practitioners. Break a big judgment into its load-bearing parts, pair each part with something observable, and pre-decide the response to each before the moment arrives (Coil, *Order It and It Comes*).

Turned outward, the same discipline breaks "am I safe" into observable, pre-specified conditions. A named fire on a named ridge. A specific evacuation level. A particular wind. Each gets a pre-decided action. The deciding happens on a cold day, when the picture is not needed and imagination is not required.

A person 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. It works precisely because it does not rely on real-time judgment, which this paper has shown is not available.

The limit is that pre-commitment has to be made when the threat is distant, and 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 real, and it is bounded.

5.2 Borrowed outcomes

If a person's own outcome library caps below the lethal regime, the only way to add an outcome at the right size is to borrow one. A distant catastrophe supplies an outcome to a

library that has none, and secondhand experience of a regime is the nearest available substitute for experience the person cannot and should not acquire directly.

The move is real and it is perishable. A borrowed outcome fades as it recedes in time. It is subject to the same discount as any distant event: filed as there and not here, held abstractly, and weakened by the person's interest in keeping it that way.

Borrowed outcomes lift the ceiling temporarily. They do not install a durable picture, 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 cognition, and this paper has spent its length showing that cognition, against this hazard, is unreliable in specific and 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 whether or not its occupant perceives the risk correctly, forms an accurate picture, or prefers to act. A hardened home does not require its owner to imagine the firestorm. It survives the ember cast whether the owner left early with a clear picture or fled at the last minute with no picture at all.

This moves the burden of protection off perception and onto the built environment, which is durable, inspectable, and indifferent to the occupant's mental model.

It follows from the analysis rather than being appended to it. If the failure is structural in perception, then a remedy operating through perception inherits the structural failure, and a remedy operating on physical exposure does not.

The recommendation is therefore not warn better. It is that warning is real, necessary, and bounded, 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.

5.4 After hardening: four levers that are not mandates

Section 5.3 is correct and it is not sufficient, and the comparison case is what shows why.

Hardening is itself the expensive action. Telling a resident that the reliable protection is a hardened home is telling them that the reliable protection costs tens of thousands of dollars they have an active interest in not believing they need. The remedy runs straight into the modulator.

Two conventional answers exist, and both have known limits. Government mandates work at the point of new construction and reach the existing stock slowly and unevenly. Fuels treatment reduces the stock and is constrained by budget, capacity, season length, and scale. Neither is wrong. Neither is fast enough on its own.

The finding in Section 2.5 points somewhere else. If the binding constraints are how concretely the hazard arrives and what it costs to act, then the useful moves do two things. They route the exposure signal into a channel the person already attends to. And they intervene at the moments when the implied action is cheapest.

Four levers follow. None is a mandate and none is fuels work.

Lever one: price the parcel, not the territory

Insurance is the only institution that prices exposure rather than outcome. That is the exact quantity the public cannot perceive, and premium is a channel every homeowner attends to annually because they pay it.

The instrument currently underperforms. Where carriers price wildfire risk at the territory level, a mitigated home and an unmaintained one on the same block can carry the same premium. A premium that cannot tell the two apart is noise, not signal.

Parcel-level underwriting changes that. Where carriers assess roofing material, vent type, vegetation clearance, and fuel proximity property by property, premium becomes a legible statement about a specific address. Two similar buildings on one block land in different pricing tiers, and the difference is legible to the owner.

That converts an invisible stock into an annual number a household cannot ignore. The mechanism does not depend on the household forming any picture of the fire.

Two corrections to that argument are required, and the second is serious.

First, the leading example is not a market movement. California has required insurers who price wildfire risk to credit documented mitigation since 2022, and the state's insurer of last resort now runs a twelve-item hardening discount schedule. The instrument is being driven by regulation. Calling it mandate-free would be wrong.

Second, the discounts are small for a reason that is not a modelling deficiency. Liao and Walls attribute the gap to genuine uncertainty about how much any one household's mitigation reduces its own loss, and to risk externalities: fire spreads structure to structure, and community-level fuel governs much of the outcome. A hardened house on an unhardened block is still exposed to its neighbours (Liao & Walls 2025).

That is a principled ceiling on this lever, not a temporary one. An insurer that priced the parcel accurately would still find that the parcel is not where much of the risk lives. The signal can be sharpened. It cannot be made to carry what the block determines.

The implication is that pricing works best where it is applied at the scale the hazard actually operates on. Community-level certification, where a whole subdivision qualifies together, matches the instrument to the externality. It also converts a private decision into a shared one, which is a different and harder problem than the one this paper set out to solve.

Section 3.10 states the costs, and they are serious. Accurate pricing produces uninsurability for some households, and the price signal identifies the problem without financing the fix. This lever is incomplete without a financing mechanism attached, and recommending it without one would be recommending half a solution.

Lever two: intervene at the transaction

The sale of a house is the single point at which the cost of acting on fire risk is lowest. The buyer has not bought yet. They can walk, negotiate the retrofit into the price, or choose differently, and none of those options requires abandoning a life already built in a place.

Every mechanism in Section 2.4 is weakest at that moment. The buyer has no local anchor to misclassify against. They have not renormalized across quiet seasons in that place. They are not yet part of the surviving population. And the cost of believing is near zero. Believing it changes a purchase decision rather than uprooting a household.

This is the closest analogue in fire to what makes the outbreak case work: a concrete fact delivered at a moment when acting on it is cheap.

It is currently among the least used levers available. Exposure information reaches buyers unevenly, late, and in formats that thin toward abstraction exactly as Section 2.4 predicts.

Lever three: use the rebuild window

After a fire, a community holds every condition this paper says is normally missing. The borrowed outcome is not borrowed, it is local. The picture is concrete. The regime has been observed rather than described.

And the cost of acting is at its structural minimum, because the house is being built anyway. The marginal cost of an ignition-resistant assembly on a structure already under construction is a fraction of the cost of retrofitting a standing home.

That window is currently the largest wasted opportunity in the system. It is also short, and it closes under pressure to rebuild quickly.

The move is to have the financing, the material supply, the design guidance, and the permitting pathway staged in advance, so that the hardened rebuild is the fast path rather than the slow one. Speed is what a rebuilding community optimizes for. A remedy that makes hardening the slower option will lose to that pressure every time.

Lever four: make the hardened option the default

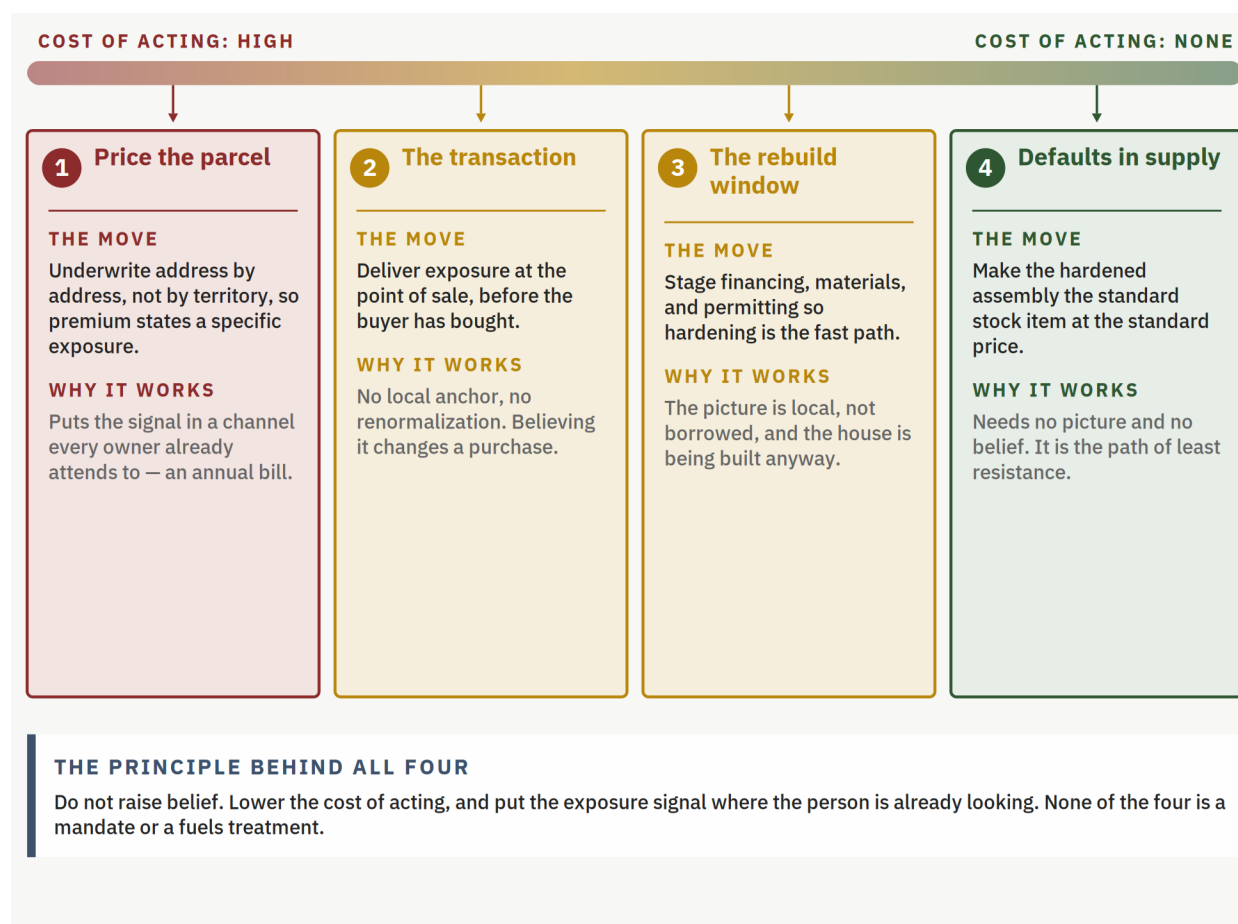
The purest form of the finding is to drive the cost of acting toward zero, at which point belief becomes unnecessary.

If the ignition-resistant vent, the ember-resistant roof assembly, and the enclosed eave are the standard stock item at the standard price, then choosing them requires no decision, no picture of the firestorm, and no belief about exposure. The protective action happens because it is the path of least resistance rather than because anyone was persuaded.

This is supply-side rather than regulatory. It runs through manufacturers, distributors, builders, and the assemblies that get specified by default in standard plans. **FIGURE 5**

Four levers, ordered by the cost of acting

Each lever either routes the exposure signal into a channel the person already attends to, or intervenes at a moment when the implied action is cheap. Cost of acting falls from left to right; none of the four is a mandate or a fuels treatment.



The levers do not try to make people believe more. They try to make believing cheaper, and in the last case unnecessary.

Defaults do not require perception. That is the whole argument of this paper, applied to procurement rather than to warning.

The cost is a higher floor on construction, in a country with a housing shortage. Section 3.10 names it. It is a real trade and it should be argued, not assumed.

THE PRINCIPLE BEHIND ALL FOUR

Do not raise belief. Lower the cost of acting, and put the exposure signal where the person is already looking.

That principle is not derived from fire. It is derived from the hazard where the public does act, and it is the reason the comparison case earns its place in this paper.

What to do: apply the test to any proposed remedy. If it asks the public to believe something expensive, it inherits the failure this paper describes.

5.5 Why the timing matters

This section makes a claim about rate, not a claim about urgency in the rhetorical sense. It should be read as evidence and held to that standard.

Three quantities are moving, and they are moving at different speeds.

The stock is still rising. Fuel accumulation and interface density are the hazard, and neither has turned. Treatment reduces the stock at a rate constrained by budget, workforce, burn windows, and season length. Development adds to it at a rate constrained by housing demand. The gap in Figure 2 widens for as long as the second exceeds the first.

The regime is arriving in places whose experience does not contain it. The fire-generated thunderstorm recorded in the Gironde in late July 2026, reported as unprecedented for that country, is one instance. Very extreme fire danger forecast over the British Isles in the same month is another. Populations at the leading edge of that expansion carry three disadvantages at once. The least experience. The least adapted building stock. The least institutional expectation that any of this applies to them. The baseline divergence is widest where the hazard is newest.

The price of exposure is repricing faster than perception is updating. California regulators approved a 29.1 percent average increase for the state's insurer of last resort, effective 15 October 2026. That plan's policy count grew 44 percent between autumn 2024 and the end of 2025, to roughly 668,600. One qualification belongs with it: quarterly growth slowed sharply in early 2026, so the trend is a large book still growing rather than a runaway one. This is the exposure signal arriving through the one channel that cannot be filed as distant, and it arrives as a bill.

The operational consequence is a sequencing claim, and it is the only directive-shaped statement in this paper.

Lever three, the rebuild window, is time-limited by definition, and it opens and closes on a schedule set by fires that have already happened. Capacity to use it has to exist before the window opens, not after. A community that begins assembling financing, material supply, and permitting pathways after a fire has already lost the window.

Lever four, defaults in the supply chain, has the longest lead time of the four, because it moves through manufacturing and distribution rather than through decisions. It produces nothing for years and then produces everything. Started late, it delivers after the stock has grown further.

Both of those are slow instruments applied against a fast quantity. That is the argument for starting them now, and it is an argument about lead times rather than an appeal to alarm.

5.6 What would show this wrong

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

The account is wrong if fire-specific measurement 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 pictures do not gap systematically below the actual regime. If resistance to forming the concrete picture 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.

It is wrong about its central new claim if public response to hazards does not track the cost of the implied action and the concreteness of the hazard. That is testable against recall compliance and evacuation compliance data, and it is the single most checkable claim in this paper.

It is wrong about its remedy if hardened structures do not outperform occupant perception as a predictor of survival.

The new levers are wrong on three separate tests. If parcel-level pricing does not change mitigation behavior where it has been introduced. If disclosure at the transaction does not change buyer behavior. If default specification does not move installed building stock faster than mandates do. Each is measurable where the instruments already exist.

Each of these would falsify a specific part of the structure rather than the whole, which is the property a decomposition is supposed to have.

5.7 Scope

This paper analyzes one clock. It treats the chronic problem, the standing misperception of a rare, latent, regime-shifting hazard by the public exposed to it. 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. The account here sits beneath them, explaining why the failure they manage persists even when they are performed well.

The claim is bounded to what has been taken apart. One updating rule fails in two structurally distinct ways against fire. The two failures are coupled through the one instrument meant to correct them. The same rule fails in the opposite direction against a hazard with the opposite properties. What the two hazards share is the rule. What separates them is how concretely the hazard arrives, and what it costs to believe 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.
- Coil, J. (2026). *Order It and It Comes*. Cogplexiti Systems & Decision Analysis. Flagstaff, AZ: Cogplexiti.
- Coil, J. (2026). *The Reveal*. Cogplexiti Systems & Decision Analysis. Flagstaff, AZ: Cogplexiti.
- Dekker, S. (2011). *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*. Farnham, UK: Ashgate.
- European Forest Fire Information System (2026). Current wildfire situation in Europe: fire danger forecast. Copernicus Emergency Management Service, European Commission Joint Research Centre. Bulletin data to 22 July 2026. [Season-to-date; provisional.]
- Liao, Y., & Walls, M. A. (2025). *From Risk to Reward: Insurance Discounts for Wildfire Mitigation (Working Paper 25-30)*. Washington, DC: Resources for the Future.
- Ministère de l'Intérieur, France (2026). National burned-area statement, 26 July 2026, as reported by France 24 and Agence France-Presse. [Ministry figure; satellite estimates differ.]
- 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.

Production note. A citation-verification pass was run on 31 July 2026 against the European fire-season and insurance-market claims in Sections 2.4, 2.7, 3.10, 5.4, and 5.5. European figures remain season-to-date and provisional; national ministry and satellite estimates differ in method and are reported separately rather than reconciled. Press-sourced claims should be re-checked against a primary source before publication. Figures 1–5 are placed at 300 DPI, authored in point-space to the Gen 3 wrapper, with a verified 7pt type floor.

*The quiet is real, and people read it as evidence.
The evidence points the wrong way.*

Two hazards reached the public in the same season. A parasite in packaged produce, and fire. People emptied the shelves of the first and did not move for the second. Same public, same week, opposite errors — and the error that cost almost nothing was the one they acted on.

**The binding constraint is not what people know.
It is what it costs them to believe it.**

This paper takes the failure apart. People update on outcome, not exposure — the error the operational literature names in commanders, moved one level down to the public. Against fire that single rule produces two structurally different failures: a baseline drift that builds across quiet years, and a representation failure that arrives in a moment, when the event turns out to belong to a regime experience cannot supply. The two are coupled through the one instrument meant to correct them, and the coupling has no correct setting. The recommendation follows the structure rather than the sentiment: route protection off perception and onto the built environment, then lower the cost of acting through four levers that are neither government mandates nor fuels treatment.

Part of the Cogplexiti Systems and Decision Analysis line, which carries its own red team and a credible dissenting alternative in every paper, and states in advance what evidence would show it wrong. It supersedes the working draft circulated as *The Foot of the Volcano*; that draft analyzed one hazard, and the second one changed the finding.

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.

