EVIDENCE

Detection-to-Decision: Benchmarking the Critical Window

Between the moment a fire becomes detectable and the moment a crew is authorized to act on it lies a window that is mostly manual, largely unmeasured, and — the evidence suggests — where the leverage actually is. This analysis maps that window, confronts the data that complicates the easy story, and estimates what is at stake.

The window, defined

Every wildfire response runs through the same chain. A fire ignites; at some point it becomes detectable; it is discovered and reported; its location and nature are confirmed; a decision is made to commit resources; those resources mobilize, travel, arrive, size up the fire, and act. The total elapsed time from ignition to effective suppression determines, along with fuel and weather, whether the fire is caught small or escapes.

This analysis concerns one segment of that chain in particular — the interval from a detectable signal to an authorized, coordinated response. Call it the detection-to-decision window. It is the part of the chain that is least about physics and most about workflow: confirmation, correlation, the decision to act, and the dispatch that follows. Three questions organize what follows. Where does the time in this window actually go? Does compressing it change outcomes? And what, in dollars, is at stake?

The intuitive story, and the evidence that complicates it

The intuitive story is simple and nearly universal in wildfire-technology marketing: detect the fire sooner, act sooner, keep it small, spend less. It has the ring of obvious truth. It is also, taken at face value, not well supported by the best econometric evidence — and a serious analysis has to say so before it says anything else.

In 2024, an analysis in *PLoS One* applied double/debiased machine learning to roughly 4,900 western-Canadian wildfires that had been successfully suppressed before reaching a critical size, isolating the effect of reporting delay on suppression cost. The finding was statistically significant but modest: reporting delay accounted for only about three percent of total suppression cost, and reducing it by a full hour reduced cost by roughly a quarter of one percent. Read plainly, that result says detection speed, by itself, is a weak lever on cost. Anyone selling minutes as money should have to reckon with it.

Why the average hides the answer

The reconciliation lies in what that estimate does and does not measure. By construction, the *PLoS One* sample consisted of fires that were already contained before they grew large. It measured the cost effect of detection timing on the fires that stayed small — and found it modest, which is exactly what one should expect, because those fires were never the expensive ones. It did not, and structurally could not, measure the effect that matters most: whether faster, better-coordinated action prevents a fire from escaping in the first place.

That distinction is the whole game, because wildfire cost is not distributed evenly across fires. It is concentrated in the few that escape. The pattern is consistent across fire-management systems: the large majority of ignitions — commonly cited at well over ninety percent, and around ninety-seven percent in a recent study of nearly 27,000 California ignitions — are contained by initial attack and stay small. The small remainder that escape account for the overwhelming majority of the damage. A study of roughly 1,200 boreal fires found that response failures, about ten percent of cases, were responsible for eighty-five percent of the area burned; broader reviews conclude that the fires which grow large account for nearly all area burned and nearly all suppression cost. Critically, the same boreal analysis found that the probability of a response failure rose with longer response times, and a 2025 California model treats response time, in minutes, as a primary predictor of initial-attack success.

So the timing effect is not absent. It is concentrated. It is negligible for the many fires that would have stayed small regardless, and decisive for the few, near the threshold of escape, that determine the entire cost distribution. Averaging a concentrated effect across the whole population dilutes it to “modest” — which is precisely what the econometrics found. The right object of analysis is not the average cost saved per minute. It is the change in the probability that a given fire is prevented from escaping.

It helps to name three categories of fire, as fire managers do. Some would stay small no matter what anyone did; call them Category A. Some will grow large no matter what anyone does, because conditions overwhelm any initial attack; call them Category C. And some — Category B — stay small only if effective suppression reaches them in time, and escape if it does not. Every dollar of value from a faster or better-coordinated response lives in Category B. It is the only category in which the decision window changes the outcome, and it sits exactly at the escape threshold where the cost distribution is decided.

Where the minutes go — and the part no one measures

If the leverage is in compressing the path from detection to effective action for Category B fires, the operational question is where that path actually spends its time. Walking the chain segment by segment reveals an uncomfortable asymmetry.

The front of the chain — detection itself — has been heavily optimized and is increasingly well measured. Geostationary satellites flag thermal anomalies in near-real time; polar-orbiting sensors add resolution on a slower revisit; and AI-analyzed camera networks, such as the more than eleven hundred cameras California has watched with machine vision since 2023, have compressed discovery time for fires within line of sight. Detection latency is shrinking, and the *PLoS One* result tells us its marginal cost effect, for fires that stay small, is modest. By any reading, detection is no longer the binding constraint — the conclusion reached independently in *The Integration Gap*.

The back of the chain — mobilization, travel, and arrival — is partly physics and partly positioning, and it, too, is increasingly measured. A 2024 study reconstructed aerial response times across the 2020 California season by matching more than seven thousand wildfire-discovery records to aircraft positions broadcast over ADS-B; Canadian work has decomposed helitack getaway and travel times into their contributing factors. Travel time can be shortened by where assets are pre-positioned, but some of it is irreducible distance.

Between these two reasonably measured ends sits the segment that is neither optimized nor measured: confirmation, correlation, the decision to commit, and the dispatch that follows. Here the published record thins out, and the reason it thins out is itself the finding. The Government Accountability Office reports that the precise location of a suspected fire often still has to be determined by trained personnel, and that data from the patchwork of public and private detection systems “may not be easily integrated and compatible for analyses and quick response.” This is the step in which a credible detection becomes an authorized, coordinated action — and it is performed, today, by people correlating across radios, phones, and incompatible consoles, agency by agency, with no shared operating picture. It is the least instrumented segment of the entire chain. We have reasonable data on how fast fires are detected and on how long aircraft take to arrive, but almost no standardized data on how long the manual decision step in the middle actually takes — because, as the GAO keeps finding, the data required to see it is not integrated. You cannot benchmark what you do not instrument, and the most improvable segment of the response chain is the one no one is measuring.

The cost-and-loss translation

What is at stake can be bounded with authoritative figures. A NIST literature survey estimates the annualized economic burden of wildfire in the United States at between $71.1 billion and $347.8 billion in 2016 dollars. Against that, federal suppression spending runs a few billion dollars a year — on the order of the $3.2 billion average the GAO cites. The Department of the Interior’s own policy analysis draws the obvious conclusion: total wildfire costs are vastly larger than what is spent on suppression, and additional investment in response could be a cost-effective way to reduce the losses. The burden is dominated not by the cost of fighting fires but by the losses fires cause — structures, infrastructure, health effects from smoke, watershed and ecological damage, and the long-tail costs of recovery — and those losses, like area burned, are concentrated in the fires that escape.

That concentration is what makes the expected-value arithmetic counterintuitive. Consider — only to illustrate the structure of the argument, not to estimate a return — a protected region with a few hundred ignitions a year, of which the great majority are contained small and a few percent escape. If the losses from an escaped fire run two to three orders of magnitude above those from a contained one, consistent with the cost concentration the literature documents, then reducing the escape rate by even a few tenths of a percentage point avoids more expected loss than eliminating detection latency entirely on the fires that were never going to escape. The arithmetic is dominated by the tail. A program optimizing for average detection speed is optimizing the wrong moment; the moment that pays is the one that tips a Category B fire from escape to containment.

For an insurer, the same logic appears on the balance sheet. Insured exposure is overwhelmingly structure and property loss, concentrated in wildland-urban interface fires that escape initial attack. The number that moves a loss ratio is the reduction in expected insured loss, and that tracks escape prevention — the Category B tip — far more than it tracks how quickly a fire was first seen.

The measurement problem is the opportunity

Put the pieces together and the conclusion is uncomfortable for the conventional emphasis of wildfire technology. The field has poured effort into the front of the response chain, where progress is visible and measurable, and where the best evidence says the marginal payoff is modest. It has underinvested in the manual decision step in the middle, where the cost-and-loss evidence says the leverage actually is — and it has underinvested there in part because that step is invisible: unmeasured, unbenchmarked, and built on data that does not connect.

The first thing the evidence calls for, then, is not a faster sensor. It is instrumentation of the decision. A system that records every step from detection to authorization — every correlation, every recommendation, every human approval or refusal, every dispatch, logged against the policy in effect and reconstructable end to end — turns the detection-to-decision window from a black box into a measurable, improvable quantity. Only then can the field state, rather than assume, how much of the critical window is spent in manual coordination, and how much of that can be compressed without removing the human from any consequential decision. The targets such a system should be held to are concrete and measurable on its own records: a large reduction in the raw-alert volume operators must process, a median time-to-triage measured in seconds rather than minutes, and a high rate of operator agreement with what the system proposes. Those are hypotheses to be validated in real operations, not results to be claimed — but they are, for the first time, hypotheses that can be measured.

The question worth funding is therefore not the one the market has been answering. It is not “can we detect fires faster,” which is largely settled and modest in payoff. It is “can we measure and compress the manual decision window that turns a detection into coordinated action” — the segment the cost-and-loss evidence identifies as the lever, and the one no one has yet instrumented well enough to benchmark. That is a question you can only answer by building the instrument.

Sources

U.S. Government Accountability Office, *Science & Tech Spotlight: Wildfire Detection Technologies* (GAO-25-108161), May 2025.

U.S. Government Accountability Office, *Wildfire Management: Technologies for Forecasting, Detection, Mitigation, and Response* (GAO-25-108589), June 2025.

D. Thomas, D. Butry, S. Gilbert, D. Webb & J. Fung, *The Costs and Losses of Wildfires: A Literature Survey*, NIST Special Publication 1215, 2017.

U.S. Department of the Interior, Office of Policy Analysis, *Review of the Economic Literature on Wildland Fire Management* (R-2023-001), 2023.

M. S. Huang & B. Wichmann, "Machine learning estimates on the impacts of detection times on wildfire suppression costs," *PLoS One* 19(11), 2024.

"Empirical models of forest fire initial attack success probabilities: the effects of fuels, anthropogenic linear features, fire weather, and management," *Canadian Journal of Forest Research*, 2006.

A. Cardil et al., "Assessing the suppression difficulty of wildland fires for initial attack response," *International Journal of Wildland Fire* 34, 2025.

"Measuring Initial Attack Suppression Effectiveness through Burn Probability," *Fire* 2(4), 2019.

"An assessment of aerial firefighting response times between agencies during the 2020 fire season in California," *International Journal of Disaster Risk Reduction*, 2024.

"Factors influencing travel time and getaway time of helitack crews," *Canadian Journal of Forest Research* (Quebec helitack study).

"Wildfire Prediction to Inform Fire Management: Statistical Science Challenges," *Statistical Science* (review).

Figures are drawn from the published sources above; the expected-value passage is an illustrative structural argument, not an estimate of return, and is labeled as such.

About Inari Watch

Inari Watch is a veteran-founded company building an integrated system of systems for regional wildfire defense. We unify the sensors, aerial assets, and data feeds a region already depends on into one coordinated capability — connected by AI at the edge and bounded agentic coordination, with humans in command of every consequential decision.