

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

SIX EXPLANATIONS, ONE GAP

An Analysis of Competing
Hypotheses on Why Prescribed
Fire Lags Its Targets

JAYSON COIL

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

WHITE PAPER

Six Explanations, One Gap

An Analysis of Competing Hypotheses on Why Prescribed Fire Lags Its Targets

JAYSON COIL

Copyright © 2026 by Jayson Coil. All rights reserved.

Published by Cogplexiti LLC · Flagstaff, Arizona · cogplexiti.com

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.

Cogplexiti · Systems & Decision Analysis

ISBN: 979-8-9959056-9-1

First edition · July 2026

Suggested citation: Coil, Jayson. 2026. *Six Explanations, One Gap: An Analysis of Competing Hypotheses on Why Prescribed Fire Lags Its Targets*. Flagstaff, AZ: Cogplexiti.

How to Read This Paper

This white paper belongs to the Cogplexiti Systems & Decision Analysis line and follows the house analytical format, with the Analysis of Competing Hypotheses matrix in Exhibit B as its central deliverable. The matrix is designed to be argued with: every rating carries its rationale, contested cells are marked, and the adversarial re-rating result is printed. A citation-verification pass against primary sources is required before publication; production insertion points for current figures are marked in gold.

Contents

How to Read This Paper 3

Contents 3

1. Executive Summary 4

2. Support and Analysis..... 5

 2.1 The Gap, Stated Honestly 5

 2.2 Six Hypotheses at Full Strength..... 6

 2.3 The Matrix 7

3. Red Team Assessment 10

4. Dissenting Alternative..... 11

5. Recommended Course of Action12

Exhibit B — The Matrix (v1)13

References.....15

1. Executive Summary

Prescribed fire accomplishment lags national targets despite the largest fuels-management appropriations in the program's history. Every serious observer agrees the gap exists. Nothing else about it is agreed, because the debate over its cause is conducted as single-cause advocacy, and every cause has a constituency: it is funding, say the appropriators and the partner organizations who built the funding case; it is workforce, say the agencies watching their qualified benches thin; it is burn windows, says the climate literature; it is liability, say the burn bosses; it is regulatory throughput, say the line officers staring at planning timelines and smoke permits; it is institutional incentives, say the researchers who watch suppression eat everything else alive every August. Each constituency argues by confirmation — marshaling the evidence consistent with its hypothesis — and the argument never resolves, because most of the evidence is consistent with several hypotheses at once. Thirty years of this have produced a debate with no adjudication mechanism.

Analysis of Competing Hypotheses inverts the method (Heuer and Pherson 2015; UFMCS 2019). Hypotheses are not proven by the evidence consistent with them; they are eliminated by the evidence inconsistent with them, and the analytical work concentrates on the inconsistency column of the matrix — the diagnostic evidence that separates explanations rather than decorating all of them.

Run that way, the matrix produces findings the confirmation debate cannot. The most diagnostic single pattern in the record — appropriations at historic highs against accomplishment that has not moved proportionally — substantially disconfirms funding as the binding constraint: necessary, not binding. The research documenting large numbers of unused in-prescription days disconfirms shrinking burn windows as a sole constraint and is diagnostic *toward* workforce and incentives — a window that opens with no one positioned to burn in it is a staffing-and-priorities signature, not a climate one. What survives the matrix is a conjunction: workforce availability inside the windows, a liability and consequence asymmetry that prices escapes at career scale and non-treatment at zero, and an institutional incentive structure in which suppression takes precedence over the program that would reduce the need for it — with regulatory throughput binding hard in specific regions and airsheds rather than everywhere.

The surviving structure carries the paper's central finding: the binding constraint is not the same constraint everywhere. It differs between the Southeast and the Sierra, between a refuge and a national forest, between April and August. That is why every national single-lever fix has failed and why the next one will too — a lever applied nationally addresses the binding constraint only where it happens to bind, which is a category error dressed as a policy. The recommended course of action follows from the finding: publish the matrix with its reasoning exposed, name the short list of missing evidence that would

separate the surviving hypotheses, give field units a one-page instrument for diagnosing which constraint binds *them*, and shape the legislative ask as a portfolio sized to what the matrix shows each component would actually buy.

2. Support and Analysis

2.1 The Gap, Stated Honestly

The pattern this paper analyzes must survive its own measurement caveats, so the caveats come first rather than waiting for critics to discover them.

National fuels-treatment accomplishment is reported in activity acres, and activity acres are a generous unit: the same footprint treated twice counts twice, mechanical treatment counts alongside fire, and "accomplishment" measures work performed, not unique acres changed or exposure reduced (the counting problems are documented in the research literature; Kolden 2019). Reported national prescribed fire totals are dominated by the Southeast, where frequent-return burning on the same footprints is both ecologically correct and statistically flattering; the western states, where the fuels crisis concentrates, contribute a fraction of the national figure, and that fraction has been roughly flat for two decades against every escalation of rhetoric and target (Kolden 2019). Against this sits the input side: the 2021–2022 appropriations cycle delivered multibillion-dollar fuels and forest-management funding through the infrastructure and reconciliation bills, and the Forest Service's ten-year strategy committed to treatment scales several multiples of historical accomplishment.

The standing benchmark entering FY2026 is the Forest Service's ten-year goal of treating up to 50 million additional acres in the West—20 million acres on National Forest System lands and supporting 30 million acres on other federal, state, Tribal, and private lands. That equates to a nominal pace of five million additional treatment acres per year, although the current FY2026 and FY2027 budget justifications do not state a discrete national acreage target for FY2026. The funding figures require the same definitional discipline as the acreage figures: the Infrastructure Investment and Jobs Act provided approximately \$5.5 billion for Forest Service lands and programs broadly, and the Inflation Reduction Act provided another \$5 billion, but the amounts specifically identified for Forest Service wildfire-risk and hazardous-fuels work were approximately \$1.4 billion and \$1.8 billion, respectively—about \$3.2 billion combined. The most recent final agency report, for FY2025, records approximately 3.1 million acres of hazardous-fuels activity across federal, state, and private lands, including 2.86 million acres on National Forest System lands and nearly 1.4 million acres treated with prescribed fire—down from 4.29

million total activity acres and 2.19 million prescribed-fire acres in FY2024. These remain activity acres, not unique landscape acres: separate mechanical, fire, maintenance, or repeat treatments may each generate an accomplishment entry on the same footprint, and the reported total does not itself demonstrate how many unique acres changed condition or how much community exposure was reduced.

The honest statement of the gap, then: measured in the units the system itself reports, and granting every caveat about those units, accomplishment has not scaled with inputs, ambition, or need — and in the regions of greatest need it has barely moved. The caveats do not soften the pattern. They sharpen it, because every measurement generosity in the reporting inflates the numerator. Whatever the true gap is, the reported gap understates it.

2.2 Six Hypotheses at Full Strength

ACH requires each hypothesis stated at its strongest — as its most informed advocate would state it — before any is tested. Each carries a constituency, and each constituency is locally rational: the hypothesis each group holds is the constraint each group actually experiences from its chair. The discipline of this corpus applies here as everywhere: naming the constituencies is diagnosis, not accusation.

H1 — Funding. Prescribed fire is labor, planning, and contingency resources, all of which cost money the program historically never had; targets were aspiration unbacked by appropriation. The gap is what underfunding looks like, and the recent appropriations are the experiment in progress. *Constituency*: appropriators, partner organizations, and agency budget offices — the coalition that built the funding case and won it.

H2 — Workforce. The people are the constraint: burn bosses and qualified module leadership are scarce, the qualified are dual-hatted into suppression, and the burn windows fall in spring and fall — precisely when seasonal workforces are onboarding or gone. Money without hands burns nothing. *Constituency*: the agencies' fire staffs, who watch fully funded units fail to fill an RXB2 order in April.

H3 — Burn windows. The climate is closing the aperture: fuels too dry, winds too erratic, seasons in which prescription parameters, smoke dispersion, and staffing align arriving less often and less predictably. The gap is atmospheric, and no policy lever reopens a window physics has closed. *Constituency*: the climate-fire research community and every burner who has watched a planned unit sit unburnable through a whole season.

H4 — Liability and consequence asymmetry. The decision environment prices the two failure modes at wildly different rates: an escape carries investigation, career damage, personal legal exposure, and national visibility; a unit left unburned carries

nothing — no review, no file, no name attached when the deferred fire arrives a decade later. Rational agents facing that asymmetry defer, and the gap is the sum of a million rational deferrals. *Constituency*: burn bosses and line officers — the people whose signatures are on the plans.

H5 — Regulatory throughput. The compliance stack — environmental planning, smoke management, air-quality permitting across multiple jurisdictions — meters the program to a trickle regardless of funding or intent, and it embeds a documented perversity: prescribed smoke is permitted, monitored, and counted against airshed attainment, while wildfire smoke is excluded as an exceptional event, so the regulatory apparatus structurally prefers the smoke nobody chose to the smoke somebody planned (Engel 2013). *Constituency*: line officers and state forestry programs navigating the stack; air regulators occupy the other chair of the same hypothesis, also rationally.

H6 — Institutional incentives. The organization measures, funds, promotes, and mourns around suppression; prescribed fire is what the system does with its residual attention. Suppression assignments strip prescribed-fire-qualified staff at need, targets are set in acres that reward easy burning, and the career path runs through the incident, not the burn unit. The gap is the revealed preference of the institution, visible in every August. *Constituency*: the research and reform community — and, privately, much of the workforce.

2.3 The Matrix

The full matrix is Exhibit B: eleven evidence items rated consistent (C), inconsistent (I), or neutral/ambiguous (N) against each hypothesis, with the rating rationale exposed per row and contested cells marked. The body carries the diagnostic findings — the inconsistency column — because that is where ACH earns its keep. Four evidence items do most of the separating.

E1 — Record inputs, flat outputs. The appropriations experiment of 2021–2022 is the closest thing the debate has to a controlled intervention on H1, and its early returns are the single most diagnostic pattern in the record: funding rose by multiples while accomplishment, in exposure-relevant regions, did not move proportionally. This is inconsistent with H1 as the *binding* constraint — if money were the wall, money should have moved the wall — while remaining fully consistent with H1 as a necessary condition that reveals the next constraint the moment it is satisfied. The rating is marked contested in one respect: implementation lag is real, and H1's advocates deserve the concession that a fair test requires several more seasons. The matrix carries the concession and notes

what it purchases: a hypothesis whose test is always several seasons away is approaching unfalsifiability, and the burden shifts a little more each flat year.

E2 — Unused windows. The retrospective burn-window research is quietly devastating to H3 as a sole explanation: analyses of actual weather records against actual prescription parameters find substantial numbers of in-prescription, dispersion-compliant days going unused — in some studied landscapes, most of them (Striplin et al. 2020). Climate change is real and the windows are shifting and will shift further; but a system that does not burn on a large fraction of the days it *could* burn is not window-limited, and the unused-window pattern is affirmatively diagnostic toward H2 and H6: windows open when the seasonal workforce is not yet hired, when the qualified are committed elsewhere, when the planning stack has not produced a burnable unit to meet the weather — that is a staffing, sequencing, and priorities signature. H3 survives as a regional and worsening constraint, disconfirmed as the national explanation.

E3 — The 2022 sequence. Hermits Peak–Calf Canyon is the most consequential single data point in the record, and the matrix reads it with the same system-level discipline this corpus applies to every case: the reviews located the failure in prescription development, planning inputs, and holdover risk assessment — the system's processes — not in rogue execution (USDA Forest Service 2022a). What the event disconfirms and confirms is institutional, and it is best read through the *response*: a program with a documented hold rate above ninety-nine percent experienced an escape, and the institutional answer was a ninety-day national pause of the entire program (USDA Forest Service 2022b) — followed, months later and in a different state, by a county sheriff arresting a federal burn boss after an escape onto private land, an event without living precedent in the community. No comparable institutional mechanism has ever paused, reviewed, or arrested anyone for the millions of acres *not* treated, whose costs arrive later, diffusely, and unattributed. The asymmetry could not be more cleanly institutionalized, and the rating follows: strongly consistent with H4 and H6 — the system behaves as if one failure mode is intolerable and the other invisible — and diagnostic against any account in which burn-boss risk aversion is dismissed as timidity. The deferrals are not a bias. They are a correct reading of the incentive structure the institution has published through its own behavior. The compensation legislation that followed the New Mexico fires, at multibillion-dollar scale, completed the lesson for every signature-holder in the country.

E4 — The natural experiments. The state-and-regional variation the national debate treats as noise is the matrix's cleanest signal. The Southeast burns an order of magnitude more than the West under the same national funding regime and the same climate direction — but under different liability doctrine (certified-burner statutes, gross-negligence standards), different smoke regulatory practice, different land tenure, and a culture in which burning is the norm rather than the exception. Where western states have moved pieces of that package — liability reform, certified burner programs, claims

funds — private and state burning has responded (the manager-survey literature documents the mechanism from the practitioner side: Quinn-Davidson and Varner 2012; Schultz, McCaffrey, and Huber-Stearns 2019; Miller, Field, and Mach 2020). This is inconsistent with any hypothesis that predicts uniform performance under uniform national policy — which is most of the single-cause versions — and is the empirical foundation of the paper's central finding: the constraint that binds is local, and the Southeast is not evidence that the system works; it is evidence of what a different constraint configuration produces.

FIGURE 1

The matrix at a glance

Eleven evidence items rated against six hypotheses. Evidence consistent with everything separates nothing; the analysis is driven by the inconsistency column. Full ratings and rationale in Exhibit B.

		H1 funding	H2 workforce	H3 windows	H4 liability	H5 regulatory	H6 incentives
E1	record inputs, flat outputs	I†	C	N	C	C	C
E2	unused compliant burn days	N	C+	I	C	N†	C+
E3	the 2022 response function	N	N	N	C+	N	C+
E4	SE/West differential	I	C	N†	C+	C	C
E5	practitioner surveys	I†	C	N	C	C	C
E6	smoke-rule asymmetry	N	N	N	C	C+	C
E7	NEPA mostly CE, rarely sued	N	N	N	N	I†	N
E8	window–workforce misalignment	N	C+	C	N	N	C
E9	2022 review: process failure	N	C	N†	C	N	C
E10	one-way suppression valve	N	C	N	N	N	C+
E11	exposure-weighted acres (thin)	N	N	N	C†	C†	C†
verdict		demoted	survives	demoted	survives	regional	survives
I	inconsistent — the diagnostic column						
C+	strongly consistent						
C	consistent						
N	neutral						

† contested cell — argued in the row rationale; findings survive adverse re-rating of every †

The question is not which explanation is right. It is which constraint binds here.

The surviving structure. Run to completion, the matrix eliminates no hypothesis outright and demotes two: H1 to necessary-not-binding, H3 to regional-and-worsening-not-sole. H2, H4, and H6 survive robustly and — the matrix's second-order finding — they interact as a self-reinforcing set: the incentive structure (H6) allocates the scarce workforce (H2) away from burning, and the consequence asymmetry (H4) makes every individual allocation decision locally correct, so no actor anywhere has to be wrong for the system to systematically underburn. H5 binds hard in specific airsheds and planning contexts and not in others — the throughput evidence is genuinely mixed across regions,

and the matrix marks those cells contested rather than forcing a verdict the data will not support (Fleischman et al. 2020 cuts against H5 as a universal binder; the smoke asymmetry of Engel 2013 cuts toward it in constrained airsheds). The question the debate has been asking — *which explanation is right?* — is therefore malformed. The answerable question is *which constraint binds here, now* — and that question has a different answer on a Georgia refuge in February than on a Sierra forest in October, which is exactly what the failure record of national single-lever fixes would look like if the conjunctive model were true. It does, and it is.

3. Red Team Assessment

3.1 Qualitative ACH risks matrix theater. A consistency rating on contested evidence is a judgment wearing a grid, and a matrix whose cells cannot be interrogated is pseudo-quantification — the exact vice this paper attributes to the acre metric. The exposure is real and the countermeasure is total transparency: every row in Exhibit B carries its rating rationale in prose, contested cells are marked and argued rather than smoothed, and the sensitivity move is performed in public — the paper's findings are re-stated under the most adverse defensible re-rating of the contested cells, and the reader is shown that the conjunctive finding survives while the single-cause verdicts do not. If a critic can flip a finding by defensibly re-rating a cell, the paper wants that critic to be able to try, in writing, against the published rationale. A matrix that survives adversarial re-rating is rigor; one that cannot be re-rated is theater. This paper stakes its method on being the first kind.

3.2 The matrix disconfirms the favorite hypotheses of natural allies. The funding narrative this analysis demotes is the narrative that agencies, partner organizations, and congressional champions used — in good faith, and correctly at the time — to win the appropriations the program needed for decades. Publishing E1's finding hands ammunition to appropriations skeptics and risks reading, to allies, as ingratitude with a matrix attached. The relationship cost is real and the paper should carry it knowingly rather than discover it in the reactions. Three mitigations are structural. The finding is *necessary-not-binding*, not *wasted* — the paper's explicit claim is that the money was required to reveal the next constraint and remains required to relieve it, which is an argument for spending the appropriation better, not for clawing it back. The portfolio COA gives every constituency a live lever, including the funders, whose money is what the workforce and liability components spend. And the alternative to this paper is worse for the allies than the paper is: several more flat years will produce the appropriations-skeptic version of this analysis without the portfolio, without the local-rationality discipline, and without the reinvestment argument. Better that the adjudication arrive from inside the coalition, carrying its own remedy.

3.3 The data may be too soft to anchor the anchor. The paper leans on the record-funding-flat-accomplishment pattern while conceding that accomplishment reporting is inconsistent, activity-acre inflated, and definitionally unstable — and a critic can fairly ask whether data too soft to trust as a numerator is somehow firm enough to carry the paper's most diagnostic finding. The answer is directional and must be stated carefully: every documented softness in the reporting *inflates* accomplishment — double-counted footprints, mechanical acres in fire totals, activity counted over outcome — so the measurement error runs in one known direction, and the true gap is bounded below by the reported one. A pattern that survives when all the noise favors the null is the strongest kind the available data can produce. The residual concession stands, though, and 3.1's discipline applies: the E1 cells are marked contested, the several-more-seasons caveat is printed, and the COA's first missing-evidence item is the reporting reform that would retire this objection permanently.

3.4 The constituency framing can read as false balance. Honoring six local rationalities risks implying six equally weighted truths, when the matrix itself concludes the weights are unequal. The discipline cuts one way in the diagnosis and another in the verdict: every constituency is right about the constraint it experiences, and the matrix is entitled to conclude — and does — that some experienced constraints bind the national outcome more than others. Local rationality is a method for reading actors without villains. It is not a promise that every hypothesis survives contact with the inconsistency column, and the paper should say so before a reviewer says it first.

4. Dissenting Alternative

The gap is the wrong object of analysis. The credible alternative refuses the paper's premise: acres are a political metric, not a risk metric, and adjudicating the causes of an acre gap dignifies a number that should never have been the target. Exposure is what matters — the right acres, near the values, in the condition classes where treatment changes outcomes — and against an exposure-weighted standard the system's performance might look far worse than the acre gap suggests (if reported acres concentrate where burning is easy, flat, and far from anything) or meaningfully better (if the modest western accomplishment is well-targeted). Worse, a system managed against an acre target is *actively incentivized* toward the easy acres: the metric rewards exactly the burning that buys the least. On this reading, the entire thirty-year debate this paper adjudicates is an artifact of a malformed objective function, the correct paper is about the metric, and this one — however disciplined its matrix — is a rigorous analysis of the wrong question.

The dissent is right about the metric and wrong that it dissolves the question, and the concessions it earns are already installed. It is right enough that the matrix carries the

exposure question as an evidence row (E11), that the measurement reform leads the missing-evidence list, and that the COA's legislative section treats metric reform as a component co-equal with liability and workforce. It does not dissolve the question, for one empirical reason the exposure lens sharpens rather than blunts: the constraint conjunction the matrix identifies — workforce, consequence asymmetry, incentives — binds *hardest* on precisely the exposure-relevant acres, where units are complex, smoke falls on communities, escapes have addresses, and the liability calculus is at its most punishing. Re-denominate the gap in exposure terms and the surviving hypotheses survive harder. The metric critique and the constraint diagnosis are not rivals; they are the same finding read at two resolutions, and the portfolio needs both.

5. Recommended Course of Action

5.1 Publish the matrix with its reasoning exposed. Exhibit B ships whole: eleven evidence rows, per-row rationale, contested cells marked and argued, and the adversarial re-rating exercise printed as an appendix rather than performed privately. The transparency is not a stylistic choice; it is the method's answer to Red Team 3.1 and the paper's credential for 3.2 — an adjudication that shows its work can be argued with, which is what distinguishes it from the thirty years of assertion it replaces.

5.2 Name the missing evidence, decision-shaped. The short list of collectable facts that would separate the surviving hypotheses, published as a research ask in the form the researcher vignette in this line demands — commissioned in the shape of the decisions it must inform (Coil 2026c). Four items lead: unique-acre, exposure-weighted accomplishment reporting to retire the activity-acre ambiguity; windows-offered-versus-windows-burned tracking as a standard metric, converting the unused-window finding from a research result into an operational statistic; a national census of prescribed-fire-qualified personnel and their actual availability inside historical window periods; and systematic tracking of professional and legal outcomes following escapes versus non-treatment, to put numbers on the asymmetry the system currently publishes only through behavior.

5.3 Field the unit-level diagnosis instrument. The paper's conversion from national argument to operational tool: a one-page binding-constraint assessment a district, refuge, or state unit runs on itself — five questions, one per surviving constraint family, each keyed to observable local facts (did funded units go unburned in open windows; what happened to the last local escape and the last decade of non-treatment; where did the qualified staff physically spend the window weeks; which permits and plans were the long pole; what does the unit's acre mix look like exposure-weighted). The instrument's output is not a score; it is the name of the local binding constraint — which is the only

information a portfolio policy needs from the field, and information no national dataset currently carries.

5.4 Shape the legislative ask as a portfolio, sized by the matrix. The lever list follows the surviving hypotheses, each component sized to what the evidence says it buys and where: liability reform on the certified-burner and gross-negligence model, paired with claims funds, targeted at the jurisdictions where E4 shows the asymmetry binding; a dedicated prescribed-fire workforce with suppression-insulated availability, targeted at the unused-window signature of E2; smoke-regulatory parity addressing the perverse preference E3's literature documents, targeted at the constrained airsheds where H5 actually binds; and metric reform per the dissent, so the portfolio's own success is measured in exposure rather than acres. The portfolio's design principle is the paper's finding stated as policy: no single lever, because there is no single constraint — and every component carries its diagnostic, so the next Congress can see which levers moved what.

End state. A debate converted into a diagnosis, a diagnosis converted into an instrument, and a policy shaped like the problem — so that the unburned bench above Pine Bench, fuels running unbroken to the ridge, stops being an argument about whose explanation is right and becomes a unit with a named constraint, a matched lever, and a window someone is staffed to use.

Exhibit B — The Matrix (v1)

Method: Analysis of Competing Hypotheses (Heuer and Pherson 2015). Ratings: C = consistent, I = inconsistent, N = neutral/ambiguous. † = contested cell, argued in rationale. Diagnosticity lives in the I column; evidence consistent with everything separates nothing. Hypotheses: H1 funding · H2 workforce · H3 windows · H4 liability asymmetry · H5 regulatory throughput · H6 institutional incentives.

****E1** — Multibillion-dollar 2021–2022 appropriations; accomplishment not proportionally moved H1 I†* · H2 C · H3 N · H4 C · H5 C · H6 C *Rationale:** the money experiment ran; the wall did not move proportionally — inconsistent with funding as binding, consistent with every hypothesis predicting a non-monetary wall behind the monetary one. †Contested on implementation lag; concession printed in 2.3, burden shifting each flat season.

E2 — Large fractions of in-prescription, dispersion-compliant days unused (Striplin et al. 2020). H1 N · H2 C+ · H3 I · H4 C · H5 N† · H6 C+ *Rationale:* a window-limited system uses its windows; unused compliant days are affirmatively diagnostic toward staffing and priorities. †H5 neutral-contested: permit timing can consume windows in constrained airsheds — regionally C.

E3 — Escape rate above 99% held; one escape triggers 90-day national program pause (USDA Forest Service 2022a, 2022b); first-in-memory arrest of a federal burn boss follows in-year; multibillion-dollar compensation act completes the signal. H1 N · H2 N · H3 N · H4 C+ · H5 N · H6 C+ *Rationale*: the institution's own response function prices the two failure modes — escape versus non-treatment — at incommensurable rates; deferral is the rational reading of the published incentive. No re-rating of this row was found defensible.

E4 — Order-of-magnitude Southeast/West accomplishment differential under uniform national funding and shared climate direction; state liability/certification reforms correlate with burning response where enacted (Kolden 2019; Quinn-Davidson and Varner 2012; Miller, Field, and Mach 2020). H1 I · H2 C · H3 N† · H4 C+ · H5 C · H6 C *Rationale*: uniform national policy producing order-of-magnitude regional variance is inconsistent with any nationally-uniform binding constraint; the variance tracks the liability/regulatory/cultural configuration. †H3: fuel types and window climatology differ regionally — partial confound, argued in row notes.

E5 — Manager and practitioner surveys rank capacity, risk, and institutional support above funding, with regional variation (Schultz, McCaffrey, and Huber-Stearns 2019; Miller, Field, and Mach 2020; Quinn-Davidson and Varner 2012). H1 I† · H2 C · H3 N · H4 C · H5 C · H6 C *Rationale*: the people closest to the constraint consistently name the non-monetary walls. †Self-report bias acknowledged; practitioners may under-name funding they've stopped expecting.

E6 — Prescribed smoke permitted and counted against attainment; wildfire smoke excluded as exceptional event (Engel 2013). H1 N · H2 N · H3 N · H4 C · H5 C+ · H6 C *Rationale*: documented structural preference for unplanned smoke over planned smoke; binds where airsheds are constrained, which correlates with exposure-relevant geography.

E7 — Most fuels projects proceed under categorical exclusions; NEPA timelines variable, litigation rare (Fleischman et al. 2020). H1 N · H2 N · H3 N · H4 N · H5 I† · H6 N *Rationale*: inconsistent with planning compliance as a universal binder. †Contested: CE availability does not equal planning capacity; workforce to produce even CE-level documentation loops back to H2.

E8 — Burn windows concentrate in spring/fall shoulder seasons, misaligned with seasonal workforce onboarding and suppression season creep (USDA Forest Service 2022b, workforce findings). H1 N · H2 C+ · H3 C · H4 N · H5 N · H6 C *Rationale*: the window-workforce misalignment is a design fact of the current employment model; lengthening suppression seasons consume the fall shoulder from the H6 direction.

E9 — 2022 national review locates failure in prescription development, planning inputs, holdover risk process — not rogue execution (USDA Forest Service 2022a). H1

N · H2 C · H3 N† · H4 C · H5 N · H6 C *Rationale*: system-process failure signature consistent with an overloaded, thinly-staffed planning apparatus (H2/H6) and with drier-than-recognized conditions (†H3 partial C: the climate is moving the prescription math faster than the process updates it).

E10 — Suppression draw-down strips prescribed-fire-qualified staff at need; no reciprocal mechanism strips suppression for burning. H1 N · H2 C · H3 N · H4 N · H5 N · H6 C+ *Rationale*: the one-way valve is the incentive structure rendered as a dispatch practice; rated on the system's standing resource-priority behavior.

E11 — Exposure-weighted accomplishment (dissent row): distribution of burned acres relative to values-at-risk and condition class — evidence currently thin by design of the reporting system. H1 N · H2 N · H3 N · H4 C† · H5 C† · H6 C† *Rationale*: what little exposure-resolved analysis exists suggests accomplishment skews toward low-complexity, low-exposure units — the acre metric's predicted fingerprint. †All cells contested pending the 5.2 reporting reform; the row exists to make the gap in the evidence itself visible.

Matrix verdicts. H1: demoted, necessary-not-binding (I in three of the four most diagnostic rows). H3: demoted, regional-and-worsening, disconfirmed as sole (E2). H2, H4, H6: survive robustly, interacting as a self-reinforcing conjunction (E2, E3, E4, E8, E10). H5: binds regionally — constrained airsheds and specific planning contexts — verdict deliberately split (E6 vs E7). Sensitivity: under the most adverse defensible re-rating of all † cells, the conjunctive finding stands; no single-cause hypothesis can be promoted to binding without un-rating an uncontested row.

References

- Coil, Jayson. 2026c. *The System Tier*. Cogplexiti Systems & Decision Analysis Series No. 1. Flagstaff, AZ: Cogplexiti.
- Engel, Kirsten H. 2013. "Perverse Incentives: The Case of Wildfire Smoke Regulation." *Ecology Law Quarterly* 40 (3): 623–672.
- Fleischman, Forrest, et al. 2020. "US Forest Service Implementation of the National Environmental Policy Act: Fast, Variable, Rarely Litigated, and Declining." *Journal of Forestry* 118 (4): 403–418.
- Heuer, Richards J., Jr., and Randolph H. Pherson. 2015. *Structured Analytic Techniques for Intelligence Analysis*. 2nd ed. Thousand Oaks, CA: CQ Press.
- Kolden, Crystal A. 2019. "We're Not Doing Enough Prescribed Fire in the Western United States to Mitigate Wildfire Risk." *Fire* 2 (2): 30.

- Miller, Rebecca K., Christopher B. Field, and Katharine J. Mach. 2020. "Barriers and Enablers for Prescribed Burns for Wildfire Management in California." *Nature Sustainability* 3: 101–109.
- Quinn-Davidson, Lenya N., and J. Morgan Varner. 2012. "Impediments to Prescribed Fire Across Agency, Landscape and Manager: An Example from Northern California." *International Journal of Wildland Fire* 21 (3): 210–218.
- Schultz, Courtney A., Sarah M. McCaffrey, and Heidi R. Huber-Stearns. 2019. "Policy Barriers and Opportunities for Prescribed Fire Application in the Western United States." *International Journal of Wildland Fire* 28 (11): 874–884.
- Striplin, Ryan, Stacy A. McAfee, Hugh D. Safford, and Michael J. Papa. 2020. "Retrospective Analysis of Burn Windows for Fire and Fuels Management: An Example from the Lake Tahoe Basin, California, USA." *Fire Ecology* 16: 13.
- University of Foreign Military and Cultural Studies (UFMCS). 2019. *The Red Team Handbook: The Army's Guide to Making Better Decisions*. Version 9.0. Fort Leavenworth, KS: TRADOC G-2.
- USDA Forest Service. 2022a. *Gallinas-Las Dispensas Prescribed Fire Declared Wildfire Review*. Washington, DC: USDA Forest Service.
- USDA Forest Service. 2022b. *National Prescribed Fire Program Review*. Washington, DC: USDA Forest Service.

No actor anywhere has to be wrong for the system to systematically underburn and for the debt to continue to accumulate.

Prescribed fire lags its targets despite the largest fuels appropriations in history, and the debate over why is conducted as single-cause advocacy — funding, workforce, windows, liability, regulation, incentives — each with a constituency, each arguing by confirmation. Thirty years of this have produced a debate with no adjudication mechanism, because most of the evidence is consistent with several explanations at once.

There is no single constraint, so there is no single lever. The question is not which explanation is right. It is which constraint binds here — and that answer differs by region, by unit, and by season.

The analysis inverts the method: Analysis of Competing Hypotheses tests explanations by disconfirmation, and the matrix's diagnostic evidence does the separating — record inputs against flat outputs, large fractions of compliant burn days unused, an institutional response function that pauses a program nationally for one escape and has never reviewed anyone for the acres not treated, and an order-of-magnitude regional differential under uniform national policy.

What survives is a conjunction — workforce, consequence asymmetry, and institutional incentives, self-reinforcing — with funding demoted to necessary-not-binding and every rating's rationale published for adversarial re-rating. The course of action follows the finding: a unit-level binding-constraint diagnosis, a decision-shaped research ask, and a legislative portfolio sized to what the matrix shows each lever buys.

Part of the Cogplexiti Systems & Decision Analysis line. Applies the local-rationality discipline of Nobody's Villain to the six constituencies; its diagnosis instrument carries the line's conversion from national argument to operational tool.

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.

