

Coarse by Design: Risk Communication and the Oregon Wildfire Map

Joakim A. Weill*

October 1, 2026

Abstract

Governments communicate risk through coarse categories. I show that when private behavior generates spillovers, a coarse but truthful risk disclosure can strictly dominate detailed disclosure, even for fully rational individuals. Coarseness is a second-best policy instrument: expanding the high-risk category draws additional individuals into private action that protects others. I apply this framework to Oregon’s 2022 Wildfire Risk Map, which assigned each of the state’s two million parcels to one of five risk categories and drew intense backlash from owners disputing its accuracy. I first show that Oregon’s free model predicts held-out fires about as well as the leading proprietary alternatives. The map’s coarse categories embed a trade-off between false alarms and misses, which I recover from the boundary of the high-risk category: over the two decades of fires the state could observe, catching one additional fire costs roughly twenty-five flagged parcels that do not burn. Given external estimates of adaptation costs and effectiveness, this permissive boundary is consistent with a planner who internalizes spillovers. Moving this boundary traces the set of maps that planners with different preferences over false alarms and misses would have drawn.

JEL Classification: Q54, H23, D62, D83, Q58

*Federal Reserve Board of Governors, email: joakim.a.weill@frb.gov. I thank Chris Dunn and Andy McEvoy for generously sharing data, and Pierre Courtois, Erik Heitfield, Jacob Gellman, Zack Grossman, Mallick Hossain, You Suk Kim, Nicolas Legrand, Nicolas Treich, and many seminar participants for valuable feedback. The views expressed in this paper are mine and do not reflect the views of the Federal Reserve Board of Governors or of any other person associated with the Federal Reserve System.

1 Introduction

Governments typically communicate risk through a few labeled categories. Flood maps focus on the high-risk floodplain, disease outbreaks such as COVID-19 are tracked in a handful of risk tiers, extreme heat advisories trigger on a temperature threshold, air quality is color-coded from good to hazardous. Coarseness is usually justified on cognitive grounds: discrete categories translate risk into formats lay audiences can interpret and act on (Slovic, 1987; Fischhoff, 1995; Mullainathan, Schwartzstein and Shleifer, 2008). The choice to coarsen is consequential because labels move behavior (Figlio and Lucas, 2004; Hino and Burke, 2021): where the thresholds are set determines who ends up in which category, and therefore who acts on what they learn.

Oregon illustrates what is at stake in this design choice. In June 2022, as mandated by Senate Bill 762, the state released a wildfire risk map assigning each of nearly two million parcels to one of five risk categories, and 80,000 property owners were notified that they fell in the top two (“high” and “extreme” wildfire risk). The map triggered an immediate backlash: thousands of owners appealed, many viewing the coarse categories as inadequate, others arguing the underlying estimates failed to reflect fine-scale local conditions. Town halls were derailed by threats of violence and viral conspiracy theories, and within a month lawmakers from both parties called for the map’s withdrawal (Profita, 2022; Davis, 2025). Figure 1 reproduces the published map.




Figure 1: The Oregon 2022 Wildfire Risk Map

Notes: The five-category wildfire risk map released in June 2022 under Senate Bill 762, covering every taxlot in the state. Owners of parcels classified High or Extreme received notification letters.

This paper asks how governments should communicate risk, and what trade-offs arise when they do so through coarse categories. I focus on the case where individuals are fully rational, ruling out the standard cognitive justification for coarseness. I then take these questions to the data, treating Oregon’s published wildfire risk map as a policy choice.

I begin with a simple model where a social planner chooses how to communicate risk. Each property is owned by a single individual and has a continuous, fixed risk type that this individual does not observe. Individuals know the population distribution of risk types and update beliefs by Bayesian inference given the information they receive. They then decide whether to invest in adaptation, taking the overall adaptation rate as given. Adaptation reduces an individual’s own expected loss and generates spillovers that lower losses for other properties. The social planner observes property risk types with noise and chooses an information structure: either disclose its property-level risk estimate to each individual, or instead disclose a coarse but truthful map of risk (a “lower risk” or “higher risk” label).

The adaptation externality and the information friction in this framework capture two central aspects of disaster risk. Most forms of disaster adaptation create spillovers: in wild-fire settings, defensible space and home hardening reduce not only a home’s own probability of ignition, but also the probability that it becomes an ember source that threatens nearby structures (Arrowsmith et al., 2021; Zamanialaei et al., 2025). Similar spillovers arise with other perils: individual protective actions that reduce emergency visits during extreme heat can alleviate hospital congestion, a channel that accounts for a large share of excess mortality (Aguilar-Gomez, Zivin and Neidell, 2025), and more broadly, property insurance coverage generates spillovers by easing individuals’ post-disaster liquidity constraints that can then support local economic recovery (You and Kousky, 2024).¹ In addition, disaster risk is unobserved and must be inferred from models built to represent rare events, so estimates are noisy, model-dependent, and sometimes inconsistent across commercial vendors (Charpentier, 2008; Dietz and Niehörster, 2021; Boomhower et al., 2024; Moore, 2025; Heitfield, 2026; Condon, 2023).

I first show that when adaptation generates spillovers, a coarse risk map can strictly dominate detailed disclosure of the planner’s property-level risk estimates, even though individuals are perfectly rational and the planner never misrepresents risk. This result occurs because under detailed disclosure, only the properties with the highest estimated risk adapt, as individuals do not internalize the benefits their investments confer on neighbors. In contrast, under a coarse (but still truthful) map, the planner can optimally set the higher-risk category threshold below the private cost–benefit cutoff. A coarse map can therefore enlarge the number of adapters, drawing some properties with moderate estimated risk into adaptation. This “benevolent cartographer” result holds for any noise level, provided the underlying risk model remains informative in a monotone sense made precise below. The result is an instance of the theory of the second best (Lipsey and Lancaster, 1956). Costless information is always weakly valuable to an individual decision-maker (Blackwell, 1953), but when disclosure moves equilibrium behavior and the externality remains unpriced, correcting the other distortion, individuals’ imperfect information, can lower welfare. Because individ-

¹The mechanism studied here can operate wherever public risk information moves private actions that benefit others, and several forms of public risk communication are already coarse: disease-risk tiers encourage vaccination, which protects others by reducing transmission; cybersecurity advisories, which bin continuous vulnerability scores into severity tiers, prompt firms to patch systems that would otherwise become staging grounds for attacks on others; and color-coded air quality alerts discourage driving on high-pollution days, lowering the pollution everyone breathes.

uals understand both the risk distribution and the noise in the planner’s signal, the value of coarsening operates only through spillovers.

Any coarse classification rule embeds a trade-off between Type I and Type II errors: classifications that induce excess or insufficient adaptation relative to the private optimum of an individual who knows the property’s true risk. The benevolent cartographer mechanism rationalizes a welfare-maximizing planner’s willingness to accept more Type I errors in exchange for fewer Type II errors: the planner intentionally expands the “higher-risk” group to draw extra individuals into adaptation that benefits their neighbors.

I apply this framework to Oregon’s 2022 Wildfire Risk Map.² The map’s policy-relevant boundary is the Moderate/High threshold: the notification letters went to owners of parcels mapped High or Extreme, and defensible-space and home-hardening requirements for these two categories were considered at one point (Oregon Building Codes Division, 2022; Oregon Wildfire Programs Advisory Council, 2022). Because “true” underlying wildfire risk is unobserved, the map’s assignments cannot be directly labeled as false positives or false negatives. I instead benchmark the map against realized fires, assembling a parcel-level dataset combining the Oregon model’s underlying continuous risk score for the universe of residential properties in the state and fire perimeters from 2000 to 2025. To assess the relative performance of the underlying Oregon model, I complement this dataset with property-level estimates from Cotality and First Street, two proprietary wildfire risk models used in the insurance industry. I split the analysis into two periods: the pre-map window (2000 to 2021) reproduces the planner’s information set at the time the map was drawn, while the post-map window (2022 to 2025) provides a held-out test against ground truth.

The first empirical finding is that Oregon’s public model predicts wildfires about as well as the leading proprietary alternatives. Evaluated on held-out 2022 to 2025 fires, the Oregon score ranks a randomly drawn burned parcel above a randomly drawn unburned one with probability 0.92, similar to Cotality’s catastrophe model (0.94) and the First Street model (0.85); pooling all three signals improves on the Oregon risk score alone by only two points. Just above the Moderate/High boundary of the Oregon map, parcels burn at a higher rate under the Oregon ranking than at the same ranks of either proprietary model.

The second finding characterizes the trade-off embedded in Oregon’s published map. At the observed Moderate/High threshold, the map classifies 10 percent of parcels as high risk; against held-out 2022 to 2025 fires, this rule correctly catches 76 percent of the parcels that subsequently burned, while incorrectly flagging about 10 percent of those that did not. The implicit weight on false positives that rationalizes the cutoff, computed from the planner’s pre-map information, is 0.039 (95 percent confidence interval 0.026 to 0.059): over the 2000 to 2021 calibration window, the map trades roughly twenty-five parcels needlessly flagged for each additional missed fire it avoids. Recovering the same weight from held-out fires on a common annualized basis yields a value of the same order of magnitude. In the benchmark without spillovers, this recovered weight is welfare-exact: it equals the marginal

²Terminology differs across fields. In the natural-hazards literature, risk is hazard weighted by exposure and vulnerability, so a map built only from fire likelihood and intensity is properly a *hazard* map, and Oregon has since renamed the map the *Oregon Wildfire Hazard Map*. I use “risk” throughout, both because it is the object’s original name and because it matches the common economic usage.

welfare cost of a false positive relative to a false negative. Under spillovers, if the published threshold is the planner’s optimum, theory places the weight strictly below the private-optimum weight. In most calibration exercises based on external estimates of adaptation costs and effectiveness, this prediction holds: Oregon’s threshold is consistent with a planner who also counts the benefits adaptation confers on neighbors.

Finally, I use the calibrated burn probabilities to draw counterfactual maps at alternative cutoffs, showing that the 2022 map is one point on a continuum of coarse disclosure rules: a conservative rule flagging only 1 percent of parcels would have missed 96.1 percent of subsequently burned parcels, while Oregon’s rule, relative to a candidate flagging 4 percent, cuts misses by 70 percent at a price of roughly 54 false positives per avoided miss.³

Taken together, these findings add a justification for coarse risk communication beyond the usual cognitive one: coarseness can have value wherever private action benefits others. Such spillovers are pervasive in risk settings because hazards propagate across neighbors and strain shared response capacity. The findings also clarify why the question that dominates the debate, whether the underlying model is accurate, would leave the policy question open even if settled. Any coarse risk communication converts a continuous estimate into categories through a threshold, and that threshold is a policy choice about which errors individuals bear. Evaluating a risk communication therefore requires recovering the trade-off it encodes.

This paper builds on three literatures. The first is the environmental economics literature on the choice and design of policy instruments to tackle externalities (Baumol, 1972; Weitzman, 1974; Keohane, Revesz and Stavins, 2019), and in particular on information-based instruments such as disclosure programs and certifications (Tietenberg, 1998; Kotchen, 2013; Cicala, Hémous and Olsen, 2022). This literature treats information as an instrument for shifting behavior, but the granularity of the information released is usually not a policy choice. The closest antecedent is Rodemeier and Löschel (2025), who find that a coarser energy-efficiency nudge can raise welfare by inflating consumers’ perceived savings, which curbs the carbon externality.⁴ I show that coarseness can itself be a policy instrument to tackle externalities, by operating on the rational pooling of beliefs. This instrument is generally second best: the optimal coarsening attains the first best, implemented by score-contingent Pigouvian taxes on non-adaptation, only when the pooled high bin weakly prefers to adapt at the first-best cutoff (Appendix B).

The second literature is information design and disclosure. In the Bayesian persuasion framework of Kamenica and Gentzkow (2011), a sender who commits to a signal structure can gain by withholding information (see also Rayo and Segal, 2010), and more precise public information can reduce welfare when actions are interdependent (Morris and Shin, 2002; Angeletos and Pavan, 2007). Theories of coarse certification explain compressed grades

³This exchange rate is measured on the four-year held-out window, whereas the twenty-five-to-one headline is a 22-year quantity; Section 4.5 places the two windows on a common annualized basis.

⁴Label granularity appears elsewhere in this literature through other channels: Houde (2018) shows that a coarse binary certification and detailed energy disclosure move demand in different ways, and in Fischer and Lyon (2019) a planner or a single certifier may prefer a binary ecolabel, whereas competition between certifiers leads them to offer multitier labels. In a related result, Houde and Myers (2026) show that an efficiency standard can dominate a Pigouvian tax when consumers misperceive energy costs heterogeneously.

through participation, incentive, or market constraints (Harbaugh and Rasmusen, 2018; Goldstein and Leitner, 2018; Inostroza and Pavan, 2025; Vatter, 2025). The planner here is restricted to a truthful, monotone coarsening of the risk score, a form that is optimal (Kolotilin, Mylovanov and Zapechelnyuk, 2022; Lyu, Suen and Zhang, 2023) or imposed (Onuchic and Ray, 2023) in related information design problems. Coarseness is valuable only because adaptation generates physical spillovers among neighbors: widening the higher-risk category substitutes for the Pigouvian instrument the planner lacks.⁵ Empirical work on disclosure estimates how labels move consumers and providers (Jin and Leslie, 2003; Dranove et al., 2003), and Barahona, Otero and Otero (2023) solve for welfare-maximizing warning-label thresholds in an equilibrium model with firm reformulation; the contribution here is instead to recover the trade-off embedded in an observed classification rule. I treat the map as a prediction policy problem (Kleinberg et al., 2015): as in Kleinberg et al. (2018), I hold an observed decision rule to realized outcomes and draw counterfactual rules from the same predictions, and as in Chan, Gentzkow and Yu (2022), I separate the predictive skill of the score from the threshold applied to it. Oregon’s rule is a single explicit threshold on a published score rather than an aggregate of implicit human decisions, so the weight it reveals is read directly off the calibrated score.

The third literature concerns disaster risk information and the models that produce it. Flood maps and wildfire hazard zones move beliefs, home values, and insurance demand (Gibson and Mullins, 2020; Bakkensen and Ma, 2020; Hino and Burke, 2021; Mulder, 2024; Lee, 2023; Ma et al., 2024; Weill, 2026), and risk labels can directly mandate adaptation through building codes (Wagner, 2022; Ostriker and Russo, 2022; Baylis and Boomhower, 2026). Evidence comparing competing risk models remains limited, in part because vendor estimates are costly to obtain (Pollack et al., 2026); existing comparisons measure disagreement across models for flood risk in Los Angeles (Schubert, Mach and Sanders, 2024), hurricane loss projections across Florida ZIP codes (Heitfield, 2026), and a small set of properties worldwide (Global Association of Risk Professionals, 2025). I provide the first large-scale comparison of property-level risk estimates benchmarked against subsequently realized disasters. In closely related work, Issler et al. (2025) develop a spatiotemporal neural network to account for weekly or even daily changes in weather conditions, predicting short-term burn probabilities for California real estate; the models I benchmark instead target burn probabilities at the multi-annual scale relevant for adaptation and disclosure policy.

The rest of this paper proceeds as follows. Section 2 presents a simple framework of information disclosure under spillovers, while Section 3 characterizes the social planner trade-off between Type I and Type II errors. Section 4 applies the framework to Oregon’s 2022 Wildfire Risk Map, and Section 5 concludes. Appendix Figure S1 summarizes the paper’s logical structure.

⁵Closest are models of disclosure with externalities among receivers. Hedlund, Hernández-Chanto and Oyarzún (2024) show that a planner may pool intermediate infection rates when protection efforts are strategic complements, and Anunrojwong, Iyer and Manshadi (2023) that coarse wait-time information can screen low-need users out of a congested service. In both, the state is a common aggregate and pooling is across its realizations; here the state is property-level risk, and a coarse category pools heterogeneous neighbors whose adaptation protects one another. See Kamenica (2019) for a survey.

2 A model of coarse disclosure

2.1 Toy example: the brick, stick, and straw houses

To build intuition, consider three houses with risk θ_i , $i \in \{1, 2, 3\}$, and with baseline expected wildfire losses $\ell(\theta_1) = 12$, $\ell(\theta_2) = 18$, and $\ell(\theta_3) = 36$. Adaptation is binary and costs $c = 16$ (think of adaptation as creating defensible space around the home). Assume adaptation fully eliminates own damages, so a homeowner who adapts pays c and otherwise faces expected loss $\ell(\theta_i)\phi(A)$, where $\phi(A)$ is the residual risk faced by non-adapters and depends on the aggregate adaptation rate A . I parameterize spillovers through $\phi(A) = 1 - \frac{A}{2}$, so greater community adaptation lowers residual risk for everyone. Each homeowner takes A as given, as in the large population of the general model in Section 2.2, and adapts if his perceived private expected loss exceeds the cost: $\mathbb{E}[\ell(\theta_i) \mid \text{information}] \cdot \phi(A) \geq c$.

If the planner communicates detailed (granular) risk, each homeowner learns his own $\ell(\theta_i)$. In equilibrium, only the highest-risk homeowner adapts. To see this, suppose $A = \frac{1}{3}$. Then $\phi(A) = 1 - \frac{1}{6} = \frac{5}{6}$, so $\ell(\theta_3)\phi(A) = 36 \cdot \frac{5}{6} = 30 > 16$ and homeowner 3 adapts, while $\ell(\theta_2)\phi(A) = 18 \cdot \frac{5}{6} = 15 < 16$, so homeowners 2 and 1 do not adapt. This implies exactly one adapter, hence $A = \frac{1}{3}$. Total costs are $c + (\ell(\theta_1) + \ell(\theta_2))\phi(\frac{1}{3}) = 16 + (12 + 18) \cdot \frac{5}{6} = 41$.

Now consider a coarse threshold disclosure where the social planner labels houses 2 and 3 as “higher risk” and house 1 as “lower risk.” Homeowners are Bayesian updaters and therefore use the average loss in their bin: $\mathbb{E}[\ell(\theta) \mid \text{low}] = 12$ and $\mathbb{E}[\ell(\theta) \mid \text{high}] = (18 + 36)/2 = 27$. In equilibrium, both higher-risk homeowners adapt. To see this, suppose $A = \frac{2}{3}$. Then $\phi(A) = 1 - \frac{1}{3} = \frac{2}{3}$. A high-bin homeowner perceives expected loss $27 \cdot \frac{2}{3} = 18 \geq 16$ and adapts, while the low-bin homeowner perceives $12 \cdot \frac{2}{3} = 8 < 16$ and does not. This yields exactly two adapters, hence $A = \frac{2}{3}$. In this setting, the coarse map leads to lower total costs: $2c + \ell(\theta_1)\phi(\frac{2}{3}) = 32 + 12 \cdot \frac{2}{3} = 40$.

Thus, with detailed disclosure the equilibrium has only the highest-risk homeowner adapt, whereas with a single-threshold, two-zone coarse map the equilibrium has the two highest-risk homeowners adapt. The difference arises because coarse disclosure changes perceived risk in the high bin: homeowner 2 does not know whether he is the moderate- or highest-risk property, only that he is not the safest, which raises his posterior expected loss enough to induce adaptation. At the same time, the resulting increase in A reduces residual risk $\phi(A)$ for everyone, amplifying the benefits of moving from one to two adapters. Throughout, the threshold map is truthful: it shifts equilibrium behavior only by omitting within-bin variation. In this illustrative example the planner observes the true risk θ directly; the general model in Section 2.2 relaxes this by giving the planner only a noisy score m of θ .

2.2 General model

Each property has a risk type $\theta \in [0, 1]$ with cumulative distribution function F , and is owned by a single individual who decides whether to adapt it. There is a large population of such properties, so that no single property is large enough to influence population-level outcomes on its own. Let $A \in [0, 1]$ denote the aggregate adaptation rate, defined as the population

share of properties that adapt. Because no single property moves A , each individual takes A as given when deciding whether to adapt. The adaptation spillovers are captured by a residual risk factor

$$\phi(A) > 0, \quad \phi \in C^1([0, 1]).$$

ϕ can be constant (no spillovers) or satisfy $\phi'(A) < 0$ for all $A \in [0, 1]$ (positive spillovers), depending on the case. I normalize $\phi(0) = 1$, so that the baseline loss defined next is the expected loss when no property adapts. I also allow for a general mapping from type to baseline loss. Let

$$\ell : [0, 1] \rightarrow (0, \infty)$$

be continuous and strictly increasing, and interpret $\ell(\theta)$ as the expected loss of a type- θ property in the absence of any adaptation and abstracting from spillovers. Here θ indexes a property's baseline hazard, fixed with respect to individual behavior.

A non-adapting type θ faces expected loss $\ell(\theta) \phi(A)$, while an adapting type θ faces expected loss $(1 - \alpha) \ell(\theta) \phi(A)$, with $\alpha \in (0, 1)$, plus adaptation cost $c > 0$. Adaptation therefore operates on losses, privately through α and socially through ϕ . Adaptation is binary; this assumption fits a wide range of actual adaptation investments and delivers cleaner results. Appendix F relaxes it.

The private gain from adapting is the difference between these expected losses,

$$\alpha \ell(\theta) \phi(A) - c,$$

which for any fixed aggregate adaptation A is strictly increasing in θ because ℓ is strictly increasing.

Neither the planner nor the individuals observe a property's realized type θ . However, the planner has access to a disaster model that produces a noisy risk score m of θ , with cumulative distribution function G and density g . The joint distribution of (θ, m) is common knowledge to the planner and the individuals. The planner observes the realized score m , while the individuals learn only what the planner communicates.⁶ I impose a monotonicity condition on the risk signal:

Assumption 1 (Monotone signal). The conditional expectation $m \mapsto \mathbb{E}[\ell(\theta) \mid m]$ is finite and strictly increasing in m .

Lemma 1 (Single-crossing in the risk score implies cutoff rules). *Fix an aggregate adaptation rate A . Under Assumption 1, any best response is weakly increasing in m and hence takes a cutoff form: there exists $x \in \mathbb{R}$ such that the individual adapts for all $m > x$ and does not adapt for all $m < x$.*

Proof. By Assumption 1, $\mathbb{E}[\ell(\theta) \mid m]$ is strictly increasing in m . Multiplying by the positive scalar $\alpha \phi(A)$ and subtracting c preserves strict monotonicity, so $\mathbb{E}[\alpha \ell(\theta) \phi(A) - c \mid m]$ crosses zero at most once. The optimal private action is therefore a cutoff rule in m . $\square$

⁶Field evidence from Oregon is consistent with this assumption: homeowners' take-up of defensible space assessments does not vary with their modeled baseline fire probability (Baylis et al., 2026).

Motivated by this monotonic private response, I focus on threshold disclosure and adaptation rules that are cutoff rules in the score m . Let x be in a fixed interval $(\underline{m}, \overline{m})$ contained in the support of m . For any such cutoff, when properties adapt if and only if $m \geq x$, the induced adaptation rate is

$$A(x) \equiv \Pr(m \geq x) = 1 - G(x).$$

Given a threshold x in score space, the per-property expected social cost it induces is

$$\begin{aligned} C(x) &= \mathbb{E}[1\{m < x\} \ell(\theta) \phi(A(x)) + 1\{m \geq x\} ((1 - \alpha) \ell(\theta) \phi(A(x)) + c)] \\ &= \mathbb{E}[\ell(\theta) \phi(A(x))] - \mathbb{E}[(\alpha \ell(\theta) \phi(A(x)) - c) 1\{m \geq x\}]. \end{aligned}$$

I impose a standard boundary condition to rule out corner outcomes in which either all individuals adapt or none do.

Assumption 2 (Interior cutoff). *Fix an interval $(\underline{m}, \overline{m})$ contained in the support of m . For each equilibrium analyzed, the expected private gain from adapting satisfies the boundary sign conditions*

$$\lim_{m \downarrow \underline{m}} \mathbb{E}[\alpha \ell(\theta) \phi(A) - c \mid m] < 0 \quad \text{and} \quad \lim_{m \uparrow \overline{m}} \mathbb{E}[\alpha \ell(\theta) \phi(A) - c \mid m] > 0.$$

Combined with Assumption 1, Assumption 2 ensures that the expected private gain crosses zero on $(\underline{m}, \overline{m})$ and therefore that equilibrium adaptation is characterized by an interior cutoff.

2.3 Detailed vs coarse disclosure

I now consider two ways the social planner can disclose the noisy risk score m : detailed disclosure, and coarse (threshold) disclosure. The disclosure rule is common knowledge to the planner and the individuals.

Under *detailed* disclosure, the planner reveals the risk score m ; by Lemma 1 each individual adapts if and only if $m \geq m_D$, where the *detailed-disclosure equilibrium cutoff* m_D is the score at which the privately indifferent individual breaks even given the equilibrium-consistent aggregate $A(m_D) = 1 - G(m_D)$. Writing $\mathcal{H}$ for the private indifference schedule,

$$\mathcal{H}(x) \equiv \alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m = x], \quad \text{the cutoff solves} \quad \mathcal{H}(m_D) = c. \quad (1)$$

The schedule $\mathcal{H}$ is strictly increasing: the conditional mean rises with the score by Assumption 1, and $\phi(A(x))$ is nondecreasing in x because $A(x)$ falls in x and ϕ is nonincreasing. Hence $\mathcal{H}(x) = c$ has at most one solution on $(\underline{m}, \overline{m})$, and Assumption 2 provides one, so the detailed-disclosure equilibrium is unique.

Under *coarse* disclosure, the planner instead announces a threshold x on the risk score and reports only whether $m \geq x$ (“high”) or $m < x$ (“low”). Individuals in a given bin are informationally identical and share the bin posterior $\mathbb{E}[\ell(\theta) \mid m \geq x]$ or $\mathbb{E}[\ell(\theta) \mid m < x]$. Unlike under detailed disclosure, there is no marginal indifferent individual within a bin, so

the equilibrium adjusts through the share of (identical) bin members who adapt rather than through a within-bin cutoff.

To compare the two regimes, write $C^D \equiv C(m_D)$ for the per-property social cost under detailed disclosure and C^T for the cost under the planner's optimal *implementable* threshold m^* : $C^T \equiv C(m^*) = \min_{x \in (\underline{m}, \overline{m})} C(x)$ subject to the two bin incentive constraints, $\alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m \geq x] \geq c$ and $\alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m < x] \leq c$, so that the high bin adapts and the low bin does not. At every such threshold the equilibrium aggregate is unique.⁷ The unconstrained minimum of $C(\cdot)$ is the first-best benchmark C^{FB} discussed below.

Restricting attention to two bins is likewise without loss. Under any interval partition of the score whose bins each satisfy its incentive constraint, bin posteriors rise across bins by Assumption 1, so the adapters form a top-set $\{m \geq x\}$, and by the expression for $C(x)$ above the social cost depends only on this set. Because the two-bin partition at x pools each side most coarsely, it relaxes both incentive constraints by the most, so any adapting set implementable by a finer partition is also implementable by a single threshold at the same cost.

Whether optimal coarse disclosure merely replicates detailed disclosure or strictly improves on it turns entirely on adaptation spillovers, as the next proposition shows. Its second part relies on a marginal perturbation of the threshold around the detailed-disclosure cutoff m_D , and therefore requires the score to be well-behaved on the score range:

Assumption 3 (Score regularity). The score m has a density g on the interval $(\underline{m}, \overline{m})$ of Assumption 2, with g continuous and strictly positive there, and $m \mapsto \mathbb{E}[\ell(\theta) \mid m]$ is continuous on $(\underline{m}, \overline{m})$.

Proposition 1 (The benevolent cartographer's rule). *Suppose Assumption 1 holds.*

- (i) (No spillovers: equivalence.) *If there are no spillovers ($\phi(A) \equiv \phi_0 > 0$ for all $A \in [0, 1]$), then detailed disclosure and the optimal coarse disclosure induce the same adaptation behavior, and therefore $C^T = C^D$.*
- (ii) (Positive spillovers: strict dominance.) *If, in addition, $\phi'(A) < 0$ for all $A \in [0, 1]$ and Assumptions 2 and 3 hold, then coarse threshold disclosure strictly dominates detailed disclosure, $C^T < C^D$, and the optimal implementable threshold lies strictly below the detailed-disclosure cutoff: $m^* < m_D$.*

Appendix C proves both parts. Part (ii) follows from perturbing the disclosure threshold around the detailed-disclosure cutoff: at m_D the marginal adapter is privately indifferent, so

⁷With ϕ decreasing, each bin's incentive to adapt falls as the aggregate rises, so the aggregate best response is nonincreasing and its fixed point unique. At thresholds violating a bin constraint the fixed point has that bin indifferent and adapting only partially. Such thresholds are never optimal, so restricting the planner's problem to thresholds satisfying both constraints is without loss. When the high bin adapts only partially, raising the threshold moves its least risky members to the low bin and raises the average risk of those who remain, so the adaptation rate rises and residual risk falls for everyone. When the low bin adapts only partially, lowering the threshold at an unchanged adaptation rate replaces a random draw of low-bin adapters with its riskiest members, who gain more from adapting.

the perturbation has no first-order private effect and only an inframarginal spillover gain; a threshold slightly below m_D remains implementable because pooling gives the high bin strict slack in its incentive constraint.

Part (ii) comes from the externality that the private equilibrium leaves unpriced. Under detailed disclosure each individual adapts only when his own posterior loss clears the cost, ignoring that his adaptation lowers residual risk $\phi(A)$ for everyone; the cutoff is pinned at the private indifference point m_D , which the planner cannot move, so it sits too high and too few properties adapt. The planner counters this under-adaptation by coarsening the map: pooling the marginal types into a single “high” bin raises their perceived loss from $\mathbb{E}[\ell(\theta) \mid m]$ to the bin average $\mathbb{E}[\ell(\theta) \mid m \geq x]$, which lets the planner sustain a strictly lower cutoff $m^* < m_D$ and tip into adaptation individuals who would not have adapted under their own score; each new adapter in turn lowers $\phi(A)$ for the rest. Coarsening the map is therefore benevolent: the planner discards within-bin information because coarseness is the only instrument it has for lowering the cutoff below its detailed-disclosure level.

Figure 2 illustrates part (ii) in a stylized example. Rearranging the adaptation condition, an individual adapts when his posterior expected loss is at least $c/(\alpha\phi(A))$, which I call the adaptation hurdle; it rises with A because non-adapters free-ride on the spillover. Panel A shows detailed disclosure: beliefs rise with the score and clear the hurdle from m_D on, and the right plot finds the equilibrium rate $A(m_D) = 0.17$ as the fixed point at which the share of properties whose belief clears the hurdle equals A . Panel B shows two threshold maps. With the threshold at m_D , the high-bin constraint holds with slack and the low bin stays below the hurdle, so the adapters are those of detailed disclosure. The optimal threshold m^* spends that slack: as the threshold falls, the high-bin posterior falls and the hurdle rises until the two meet, the high-bin constraint binds, and adaptation rises to $A(m^*) = 0.27$. Social cost falls: the new adapters pay c for a private benefit below c , but they lower residual risk $\phi(A)$ on every property.

A natural benchmark holds the planner’s information fixed but adds an instrument: a planner who conditions on the *same* noisy score m , but who can levy a score-contingent Pigouvian tax on non-adaptation, so that each individual internalizes the spillover. Its first-best cutoff m_{FB} , the largest minimizer of $C(x)$, attains $C^{FB} \leq C^T$. This minimizer is assumed interior, $m_{FB} \in (\underline{m}, \bar{m})$, as Assumption 2 ensures for the detailed-disclosure cutoff m_D . The low bin’s incentive constraint never binds (Lemma 2 in Appendix B), so the threshold either attains the first best or, when the pooled high bin would not adapt at m_{FB} , rises above it and induces under-adaptation, $m^* > m_{FB}$. Appendix B gives the full welfare-gap decomposition and Figure S2 illustrates the under-adaptation case.

The model is deliberately simple: even with fully rational, risk-neutral individuals, coarse disclosure can be strictly preferred whenever adaptation generates spillovers.⁸ The appendix

⁸Risk neutrality is assumed here for simplicity (payoffs are expected monetary costs), but risk aversion would leave the mechanism intact. If a fire reaches a property with probability θ and the loss it then causes has a distribution that does not depend on θ , expected utility is affine in θ ; posteriors then enter behavior only through their mean, best responses remain cutoff rules in m , and the coarsening logic of Proposition 1 goes through. Concave utility instead relocates the detailed-disclosure cutoff m_D : it raises the utility stakes of the rare large loss relative to the certain cost c , operating like a uniform increase in ℓ and inducing




Figure 2: Beliefs and equilibrium adaptation under detailed disclosure and two-bin maps

Notes: Stylized example: $\theta \sim \text{Beta}(2, 5)$, $\ell(\theta) = 100\theta$, $m \mid \theta \sim N(\theta, 0.08^2)$, $\alpha = 0.8$, $c = 30$, and $\phi(A) = 1 - 0.8A$. Left plots: beliefs against the score m ; the dashed hurdle $c/(\alpha\phi(A))$, at the equilibrium rate, is the belief needed to adapt, and shading marks adapters (Panel B: under m^*). Right plots: colored curves redraw the left-plot beliefs against $A(x)$, the share of properties scoring at least x , and thus give the share clearing each hurdle. The dashed hurdle rises with A because non-adapters free-ride. Panel B: the light blue map (threshold m_D , open circle) keeps the adapters of detailed disclosure. The dark blue map lowers the threshold to m^* : arrows move Panel A beliefs (gray) to bin averages, and the high-bin belief falls onto the rising hurdle.

develops three extensions of this model. First, when losses scale up uniformly the optimal cutoff falls, and when the score sorts losses more sharply the cutoff falls as well, provided the high-bin constraint binds under the sharper score (Appendix E). Second, binary adaptation with homogeneous cost fits settings where the prescribed action is standardized, such as clearing defensible space, and both assumptions can be relaxed: pooling then shifts the adaptation of every pooled property rather than only the marginal ones, and when a ceiling on adaptation binds for the highest-risk properties, pooling the highest scores and reporting

adaptation at lower scores.

the rest in detail strictly beats detailed disclosure for any positive spillover (Appendices F and G). Third, wildfire spreads spatially (Rothermel, 1972; Finney, 2002; Shafran, 2008), so the population is best read at the scale over which fires spread: A is a community’s adaptation share, $\phi(A)$ the reduced form of a local spread process, and a statewide map applies the analysis community by community; the two forces behind Proposition 1 do not depend on this aggregation.

3 Type I and Type II errors

A coarse map classifies properties into those that should adapt and those that should not.⁹ I judge each classification against the *private optimum*: the action an owner who knew the property’s type θ would choose at the adaptation rate A the map induces. A map can induce too much adaptation relative to this benchmark (Type I errors) or too little (Type II errors). With spillovers, the optimal map sets its threshold below the detailed-disclosure cutoff (Proposition 1(ii)), so it flags parcels whose owners would not adapt if shown their own scores and accepts more Type I errors to cut Type II errors. Noise in the score adds errors of both types at any threshold. This section inverts the design problem of Section 2: given an *observed* coarse risk communication, what does the threshold choice reveal about the implied trade-off between Type I and Type II errors?

3.1 Empirical counterpart: realized fires

Judging classifications against the private optimum requires the type θ and the primitives (α, c, ϕ) , none of which is observed. Taking the model to data instead calls for one piece of structure beyond Section 2: a link between the latent type and an outcome I can observe. Interpret the type $\theta \in [0, 1]$ as the parcel’s latent probability that a wildfire reaches it (the landscape hazard), and let $Y \in \{0, 1\}$ denote its ex-post realization, a fire reaching the parcel, so that $Y \mid \theta \sim \text{Bernoulli}(\theta)$.¹⁰ The score is a noisy reading of θ and carries no further information about the realization ($Y \perp m \mid \theta$), so by the law of iterated expectations the score-conditional fire probability equals the expected type given the score:

$$\pi(m) \equiv \Pr(Y = 1 \mid m) = \mathbb{E}[\Pr(Y = 1 \mid \theta, m) \mid m] = \mathbb{E}[\theta \mid m]. \quad (2)$$

Section 4 estimates this score-conditional probability from realized fires and evaluates classification errors against Y .¹¹ Each score threshold x has a false-positive rate (FPR) and

⁹In the model, adaptation is voluntary. When the map instead triggers a mandate, so that properties in the high bin must adapt, the high-bin incentive constraint disappears. The low-bin constraint is slack at the first-best cutoff (Lemma 2), so a two-bin map with threshold m_{FB} implements the first best.

¹⁰Both θ and Y are defined relative to a fixed observation window: Y indicates at least one fire reaching the parcel within the window, and θ is the corresponding window-length burn probability. Every object built from them is therefore indexed by the window length; Section 4 places estimates from windows of different lengths on a common annualized basis to compare them.

¹¹Evaluating against Y rather than the type θ adds a third source of error beyond the threshold choice

a false-negative rate (FNR) against realized fires:

$$\text{FPR}(x) \equiv \Pr(m \geq x \mid Y = 0), \quad (3)$$

$$\text{FNR}(x) \equiv \Pr(m < x \mid Y = 1). \quad (4)$$

A false positive (Type I error) is a flagged parcel, one classified higher-risk, that did not burn; a false negative (Type II error) is an unflagged parcel that did. Raising the threshold flags fewer parcels, which lowers the false-positive rate and raises the false-negative rate. As x varies, the pair $(\text{FPR}(x), \text{FNR}(x))$ therefore traces a downward-sloping locus, which I refer to as the fire-based error frontier.

Lowering the threshold slightly from x flags the parcels with score x . Each burns with probability $\pi(x)$, so it removes a false negative with probability $\pi(x)$ and adds a false positive with probability $1 - \pi(x)$. FNR and FPR are rates within burned and unburned parcels, whose population shares are $\Pr(Y = 1)$ and $\Pr(Y = 0)$, so the slope of the frontier at x is

$$\frac{d\text{FNR}}{d\text{FPR}} = -\frac{\pi(x)}{1 - \pi(x)} \cdot \frac{\Pr(Y = 0)}{\Pr(Y = 1)}.$$

In absolute value, the slope is the odds that a parcel at the threshold burns, scaled by the ratio of unburned to burned parcels.¹² When the baseline loss $\ell(\theta)$ is proportional to θ , Assumption 1 makes π strictly increasing in the score. The odds then rise with the threshold, so the frontier is steeper at higher thresholds and therefore convex.

3.2 The recovered Type I/Type II trade-off

A planner who publishes a map with score threshold x selects the point $(\text{FPR}(x), \text{FNR}(x))$ on the fire-based error frontier. The point selected reflects how the planner weighs the two errors. Let $\lambda > 0$ denote the weight the planner places on a false positive relative to a false negative. Across parcels, the expected classification loss of a map with threshold x is then $\lambda \Pr(Y = 0) \text{FPR}(x) + \Pr(Y = 1) \text{FNR}(x)$, which is constant along lines of slope $-\lambda \Pr(Y = 0) / \Pr(Y = 1)$ in (FPR, FNR) space, the iso-loss lines. When π is strictly increasing, the planner lowers the threshold as long as the frontier is steeper than these lines and stops where an iso-loss line is tangent to the frontier. At the tangency the two slopes are equal, so the odds that a parcel at the threshold burns equal the weight, $\pi(x)/(1 - \pi(x)) = \lambda$. The planner therefore flags the parcels whose fire probability reaches the cutoff $\kappa(\lambda)$:

$$\pi(m) \geq \kappa(\lambda), \quad \kappa(\lambda) \equiv \frac{\lambda}{1 + \lambda}. \quad (5)$$

and score noise: *realization noise*, the gap between a parcel's burn probability and its realized 0/1 outcome, present even when $m = \theta$ because Y is a Bernoulli draw. Realization noise inflates the level of the error rates defined below, but does not distort the trade-off between them: the error counts are linear in Y , so across parcels they converge to expectations that depend on the primitives only through $\pi(m)$; see Proposition 2(i).

¹²By the law of iterated expectations, wherever the score density g and π are continuous, FPR and FNR have derivatives $-(1 - \pi(x))g(x)/\Pr(Y = 0)$ and $\pi(x)g(x)/\Pr(Y = 1)$, and the slope is their ratio when $g(x) > 0$ and $\pi(x) < 1$. On $(\underline{m}, \bar{m})$, Assumption 3 makes g continuous and positive, and makes π continuous when $\ell(\theta)$ is proportional to θ .

A larger λ raises the cutoff and the threshold, and shrinks the flagged “higher-risk” group.

Read in reverse, rule (5) assigns a weight to any map. With π strictly increasing, a map that flags the scores $m \geq \hat{m}$ flags exactly the parcels with $\pi(m) \geq \pi(\hat{m})$, so it is the map that a planner with cutoff $\hat{\kappa} \equiv \pi(\hat{m})$ and weight $\hat{\lambda} \equiv \hat{\kappa}/(1 - \hat{\kappa})$ would draw. Proposition 2 compares $\hat{\kappa}$ with the private optimum’s cutoff.

Proposition 2 (Structural content of the recovered $\hat{\lambda}$). *Let $\ell(\theta) = \theta \bar{L}$ with per-event loss $\bar{L} > 0$, and let Assumption 1 hold, so that π is strictly increasing. Let $\hat{\kappa} = \pi(\hat{m})$ and $\hat{\lambda} = \hat{\kappa}/(1 - \hat{\kappa})$ be the cutoff and weight revealed by the planner’s threshold $\hat{m}$, and let $\kappa^{\text{priv}}(A) \equiv c/(\alpha \phi(A) \bar{L})$ be the private cutoff.*

- (i) (Equivalence of fire-based and private objectives.) *The expected classification loss with weight λ and the per-parcel private cost with the aggregate adaptation rate held fixed at A differ only by a positive scale factor and an additive constant, and hence rank thresholds identically, if and only if $\kappa(\lambda) = \kappa^{\text{priv}}(A)$.*
- (ii) (No spillovers: welfare-exact.) *If ϕ is constant (Proposition 1(i)), the cutoff m_D under detailed disclosure, the planner’s optimal implementable threshold m^* , and the first-best cutoff m_{FB} coincide; if $\hat{m} = m^*$, $\hat{\lambda}$ equals the welfare cost of a false positive relative to a false negative.*
- (iii) (Spillovers: sign result.) *Suppose $\phi'(A) < 0$ for all $A \in [0, 1]$; Assumptions 2 and 3 hold, the former at the equilibrium induced by $\hat{m}$; and $\hat{m} \in (\underline{m}, \bar{m})$ is either the first-best cutoff m_{FB} or the planner’s optimal implementable threshold m^* of Section 2. Let $A = A(\hat{m})$ be the induced adaptation rate, $L = \mathbb{E}[\ell(\theta)]$, and $L_{\text{hi}}(\hat{m}) = \mathbb{E}[\ell(\theta)1\{m \geq \hat{m}\}]$ the expected loss carried by the above-cutoff (“high”) bin. Then*

$$\underbrace{\frac{c}{\alpha \phi(A) \bar{L}}}_{\kappa^{\text{priv}} \text{ (private cutoff)}} + \underbrace{\frac{\phi'(A) (L - \alpha L_{\text{hi}}(\hat{m}))}{\alpha \phi(A) \bar{L}}}_{w(\hat{m}) \text{ (spillover wedge)} < 0} \leq \hat{\kappa} < \kappa^{\text{priv}}, \quad (6)$$

with equality on the left at $\hat{m} = m_{FB}$ and whenever the high-bin incentive constraint is slack at $\hat{m} = m^$. Hence $\hat{\lambda} < \kappa^{\text{priv}}/(1 - \kappa^{\text{priv}})$.*

Appendix D proves the proposition. A false positive wastes the adaptation cost; a false negative forgoes the loss that adaptation would have avoided, net of that cost. Weighting the two by their ratio, the private-optimum weight $\kappa^{\text{priv}}/(1 - \kappa^{\text{priv}})$, reproduces the owner’s calculation (part (i)), which is also the planner’s when there are no spillovers (part (ii)). With spillovers, an adapter lowers residual risk for everyone else, whether or not the parcel burns, so the planner weights a false positive less than the owner does (part (iii)). At $\hat{m} = m^*$, the upper bound in part (iii) is the inequality $m^* < m_D$ of Proposition 1(ii), expressed in fire-probability units.

4 An Application to Oregon’s 2022 Wildfire Risk Map

Sections 2 and 3 established three results that the empirical work below builds on: a planner facing adaptation spillovers optimally publishes a binary lower-risk / higher-risk map defined by a threshold on a continuous risk score; that threshold is a classification rule whose errors, measured against realized fires, weigh in expectation the same trade-off as errors against the private optimum when losses are linear in the risk type (Proposition 2); and the chosen threshold reveals the implicit weight $\lambda > 0$ on a false positive relative to a false negative, which internalizing spillovers pushes below the private-optimum weight.

This section applies that framework to Oregon’s 2022 Wildfire Risk Map: I describe the map and the data, validate the three risk models against held-out fires, recover the implicit trade-off embedded in the published cutoff, and draw counterfactual maps at alternative cutoffs.

4.1 Setup: Oregon’s signal, cutoff, and the model’s primitives

The Oregon Department of Forestry developed the 2022 Wildfire Risk Map in fulfillment of Senate Bill 762 (SB 762), in collaboration with researchers at Oregon State University who carried out the underlying modeling (Oregon Legislative Assembly, 2021). The map covers every parcel in the state and is constructed in three layers (Dunn and McEvoy, 2022). First, a stochastic fire-spread model (FSim) simulates plausible fire seasons from local weather records, landscape conditions, and historical patterns of fire occurrence and size, and reports a parcel-level burn probability (Finney, 2002; Scott, Short and Finney, 2018). Second, a fire-behavior analysis model (FlamMap) estimates flame length and fireline intensity under user-specified weather and fuel-moisture conditions (Finney, 2006). Third, the burn-probability and flame-length layers are combined into a single continuous *risk score*, which is coarsened into five ordinal categories before public release: “No risk”, “Low”, “Moderate”, “High”, and “Extreme”.¹³ The category thresholds were fixed in administrative rule adopted by the Oregon Board of Forestry, the policy body that SB 762 charged with establishing the map’s criteria (Oregon Board of Forestry, 2022).

I take the Oregon model’s composite continuous score as the planner’s signal m in the model of Section 2: an observable, monotone summary of latent risk θ . The published 5-category map partitions m at four thresholds; the policy-relevant one is the Moderate/High boundary, since SB 762 implementation discussions contemplated the “High” and “Extreme” classifications as triggers for both defensible-space requirements and more stringent wildfire-hardening building-code requirements (Oregon Building Codes Division, 2022; Oregon Wildfire Programs Advisory Council, 2022).¹⁴ This Moderate/High boundary on m is the thresh-

¹³This format is common across countries and perils: Appendix H catalogs the public flood, wildfire, and hurricane risk information available in the G20 countries and classifies each source by disclosure form; most of it is released as mapped zones or discrete categories.

¹⁴The Office of the State Fire Marshal was “responsible for developing a defensible space code applied to high and extreme risk classes inside the wildland-urban interface state-wide,” and the Department of Consumer and Business Services for “adopting fire hardening building code standards for new development

old $\hat{m}$ of Section 3.2: parcels classified High or Extreme form the map’s higher-risk group, and all other parcels its lower-risk group. I read $\hat{m}$ as the empirical counterpart of the planner’s optimal cutoff m^* from Section 2: in the spirit of revealed preference, I take $\hat{m}$ as given and recover the implicit error weight $\hat{\lambda}$ at which it would be the planner’s optimum; Section 4.5 clarifies what this recovery assumes.

Panel A of Figure 3 repeats the published 5-category map of Figure 1. Panels B and C show the continuous layers that underlie the risk model (FSim burn probability, and FlamMap flame length), while Panel D shows the composite continuous risk score m , which the planner observes but coarsens into Panel A. The empirical exercise that follows asks (i) how informative m is about future wildfire occurrence, (ii) what classification cutoff $\hat{m}$ is implicit in moving from Panel D to Panel A, and (iii) what relative weight $\hat{\lambda}$ on Type I versus Type II errors this cutoff reveals.

4.2 Data: parcel-level risk signals and realized fires

The merged sample combines four parcel-level data sources at the property level: (i) the Oregon map (both the underlying continuous risk score m , and its 5-category policy coarsening); (ii) Cotality’s wildfire Average Annual Loss (AAL) s_C , a proprietary catastrophe-model output (CoreLogic, Inc., 2025); (iii) parcel-level wildfire metrics from First Street Foundation (FSF), of which I use the burn probability s_F as the canonical signal (it more closely targets realized burning, the outcome of interest, and includes fewer missing values than the other variables);¹⁵ and (iv) realized fire occurrence 2000–2025 from the National Interagency Fire Center (NIFC): the Interagency Fire Perimeter History for 2000–2019 and the Interagency Fire Perimeters compiled by the Wildland Fire Interagency Geospatial Services group for 2020–2025.

The three models all estimate wildfire risk, but they measure slightly different objects. The Oregon model is structurally closest to the First Street Foundation model: both predict physical wildfire outcomes, and both build on Rothermel’s fire-behavior equations. Cotality’s AAL instead combines physical fire modeling with property-level characteristics and expresses the result as the expected annual share of total insurable value lost to wildfire. Property survival conditional on wildfire exposure is not negligible and rises with fire-resistant construction and defensible space (Syphard, Brennan and Keeley, 2014; Baylis and Boomhower, 2026), and this conditional destruction probability is presumably where Cotality’s property-level characteristics enter. The comparisons below evaluate all three scores against the same outcome: whether a fire reaches the parcel. The common per-event loss of Proposition 2 requires that this conditional destruction probability not vary systematically with the score. Where it does, the calibration of Section 4.7 still bounds the per-event loss

and significant home updates in high and extreme risk areas in the wildland-urban interface” (Oregon Wildfire Programs Advisory Council, 2022). Read as triggers for these requirements, the High and Extreme classifications make adaptation mandatory above $\hat{m}$, which removes the high-bin incentive constraint of Section 2. These building code standards were never imposed.

¹⁵The Cotality scores are the model output as of December 30, 2021. The First Street scores come from version 3.0 of its wildfire model, published on July 31, 2024 (First Street, 2024).




Figure 3: Oregon 2022 Wildfire Risk Map and its underlying continuous layers.

Notes: Panel A reproduces the 5-category map released in June 2022. Panels B and C show FSim burn probability and FlamMap flame length. Panel D shows the composite continuous risk score m that the Oregon model constructs from B and C and then coarsens into Panel A. Panels B–D are property-level outputs of the Oregon model that are not communicated to property owners; this paper observes all three.

of the marginal parcel from above. The modeling underlying Cotality’s measure is proprietary, so I cannot assess how its assumptions relate to those of the other two models (a transparency concern discussed in Condon (2023)).

To place the three risk signals on a common scale, I compare them by their within-sample ranks. After nearest-neighbor matching of the Cotality and FSF properties to Oregon’s parcels and dropping parcels without properties or with missing model data, the working sample contains 1,860,526 observations (nearly every residential parcel in the state). Appendix Figures S4, S5, and S6 compare the distribution and the spatial pattern of the three risk signals.

Table 1 summarizes the three signals and the realized-fire outcomes. Throughout, a parcel is recorded as burned when it falls within a NIFC fire perimeter. Perimeters are operational containment polygons and can enclose unburned islands, so this outcome measures whether a fire reached the parcel rather than whether the structure burned; Appendix J repeats

	N	p50	p90	p99
<i>Panel A. Parcel-level risk signals (native units)</i>				
Oregon score m (score)	1,860,526	0.0003	0.1397	1.2359
Cotality AAL s_C (% of value)	1,860,526	0.0003	0.0383	0.2946
First Street s_F (% burn prob.)	1,860,526	0.01	0.77	2.18
<i>Panel B. Sample and realized fires</i>				
	Count	Rate (%)		
Working sample (parcels)	1,860,526	—		
burned 2000–2021 (pre-map)	20,386	1.10		
Clean out-of-sample (no 2000–2021 fire)	1,840,140	—		
burned 2022–2025 (held-out)	3,673	0.20		

Table 1: Summary statistics: risk signals and realized fires in the Oregon working sample.

Notes: Working sample of 1,860,526 Oregon parcels, after nearest-neighbor matching of the Cotality and First Street properties and dropping parcels with missing data. Panel A reports percentiles of each risk signal in its native units: the Oregon composite score m in score units, Cotality’s average annual loss s_C as a percentage of insurable value, and First Street’s burn probability s_F as a percentage. All comparisons in Section 4.3 onward instead place the three signals on a common within-sample rank scale, since their native units are not comparable. The three signals are strongly right-skewed, with near-zero medians and mass concentrated in the upper tail. Panel B reports sample sizes and realized-fire counts against NIFC interagency fire perimeters: pre-map fires over 2000–2021 (the planner’s information set) and held-out fires over 2022–2025, the latter on the “clean” out-of-sample subset that drops parcels with any 2000–2021 fire. Rates are the share of the relevant sample.

the calibration with satellite-mapped burned area and obtains very similar results. The pre-period (2000–2021) is the planner’s information set when the 2022 map was drawn; I use it to calibrate the score-conditional burn probability $\hat{\pi}_{\text{pre}}(m)$ that recovers the implicit trade-off $\hat{\lambda}$. Over this window 20,386 parcels (1.10% of the sample) fall within a recorded fire perimeter. On the other hand, the post-period (2022–2025), which Oregon could not have used in constructing the map, is a held-out window that I use to validate that signal against realized fires (Sections 4.3–4.4). The post-period contains 3,673 parcels within a perimeter in the “clean” out-of-sample (OOS) subset (parcels with no recorded fire in 2000–2021, total $N = 1,840,140$).

4.3 Validation of risk models against realized fires

Table 2 reports the area under the receiver operating characteristic curve (AUC) for four logit models, each trained on 2000–2021 NIFC perimeter occurrence and evaluated on the disjoint 2022–2025 window restricted to parcels not burned in the training window. The AUC is the probability that a randomly drawn parcel that burned receives a higher predicted risk than a randomly drawn parcel that did not; it equals 0.5 for a random ranking and 1 for a perfect one. The Oregon model achieves AUC 0.923 on the held-out window, two points below the Cotality model (0.940) and above the FSF model’s burn probability model (0.848).

Model	AUC (train, 2000–2021)	AUC (test, 2022–2025)
Oregon m alone	0.823	0.923
Cotality s_C alone	0.809	0.940
First Street burn prob.	0.721	0.848
Joint (m, s_C, s_F)	0.833	0.943

Table 2: Held-out predictive performance: train on 2000–2021, test on 2022–2025 NIFC fires.

Notes: Each model is a logit of 2000–2021 NIFC burn on rank-transformed signals, fit on all parcels. AUC is computed on 2022–2025 NIFC outcomes for parcels with no 2000–2021 NIFC burn. The public Oregon model lands within two AUC points of the proprietary Cotality model; the proprietary FSF burn probability is the weakest of the three. Appendix J repeats the calibration of the Oregon score with a satellite record of burning to assess robustness.

A joint model that pools all three signals achieves AUC 0.943, a marginal improvement of two percentage points over m alone. Overall, the free, public Oregon risk signal performs in the same neighborhood as leading proprietary commercial alternative.

The model in Section 2 requires that $m \mapsto \mathbb{E}[\ell(\theta) \mid m]$ be strictly increasing (Assumption 1), so that the individual’s optimal action is a cutoff rule in m . Panel A of Figure 4 plots held-out burn rates across the percentiles of each signal, zooming into the 75th–99th region where the Moderate/High map cutoff is located (Appendix Figure S7 shows a similar figure across the full distribution). For all three signals, the burn rate rises with the signal through the cutoff region. Wildfires are rare and spatially clustered, so bin-level burn rates are noisy and show local reversals. A bootstrap test that resamples whole fires does not reject the null hypothesis that the held-out burn rate is weakly increasing in each signal ($p = 0.16, 0.17$, and 0.73 for m , s_C , and s_F ; Appendix Figure S7).

Interestingly, while burn probability rises with all three risk models, the signals differ in where they separate burned from unburned parcels. Cotality dominates at the extreme top: above roughly the 97th percentile its bin-mean burn rates climb steeply past the other two signals, reaching 3.3% in its top bin. Below the 95th percentile, though, its rank-conditional burn rate is nearly flat: parcels in Cotality’s 75th–95th band burned at 0.24%, against 0.69% for the corresponding Oregon band. At the Moderate/High cutoff, which sits at the 90th percentile, parcels burn at a higher rate under Oregon’s ranking than at the same ranks of either proprietary alternative. Because the First Street vintage postdates the start of the held-out window (Section 4.2), part of the 2022–2025 fire record was available to that model when it was built. Appendix Figure S8 repeats Panel A with the held-out window starting in 2023 and in 2024, the latter after the First Street release. The comparison is mostly unchanged: through the 75th–95th percentile band the Oregon and First Street burn rates track each other while Cotality’s stays flat, and just above the cutoff the Oregon signal’s parcels burn at the highest rate of the three.

I then estimate the score-conditional fire probability on the held-out window, $\pi_{\text{OOS}}(m) = \Pr(\text{2022–2025 fire} \mid m)$, by calibrating the Oregon score to held-out realized fires. I use a shape-constrained monotone-increasing P-spline logit fit on $\sqrt{m}$, which spreads the strongly




Figure 4: Calibration against held-out 2022–2025 fires: all three signals in the policy-relevant region (Panel A) and the Oregon calibration $\hat{\pi}_{OOS}(m)$ (Panel B).

Notes: Held-out 2022–2025 NIFC burn rates. The shaded band in both panels is the 95% fire-cluster bootstrap confidence band for the fitted calibration $\hat{\pi}_{OOS}(m)$, refitting the monotone spline in each replicate with its smoothing parameter held at the full-sample value ($B = 1,000$; cluster = 2022–2025 NIFC fire perimeter for burned parcels, 5km grid cell otherwise); every replicate curve is monotone, so the band contains only weakly increasing calibrations. Bin-mean burn rates are shown as points without intervals: fires arrive as contiguous events, so a single bin’s four-year burn rate rests on a handful of fire events and its fire-clustered interval spans much of the plot. Panel A: 50 fine bins per signal in the 75th–99th percentile band of each signal’s rank distribution; the black curve is the $\hat{\pi}_{OOS}(m)$ spline of Panel B mapped to rank space. Panel B: centile-mean burn rates (blue points, size proportional to bin count) and the monotone-increasing P-spline logit on $\sqrt{m}$ (black); $\hat{\pi}_{OOS}(m)$ in black is the estimate of $\Pr(2022\text{--}2025 \text{ fire} \mid m)$ used in Sections 4.4–4.6; the horizontal axis is on a $\sqrt{m}$ scale (m is strongly right-skewed). Dashed lines mark Oregon’s Moderate/High cutoff $\hat{m}$ (in rank space in Panel A, in m -space in Panel B). Cotality’s calibration advantage is confined to the extreme upper tail (above roughly the 97th percentile).

right-skewed score. Panel B of Figure 4 plots the resulting estimate $\hat{\pi}_{OOS}(m)$ over the non-parametric centile mean burn rates on a $\sqrt{m}$ axis, and the black curve in Panel A maps the resulting spline into rank space over the cutoff region.¹⁶ Appendix K shows the main results are robust to using alternative spline specifications, as well as a step function.

4.4 The empirical Type I/Type II frontier

Each candidate cutoff x on m generates the pair of error rates $(FPR(x), FNR(x))$ of equations (3)–(4), computed against realized fires on the held-out sample. Sweeping x across the

¹⁶The estimation of the implied weight $\hat{\lambda}$ in Section 4.5 instead uses the identically constructed pre-map calibration $\hat{\pi}_{pre}(m)$, the estimate of $\Pr(2000\text{--}2021 \text{ fire} \mid m)$ that captures the planner’s information set; the two are compared directly there.




Figure 5: Empirical (FPR, FNR) frontiers across the three wildfire models

Notes: Each frontier traces (FPR, FNR) by varying cutoffs on the respective signal against held-out 2022–25 NIFC fire outcomes (lower-left is better). The red dot is Oregon’s map actual cutoff $\hat{m}$ on the Oregon m frontier. The gray line is the iso-loss line tangent to this realized-fire (2022–2025) frontier at Oregon’s point, with slope $= -\hat{\lambda}_{\text{OOS}} \Pr(Y = 0) / \Pr(Y = 1)$; the rare-event base rate over a four-year window makes $\hat{\lambda}_{\text{OOS}}$ small in absolute terms.

support of m traces out the fire-based error frontier of Section 3.1. Figure 5 plots this frontier for the three risk models, and marks Oregon’s actual cutoff $\hat{m} = 0.1379$, which sits at the Moderate/High boundary in m -space. At $\hat{m}$, Oregon’s policy generates a false-positive rate of 0.0973 and a false-negative rate of 0.2352 against realized 2022–2025 fires: roughly 10% of unburned parcels are classified High or Extreme, and roughly 24% of parcels that burned during the window are missed. Oregon sits at an interior point of the frontier, where a small improvement in the false-positive rate would cost a substantial increase in the false-negative rate.¹⁷ For cutoffs stricter than Oregon’s, to the upper left of the red dot, the Oregon frontier

¹⁷These error rates are measured over the four-year held-out window: a false positive is a parcel above the cutoff that did not burn in 2022–2025, and it may still burn after the window closes.

bends the wrong way, because the parcels with the very highest scores burned less often than parcels closer to $\hat{m}$ (one of the local reversals noted in Section 4.3). This stretch lies entirely above $\hat{m}$ and does not affect the trade-off at Oregon’s cutoff.

In Figure 5, the Oregon and Cotality frontiers cross at a cutoff slightly more permissive than Oregon’s. At Oregon’s own false-positive rate and for stricter rules, the Cotality frontier lies below, which reflects its advantage in the extreme upper tail of its score (Figure 4); for more permissive rules, the Oregon frontier offers lower error rates. The First Street frontier lies above the other two at nearly every cutoff.

4.5 Oregon’s coarsening and its implicit trade-off

By choosing a threshold $\hat{m}$ on its model’s risk score to draw the 2022 Wildfire Risk Map, the Oregon planner fixed a boundary that can be rationalized by an implicit weight λ on false positives relative to false negatives.¹⁸ In the spirit of revealed preference, I model the planner as if it optimizes the classification objective of Section 3.2 using the empirical historical burn probabilities, and recover λ from that rule:

$$\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m}), \quad \hat{\lambda} = \frac{\hat{\kappa}}{1 - \hat{\kappa}}. \quad (7)$$

Here $\hat{\pi}_{\text{pre}}$ is the empirical calibration of the pre-map burn risk, the score-conditional probability $\pi_{\text{pre}}(m) = \Pr(\text{2000–2021 fire} \mid m)$. I estimate it as in Section 4.3 by calibrating the score to realized fires, but on the pre-map (2000–2021) window alone; Appendix Figure S10 plots the resulting $\hat{\pi}_{\text{pre}}$ against the pre-map burn rates, alongside the analogous calibrations of the two proprietary models. Reading $\hat{\lambda}$ off $\hat{\pi}_{\text{pre}}$ treats Oregon’s cutoff on the score as a cutoff on the pre-map burn probability.¹⁹ Appendix K reports additional specification checks, and Appendix J repeats the recovery with the satellite record of burning from the Monitoring Trends in Burn Severity (MTBS) program in place of the NIFC perimeters.

At Oregon’s cutoff the recovered trade-off is $\hat{\lambda} = 0.039$, with a 95% cluster-bootstrap confidence interval (CI) of $[0.026, 0.059]$ (first column of Table 3). It is far below one: per false negative, Oregon’s cutoff tolerates roughly $1/\hat{\lambda} \approx 25$ parcels classified High that did not burn in order to catch the ones that do.

Both the estimate and its units depend on the observation window: a “false positive” here is a parcel that did not burn within the 22-year pre-map window, and the label is provisional

¹⁸Isakov, Lo and Montazerhodjat (2019) derive optimal significance levels for Food and Drug Administration approval from disease-specific error costs and impute the costs implicit in the conventional 2.5% level; Tetenov (2012) shows that, under normality, a test at a given level is the minimax-regret rule for a specific weight on Type I relative to Type II regret. Both infer the weight from a test’s design parameters, whereas I infer it from realized outcomes.

¹⁹The cutoff-in- m /cutoff-in- π equivalence behind equation (7) requires weak monotonicity of the true calibration π_{pre} , which the shape-constrained estimator imposes. The binned pre-map burn rates move around the fitted curve because rare fires arrive as spatially contiguous events: four fires (Beachie Creek, Holiday Farm and Alameda Drive in September 2020, Bootleg in 2021) account for four-fifths of the burned parcels between the 68th and 78th rank percentiles. A bootstrap that resamples fire-level clusters fails to reject weak monotonicity of π_{pre} ($p = 0.23$).

	Pre-map (headline) $\hat{\pi}_{\text{pre}}, 2000\text{--}2021$	Held-out (check) $\hat{\pi}_{\text{OOS}}, 2022\text{--}2025$
Burn probability at $\hat{m}$	0.0375	0.0120
Implied weight λ	0.0389	0.0121
95% CI	[0.0257, 0.0590]	[0.0064, 0.0188]
Implied weight, annualized (1-yr basis)	0.00174	0.00302
95% CI	[0.00116, 0.00261]	[0.00160, 0.00467]
Annualized difference, held-out – pre-map	0.00128	
95% CI	[–0.00040, 0.00309]	
Window length (years)	22	4
Sample	All parcels	No 2000–2021 fire
Parcels	1,860,526	1,840,140
Bootstrap clusters	8,827	8,187

Table 3: The weight $\hat{\lambda}$ implied by Oregon’s cutoff $\hat{m}$: pre-map recovery and held-out check.

Notes: Both columns evaluate a monotone P-spline calibration of the Oregon score at the Moderate/High cutoff $\hat{m} = 0.1379$. The pre-map column fits $\hat{\pi}_{\text{pre}}$ to 2000–2021 fires on the full sample and reports $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$, $\hat{\lambda} = \hat{\kappa}/(1 - \hat{\kappa})$, and $\hat{\lambda}_{\text{pre}}^{\text{ann}}$. The held-out column fits $\hat{\pi}_{\text{OOS}}$ to 2022–2025 fires on the clean sample of parcels with no 2000–2021 fire and reports $\hat{\pi}_{\text{OOS}}(\hat{m})$, $\hat{\lambda}_{\text{OOS}}$, and $\hat{\lambda}_{\text{OOS}}^{\text{ann}}$. Confidence intervals are cluster bootstraps with $B = 1,000$, refitting the calibration spline in each replicate with its smoothing parameter held at the full-sample value; the pre-map column resamples pre-2022 fire perimeters and 5km-grid cells, the held-out column 2022–2025 fire perimeters and 5km-grid cells. The annualized rows convert each bootstrap draw to a one-year burn probability under a constant annual hazard over the window length shown. The difference row pairs the two sets of draws by replicate as independent draws, since the two bootstraps resample different clusters.

because an unburned parcel can still burn after the window closes while a burned one cannot revert. Under an approximately constant annual hazard h , the window- T burn probability is $1 - (1 - h)^T \approx hT$ at these rare-event scales, so $\hat{\lambda}$ scales with the length of the window. Placing the 22-year pre-map and 4-year post-map recoveries on a common annualized basis makes them comparable: $\hat{\lambda}_{\text{pre}}^{\text{ann}} = 0.00174$ versus $\hat{\lambda}_{\text{OOS}}^{\text{ann}} = 0.00302$, the latter equivalent to $\hat{\lambda}_{\text{OOS}} = 0.012$ read directly off the held-out 2022–2025 frontier of Section 4.4 (second column of Table 3). The two are not statistically different from each other at the 5% level: the annualized difference is 0.00128, with 95% confidence interval $[-0.00040, 0.00309]$.

The interpretation of these estimates depends on the assumptions one is willing to make, as clarified in Proposition 2. Without further assumptions $\hat{\lambda}$ is *descriptive*: it is the rate at which Oregon’s cutoff sits on the realized-fire frontier, with (7) the estimator implied by the cutoff rule (5). In the absence of spillovers, $\hat{\lambda}$ is *welfare-exact*. Under spillovers, once Oregon’s threshold is read as the planner’s optimal implementable threshold m^* , theory delivers only a *sign*: $\hat{\kappa}$ lies strictly below the private cutoff κ^{priv} , by at most the spillover wedge of (6), so $\hat{\lambda}$ lies strictly below the private-optimum weight $\kappa^{\text{priv}}/(1 - \kappa^{\text{priv}})$. In none of the three do I recover the primitives (α, c, ϕ) . The direction $\hat{\lambda} \ll 1$ is therefore consistent

with, but does not identify, the prediction of Proposition 1(ii) that a planner internalizing positive spillovers sets the disclosure threshold permissively to induce additional adaptation.

4.6 Counterfactual maps at alternative cutoffs

Holding the calibration $\hat{\pi}_{\text{OOS}}(m)$ fixed, varying the cutoff κ traces out a family of counterfactual classification rules $1\{\hat{\pi}_{\text{OOS}}(m_i) \geq \kappa\}$; each κ corresponds to a point on the empirical frontier of Section 4.4 and to an implicit weight $\lambda = \kappa/(1 - \kappa)$. Figure 6 renders five such maps (Panels A–E), ordered from stringent to permissive FPR, and Panel F locates each one on the held-out (FPR, FNR) frontier. Because per-parcel burn probabilities are small under a realized-fire window, I anchor the panels on policy-meaningful target shares of parcels classified High, $\{1\%, 4\%, 10\%, 25\%, 50\%\}$, and I then back out the implied (κ, λ) pairs from the calibration distribution. The most stringent rule flags 1% of the state’s parcels but misses 96.1% of the parcels that burned in 2022–2025 (Panel A), while the most permissive flags half of all parcels and misses only 0.1% (Panel E). Panel C corresponds to Oregon’s empirical share flagged high of 10.08%.

Oregon’s geography shows through in all panels: the wet Pacific coast and Coast Range, the Willamette Valley, and the Portland metropolitan area remain blue, while “High” classification concentrates in the fire-prone Cascades, the Klamath-Siskiyou region of the southwest, and the drier interior of central and eastern Oregon. Moving along the frontier shifts the classification boundary in transition zones: more permissive cutoffs pull additional foothill and wildland-urban-interface parcels into the High zone, while more stringent cutoffs retract the boundary to the highest-risk cores.

4.7 Calibrating the private benchmark

Under spillovers, and with Oregon’s threshold read as the planner’s optimum, the bound (6) places the map cutoff strictly below the private cutoff $\kappa^{\text{priv}} = c/(\alpha \phi(A) \bar{L})$, by at most the spillover wedge, but the map alone identifies neither the primitives (α, c, ϕ) nor the size of the gap. This section calibrates κ^{priv} from external sources and asks whether even its lowest plausible value could rationalize Oregon’s cutoff choice *without* spillovers. Because ϕ is decreasing with $\phi(0) = 1$, evaluating κ^{priv} at $\phi = 1$ minimizes it over every adaptation rate: if Oregon’s cutoff lies below the calibrated range at $\phi = 1$, then it lies below $\kappa^{\text{priv}}(A)$ for any A , the direction the bound predicts under spillovers and rules out without them.

The comparison requires a common time horizon. The recovered cutoff $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$ is a 22-year window burn probability, while adaptation combines upfront investments with recurring maintenance, so both sides are converted to annual flows. On the cutoff side, the annualized cutoff $\hat{\kappa}^{\text{ann}} = 1 - (1 - \hat{\kappa})^{1/22}$ is the constant annual hazard that reproduces $\hat{\kappa}$ over the 22-year window, as in Section 4.5. It equals 1.74×10^{-3} , with 95% interval $[1.16, 2.61] \times 10^{-3}$. The cost side converts each upfront component into an annuity at real rate r over the component’s useful life and adds annual maintenance. The resulting annual cost c^{ann} gives the private benchmark $\kappa^{\text{priv,ann}} = c^{\text{ann}}/(\alpha \bar{L})$, the annual counterpart of κ^{priv} at $\phi = 1$.




Figure 6: Counterfactual wildfire risk maps of Oregon at alternative cutoffs on the frontier.

Notes: Panels A–E apply $1\{\hat{\pi}_{\text{OOS}}(m_i) \geq \kappa\}$ to each of the 1,860,526 residential parcels in the working sample and aggregate the share classified High to a 2km grid for rendering. Grid cells containing fewer than five residential parcels are left blank, so white areas indicate sparse or absent housing (national forest, wilderness, rangeland), not low risk; pale blue marks populated cells in which no parcel is classified High, and colors run through yellow and orange to dark red as the cell share rises to one. Panels are ordered from stringent (A, 1% of parcels flagged High) to permissive (E, half of all parcels); each panel title reports the share of held-out 2022–2025 burned parcels its rule misses, computed on the clean out-of-sample subset of Section 4.2. Panel C (outlined) is Oregon-equivalent: its 10% target matches the published map’s share-High of 10.08%. Panel F places each cutoff on the held-out (FPR, FNR) frontier of Figure 5; the red point marks the Oregon-equivalent cutoff, and blue points mark the four counterfactual alternatives, with letters colored to match the map panels. The implied probability cutoffs are A: $\kappa = 1.5 \times 10^{-2}$; B: $\kappa = 1.5 \times 10^{-2}$; C: $\kappa = 1.2 \times 10^{-2}$; D: $\kappa = 1.5 \times 10^{-4}$; E: $\kappa = 1.4 \times 10^{-5}$; the Oregon-equivalent value (C) coincides with the ground-truth check $\hat{\lambda}_{\text{OOS}}$ of Table 3.

The primitives cover two nested adaptation bundles considered under SB 762: defensible space alone, and defensible space plus ember-resistant hardening. Costs come from state assistance programs and abatement benchmarks (California Legislative Analyst’s Office, 2021) and from component-level retrofit estimates for a 2,000 square-foot home (Barrett and Quarles, 2024), all expressed in 2024 dollars. The effectiveness α depends on the adaptation measures: for defensible space alone, destruction relative risks among exposed structures imply




Figure 7: Calibrated private benchmarks $\kappa^{\text{priv,ann}}$ against Oregon’s annualized cutoff $\hat{\kappa}^{\text{ann}}$.

Notes: Each point is the private benchmark $\kappa^{\text{priv,ann}} = c^{\text{ann}}/(\alpha\bar{L})$ implied by one of the 162 parameter combinations of Appendix Table S6. It is the annual counterpart of the private cutoff κ^{priv} , evaluated at $\phi = 1$ (the zero-adaptation residual risk $\phi(0) = 1$, which minimizes κ^{priv}). Rows separate the two adaptation bundles, and diamonds mark the all-central combination. The vertical line is Oregon’s annualized cutoff $\hat{\kappa}^{\text{ann}} = 1.74 \times 10^{-3}$, with its 95% cluster-bootstrap interval shaded. The horizontal axis is logarithmic. Effectiveness values differ by bundle, so the combined bundle’s effectiveness is never paired with the costs of defensible space alone.

reductions of 16–31% (Syphard, Brennan and Keeley, 2014), and for the combined bundle, comparisons across construction eras in the Camp Fire and simulated counterfactuals imply reductions of roughly 31–52% (Knapp et al., 2021; Zamanialaei et al., 2025). The calibration widens both ranges, placing its highest effectiveness values above the largest of these estimates; since $\kappa^{\text{priv,ann}}$ falls as α rises, the widening works against finding $\hat{\kappa}^{\text{ann}}$ below the calibrated range. The per-event loss $\bar{L}$ ranges from the structure-only value of the median Oregon home to its full property value, an upper bound on the loss an arriving fire can cause, and the real rate r spans 2–5% (Office of Management and Budget, 2023). Appendix Table S6 consolidates the low, central, and high values with their sources. Crossing the three levels of cost, effectiveness, per-event loss, and discount rate within each bundle yields a parameter grid of 162 combinations; the corner of the grid that combines the cheapest costs, the highest effectiveness, the full property value, and the lowest rate minimizes $\kappa^{\text{priv,ann}}$.

Figure 7 plots the private benchmark $\kappa^{\text{priv,ann}}$ implied by each combination, one row per bundle, against Oregon’s annualized cutoff $\hat{\kappa}^{\text{ann}}$, drawn as a vertical line with its 95% interval shaded. When the bundle includes hardening, every combination lies to the right of the line, and all but the bundle’s minimizing corner lie to the right of the shaded interval. At central parameter values, marked by a diamond, the benchmark is 4.0 times the cutoff. Because the benchmark is evaluated at $\phi = 1$, Oregon’s cutoff then lies below the private cutoff at any adaptation rate, the direction that the bound (6) predicts under spillovers. In dollars, rationalizing the cutoff as privately optimal without spillovers requires the bundle to cost at most $\hat{\kappa}^{\text{ann}} \alpha \bar{L}$ per year. At the bundle’s highest effectiveness and the full property value, this

ceiling is \$496 per year, while defensible space and hardening together cost at least \$684 per year. Defensible space alone costs less. Its row contains the lowest benchmarks, and 6 of its combinations lie to the left of the line. Across the two bundles, the benchmark exceeds the cutoff in 96% of the combinations. The cutoff Oregon drew is therefore consistent with the choice of a planner pricing more than the flagged owners' private returns.

5 Conclusion

Governments communicate risk through coarse categories, and the usual justification is cognitive. This paper shows that coarseness can be optimal even when individuals are fully rational. When adaptation protects others as well as the adapter, a coarse but truthful map can strictly dominate detailed disclosure, because widening the higher-risk category draws into adaptation individuals whose own risk alone would not justify the cost. Any coarse map is a threshold on a risk score, and that threshold implies a trade-off between two errors: the map induces too much adaptation for some individuals and too little for others. Given the published threshold and the outcomes realized afterwards, the weight the planner placed on this trade-off can be recovered.

I apply the framework to Oregon's 2022 Wildfire Risk Map, whose release triggered intense controversy over the model's accuracy. Benchmarked against realized fires, Oregon's free public model predicts which parcels burned in 2022 to 2025 about as well as two leading proprietary catastrophe models. Settling whether the model is accurate, the question that dominated the public debate, still leaves open the choice the map itself reveals: where to draw the line between lower and higher risk categories. Over the planner's two-decade information window, the map accepts roughly twenty-five parcels flagged as high risk that do not burn to avoid missing one that does. This permissive direction is consistent with, though not unique to, a planner internalizing adaptation spillovers. Counterfactual maps drawn at alternative cutoffs show that the published map is one point on a frontier of feasible policy options, each with its own balance of false positives and false negatives.

The argument extends beyond wildfires to protective actions that pay off more for those at higher risk and also protect others. Hazards that spread between units have this structure. Vaccination protects the vaccinated person and slows transmission to others; treating one's own land against a crop pest or invasive species reduces own damages and slows spread to neighboring land; and securing one's own computer systems lowers the risk of a breach and removes a point from which attacks can spread across a network. Where coarse risk categories rest on a continuous risk score and outcomes are later observed, the exercise applied here to Oregon recovers the implicit weight on false positives that rationalizes the chosen threshold.

The welfare comparisons in this paper use each individual's true risk and realized adaptation, so beliefs matter only instrumentally. This abstracts from a literature in which beliefs enter utility directly (Loewenstein, 1987; Caplin and Leahy, 2001; Brunnermeier and Parker, 2005; Bénabou and Tirole, 2016), and from the question of whether welfare should be assessed under individuals' beliefs or objective probabilities (Bernheim and Rangel, 2009).

References

- Aguilar-Gomez, Sandra, Joshua S Graff Zivin, and Matthew J Neidell.** 2025. “Killer congestion: Temperature, healthcare utilization and patient outcomes.” National Bureau of Economic Research.
- Angeletos, George-Marios, and Alessandro Pavan.** 2007. “Efficient use of information and social value of information.” *Econometrica*, 75(4): 1103–1142.
- Anunrojwong, Jerry, Krishnamurthy Iyer, and Vahideh Manshadi.** 2023. “Information design for congested social services: Optimal need-based persuasion.” *Management Science*, 69(7): 3778–3796.
- Arrowsmith, Ellie, Frederick Dube Fortier, Anne D. Cope, et al.** 2021. “Wildfire Fuel Management and Risk Mitigation: Where to Start?” Insurance Institute for Business & Home Safety (IBHS) and Zesty.ai Report. Accessed 2025-12-27.
- Bakkensen, Laura A, and Lala Ma.** 2020. “Sorting over flood risk and implications for policy reform.” *Journal of Environmental Economics and Management*, 104: 102362.
- Barahona, Nano, Cristóbal Otero, and Sebastián Otero.** 2023. “Equilibrium effects of food labeling policies.” *Econometrica*, 91(3): 839–868.
- Barrett, Kimiko, and Stephen L. Quarles.** 2024. “Retrofitting a Home for Wildfire Resistance: Costs and Considerations.” Headwaters Economics Report. Accessed 2026-07-05.
- Baumol, William J.** 1972. “On taxation and the control of externalities.” *The American Economic Review*, 62(3): 307–322.
- Baylis, Patrick, Judson Boomhower, Bob Horton, and Chris Mulverhill.** 2026. “Incentivizing Self-Protection from Wildfires.” National Bureau of Economic Research 35764.
- Baylis, Patrick W., and Judson Boomhower.** 2026. “Mandated versus Voluntary Adaptation to Natural Disasters: The Case of US Wildfires.” *Journal of Political Economy*, 134(3): 895–948.
- Bénabou, Roland, and Jean Tirole.** 2016. “Mindful Economics: The Production, Consumption, and Value of Beliefs.” *Journal of Economic Perspectives*, 30(3): 141–164.
- Bernheim, B. Douglas, and Antonio Rangel.** 2009. “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics.” *Quarterly Journal of Economics*, 124(1): 51–104.
- Blackwell, David.** 1953. “Equivalent comparisons of experiments.” *The annals of mathematical statistics*, 265–272.

- Boomhower, Judson, Meredith Fowlie, Jacob Gellman, and Andrew Plantinga.** 2024. “How Are Insurance Markets Adapting to Climate Change? Risk Classification and Pricing in the Market for Homeowners Insurance.” National Bureau of Economic Research.
- Brunnermeier, Markus K., and Jonathan A. Parker.** 2005. “Optimal Expectations.” *American Economic Review*, 95(4): 1092–1118.
- California Legislative Analyst’s Office.** 2021. “Reducing the Destructiveness of Wildfires: Promoting Defensible Space in California.” California Legislative Analyst’s Office Report. Accessed 2026-07-05.
- Caplin, Andrew, and John Leahy.** 2001. “Psychological Expected Utility Theory and Anticipatory Feelings.” *Quarterly Journal of Economics*, 116(1): 55–79.
- Chan, David C., Matthew Gentzkow, and Chuan Yu.** 2022. “Selection with Variation in Diagnostic Skill: Evidence from Radiologists.” *Quarterly Journal of Economics*, 137(2): 729–783.
- Charpentier, Arthur.** 2008. “Insurability of climate risks.” *The Geneva Papers on Risk and Insurance-Issues and Practice*, 33(1): 91–109.
- Cicala, Steve, David Hémous, and Morten G Olsen.** 2022. “Adverse selection as a policy instrument: unraveling climate change.” National Bureau of Economic Research.
- Condon, Madison.** 2023. “Climate services: The business of physical risk.” *Ariz. St. LJ*, 55: 147.
- CoreLogic, Inc.** 2025. “U.S. Wildfire Catastrophe Model: Quantifying the Loss Potential of Wildfires in High-Risk States.” CoreLogic, Inc. Technical update. RAA update; overview of the model methodology (hazard, ignition, burn and suppression modules, smoke plume, vulnerability, and financial modeling).
- Davis, Rob.** 2025. “How the Rapid Spread of Misinformation Pushed Oregon Lawmakers to Kill the State’s Wildfire Risk Map.” *ProPublica*.
- Dietz, Simon, and Falk Niehörster.** 2021. “Pricing ambiguity in catastrophe risk insurance.” *The Geneva Risk and Insurance Review*, 46(2): 112–132.
- Dranove, David, Daniel Kessler, Mark McClellan, and Mark Satterthwaite.** 2003. “Is more information better? The effects of “report cards” on health care providers.” *Journal of Political Economy*, 111(3): 555–588.
- Dunn, Chris, and Andy McEvoy.** 2022. “Technical Guide: Mapping Wildfire Risk to Structures and Other Human Development to Support Implementation of Oregon’s 2021 Senate Bill 762.” Oregon State University, College of Forestry. Secondary archive link provided as original state URLs are redirected to 2025 hazard map documentation.

- Figlio, David N., and Maurice E. Lucas.** 2004. “What’s in a grade? School report cards and the housing market.” *American Economic Review*, 94(3): 591–604.
- Finney, Mark A.** 2002. “Fire growth using minimum travel time methods.” *Canadian Journal of Forest Research*, 32(8): 1420–1424.
- Finney, Mark A.** 2006. “An Overview of FlamMap Fire Modeling Capabilities.” *Proceedings RMRS-P-41*, 213–220. Fort Collins, CO:U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station. Conference held March 28–30, 2006, Portland, OR.
- First Street.** 2024. “First Street Wildfire Model (FS-WFM) Technical Methodology, Version 3.0.” First Street. Published 2024-07-31. Version 3.0 fuels, ignition, and calibration updates dated 2024-08-01. TODO: confirm the cover title against the PDF (cover is an image).
- Fischer, Carolyn, and Thomas P. Lyon.** 2019. “A theory of multitier ecolabel competition.” *Journal of the Association of Environmental and Resource Economists*, 6(3): 461–501.
- Fischhoff, Baruch.** 1995. “Risk perception and communication unplugged: Twenty years of process.” *Risk Analysis*, 15(2): 137–145.
- Gibson, Matthew, and Jamie T Mullins.** 2020. “Climate risk and beliefs in New York floodplains.” *Journal of the Association of Environmental and Resource Economists*, 7(6): 1069–1111.
- Global Association of Risk Professionals.** 2025. “A Risk Professional’s Guide to Physical Risk Assessments: A GARP Benchmarking Study of 13 Vendors.” Financial Conduct Authority (FCA). Accessed: 2026-01-22.
- Goldstein, Itay, and Yaron Leitner.** 2018. “Stress tests and information disclosure.” *Journal of Economic Theory*, 177: 34–69.
- Harbaugh, Rick, and Eric Rasmusen.** 2018. “Coarse grades: Informing the public by withholding information.” *American Economic Journal: Microeconomics*, 10(1): 210–235.
- Hedlund, Jonas, Allan Hernández-Chanto, and Carlos Oyarzún.** 2024. “Contagion management through information disclosure.” *Journal of Economic Theory*, 218: 105837.
- Heitfield, Erik.** 2026. “Model Uncertainty and the Pricing of Hurricane Risk in Florida.” Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 2026-016.
- Hino, Miyuki, and Marshall Burke.** 2021. “The effect of information about climate risk on property values.” *Proceedings of the National Academy of Sciences*, 118(17): e2003374118.

- Houde, Sébastien.** 2018. “How consumers respond to product certification and the value of energy information.” *The RAND Journal of Economics*, 49(2): 453–477.
- Houde, Sébastien, and Erica Myers.** 2026. “Heterogeneous (Mis)Perceptions of Energy Costs: Implications for Measurement and Policy Design.” *Journal of Political Economy Microeconomics*, 4(3): 906–955.
- Inostroza, Nicolas A., and Alessandro Pavan.** 2025. “Adversarial coordination and public information design.” *Theoretical Economics*, 20(2): 763–813.
- Isakov, Leah, Andrew W. Lo, and Vahid Montazerhodjat.** 2019. “Is the FDA too conservative or too aggressive? A Bayesian decision analysis of clinical trial design.” *Journal of Econometrics*, 211(1): 117–136.
- Issler, Paulo, Richard Stanton, Nancy Wallace, and Yao Zhao.** 2025. “Measuring the Wildfire Risk of California Real Estate with Spatiotemporal Convolutional Neural Networks.” *Available at SSRN 5168940*.
- Jin, Ginger Zhe, and Phillip Leslie.** 2003. “The effect of information on product quality: Evidence from restaurant hygiene grade cards.” *The Quarterly Journal of Economics*, 118(2): 409–451.
- Kamenica, Emir.** 2019. “Bayesian persuasion and information design.” *Annual Review of Economics*, 11(1): 249–272.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. “Bayesian persuasion.” *American Economic Review*, 101(6): 2590–2615.
- Keohane, Nathaniel O, Richard L Revesz, and Robert N Stavins.** 2019. “The choice of regulatory instruments in environmental policy.” *Environmental law*, 491–545.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan.** 2018. “Human decisions and machine predictions.” *The Quarterly Journal of Economics*, 133(1): 237–293.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer.** 2015. “Prediction policy problems.” *American Economic Review*, 105(5): 491–495.
- Knapp, Eric E., Yana S. Valachovic, Stephen L. Quarles, and Nels G. Johnson.** 2021. “Housing arrangement and vegetation factors associated with single-family home survival in the 2018 Camp Fire, California.” *Fire Ecology*, 17(1): 25.
- Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk.** 2022. “Censorship as optimal persuasion.” *Theoretical Economics*, 17(2): 561–585.
- Kotchen, Matthew J.** 2013. “Voluntary-and information-based approaches to environmental management: A public economics perspective.” *Review of Environmental Economics and Policy*.

- Lee, Seunghoon.** 2023. “Adapting to natural disasters through better information: Evidence from the home seller disclosure requirement.” *MIT Center for Real Estate Research Paper*, , (21/17).
- Lipsey, Richard G., and Kelvin Lancaster.** 1956. “The General Theory of Second Best.” *The Review of Economic Studies*, 24(1): 11–32.
- Loewenstein, George.** 1987. “Anticipation and the Valuation of Delayed Consumption.” *Economic Journal*, 97(387): 666–684.
- Lyu, Qianjun, Wing Suen, and Yimeng Zhang.** 2023. “Coarse information design.” *arXiv preprint arXiv:2305.18020*.
- Ma, Lala, Margaret Walls, Matthew Wibbenmeyer, and Connor Lennon.** 2024. “Risk disclosure and home prices: Evidence from California wildfire hazard zones.” *Land Economics*, 100(1): 6–21.
- Moore, Frances C.** 2025. “Learning, Catastrophic Risk, and Ambiguity in the Climate-Change Era.” *Environmental and Energy Policy and the Economy*, 6(1): 140–168.
- Morris, Stephen, and Hyun Song Shin.** 2002. “Social value of public information.” *American Economic Review*, 92(5): 1521–1534.
- Mulder, Philip.** 2024. “Mismeasuring risk: The welfare effects of flood risk information.” *Available at SSRN 4966795*.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer.** 2008. “Coarse thinking and persuasion.” *The Quarterly journal of economics*, 123(2): 577–619.
- Office of Management and Budget.** 2023. “Circular A-4: Regulatory Analysis.” Office of Management and Budget, Executive Office of the President Circular. Accessed 2026-07-05.
- Onuchic, Paula, and Debraj Ray.** 2023. “Conveying value via categories.” *Theoretical Economics*, 18(4): 1407–1439.
- Oregon Board of Forestry.** 2022. “Wildfire Risk Classes, Oregon Administrative Rule 629-044-1021.” *Oregon Secretary of State, Administrative Rules Chapter 629, Division 44*, Adopted effective June 14, 2022; amended in 2024 to three hazard classes following Senate Bill 80 (2023).
- Oregon Building Codes Division.** 2022. “Senate Bill 762 Wildfire Risk Map Fact Sheet.” State of Oregon, Building Codes Division Fact sheet / update 440-5793.
- Oregon Legislative Assembly.** 2021. “Senate Bill 762.” *81st Legislative Assembly, 2021 Regular Session. Enacted as Oregon Laws 2021, chapter 592*.
- Oregon Wildfire Programs Advisory Council.** 2022. “Reducing Wildfire Risk to Oregon’s Homes and Communities: Fact Sheet.” State of Oregon Fact sheet.

- Ostriker, Abigail, and Anna Russo.** 2022. “The effects of floodplain regulation on housing markets.” *Working Paper*.
- Pollack, Adam B, Lisa Auermuller, Casey D Burleyson, Jentry Campbell, Madison Condon, Courtney Cooper, Matteo Coronese, Sönke Dangendorf, James Doss-Gollin, Prabhat Hegde, et al.** 2026. “Unlocking the benefits of transparent and reusable science for climate risk management.” *Proceedings of the National Academy of Sciences*, 123(3): e2422157123.
- Profita, Cassandra.** 2022. “Swamped by public outcry, Oregon withdraws controversial wildfire risk map.” *OPB*.
- Rayo, Luis, and Ilya Segal.** 2010. “Optimal information disclosure.” *Journal of political Economy*, 118(5): 949–987.
- Rodemeier, Matthias, and Andreas Löschel.** 2025. “Information nudges, subsidies, and crowding out of attention: Field evidence from energy efficiency investments.” *Journal of the European Economic Association*, 23(4): 1397–1430.
- Rothermel, Richard C.** 1972. “A mathematical model for predicting fire spread in wildland fuels.” U.S. Department of Agriculture, Forest Service, Intermountain Forest and Range Experiment Station, Ogden, UT.
- Schubert, Jochen E, Katharine J Mach, and Brett F Sanders.** 2024. “National-scale flood hazard data unfit for urban risk management.” *Earth’s future*, 12(7): e2024EF004549.
- Scott, Joe H., Karen C. Short, and Mark A. Finney.** 2018. “FSim: The Large-Fire Simulator: Guide to Best Practices.” Pyrologix LLC and USDA Forest Service, Rocky Mountain Research Station. Version 0.3.1.
- Shafran, Aric P.** 2008. “Risk externalities and the problem of wildfire risk.” *Journal of Urban Economics*, 64(2): 488–495.
- Slovic, Paul.** 1987. “Perception of risk.” *Science*, 236(4799): 280–285.
- Syphard, Alexandra D., Teresa J. Brennan, and Jon E. Keeley.** 2014. “The role of defensible space for residential structure protection during wildfires.” *International Journal of Wildland Fire*, 23(8): 1165–1175.
- Tetenov, Aleksey.** 2012. “Statistical treatment choice based on asymmetric minimax regret criteria.” *Journal of Econometrics*, 166(1): 157–165.
- Tietenberg, Tom.** 1998. “Disclosure strategies for pollution control.” *Environmental and resource Economics*, 11(3): 587–602.
- Vatter, Benjamin.** 2025. “Quality disclosure and regulation: Scoring design in medicare advantage.” *Econometrica*, 93(3): 959–1001.

- Wagner, Katherine RH.** 2022. “Adaptation and adverse selection in markets for natural disaster insurance.” *American Economic Journal: Economic Policy*, 14(3): 380–421.
- Weill, Joakim A.** 2026. “Flood risk mapping and the distributional impacts of climate information.” *American Economic Journal: Economic Policy*, 18(3): 275–305.
- Weitzman, Martin.** 1974. “Prices vs. Quantities.” *The Review of Economic Studies*, 41(4): 477–491.
- You, Xuesong, and Carolyn Kousky.** 2024. “Improving household and community disaster recovery: Evidence on the role of insurance.” *Journal of Risk and Insurance*, 91(2): 299–338.
- Zamanialaei, Maryam, Daniel San Martin, Maria Theodori, Dwi Marhaendro Jati Purnomo, Ali Tohidi, Chris Lautenberger, Yiren Qin, Arnaud Trouvé, and Michael Gollner.** 2025. “Fire risk to structures in California’s Wildland-Urban Interface.” *Nature Communications*, 16(1): 8041.

A Logical structure of the paper

Figure S1 summarizes the logical structure of the paper.




Figure S1: Logical structure of the paper.

Notes: Top tier: theory (Section 2). Middle tier: realized fires replace the unobservable private benchmark (Section 3). Bottom tier: empirical steps of Section 4, each tagged with the fire window it uses.

B First-best benchmark and welfare gap

Restrict attention to cutoff rules with adaptation rate $A(x) = \Pr(m \geq x)$, and define $L \equiv \mathbb{E}[\ell(\theta)]$ and $L_{\text{hi}}(x) \equiv \mathbb{E}[\ell(\theta) \mathbf{1}\{m \geq x\}]$. The per-property social cost simplifies to

$$C(x) = c A(x) + \phi(A(x))(L - \alpha L_{\text{hi}}(x)). \quad (8)$$

A social planner with score-contingent Pigouvian taxes on non-adaptation chooses m_{FB} , the largest minimizer of $C(x)$, internalizing the adaptation externality and attaining

$$C^{FB} = c A^{FB} + \phi(A^{FB})(L - \alpha L_{\text{hi}}^{FB}), \quad (9)$$

where $A^{FB} \equiv A(m_{FB})$ and $L_{\text{hi}}^{FB} \equiv L_{\text{hi}}(m_{FB})$. The welfare gap $\Delta C^T \equiv C^T - C^{FB}$ between the optimal threshold disclosure rule (Proposition 1(ii)) and the first best decomposes as

$$\Delta C^T = \underbrace{c(A^T - A^{FB})}_{\text{aggregate adaptation gap}} + \underbrace{[\phi(A^T) - \phi(A^{FB})](L - \alpha L_{\text{hi}}^T)}_{\text{residual-risk gap}} - \underbrace{\alpha \phi(A^{FB})(L_{\text{hi}}^T - L_{\text{hi}}^{FB})}_{\text{composition effect}}, \quad (10)$$

where $A^T \equiv A(m^*)$ and $L_{\text{hi}}^T \equiv L_{\text{hi}}(m^*)$.

Define bin-level expected private gains

$$I_{\text{high}}(x) \equiv \mathbb{E}[\alpha \ell(\theta) \phi(A(x)) - c \mid m \geq x], \quad I_{\text{low}}(x) \equiv \mathbb{E}[\alpha \ell(\theta) \phi(A(x)) - c \mid m < x].$$

A separating equilibrium at cutoff x requires $I_{\text{high}}(x) \geq 0$ and $I_{\text{low}}(x) \leq 0$. Since in every regime the adapters at cutoff x form the top set $\{m \geq x\}$ (Lemma 1), all regimes share the same aggregate $A(x) = 1 - G(x)$, and each regime's cutoff condition compares a single schedule with the adaptation cost c :

$$\begin{aligned} \mathcal{H}^D(x) &\equiv \alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m = x], \\ \mathcal{H}^T(x) &\equiv \alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m \geq x], \\ SMB(x) &\equiv \alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m = x] - \phi'(A(x))(L - \alpha L_{\text{hi}}(x)). \end{aligned}$$

$\mathcal{H}^D$ is the private indifference schedule of detailed disclosure (the schedule $\mathcal{H}$ of equation (1)): m_D solves $\mathcal{H}^D(x) = c$. The social marginal benefit SMB adds the spillover term: by Assumption 3, differentiating (8) gives $C'(x) = g(x)[SMB(x) - c]$, so an interior m_{FB} solves $SMB(x) = c$. $\mathcal{H}^T$ is the pooled marginal benefit of the high bin, with $I_{\text{high}}(x) = \mathcal{H}^T(x) - c$. Both $\mathcal{H}^D$ and $\mathcal{H}^T$ are strictly increasing, the conditional means by Assumption 1 and $\phi(A(x))$ because $A(x)$ falls in x and ϕ is decreasing, so each crossing is unique. Because the bin's tail average exceeds the marginal type's posterior, $\mathbb{E}[\ell(\theta) \mid m \geq x] > \mathbb{E}[\ell(\theta) \mid m = x]$, the schedule $\mathcal{H}^T$ lies strictly above $\mathcal{H}^D$ at every x : pooling replaces the marginal type's own posterior with the bin average, which is why coarse disclosure sustains a lower cutoff, $m^* < m_D$.

Lemma 2 (The low-bin constraint never binds). *Suppose Assumption 1 holds, $\phi'(A) < 0$ for all $A \in [0, 1]$, and on $(\underline{m}, \overline{m})$ the score has a density g , with g and $\mathbb{E}[\ell(\theta) \mid m]$ continuous and $g > 0$. Then $I_{\text{low}} < 0$ at an interior first-best cutoff m_{FB} and at any interior optimal implementable threshold m^* . Consequently, either $I_{\text{high}}(m_{FB}) \geq 0$, the first best is implementable, and $C^T = C^{FB}$, or $I_{\text{high}}(m_{FB}) < 0$, $m^* > m_{FB}$, and $C^T > C^{FB}$.*

Proof. At m_{FB} , $SMB(m_{FB}) = c$ gives $\alpha \phi(A^{FB}) \mathbb{E}[\ell(\theta) \mid m = m_{FB}] = c + \phi'(A^{FB})(L - \alpha L_{hi}^{FB}) < c$, and $\mathbb{E}[\ell(\theta) \mid m < m_{FB}] < \mathbb{E}[\ell(\theta) \mid m = m_{FB}]$ by Assumption 1 and the positive mass below m_{FB} , so $I_{low}(m_{FB}) < 0$. At m^* , suppose $I_{low}(m^*) = 0$. Then $\alpha \phi(A(m^*)) \mathbb{E}[\ell(\theta) \mid m = m^*] > c$, so $SMB(m^*) > c$ and $C'(m^*) > 0$. The function $x \mapsto \alpha \phi(A(x)) \mathbb{E}[\ell(\theta) \mid m < x]$ is nondecreasing, I_{high} is continuous, and $I_{high}(m^*) > I_{low}(m^*) = 0$, so thresholds slightly below m^* remain feasible and carry strictly lower cost, contradicting the optimality of m^* . For the consequence, if $I_{high}(m_{FB}) \geq 0$ then m_{FB} is feasible and attains the unconstrained minimum of C , so $C^T = C^{FB}$; if $I_{high}(m_{FB}) < 0$, every feasible x has $I_{high}(x) \geq 0 > I_{high}(m_{FB})$, and I_{high} is increasing, so $x > m_{FB}$; since m_{FB} is the largest minimizer of C , $C(x) > C^{FB}$ at every feasible x , and in particular $C^T = C(m^*) > C^{FB}$. $\square$

Figure S2 draws the under-adaptation case: SMB reaches c at m_{FB} before $\mathcal{H}^T$ does, so the high-bin constraint fails at m_{FB} and m^* solves $\mathcal{H}^T(x) = c$. Had c lain below the level at which the two schedules cross in the figure, $\mathcal{H}^T$ would reach c first and the first best would be implementable.




Figure S2: Adaptation cutoffs under the first best, optimal threshold disclosure, and detailed disclosure.

Notes: Stylized schedules. The horizontal line at c is the cost of adapting. The cutoffs m_{FB} and m_D solve $SMB(x) = c$ and $\mathcal{H}^D(x) = c$. In the case drawn, as in the example of Figure 2, the high-bin constraint fails at m_{FB} , so m^* solves $\mathcal{H}^T(x) = c$. SMB is drawn increasing, which holds when ϕ is weakly convex (Appendix E). Horizontal bars mark the properties that adapt in each regime, those with scores at or above the regime's cutoff.

C Proof of Proposition 1

Proof. Part (i): no spillovers. Under detailed disclosure, individuals observe m and adapt if and only if

$$\mathbb{E}[\alpha \ell(\theta) \phi_0 - c \mid m] \geq 0.$$

Since this expression is strictly increasing in m (by Assumption 1), there exists a cutoff m_D such that the expression is negative for $m < m_D$ and positive for $m > m_D$ (with equality at $m = m_D$). Hence the best response takes the cutoff form

$$a^D(m) = 1\{m \geq m_D\},$$

where $a^D(m) = 1$ if the individual adapts, with an arbitrary tie-break at $m = m_D$.

Threshold disclosure. Consider the threshold disclosure rule that reveals only whether m exceeds the same cutoff, i.e. the message $1\{m \geq m_D\}$. Because the disclosure rule is publicly known, observing the high message implies $m \geq m_D$, so

$$\mathbb{E}[\alpha \ell(\theta) \phi_0 - c \mid m \geq m_D] \geq 0,$$

and adapting is optimal. Likewise, observing the low message implies $m < m_D$, so

$$\mathbb{E}[\alpha \ell(\theta) \phi_0 - c \mid m < m_D] \leq 0,$$

and not adapting is optimal. Therefore the action induced under threshold disclosure is

$$a^T(m) = 1\{m \geq m_D\} = a^D(m),$$

again up to tie-breaking at $m = m_D$.

Since the two disclosure regimes induce the same adaptation decisions, expected social cost is the same: $C^T = C^D$. To see that m_D is the unique minimizer of $C(x)$, note that with ϕ constant, $C'(x) = g(x) [\phi_0 \alpha \mathbb{E}[\ell(\theta) \mid m = x] - c]$. By Assumption 1 this is negative for $x < m_D$ and positive for $x > m_D$, so $C(x)$ is strictly decreasing then strictly increasing, with a unique minimum at m_D . Hence no threshold rule achieves cost below $C(m_D) = C^D$, confirming $C^T = C^D$.

Part (ii): positive spillovers. By Lemma 1 both regimes induce score cutoffs; let m_D be the (unique) detailed-disclosure cutoff, with $A^D = A(m_D)$ and private indifference

$$\mathbb{E}[\alpha \ell(\theta) \phi(A^D) - c \mid m = m_D] = 0. \quad (11)$$

By Assumption 3, $A(x) = 1 - G(x)$ is differentiable near m_D with $A'(m_D) = -g(m_D) < 0$. Differentiating the cost $C(x)$ defined in Section 2 and evaluating at $x = m_D$, the term generated by the moving indicator boundary is proportional to the private net gain at $m = m_D$ and hence vanishes by (11); only the inframarginal spillover channel survives:

$$\left. \frac{dC(x)}{dx} \right|_{x=m_D} = \phi'(A^D) A'(m_D) \left\{ \mathbb{E}[\ell(\theta)] - \mathbb{E}[\alpha \ell(\theta) 1\{m \geq m_D\}] \right\}.$$

Both factors are negative, $\phi'(A^D) < 0$ by assumption and $A'(m_D) < 0$ by Assumption 3, and the bracket is strictly positive: it is $\mathbb{E}[\ell(\theta)]$ minus the loss accruing to adapters scaled by $\alpha \in (0, 1)$, which is strictly smaller. Hence $dC/dx|_{m_D} > 0$, so lowering the cutoff slightly below m_D strictly reduces social cost.

It remains to check that a threshold near m_D is implementable, i.e. the high bin weakly prefers to adapt and the low bin weakly prefers not to. At $x = m_D$, the law of iterated expectations and Assumption 1 give $\mathbb{E}[\ell(\theta) \mid m \geq m_D] > \mathbb{E}[\ell(\theta) \mid m = m_D] \geq \mathbb{E}[\ell(\theta) \mid m < m_D]$, the first inequality strict because $\mathbb{E}[\ell(\theta) \mid m]$ is strictly increasing (Assumption 1) and m_D is interior with positive mass above it (Assumptions 2 and 3). Hence by (11) the high-bin constraint $\alpha\phi(A^D)\mathbb{E}[\ell(\theta) \mid m \geq m_D] \geq c$ holds with strict slack and the low-bin constraint $\alpha\phi(A^D)\mathbb{E}[\ell(\theta) \mid m < m_D] \leq c$ holds. Lowering the threshold slightly relaxes the low-bin constraint (the low bin sheds its highest scores), while the high bin's strict slack survives a small perturbation by continuity, so both constraints hold for x slightly below m_D . Therefore there exists $\delta > 0$ such that

$$C^T = C(m^*) \leq C(m_D - \delta) < C(m_D) = C^D.$$

Moreover $m^* < m_D$: by Lemma 2 in Appendix B, at m^* either $C'(m^*) = 0$ or the high-bin constraint binds, and in both cases the private indifference schedule $\mathcal{H}$ of (1) lies strictly below c at m^* , while $\mathcal{H}(m_D) = c$ and $\mathcal{H}$ is strictly increasing. $\square$

D Proof of Proposition 2 (Structural content of the recovered $\hat{\lambda}$)

Proof. Because $\mathbb{E}[\ell(\theta) \mid m] = \bar{L} \mathbb{E}[\theta \mid m] = \bar{L} \pi(m)$, Assumption 1 makes π strictly increasing, so $\{m \geq x\} = \{\pi(m) \geq \pi(x)\}$ and the threshold $\hat{m}$ implements the cutoff $\hat{\kappa} = \pi(\hat{m})$. For part (i), write the two objectives at threshold x as

$$\begin{aligned} R_\lambda(x) &\equiv \mathbb{E}[\lambda(1-Y) 1\{m \geq x\} + Y 1\{m < x\}], \\ C_A(x) &\equiv \mathbb{E}[\ell(\theta)] \phi(A) - \mathbb{E}[(\alpha \ell(\theta) \phi(A) - c) 1\{m \geq x\}]. \end{aligned}$$

Write $a \in \{0, 1\}$ for the classification of a parcel, with $a = 1$ when it is flagged. Conditioning on m ,

$$\mathbb{E}[\lambda(1-Y)a + Y(1-a) \mid m] = \lambda(1-\pi(m))a + \pi(m)(1-a) = \pi(m) + (1+\lambda)a(\kappa(\lambda) - \pi(m)),$$

and taking expectations over m with $a = 1\{m \geq x\}$ gives

$$R_\lambda(x) = \mathbb{E}[\pi(m)] + (1+\lambda) \mathbb{E}[(\kappa(\lambda) - \pi(m)) 1\{m \geq x\}].$$

Similarly, at a fixed adaptation rate A the per-parcel private cost of action a is

$$\begin{aligned} (1-a) \mathbb{E}[\ell(\theta) \mid m] \phi(A) + a((1-\alpha) \mathbb{E}[\ell(\theta) \mid m] \phi(A) + c) \\ = \bar{L} \phi(A) \pi(m) + \alpha \phi(A) \bar{L} a (\kappa^{\text{priv}}(A) - \pi(m)), \end{aligned}$$

which, after taking expectations with $a = 1\{m \geq x\}$ and using $\alpha \phi(A) \bar{L} \kappa^{\text{priv}}(A) = c$, yields

$$C_A(x) = \bar{L} \phi(A) \mathbb{E}[\pi(m)] + \alpha \phi(A) \bar{L} \mathbb{E}[(\kappa^{\text{priv}}(A) - \pi(m)) 1\{m \geq x\}].$$

Both objectives thus take the form $\beta_0 \mathbb{E}[\pi(m)] + \beta_1 \mathbb{E}[(\kappa - \pi(m)) 1\{m \geq x\}]$ with $\beta_1 > 0$. If $\kappa(\lambda) = \kappa^{\text{priv}}(A)$, the two differ only by the positive scale β_1 and an additive constant. Conversely, $R_\lambda(x) = u + v C_A(x)$ for all x with $v > 0$ requires matching the coefficients on $\Pr(m \geq x)$ and $\mathbb{E}[\pi(m) 1\{m \geq x\}]$ separately; these two functions of x are not proportional when π is nonconstant (Assumption 1), so matching forces $(1+\lambda) = v \alpha \phi(A) \bar{L}$ and $(1+\lambda) \kappa(\lambda) = v \alpha \phi(A) \bar{L} \kappa^{\text{priv}}(A)$, hence $\kappa(\lambda) = \kappa^{\text{priv}}(A)$.

For part (ii), let $\phi \equiv \phi_0$. Appendix C shows that C is then strictly decreasing below m_D and strictly increasing above it, so $m_{FB} = m_D$, and that the threshold m_D is implementable, so $m^* = m_D$. At this common cutoff $\pi(m_D) = \kappa^{\text{priv}}$, and with ϕ constant the private cost C_A coincides with the social cost C , so by part (i) the planner equates the λ -weighted loss (5) and $\hat{\lambda} = \hat{\kappa}/(1-\hat{\kappa})$ is the welfare trade-off.

For part (iii), write the per-property social cost as in (8), $C(x) = \phi(A(x)) [L - \alpha L_{\text{hi}}(x)] + c A(x)$ with $A(x) = 1 - G(x)$. On $(\underline{m}, \bar{m})$, $A'(x) = -g(x)$ and $L'_{\text{hi}}(x) = -\bar{L} \pi(x) g(x)$, so C is differentiable with

$$C'(x) = g(x) \alpha \phi(A(x)) \bar{L} [\pi(x) - \kappa^{\text{priv}}(A(x)) - w(x)],$$

where $w(x)$ is the spillover wedge of (6) evaluated at x . This is the derivative $C'(x) = g(x)[SMB(x) - c]$ of Appendix B with $\mathbb{E}[\ell(\theta) \mid m = x] = \bar{L}\pi(x)$. Since $g > 0$, $C'(x)$ has the sign of the bracket, and $w(x) < 0$ because $\phi' < 0$ and $L - \alpha L_{\text{hi}}(x) \geq (1 - \alpha)L > 0$.

If $\hat{m} = m_{FB}$, it is an interior minimizer of C , so $C'(m_{FB}) = 0$ and $\hat{\kappa} = \kappa^{\text{priv}} + w(\hat{m}) < \kappa^{\text{priv}}$.

Let now $\hat{m} = m^*$, which minimizes C over the thresholds $x \in (\underline{m}, \bar{m})$ satisfying the high-bin and low-bin incentive constraints

$$\alpha \phi(A(x)) \bar{L} \mathbb{E}[\pi(m) \mid m \geq x] \geq c \quad \text{and} \quad \alpha \phi(A(x)) \bar{L} \mathbb{E}[\pi(m) \mid m < x] \leq c.$$

Both constraint functions are continuous and nondecreasing in x : $\phi(A(x))$ rises as $A(x)$ falls, and each bin's mean of π rises as the high bin sheds its lowest scores or the low bin gains higher ones. At any $x \in (\underline{m}, \bar{m})$ the high bin has positive mass and π is strictly increasing, so $\pi(x) < \mathbb{E}[\pi(m) \mid m \geq x]$. The low-bin constraint is strictly slack at m^* by Lemma 2, whose hypotheses hold here because $\mathbb{E}[\ell(\theta) \mid m] = \bar{L}\pi(m)$ is continuous and strictly increasing.

Two cases remain. If the high-bin constraint is slack at m^* , a neighborhood of m^* is feasible, so m^* is a local minimizer of C , $C'(m^*) = 0$, and $\hat{\kappa} = \kappa^{\text{priv}} + w(m^*) < \kappa^{\text{priv}}$. If it binds, $\mathbb{E}[\pi(m) \mid m \geq m^*] = \kappa^{\text{priv}}$, and $\hat{\kappa} = \pi(m^*) < \kappa^{\text{priv}}$. For the lower bound, thresholds slightly above m^* remain feasible, because the high-bin constraint function is nondecreasing and the low-bin constraint is strictly slack; optimality of m^* therefore requires $C'(m^*) \geq 0$, that is, $\hat{\kappa} \geq \kappa^{\text{priv}} + w(m^*)$.

In every case $\kappa^{\text{priv}} + w(\hat{m}) \leq \hat{\kappa} < \kappa^{\text{priv}}$. The upper boundary condition of Assumption 2 gives $\alpha \phi(A) \bar{L} \pi(m) > c$ for m close to $\bar{m}$, so $\kappa^{\text{priv}} < 1$, and $\hat{\lambda} < \kappa^{\text{priv}} / (1 - \kappa^{\text{priv}})$ follows because $\kappa \mapsto \kappa / (1 - \kappa)$ is increasing on $[0, 1)$. $\square$

E Location of the optimal disclosure threshold

This appendix derives two comparative statics for the optimal disclosure cutoff m^* : the first with respect to the overall level of losses (loss scaling), the second with respect to how sharply the score sorts losses (score polarization).

Loss scaling

Suppose the loss function takes the form $\ell(\theta; \eta) = \eta \ell_0(\theta)$ for a scale parameter $\eta > 0$, holding the baseline profile ℓ_0 and the joint distribution of (θ, m) fixed. Because the score m is a signal of θ alone and does not depend on ℓ , the score distribution G and the adaptation rate $A(x) = 1 - G(x)$ are unaffected by η . I also require the schedule SMB of Appendix B to be strictly increasing on $(\underline{m}, \overline{m})$. This holds whenever ϕ is weakly convex, and in particular for linear ϕ : $\mathcal{H}^D$ is strictly increasing, $x \mapsto -\phi'(A(x))$ is positive and nondecreasing because ϕ' is nondecreasing and A is nonincreasing, and $x \mapsto L - \alpha L_{hi}(x)$ is positive and strictly increasing.

Proposition 3 (Loss scaling). *Suppose the hypotheses of Lemma 2 hold and SMB is strictly increasing on $(\underline{m}, \overline{m})$. On any range of η over which $m^*(\eta)$ is interior, $m^*(\eta)$ is strictly decreasing in η .*

Proof. Under loss scaling, $\mathcal{H}^T(x; \eta) = \eta \Gamma(x)$ and $SMB(x; \eta) = \eta \Pi(x)$, where, writing L_0 and $L_{hi,0}$ for the objects L and L_{hi} at $\eta = 1$,

$$\begin{aligned}\Gamma(x) &\equiv \alpha \phi(A(x)) \mathbb{E}[\ell_0(\theta) \mid m \geq x], \\ \Pi(x) &\equiv \alpha \phi(A(x)) \mathbb{E}[\ell_0(\theta) \mid m = x] - \phi'(A(x))(L_0 - \alpha L_{hi,0}(x)),\end{aligned}$$

neither of which depends on η . Both are continuous on $(\underline{m}, \overline{m})$ and strictly increasing: Γ by the argument of Appendix B ($\phi(A(x))$ rises as $A(x)$ falls, and the tail mean rises strictly by Assumption 1), and Π by hypothesis. Let $m^T(\eta)$ denote the root of $\Gamma(x) = c/\eta$ on $(\underline{m}, \overline{m})$, and $m_{FB}(\eta)$ the root of $\Pi(x) = c/\eta$, with the convention that a root equals $\underline{m}$ when the schedule exceeds c/η on the whole interval. Each root is unique, and because c/η is strictly decreasing in η while Γ and Π are strictly increasing in x , each interior root is strictly decreasing in η .

I claim $m^*(\eta) = \max\{m_{FB}(\eta), m^T(\eta)\}$. Since $C'(x) = g(x)[\eta \Pi(x) - c]$ with $g > 0$, the cost $C(\cdot; \eta)$ is strictly decreasing on $(\underline{m}, m_{FB}]$ and strictly increasing on $[m_{FB}, \overline{m})$. The high-bin constraint $\eta \Gamma(x) \geq c$ holds exactly on $[m^T, \overline{m})$. The low-bin constraint is slack at both candidates: at m_{FB} by Lemma 2, and at m^T because $\mathcal{H}^D(m^T) < \mathcal{H}^T(m^T) = c$ and $\mathbb{E}[\ell(\theta) \mid m < m^T] < \mathbb{E}[\ell(\theta) \mid m = m^T]$. If $m_{FB} \geq m^T$, then m_{FB} is feasible and attains the unconstrained minimum of C . If $m^T > m_{FB}$, every feasible threshold lies in $[m^T, \overline{m})$, where C is strictly increasing, so m^T is the unique constrained minimizer.

Finally, let $\eta_1 < \eta_2$ lie in a range over which m^* is interior, and let r denote whichever of m_{FB} and m^T attains the maximum at η_2 . Then $r(\eta_2) = m^*(\eta_2)$ is interior, so r is strictly decreasing between η_1 and η_2 , and $m^*(\eta_2) = r(\eta_2) < r(\eta_1) \leq m^*(\eta_1)$. $\square$

Score polarization

Index the joint distribution of (θ, m) by $s \geq 0$ and write $\mathbb{E}_s$ for expectations under it. Along the family the mean loss L and the score distribution G , hence $A(x) = 1 - G(x)$, stay fixed; what varies is how sharply the score sorts losses. Fixing G is a normalization: every object in the planner's problem depends on the score only through its rank, so scores with different distributions can each be expressed in quantile units, and a lower threshold then means a higher adaptation rate.

Assumption 4 (Score polarization). *For every interior x in the support of m , $\mathbb{E}_s[\ell(\theta) \mid m \geq x]$ is strictly increasing in s .*

With G and L fixed, the low-bin mean $\mathbb{E}_s[\ell(\theta) \mid m < x]$ then falls in s : raising s concentrates high losses at high scores and low losses at low scores.

Proposition 4 (Score polarization lowers a binding threshold). *Suppose the hypotheses of Lemma 2 hold at every s , together with Assumption 4, and let $s_1 < s_2$. If the high-bin incentive constraint binds at $m^*(s_2)$, then $m^*(s_2)$ lies strictly below every threshold implementable at s_1 . In particular, $m^*(s_2) < m^*(s_1)$ and $A(m^*(s_2)) > A(m^*(s_1))$.*

Proof. Write $\mathcal{H}^T(x; s) = \alpha \phi(A(x)) \mathbb{E}_s[\ell(\theta) \mid m \geq x]$ for the high-bin schedule under $\mathbb{E}_s$. Because $A(x)$ does not vary with s , Assumption 4 and the binding constraint at s_2 give

$$\mathcal{H}^T(m^*(s_2); s_1) < \mathcal{H}^T(m^*(s_2); s_2) = c.$$

Since $\mathcal{H}^T(\cdot; s_1)$ is strictly increasing (Appendix B), $\mathcal{H}^T(x; s_1) < c$ at every $x \leq m^*(s_2)$, so the high-bin constraint fails at all such x and every threshold implementable at s_1 lies strictly above $m^*(s_2)$. The ranking of adaptation rates follows because $g > 0$ makes A strictly decreasing. $\square$

F Continuous adaptation

The baseline model of Section 2.2 treats adaptation as binary: each property either fully adapts ($a = 1$, reducing its own expected loss by the fraction α at cost c) or does not ($a = 0$). This appendix relaxes that assumption and lets adaptation be a continuous intensity $a \in [0, \bar{a}]$. There are two findings. First, the dominance result of Proposition 1(ii) ($C^T < C^D$) does not hold unconditionally once adaptation is continuous: coarse disclosure now distorts the adaptation intensity of the properties it pools, a *composition cost* that must be weighed against the spillover gain. Second, when a ceiling on adaptation binds for the highest-risk properties, a rule that pools the highest scores into one cell and reports all other scores exactly strictly improves on detailed disclosure for any positive spillover.

Continuous-intensity environment

Using the same notation as Section 2.2: types $\theta \in [0, 1]$ with cumulative distribution function F ; loss $\ell : [0, 1] \rightarrow (0, \infty)$ continuous and strictly increasing; score m with cumulative distribution function G and density g , satisfying Assumption 1; residual-risk factor $\phi(A) > 0$, $\phi \in C^1$, and (in the spillover case) $\phi' < 0$; and the loss-reduction parameter $\alpha \in (0, 1)$.

I modify only the action set: a property now chooses an adaptation *intensity* $a \in [0, \bar{a}]$ with $\bar{a} \in (0, 1]$, faces residual own expected loss $(1 - \alpha a) \ell(\theta) \phi(A)$, and pays the adaptation cost $c(a)$. I always write the cost function with its argument or a prime, so that c alone keeps its main-text meaning. The aggregate A is now the population average intensity.

Assumption 5 (Convex adaptation costs). The cost function $c : [0, \bar{a}] \rightarrow [0, \infty)$ is twice continuously differentiable with $c(0) = 0$, $c' > 0$ on $(0, \bar{a}]$, and $c'' > 0$ on $[0, \bar{a}]$.

The binary model of the main text is the special case $\bar{a} = 1$ with linear cost, in which $c(a)$ equals a times the main-text cost c . The private objective is then linear in a , so the individual chooses a corner: he adapts fully when $\alpha \mathbb{E}[\ell(\theta) \mid \text{info}] \phi(A) \geq c$ and does not adapt otherwise. This is the cutoff rule of Lemma 1, with indifferent individuals adapting as in the main text. Assumption 5 excludes this case: strictly convex cost makes partial adaptation optimal for a range of beliefs.

Because the private objective is linear in $\ell(\theta)$, an individual's decision depends on his information only through the posterior mean loss

$$\varepsilon \equiv \mathbb{E}[\ell(\theta) \mid \text{info}].$$

A disclosure rule is a partition $\mathcal{B}$ of the scores into cells, and it reports to each property the cell that contains its score. Under detailed disclosure each cell is a single score, so $\varepsilon = \mathbb{E}[\ell(\theta) \mid m]$, which is finite and strictly increasing in m by Assumption 1. Under a coarser rule, a property whose score lies in cell B holds the belief $\varepsilon = \mathbb{E}[\ell(\theta) \mid m \in B]$, which by the law of iterated expectations is the average of the detailed beliefs $\mathbb{E}[\ell(\theta) \mid m]$ over the scores in B . Threshold disclosure at cutoff x , the policy instrument of the main text, is the two-cell partition $\{m < x\}, \{m \geq x\}$.

Private optimum and aggregate equilibrium

Lemma 3 (Best response). *Under Assumption 5, for any belief $\varepsilon \geq 0$ and any aggregate $A \in [0, 1]$, the privately optimal intensity is unique and equals*

$$a^*(\varepsilon, A) = \begin{cases} 0 & \text{if } \alpha \varepsilon \phi(A) \leq c'(0), \\ (c')^{-1}(\alpha \varepsilon \phi(A)) & \text{if } c'(0) < \alpha \varepsilon \phi(A) < c'(\bar{a}), \\ \bar{a} & \text{if } \alpha \varepsilon \phi(A) \geq c'(\bar{a}). \end{cases}$$

It is continuous in (ε, A) and non-decreasing in ε , and it is non-increasing in A whenever $\phi' < 0$.

Proof. The individual minimizes the perceived private cost $(1 - \alpha a) \varepsilon \phi(A) + c(a)$ over $a \in [0, \bar{a}]$. Its derivative in a , $c'(a) - \alpha \varepsilon \phi(A)$, is continuous and strictly increasing on $[0, \bar{a}]$ by Assumption 5. If this derivative is non-negative at $a = 0$, it is positive on $(0, \bar{a}]$, so the cost is strictly increasing and the minimizer is 0. If it is non-positive at $a = \bar{a}$, it is negative on $[0, \bar{a})$, so the cost is strictly decreasing and the minimizer is $\bar{a}$. Otherwise the derivative is negative at 0 and positive at $\bar{a}$, so it changes sign exactly once, at the unique $a \in (0, \bar{a})$ with $c'(a) = \alpha \varepsilon \phi(A)$, and that point is the minimizer.

Because c' is continuous and strictly increasing, so is its inverse, which maps $c'(0)$ to 0 and $c'(\bar{a})$ to $\bar{a}$. Hence a^* is a continuous, non-decreasing function of the product $\alpha \varepsilon \phi(A)$, which is continuous in (ε, A) , increasing in ε , and non-increasing in A when $\phi' < 0$. $\square$

Where adaptation is interior, $a^*(\varepsilon, A) = (c')^{-1}(\alpha \varepsilon \phi(A))$. The inverse of an increasing convex function is concave, so the response is *concave* in beliefs where marginal cost c' is convex, *linear* when c is quadratic, and *convex* where c' is concave. The ceiling adds a concave kink at the belief where $\alpha \varepsilon \phi(A) = c'(\bar{a})$, above which the response is flat, and when $c'(0) > 0$ the floor at zero adds a convex kink where $\alpha \varepsilon \phi(A) = c'(0)$.

Lemma 4 (Existence and uniqueness of the adaptation equilibrium). *Fix a disclosure rule $\mathcal{B}$. An aggregate equilibrium is a fixed point of the response map*

$$\mathcal{T}(A) \equiv \mathbb{E}[a^*(\varepsilon, A)],$$

where the expectation is over scores $m \sim G$ and ε is the belief that $\mathcal{B}$ assigns to m . Under Assumption 5 with $\phi' < 0$, $\mathcal{T}$ is continuous and non-increasing in A , so the equilibrium aggregate $A(\mathcal{B})$ exists and is unique.

Proof. For every belief ε , $a^*(\varepsilon, A)$ is continuous and non-increasing in A (Lemma 3) and lies in $[0, \bar{a}]$, so $\mathcal{T}$ is continuous and non-increasing, and any fixed point lies in $[0, \bar{a}]$. On that interval $\mathcal{T}(A) - A$ is continuous and strictly decreasing, non-negative at $A = 0$ and non-positive at $A = \bar{a}$, so it has exactly one zero. $\square$

Social cost

The per-property expected social cost of a disclosure rule $\mathcal{B}$ with equilibrium aggregate $A = A(\mathcal{B})$ is

$$C(\mathcal{B}) = \mathbb{E}[\psi(\varepsilon, A)], \quad \psi(\varepsilon, A) \equiv \min_{a \in [0, \bar{a}]} \{(1 - \alpha a) \varepsilon \phi(A) + c(a)\}, \quad (12)$$

with the expectation again over $m \sim G$. Equation (12) uses the law of iterated expectations: conditional on the information an individual observes, his action a^* is already determined and his expected loss is ε , so his expected social cost $\mathbb{E}[(1 - \alpha a^*) \ell(\theta) \phi(A) + c(a^*) \mid \text{info}]$ equals his minimized private cost $\psi(\varepsilon, A)$. The only externality runs through the aggregate A . Write C^D and A^D for the social cost and the equilibrium aggregate under detailed disclosure.

The results below use two properties of ψ . First, for each a the expression in braces is affine in ε , so $\psi(\cdot, A)$, a minimum of affine functions, is concave. Second, when $\phi' < 0$ the expression in braces is strictly decreasing in A for each a , because $(1 - \alpha a) \varepsilon > 0$, so $\psi(\varepsilon, \cdot)$ is strictly decreasing: a higher aggregate lowers everyone's cost.

Decomposing the welfare effect of coarsening

The welfare effect of moving from detailed disclosure to a coarser rule splits into two steps: change the information while holding the aggregate at A^D , then let the aggregate adjust.

Proposition 5 (Composition–spillover decomposition). *Suppose Assumptions 1 and 5 hold and $\phi' < 0$. For any disclosure rule $\mathcal{B}$ with equilibrium aggregate $A = A(\mathcal{B})$,*

$$C(\mathcal{B}) - C^D = \underbrace{\mathbb{E}[\psi(\varepsilon, A^D)] - \mathbb{E}[\psi(\mathbb{E}[\ell(\theta) \mid m], A^D)]}_{\text{composition cost}} + \underbrace{\mathbb{E}[\psi(\varepsilon, A)] - \mathbb{E}[\psi(\varepsilon, A^D)]}_{\text{spillover gain}}. \quad (13)$$

The composition cost is non-negative. The spillover gain is non-positive (it weakly lowers social cost) if and only if coarsening weakly raises the aggregate, $A \geq A^D$, which holds whenever $a^(\cdot, A^D)$ is concave over the range of detailed beliefs that each cell of $\mathcal{B}$ spans.*

Proof. Adding and subtracting $\mathbb{E}[\psi(\varepsilon, A^D)]$ gives (13). For the composition cost, the belief in each cell B is the average of the detailed beliefs over B and $\psi(\cdot, A^D)$ is concave, so by Jensen's inequality $\psi(\varepsilon, A^D) \geq \mathbb{E}[\psi(\mathbb{E}[\ell(\theta) \mid m], A^D) \mid m \in B]$; averaging over cells gives a non-negative term. For the spillover gain, $\psi(\varepsilon, \cdot)$ is strictly decreasing, so the term is non-positive if $A \geq A^D$ and positive if $A < A^D$. Finally, $\mathcal{T}(A') - A'$ is strictly decreasing in A' with its zero at A (Lemma 4), so $A \geq A^D$ if and only if $\mathcal{T}(A^D) \geq A^D$. If $a^*(\cdot, A^D)$ is concave over the range of detailed beliefs in each cell, Jensen's inequality applied cell by cell gives $\mathcal{T}(A^D) = \mathbb{E}[a^*(\varepsilon, A^D)] \geq \mathbb{E}[a^*(\mathbb{E}[\ell(\theta) \mid m], A^D)] = A^D$, where the equality is the fixed-point condition under detailed disclosure. $\square$

Proposition 5 is the continuous-adaptation counterpart of the proof of Proposition 1(ii). In that proof, lowering the threshold below m_D adds adapters at the boundary whose private

net gain is zero, so the term they generate vanishes and only the spillover channel survives. Continuous adaptation extends the reach of the composition cost. In the binary model $\psi(\cdot, A)$ is linear on each side of the adaptation hurdle, so pooling costs nothing unless a cell straddles the hurdle. With continuous adaptation $\psi(\cdot, A)$ is strictly concave wherever adaptation is interior (its slope in ε , $\phi(A)(1 - \alpha a^*)$, falls as a^* rises), so pooling distorts the intensity of every pooled property whose adaptation is interior. Quadratic cost makes this sharp. With $c(a) = a^2/2$ (the coefficient plays no role) and interior adaptation at every detailed belief, $a^*(\varepsilon, A^D) = \alpha \varepsilon \phi(A^D)$ is linear in beliefs, so every coarse rule leaves the equilibrium aggregate at A^D and yields no spillover gain, while $\psi(\varepsilon, A^D) = \varepsilon \phi(A^D) - (\alpha \varepsilon \phi(A^D))^2/2$ is strictly concave, so every coarse rule that pools a positive mass of scores strictly raises social cost. In particular, every threshold rule has strictly higher social cost than detailed disclosure.

Binding ceiling

Proposition 6 (Binding ceiling). *Suppose Assumptions 1, 3, and 5 hold with $\phi' < 0$, and that the belief ε_k at which the ceiling starts to bind, $\alpha \varepsilon_k \phi(A^D) = c'(\bar{a})$, is the detailed belief $\mathbb{E}[\ell(\theta) \mid m]$ of some score m in the interval of Assumption 3. Consider the rule $\mathcal{B}_\delta$ that pools all scores whose detailed belief is at least $\varepsilon_k - \delta$ into one high cell and reports every other score exactly. For small enough $\delta > 0$, this rule has strictly lower social cost than detailed disclosure.*

Proof. For small δ the belief in the high cell exceeds ε_k , because as $\delta \rightarrow 0$ it tends to the average of the detailed beliefs above ε_k . At A^D every property in the high cell then adapts at $\bar{a}$, so the rule changes intensities only for properties with detailed beliefs in $[\varepsilon_k - \delta, \varepsilon_k]$, which rise to $\bar{a}$. These properties have positive mass (Assumptions 1 and 3), so the average intensity at A^D rises by some $\Delta > 0$.

The composition cost in (13) is at most $\alpha \phi(A^D) \delta \Delta$. Pooling properties already at the ceiling costs nothing, because at a fixed intensity expected social cost is affine in the belief. A property with detailed belief $\varepsilon \in [\varepsilon_k - \delta, \varepsilon_k]$ that moves from a^* to $\bar{a}$ raises its expected social cost by $c(\bar{a}) - c(a^*) - \alpha \varepsilon \phi(A^D)(\bar{a} - a^*)$, which is at most $\alpha \phi(A^D) \delta (\bar{a} - a^*)$ because $c' \leq c'(\bar{a}) = \alpha \varepsilon_k \phi(A^D)$ on $[a^*, \bar{a}]$.

The spillover gain is at most a negative multiple of Δ that does not depend on δ . Because c'' is bounded away from zero on $[0, \bar{a}]$ and beliefs are bounded, intensities fall with the aggregate at a bounded rate, at most some K per unit. The rule's equilibrium aggregate A_δ exceeds A^D (Lemma 4), and it equals the rule's average intensity at A_δ , which is at most $K(A_\delta - A^D)$ below its value $A^D + \Delta$ at A^D ; hence $A_\delta - A^D \geq \Delta/(1 + K)$. As the aggregate rises, each property could keep its intensity, so its cost falls at a rate of at least $(1 - \alpha \bar{a}) \ell(0) \min_{[0,1]} |\phi'|$. The spillover gain is therefore at most $-(1 - \alpha \bar{a}) \ell(0) \min_{[0,1]} |\phi'| \Delta/(1 + K)$, while the composition cost is at most $\alpha \phi(A^D) \delta \Delta$, so their sum is negative for small enough δ . $\square$

This rule reveals the low region in detail. Whether the two-message threshold of the main text, which also pools the low region, dominates as well depends on the composition

cost of that lower cell, which this argument does not control. A binding ceiling is a natural assumption for disaster adaptation: defensible-space clearing, home retrofits, or elevation cannot reduce own losses by more than a technological maximum. The highest-risk properties bunch at that maximum.

More than two messages

Continuous adaptation also changes the number of messages a map needs. In the binary model, restricting attention to a single threshold is without loss: social cost depends on the partition only through the set of adapters, so any finer partition is replicated by one cutoff (Section 2). Here within-bin intensities respond to the bin posterior, so that argument fails. Pooling the top of the map lowers social cost when the ceiling binds (Proposition 6), while pooling beliefs over which $a^*(\cdot, A)$ is linear and increasing at the prevailing equilibrium aggregate A raises it, by the argument of the quadratic example above. In particular, a two-message rule whose low bin spans only such beliefs is strictly improved on by the rule that keeps its high bin and reports the scores in its low bin exactly. Characterizing the optimal partition is an information-design problem beyond the scope of this appendix.

G Heterogeneous adaptation costs

The baseline model gives every property the same adaptation cost c . This appendix lets the cost vary across properties, $c_i \sim H$. Under cost–score independence, heterogeneous costs are an exact relabeling of the continuous-adaptation model of Appendix F, with the share of individuals who adapt in place of the adaptation intensity, so the results of that appendix carry over once the two social costs are shown to coincide. The binary model of the main text is the limit in which H concentrates on a single cost.

Environment

I retain all primitives of Section 2.2 ($\theta \sim F$, ℓ increasing, $m \sim G$ satisfying Assumption 1, $\phi \in C^1$ with $\phi' < 0$ in the spillover case, $\alpha \in (0, 1)$) and binary adaptation $a_i \in \{0, 1\}$. I modify only the cost.

Assumption 6 (Cost heterogeneity). The costs $\{c_i\}$ are independent draws from a distribution H , independent of (θ, m) , whose density h is continuous and positive on its support $[\underline{c}, \bar{c}] \subset (0, \infty)$. Each individual privately observes his own c_i ; the planner knows H .

As in Appendix F, an individual’s decision depends on his information only through his posterior mean loss $\varepsilon \equiv \mathbb{E}[\ell(\theta) \mid \text{info}]$, which under a disclosure rule $\mathcal{B}$ is the average of the detailed beliefs $\mathbb{E}[\ell(\theta) \mid m]$ over the cell that contains his property’s score. An individual with belief ε and cost c adapts if and only if $\alpha \varepsilon \phi(A) \geq c$, so the share of individuals with belief ε who adapt is

$$q(\varepsilon, A) \equiv H(\alpha \varepsilon \phi(A)). \quad (14)$$

Reduction to the continuous-adaptation model

Lemma 5 (Equivalence with continuous adaptation). *Under Assumption 6, the per-property social cost of a disclosure rule $\mathcal{B}$ with equilibrium aggregate A is $\mathbb{E}[\psi(\varepsilon, A)]$, where ψ is the social cost (12) of the continuous-adaptation model with ceiling $\bar{a} = 1$ and cost function $c(a) = \int_0^a H^{-1}(s) ds$. This cost function satisfies Assumption 5, and its best response is $a^* = q$.*

Proof. Integrating the binary social cost over the cost draw and over $\theta \mid \text{info}$ (using independence and $\mathbb{E}[\ell(\theta) \mid \text{info}] = \varepsilon$) gives

$$\begin{aligned} \psi(\varepsilon, A) &= \int \min \{ \varepsilon \phi(A), (1 - \alpha) \varepsilon \phi(A) + c \} dH(c) \\ &= \phi(A)(1 - \alpha q(\varepsilon, A))\varepsilon + \int_0^{\alpha \varepsilon \phi(A)} c dH(c). \end{aligned} \quad (15)$$

The cost function has $c(0) = 0$ and marginal cost $c'(a) = H^{-1}(a) > 0$, the a -th quantile of the cost distribution, and $c''(a) = 1/h(H^{-1}(a))$ is continuous and positive on $[0, 1]$ by

Assumption 6, so it satisfies Assumption 5. Since $c'(0) = \underline{c}$ and $c'(1) = \bar{c}$, Lemma 3 gives $a^*(\varepsilon, A) = H(\alpha \varepsilon \phi(A)) = q(\varepsilon, A)$, and the minimized cost $(1 - \alpha q) \varepsilon \phi(A) + \int_0^q H^{-1}(s) ds$ equals the second line of (15) after the change of variable $s = H(c)$. $\square$

Lemmas 3 and 4 and Propositions 5 and 6 therefore apply with the headcount q in place of the intensity a^* . With $\phi' < 0$, the equilibrium is unique, coarsening carries a non-negative composition cost, and it weakly raises aggregate adaptation when $q(\cdot, A^D)$ is concave over the range of detailed beliefs that each cell spans. Because $q = H(\alpha \varepsilon \phi)$, this holds where the cost density is non-increasing (taking $h = 0$ outside $[\underline{c}, \bar{c}]$), the counterpart of a convex marginal cost in Appendix F.

Implications

The upper bound on costs plays the role of the ceiling: once $\alpha \varepsilon \phi(A) \geq \bar{c}$, every individual adapts whatever his cost draw, so the highest-risk properties all adapt.

Corollary 1 (Pooling the highest-risk properties). *Suppose Assumptions 1, 3, and 6 hold with $\phi' < 0$, and that the belief ε_k defined by $\alpha \varepsilon_k \phi(A^D) = \bar{c}$ is the detailed belief of some score in the interval of Assumption 3. For small enough $\delta > 0$, the rule $\mathcal{B}_\delta$ of Proposition 6 has strictly lower social cost than detailed disclosure.*

Proof. By Lemma 5, this is Proposition 6 for the continuous-adaptation model with $c'(a) = H^{-1}(a)$, whose ceiling level is $c'(\bar{a}) = H^{-1}(1) = \bar{c}$. $\square$

The corollary restricts the shape of H only through Assumption 6. With uniformly distributed costs, for instance, $q(\cdot, A^D)$ is linear where adoption is interior, so pooling such beliefs leaves aggregate adaptation unchanged, yet pooling the top of the map lowers social cost.


Remark 1 (Heterogeneous costs without spillovers). *If $\phi' \equiv 0$, the spillover gain vanishes, and a coarsening strictly raises social cost whenever it pools a positive mass of beliefs with different adoption rates, because the slope of ψ in ε , $\phi(1 - \alpha q)$, falls between them. The equality $C^T = C^D$ of Proposition 1(i) thus becomes $C^T > C^D$ whenever a positive mass of properties adopts at a rate strictly between zero and one under detailed disclosure, since every threshold rule then pools such beliefs.*

Remark 2 (Recovered weights). *The recovery of $\hat{\lambda}$ in Section 4.5 uses only the monotonicity of π and the observed cutoff, through equation (7), so it still describes where the map draws its line. The link to the private objective in Proposition 2(i) and the welfare readings in parts (ii) and (iii) rest on homogeneous costs, under which the line separates adapters from non-adapters; with dispersed costs, individuals in both bins adapt at rates set by their bin's posterior, and the fire probability at the line no longer summarizes the welfare trade-off.*


H Disaster risk communication in the G20 countries

Tables S1 to S3 list, for each G20 member and the European Union, the main public source of flood, wildfire, and hurricane risk information. The column “Disclosure form” records how the source releases risk. “Coarse” means the public product pools risk into mapped zones, discrete classes, or categorical alert levels. “Detailed” means it returns a continuous value or a property-level result. “Both” means a coarse public classification sits on top of accessible continuous values. Countries whose cited national source does not establish the disclosure form, for instance because it shows only a forecast track or a catalogue of studies, are omitted from the relevant table and listed in its notes. Figure S3 maps the disclosure form of each listed product by country and peril. Of the forty-five products listed, forty-three release a zoned or categorical classification, ten of them alongside continuous values, and two release continuous values only. Several sources are weather warning services or short-horizon danger forecasts rather than long-run risk maps; the description column says so where this applies. For federal countries, the description also notes whether state or provincial governments publish their own products, where this could be verified; in the United States, state wildfire hazard maps typically use three or four classes. Each classification was checked against the cited page or the agency’s own product documentation.


Floods



Wildfires



Hurricanes and typhoons



■ Coarse ■ Detailed ■ Both ■ Not established

Figure S3: Disclosure of public flood, wildfire, and hurricane risk information in G20 countries.

Notes: Each country is colored by the disclosure form of the product listed for it in Tables S1 to S3. Lighter shades mark European Union member states without their own row, colored by the European Union product (floods and wildfires). Darker grey marks G20 members omitted from a table because the cited source does not establish the disclosure form; light grey marks all other countries. France's hurricane row covers La Réunion, which is too small to show at this scale; French Guiana takes France's color in all panels. Country boundaries from Natural Earth at 1:110m scale.

Table S1: Public sources for flood risk information (G20 countries).

Country	Product and format	Disclosure form	Source
Brazil	Municipal susceptibility charts in three classes (high, medium, low) at 1:25,000	Coarse	Serviço Geológico do Brasil, 2025
Canada	Continuous national flood susceptibility index at 30 m resolution, not intended for street-level use. Provincial level: floodway and flood fringe zones with clickable depths (Alberta), four intensity classes (Quebec)	Detailed	Government of Canada, 2025
China	Historical national statistical flood maps at small scale (damage, affected population, frequency), not current hazard zoning	Coarse	International Knowledge Centre for Engineering Sciences and Technology, 2025
France	Regulatory flood zones (from flood risk prevention plans) and flood risk area maps in three hazard classes; address-level report returns the zone category	Coarse	Géorisques, 2025
Germany	Inundation extents for three scenarios (frequent, medium, extreme) plus mapped water depths. State level: the Länder produce the underlying maps, with depths continuous (North Rhine-Westphalia) or in five classes (Bavaria)	Both	Bundesanstalt für Gewässerkunde, 2025
India	Flood forecasting service: river gauge levels in metres with four colour flood categories (warning service, not a hazard map)	Both	India Water Resources Information System, 2025
Indonesia	InaRISK online hazard and risk layers in three classes (low, medium, high)	Coarse	Badan Nasional Penanggulangan Bencana, 2025 ^b
Italy	Flood hazard maps for three probability scenarios and four risk classes, with depths and velocities	Both	Autorità di Bacino Distrettuale del Fiume Po, 2025
Japan	Inundation assumption maps with depth class polygons; point queries return maximum depth and arrival time	Both	Geospatial Information Authority of Japan, 2025
Mexico	National Risk Atlas municipal flood hazard grades plus city-scale inundability indices	Coarse	Centro Nacional de Prevención de Desastres, 2025
Russia	Registered flood and inundation zone polygons in the public cadastral map; parcel lookup returns the zone name and restrictions	Coarse	Federal Service for State Registration, Cadastre and Cartography , Rosreestr
Saudi Arabia	Weather early warnings in three colour levels by region and governorate (warning service, no flood map)	Coarse	National Center for Meteorology, Saudi Arabia, 2025
South Africa	Impact-based severe weather warnings by colour and level over municipal areas (warning service, no flood map)	Coarse	South African Weather Service, 2025 ^b
South Korea	National flood risk map: inundation extent by return period with five depth classes	Coarse	Government of the Republic of Korea, 2025
Turkey	Province-level meteorological alerts in four colour levels (warning service, no flood map)	Coarse	Turkish State Meteorological Service, 2025

Table S1 (continued)

Country	Product and format	Disclosure form	Source
United Kingdom	Flood map for planning: zones 1, 2, 3a, 3b; detailed modelled data only on request	Coarse	Environment Agency, 2025
United States	FEMA flood zone letters at address level, with base flood elevations where studied. State level: some states add statewide modelled inundation rasters to the FEMA zones (Texas, Iowa)	Both	Federal Emergency Management Agency, 2025
European Union	Floods Directive viewer: polygons of areas of potentially significant flood risk, linking to national maps	Coarse	European Environment Agency, 2025

Notes: Argentina and Australia are omitted because the cited national sources (a research programme page and a catalogue of flood studies) do not establish the disclosure form. State-level products were verified for the United States, Canada, and Germany only.

Table S2: Public sources for wildfire risk information (G20 countries).

Country	Product and format	Disclosure form	Source
Argentina	Daily fire danger forecast (24 to 72 hours) in five categorical danger classes on a national map	Coarse	Servicio Nacional de Manejo del Fuego, 2025
Australia	Fire danger ratings issued per fire weather district in four categories; the underlying Fire Behaviour Index is continuous	Coarse	Australian Fire Danger Rating System, 2025
Brazil	Active fire hotspot monitoring plus a fire risk forecast in five classes on a roughly 25 km grid	Both	Instituto Nacional de Pesquisas Espaciais, 2025
Canada	Continuous Fire Weather Index system values mapped on national grids	Detailed	Canadian Wildland Fire Information System, 2025
China	Daily national forest fire danger forecast map (24 hours) delineating zones at the top three of five national danger levels	Coarse	China Meteorological Administration Public Meteorological Service Center, 2026
France	Regulatory wildfire hazard zoning with address-level and parcel-level risk reports returning the zone category	Both	Géorisques, 2025
Germany	Daily forest fire danger index shown in five discrete levels	Coarse	Deutscher Wetterdienst, 2025
India	Near real-time fire alerts plus a fire danger rating in five classes on a 5 km grid	Both	Forest Survey of India, 2025
Indonesia	InaRISK forest and land fire hazard layer in three classes (low, medium, high)	Coarse	Badan Nasional Penanggulangan Bencana, 2025a
Italy	Wildfire danger forecast in five levels for six regional macro-areas (regionally administered; Piedmont shown as representative example)	Coarse	Regione Piemonte, 2025

Table S2 (continued)

Country	Product and format	Disclosure form	Source
Japan	Dry-air advisory issued for public forecast areas, used here as a proxy; no wildfire-specific public risk product identified	Coarse	Japan Meteorological Agency, 2025 ^b
Russia	Weather-based fire danger in five classes reported by federal subject and mapped	Coarse	Avialesookhrana, 2025
South Africa	Fire danger index in five colour-coded classes defined over numeric index ranges	Coarse	South African Weather Service, 2025 ^a
South Korea	Fire danger forecast in four grades by district plus a numeric risk index from 1 to 100	Both	National Institute of Forest Science, 2025
Turkey	No national product cited; covered by the pan-European fire danger classes of the European Forest Fire Information System	Coarse	Copernicus Emergency Management Service, 2025
United Kingdom	Daily Fire Severity Index in five levels displayed on Natural England open-access mapping	Coarse	Met Office, 2025
United States	Monthly outlooks delineating areas of above, near, and below normal significant fire potential. State level: states typically publish coarse hazard maps in three or four classes (California, Oregon, Washington, Texas); Colorado also releases continuous 30 m burn probability	Coarse	National Interagency Fire Center, 2025
European Union	European Forest Fire Information System current situation viewer: six Fire Weather Index danger classes on an 8 to 10 km grid plus hotspot and perimeter polygons	Both	Copernicus Emergency Management Service, 2025

Notes: Mexico is omitted because the National Risk Atlas wildfire application could not be loaded at verification, so the disclosure form is not established. State-level products were verified for the United States only; Australia’s ratings are issued by state agencies under a national system, and Italy’s forecasts are regionally administered.

Table S3: Public sources for hurricane and typhoon risk information (G20 countries).

Country	Product and format	Disclosure form	Source
Australia	Cyclone watch and warning zones defined as coastal segments between named towns, plus intensity categories 1 to 5. State level: state agencies issue colour-coded cyclone warnings in three levels (Western Australia)	Coarse	Australian Bureau of Meteorology, 2025
China	Typhoon warnings in four colour levels issued for named coastal seas and provinces	Coarse	National Meteorological Center, 2025
France	Vigilance map for La Réunion with ordered cyclone alert levels (pre-alert, orange, red, violet) applied to all or part of the island	Coarse	Météo-France, 2025

Table S3 (continued)

Country	Product and format	Disclosure form	Source
India	Four-stage cyclone warnings and colour codes mapped by district and subdivision. State level: state disaster authorities publish static cyclone-prone-area maps (Odisha)	Coarse	India Meteorological Department, 2025
Indonesia	Tropical cyclone outlook: suspect areas rated small, medium, or large development possibility within a defined region	Coarse	BMKG, 2025
Japan	Municipality-level typhoon warnings, typhoon intensity and size classes, and a 70 percent probability circle for the track	Coarse	Japan Meteorological Agency, 2025a
Mexico	National Risk Atlas municipal-level hazard and vulnerability degrees, including tropical cyclone layers. State level: state risk atlases exist but their cyclone layers were not verified	Coarse	Centro Nacional de Prevención de Desastres, 2025
Saudi Arabia	Weather early warnings in three colour levels by region and governorate (warning service, no cyclone-specific product)	Coarse	National Center for Meteorology, Saudi Arabia, 2025
United States	Coastal and inland watch and warning zones set by breakpoints, plus the forecast track cone. State level: coastal states assign addresses to lettered evacuation zones (Florida, North Carolina, South Carolina)	Coarse	NOAA National Hurricane Center, 2025

Notes: Canada and South Korea are omitted because the cited national sources show forecast tracks without a zoned or categorical classification, so the disclosure form is not established. State-level products were verified for the United States, Australia, and India only.

I Comparing the three risk signals

This appendix compares the distribution and the spatial pattern of the three risk signals of Section 4.2 (Figures S4 to S6). It then repeats Panel A of Figure 4 over the full distribution of each signal (Figure S7) and with the held-out window starting in 2023 and in 2024 (Figure S8).




Figure S4: The Oregon map vs the FSF model

Notes: Distribution of the FSF model's wildfire burn probability within each of the Oregon map's risk categories ("No risk" in the top row, to "Extreme" risk in the bottom row). For clarity, the mass of properties with a burn probability of 0 is depicted separately on the left for each Oregon map category, and burn probabilities are top-coded at 5%.

Figure S8 addresses the vintage mismatch noted in Section 4.2: the First Street model was published in July 2024, after part of the held-out window. Its ignition density draws on the 2022 edition of the Fire Occurrence Database, and its fuel layers incorporate wildfires through 2023 as fuel disturbances (First Street, 2024), so fires in 2022 and 2023 are not entirely out of sample for that signal. Each panel repeats Panel A of Figure 4 with a later window start, dropping from the sample any parcel burned before that start, so Panel C uses only fires the First Street model could not have seen. The three signals keep their ordering in every window.




Figure S5: The Oregon map vs the Cotality model

Notes: Distribution of the Cotality model's wildfire average annual loss (AAL) within each of the Oregon map's risk categories ("No risk" in the top row, to "Extreme" risk in the bottom row). For clarity, the mass of properties with an AAL of 0 is depicted separately on the left for each Oregon map category, and AAL is top-coded at 0.2% of insurable value.




Figure S6: Property-level risk quintiles under the three models

Notes: Each panel maps the 1.86 million matched Oregon properties, colored by the within-model quintile of their risk score: the Oregon model (m) in Panel A, the Cotality model (s_C) in Panel B, and the FSF model (burn probability) in Panel C. Quintiles are computed from percentile ranks with tied scores assigned to the lowest group, so properties with identical scores receive identical colors; because each model places a large mass of properties at exactly zero (24% for Oregon, 15% for Cotality, 43% for FSF), the bottom group can exceed 20% of properties. For the FSF model the zero mass extends past the 40th percentile, so no property falls in the second quintile. Colors run from the lowest quintile (Q1) to the highest (Q5), and points are plotted in ascending quintile order, so where properties of adjacent parcels fall in different quintiles the higher-risk color is drawn on top. In densely developed areas this can visually mask the low-risk (Q1) properties beneath, most notably the zero-risk mass in the Willamette Valley in Panel A.




Figure S7: Empirical test of Assumption 1: held-out burn rate by signal vigintile.

Notes: Held-out 2022–2025 burn rate at each vigintile of the Oregon model (m), the Cotality model (s_C), and the FSF model (burn probability s_F) on the clean OOS sample, with 95% fire-cluster bootstrap bands ($B = 1,000$; cluster = 2022–2025 NIFC fire perimeter for burned parcels, 5km grid cell otherwise). Parcels with tied signal values at the bottom of each distribution fall into a single vigintile, so the three series start at different ranks. The same cluster bootstrap does not reject the null hypothesis that the burn rate is weakly increasing across vigintiles ($p = 0.16$ for m , 0.17 for s_C , and 0.73 for s_F).




Figure S8: Calibration in the policy-relevant region under later held-out windows.

Notes: Held-out NIFC burn rates by signal rank percentile in the 75th–99th band, as in Panel A of Figure 4, for windows starting in 2022 (Panel A, 3,673 burned parcels), 2023 (Panel B, 3,404) and 2024 (Panel C, 2,668). Each window’s sample drops parcels burned in any earlier year, so Panel A reproduces Panel A of Figure 4 exactly. The black curve is the monotone P-spline calibration $\hat{\pi}(m)$ refit on each window and mapped to rank space, with the smoothing parameter held at the value of the 2022 window; the shaded band is its 95% fire-cluster bootstrap band ($B = 1,000$; cluster = held-out fire perimeter for burned parcels, 5km grid cell otherwise). Burn rates in the 75th–95th band are 0.69, 0.67 and 0.48% for Oregon, 0.64, 0.59 and 0.42% for First Street, and 0.24, 0.21 and 0.16% for Cotality. The y -axis is capped per panel; Cotality’s 99th-percentile bin reaches 3.3, 3.1 and 2.9% and lies above the displayed range in Panel C. Dashed line: Oregon’s Moderate/High cutoff $\hat{m}$ in rank space.

J Robustness to the fire record

The realized-fire outcomes of Section 4 come from NIFC interagency fire perimeters: operational polygons reported by the managing agencies for incidents of every size. This appendix repeats the calibration and the recovery of Section 4.5 with an independent record of burning, the burned-area boundaries of the Monitoring Trends in Burn Severity (MTBS) program (Eidenshink et al., 2007). MTBS maps the burned extent of each fire from Landsat imagery, so its polygons measure where vegetation burned rather than where an incident was contained, but it covers only fires of at least 1,000 acres in the western United States and is released with a lag: its last complete year in Oregon is 2024, with a single polygon recorded for 2025. Figure S9 maps the two records over 2000–2025, the NIFC perimeters in Panel A and the MTBS burned-area boundaries in Panel B. The two records disagree in both directions over the pre-map window: 4,148 parcels burned under MTBS but under no NIFC perimeter, and 2,165 the reverse, so 89% of the NIFC-burned parcels are also MTBS-burned and 81% of the MTBS-burned parcels are also NIFC-burned. Prescribed burns remain in both records, as in the main text.




Figure S9: Wildfires in Oregon under the two fire records, 2000–2025

Notes: Panel A maps the 5,236 NIFC fire perimeters and Panel B the 800 MTBS burned-area boundaries recorded in 2000–2025, each colored by fire year. Later fires are drawn above earlier ones, so where perimeters overlap the most recent fire is visible.

Because MTBS differs from the NIFC record in its size floor and its end year as well as in its source, Table S4 reports three definitions of the realized-fire outcome. The first is the headline. The second applies MTBS’s restrictions to the NIFC record itself, keeping only perimeters whose polygon area is at least 1,000 acres and ending the held-out window in

2024, so that the gap between the first two columns reflects the size floor and the shorter window while the gap between the last two reflects the source. Everything else is held fixed across columns: the parcels, the score m , the cutoff $\hat{m}$, the calibration spline, and the bootstrap design. The one change to the bootstrap is that the fire clusters are rebuilt from each record’s own polygons, so the interval in each column resamples the footprints of the record it evaluates.

	NIFC (headline)	NIFC, 1,000+ acres	MTBS
<i>Panel A. Fire record</i>			
Record type	Perimeters	Perimeters	Satellite burn extent
Size floor	None	1,000 acres	1,000 acres
Polygons, 2000–2021	3,829	830	717
Parcels burned, 2000–2021	20,386 (1.10%)	19,222 (1.03%)	22,369 (1.20%)
Held-out window	2022–2025	2022–2024	2022–2024
Polygons, held-out window	1,407	131	83
Parcels burned, held-out (clean sample)	3,673 (0.20%)	2,924 (0.16%)	2,546 (0.14%)
Bootstrap clusters, pre-map	8,827	8,596	8,623
<i>Panel B. Recovery at Oregon’s cutoff $\hat{m}$</i>			
$\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$	0.0375	0.0350	0.0395
$\hat{\lambda} = \hat{\kappa}/(1 - \hat{\kappa})$	0.0389	0.0363	0.0411
95% CI	[0.0257, 0.0590]	[0.0226, 0.0559]	[0.0245, 0.0662]
Δ vs. headline	—	-6.9%	+5.7%
$\hat{\lambda}_{\text{pre}}^{\text{ann}}$ (1-yr basis)	0.00174	0.00162	0.00183
95% CI	[0.00116, 0.00261]	[0.00102, 0.00248]	[0.00110, 0.00292]

Table S4: The recovered weight $\hat{\lambda}$ under three definitions of the realized-fire outcome.

Notes: Panel A describes each fire record on the working sample of 1,860,526 parcels; the held-out counts are on each record’s clean sample, the parcels with no pre-map fire under that record, and their levels are comparable only between the two 2022–2024 columns. Panel B recovers $\hat{\lambda} = \hat{\kappa}/(1 - \hat{\kappa})$ with $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$ at the fixed Moderate/High cutoff $\hat{m}$, where $\hat{\pi}_{\text{pre}}$ is the monotone P-spline calibration of Section 4.5 fit to that column’s 2000–2021 outcome on the full sample, with the smoothing parameter selected automatically and checked against a fixed-penalty grid as in the headline. The 95% confidence intervals are cluster bootstraps ($B = 1,000$ replicates, smoothing parameter held at the full-sample value) over each record’s own pre-map fire polygons and 5km-grid cells (the cluster counts in Panel A). The annualized rows apply the constant-hazard conversion of Section 4.5 to the 22-year window; the Δ row reports the percent deviation of each point estimate from the headline.

Panel B of Table S4 reads the recovered weight off each calibration at Oregon’s cutoff. Restricting the NIFC record to large perimeters moves $\hat{\lambda}$ by -6.9%, to 0.036, and replacing those perimeters with the MTBS burn extents moves it by +5.7%, to 0.041; read across the two 1,000-acre columns, the source alone accounts for +13.5%. Both alternatives lie inside the headline confidence interval, and the annualized recoveries move by nearly the same proportions. The implied trade-off ranges from 24 to 28 unburned parcels classified High per parcel that burns and is caught, against 26 in the headline.

K Robustness to the calibration and the cutoff

Table S5 re-runs the recovery of Section 4.5 under a battery of alternative specifications. The headline calibration $\hat{\pi}_{\text{pre}}(m)$ is a monotone-increasing penalized spline, a shape-constrained P-spline (`scam`, basis `bs="mpi"`, dimension $k = 10$) fit by penalized likelihood to the 100 equal- N centile burn rates on the $\sqrt{m}$ axis, with the effective smoothness selected by the penalty rather than fixed. The monotonicity constraint enforces the weakly increasing calibration that the cutoff-in- m / cutoff-in- π equivalence (Section 3.2) relies on, and it removes the spurious downturn in the sparse upper tail that an unconstrained fit produces. From this calibration, $\hat{\lambda} = \hat{\pi}_{\text{pre}}(\hat{m}) / (1 - \hat{\pi}_{\text{pre}}(\hat{m}))$ is read off at the Moderate/High cutoff $\hat{m}$, with a cluster bootstrap over fire-perimeter and 5km-grid clusters for the confidence interval. Each row of Table S5 changes one ingredient of the calibration form and re-recovers $\hat{\lambda}$; the headline row reproduces the point estimate of Table 3 exactly. Appendix J instead varies the realized-fire record itself.

Calibration form	$\hat{\kappa}$	$\hat{\lambda}$	95% CI	Δ vs. headline
Monotone P-spline, $k = 10$, $\sqrt{m}$ axis (headline)	0.0375	0.0389	[0.0257, 0.0590]	—
Vigintile-mean step function	0.0265	0.0272	[0.0183, 0.0398]	-30.0%
Monotone P-spline, $k = 6$	0.0414	0.0432	[0.0301, 0.0619]	+10.9%
Monotone P-spline, $k = 15$	0.0343	0.0355	[0.0253, 0.0547]	-8.8%
Monotone P-spline, raw- m axis	0.0319	0.0330	[0.0247, 0.0439]	-15.3%
Unconstrained P-spline (no monotonicity)	0.0359	0.0373	[0.0250, 0.0537]	-4.2%

Table S5: Robustness of the recovered false-positive/false-negative weight $\hat{\lambda}$.

Notes: Each row recovers $\hat{\lambda} = \hat{\kappa} / (1 - \hat{\kappa})$ with $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$ at the fixed Moderate/High cutoff $\hat{m}$, changing one ingredient of the calibration form at a time relative to the headline (first row): a nonparametric vigintile-mean step function, the spline basis dimension k , the raw- m versus $\sqrt{m}$ axis, and dropping the monotonicity constraint. The 95% confidence interval is a cluster bootstrap ($B = 1,000$ replicates over 8,827 fire-perimeter and 5km-grid clusters, with each spline specification’s smoothing parameter held at its full-sample value); the last column reports the percent deviation of the point estimate from the headline.

A fully nonparametric vigintile-mean step function, the spline basis dimension swept from $k = 6$ to $k = 15$, and the raw- m axis all deliver $\hat{\lambda}$ within roughly a third of the headline, with the headline monotone spline in the upper part of that range. Dropping the monotonicity constraint changes $\hat{\lambda}$ at the cutoff by about four percent: the constraint disciplines the fit in the sparse tail and has little effect at the policy-relevant boundary.

Figure S10 plots the recovery input $\hat{\pi}_{\text{pre}}$ directly, against the empirical pre-map (2000–2021) burn rates it is fit to, together with the analogously constructed pre-map calibrations of the two proprietary signals. The bin-level burn rates are noisy and locally non-monotone because rare fires arrive as spatially contiguous events; a single fire complex can dominate several adjacent rank bins. The shaded bands are therefore 95% cluster-bootstrap confidence bands for each fitted calibration, refitting the monotone spline within each replicate, with its smoothing parameter held at the full-sample value, under the same fire-perimeter/5km-

grid resampling as the $\hat{\lambda}$ confidence interval; the Oregon band evaluated at $\hat{m}$ reproduces the $\hat{\kappa}$ interval behind Table 3. Against that clustered uncertainty the data are consistent with the weakly increasing calibration the recovery requires: a cluster-bootstrap test of weak monotonicity, based on the distance between the bin rates and their isotonic projection, gives $p = 0.23$ for the Oregon score and $p = 0.57$ for Cotality. Through the policy-relevant region around the Moderate/High cutoff (rank ≈ 0.90) the three calibrations sit at comparable levels, and the recovered $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$ (black diamond), read where the Oregon calibration crosses the cutoff, sits close to the raw burn rate of 3.3% within four rank percentiles of the cutoff. Cotality turns up steeply only in the extreme upper tail, the same tail advantage visible out of sample in Figure 4.




Figure S10: Pre-map calibration $\hat{\pi}_{\text{pre}}$ against 2000–2021 fires, three models.

Notes: Pre-map (2000–2021) calibration $\Pr(\text{fire} \mid \text{signal})$ for the three risk models, on a common within-signal rank axis. Points are 50 equal- N rank-bin burn rates; curves are monotone-increasing P-spline logits (`scam`, `bs="mpi"`, $k = 10$) fit on $\sqrt{\text{signal}}$, identical in construction across the three models, with 95% fire-cluster bootstrap confidence bands obtained by refitting the monotone spline in each replicate with its smoothing parameter held at the full-sample value ($B = 1,000$; cluster = 2000–2021 NIFC fire perimeter for burned parcels, 5km grid cell otherwise). The Oregon curve is the saved calibration $\hat{\pi}_{\text{pre}}$ that recovers $\hat{\lambda}$ in Section 4.5, reused here unchanged. Dashed line: Oregon’s Moderate/High cutoff (rank ≈ 0.90); black diamond: $\hat{\kappa} = \hat{\pi}_{\text{pre}}(\hat{m})$, the recovery input. The y -axis is capped at 8%; Cotality’s top-centile pre-map burn rate reaches roughly 10.5% (off-scale), the extreme-tail advantage discussed in the text.

Across all 6 specifications $\hat{\lambda}$ lies between 0.027 and 0.043, so the implied trade-off $1/\hat{\lambda}$ ranges from roughly twenty-three to thirty-seven unburned parcels classified High per parcel that burns and is caught. The cutoff $\hat{m}$ is the SB 762 Moderate/High regulatory boundary,

and the boundary is near-sharp: pinning it at the lower edge $\max(m \mid \text{Moderate})$ or the upper edge $\min(m \mid \text{High})$ instead of the midpoint leaves $\hat{\lambda}$ at 0.039 and 0.039 respectively, unchanged to three decimals.

Realized fires could in principle depend on adaptation undertaken in response to the 2022 map itself. The headline $\hat{\lambda}$ is calibrated on pre-map (2000–2021) fires, which a map drawn in 2022 cannot have affected. The concern therefore bites only on the held-out 2022–2025 consistency check $\hat{\lambda}_{\text{OOS}}$. Given that the map was repealed a few weeks after its disclosure, adaptation responses are unlikely.

L Calibration of the private benchmark

Table S6 consolidates the parameter ranges behind the calibration of the private cutoff $\kappa^{\text{priv}} = c/(\alpha \phi(A) \bar{L})$ in Section 4.7, with the source for each range.

	Low	Central	High	Source
Defensible space, initial (\$)	1,160	3,470	4,730	California Legislative Analyst’s Office (2021); industry
Defensible space, maint. (\$/yr)	95	284	568	Industry quotes
Ember hardening, upfront (\$)	10,300	12,900	15,400	Barrett and Quarles (2024)
Component life (years)	10 (space), 30 (hardening)			Assumption
α , defensible space only	0.15	0.25	0.35	Syphard, Brennan and Keeley (2014)
α , space + hardening	0.30	0.45	0.60	Knapp et al. (2021); Zamanialaei et al. (2025)
$\bar{L}$, per-event loss (\$)	290,000	435,000	476,000	ACS; FHFA; Zillow (see notes)
Real discount rate r	2%	3%	5%	Office of Management and Budget (2023)

Table S6: Calibrated ranges for the primitives of the private cutoff κ^{priv} .

Notes: Primitives of the private benchmark $\kappa^{\text{priv,ann}} = c^{\text{ann}}/(\alpha \bar{L})$, the annual counterpart of the private cutoff κ^{priv} evaluated at $\phi = 1$; because $\phi(A) \leq 1 = \phi(0)$, this choice minimizes κ^{priv} and makes the comparison with $\hat{\kappa}^{\text{ann}}$ conservative. Upfront costs are converted to annual flows as an annuity at real rate r over the component life and added to annual maintenance; the recovered cutoff is annualized with the constant-hazard conversion of Section 4.5. Cost figures are for a 2,000 square-foot home. All dollar values are expressed in 2024 dollars: each input is converted from the price year of its source with the annual average of the all-items consumer price index for urban consumers, and the common base year cancels in the ratio $c^{\text{ann}}/(\alpha \bar{L})$. The source price years are 2021 for the defensible-space program benchmarks, 2026 for the industry quotes, 2023 for the retrofit estimates (their RSMeans cost basis), 2022 for the median home value (the final year of the American Community Survey (ACS) period), and 2026 for the Zillow index. Effectiveness α is the share of expected loss eliminated conditional on fire reaching the parcel and is bundle-specific: the grid never pairs combined-bundle effectiveness with defensible-space-only costs. The effectiveness ranges extend the published estimates (16–31% for defensible space alone, 31–52% for the combined bundle) at both ends. $\bar{L}$: the low value is the structure-only share of the median Oregon owner-occupied home value (\$423,100 in 2022 dollars, ACS 2018–2022, table B25077), using one minus Oregon’s land share of 0.361 from the June 2024 release of the Federal Housing Finance Agency (FHFA) residential land-price data of Davis et al. (2021); the central value adds contents at the homeowners-insurance convention of roughly half of dwelling coverage; the high value is the full property value (Zillow Home Value Index for Oregon, accessed July 2026). Defensible-space maintenance costs are contractor-quote industry figures rather than government or academic estimates, and the component lives are stated assumptions.

References

- Australian Bureau of Meteorology.** 2025. “Tropical cyclone information.” <https://www.bom.gov.au/weather-and-climate/specialised-forecasts-and-observations/tropical-cyclone>, Accessed 2025-12-01.
- Australian Fire Danger Rating System.** 2025. “Fire Danger Ratings map.” <https://fdv.afdrs.org.au/>, Accessed 2025-12-01.
- Autorità di Bacino Distrettuale del Fiume Po.** 2025. “PGRA maps (flood hazard and risk maps).” <https://pianoalluvioni.adbpo.it/pgra-mappe/>, Accessed 2025-12-01.
- Avialesookhrana.** 2025. “Meteoinformation (fire danger and related products).” <https://aviales.ru/default.aspx?textpage=63>, Accessed 2025-12-01.
- Badan Nasional Penanggulangan Bencana.** 2025a. “InaRISK (portal).” <https://inarisk.bnpb.go.id/>, Accessed 2025-12-01.
- Badan Nasional Penanggulangan Bencana.** 2025b. “InaRISK WebGIS.” <https://inarisk.bnpb.go.id/webgis/>, Accessed 2025-12-01.
- Barrett, Kimiko, and Stephen L. Quarles.** 2024. “Retrofitting a Home for Wildfire Resistance: Costs and Considerations.” Headwaters Economics Report. Accessed 2026-07-05.
- BMKG.** 2025. “Tropical cyclone outlook.” <https://tropicalcyclone.bmkg.go.id/content/tropical-cyclone-outlook>, Accessed 2025-12-01.
- Bundesanstalt für Gewässerkunde.** 2025. “HWRM maps (flood risk management maps viewer).” https://geoportal.bafg.de/karten/HWRM_Aktuell, Accessed 2025-12-01.
- California Legislative Analyst’s Office.** 2021. “Reducing the Destructiveness of Wildfires: Promoting Defensible Space in California.” California Legislative Analyst’s Office Report. Accessed 2026-07-05.
- Canadian Wildland Fire Information System.** 2025. “Fire weather maps.” <https://cwfis.cfs.nrcan.gc.ca/maps/fw>, Accessed 2025-12-01.
- Centro Nacional de Prevención de Desastres.** 2025. “Atlas Nacional de Riesgos.” <https://www.atlasnacionalderiesgos.gob.mx/>, Accessed 2025-12-01.
- China Meteorological Administration Public Meteorological Service Center.** 2026. “National forest fire danger weather forecast map.” <http://www.nmc.cn/publish/environment/forestfire-doc.html>, Accessed 2026-09-24. Danger levels defined in national standard GB/T 36743-2018.

- Copernicus Emergency Management Service.** 2025. “European Forest Fire Information System (EFFIS): current situation.” <https://forest-fire.emergency.copernicus.eu/apps/effis.csv/>, Accessed 2026-09-03.
- Davis, Morris A., William D. Larson, Stephen D. Oliner, and Jessica Shui.** 2021. “The price of residential land for counties, ZIP codes, and census tracts in the United States.” *Journal of Monetary Economics*, 118: 413–431.
- Deutscher Wetterdienst.** 2025. “Waldbrandgefahrenindex (forest fire danger index) map product.” https://www.dwd-geoportal.de/products/G_EBM/, Accessed 2025-12-01.
- Eidenshink, Jeff, Brian Schwind, Ken Brewer, Zhi-Liang Zhu, Brad Quayle, and Stephen Howard.** 2007. “A project for monitoring trends in burn severity.” *Fire Ecology*, 3(1): 3–21.
- Environment Agency.** 2025. “Flood map for planning.” <https://flood-map-for-planning.service.gov.uk/>, Accessed 2025-12-01.
- European Environment Agency.** 2025. “Floods viewer (EU Floods Directive map viewer).” <https://discomap.eea.europa.eu/floodviewer/>, Accessed 2025-12-01.
- Federal Emergency Management Agency.** 2025. “Flood Map Service Center (MSC).” <https://msc.fema.gov/portal/search>, Accessed 2025-12-01.
- Federal Service for State Registration, Cadastre and Cartography (Rosreestr).** 2025. “Public Cadastral Map (Publichnaya Kadastrovaya Karta).” <https://pkk.rosreestr.ru/>, Accessed 2025-12-01. Flood and inundation zones are accessible as “zones with special conditions of use” (ZUIT) layers in the national cadastral viewer.
- First Street.** 2024. “First Street Wildfire Model (FS-WFM) Technical Methodology, Version 3.0.” First Street. Published 2024-07-31. Version 3.0 fuels, ignition, and calibration updates dated 2024-08-01. TODO: confirm the cover title against the PDF (cover is an image).
- Forest Survey of India.** 2025. “Forest fire portal.” <https://fsiforestfire.gov.in/>, Accessed 2025-12-01.
- Géorisques.** 2025. “Cartes interactives (aléas et risques).” <https://www.georisques.gouv.fr/>, Accessed 2026-09-03.
- Geospatial Information Authority of Japan.** 2025. “Suibou Map (flood inundation assumption map).” <https://suiboumap.gsi.go.jp/en/index.html>, Accessed 2025-12-01.

- Government of Canada.** 2025. “Flood susceptibility index.” <https://geo.ca/flood-mapping/flood-map-gallery/flood-susceptibility-index/>, Accessed 2026-09-03.
- Government of the Republic of Korea.** 2025. “Flood risk map portal (Hongso Wiheom Jido).” <https://www.floodmap.go.kr/>, Accessed 2025-12-01.
- India Meteorological Department.** 2025. “Mausam (public weather portal).” <https://mausam.imd.gov.in/>, Accessed 2025-12-01.
- India Water Resources Information System.** 2025. “Flood Forecasting and Warning System (FFS).” <https://ffs.india-water.gov.in/>, Accessed 2025-12-01.
- Instituto Nacional de Pesquisas Espaciais.** 2025. “Programa Queimadas (portal and maps).” <https://terrabrasilis.dpi.inpe.br/queimadas/portal/>, Accessed 2025-12-01.
- International Knowledge Centre for Engineering Sciences and Technology.** 2025. “Flood disaster risk resources (portal listing).” https://ikcest-drr.data.ac.cn/list/dis_flood_disasters/10, Accessed 2025-12-01.
- Japan Meteorological Agency.** 2025a. “Tropical cyclone information.” <https://www.data.jma.go.jp/multi/cyclone/index.html?lang=en>, Accessed 2025-12-01.
- Japan Meteorological Agency.** 2025b. “Weather warning / advisory (map interface).” <https://www.data.jma.go.jp/multi/warn/index.html?lang=en>, Accessed 2025-12-01.
- Knapp, Eric E., Yana S. Valachovic, Stephen L. Quarles, and Nels G. Johnson.** 2021. “Housing arrangement and vegetation factors associated with single-family home survival in the 2018 Camp Fire, California.” *Fire Ecology*, 17(1): 25.
- Météo-France.** 2025. “Vigilance: La Réunion.” <https://vigilance.meteofrance.fr/fr/la-reunion>, Accessed 2025-12-01.
- Met Office.** 2025. “Fire Severity Index (government service page).” <https://www.metoffice.gov.uk/services/government/environmental-hazard-resilience/fire-severity-index>, Accessed 2025-12-01.
- National Center for Meteorology, Saudi Arabia.** 2025. “Early warnings and weather alerts.” <https://www.ncm.gov.sa>, Accessed 2025-12-01.
- National Institute of Forest Science.** 2025. “Forest fire information service.” <https://forestfire.nifos.go.kr/>, Accessed 2025-12-01.

- National Interagency Fire Center.** 2025. “Predictive Services outlooks.” *https://www.nifc.gov/nicc/predictive-services/outlooks*, Accessed 2025-12-01.
- National Meteorological Center.** 2025. “Typhoon information.” *https://eng.nmc.cn/typhoon/*, Accessed 2025-12-01.
- NOAA National Hurricane Center.** 2025. “Hurricane graphics (forecast cone and related products).” *https://www.nhc.noaa.gov/graphics_at5.shtml?start=*, Accessed 2025-12-01.
- Office of Management and Budget.** 2023. “Circular A-4: Regulatory Analysis.” Office of Management and Budget, Executive Office of the President Circular. Accessed 2026-07-05.
- Regione Piemonte.** 2025. “Wildfire danger forecasting and evaluation.” *https://www.regione.piemonte.it/web/temi/protezione-civile-difesa-suolo-opere-pubbliche/protezione-civile/incendi-boschivi/previsione-valutazione-pericolo-incendi*, Accessed 2025-12-01. Italy’s wildfire warning systems are managed at the regional level; Piedmont (Piemonte) is shown as a representative example.
- Servicio Nacional de Manejo del Fuego.** 2025. “Fire danger map (evaluación de peligro y alerta temprana).” *https://www.argentina.gob.ar/seguridad/servicio-nacional-de-manejo-del-fuego/evaluacion-de-peligro-y-alerta-temprana/mapa-de*, Accessed 2025-12-01.
- Serviço Geológico do Brasil.** 2025. “Geoportal de desastres.” *https://geoportal.sgb.gov.br/desastres/*, Accessed 2025-12-01.
- South African Weather Service.** 2025a. “Fire danger indicator.” *https://www.weathersa.co.za/home/fireindex*, Accessed 2025-12-01.
- South African Weather Service.** 2025b. “Warnings.” *https://www.weathersa.co.za/home/warnings*, Accessed 2025-12-01.
- Syphard, Alexandra D., Teresa J. Brennan, and Jon E. Keeley.** 2014. “The role of defensible space for residential structure protection during wildfires.” *International Journal of Wildland Fire*, 23(8): 1165–1175.
- Turkish State Meteorological Service.** 2025. “Meteorological warnings map (Turkey).” *https://www.mgm.gov.tr/metouyari/turkiye.aspx*, Accessed 2025-12-01.
- Zamanialaei, Maryam, Daniel San Martin, Maria Theodori, Dwi Marhaendro Jati Purnomo, Ali Tohidi, Chris Lautenberger, Yiren Qin, Arnaud Trouvé, and Michael Gollner.** 2025. “Fire risk to structures in California’s Wildland-Urban Interface.” *Nature Communications*, 16(1): 8041.