Wildfire Risk Durably Falls After Timber Harvests:
Evidence from 40 Years of Harvest and Fire in Canada

Jeffrey E. Sun*

Preliminary draft. All errors are my own.

Link to Latest Version

July 24, 2026

Abstract

I estimate the association between timber harvest and subsequent wildfire risk nationally across Canada's managed forest using a staggered-adoption event study on 447935 thirty-meter stands which cover the universe of Canada's managed forest. Smaller studies have shown mixed results, often that harvested stands burn more, but they are typically confined to a single region and to plot- or severity-scale ecological evidence. This is the first national, continental-scale quasi-experimental estimate of the harvest–fire relationship in Canada.Harvested stands face a first-fire hazard about 28% lower than comparable unharvested neighbors. The protective effect is persistent to at least 29 years after harvest. This paper's preferred specification controls by strictly pre-treatment siting and terrain covariates, ecozone-by-year and fireshed fixed effects. The result is robust to clustering scheme, to dropping the extreme 2023 season, and to choice of covariates up to dropping all covariates: the raw association is already protective and attenuates under adjustment. The result highlights the potential benefits of co-management of timber harvest and wildfire risk, and suggests that the current separate-management regime in Canada may be imposing very large economic, environmental, and human costs.

1 Introduction

In the record 2023 Canadian fire season, wildfires in Canada released approximately 2.4 billion tonnes of CO_2 into the atmosphere: roughly 6% of global fossil-fuel emissions, and more than four times the emissions from all economic activity in Canada combined in that year (Byrne et al., 2024). While these are gross biogenic emissions, partially reabsorbed by regrowth over subsequent years, and the 2023 season does not represent recurring annual emissions, the scale of the externality is nevertheless enormous.

The land that burns is, to a large extent, the same land that Canada manages for timber. Yet in every province, forest and fire management are administered by operationally independent

*Department of Economics, University of Toronto

and often largely non-coordinating agencies.[1] Wildfire suppression is conducted by provincial and territorial agencies, while timber production is planned decades ahead through licensed tenures, often implemented by non-governmental forest management companies (Tymstra et al., 2020).

The potential for timber harvest to raise or lower a stand's subsequent fire risk represents one of potentially many reasons why separate management may impose large economic, environmental, and human costs. The empirical question of the direction of the harvest-fire relationship is not settled.

I estimate the effect of past timber harvest on a stand's subsequent *first-fire hazard*, nationally across Canada's managed forest. This presents two major challenges, which my design addresses to a greater extent than previous studies (though completely solving these may be impossible). First, harvest location is endogenous, and depends on access, stand age, and operability, all of which may be correlated with fire risk. Second, burned areas are endogenous to fire suppression efforts, which are not only targeted, but targeted partly at merchantable timber.

I employ a staggered-adoption discrete-time (complementary log-log) hazard event design fit to the national managed-forest panel. I use separate datasets for treatment and outcome. This avoids the mechanical correlation between the two due to the shared classifier used to detect both harvest and fire.My outcome data is an independent record of fire perimeters;the covariates are pre-treatment forest structure and terrain, measured before the harvest year; the fixed effects are additive ecozone-by-year and coarsened-fireshed absorbers; the controls are never- and not-yet-treated stands; and inference is clustered on fire-years against a $t(G-1)$ reference. The ecozone-by-year fixed effects control for the common fire weather of a given season in a given biome (e.g. 2023 in the boreal). The fireshed fixed effects are spatial tiles on the order of a hundred kilometers and focus the analysis to comparisons among neighbors that share a fire environment rather than across the continent. The estimating panel covers 447935 stands and about 12.18 million stand-years,[2] spanning 39 fire-years. I defer the design to Section 3.

Conditioning on siting through observables and fixed effects, harvested stands show a robust protective association with subsequent first-fire hazard. The pooled estimate is a cloglog coefficient of -0.3332—a hazard ratio of 0.7166 (95% CI [0.565, 0.908], $p \approx 0.007$), a first-fire hazard roughly 28% below that of comparable unharvested stands. The association is protective at 27 of 29 post-treatment horizons and survives dropping the extreme 2023 fire year. The raw, unadjusted association is *already* protective with a hazard ratio near 0.55 before any adjustment. Conditioning attenuates it toward 0.7166. Currently, however, pre-trend and placebo-lead analysis is challenging due to immortal-time survivorship compounded by salvage-window contamination (Section B), and I hope to resolve these issues in future versions.

This paper makes three contributions. First, it provides the first national, continental-scale

[1] British Columbia has moved furthest toward closing this gap: its Forest Practices Board explicitly recommends transitioning from separate management to "landscape fire management," in which wildfire objectives are set across land-use zones and fuel treatment is coordinated with forest management (British Columbia Forest Practices Board, 2023).

[2] About 8.90 million stand-years enter the preferred specification once fully separated zero-fire cells are absorbed, see Section 3.

quasi-experimental estimate of the harvest–fire relationship in Canada: a staggered-adoption first-fire-hazard event study fit to the universe of the managed forest, rather than the plot-scale or single-region evidence that dominates the existing literature.Second, it measures the fire outcome independently of the harvest treatment—an NBAC fire-perimeter record produced by a pipeline separate from the CanLaD disturbance classifier—and clusters inference on fire-years, so that the estimate rests on measurement independence and on inference appropriate to spatially and serially correlated fire rather than on the pixel-level standard errors of a shared-classifier design. Third, (in ongoing work, not yet present in this draft) I connect the reduced-form estimate to the economics of fire–forest co-management (Section 5).

1.1 Related Literature

There is a large forest-ecology literature on the effect of harvest on fire risk, and this literature review will necessarily omit many important and foundational contributions, providing instead only a small sample. Lindenmayer et al. (2009) find that logging elevates fire proneness in moist forests. A long line of stand- and severity-scale studies debates whether the young, even-aged regeneration that follows a cut carries fire more readily than the standing biomass it replaced. My contribution against this body is one of design and scale: a national, quasi-experimental hazard estimate built on an independent fire outcome, rather than plot-scale or severity-conditional ecological evidence, and one that estimates a sign for clearcut-dominated harvest as practiced rather than for an idealized treatment.

A related strand studies deliberate fuel management, where the mechanism my event-time profile is meant to adjudicate is stated cleanly. Agee and Skinner (2005) set out the principles of fuel-reduction treatment, in which removing ladder and surface fuel lowers fire behavior; Thompson et al. (2020), by contrast, find in an experimental boreal crown fire that recent crown thinning reduced neither the rate of spread nor total fuel consumption, a caution that even a deliberate treatment need not lower the hazard. However, this does not resolve or even necessarily directly address the potential effect of clearcuts, which are not a designed fuel treatment. Indeed, I provide evidence that the effect of a clearcut is qualitatively different.

It must be strongly emphasized that forest and fire management are not the primary drivers of the fire trend in Canada. That is almost certainly climate change. Hanes et al. (2019) document fire-regime intensification in Canada over the last half century, and Ellis et al. (2022) the climate-driven decline in fuel moisture behind it. Indeed, while harvest is a highly direct potential lever, its potential impact is inescapably dwarfed by the climatically driven trend. On the policy side, Tymstra et al. (2020) review the Canadian wildfire-management status quo and the potential for co-management of forestry and wildfire.

Methodologically, the estimating approach consists of largely off-the-shelf tools from the modern heterogeneity-robust difference-in-differences literature—Callaway and Sant'Anna (2021), Sun and Abraham (2021), and Borusyak et al. (2024) for estimation under staggered adoption; Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020) for the negative-weighting critique of naive two-way fixed effects; and Roth et al. (2023) for the pre-trend and robustness practice I follow. However, this study represents, to my knowledge, the most sophisticated application of panel data econometrics, and the use of the largest dataset, to the estimation of the effect of timber harvests on wildfire risk.

2 Data

2.1 Sources

Treatment (timber harvests) comes from the Canadian Landsat Disturbance dataset (CanLaD), a set of annual 30 m disturbance rasters covering 1985–2024 that classify each pixel-year into one of a small set of disturbance types—`wildfire`, `harvest`, `windthrow`, `water`, `pest`, and `defoliation` (Perbet et al., 2022). The `harvest` variable is my treatment variable. One limitation of CanLaD is that it is a product of Landsat change-detection rather than a ground census. In particular, it labels the disturbance type but not the intent behind it, so that harvest of salvage timber after a fire is coded identically to routine commercial harvests. In Section B, however, I argue that this likely does not matter for the main estimates. I use CanLaD's `wildfire` code to study how the use of a contaminated fire measure may bias the analysis. I use the `windthrow` code for a natural-disturbance comparison in a falsification exercise.

The outcome is measured by a different instrument: the National Burned Area Composite (NBAC), which assembles fire perimeters from agency fire mapping and satellite imagery through a pipeline entirely separate from the CanLaD disturbance classifier (Canadian Forest Service, 2024). This is important: the fire outcome and harvest treatment, if sourced from the same classifier, would be linked mechanically and introduce correlated measurement error.

I transform the NBAC perimeters into the model outcome: a stand's first fire. For each stand and each post-treatment year I record whether this is the first year that an NBAC perimeter intersects the stand. A stand that has not yet burned remains in the risk set, and once it burns it leaves. The event study is thus a discrete-time model of the time to first fire, and the object of interest is how harvest shifts that hazard. Importantly, reburns—second and later fires on the same ground—are not modeled. The set of already-burned stands is already small, and the set of stands with second burns is smaller still; they would also represent a plausibly qualitatively different forest state into the model and muddy the identification. The survival panel is right-censored: a stand leaves the risk set when the observation window ends without a fire. The estimator accounts for both exits.

Pre-treatment covariates describe the land as it was before the cut. Each is chosen because it plausibly influences operators' harvest location decision: the siting margin the exercise must address. Forest structure variables are aboveground biomass, canopy height, and canopy closure from SCANFI, a spatialized rendering of the Canadian National Forest Inventory built on dense Landsat time series (Guindon et al., 2024). These three variables proxy for the level of development and current fuel load of a stand; I expect them to be highly correlated with merchantable value. Terrain variables are elevation, slope, and a terrain-ruggedness index from the Canadian Digital

Elevation Model; it serves as a proxy for (1) ease of access for harvest operations and (2) the physical determinants of fire spread. A prior-fire binary variable indicates stands whose earliest recorded NBAC fire occurs on or before the 1985 baseline. It is excluded from the preferred specification and enters only in the robustness tests in Section B, Appendix E, row R15. Each covariate is measured before harvest, as structure or terrain variables affected by the harvest might otherwise serve as a mediator of the effect of harvest on fire (Section 3).

Cells which are non-forest, water, or off-Canada are masked out using the VLCE2 land-cover product. The disturbance and fire records cover the entirety of Canada's managed forest. The study population post-filtering is 447935 stands (221830 treated and 226105 control) distributed across 201 firesheds spanning the national managed forest. Firesheds are tiles on the order of 100 km; they serve as the level of geographic fixed effects, so that comparisons are made among "neighbors" rather than across the continent (Section 3).

Descriptive statistics are reported in Table 1. Harvested stands average 126 tonnes per hectare of aboveground biomass compared with 74 for controls; they average 17.0 meters tall compared with 11.0; and they close 62% of canopy compared with 46%. Standardized differences (differences as share of pooled standard deviation) are 0.71, 0.92, and 0.89, respectively. Terrain is more similar across treated and control groups: standardized differences are −0.14 for slope, +0.06 for elevation, and −0.15 for the terrain-ruggedness index. These imbalances are what the covariate adjustment and fixed effects of Section 3 are intended to remove.

	All	Treated	Control	Std. diff.
Aboveground biomass (t/ha)	99.5 (78.1)	125.7 (81.7)	73.7 (64.8)	+0.71
Canopy height (m)	13.9 (7.2)	17.0 (6.2)	11.0 (6.8)	+0.92
Canopy closure (%)	54.2 (19.8)	62.2 (13.8)	46.2 (21.6)	+0.89
Slope (degrees)	5.1 (6.9)	4.6 (5.7)	5.6 (7.9)	−0.14
Elevation (m)	536.5 (378.1)	548.7 (365.7)	524.5 (389.5)	+0.06
Terrain ruggedness index	7.1 (10.0)	6.4 (7.9)	7.9 (11.6)	−0.15
Stands	447,935	221,830	226,105	
First-fire events	55,655	12,848	42,807	
Stand-years	12,402,337	4,302,342	8,099,995	
Stand-years used	12,180,507	4,080,512	8,099,995	

Table 1: Descriptive statistics for treated (harvested) and control (never- and not-yet-harvested) stands. Upper block cells report means with standard deviations in parentheses. Lower block cells report counts.

2.2 Grid Assembly and Spatial Harmonization

Every source layer is warped onto a common 30 m reference grid, which is then sampled from. The unit of analysis is a 30 m sampling cell drawn from a strided lattice over this reference grid. A panel over all 178,400 × 119,100 pixels would carry on the order of 10^{10} observations with extremely high local spatial correlation. I use striding instead of aggregation because of the potentially genuinely spatially fine scale of effects. I subsample the national grid onto a spatial lattice whose sampled columns sit about 3 km apart, thinning the read bands as well,so the lattice remains an unbiased spatial sample of the managed forest while removing near-duplicate neighbors. Details are recorded in Appendix C.

The wide ∼3 km column spacing is chosen so that two cells cut in one harvest operation are essentially never both sampled along a row, so the lattice does not stack near-identical neighboring pixels as though they were independent draws. Rows within a sampled band are read at full 30 m resolution, so some residual spatial dependence remains; that is handled at the inference stage, where standard errors are clustered on fire-years rather than on cells (Section 3). Cells enter the hazard likelihood unweighted, so the estimand is the average treatment effect on the average retained cell rather than on the average hectare or the average harvest operation. I refer to these sampling cells as "stands" throughout the text.

In terms of timing, CanLaD dates a harvest to the year it is first detected. This can occur at most one year after the harvest is conducted. In light of the potential up-to-one-year lag in harvest detection, I quarantine event time $\tau = 1$ and trust the panel only from $\tau \geq 2$ to $\tau = 30$ years after treatment. Beyond that, the treated sample is thin, and even the longest-horizon estimates are based on comparatively few stands.

3 Design and Identification

3.1 Estimand and Estimating Equation

For a stand harvested in calendar year c, write $\tau = t - c$ for the years elapsed since the harvest. The object of interest is $\mathrm{ATT}(\tau)$, the effect of having been harvested on the stand's first-fire hazard τ years later, relative to unharvested but otherwise comparable stands. I denote the first-fire hazard of stand i in year t as h_{it}, the probability that stand i records its first fire in year t conditional on survival to that year without fire. The profile $\{\mathrm{ATT}(\tau)\}$ traces the effect of harvest on that hazard across the post-treatment horizon.

I fit the profile with a staggered-adoption discrete-time complementary-log-log (cloglog) hazard,

$$\mathrm{cloglog}\{h_{it}\} \;=\; \sum_{\tau=2}^{30} \beta_\tau\, D_{it}^{\tau} \;+\; x_i'\gamma_1 \;+\; (x_i^2)'\gamma_2 \;+\; \alpha_{e(i),\,t} \;+\; \phi_{f(i)}, \tag{1}$$

where D_{it}^{τ} is an indicator that stand i survives τ years past its own harvest in year t, x_i are structure and terrain covariates described in Section 2, entered as linear and quadratic terms; $\alpha_{e,t}$ is an ecozone-by-year fixed effect, and ϕ_f is a fireshed fixed effect.[3] For the pooled headline I replace the D_{it}^{τ} with a τ-independent post-treatment indicator $\mathbf{1}\{\tau \geq 2\}$.

[3] I estimate Equation (1) in R with `fixest::feglm` under the `cloglog` link, absorbing the two high-dimensional fixed effects by alternating projection; the sufficient-statistic aggregation and fixed-effect details are described in Section D.

The controls are never-treated and not-yet-treated stands only. Treated stands are removed from the dataset 30 years after treatment and enter as neither treated observations nor controls.[4]

The ecozone-by-year fixed effects partial out the common fire weather of a given season in a given biome (e.g. the boreal in 2023). The fireshed fixed effects restrict comparisons to be made among neighbors: within a fireshed, harvested and unharvested stands face similar terrain, regional climate, and suppression jurisdictions.

3.2 Identifying Assumptions and Threats

The identifying assumption is conditional parallel trends of first-fire hazard of treated and untreated stands within a fireshed and ecozone-year, after adjusting for pre-treatment forest structure and terrain. I report the pre-treatment lead profile. I report that check and its limitations in Section B.5. I address the two most significant threats to this assumption below.

The first and most consequential is siting: operators cut where access, stand age, and operability are favourable, and that same land tends to burn less readily than the remote, rugged, or unmanaged forest that is never cut. The pre-treatment covariates and the fireshed fixed effect exist to absorb exactly this imbalance, but may not absorb it fully if this confounding occurs within firesheds and ecozone-years. The partial success of the addition of fixed effects can be seen in that: the raw association is already protective, and conditioning moves it *toward* the null (Section B.4). Figure 1 shows the same pattern with respect to the covariates: the raw standardized differences on forest structure run from 0.71 to 0.92, and demeaning within the ecozone-year and fireshed fixed effects reduced the to around 0.28 without reducing them to zero. That the within-fixed-effect gap does not close to zero is the visible sign that the adjustment is partial: a same-signed siting confound not captured by structure, terrain, or fireshed would leave a residual protective bias. In Section B.5 and Section B.6 I return to this issue and bound the residual bias.

[4]This addresses the negative-weighting issue of naive two-way fixed effects under staggered timing described in Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020).

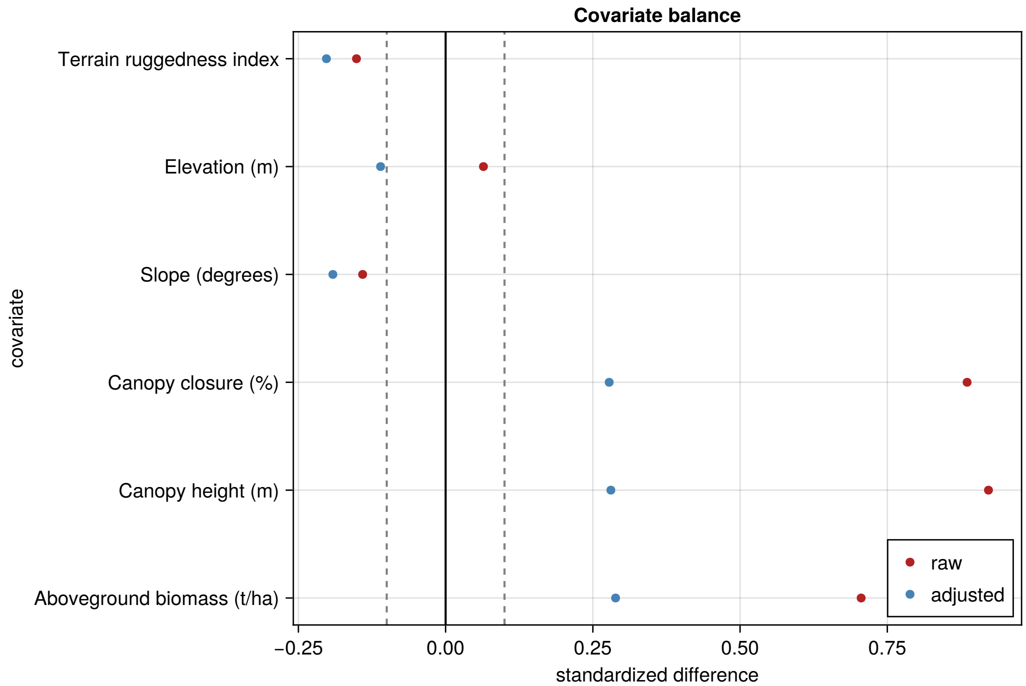

Figure 1: Standardized differences in pre-treatment covariates between treated (harvested) and control stands. Each row is a covariate. The raw markers give the standardized mean difference (treated minus control over the pooled standard deviation) in the unadjusted stand-level sample; the adjusted markers give the within-fixed-effect difference after demeaning by ecozone×year and coarsened fireshed on the estimation sample.

The second threat is suppression. Fire suppression is itself targeted. Worse, it is partly targeted at merchantable timber specifically. That is, the protection that can be read off the hazard could in part reflect the fire agency's choice to protect merchantable (and harvested) areas.[5] Ideally, management-zone fixed effects would be added, but a national management-zone raster does not exist to my knowledge. I return to this gap in Section B and Section 6.

There is a third potential threat to identification which I solve, however. If I were to source the fire outcome from the same disturbance classifier that defines the harvest treatment, treatment and outcome would be correlated mechanically. By drawing the outcome from NBAC (a fire record produced by an entirely separate pipeline), I avoid this by construction.

This represents as far as I am currently able to go in addressing siting selection and selective suppression. Notably, I cannot reliably compute pre-trends as an immortal-time survivorship gradient in the never-treated obstructs this, and is compounded by salvage-window contamination (Section B.5). Conditional parallel trends is assumed here rather than verified. A same-signed siting confound that observable adjustment cannot reach cannot be excluded. I am nonetheless confident in the protective *direction* if not a identified point estimate of the *effect*. In Section B.5, I describe how progress might be made on this issue.

[5]The countervailing effect is also possible: harvest could bring road and ignition access as well.

3.3 Inference

I cluster the standard errors on fire-years. I interpret the t statistic using a $t(G-1)$ distribution with $G = 39$, the conservative small-cluster reference.

As an alternative clustering level, I use firesheds, (the 100 km spatial tiles intersected with ecozone which I use as the spatial fixed effect). Fireshed-clustered and two-way (fire-year $\times$ fireshed) standard errors are reported as robustness in Section B.1. All three preserve significance (intervals exclude a hazard ratio of one).

3.4 The Estimand's Support

The fireshed and ecozone-year fixed effects introduce fine spatial control but reduce power and drop 26.9% of observations (fixed-effect groups that contain no fire at all), with 127 of 390 ecozone-year groups (32.6%) and 86 of 296 fireshed groups (29.1%) fully separated by the absence of fire. Among treated rows, 29.2% are dropped for the same reason.[6].

Within that population, nearly all of the treated stock sits on common within-fireshed support.Of the 201 firesheds in the overlap diagnostic, 197 contribute to identification (98.0%), holding 99.5% of treated stands on within-fireshed support.

4 Results

4.1 Primary Event-Study Estimates

Table 2 presents the estimates, both dependent on horizon and pooled. Conditioning on siting through observables and fixed effects, harvested stands in Canada's managed forest show a robust protective association with subsequent first-fire hazard. The pooled post-treatment estimate is a complementary-log-log coefficient of -0.3332, a hazard ratio of 0.7166 (95% CI [0.565, 0.908]), a first-fire hazard roughly 28% below that of comparable unharvested stands. The sign is negative for 27 of 29 post-treatment horizons, but individual horizon-level estimates are imprecise, with only 9 out of 29 individually significant.

The more highly-powered pooled estimate is a hazard ratio of 0.7166 (cloglog -0.3332, SE 0.1171, $t = -2.84$, $p = 0.0071$ over $G = 39$ fire-years), with a 95% confidence interval of [0.565, 0.908]. That is, harvested stands face a first-fire hazard about 28% lower than otherwise-comparable unharvested neighbors in the same fireshed and ecozone-year. [7]

The profile in the upper rows of Table 2, plotted as Figure 2, is protective almost everywhere. Across the 29 horizons from $\tau = 2$ to $\tau = 30$, the hazard ratio lies below one at 27, the two exceptions being $\tau = 9$ (HR 1.18) and $\tau = 29$ (HR 1.14), and both have confidence bands that comfortably cover one. The near horizons are protective but individually imprecise: $\tau = 2$ through $\tau = 8$ all sit below one, between roughly 0.66 and 0.83, but each carries a band wide enough to cover it, so the

[6] The estimand is therefore the ATT *among fire-experiencing ecozones and firesheds*

[7] An alternative aggregation that averages the per-horizon estimates ratios rather than fitting a single pooled estimate delivers 0.6818.

early profile is a gentle protective drift rather than a sharp on-impact drop. There is, in particular, no early positive bump: there is no horizon at which harvested stands burn detectably *more*, which I discuss further in Section 5. The protection, if anything, appears to deepen at longer horizons.

Power diminishes at the longest horizons: stands are progressively censored by their own first fires and by the edge of the observation window. I shade the horizons beyond $\tau \approx 20$ as low-power in Figure 2.

τ	HR $= e^{\beta}$	95% CI (HR)
2	0.81	[0.57, 1.13]
3	0.76	[0.56, 1.02]
4	0.66	[0.35, 1.22]
5	0.67	[0.43, 1.05]
6	0.73	[0.40, 1.34]
7	0.75	[0.42, 1.34]
8	0.83	[0.58, 1.19]
9	1.18	[0.66, 2.11]
10	0.82	[0.41, 1.66]
11	0.55	[0.32, 0.94]
12	0.82	[0.57, 1.17]
13	0.97	[0.59, 1.60]
14	0.79	[0.57, 1.08]
15	0.78	[0.49, 1.25]
16	0.85	[0.49, 1.46]
17	0.58	[0.35, 0.98]
18	0.66	[0.38, 1.15]
19	0.45	[0.23, 0.91]
20	0.61	[0.38, 0.99]
21	0.63	[0.40, 0.98]
22	0.61	[0.44, 0.85]
23	0.59	[0.29, 1.20]
24	0.48	[0.31, 0.74]
25	0.40	[0.24, 0.68]
26	0.52	[0.32, 0.85]
27	0.63	[0.30, 1.29]
28	0.33	[0.21, 0.52]
29	1.14	[0.55, 2.36]
30	0.95	[0.64, 1.42]
Pooled	**0.72**	[0.57, 0.91]

Table 2: Primary event-study estimates of the association between past harvest and subsequent first-fire hazard, by horizon τ and pooled. $\tau = 1$ arm is dropped due to harvest-year timing ambiguity. Each row reports the hazard ratio $\exp(\beta)$ from the complementary-log-log first-fire hazard and the 95% confidence interval.

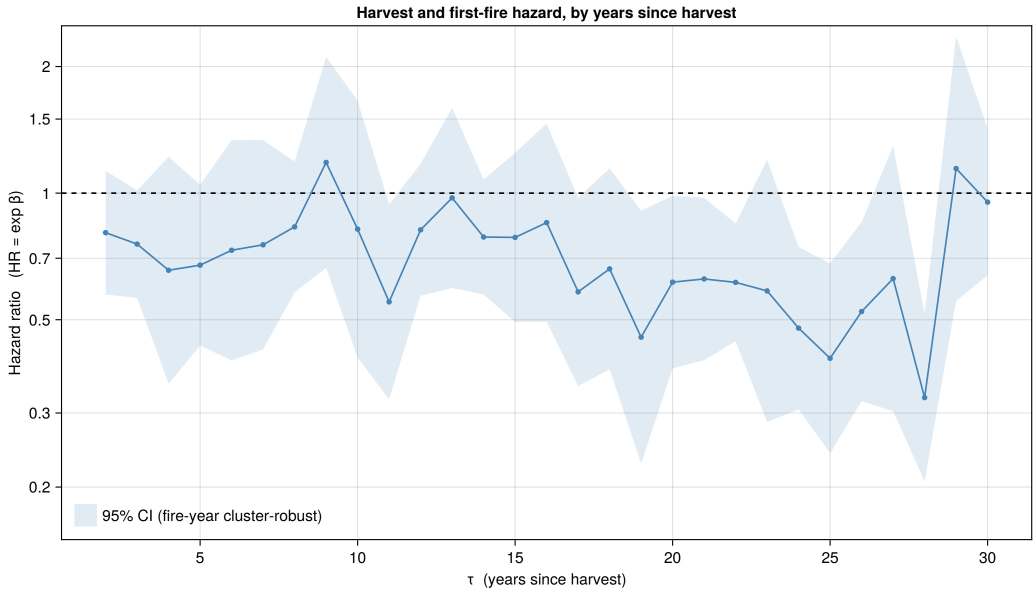

Figure 2: Event-study profile of the harvest–wildfire association. Each point is the hazard ratio $\exp(\beta_\tau)$ for a first fire at τ years since harvest, relative to never- and not-yet-harvested stands. The shaded band is the 95% fire-year cluster-robust confidence interval. The vertical axis is on a log scale. Points below the line indicate a lower subsequent fire hazard for harvested stands.

4.2 Progressive Specification Analysis

Table 3 shows the result of progressively adding controls to the raw contrast. Part of its purpose is to demonstrate the conditionality of the results: to the extent that siting endogeneity is imperfectly addressed by the final specification, the true result may be further attenuated. Column (C0) is the raw, unadjusted contrast: a hazard ratio of 0.551. Column (C1) adds the Tier-A siting confounders—SCANFI biomass, canopy height, canopy closure, and DEM slope—and the estimate moves to 0.658. Column (C2) adds the terrain precision terms, elevation and the terrain-ruggedness index, and the estimate holds at 0.658—the terrain terms add essentially nothing once the SCANFI structure covariates are in, which locates the confounding work in stand structure and canopy rather than in elevation and ruggedness. Column (C3), the preferred specification, adds the ecozone×year and coarsened-fireshed fixed effects and moves the estimate to 0.7166. The estimate moves *up* (toward one) at every step that adds a control: adjustment attenuates the raw protection rather than deepening it. The estimation sample of stand-years does contract at the fixed-effects step—from about 12.18 to 8.90 million as the absorbers separate off zero-fire cells (Section 3.4)—so the same fires are explained on a sample restricted to fire-experiencing cells, with progressively more of the siting structure held fixed.

Only the fixed-effects column, (C3), however, has a confidence interval that excludes a hazard ratio of one; the raw and covariate-only columns (C0) through (C2) have wider intervals that cover one, at p-values of 0.067, 0.109, and 0.113 respectively. That is, the fixed effects both attenuate the effect and sharpen the estimate to conventional significance. The raw contrast, while protective, is not on its own distinguishable from one.

	(C0) Raw	(C1) + Confounders	(C2) + Terrain	(C3) + FE (preferred)
Pooled harvest HR	0.551	0.658	0.658	**0.717**
95% CI (HR)	[0.29, 1.04]	[0.39, 1.10]	[0.39, 1.11]	**[0.57, 0.91]**
p-value	0.067	0.109	0.113	**0.007**
Structure + slope (Tier-A)	—	✓	✓	✓
Terrain precision (elev., TRI)	—	—	✓	✓
Ecozone×year FE	—	—	—	✓
Fireshed FE	—	—	—	✓
First-fire events (both arms)	54,694	54,694	54,694	54,694
Estimation stand-years	12.18M	12.18M	12.18M	8.90M
Fire-year clusters (G)	39	39	39	39

Table 3: Progressive specifications of the pooled harvest–wildfire association. Columns add controls cumulatively from left to right: (C0) raw / unadjusted; (C1) + biomass, height, canopy closure, DEM slope; (C2) + elevation, terrain-ruggedness index; (C3) + ecozone×year and coarsened-fireshed fixed effects. The estimate attenuates toward one as siting confounders and fixed effects enter, the pattern expected when harvest is sited on lower-fire-risk land.

5 Mechanisms

The forest-ecology literature emphasizes two channels with opposite directions:the young-fuel channel (a cut is followed by dense regeneration and a flush of more flammable surface fuel) and the reduced-load channel (removing standing biomass and therefore fuel) which motivates deliberate fuel-reduction treatment (Agee and Skinner, 2005; Brodie et al., 2024). The overall effect is an empirical question, and the profile of Section 4.1 suggests that it is the reduced-load channel.

The profile shows a protective association almost everywhere and no early positive bump. Were the young-fuel channel the governing force, its natural signature would be an early interval of elevated hazard that fades as the regeneration matures and self-thins. The data show the reverse: a gentle protection early that deepens rather than reverses.

Three substantial sources of uncertainty remain, however. The first is that the absence of an early positive bump is imprecise enough that the confidence bands cover a possible modest young-fuel spike.The second is the potential residual siting confound (Section B.5). The third is that a clearcut is not a designed fuel treatment. Clearcutting, which dominates Canadian harvest (Appendix A.1), takes the merchantable stems and leaves slash behind; its net effect on flammability is not the typical case covered by fuel-treatment theory. Indeed, in other work crown thinning has been found not to slow a fire's rate of spread or reduce its total fuel consumption (Appendix A.2). I describe next steps in addressing these issues in Section G.1.

6 Conclusion

Across Canada's managed forest, stands that have been harvested carry a lower subsequent first-fire hazard than comparable unharvested stands. The pooled estimate is a hazard ratio of 0.7166: a first-fire hazard roughly 28% below the unharvested baseline. The point estimate is protective at 27 of 29 event-time horizons, and survives observable adjustment, multiple clustering schemes, and the removal of the extreme 2023 fire season.

However, challenges to identification remain: possible residual bias due to siting endogeneity, obstructions to pre-trend testing, and low power. Nonetheless, I find a robust protective *direction*, though uncertainty remains in the true causal effect.

This is the first national, quasi-experimental estimate of the effect of harvest on first-fire hazard in Canada. Its design furthermore eliminates a number of confounds that have likely affected prior studies. Further work is needed to determine the robustness of the protective effect and its external validity. The reduced-form protective estimate speaks to the most contested premise of the fire–forest co-management case (Appendix A.2): whether harvest as practiced leaves land more or less likely to burn. It demonstrates that where stands permit, commercial or partial harvest could in principle deliver fuel reduction while generating revenue that defrays its own cost. This is in contrast to pay-to-dispose fuel cuts which lowers load but at high cost. This result therefore motivates further study into the program of fire–forest co-management.

Four limitations still remain. First, residual siting and suppression selection remains the open challenge to identification: the design conditions on observed structure, terrain, and fireshed, but an unobserved same-signed confound would cause a protective bias to remain. Second, the first-class response to the suppression channel, a management-zone split, is limited by the absence of a national management-zone raster. Third, treatment and outcome are measured, not observed: CanLaD dates harvest to detection and NBAC records perimeters, and both carry timing and edge imprecision. Fourth, fireshed and ecozone-year fixed effects restrict the estimand to fire-experiencing cells on marginal common support.

References

Agee, James K. and Carl N. Skinner, "Basic principles of forest fuel reduction treatments," *Forest Ecology and Management*, 2005, *211* (1–2), 83–96.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: Robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

British Columbia Forest Practices Board, "Forest and Fire Management in BC: Toward Landscape Resilience," Special Report SR61, Forest Practices Board, Victoria, BC 2023. `https://www.bcfpb.ca/wp-content/uploads/2023/06/SR61-Landscape-Fire-Management.pdf`.

Brodie, Emily G., Eric E. Knapp, Wesley R. Brooks, Stacy A. Drury, and Martin W. Ritchie, "Forest thinning and prescribed burning treatments reduce wildfire severity and buffer the impacts of severe fire weather," *Fire Ecology*, 2024, *20* (1), 17.

Byrne, Brendan et al., "Carbon emissions from the 2023 Canadian wildfires," *Nature*, 2024, *633* (8031), 835–839.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Canadian Forest Service, "National Burned Area Composite (NBAC)," Canadian Wildland Fire Information System, Natural Resources Canada 2024. `https://cwfis.cfs.nrcan.gc.ca/datamart`.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Ellis, Tegan M., David M. J. S. Bowman, Piyush Jain, Mike D. Flannigan, and Grant J. Williamson, "Global increase in wildfire risk due to climate-driven declines in fuel moisture," *Global Change Biology*, 2022, *28* (4), 1544–1559.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Guindon, Luc, Philippe Villemaire, Michael Béland et al., "A new approach for spatializing the Canadian National Forest Inventory (SCANFI) using Landsat dense time series," *Canadian Journal of Forest Research*, 2024, *54* (7), 793–815.

Hanes, Chelene C., Xianli Wang, Piyush Jain, Marc-André Parisien, John M. Little, and Mike D. Flannigan, "Fire-regime changes in Canada over the last half century," *Canadian Journal of Forest Research*, 2019, *49* (3), 256–269.

Lindenmayer, David B., Malcolm L. Hunter, Philip J. Burton, and Philip Gibbons, "Effects of logging on fire regimes in moist forests," *Conservation Letters*, 2009, *2* (6), 271–277.

Perbet, Pierre, Luc Guindon, David L. P. Correia, Philippe Villemaire, Omid Reisi Gahrouei, and Rémi St-Amant, "CanLaD: The Canadian Landsat Disturbance dataset (annual, 30 m)," Natural Resources Canada / Open Government 2022. https://open.canada.ca/data/en/dataset/902801fd-4d9d-4df4-9e95-319e429545cc.

Plantinga, Andrew J., Randall Walsh, and Matthew Wibbenmeyer, "Priorities and effectiveness in wildfire management: Evidence from fire spread in the western United States," *Journal of the Association of Environmental and Resource Economists*, 2022, *9* (4), 603–639.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Thompson, Dan K., Dave Schroeder, Sophie L. Wilkinson, Quinn Barber, Greg Baxter, Hilary Cameron, Rex Hsieh, Ginny Marshall, Brett Moore, Razim Refai, Chris Rodell, Tom Schiks, Gregory J. Verkaik, and Jessica Zerb, "Recent crown thinning in a boreal black spruce forest does not reduce spread rate nor total fuel consumption: results from an experimental crown fire in Alberta, Canada," *Fire*, 2020, *3* (3), 28.

Tymstra, Cordy, Brian J. Stocks, Xinli Cai, and Mike D. Flannigan, "Wildfire management in Canada: Review, challenges and opportunities," *Progress in Disaster Science*, 2020, *5*, 100045.

VanderWeele, Tyler J. and Peng Ding, "Sensitivity analysis in observational research: Introducing the E-value," *Annals of Internal Medicine*, 2017, *167* (4), 268–274.

A Institutional Background: Fire and Forest Management in Canada

A.1 Two Separate Bureaucracies

The policy premise of this paper is a division of labor. Across Canada, wildland fire and timber production are run as operationally distinct programs, with different delivery channels, budgets, and planning horizons, even where they act on the same forest and answer, formally, to the same government. In Ontario, for instance, both functions sit inside the provincial natural-resources ministry, but fire suppression is delivered directly and in-house through a dedicated aviation-and-fire branch, while timber management is delegated to licence holders—forest-products firms, co-operatives, and arm's-length Crown corporations—who prepare and implement the forest-management plans on their units under a Registered Professional Forester's sign-off and Crown approval.[8] Wildland fire is provincial and territorial jurisdiction throughout: there is no federal fire agency, and the national interagency center brokers resource-sharing and mutual aid without holding command over the agencies that actually fight fires (Tymstra et al., 2020). Whatever else it is, the split is durable, and it is the wedge this paper's question exploits.

The split is also old. Fire protection was institutionalized as a distinct administrative branch roughly a century ago, in the provincial forest-fire statutes of the early 1900s, and has stayed operationally separate from timber administration ever since. I read that durability as institutional path-dependence—a boundary drawn when aerial detection and organized suppression were the frontier of fire policy, and left in place as both functions professionalized around it—though I am aware of no source that labels it so, and offer the reading as mine rather than as an established finding. The label matters less than the fact: the boundary is stable, and it did not arise from any judgment that fire and timber are best managed apart.

A second force plausibly reinforces the separation, through the economics of trade rather than the history of administration. The Canada–US softwood-lumber dispute, one of the longest-running trade conflicts between the two countries, turns on the allegation that Canadian provinces subsidize lumber by charging below-market stumpage on Crown land, which supplies the overwhelming majority of the country's fiber; it has produced combined countervailing and antidumping duties on the order of a third of producers' value, and recurrent litigation. This is not a distant sensitivity to fire policy. In a 2026 filing the US industry coalition named Canadian wildfire fuel-management spending and forest-access-road funding among the programs it wants treated as countervailable subsidies. From that I draw only a conjecture, and label it as mine: a government wary of triggering countervailing-duty action has reason to avoid anything that could be cast as subsidizing timber producers, which is exactly what publicly funded fuel-reduction thinning that yields marketable wood would look like. I have found no source that establishes this chilling effect—the documented pattern so far runs the other way, with Canada funding such programs and the US then attacking

[8] The delegated-licence architecture is Ontario's, not Canada's. Quebec, under its *Sustainable Forest Development Act*, has the province itself prepare the management plans and auctions a fifth to a quarter of public timber through a state marketing board, allocating the balance through guaranteed supply agreements. "Delegated to licence holders" is therefore a provincial fact, not a national one.

them—so I offer it as a plausible reinforcing pressure on the separation, not as a demonstrated one.

The two programs run on incompatible clocks, and that is where the separation bites. Timber harvest on public land proceeds through long-horizon forest-management plans—in Ontario, a ten-year planning term layered on tenures that run longer still—prepared years in advance and revised on a decadal cycle. Fire suppression, by contrast, is dispatched operationally and reactively, its horizon a fire season or a single afternoon's weather. Fuel loads and harvest schedules are settled inside the slow planning process; the decision to defend a given stand is taken inside the fast one. Because the two horizons never meet, harvest and fuel are not co-optimized, and there is no institutional venue in which a cut might be timed or shaped to serve a fire objective, or a fire objective allowed to reshape a cut. The separation is not merely administrative tidiness. It is the reason the joint problem is, at present, nobody's problem.

Harvest practice sharpens the point. Clearcutting is the predominant silvicultural system in the Canadian managed forest; partial cuts and commercial thinning are the minority. This matters for what my estimate can mean, because a clearcut is not a fuel-reduction treatment—if anything it is close to the reverse. Deliberate fuel management targets the small-diameter, non-merchantable material that drives fire behavior—ladder fuels, understory, and surface fuels below the roughly 7–12 cm merchantability threshold—in labor-intensive, low-mechanization work that yields little or no timber revenue and is therefore expensive, with published costs spanning from tens to well over a thousand US dollars per acre. The material it removes has, in the typical case, no buyer: it is pile-burned, masticated and left on site, or broadcast-burned, because handling and haul cost exceed the value of dispersed small wood.[9] A clearcut inverts each of these features. It takes the merchantable stems, generates revenue, and leaves behind, over the following years, a dense even-aged regeneration rather than an opened, fuel-poor stand. The sign of the harvest–fire relationship is thus ambiguous on the underlying forest science—young regenerating fuel arguing one way, removed standing biomass the other—and a protective finding, if it holds, cannot be waved through as obvious fuel reduction. It would be an empirical statement about clearcut-dominated harvest as it is actually practiced, not about an idealized treatment that Canadian forestry does not, for the most part, perform.

A.2 Why Wildfire is a Material Economic and Carbon Problem

The stakes that make the question worth answering are, first, atmospheric. The 2023 season, whose 647 million-tonne carbon release opened this paper, amounted to roughly 6% of global fossil-fuel CO_2 that year—set against a global fossil total near 37 billion tonnes—and to more than four times Canada's own fossil emissions (Byrne et al., 2024). A single country's forests, burning, registered at the scale of a major emitting economy. The figure carries two qualifications that I state rather than bury: it is a gross biogenic release that regrowth reabsorbs in part over subsequent decades, so it is not interchangeable with a permanent fossil-fuel tonne, and 2023 was an extreme outlier, on the

[9]Biomass-energy utilization is a genuine exception, but it is confined to the places where the collection and conversion infrastructure already exists.

order of seven times the typical year's area burned. Both bear on interpretation, and neither shrinks the season to insignificance. That singular character is also why the year earns its own robustness check: because 2023 is one cross-fireshed mega-shock rather than many independent events, I ask throughout whether the protective association is merely its shadow, and drop it to find out (Section B.2).

The season was extreme, but the baseline it departed from is itself rising. Fire-regime intensification across Canada over the last half century (Hanes et al., 2019), driven in part by the climate-linked decline in fuel moisture that raises flammability continent-wide (Ellis et al., 2022), means the relevant counterfactual is not a stationary risk but a trend. Harvest, against that backdrop, is at most a marginal lever: no plausible change in cutting practice offsets the climatic forcing, and any harvest signal must be identified against fire-years that move together strongly. That correlation structure is not a nuisance to be conditioned away but a first-order feature of the data, and it dictates how I conduct inference—clustered on fire-years, and stress-tested by removing the years that dominate the record.

It is the conjunction of the two—a rising externality and two programs that do not coordinate—that gives the question its economic edge. The efficiency argument for managing fire and forest together is easy to state: where stands permit, commercial or partial harvest could achieve fuel reduction while generating revenue that defrays the cost of treatment, in place of both the clearcut-plus-dense- regeneration pattern that may raise flammability and the pay-to-dispose fuel cut that raises no revenue at all. Whether those joint gains are large, or even reliably positive, is genuinely contested rather than a free lunch. Cost recovery from thinning is usually incomplete; treatment efficacy is overwhelmed by the extreme weather that drives the worst boreal fires; at least one controlled boreal experiment found that crown thinning did not slow a fire's rate of spread, even as it cut fire intensity to roughly a third of the untreated control (Thompson et al., 2020); and a substantial ecological literature holds that logging can raise, not lower, fire severity (Lindenmayer et al., 2009). I raise the co-management case here as motivation, and hold it there. The reduced-form protective direction this paper recovers speaks directly to the last, most contested premise—whether harvest as practiced leaves land more or less likely to burn—but it does not by itself settle the policy question, which would turn on a welfare and management model I defer to future work (Section 6). What the separation and the stakes together establish is narrower and sufficient: the sign of the harvest–fire relationship is worth knowing, and no institution is currently charged with learning it.

B Robustness and Falsification

This section tests: whether the protective association survives the inference regime, the extreme fire-years, the support restriction, and the progressive adjustment; whether the pre-treatment leads permit a parallel-trends test; how large an unobserved confound would have to be to overturn the result; and where across the country the association concentrates.

Apart from the pre-trend test, which appears to be infeasible rather than feasible and bad news, the answers are mostly reassuring.

B.1 Inference and Clustering

The pooled point estimate does not move with the choice of clustering; only the interval does. Table 4 reports the preferred specification under three schemes. Under fire-year clustering, the primary, the hazard ratio is 0.7166 with $G = 39$ and $p = 0.0071$. Under fireshed clustering it is the same 0.7166 with $G = 296$ and a tighter $p = 0.0046$; under two-way fire-year $\times$ fireshed clustering it is again 0.7166 with $p = 0.0127$. All three intervals exclude a hazard ratio of one. I lead with the fire-year number and do not promote the tighter fireshed p-value, because the season-wide shock—one weather system lighting up a whole biome—is the correlation structure most likely to be underestimated, and fire-year is the conservative block that respects it. The two-way scheme, which allows correlation along both the temporal and the spatial margin at once, is the most demanding of the three, and the estimate clears it as well.

Clustering	HR	95% CI (HR)	SE	p	G
Fire-year	0.717	[0.57, 0.91]	0.117	0.007	39
Fireshed	0.717	[0.57, 0.90]	0.117	0.005	296
Fire-year $\times$ fireshed	0.717	[0.55, 0.93]	0.127	0.013	39

Table 4: Sensitivity of the pooled harvest–wildfire estimate to the choice of clustering. Rows report the same preferred complementary-log-log specification with cluster-robust standard errors computed on fire-years ($G = 39$), on coarsened firesheds ($G = 296$), and two-way on fire-year$\times$fireshed. Columns report the hazard ratio, its 95% confidence interval, the standard error, the two-sided p-value, and the number of clusters G. The point estimate (HR 0.717) is identical across schemes by construction; only the standard error and interval width change, and all three intervals exclude HR = 1. Fire-year clustering is the primary specification (the conservative block for season-wide fire shocks); the others are reported as disclosure.

B.2 Leave-Out Robustness

The record is dominated by a handful of extreme seasons, and 2023 above all: one cross-fireshed mega-shock that burned on the order of seven times a typical year's area (Appendix A.2). A natural worry is that the protective association is merely that season's shadow—that harvested stands happened to sit off the path of the 2023 fires and the estimate is reading geography, not harvest. The leave-2023 test in Table 5 answers it directly. Dropping 2023 entirely and re-estimating on the remaining 38 fire-years leaves the association not weaker but stronger, at a hazard ratio of 0.6307 (cloglog -0.4609, $p = 0.0057$). The sign is not a 2023 artifact; if anything, the single most extreme season was pulling the pooled estimate toward one, not away from it.

A wider leave-out, dropping the entire 2018–2024 high-fire window, moves the estimate further still, to a hazard ratio of 0.554, which I report and read asymmetrically in the appendix (Section E, Table 7): that window removes about a fifth of the fire-years, so its strengthened point is consistent with the association without independently confirming it. I resist reading the monotone leave-out ordering as a dose-response in harvest; it is in part a story about which fire-years remain in the sample and how their composition shifts the pooled contrast.

Specification	HR	95% CI (HR)	t
Full sample (preferred)	0.717	[0.57, 0.91]	-2.84
Leave out 2023	0.631	[0.46, 0.87]	-2.93

Table 5: Leave-out robustness of the pooled harvest–wildfire estimate. Rows report the preferred specification on the full sample (the anchor) and after dropping 2023, the single extreme cross-fireshed fire season (R1). Columns report the hazard ratio, its 95% fire-year cluster-robust confidence interval, and the t statistic. The protective sign persists and strengthens when 2023 is removed, so the estimate is not a 2023 artifact. Further leave-outs (R2, dropping the 2018–2024 high-fire window; R15, +prior_fire) are reported in the appendix (Section E, Table 7).

B.3 Within-Fireshed Common Support

The estimand rests on within-fireshed comparisons (Section 3.4), so the question is whether enough of the treated stock has unharvested neighbors to compare against. It does: 197 of the 201 firesheds in the overlap diagnostic contribute (98.0%), holding 99.5% of treated stands on within-fireshed support, and the four non-contributing firesheds are tiny treated-only or zero-on-support cells with at most seven treated stands each, immaterial to the estimand. The fireshed fixed effect therefore identifies the ATT off essentially the whole fire-experiencing treated population, and I report the headline as a national estimate on that basis. The qualification set out in Section 3.4 still holds: this is a *marginal* bounding-box test, which certifies that each treated stand's covariate ranges lie inside its fireshed's control ranges but cannot certify joint support or detect holes in the interior of the covariate space, so I call the result marginal within-fireshed support rather than claim positivity is a non-issue. The full overlap scatter is in the appendix (Section F, Figure 5).

B.4 The Coefficient Path: Confounding-Removal vs. Overcontrol

The single most informative object for identification is the path the estimate traces as controls enter, and it is worth reading Table 3 a second time with that in view. The pattern is the one Plantinga, Walsh, and Wibbenmeyer read off their own column movements (Plantinga et al., 2022): the coefficients diminish in magnitude as physical controls are added, and the direction of that diminution tells you which way the omitted variables were biasing the raw estimate. Here the raw hazard ratio is already protective at 0.55, and adding the siting confounders and then the fixed effects moves it *toward* the null, to 0.7166, plotted in Figure 3. The estimate travels up, not down, as adjustment proceeds.

That direction is the crux. When harvest is sited on land that is intrinsically less fire-prone—flatter, more accessible, more merchantable, and for those same reasons less likely to burn—the raw comparison *overstates* the protection, because it credits harvest with a gap that was there in the land before any cut. Conditioning on the structure and terrain that drive both the siting decision and the fire propensity removes part of that gap, and the estimate attenuates toward one. This is the ordinary signature of confounding being removed, and it is the *opposite* of the pattern that would worry a reader most: an effect "manufactured by adjustment," in which the protective sign is absent in the raw data and only *emerges* as controls are piled on, would show the estimate moving away from one as it is adjusted. It moves toward one. The protective sign is a feature of the raw data that adjustment shrinks, not a feature that adjustment conjures.

Attenuation-toward-null under the observed confounders does not, by itself, prove that the residual protection is causal, and I do not read it that way. An unobserved same-signed siting confound—favorable siting not captured by SCANFI structure, terrain, or the fireshed fixed effect—would leave a residual protective bias that no amount of the *observed* controls can reach, and the coefficient path is silent about it by construction. What the path does rule out is the manufacture story; what it does not rule out is a residual unobserved confound, which is why the bounding exercise of Section B.6 and the pre-trend discussion of Section B.5 carry the rest of the weight. One reassurance the performed battery does offer: adding a prior-fire indicator, the covariate most likely to act as a collider and pull the estimate the other way, leaves it essentially unchanged at 0.719 (Section E), so the estimate is not balanced on a knife-edge of that one control.

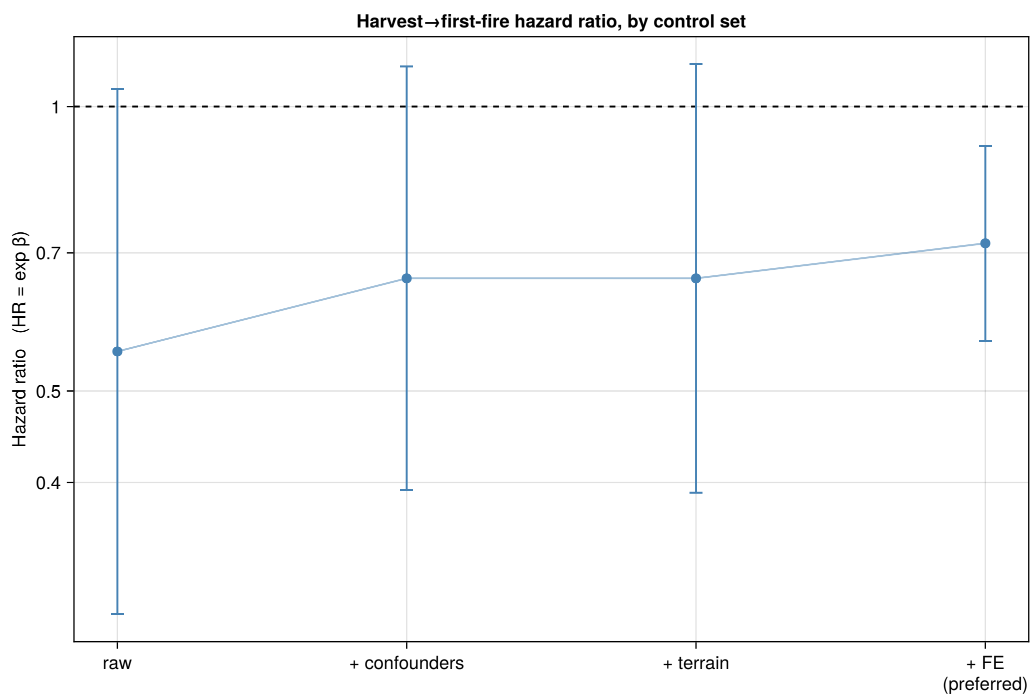

Figure 3: The coefficient path of the pooled harvest–wildfire hazard ratio as controls enter cumulatively: raw / unadjusted, + Tier-A siting confounders, + terrain precision terms, and + ecozone×year and coarsened-fireshed fixed effects (the preferred specification). Points are the hazard ratio $\exp(\beta)$; whiskers are the 95% fire-year cluster-robust confidence interval; the dashed line marks HR = 1 (no association). The vertical axis is on a log scale. The estimate moves *up*, toward the null, as siting confounders and fixed effects enter.

B.5 Pre-Trend and Placebo Leads: the Identification Crux

A parallel-trends design lives or dies on the pre-treatment leads, and this is where I am most exposed. The clean, closest lead is the reportable pre-trend statistic. The $\tau = -1$ lead coefficient is 0.204 on the cloglog scale, a hazard ratio of 1.23 with a 95% confidence interval of [0.77, 1.95], which does not reject conditional parallel trends at $\tau = -1$: the year before harvest, stands that will be cut are not detectably diverging from those that will not. On its own terms that is the check one wants, and it comes back clean (Roth et al., 2023).

The leads at longer horizons cannot be read the same way, and I am explicit that they are an artifact rather than a test. The far-lead coefficients run monotonically from a hazard ratio near 0.007 at $\tau = -5$ to near 0.13 at $\tau = -2$ —absurdly, mechanically protective—and they are so for two compounding reasons that have nothing to do with an economic pre-trend. The first is immortal-time survivorship. A stand enters the treated-lead arm only if it is in fact harvested at its eventual date and survives fire-free until then, so the lead window is populated by definitional survivors; and because the leads panel's reference arm is never-treated stands only, carrying no event time of its own, it cannot difference out the resulting survivor gradient. The second is salvage-window contamination. The five-year salvage exclusion that keeps post-fire salvage harvest out of the treatment is coextensive with the lead window, so it mechanically strips pre-harvest fires from the treated arm; and because the salvage filter keys on CanLaD alone, a small share of NBAC pre-harvest fires leaks into the treated arm as well. Together these produce the steep, monotone far-lead gradient, which I disclose as the artifact it is and do not offer as a joint parallel-trends test. The clean $\tau = -1$ lead sits outside the worst of the gradient and is the statistic to read.

It is worth being clear about what this artifact does and does not contaminate. The immortal-time survivorship that corrupts the leads is a property of the leads' reference arm, not of the pooled estimate. The post-treatment contrast draws its controls from not-yet-treated stands, which face the same requirement to survive fire-free to their own eventual harvest date, so the survivor gradient enters both arms and differences out. The leads panel cannot borrow that cancellation: its reference is never-treated stands only, carrying no harvest date and so no event time against which a matching survival window could be defined. The problem is confined to the pre-trend test; it does not travel to the headline hazard ratio.

What this leaves me is a claim held deliberately at a direction. The protective association survives observable adjustment (Section 4.2), every clustering scheme (Section B.1), and the removal of the extreme fire-years (Section B.2), and the clean $\tau = -1$ lead does not reject conditional parallel trends. But because the far-lead immortal-time gradient blocks a genuine *joint* pre-trend test, the same-signed siting confound of Section 3.2—whose net sign is itself ambiguous and potentially a material share of the estimated protection—cannot be excluded. This is the boundary set out there as scope: a robust protective direction, not yet a fully identified effect.

The path from one to the other is a definite program—a symmetric not-yet-treated leads reference, a salvage-filter fix, and re-estimation of a clean, powered, out-of-window, immortal-time-cancelled lead—which I set out as a forward agenda rather than apologize for its absence, and develop in full in Section G.1. Until that runs, the ceiling stands where I have set it.

B.6 Sensitivity and Bounding

The natural way to make the residual-confound worry concrete is to ask how strong an unobserved confounder would have to be to overturn the result, and the E-value of VanderWeele and Ding (VanderWeele and Ding, 2017) answers exactly that. On the pooled estimate, an unmeasured confounder would have to be associated with both harvest siting and first-fire hazard by a risk ratio of at least 2.14—over and above everything the observed siting covariates already capture—to explain the point estimate away entirely, and by at least 1.43 to move the confidence bound to the null. I read these as bounding targets rather than as vindication. The E-value is the appropriate sensitivity statistic here because it is valid for a rare-event outcome, which a first fire in a given stand-year is. A threshold of 2.14 sets a concrete hurdle: an unmeasured confounder would have to move both harvest siting and first-fire hazard by a risk ratio above two, over and above everything the observed covariates already absorb, to overturn the estimate. That is a substantial association—on the order of the strongest siting differences the design already conditions on, the raw structure imbalances that run to standardized differences near 0.9 (Figure 1)—not a bar a reader should assume is trivially cleared. But it is not exotic either, which is precisely why a same-signed siting confound that observable adjustment cannot reach—potentially a material share of the estimated protection—cannot be ruled out (Section B.5): a confound of that strength is conceivable, and the design has not yet excluded it. The bound and the pre-trend verdict agree, and they set the paper's honest ceiling together.

Taken one at a time, none of the checks in this section identifies the effect, and I have been explicit about the one that would—the joint pre-trend test the leads cannot yet deliver. Taken together they are still worth weighing as a conjunction. The protective sign is what survives observable adjustment, every clustering scheme, the removal of the extreme fire-years, a clean closest lead, and a confounder hurdle of 2.14; a siting story that explains the result away has to clear all of them at once, in the same direction, without tripping the one pre-trend check that does come back clean. That is not identification, and I do not dress it as such. It is why the direction is robust enough to reorient the question even while the effect remains unidentified—the ceiling I set in Section B.5 and hold to here.

B.7 Heterogeneity by Ecozone

A single national hazard ratio invites the question of whether it describes the whole country or one corner of it, and the natural way to look is to refit the preferred specification inside each ecozone and read the map. Doing so is more disclosure than test, because the fire record thins quickly once it is cut ten ways: of the ten ecozones, only four carry enough treated-and-burned events to be estimated reliably, and I report those four in Figure 4 and set the rest aside rather than read noise into them. The one ecozone estimated with real precision is the Boreal Shield, the dominant boreal fire regime, at a hazard ratio of 0.545 with a 95% confidence interval of [0.360, 0.826] —below the national pooled value and well clear of one—resting on roughly 15,400 first-fire events, the largest count of any ecozone. That the national number and the Boreal Shield number sit in the same protective range is not a coincidence: the boreal is where most of the country's fire and most of its harvest coincide, so the pooled estimate is anchored by, and largely a reflection of, this one regime.

The other three estimable ecozones lean the way the national result does without resolving on their own. Montane Cordillera and Boreal Plains come in protective at hazard ratios of 0.82 and 0.85, and Taiga Plains nominally adverse at 1.17, but all three intervals include one, so none of them individually adjudicates anything. What matters for reading the national estimate is the shape of that spread rather than any one point in it: no ecozone reverses the sign with precision, and in particular nothing adverse clears its own interval. The protective effect neither holds uniformly across the map with precision—most non-boreal ecozones are simply too fire-sparse at the harvested margin to say—nor reverses cleanly anywhere; it is concentrated in and carried by the boreal, and the ecozone cut is texture on the national result rather than a revision of it.

The six ecozones I exclude are excluded by a rule fixed in advance, so that their omission is not mistaken for the suppression of an inconvenient pattern. Two—Taiga Cordillera and Hudson Plains—are quasi-separation artifacts: within their fixed-effect cells harvested stands essentially never burn inside the window, so the cloglog likelihood runs the log-hazard-ratio to the boundary and the fitted hazard ratio collapses toward zero. I flag any fit with $|\log \mathrm{HR}| > 3$ as separated and read neither the collapsed point nor its interval as an effect. The other four—Taiga Shield, Boreal Cordillera, and the Pacific and Atlantic Maritime zones—have standard errors on the log-hazard-ratio scale above 0.5, wide enough that their intervals span an order of magnitude, so I set them aside as too imprecise to interpret. Their exclusion is, if anything, conservative for the national reading: the two separated zones collapse in the protective direction, and of the four wide-interval zones two sit nominally protective with intervals entirely below one (Taiga Shield at a hazard ratio of 0.08, Boreal Cordillera at 0.05) while only two straddle it. A note on the sparsity: fire-year clustering yields thirty-nine clusters in every ecozone, because the calendar span is full whichever ecozone one conditions on, so the binding constraint on precision is the count of treated-and-burned events and the incidence of fixed-effect separation, never the number of clusters.

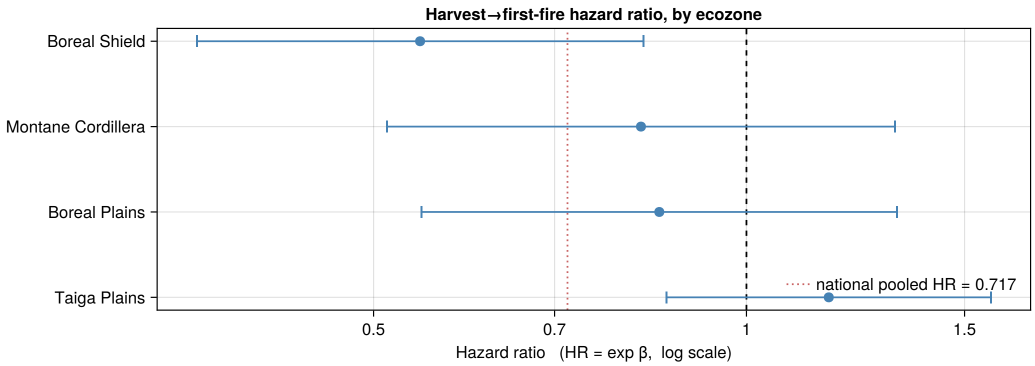

Figure 4: Harvest–wildfire hazard ratio estimated separately within each ecozone, for the four ecozones with a reliably estimable fit. Each point is the ecozone-specific hazard ratio from the preferred complementary-log-log specification refit on that ecozone alone; the horizontal bar is its 95% fire-year cluster-robust confidence interval; the horizontal axis is on a log scale. The dashed vertical line marks HR = 1 and the dotted line the national pooled hazard ratio (0.717). Boreal Shield, the dominant boreal fire regime, is the single precisely estimated ecozone (HR 0.545, interval [0.360, 0.826]) and sits below the national value; the other three credible ecozones lean protective or mildly adverse with intervals that include one. Six further ecozones are omitted as non-interpretable—four with standard errors large enough that their intervals span an order of magnitude, and two with quasi-separation artifacts in which harvested stands almost never burn inside the window within their fixed-effect cells. The omitted zones, if anything, lean protective.

C Data Construction Details

The cell-level panel the event study consumes is built in two steps that Section 2 describes in outline and that I record here in full. The first draws the sample. Rather than carry the full national pixel grid, I subsample it onto a strided spatial lattice: within each read row-band I keep every hundredth pixel column, placing sampled columns about 3 km apart, and I read only a periodic subset of the row-bands. The result is an unbiased spatial sample of the managed forest in which the sampled cells sit far enough apart along a row that two cells cut in one harvest operation are essentially never both drawn, so the lattice does not stack near-identical neighboring pixels as though they were independent observations. Cells enter the estimation unweighted, so the estimand is the average treatment effect on the average retained cell.

The second step brings the covariate layers into a common frame. Treatment (CanLaD), the independent fire outcome (NBAC), forest structure (SCANFI), terrain (the Canadian Digital Elevation Model), and land cover (VLCE2) are native to different projections, grid origins, and resolutions, so a single point on the ground lands at different pixel coordinates in each. I warp every layer onto the CanLaD 30 m national reference grid—178 400 columns by 119 100 rows, roughly 21 billion pixels (178,400 × 119,100) spanning forested Canada—with GDAL's `gdalwarp`, resampling continuous fields such as elevation bilinearly and categorical fields such as land cover by nearest neighbor, so that a given pixel index denotes the identical patch of ground in every layer. Warping the full national grid is a one-time but heavy operation, which is why it dominates the pipeline's setup cost. The structure covariates are read from a single fixed epoch: the SCANFI 1985 layer, which is strictly earlier than every cohort's harvest year (cohorts begin in 1986 or later), so it is pre-treatment for all units by construction. Terrain is time-invariant. One caveat belongs on the record: SCANFI is a modeled product whose native base is a recent inventory, so its 1985 layer is a Landsat-time-series *backcast* of 1985 structure rather than a direct 1985 observation; I treat it as a fixed pre-treatment baseline while noting that it is an estimate of the past, not a measurement of it. The warped, cell-level panel on that fixed pre-treatment baseline is the object Equation (1) is fit to.

D Estimator Details

The headline is a staggered-adoption discrete-time complementary-log-log first-fire hazard with covariate adjustment and two high-dimensional fixed effects, estimated in R with `fixest::feglm` under the `cloglog` link. The first-fire structure lets me aggregate the stand-year risk set to sufficient statistics—counts of at-risk stand-years and first-fire events within each covariate and fixed-effect cell—rather than carry the full 10^{10}-row pixel panel, which is what makes a national fit tractable; the two fixed effects, ecozone-by-year and coarsened fireshed, are absorbed by alternating projection, and the fit converges. An earlier version of the project used a bespoke sparse aggregation device; it has been retired in favor of this standard estimator, against which the aggregation was checked for numerical equivalence on the shared cells.

The fixed effects buy their spatial control at a cost the estimand pays openly. A fixed-effect group that contains no fire at all is perfectly separated in a hazard model: it carries no within-cell outcome variation and contributes nothing to a first-fire likelihood, so it is dropped. In a country where most stand-years never burn, that is a large share of the panel. Table 6 reports the accounting. The ecozone-by-year absorber separates 127 of its 390 groups and the fireshed absorber 86 of its 296, and together they drop 26.9% of stand-years and 29.2% of treated rows from the estimation sample. This is the basis for the "ATT among fire-experiencing cells" scope statement of Section 3.4, and I document it as scope rather than as a defect: the dropped cells are the ones over which a within-fireshed first-fire contrast is not even defined. Inference throughout is cluster-robust on fire-years with a $t(G-1)$ reference ($G = 39$), the conservative small-cluster block for season-wide fire shocks; the fireshed and two-way alternatives of Section B.1 leave the point estimate unchanged and move only the interval width. The wild-cluster bootstrap and national CR2 correction that would further stress the modest cluster count are part of the forward program (Section G.1).

Table 6: Fixed-effect separation and dropped observations. Rows report, for each fixed effect (ecozone×year and coarsened fireshed) and their combination, the number of groups, the number of fully-separated (zero-fire) groups, and the resulting share of stand-year rows dropped from the estimation sample, with the share of *treated* rows dropped. Perfectly-separated cells—those with no fire and hence no within-cell outcome variation—are removed by the fixed-effect estimator; the combined drop of roughly 27% of stand-years defines the estimand as the ATT among fire-experiencing cells (Section 3.4).

FE dimension	Groups	Separated	Rows dropped	Treated dropped
ecozone×year	390	127	16.7%	19.3%
fireshed	296	86	15.7%	18.4%
combined (lower bound)	—	—	26.9%	29.2%

E Full Robustness Battery

The body carries the decisive leave-out (dropping 2023) and the leave-out that is read asymmetrically is reported here, alongside the one control perturbation that probes the estimate's dependence on a single covariate. Table 7 reports both. Row R2 drops the entire 2018–2024 high-fire window and moves the estimate to a hazard ratio of 0.554; I read it asymmetrically, because dropping that window removes about a fifth of the fire-years and cuts G from 39 to 32, so a *null* in that thinned sample would have been indeterminate rather than a failure, and the strengthened point is consistent with the association without independently confirming it. Row R15 adds the Tier-B prior-fire indicator—the covariate most likely to act as a collider and pull the estimate the other way—and the hazard ratio is essentially unchanged at 0.719, so the collider-risk control is not load-bearing and the headline does not balance on it. The pending falsification and corroboration rows (suppression-gradient rebuttal, windthrow placebo, independent estimators) are named once in Section G.1 and are not tabulated here.

Table 7: Appendix leave-out and control-perturbation robustness of the pooled harvest–wildfire estimate. Rows report the preferred specification after dropping the 2018–2024 high-fire window (R2) and after adding the Tier-B prior-fire indicator (R15). Columns report the hazard ratio, its 95% fire-year cluster-robust confidence interval, and the t statistic. R2 removes roughly one-fifth of the fire-years and a large share of treated fires, so its strengthened point is read asymmetrically—indeterminate rather than corroborating (Section B.2); R15 leaves the estimate essentially unchanged, so the collider-risk prior-fire control is not load-bearing (Section B.4).

Specification	HR	95% CI (HR)	t
Leave out 2018–2024	0.554	[0.37, 0.83]	-3.00
+ prior-fire indicator	0.719	[0.57, 0.91]	-2.83

F Additional Diagnostics

The within-fireshed common-support claim of Section B.3 rests on the fireshed-by-fireshed overlap in Figure 5. Each point is a fireshed, placed by the number of treated stands it holds against the share of those stands on marginal within-fireshed support, and the mass of the distribution sits at the top: 197 of 201 firesheds contribute, holding 99.5% of treated stands, with the four non-contributing firesheds tiny treated-only cells of at most seven stands each. The figure also makes the diagnostic's limit legible. It is a marginal, bounding-box test—it certifies that each treated stand's covariate *ranges* lie inside its fireshed's control ranges, a necessary condition for a defined contrast, but it cannot see holes in the interior of the joint covariate space—which is why I describe the result as marginal within-fireshed support and stop short of declaring positivity a non-issue. The companion pre-treatment covariate-balance figure is carried in the design section (Figure 1), where it motivates the conditioning strategy.

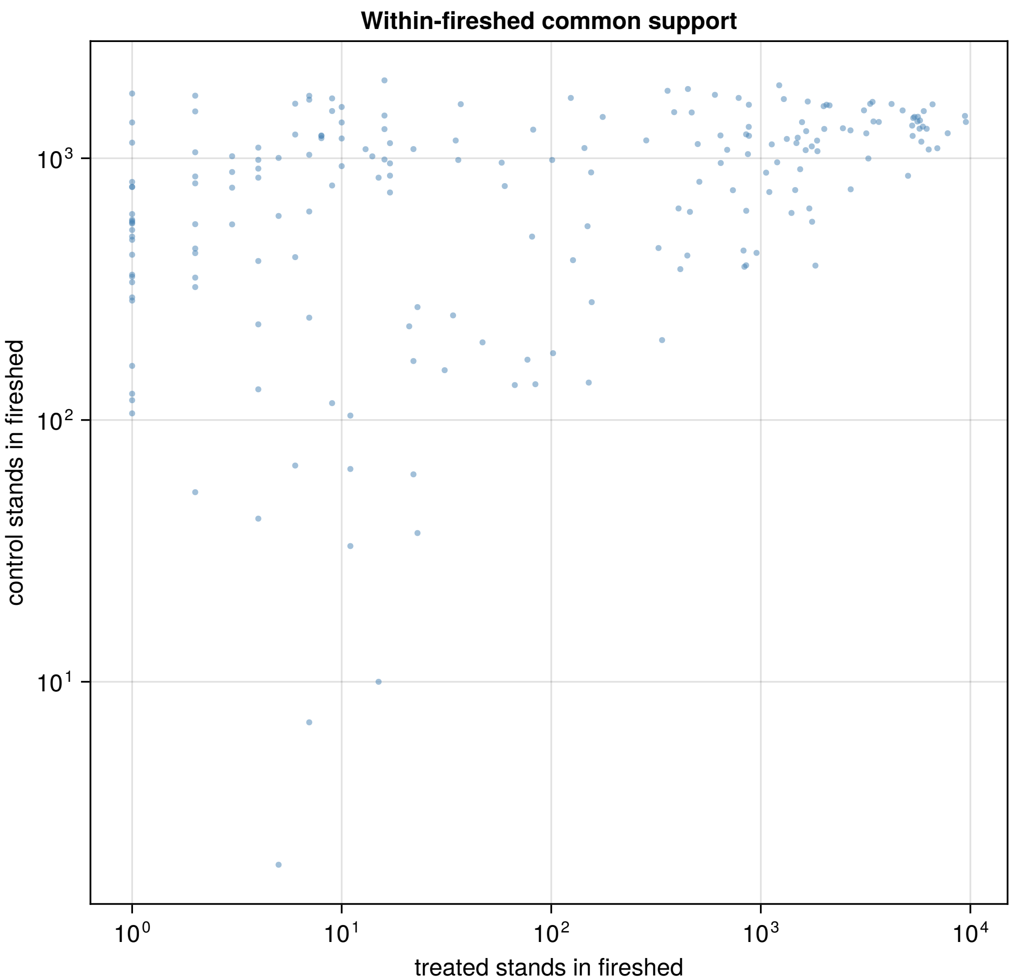

Figure 5: Marginal within-fireshed common support. Each point is a fireshed; the horizontal axis is the number of treated stands and the vertical axis the share of those treated stands whose pre-treatment covariate values fall within the covariate *ranges* (bounding box) spanned by control stands in the same fireshed. 197 of 201 firesheds contribute treated stands on marginal support, covering 99.5% of treated stands; the four non-contributing firesheds hold at most seven treated stands each. This is a marginal test of overlap: it certifies per-covariate range containment, not joint support, and cannot detect joint-support holes.

G AI Usage Disclosure

Large language models were used to assist in the coding and drafting of this paper, including but not necessarily limited to writing code, providing feedback on empirical specifications, summarizing the literature, and providing feedback on and editing the written draft. I take full and unconditional responsibility for all content of the paper and code. This paper is preliminary and has not been peer-reviewed.

G.1 Planned Extensions