

A Study of the Drossel–Schwabl Forest-Fire Model: Local Connectivity and Fire-Size Statistics

Aryan Mago
aryanmago05@gmail.com

ABSTRACT

Large wildfires impose substantial ecological and societal costs, motivating simplified models that isolate when small ignitions escalate into large burns. We study a two-dimensional stochastic forest-fire model with regrowth and lightning ignition, in an event-based formulation where each lightning event is preceded by θ random regrowth attempts, so θ controls pre-ignition density and connectivity. Beyond global density we evaluate a local predictor: tree density in the 9×9 neighbourhood around the strike. Across 72 independently seeded runs the median fire size versus local density curve shows an uptick threshold at $\rho^*_{\text{loc}} = 0.606 \pm 0.006$, resolving to one of two adjacent bins in every run and corresponding to roughly 48 of the 80 cells neighbouring the ignition site. The threshold is stable across regrowth rate, layout and system size, while its steepness grows 4.2-fold from $\theta = 250$ to 3000 and large-fire probability per lightning event rises from 0.028 to 0.196. Roughly 91% of that rise comes from fires becoming more severe once started rather than from lightning striking fuel more often. Initial layout has no detectable effect on steady-state risk once each run has an independent random stream, with between-layout variation three to nine times smaller than seed-to-seed variation. We trace this to a memory-erasure timescale that collapses on $\phi = \theta/N^2$, falling from beyond 20,000 events at $\phi \leq 0.031$ to under 100 at $\phi = 0.183$. The same variable governs finite-size comparison, where discrepancies up to $54.6\times$ at fixed θ collapse to $1.00\text{--}1.11\times$ at matched ϕ .

Keywords: forest-fire model; self-organized criticality; percolation; connectivity threshold; heavy-tailed events; finite-size scaling

1. INTRODUCTION

Wildfire risk exhibits strong nonlinearities: many ignitions self-extinguish while a small fraction become large, difficult-to-control events. The scientific question is therefore not only how often ignition occurs, but under what conditions an ignition accesses a sufficiently connected fuel network to produce a large burn. Simplified computational models provide a controlled setting in which such threshold mechanisms can be isolated and measured. Empirical burned-area distributions have often been reported as heavy-tailed, which motivates distributional summaries such as complementary cumulative distribution functions while cautioning against over-interpreting any single functional form (Malamud et al., 1998; Millington et al., 2006; Clauset et al., 2009; Newman, 2005). Models of this class are commonly studied within the broader framework of self-organized criticality and scale-free event statistics (Bak et al., 1987; Bak, 1996; Jensen, 1998).

Percolation theory formalizes a closely related idea: when occupied sites on a lattice exceed a critical density, system-spanning connected clusters become likely (Stauffer & Aharony, 1994; Grimmett, 1999; Saberi, 2015). Under a wildfire interpretation, connected clusters act as connected fuel. Percolation results, however, are typically stated in terms of a single global occupation density on an idealized lattice, whereas real landscapes are patchy and spatially heterogeneous, and a fire begins at one specific point rather than everywhere at once. This raises a question that global density alone cannot answer: does the fuel configuration in the immediate neighbourhood of an ignition predict the size of the resulting event, and if so, is there a reproducible threshold?

The present work addresses that question using a stochastic forest-fire model with regrowth and lightning, building on the Drossel–Schwabl model and subsequent analyses (Drossel & Schwabl, 1992; Clar et al., 1994; Henley, 1993; Pruessner & Jensen, 2002; Grassberger, 2002). Much of that literature is concerned with the asymptotic critical behaviour of the model and with whether its scaling is genuinely critical; Pruessner and Jensen (2002) and Grassberger (2002) in particular report broken or non-universal scaling in the asymptotic limit. Our aim is deliberately different and more modest. Rather than adjudicating asymptotic criticality, we ask what a finite, observable, locally measurable quantity, namely the density of fuel immediately surrounding an ignition point, predicts about event size at accessible system sizes.

This paper makes four specific contributions. *First*, we identify a local-density threshold at ignition, $\rho^*_{\text{loc}} = 0.606 \pm 0.006$ in a 9×9 window, and show that it is stable across regrowth scales, initial layouts and system sizes, while its steepness is not. *Second*, we separate the two distinct ways in which the regrowth parameter θ raises risk, namely the probability that lightning finds any fuel and the probability that a started fire becomes large, and quantify each. *Third*, we measure the timescale on which the model erases memory of its initial condition, show that it collapses on $\phi = \theta/N^2$, and use it to explain why initial spatial layout has no detectable effect on steady-state risk. *Fourth*, we show that the practice of comparing forest-fire results at fixed θ across system sizes is invalid for this regrowth mechanism, and demonstrate a collapse onto ϕ that reduces cross-size discrepancy from as much as $54.6\times$ to $1.00\text{--}1.11\times$.

2. METHODS

2.1 Lattice, states, and neighbourhood

The forest is represented by a two-dimensional $N \times N$ square lattice. Each site $x = (i, j)$ occupies one of four states: EMPTY = 0, TREE = 1, BURNING = 2, and BURNT = 3. In the baseline model, burning is treated as instantaneous at ignition: if lightning strikes a tree, the entire four-neighbour connected component of TREE sites containing the ignition site is converted to BURNT. Connectivity is defined by a von Neumann (four-neighbour) rule with open boundaries and no periodic wrapping unless otherwise stated. Figure 1 summarizes the full event cycle. Instantaneous cluster removal is a standard simplification in this model class and is what distinguishes it from cellular-automaton fire-spread models, which resolve the propagating front in time (Karafyllidis & Thanailakis, 1997).

2.2 Event-based regrowth and lightning

Time is indexed by lightning events $\ell = 1, 2, \dots, L$. Each lightning event consists of two stages. In the regrowth stage, θ independent regrowth attempts are performed; each attempt selects a uniformly random site, and if that site is EMPTY or BURNT it is set to TREE, otherwise it is left unchanged. In the lightning stage, one uniformly random site (i_0, j_0) is selected; if it is TREE, the entire four-neighbour connected tree cluster containing it is burned and those sites are set to BURNT, and if it is not TREE the fire size is recorded as $S = 0$.

This event-based formulation differs from the original micro-step p/f formulation of Drossel and Schwabl (1992), but the mapping is direct. Per lightning event, the probability that a given cell is selected for regrowth in at least one of the θ attempts is approximately θ/N^2 when $\theta \ll N^2$. This motivates the scaled regrowth variable

$$\varphi = \theta / N^2$$

which is used for all cross-size comparisons in Section 3.4. Its adequacy is tested rather than assumed.




Figure 1. Schematic of the event-based forest-fire model. Each lightning event is preceded by θ random regrowth attempts, then a random ignition. If ignition hits a tree, the entire four-neighbour connected tree cluster is removed (burnt), producing event size S . Local density ρ_{loc} is measured in a 9×9 neighbourhood around the ignition site.

2.3 Local density at ignition

Let $W(i_0, j_0)$ denote the 9×9 window centred on the lightning site, clipped at boundaries for open edges. The ignition cell itself is excluded, so the density is computed over the surrounding cells only — 80 of them in the interior, fewer at a clipped boundary:

$$\rho_{loc}(i_0, j_0) = [\sum \mathbf{1}\{X_{ij} = \text{TREE}\} - \mathbf{1}\{X_{i_0 j_0} = \text{TREE}\}] / (|W(i_0, j_0)| - 1)$$

Excluding the ignition cell matters for interpretation: ρ_{loc} measures the fuel available around the strike, not partly the strike site itself, so it stays uncorrelated with whether a fire starts at all. The quantity lies in $[0, 1]$ and serves as a proxy for local connectivity — higher ρ_{loc} increases the likelihood that the ignition site is embedded in a large connected tree cluster. The 9×9 window is a design choice; it is large enough to capture cluster structure beyond the immediate four neighbours while remaining local relative to the lattice, and we did not tune it against outcomes.

2.4 Initial layouts

Simulations are initialized with a target global tree density t and one of four spatial layouts. The uniform random layout sets each site to TREE independently with probability t . The clustered layout thresholds a smooth random field built from Gaussian blobs to match t . The stripes layout alternates dense and sparse vertical stripes adjusted to match t . The gradient layout imposes a left-to-right tree-probability gradient adjusted to match t . Because regrowth is spatially uniform, these layouts are expected to lose memory

over long times; whether any layout effect survives into steady state is tested in Section 3.3 rather than assumed.

2.5 Parameter sweeps and logging

Each lightning event produces a record containing the global density before and after lightning, the local density ρ_{loc} in the 9×9 window, and the fire size S , defined as the number of trees burned in the connected cluster. Additional geometry metrics (perimeter, radius of gyration, bounding-box aspect ratio, spanning flags) were logged for extensibility but are not analysed here. The first 20% of events in every run are discarded as transient; all results reported below are post-transient.

The main sweep uses system sizes $N = 128$ plus $N = 256$ for finite-size comparisons; regrowth scales $\theta \in \{250, 500, 1000, 2000, 3000\}$; initial densities $t = 0.50, 0.55$ and 0.60 ; and all four layouts. The $N = 128$ sweep covers all three initial densities, while $N = 256$ was run at $t = 0.55$ only, giving 80 configurations. Each configuration comprises two runs at different seeds, of 200,000 lightning events each, so 40,000 events are discarded per run and each configuration contributes 320,000 post-transient events, of which 126,594 to 154,291 produce a non-zero fire. Binomial standard errors are therefore below 0.0009 for p_{fire} and below 0.0008 for the unconditional exceedance probabilities. Conditional exceedance probabilities are estimated over the smaller fire-only subsample and carry standard errors up to 0.0013; this is small relative to every effect discussed below, but it is not below 0.001 and we do not claim otherwise.

2.6 Random-number streams and the layout re-run

In the main sweep above, a single random stream per run serves both the construction of the initial lattice and the subsequent event loop. This is not a safe arrangement here. The uniform, stripes and gradient initializers each draw one variate per lattice site, whereas the clustered initializer draws only twelve, so runs that differ only in layout enter the event loop at different — and in three of four cases identical — positions in the stream. Three of the four layouts therefore receive the same sequence of regrowth sites and lightning coordinates. Combined with the memory-erasure effect, measured by the procedure of Section 2.7 and quantified in Section 3.3, this makes those three runs converge onto a single trajectory, so a comparison between them measures nothing.

All layout comparisons reported in Section 3.3, and the threshold statistics reported in Section 3.1, therefore come from a separate, corrected set of runs in which the initialization stream and the event stream are drawn independently, and the event stream is additionally keyed on the layout, so that no two runs can share it. That set comprises $N = 128$, $t = 0.55$, $\theta \in \{250, 1000, 3000\}$, all four layouts and six independent seeds per layout: 72 runs of 40,000 events each, of which the first 20% are discarded, giving 32,000 post-transient events per run and 192,000 per (layout, θ) cell. Six seeds per cell provide a within-layout variance against which any between-layout difference can be tested — the quantity the original design lacked. Where the two sets overlap they agree closely: the corrected runs reproduce the

Table 1 values for p_{fire} and the exceedance probabilities to within 0.002 at all three θ , and to within 0.001 at $\theta = 250$ and $\theta = 1000$, so the risk results of Sections 3.2 and 3.4 stand as reported.

One quantity does change. The main sweep returns an identical threshold in every one of its runs, with zero variance, which we now attribute to the degeneracy rather than to the physics; the corrected runs return $\rho^*_{\text{loc}} = 0.606 \pm 0.006$, spread over two adjacent bins. Section 3.1 reports the corrected value. Figure captions state which set each panel is drawn from, since the two coexist in the paper.

2.7 Measuring the memory-erasure timescale

To quantify how long the model retains information about its initial condition, we initialize two lattices with different layouts, drive both with a single shared random stream so that they receive identical regrowth sites and identical lightning coordinates at every event, and record the first event at which the two lattices become bit-identical. Once that happens the two runs have converged onto a single trajectory and remain identical thereafter, so the crossing event is a direct measurement of the erasure time. Using a common stream is essential: it removes the driving noise as a source of difference, so any remaining difference between the lattices is attributable to the initial condition alone.

We measure this for three layout pairs, each comparing uniform against one of stripes, gradient and clustered, at five independent seeds per pair, across seven values of θ at $N = 128$ and seven at $N = 256$, chosen to span overlapping ranges of $\phi = \theta/N^2$. This gives 210 runs in total. Lattices are compared at every event up to event 30 and every tenth event thereafter, so a reported erasure time beyond event 30 is an upper bound accurate to ten events, which is immaterial at the timescales observed. Runs are capped at 20,000 events; a run reaching the cap without converging is recorded as censored, and we report such configurations as bounded below rather than assigning them a value.

2.8 Risk summary metrics

We summarize how often lightning causes any fire, and how often fires become large, using $p_{\text{fire}} = P(S > 0)$, and, for $\alpha \in (0, 1)$ with $S_\alpha = \alpha N^2$, the unconditional and conditional exceedance probabilities $P(S \geq S_\alpha)$ and $P(S \geq S_\alpha \mid S > 0)$. The unconditional quantity is the operationally meaningful one, being a probability per lightning event, while the conditional quantity separates how often lightning finds fuel from how severe a fire becomes once started. Reporting both is what allows the two mechanisms to be disentangled in Section 3.2.

2.9 Threshold extraction

Events are binned by ρ_{loc} over $[0, 1]$ using $B = 60$ equal-width bins, giving a bin width of $1/60 = 0.0167$ and an estimator resolution of ± 0.008 . The bin count is a balance between two competing pressures: too few bins would blur the threshold below the precision we report, while too many would leave individual bins with too few events to yield a stable median. At $B = 60$ the populated bins near the threshold each hold thousands of events, so the median is well determined within a bin, and the 0.0167 width is the

resulting resolution rather than a free parameter. We confirmed the extracted threshold is unchanged at $B = 50$ and $B = 80$, so the reported location is not an artefact of this choice; the resolution simply scales as $1/B$. Bins holding fewer than 200 events are discarded as too sparse for a stable median, which leaves 50 to 54 populated bins per run in the main sweep and 41 to 45 in the shorter independently seeded runs; the discarded bins lie at the extremes of the ρ_{loc} range, where few ignitions occur. For each retained bin centre we compute the median fire size and the interquartile range.

The uptick threshold ρ^*_{loc} is extracted by a baseline-crossing rule. A baseline is taken as the median of the binned medians over the low-density region $\rho_{\text{loc}} \leq 0.4$, where fire size is flat. A crossing tolerance is set to $\delta = \max(50, 3 \times \text{IQR})$ trees, where the IQR is that of the same baseline bins, so the tolerance widens if the baseline is noisy. The threshold is then the centre of the first bin, scanning upward in ρ_{loc} , whose median exceeds $\text{baseline} + \delta$, subject to at least two valid bins on either side. The knee strength is the slope of the median curve across the retained bins two positions to either side of the threshold in the list of retained bins, $(S_{\{b+2\}} - S_{\{b-2\}}) / (x_{\{b+2\}} - x_{\{b-2\}})$, and is reported in trees per unit local density.

Two properties of this estimator should be kept in view when reading Section 3.1. It locates the onset of the rise rather than its steepest point, so it is deliberately conservative and returns the lowest density at which fire size departs from baseline. And because it is extracted from binned medians, ρ^*_{loc} is discretized to the bin step, so the threshold can only be resolved to ± 0.008 regardless of how many events are collected.

3. RESULTS

3.1 A stable local threshold in ignition neighbourhoods

Across the 72 independently seeded runs, the uptick in median fire size versus local density is located at

$$\rho^*_{\text{loc}} = 0.606 \pm 0.006$$

where the quoted figure is the mean over runs and the uncertainty is the standard deviation across them. The estimate resolves to one of two adjacent bin centres in every run — 0.608 in 86% of runs and 0.592 in the remaining 14% — so the threshold varies by at most one bin under any manipulation applied, including changes of regrowth rate, initial layout and seed. Because ρ_{loc} excludes the ignition cell and is computed over the 80 surrounding cells, this corresponds to roughly $0.606 \times 80 \approx 48$ trees adjacent to the strike. Mechanistically, below about 48 neighbouring trees typical burns remain small, whereas above that level ignition increasingly accesses connected fuel clusters capable of much larger burns.

Two cautions apply. First, the resolution floor is the bin width of 0.0167, so the observed one-bin spread is the smallest non-zero variation this estimator can express; the data bound the variation rather than demonstrate exactness, and a finer binning study or a parametric fit to the median curve would be needed to resolve the threshold further. Second, the estimator locates the onset of the rise, not its steepest point

(Section 2.9), so ρ^*_{loc} should be read as the lowest local density at which fire size departs from baseline rather than as the midpoint of the transition. No bootstrap interval is claimed; the reported dispersion is across independent runs, which is the more informative quantity in any case.




Figure 2. Median fire size versus local density for $\theta = 250$ (post-transient, $N = 128$, $t = 0.55$, uniform layout). The dashed line indicates the extracted uptick threshold ρ^*_{loc} . Drawn from the main sweep of Section 2.5.




Figure 3. Median fire size versus local density for $\theta = 3000$ (post-transient, $N = 128$, $t = 0.55$, uniform layout). The threshold location is unchanged relative to Figure 2, while the post-threshold rise is markedly steeper. Drawn from the main sweep of Section 2.5.




Figure 4. Extracted threshold ρ_{loc}^* for all 72 independently seeded runs ($N = 128$, $t = 0.55$). Each point is one run, jittered horizontally; the short horizontal bars are layout means. Every run returns one of just two adjacent bin centres, 0.592 or 0.608; at $\theta = 1000$ all 24 runs return 0.608. The dotted line is the overall mean of 0.606 and the dashed line the square-lattice site-percolation threshold $p_c \approx 0.593$. Drawn from the corrected runs of Section 2.6.

Although the threshold location is stable, its steepness is not. Across the independently seeded runs the knee strength rises monotonically from 1952 ± 368 at $\theta = 250$ to 3893 ± 841 at $\theta = 1000$ and 8145 ± 4469 at $\theta = 3000$, a factor of 4.2 over the range, in trees per unit local density. The dispersion grows sharply with θ : at $\theta = 3000$ the standard deviation is more than half the mean, because once the lattice is densely connected the median curve rises almost vertically and the slope estimate becomes sensitive to which bins clear the count threshold. The direction and approximate magnitude of the effect are secure; the precise factor is not, and we quote it to one decimal place only. The transition therefore sharpens substantially as regrowth increases even though it does not move, which is visible in the contrast between Figures 2 and 3, while Figure 4 confirms that the threshold location itself is flat across the same θ range. This separation between the location of a threshold and the severity of the transition across it is the central empirical pattern of the paper.

3.2 Regrowth raises risk through two distinguishable channels

As θ increases, lightning is more likely to strike a tree at all, and a fire that starts is more likely to become large. These are separable effects, and they are not of comparable magnitude. Over the range $\theta = 250$ to $\theta = 3000$ at $N = 128$, averaged across layouts, the probability that a lightning event produces any fire rises

only modestly, from $p_{\text{fire}} \approx 0.405$ to 0.482 , an increase of 19%. Over the same range the unconditional probability of a large fire burning at least 10% of the lattice rises from 0.028 to 0.196, a factor of 7.0 (Figure 5). The conditional probability $P(S \geq 0.10N^2 \mid S > 0)$ rises from 0.069 to 0.407, a factor of 5.9. Repeating the analysis at the lower large-fire threshold $\alpha = 0.05$ gives the same qualitative picture (Figure 6), indicating the result is not an artefact of the particular cut-off chosen.

The arithmetic is informative. Because the unconditional probability is the product of p_{fire} and the conditional probability, the total 7.0-fold rise decomposes multiplicatively into a factor of 1.19 from p_{fire} and a factor of 5.9 from the conditional probability. Taking logarithms to convert this into additive shares, $\ln(5.90)/\ln(7.01) = 0.91$, so roughly 91% of the increase in large-fire risk is attributable to fires becoming more severe once started and only about 9% to lightning finding fuel more often. Increasing regrowth does not primarily make ignition more likely; it makes the fuel network that an ignition reaches more connected. This is consistent with the threshold result of Section 3.1, in which a fixed local-density threshold is crossed more frequently and more decisively as θ grows. Table 1 reports the full sweep.

Table 1. Risk summary for $N = 128$ (post-transient, averaged across the four layouts and all three initial densities). Large-fire thresholds use $S_{\alpha} = \alpha N^2$. Each cell aggregates 320,000 post-transient lightning events per configuration. Values are identical to three decimal places whether $t = 0.50, 0.55$ or 0.60 is used alone, except p_{fire} at $\theta = 500$, which differs by 0.001. Because of the stream degeneracy described in Section 2.6, the rows at $\theta \geq 1000$ average two distinct trajectories rather than four; the independently seeded runs of Section 2.6 reproduce the $\theta = 1000$ and $\theta = 3000$ rows to within 0.002; $\theta = 2000$ was not re-run.

θ	p_{fire}	$P(S \geq 0.05N^2)$	$P(S \geq 0.10N^2)$	$P(S \geq 0.10N^2 \mid S > 0)$
250	0.405	0.061	0.028	0.069
500	0.418	0.098	0.065	0.155
1000	0.434	0.134	0.107	0.246
2000	0.457	0.175	0.155	0.340
3000	0.482	0.211	0.196	0.407




Figure 5. Unconditional exceedance probability per lightning event for large fires at $\alpha = 0.10$, that is $P(S \geq 0.10N^2)$, versus θ (post-transient, $N = 128$, $t = 0.55$). Drawn from the main sweep of Section 2.5.




Figure 6. Unconditional exceedance probability at $\alpha = 0.05$, that is $P(S \geq 0.05N^2)$, versus θ (post-transient, $N = 128$, $t = 0.55$). The qualitative shape is unchanged from Figure 5, indicating the result is not an artefact of the particular large-fire threshold chosen. Drawn from the main sweep of Section 2.5.

3.3 Initial layout has no detectable effect, and the reason is measurable

Under independent seeding, initial spatial layout has no detectable effect on steady-state risk. Figure 7(a) shows the four layouts at each θ with seed-to-seed error bars; they are indistinguishable. Quantitatively, the standard deviation of the four layout means is smaller than the mean within-layout standard deviation across seeds at every θ tested: 0.00005 against 0.00047 at $\theta = 250$, 0.00031 against 0.00098 at $\theta = 1000$, and 0.00021 against 0.00124 at $\theta = 3000$, giving ratios of 0.11, 0.32 and 0.17 respectively. A one-way analysis of variance on $P(S \geq 0.10N^2)$ across layouts returns $F(3, 20) = 0.05, 0.55$ and 0.17 at the three θ values, all far below the value of 1 that would indicate between-group and within-group variation of comparable size. The threshold location is likewise unaffected: ρ^*_{loc} has the same mean across all four layouts to within one bin.

This is a genuine null result rather than an absence of resolution. Six independent seeds per layout, each contributing 32,000 post-transient events, resolve the exceedance probability to a pooled standard error below 0.001, almost two orders of magnitude smaller than the differences the θ sweep produces. A layout effect of the size θ produces would have been detected easily.

The mechanism behind the null is measurable directly, and it is more informative than the null itself. Using the procedure of Section 2.7, we take two lattices initialized with different layouts, drive them with a common random stream, and record the event at which they become bit-identical — the point at which the model has erased all memory of its initial condition. Figure 7(b) shows this erasure time against $\phi = \theta/N^2$. Three features stand out. It collapses on ϕ rather than on θ or N : at matched $\phi = 0.0916$ the erasure times are 360 events at $N = 128$ and 310 at $N = 256$, and at $\phi = 0.1831$ they are 60 and 70. It falls extremely steeply, from beyond 20,000 events at $\phi \leq 0.031$ to about 60 to 70 events at $\phi = 0.183$. And it does not depend on which pair of layouts is compared, with median erasure times of 200 to 275 events across the three pairs tested, confirming that the timescale is a property of the dynamics rather than of any particular initial condition.

The transition is sharp and lies near $\phi \approx 0.05$. Below $\phi = 0.031$ no run erased its initial condition within 20,000 events; at $\phi = 0.0458$ between a third and 40% of runs did; at $\phi \geq 0.061$ all runs did. This has a direct methodological consequence for any study of this model that discards a fixed transient. A transient long enough to reach steady state at high ϕ is long enough to destroy the very initial-condition information a layout comparison is trying to measure. In our main sweep, which discards 40,000 events, every configuration with $\phi \geq 0.061$ has erased its initial condition within the first 8,100 events, so no post-transient measurement there can carry any layout signal even in principle. The layout question is therefore not merely hard to answer at high ϕ ; it is ill-posed, and the null result above should be read as confirming that the model behaves as designed rather than as evidence that spatial structure is unimportant to wildfire risk. Answering the latter requires heterogeneity that persists — a fixed fertility mask or spatially heterogeneous regrowth probability — which this model does not provide.




Figure 7. (a) Unconditional exceedance probability by layout at each θ , with error bars showing the standard deviation across six independent seeds; the layouts are indistinguishable. (b) Events until two lattices initialized with different layouts, driven by a common random stream, become identical — the memory-erasure time — plotted against $\phi = \theta/N^2$. Filled markers show the median over non-censored runs with bars spanning their full range; open triangles mark ϕ values at which the majority of runs reached the

20,000-event cap without converging, and are lower bounds on the erasure time. The shaded band marks the transition region near $\phi \approx 0.05$. Both panels are drawn from the corrected runs of Sections 2.6 and 2.7.

3.4 Cross-size comparison requires $\phi = \theta/N^2$

At fixed θ , large-event probabilities are substantially smaller at $N = 256$ than at $N = 128$. This follows directly from the regrowth implementation: each lightning event distributes θ regrowth attempts over N^2 sites, so the effective per-cell regrowth rate is proportional to θ/N^2 , and quadrupling the lattice area at fixed θ quarters the per-cell regrowth. Comparing system sizes at fixed θ therefore compares systems in materially different dynamical regimes.

The magnitude of the resulting error is large, as Table 2 sets out. At $\alpha = 0.10$ and fixed $\theta = 250$, the $N = 128$ and $N = 256$ estimates differ by a factor of 54.6, and even at $\theta = 3000$, where both systems are well into the high-risk regime, they still differ by a factor of 2.3. When instead the two system sizes are compared at matched ϕ , the agreement is close to exact: at $\phi = 0.0305$ the two estimates are 0.0646 and 0.0647, a ratio of 1.00, and at $\phi = 0.0153$ they are 0.0280 and 0.0310, a ratio of 1.11. The same pattern holds at $\alpha = 0.05$, where matched- ϕ ratios are 1.01 and 1.07 against fixed- θ discrepancies of up to 8.1. We note that the fixed- θ rows at $N = 128$ with $\theta \geq 1000$, and the $N = 256$ entry at $\theta = 3000$, draw on the degenerate runs identified in Section 2.6, the latter because $\phi = 0.046$ there sits inside the transition region of Section 3.3. The matched- ϕ comparisons, which are the basis of the claim, use $\theta = 250$ and 500 at $N = 128$ and $\theta = 1000$ and 2000 at $N = 256$, and all four show four distinct per-layout trajectories.

Table 2. Cross-size agreement in $P(S \geq 0.10N^2)$ under two comparison schemes ($t = 0.55$, averaged across layouts). Matching on $\phi = \theta/N^2$ reduces the discrepancy between $N = 128$ and $N = 256$ from as much as $54.6\times$ to at most $1.11\times$.

Comparison held fixed	N = 128	N = 256	Ratio
$\phi = 0.0153$ (matched)	0.0280	0.0310	1.11×
$\phi = 0.0305$ (matched)	0.0646	0.0647	1.00×
$\theta = 250$	0.0280	0.0005	54.6×
$\theta = 500$	0.0646	0.0079	8.2×
$\theta = 1000$	0.1068	0.0310	3.4×
$\theta = 2000$	0.1551	0.0647	2.4×
$\theta = 3000$	0.1961	0.0851	2.3×




Figure 8. Finite-size scaling collapse. Panel (a) plots the unconditional exceedance probability against θ , where the $N = 128$ and $N = 256$ curves are clearly separated. Panel (b) plots the same data against $\phi = \theta/N^2$, where the curves lie close to a single master curve at both $\alpha = 0.05$ and $\alpha = 0.10$. Dotted vertical lines mark the two ϕ values at which the two system sizes overlap directly. Drawn from the main sweep of Section 2.5; the matched- ϕ points are unaffected by the degeneracy of Section 2.6.

Figure 8 shows this collapse directly. The two overlapping ϕ values available in the present sweep are a limited but genuine test: the sweep was designed as a θ grid rather than a ϕ grid, so only two ϕ values coincide between the two system sizes, and a dedicated ϕ -matched sweep across three or more system sizes would be needed to establish the collapse as a scaling law rather than a two-point agreement. What the data do establish is narrower and still useful, namely that ϕ and not θ is the correct control variable for cross-size comparison under this regrowth rule.




Figure 9. Extracted threshold location versus θ for both system sizes (post-transient, $t = 0.55$). Unlike the risk probabilities, ρ^*_{loc} requires no rescaling: it is a local quantity measured in a fixed 9×9 window and

is therefore unaffected by lattice size. Drawn from the main sweep of Section 2.5, in which the threshold estimate is degenerate across three of the four layouts (Section 2.6); the independently seeded runs cover $N = 128$ only and so cannot corroborate the cross-size comparison directly; they do confirm that the estimate is stable to within one bin at that size.

Notably, the local threshold itself needs no such correction. Figure 9 shows ρ^*_{loc} taking the same value at $N = 128$ and $N = 256$ at every θ , though we note this comparison is drawn from the main sweep, where the estimate is degenerate across layouts, so it demonstrates the absence of a size effect rather than measuring its magnitude. This asymmetry is instructive: quantities defined relative to the whole lattice, such as $P(S \geq \alpha N^2)$, inherit a system-size dependence from the regrowth mechanism and must be rescaled, whereas a quantity defined in a fixed local window does not. The invariance of ρ^*_{loc} across system sizes is thus not an additional coincidence but a consequence of how it is constructed, which strengthens rather than weakens its interpretation as a local mechanism.

4. DISCUSSION

The results support a two-part account of when ignitions become large events in this model. The onset of large fires is governed by a local connectivity condition that is stable under every manipulation tested: the median fire size begins rising once roughly 48 of the 80 cells surrounding the ignition point hold fuel, at $\rho^*_{\text{loc}} = 0.606 \pm 0.006$. The severity of the response beyond that condition is governed by the regrowth rate, which sharpens the transition roughly fourfold across the θ range without moving it. Location and steepness are controlled by different things.

This division is what makes the local-density measurement useful, and it is worth stating where it goes beyond a restatement of percolation. Classical percolation predicts that connectivity has a threshold in global occupation density (Stauffer & Aharony, 1994; Grimmett, 1999). The present result is a different and complementary statement: a locally observable quantity, measured in a small window around a specific point at a specific time, carries threshold information about the event initiated at that point. That is a stronger practical claim than a global one, because a global density is not something an observer at an ignition site has access to, whereas a local fuel density is. The value 0.606 is specific to the 9×9 window and the four-neighbour connectivity rule used here and should not be read as a universal constant. The instructive comparison is that it sits just above the site-percolation threshold for the square lattice, $p_c \approx 0.593$, and less than one bin width above it. That proximity is suggestive rather than conclusive at the present resolution, but it is consistent with the reading that the local threshold is a finite-window echo of the bulk percolation transition, displaced slightly upward because a small window sampled around one point needs somewhat more local fuel than the asymptotic bulk value to reliably sit inside a spanning cluster. Testing that reading properly would require varying the window size and checking whether ρ^*_{loc} converges toward p_c as the window grows — a clean experiment this design supports but does not perform.

The decomposition in Section 3.2 also refines a claim that is easy to state loosely. It is true that more regrowth means more fire, but the mechanism is lopsided: roughly 91% of the sevenfold increase in large-fire probability comes from fires being more severe once started, and only about 9% from lightning striking fuel more often. Framing regrowth as raising ignition frequency would misattribute the dominant effect almost entirely. Under a wildfire reading, the analogue is that fuel accumulation matters less because it gives lightning more targets than because it links the targets it already had into larger connected masses, though this analogy should be treated as interpretive rather than predictive given the model's simplifications.

The finite-size result carries a methodological implication beyond this paper. Because the effective regrowth per cell per lightning event scales as θ/N^2 , any comparison of this model class across system sizes at fixed θ will misstate the size dependence, and by a factor that can exceed 50 in the low- θ regime where it is most tempting to run larger lattices cheaply. The collapse onto ϕ in Section 3.4 is the correction. This is a property of the event-based regrowth formulation rather than a claim about the Drossel–Schwabl model in general; in the original p/f formulation the per-cell regrowth probability is specified directly and no such rescaling arises. The point is that when a model is reformulated for computational convenience, the control parameter may cease to be the invariant one, and the present data show what that costs if it goes unnoticed.

More broadly, the two methodological tools used here are not specific to the forest-fire model and transfer directly to other lattice-based stochastic systems. The first is the use of a locally measured observable — here the fuel density in a fixed window around an event — as a predictor of that event's outcome, rather than a global order parameter; the same construction applies wherever events are triggered at a point on a lattice and one wants a locally observable early indicator, for example activation sites in neural-network or epidemic-spreading models on lattices, or nucleation in Ising-type dynamics. The second is the common-random-numbers erasure measurement of Section 2.7: driving two differently-initialized copies of a system with a shared random stream and recording when they converge is a general and inexpensive way to measure how long any stochastic lattice model retains memory of its initial condition, and to identify the control variable on which that timescale collapses. Both tools require only that the model be reproducible under a fixed seed, which most cellular-automaton and interacting-particle models are, and both would be applied in exactly the form used here.

The same variable governs a second and less obvious quantity: how long the model remembers where it started. That ϕ sets both the risk curve and the memory-erasure timescale is the more interesting half of Section 3.3, and it carries a practical warning for anyone running this model class. Discarding a fixed transient is standard practice and is necessary to reach steady state, but the transient that guarantees stationarity at high ϕ also guarantees that every trace of the initial condition is gone — in our sweep, within the first 8,100 of 40,000 discarded events whenever $\phi \geq 0.061$. Any experiment that varies the initial condition and measures post-transient is therefore measuring nothing at high ϕ , and will return a null result whose apparent cleanliness is an artefact of the design rather than evidence about the model. The two requirements are in direct tension, and the erasure timescale is what tells you which regime you

are in. We would suggest measuring it before designing any initial-condition study in this model class, and treating $\phi \approx 0.05$ as the boundary below which such a study is meaningful at all.

These findings should be positioned carefully against the existing literature. A substantial body of work examines whether the Drossel–Schwabl model is genuinely critical in the asymptotic limit, with Pruessner and Jensen (2002) and Grassberger (2002) reporting broken scaling that complicates the simple self-organized-criticality reading advanced in earlier work (Drossel & Schwabl, 1992; Clar et al., 1994). We take no position on that question, and nothing here should be read as evidence for or against asymptotic criticality: our system sizes are far too small and our observable is not a critical exponent. What we report is a finite-size, locally measurable regularity, and its interest lies in being robust and mechanistically interpretable at accessible scales, not in what it implies about the $N \rightarrow \infty$ limit.

4.1 Limitations

This model is intentionally simplified, and several of its simplifications bear directly on how the results should be read. Burning is instantaneous at ignition, implemented as removal of the connected cluster, so there is no fire front, no burn duration, and no possibility that a fire is extinguished partway; consequently the model cannot speak to spread rate or suppression, which are the central concerns of operational fire models (Rothermel, 1972; Finney, 1998; Sullivan, 2009). There is no wind, topography, or fuel moisture, so directional and terrain-driven spread is absent by construction. The threshold estimate is limited by the binning resolution of $0.0167 = 1/60$ (the origin of this value, and the robustness of the threshold to the bin count, are set out in Section 2.9) and is reported to that precision only, which is also why the comparison against the percolation threshold is suggestive rather than decisive. The knee strength is imprecisely determined at high θ , where its standard deviation exceeds half its mean. The finite-size collapse rests on two overlapping ϕ values across two system sizes, which supports the claim that ϕ is the correct comparison variable but falls short of establishing a scaling law; a ϕ -matched sweep across three or more system sizes is the natural next step. The layout null in Section 3.3 covers three θ values at a single system size and initial density, and the memory-erasure times at $\phi \leq 0.031$ are censored at 20,000 events, so the low- ϕ end of that curve is bounded rather than measured. Finally, the analysis addresses connectivity thresholds and event statistics in a stylized setting and is not intended to forecast specific real fires; the wildfire language used throughout is interpretive framing for a lattice model, not a claim of predictive validity.

ACKNOWLEDGEMENTS

This project began when the author approached Ninaad Adhvaryu to explore whether simple physics-inspired models could help explain why some wildfire ignitions remain small while others become large events. The research questions, model design, and analysis plan were developed through regular discussions under his mentorship, within the Pangea Society research mentorship programme. The author implemented the simulation and data pipeline, produced the analyses and figures, and iterated on them with feedback on modelling assumptions and interpretation. The work was carried out without

Sep 2026

Vol 11. No 5.

formal institutional or laboratory support; computational runs and analysis were performed on personal computing resources and a privately available workstation. The author also thanks those who provided informal feedback on clarity and presentation during drafting.

DECLARATIONS

Author contribution: A.M. conceived the study with mentor guidance, implemented the simulation and analysis code, produced all figures, and wrote the manuscript. Research mentorship was provided by N.A. as described above.

Conflicts of interest: None declared.

Data and code availability: The simulation code, the corrected independently seeded runs of Section 2.6, the memory-erasure measurements of Section 2.7, and the aggregated result files underlying every figure and table are deposited in a permanent public repository at <https://github.com/aryan-m-div/forest-fire-model>.

Use of artificial intelligence: No generative AI tools were used to write, draft, or generate any portion of this manuscript or its analysis.

Human or animal participants: None. This study involves computational simulation only.

REFERENCES

- Bak, P. (1996). *How nature works: The science of self-organized criticality*. Springer.
- Bak, P., Tang, C., & Wiesenfeld, K. (1987). Self-organized criticality: An explanation of $1/f$ noise. *Physical Review Letters*, 59(4), 381–384.
- Clar, S., Drossel, B., & Schwabl, F. (1994). Scaling laws and simulation results for the self-organized critical forest-fire model. *Physical Review E*, 50(2), 1009–1018.
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703.
- Drossel, B., & Schwabl, F. (1992). Self-organized critical forest-fire model. *Physical Review Letters*, 69(11), 1629–1632.
- Finney, M. A. (1998). FARSITE: Fire area simulator — Model development and evaluation (Research Paper RMRS-RP-4). USDA Forest Service.
- Grassberger, P. (2002). Critical behaviour of the Drossel–Schwabl forest fire model. *New Journal of Physics*, 4, 17. <https://doi.org/10.1088/1367-2630/4/1/317>
- Grimmett, G. (1999). *Percolation* (2nd ed.). Springer.

- Henley, C. L. (1993). Statistical mechanics of the "forest-fire" model. *Physical Review Letters*, 71(17), 2741–2744.
- Jensen, H. J. (1998). *Self-organized criticality: Emergent complex behavior in physical and biological systems*. Cambridge University Press.
- Karafyllidis, I., & Thanailakis, A. (1997). A model for predicting forest fire spreading using cellular automata. *Ecological Modelling*, 99(1), 87–97.
- Malamud, B. D., Morein, G., & Turcotte, D. L. (1998). Forest fires: An example of self-organized critical behavior. *Science*, 281(5384), 1840–1842.
- Millington, J. D. A., Perry, G. L. W., & Malamud, B. D. (2006). Models, data and mechanisms: Quantifying wildfire regimes. Geological Society, London, Special Publications, 261, 155–167.
- Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46(5), 323–351.
- Pruessner, G., & Jensen, H. J. (2002). Broken scaling in the forest-fire model. *Physical Review E*, 65(5), 056707.
- Rothermel, R. C. (1972). A mathematical model for predicting fire spread in wildland fuels (Research Paper INT-115). USDA Forest Service.
- Saberi, A. A. (2015). Recent advances in percolation theory and its applications. *Physics Reports*, 578, 1–32.
- Stauffer, D., & Aharony, A. (1994). *Introduction to percolation theory* (2nd ed.). Taylor & Francis.
- Sullivan, A. L. (2009). Wildland surface fire spread modelling, 1990–2007. 3: Simulation and mathematical analogue models. *International Journal of Wildland Fire*, 18(4), 387–403.