How the Fireline Colorado simulation works

Model parameters, research, calculations, and recorded gameplay

Version 6.1 research update | 4 November 2026 | Documentation for instructor review

The game has been updated. I incorporated the research into version 6 before revising this report, including fuel-specific crew production, a water-demand calculation based on published fire-suppression equations, and separate roof, vent, and nearby-fuel pathways for home ignition.

Fireline Colorado is a decision-making exercise built around a probabilistic fire model. The player builds line, directs helicopter water, follows spot fires, and protects homes. This report answers the question: “Could you share documentation? I'm curious what some of the key assumptions are about the parameters, like prob a line holds, drops, etc.”

The central answer is that line holding and water-drop success depend on exposure. A completed line can block surface spread while an ember crosses it. A drop can cool part of a fire without extinguishing every burning cell. Homes respond separately to nearby heat, roof exposure, ember entry through vents, and combustible material close to the building. These distinctions follow the suppression and structure-exposure literature. Gannon et al., 2020 Penney et al., 2019 Manzello et al., 2011



Recorded gameplay from the updated build. The home inspector identifies building features and the separate exposure pathways used by the simulation.

Play the updated game | Open the videos

1. Reading the model and its evidence

I organized the report around the decisions the game asks the player to make: anticipate spread, complete line before contact, place water where it can cool fuel, follow embers, and protect buildings. Each chapter connects a rule to its inputs, explains what its numbers mean, and identifies the research that supports it.

The percentages have different denominators. Ordinary spread is the chance that one eligible receiving cell ignites during a simulation step. Water response is the chance that one burning cell goes out after receiving a specified effective water dose. The home calculation is an interval ignition probability based on accumulated exposure. None is an unconditional probability that the entire incident will be controlled.

The game uses a 6 m landscape grid and time advances of up to 5 minutes internally. Houses also carry attributes below that grid scale. This is how a 1.524 m noncombustible zone can affect a house without pretending the terrain mesh resolves a five-foot strip. The 5-foot dimension comes from IBHS guidance; the subgrid representation is my implementation choice. IBHS, 2025

Figures from gameplay are embedded directly in the PDF. Video posters open public playback pages, and the three MP4 files are also embedded as attachments. Browser PDF viewers normally use the clickable playback link; attachment-capable readers can extract the embedded files. The videos are credited simply as recorded gameplay demonstrations.

The appendix summarizes the updated parameters. The companion files contain the exact build identity, recorded run results, and test outcomes so the report can be checked against the supplied executable game.

2. Where the numbers come from

I use three kinds of numerical input. Published quantities are copied with their units and conditions. Derived values come from a cited equation or a stated unit conversion. Selected simulation values translate a mechanism into this game's spatial and temporal resolution when the literature does not provide a matching coefficient.

Input	Evidence	Implementation
Crew production	NWCG sustained Type I indirect rates	627, 323, or 455 ft/h converted to m/h
Critical water flow	Penney et al., equations 1-2	0.0129 L/m²/s baseline; 0.70 efficiency; 2,640 kJ/kg water enthalpy
Bucket capacity	SEI BB2024 capacity table	910 L nominal reference payload
Noncombustible zone	IBHS technical standard	5 ft = 1.524 m at each hardened home
Vent openings	NIST tests; IBHS guidance	Distinct 6 mm, 3 mm, and ember-resistant profiles
Response probabilities	Selected simulation mapping	Smooth water dose response and separate home pathway hazards

NWCG rates describe a particular crew category and vegetation setting. Penney's water calculation is a theoretical critical-flow model, evaluated over 540 constructed forest-fire scenarios. NIST's vent experiments report ignition of particular receiving materials under controlled firebrand exposure. Keeping those conditions attached to the numbers avoids turning a study-specific result into an unrelated universal probability. NWCG, 2021 Penney et al., 2019 Manzello et al., 2011

The research establishes the mechanisms and several dimensional inputs. The remaining probability coefficients are explicitly listed as selected values because these publications do not estimate matching probabilities for this game's six-meter cells and fifteen-minute home intervals. This classification applies throughout the report; it is not repeated as a disclaimer beside every equation.

This release is an educational simulation. Its checks establish reproducible implementation and expected response directions; they do not make it a field prediction or operational decision system.

3. The research framework

The cellular-automaton literature provides the model's basic structure: divide a landscape into cells, store each cell's state, and change that state in response to neighboring fire and environmental conditions. Ghosh and colleagues describe probabilistic transitions influenced by fuel and environmental factors. Zhou and colleagues explore learned transition rules. Fireline Colorado uses explicit equations, which make its decisions inspectable during a classroom exercise. Ghosh et al., 2024 Zhou et al., 2025

Rothermel's surface-fire model and Andrews' explanation provide a physical reference for why fuel, moisture, wind, and slope matter. Their work also clarifies the difference between spread rate, heat release, fireline intensity, and flame length. The revised water calculation uses a dimensional flame-length-to-intensity relationship, while the landscape spread engine retains its simpler transition hazard. Rothermel, 1972 Andrews, 2018

For suppression, NWCG's crew tables inform the revised production rates. Penney and colleagues provide the critical water-flow equations. Cup-and-grid drop testing explains why payload must be distributed over a ground footprint. Gannon and colleagues explain how line effectiveness depends on whether a line was engaged and what spatial outcome is being measured. NWCG, 2021 Penney et al., 2019 Suter, 2000 Gannon et al., 2020

For homes, NIST and IBHS supply the strongest direct basis for representing vents, roof assemblies, debris, and nearby combustibles separately. The updated game carries these attributes on each home and records the pathway involved in a simulated loss. The report cites the experiments and the guidance at the point each feature is discussed. Manzello et al., 2011 IBHS, n.d. IBHS, 2025

These sources have now been used to revise the game. The report describes their role in this update rather than assigning a reconstructed development story to earlier coefficients.

4. Grid scale and the generated map

Version 6 stores a 288 by 156 grid. With a cell side of 6 m, each cell represents 36 m². The grid is 1,728 m wide and 936 m high. Its total area is a model footprint, not the area of the historical incident named on a scenario card.

288 × 156 = 44,928 cells

44,928 × 36 = 1,617,408 m²

1,617,408 / 4,046.8564224 = 399.6702208 acres

The grid dimensions and six-meter cell side are prototype spatial choices. The cell side ties directly to the one-cell line and concentrated water action. A full-resolution screenshot can show much finer rendering than the simulation grid; more display pixels do not create additional physical precision.

Fewer cells reduce memory use and repeated neighborhood work. That is a computational consequence, not evidence that a measured browser performance limit selected exactly 288 by 156. No benchmark-derived selection of those dimensions was established. The model footprint should be treated as the scale of this exercise.

Cell size affects every spatial operation. A one-cell line is nominally 6 m wide; diagonal center-to-center movement is 8.485 m. The core drop includes cell centers near a target; its area changes as the center shifts. Ember landings are rounded to the receiving grid cell. Increasing resolution would require revisiting these rules rather than merely drawing more pixels.

5. Generated fuel cover and terrain values

The current landscape is a procedural exercise map. initFireGrid mixes broad noise and cell-level noise: cover = 0.62 coarse + 0.38 fine. Fuel is assigned only when cover exceeds 0.07 in grass, 0.13 in mixed fuel or 0.19 in timber, and the cell is outside the generated road and river corridors. Increasing a threshold creates more gaps; increasing the coarse share makes patches more spatially coherent.

density = round[255 × clamp(base + 0.24 cover, 0.48, 1)]

The base is 0.72 for grass and 0.80 for mixed or timber. For cover = 0.5, grass density is round(255 × 0.84) = 214, while mixed density is round(255 × 0.92) = 235. Those values become 214/255 and 235/255 in probability equations. They are normalized fuel-continuity proxies, not kilograms per square meter.

Elevation starts on a grid spaced four simulation cells apart. Broad and finer noise components are multiplied by scenario slope with coefficients 1.8 and 0.4, then bilinearly interpolated. Cell noise adds a smaller 0.025 slope-scaled perturbation. Thus the local rise used in spread is generated from a constructed surface, not read from a surveyed terrain model.

The road corridor excludes fuel within three cells of its centerline and the river within two. These raster exclusions are distinct from the rendered width of the drawn road or stream. They interrupt ordinary fuel continuity, while embers can cross them and land on receptive fuel farther away.

The exact noise proportions, thresholds and elevation scales have no established field derivation. Their purpose is a varied exercise landscape. LANDFIRE fuels and USGS elevation are possible replacements for a future mapped case, but would require explicit georeferencing, resolution and preprocessing decisions. LANDFIRE, n.d. USGS, n.d.

6. The simulation clock and random transitions

The default player turn is 15 minutes, processed as three steps of at most five minutes. For each step, weather is evaluated at its midpoint. Construction progresses first, existing fire supplies candidate ignition probabilities, draws establish new ignitions, existing burning cells age and cool, hotspots decay and may ignite a neighbor, water dries, and ember opportunities are processed.

New cells do not immediately run the ordinary neighboring-spread loop in the same step. They join the active list for later processing. Consequently, a three-step turn can move through several successive neighbors, while a single calculation applied only to the starting cells cannot represent the same spatial progression.

For an unchanged exposure with no cap, three five-minute survival probabilities equal one 15-minute survival probability:

$$[\exp(-R \times 5 / 60)]^3 = \exp(-R \times 15 / 60)$$

That identity does not make the full simulations equivalent. The wind, fuel state, line completion, source intensity, wetness and opportunities for spread can change between steps. Probability caps and linear crossing rules introduce further differences.

The engine's random generator advances an unsigned 32-bit state with multiplier 1,664,525 and increment 1,013,904,223, then divides by 4,294,967,296. These are computational constants for a linear congruential generator, not fire parameters. A fixed starting seed and unchanged sequence of calls support repeatable runs.

$$\text{seed}_{\text{next}} = (1{,}664{,}525 \times \text{seed} + 1{,}013{,}904{,}223) \bmod 2^{32}$$

Changing actions changes the sequence of draws. Reusing a seed is useful for testing but does not guarantee that two different tactics receive identical random events in every cell. Weather has a related seeded stream, and display animation is not a measurement of physical flame motion. The three retained demonstrations identify seed 6142026, control mode and weather mode so their outcomes can be interpreted correctly.

7. Weather inputs and estimated fuel moisture

The scenario supplies initial wind and humidity arrays, temperature, fuel class, slope and wind direction. The weather builder extends those into generated hourly conditions, adds daily variation and scheduled shifts, and interpolates between them. It does not ingest observed hourly weather. The word original in an original challenge refers to its preset rules.

The moisture proxy first calculates an equilibrium target using humidity, temperature and fuel class. The fuel offset is 0 for grass, 1.5 for mixed fuel and 4 for timber. Initial moisture is clamped to 1.2-24; later equilibrium targets allow 1.2-30.

M_eq = clamp[1.8 + 0.20 RH + 0.09(85 - T) + fuelOffset, 1.2, 30]

M_next = lag × M_previous + (1 - lag) × M_eq

Here RH is percentage points and T is °F. The lags are 0.55, 0.72 and 0.84 for grass, mixed and timber. Larger lag means a slower response. Starting from M = 18 with a dry target RH = 10 and T = 90 gives a grass target of 3.35 and a next value of 11.4075. Mixed fuel gives a target of 4.85 and next value 14.318; timber gives 7.35 and 16.296.

With RH = 60 and T = 70, the grass target is 15.15. Starting at 5 gives 0.55 × 5 + 0.45 × 15.15 = 9.5675 after an hourly update. These are calculated proxy values, not sampled fuel moisture.

Wotton's grass-moisture work explains why fuel exposure, solar effects and shelter matter, and why a litter-moisture model is not automatically a grass model. It supports differentiated moisture dynamics, not my exact offsets or lags. Wotton, 2009 The current proxy omits precipitation, explicit solar loading, fuel size classes and a measured moisture observation model.

8. Weather restrictions and generated practice conditions

The wind is bounded between 0 and 90 mph and humidity between 5% and 95%. Gusts are constructed above sustained wind and capped at 115 mph. Turbulence begins at 0.12 and rises around programmed events. Visibility is another proxy: 12 minus five times turbulence, with an additional 2 km night penalty, bounded to 2-15 km before smoke is considered.

$$\text{gust} = \text{clamp}[w + 2 + w(0.12 + 0.18U) + 8\,\text{turbulence}, w, 115]$$

U is a uniform random draw. The 2 mph offset creates a gust increment even at low sustained wind; the random term spreads that increment; the turbulence term couples event severity to gusts. These constants define variation in the exercise and have no recovered station-data fit.

Guided practice caps sustained wind at 6 mph, gusts at 12 mph and temperature at 78°F; it raises RH to at least 45% and moisture to at least 18. Turbulence is capped at 0.20 and visibility raised to at least 8 km. The development record explains the purpose: provide a workable learning setting while retaining the harder original presets. It does not provide a physical derivation for each limit.

Scheduled weather events occur at hours 7, 18, 31 and 54 when the selected duration includes them. The tutorial instead has early events at hours 1, 2 and 3. These clocks are lesson design. They should not be presented as the observed timing of a historical front.

The Marshall event summary from the National Weather Service supplies real incident context and high-wind observations. It does not turn the game's hourly arrays, generated terrain, or one-cell starting ignition into a reconstruction. National Weather Service, n.d. I keep those categories separate because an impressive-looking named scenario can otherwise imply more historical accuracy than its inputs support.

9. The full generated weather cycle

The generator uses sun = cos[2π(localHour - 15)/24]. After the supplied hourly wind array ends, day strength is 0.75 + 0.18 sin(1.7 day + 0.001 seed). Wind combines scenario average, day strength, a night multiplier of 0.46 or a daylight term 0.75 + 0.25 max(0,sun), and centered random variation. These constants produce a repeating exercise pattern; they are not a fitted meteorological forecast.

Temperature is base T + (sun - 1) × 10 + (U - 0.5) × 2.5, with 4 replacing 10 in the beginner tutorial. Before later restrictions, the sinusoidal term ranges from zero to a 20°F reduction in regular scenarios, while the random offset ranges from -1.25 to +1.25°F. Direction starts at the preset compass angle with ±9° random jitter.

Regular events at hours 7, 18, 31 and 54 shift direction by +70°, +40°, -45° and +40°. Their wind pulses rise and fall within a two-hour distance from the event. Typical strength is 6 mph; the hour-31 pulse uses min(14, 0.30 × peak preset wind). Event turbulence adds up to 0.45. The hour-18 front can add 18 humidity points and lower temperature 8°F at peak influence. These are generated event rules.

The Marshall preset applies a later wind multiplier of 0.68 from hour 7, with humidity recovery up to 30 percentage points at 1.5 points per later hour and cooling up to 10°F at 0.4°F per hour. The observed event's evolving weather motivates checking a time-varying scenario, but this specific schedule is not an observed reconstruction. National Weather Service, n.d.

The tutorial has additional wind 2-9 mph, RH 28-48% and temperature 66-78°F bounds before guided restrictions. Weather interpolation follows the shortest angular difference across north. A shift from 350° to 10° therefore interpolates through north, not around the long way. These details affect the conditions supplied to each five-minute step.

10. From weather to flame length and spread potential

Before individual cells interact, fireConditions produces scenario-scale flame length FL, a spread-potential index ROS, and a spotting index. The fuel constants for FL are 0.8 grass, 1.0 mixed and 1.3 timber. For ROS they are 0.5, 0.4 and 0.3 respectively.

D = clamp(1.25 - M / 55, 0.65, 1.22)

F_T = clamp[1 + (T - 70) × 0.0025, 0.8, 1.15]

FL = f_FL × (0.15w + 3) × (1 + slope / 150) × D × F_T

ROS = f_ROS × (1 + w / 8) × (1 + slope / 120) × D × F_T

FL is displayed in feet. ROS is an empirical model index: it is not a rate in mph, m/min, or chains/hour. M is the moisture proxy, slope is percent, w is mph and T is °F. The constants therefore include implicit units; these are not interchangeable dimensionless physical laws.

For mixed fuel, w = 8 mph, slope = 0, M = 9 and T = 70°F, D = 1.08636 and F_T = 1. FL becomes 4.56273 ft and ROS becomes 0.86909. At w = 24 with the other inputs unchanged, FL becomes 7.17 ft and ROS 1.73818. Wind doubles the index here because 1 + w/8 changes from 2 to 4.

Rothermel's physical rate requires fuel loading, surface-area-to-volume ratio, packing, heat content, moisture and other inputs that this formula does not contain. Its output cannot be assigned to my ROS merely because both use the word spread. Rothermel, 1972 Andrews also distinguishes fireline intensity and flame-length relationships; my normalized cell intensity is not Byram intensity in kW/m. Andrews, 2018

These precursor factors are applied again through local moisture, wind and slope terms. That creates deliberately strong sensitivity, but without calibration it may also double-count environmental effects. The register flags each coefficient separately rather than treating the entire precursor formula as research-derived.

11. What a physical spread equation would require

Rothermel's reference equation relates energy moving into unburned fuel to the energy needed to ignite that fuel. Written in its customary form, it makes the missing physical measurements in my prototype easier to identify. The equation below is a research comparison; version 6 does not evaluate it. Rothermel, 1972 Andrews, 2018

$$R_{physical} = I_R \times \xi \times (1 + \varphi_w + \varphi_s) / (\rho_b \times \varepsilon \times Q_{ig})$$

I_R is reaction intensity, ξ the propagating-flux ratio, φ_w and φ_s the wind and slope factors, ρ_b fuel-bed bulk density, ε the effective heating number, and Q_{ig} heat of preignition. Using I_R in Btu/ft²/min, ρ_b in lb/ft³ and Q_{ig} in Btu/lb produces a spread rate in ft/min. The ratios and factors are dimensionless. The fuel descriptors and moisture determine intermediate quantities rather than entering through one interchangeable fuel label. Rothermel, 1972 Andrews, 2018

For an arithmetic illustration only, choose I_R = 1,000 Btu/ft²/min, ξ = 0.05, φ_w = 1, φ_s = 0, ρ_b = 1 lb/ft³, ε = 0.5 and Q_{ig} = 100 Btu/lb. These are selected example inputs, not measurements from a Colorado incident or a calibrated fuel model. The equation, rather than these example values, comes from the reference. Rothermel, 1972

$$R_{physical} = 1{,}000 \times 0.05 \times 2 / (1 \times 0.5 \times 100) = 2 \text{ ft/min}$$

Andrews also gives the relationship between heat per unit area H_A, spread rate and Byram fireline intensity I_B. If H_A is Btu/ft² and $R_{physical}$ is ft/min, division by 60 converts minutes to seconds. The flame-length expression uses those specified customary units. Andrews, 2018

$$I_B = H_A \times R_{physical} / 60$$

$$\text{flame length in ft} = 0.45 \times I_B^{0.46}$$

With a separately chosen H_A = 1,000 Btu/ft² and the illustrative 2 ft/min rate, I_B is 33.3333 Btu/ft/s and flame length is 2.2581 ft. This example shows the chain from dimensional inputs to physical outputs. It does not assign a physical heat release to the game's normalized intensity of 1. Andrews, 2018

The game's FL and ROS indices instead come from its simpler weather formulas. Its local transition hazard then adds another set of multipliers. Converting a researched physical rate into a six-meter cell arrival process would require a stated travel-time rule, directional geometry and independent progression data. Borrowing the name ROS or the shape of a wind multiplier is not sufficient to perform that conversion.

12. Why the spread equation is multiplicative

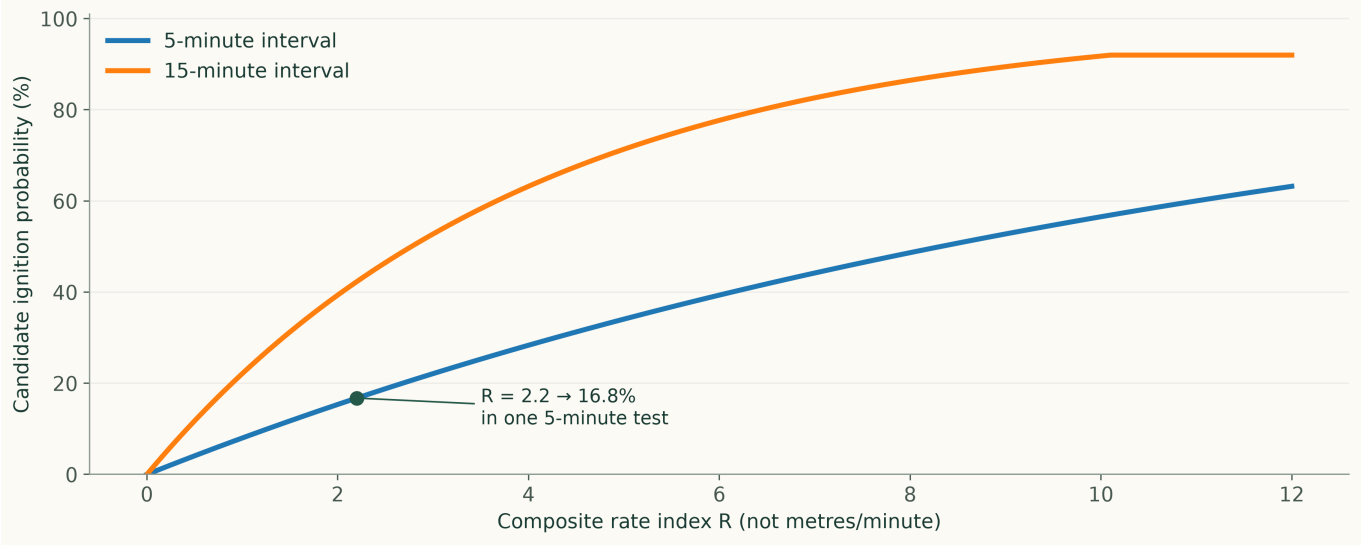
The local engine combines a baseline rate scale with environmental and state modifiers. Multiplication makes a zero fuel factor shut off ordinary spread, makes a wetting factor attenuate exposure continuously, and makes wind affect an otherwise comparable receiver according to direction.

R = 2.2 × clamp(ROS, 0.2, 8) × S × M_f × W_f × F × T_f × I × V × N_f

P = clamp[1 - exp(-R × dt / 60), 0, 0.92]

S is the slope factor, M_f moisture, W_f wetting, F receiver fuel density divided by 255, T_f fuel type, I source intensity, V directional wind and N_f burning-neighbor influence. dt is simulated minutes. R must be interpreted as an effective hourly transition hazard for dimensional consistency, even though it is built from model indices rather than a measured physical velocity.

The factor 2.2 sets the overall time scale. With all other factors one, ROS = 1 and dt = 5, the exponent magnitude is 2.2 × 5/60 = 0.183333. The ignition probability is 1 - exp(-0.183333) = 0.167509, or 16.75%. A rate scale of 1.1 would give 8.76%; 4.4 would give 30.70% for the same cell comparison.



Analytical probability conversion. Curves describe fixed-rate examples; the running engine updates state between steps.

The recovered notes already contain 2.2 and identify hand-set assumptions. They do not document how 2.2 was first selected. There is no equivalent published constant that can simply be substituted: a coefficient belongs to its equation, units, time step and spatial definition. A physical spread speed would need a documented conversion to crossing a 6 m cell and validation against observed progression.

13. Local slope moisture and water factors

The slope term uses the elevation difference between source and receiver, so uphill and downhill directions can differ within a scenario. The scenario slope then scales that directional difference.

S = clamp[1 + tanh(0.18 × rise) × slope / 100, 0.5, 1.7]

M_f = clamp[1.3 - (M + 3L) / 30, 0.18, 1.18]

W_f = exp(-1.1L)

Rise is generated elevation difference in model meters, slope is percent, and L is effective water load in L/m². The 0.18 controls how quickly tanh saturates with rise. The lower and upper bounds stop this factor reversing sign or increasing indefinitely. At a 3 m rise and 40% scenario slope, tanh(0.54) is about 0.493; S is about 1.197. A negative rise reverses the directional adjustment before clamping.

At M = 9 and L = 0, M_f = 1. At L = 1 it becomes 0.9 and W_f = 0.332871. Their product is 0.299584. Holding ROS and the other factors fixed, the baseline five-minute probability falls from 16.75% to 5.34%. At L = 2, the product is 0.8 × exp(-2.2) = 0.088643 and P is 1.61%.

At M = 18 with no water, M_f = 0.7 and the same fixed-input probability is 12.04%. The lower clamp means even very high moisture does not alone force this factor to zero. Ordinary spread can still become zero through nonfuel cells, whereas water approaches zero probability asymptotically.

Moisture impeding ignition and upslope heating increasing spread have physical precedents. Rothermel, 1972 Wotton, 2009 The exact tanh form, 3L moisture adjustment, exponential coefficient 1.1 and clamp values are not established measurements. They create a controllable response with safeguards. To calibrate them, I would need receiver ignition observations with matched fuels, slope, moisture, applied water and elapsed time.

14. Wind fuel and neighboring cells

The wind vector uses meteorological direction: the code converts a direction in degrees into [-sin(direction), cos(direction)] on the screen grid. Alignment is the dot product of that downwind vector with the unit direction from source to receiver. It is 1 directly downwind, 0 across wind and -1 upwind.

V = clamp[1 + (w / 38) × alignment × 0.8 + (gust / 100) × max(0, alignment), 0.15, 3.2]

F = receiverFuelDensity / 255

T_f = 1.25 grass, 1.00 mixed, 0.80 timber

I = max(0.04, sourceIntensity)

N_f = 0.7 + 0.15 × min(5, burningNeighbors)

With 30 mph sustained wind and 40 mph gusts, the downwind factor is 2.03158. Across wind it is 1; directly upwind it is 0.36842. Keeping the other factors at one gives five-minute probabilities of about 31.10%, 16.75% and 6.53%. This is a directional sensitivity test, not a forecast from a measured 30 mph fire.

The density value is a byte-scale fuel-continuity proxy. Dividing by 255 maps its full range into 0-1; it does not convert kilograms of fuel into a density. The fuel-type factor changes the baseline local probability to 20.48% for grass and 13.64% for timber when ROS is held at one. Actual scenario comparisons also change precursor ROS, flame length, residence time, and moisture.

The neighbor factor is 0.85 for one burning neighbor, 1.00 for two, and 1.45 for five or more. The engine examines each source but retains the largest candidate probability for the receiver, followed by one random draw. It does not calculate 1 minus the product of every neighbor's survival probability. This avoids independent duplicate trials but discards some combined exposure information.

The 0.04 intensity floor keeps weak burning sources contributing. These numerical choices control the simplified landscape response. Wind alignment and local neighborhoods have support in the supplied cellular-automaton papers, but those publications do not establish these constants. Ghosh et al., 2024 Zhou et al., 2025

15. Probability caps and what the tests establish

The exponential conversion is useful because it maps a nonnegative rate to a probability below one. For small R dt/60 it is approximately linear; at larger exposure it saturates. The 0.92 cap adds a further restriction: a candidate receiver retains at least an 8% chance of avoiding that ordinary transition in a step, even if the uncapped calculation is nearly certain.

The cap becomes active at $R = -\ln(0.08) \times 60/5 = 30.3087$ for a five-minute step. This number is calculated from the cap; it is not another independent parameter. For three identical capped opportunities, the probability of at least one ignition is $1 - 0.08^3 = 99.9488\%$, provided the receiver remains exposed and eligible. The cap therefore does not mean a cell can never have more than 92% cumulative ignition probability.

ROS with other factors one	Five-minute probability	Uncapped 15-minute probability
0.2	3.6003%	10.4166%
0.5	8.7591%	24.0428%
1	16.7509%	42.3050%
2	30.6959%	66.7129%
4	51.9695%	88.9197%
8	76.9307%	98.7723%

The last column is a mathematical reference; the running game uses three updated steps and would cap a single 15-minute ordinary candidate at 92% if called that way. It is not the actual result of a 15-minute turn.

The controlled calculations compare exported development functions with an independent implementation of the equations. They verify arithmetic and direction of response. They do not establish that 2.2 produces the intended physical spread rate. That would require georeferenced arrival times and an agreed error metric across multiple fires.

Probability caps also complicate calibration: different combinations of large coefficients can all produce the same capped output. I would fit and inspect uncapped hazards as well as probabilities, and test whether apparent agreement occurs only because the model saturates.

16. Building line over time

A line starts as a plan. The game converts the drawn path into connected sections, credits crew work as simulation time advances, and activates a section only when enough work has been completed. A horizontal or vertical section costs 6 m; a diagonal section costs $6\sqrt{2}$ m. The first section counts as 6 m.

The updated hand-crew rate depends on fuel. I adopted the sustained Type I indirect values from the NWCG table because the core action normally constructs line ahead of the active front. Short grass uses 627 ft/h, brush uses 323 ft/h, and the grouped timber-litter categories use 455 ft/h. Mixed terrain maps to brush as a stated category approximation. NWCG, 2021

$R_{crew} = R_{NWCG} \times 0.3048$ meters per foot

$\text{capacity} = (R_{crew} / 60) \times dt \times \text{clamp}(1 - \text{slope} / 130, 0.25, 1)$

$q = \text{clamp}(0.90 - 0.0008 \times \text{slope}, 0.60, 0.99)$

At flat terrain, the resulting rates are 191.1096 m/h in grass, 98.4504 m/h in mixed fuel, and 138.6840 m/h in timber. Over 15 minutes the capacities are 47.7774 m, 24.6126 m, and 34.6710 m. Only whole sections become completed line; remaining credit is carried forward.

Core projects share a single crew capacity. Drawing additional projects changes the work queue rather than creating extra crews. The slope multiplier reduces production, and nearby active fire may pause work. Those adjustments express the scenario's work constraints; the NWCG numbers supply the starting rate.



The updated line demonstration shows partial completion after time advances. The counter represents completed sections, not the whole drawn plan.

PLAY VIDEO 01 - Fuel-specific crew production and line construction

17. Research behind the crew rates

NWCG's 2021 memorandum preserves production tables assembled from earlier work and distinguishes sustained work from initial action. I use the sustained 20-person Type I indirect category. The table's short-grass and open-timber/grass grouping lists 627 ft/h with an expected range of 508-746 ft/h. Brush lists 323 ft/h with 244-403 ft/h. The timber-litter grouping lists 455 ft/h with 396-515 ft/h. NWCG, 2021

Game fuel mapping	Adopted rate	Published expected range, converted
Grass: short grass/open timber-grass	191.1096 m/h	154.8384-227.3808 m/h
Mixed: brush proxy	98.4504 m/h	74.3712-122.8344 m/h
Timber: grouped litter fuels	138.6840 m/h	120.7008-156.9720 m/h

These conversions are exact; the mapping of a broad game fuel class to a table row is an adaptation. The reported ranges describe production variability in the source, not statistical confidence intervals for this game. A simulated slope penalty is applied after selecting the nominal rate. NWCG, 2021

For an 18% slope in mixed fuel, the terrain multiplier is 1 - 18/130 = 0.861538. Fifteen minutes supplies 98.4504 × 0.25 × 0.861538 = 21.2047015 m of work. Three 6 m sections complete 18 m and retain 3.2047015 m credit. A longer path may require several turns, depending on safety pauses and other queued work.

The 6 m map treatment corridor is an abstraction. The NWCG table does not establish that a crew clears a continuous six-meter mineral-soil strip at these rates. The code uses the corridor to represent an effective break in fuel continuity. A finer engineering model would distinguish mineral strip, cleared vegetation width, and holding resources.

The old uniform 180 m/h value has been replaced in both the browser engine and its background simulation. Tests confirm that the two execution paths give the same result for an identical seed and sequence of actions.

18. Line quality width and construction restrictions

Quality changes two things. It enters the line-stress equation, and it controls how much original fuel remains when construction finishes. In grass, quality above 0.85 clears fuel completely; quality above 0.95 clears any fuel class completely. Otherwise the retained fraction is clamp[(1 - q) × 0.35, 0, 0.15].

At q = 0.8856 in mixed fuel, the retained fraction is 0.04004. A cell originally at 255 becomes round(255 × 0.04004) = 10. That is a 10/255 fuel proxy, not a field observation that 3.92% of combustible mass remains. At q = 0.90 in grass, the special branch leaves zero ordinary spread fuel.

The six-meter core width is a one-cell representation. A real handline's scraped width and the surrounding treatment or fuelbreak width are distinct quantities. Neither the production table nor the four-foot flame guidance establishes that a 20-person crew continuously clears a 6 m mineral-soil strip at the game's speed.

The hand method disallows scenario slope above 75%, local flame length above 4 ft, or gusts above 45 mph. The flame estimate is c.FL × √I for nearby active cells. The core interface adds a fixed radius of three cells, or 18 m, around active flames. The dynamic search radius is bounded to 2-8 cells and expands with flame length and gusts.

NWCG's surface-fire interpretation describes the difficulty of direct hand-tool attack as flame length increases beyond roughly 4 ft. It supports a tactical distinction, not this exact universal stop rule, the 18 m buffer or the 75% slope threshold. NWCG, 2021, PMS 437-1 These are exercise gates and do not compute escape routes or a safety zone.

Expanded controls retain a dozer method at 12 m, 720 m/h and quality 0.96, with a 40% slope gate and 8 ft flame gate; wide control line uses 18 m, 450 m/h and 0.98, with 30% and 8 ft gates. Their terrain divisor is 65 and quality penalty 0.001 per slope percentage point. They are separate prototype categories, not the default action.

19. What it means for a line to hold

An unengaged line has not been reached by fire. An engaged line may hold an active edge, be burned over, or be bypassed by a different pathway. A treated cell can ignite through residual fuel; heat can reach a receiver beyond the line; or an ember can land beyond it. An entire incident is controlled only if all relevant pathways stop and the model's cooling and observation requirements are met.

Gannon and colleagues studied spatial fireline and final-perimeter records for 33 western United States fires from 2017-2018, selected for data completeness rather than random sampling. They cleaned line records and classified segments relative to the final fire area and perimeter buffer. Held, burned-over and unengaged lengths supplied engagement and holding ratios. Their analysis found substantial incident variation; handline's held-to-engaged ratio was 0.76, compared with 0.53 for dozer line. The authors caution that placement and fire behavior, rather than intrinsic line quality alone, may explain that contrast. Methods sections 2.1-2.3 and results section 3.3 provide the relevant scope. Gannon et al., 2020

Those ratios do not give my next-cell ignition probability. They combine many exposures, choices of location, weather changes and lengths of line. Unengaged line must be excluded from an engagement-conditioned holding denominator, but may still matter when evaluating resource use.

engagement ratio = (held length + burned-over length) / total line length

held-to-engaged ratio = held length / (held length + burned-over length)

Gannon and colleagues classified held line with a 50 m buffer around the final perimeter and tested 25, 50, 75 and 100 m alternatives. That spatial tolerance is part of the measurement definition. It cannot be treated as an exact observation that a six-meter model segment stopped a front. Their reported ratios also reflect where crews chose to build and the conditions those lines encountered. Gannon et al., 2020

The Forest Service dashboard extends spatial assessment across hundreds of large wildfires. Its incident and line views offer a potential evaluation dataset, with the same need to understand mapping quality and final-outcome definitions. USDA Forest Service, n.d.

Version 6's counters are different: lineEncounters counts repeated source-neighbor checks, and direct breaches count successful classified receiver ignitions. They are not independent trials on equal line segments. Dividing breaches by encounters and subtracting from one would not reproduce the geospatial held-to-engaged ratio, nor justify calling the result a universal chance that a line holds.

20. Line stress and interval probability

The line-exposure function compares a heat-reach proxy with completed line width. Longer flames, wind aligned toward the receiver, and steeper scenario slopes increase reach. Better line quality, higher humidity, and water at the receiving cell reduce stress. These are simplified representations of exposure and fuel continuity; the geospatial line-effectiveness studies motivate the variables and outcome distinctions. Gannon et al., 2020

reach = FL × 0.3048 × √I × (1.1 + H × gust / 24) × (1 + slope / 180)

stress = (reach / width) × (1.15 - 0.55q) × F_t × (1 - RH / 180) × [2 / (L + 2)]

r15 = clamp(0.20 × stress + B, 0.002, 0.65)

P_cross = 1 - (1 - r15)^[(dt / 15) × exp(-L) × F]

H is nonnegative downwind alignment, q is completed quality, F_t is 1.35 for timber and 1 otherwise, L is receiving-cell water load, and F is receiving fuel density divided by 255. B is 0.055 when the receiver is more than 0.8 m below the source and scenario slope exceeds 25%; it approximates rolling material.

The update interprets r15 as a reference fifteen-minute crossing probability. The logarithmic survival conversion then calculates the probability for the actual step length. This replaces linear multiplication by dt/15, which did not preserve the stated reference probability over repeated steps.

The coefficients controlling reach, quality, and rolling material are selected response values. Research does not provide a matching local probability for this particular grid. The equation preserves the relevant directions of response and keeps a clear denominator: an eligible receiving cell beyond an engaged line.

The engine searches two to seven cells beyond an adjacent completed line, depending on flame and wind conditions. A candidate crossing competes with other ignition pathways to that receiving cell. Ember transport is evaluated separately and can bypass this local search altogether. Albini, 1979

21. A complete line crossing calculation

Consider a dry receiving cell with full fuel density, a 6 m line, 4 ft flames, source intensity 1, 24 mph gusts aligned with the receiver, 18% slope, 30% relative humidity, and quality 0.8856. There is no downhill rolling increment.

reach = 4 × 0.3048 × (1.1 + 24/24) × (1 + 18/180) = 2.816352 m

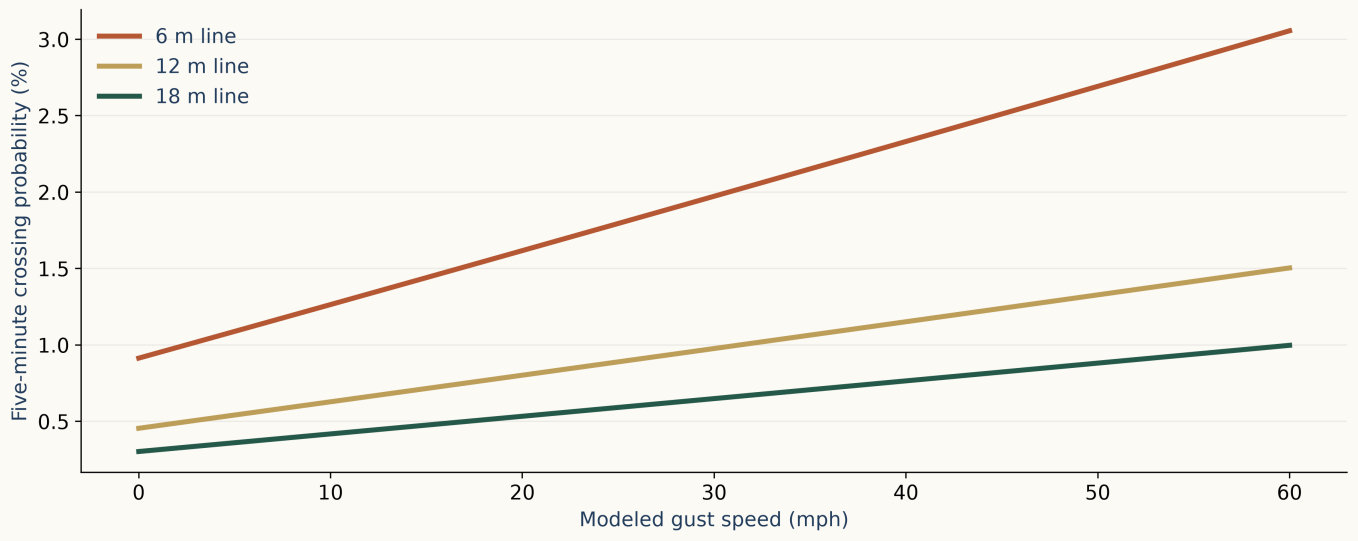
stress = (2.816352 / 6) × (1.15 - 0.55 × 0.8856) × (1 - 30/180) = 0.2593077872

r15 = 0.20 × 0.2593077872 = 0.05186155744

P_cross(5 min) = 1 - (1 - 0.05186155744)^(1/3) = 0.01759495 approximately

The five-minute candidate is about 1.76%. For three identical eligible exposures, the cumulative probability is exactly the reference 5.1862% because survival multiplies: $(1 - P5)^3 = 1 - r15$. Actual game turns update fire, fuel, weather, and water between steps, so the three probabilities usually differ.

At the maximum reference r15 = 0.65, the five-minute probability is 1 - 0.35^(1/3) = 29.5270%. Three unchanged steps return 65%, rather than the approximately 51.93% produced by the old linear time scaling. The update makes the time interpretation internally consistent.



Updated reference-to-step conversion. These are receiving-cell probabilities under fixed conditions, not measured whole-line success rates.

Gannon's held-to-engaged ratios concern spatial lengths of line after an incident. Their denominators cannot be substituted for this single receiving-cell calculation. A long line encounters many sources and can be crossed by multiple pathways. Gannon et al., 2020

22. Helicopter payload and availability

The updated utility helicopter carries a nominal 910 L bucket, matching the listed metric capacity of SEI's BB2024 reference bucket. The manufacturer also lists 240 US gallons; the two printed capacities are rounded nominal values. This provides a concrete payload reference without claiming the generic aircraft is certified for that load in every setting. SEI Industries, 2021 NWCG, 2019

Setting	Utility bucket	Heavy tank
Nominal payload	910 L	2,400 L
Geometry speed	45 mph	65 mph
Sustained-wind exercise gate	28 mph	36 mph
Gust exercise gate	38 mph	47 mph
Minimum modeled visibility	1.6 km	1.2 km
Fall time for drift	2.4 s	2.0 s
Timber canopy interception	0.28	0.20

The heavy tank and the flight gates remain generic training configurations. Aircraft loading in practice depends on the specific aircraft, density altitude, fuel, crew, equipment, and mission. The table's operating gates are scenario rules rather than operating authorizations. NWCG helicopter guidance explains the load and performance considerations. NWCG, 2019

Smoke reduces modeled visibility by $0.035 \times \sqrt{\text{activeCells}} \times \text{meanIntensity}$, subject to a 0.2 km floor. The availability check can refuse a drop before any payload is spent. In severe wind, that restriction makes the player confront the limited availability of aerial suppression.

The core exercise replenishes five drop opportunities each fifteen-minute turn after mobilization. This is a pacing budget, not a calculated helicopter turnaround time. Payload is physically accounted for inside each accepted drop; the dispatch schedule remains a classroom simplification.

23. Drift corrected aiming and footprint shape

The wind-drift calculation uses the modeled time between release and landing. One mph is 1,609.344 m / 3,600 s = 0.44704 m/s. Multiplying wind speed by seconds gives meters; dividing by the 6 m cell side gives cells.

drift_cells = w_mph × 0.44704 × fallSeconds / 6

At 20 mph the utility drift is 21.45792 m, or 3.57632 cells. The heavy drift is 17.8816 m, or 2.98027 cells. At 30 mph those distances are 32.18688 and 26.8224 m. These are transport calculations even if the dispatch gate would prevent a utility drop under those sustained winds.

Manual aim releases at the clicked location and lets the center drift downwind. Corrected aim subtracts that drift from the release point, so the mean landing center returns to the clicked location. The random crosswind term remains:

cross_cells = (U1 + U2 - 1) × [1 + 0.12(gust - w) + 3 turbulence] / 6

The sum of two uniforms produces a triangular distribution concentrated near zero. It is not a measured turbulence distribution. The exact coefficients control scatter in the exercise.

Core mode fixes both footprint semi-axes to one cell. Other patterns use speed, wind and class to determine an elongated ellipse: the trail length proxy uses speed × 0.44704 × 0.85, while compact starts at 13 m; heavy applies 1.18 and wind adds 0.12w. Dividing by 12 yields the major semi-axis with a minimum of 1.1 cells. Width starts at 12 m utility or 18 m heavy, adds 0.13(gust - w), and is divided by 12 with a minimum of 0.8.

Those additional patterns are retained experiments. In the default core footprint, changing geometry speed does not change the one-cell axes. That distinction prevents an aircraft table from implying that every displayed specification affects every mode.

24. Allocating the payload across cells

For each nearby cell center, the drop routine rotates displacement into the ellipse's along and cross directions and scales those coordinates by the semi-axes. The variable called r is the sum of their squares. Its cutoff is $r \leq 1.2$, so the accepted distance in scaled coordinates is $\sqrt{1.2}$, about 1.0954, rather than 1.2.

$$r = [(dx \cos\theta + dy \sin\theta)/a]^2 + [(-dx \sin\theta + dy \cos\theta)/b]^2$$

$$weight_i = \exp(-1.9 r_i) \times (0.82 + 0.36 U_i)$$

$$allocatedLitres_i = payload \times weight_i / sum(all\ weights)$$

The exponential favors the center. The random factor ranges from 0.82 up to 1.18 and changes relative allocation. Normalization then forces allocations over included on-map and off-map centers to sum to the payload. The 1.9 shape coefficient and scatter range have no established measurement-derived origin.

Off-map weights remain in the denominator. Their liters are recorded as escaped rather than being redistributed onto the map. This is important near boundaries: excluding them before normalization would artificially concentrate the whole payload within the visible grid.

With core axes a = b = 1, a nearly centered drop often includes the center and four orthogonal neighbors. The nominal included cell area is then $5 \times 36 = 180$ m². That is the sum of whole included grid cells, not a measured continuous wet polygon. Fractional shifts of the center can alter membership and make the footprint discontinuous.

Actual drop-testing methods measure spatial coverage rather than assuming equal volume everywhere. The Forest Service's cup-and-grid account describes weighed collection cups and interpolated coverage contours, including measurement errors from evaporation and splash. Suter, 2000 This supports measuring and normalizing distribution, but does not validate the game's circular core, the cutoff or the weight curve. A future implementation could use measured drop rasters resampled to the game's grid while conserving liters.

25. Delivered water and effective water load

The routine records several different quantities. Payload is water assigned to the aircraft class. Delivered liters are allocated to included cells on the map, before canopy and atmospheric losses. Intercepted liters are the modeled share caught by canopy. Effective load is the remaining water per cell area after both reductions. The interface's delivered-liter total should therefore not be read as water measured at combustible material.

$E = \text{clamp}[0.003(T - 55) + 0.001(100 - RH) + 0.004w, 0.03, 0.45]$

$L_i = \text{allocatedLitres}_i / 36 \times (1 - \text{canopyLoss}) \times (1 - E)$

T is °F, RH is percentage points, w is mph and L is L/m². One L/m² has the geometric depth of one millimeter if uniformly spread. Here it is an effective load used by several response rules, not measured soil moisture or a measured change in fuel moisture.

For 90°F, RH = 10% and w = 30 mph, $E = 35 \times 0.003 + 90 \times 0.001 + 30 \times 0.004 = 0.315$. No clamp is active. In timber, the utility class retains $0.72 \times 0.685 = 0.4932$ of allocated water after both reductions. If a cell receives 180 L, its effective load is $180/36 \times 0.4932 = 2.466$ L/m².

For grass, canopy loss is zero. For mixed fuel, it is 45% of the aircraft's timber value: 0.126 utility and 0.09 heavy. With the same atmospheric conditions and 180 L allocation, mixed-fuel utility load is $5 \times 0.874 \times 0.685 = 2.99345$ L/m².

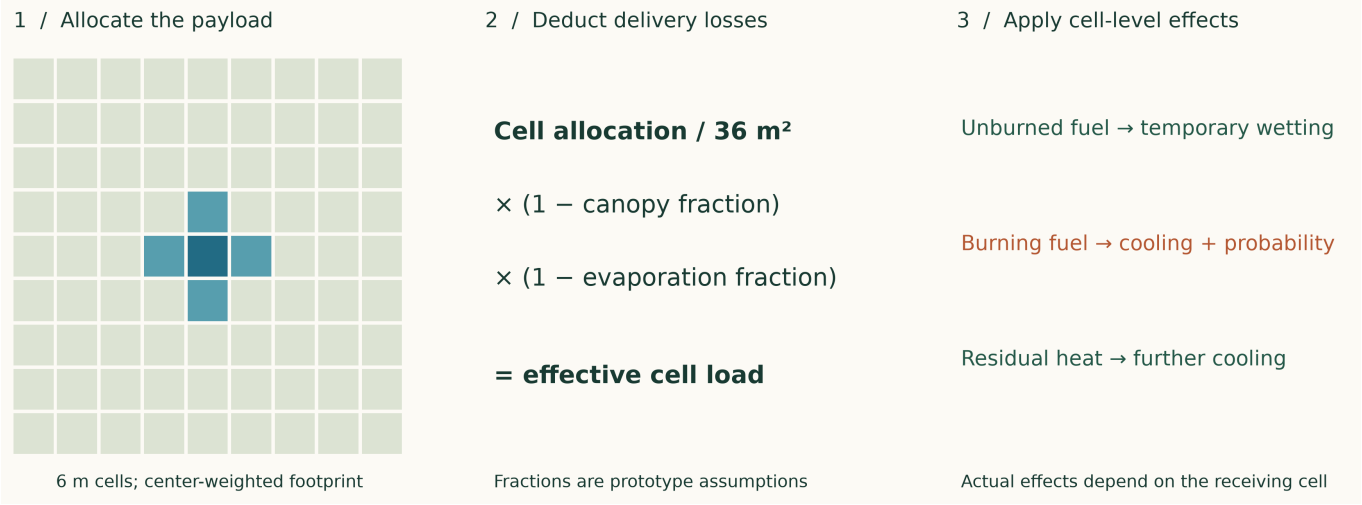
The signs of the atmospheric terms express a plausible response to heat, dryness and wind. Their linear form and exact percentages are not a droplet-evaporation law. They omit drop height, droplet spectrum, entrainment and flight through the fire plume. The canopy fractions likewise lack a recovered field fit. The word evaporation is the code's label for a simplified effective-delivery loss.

26. What measurements say about drop footprints

Johnson and Jordan's Forest Service guide tested a 324-US-gallon Bambi helibucket using collection bowls. Test airspeeds were 19-89 knots and release heights 25-120 ft. Water tests in Table 2 produced the longest line at coverage level 1 gallon per 100 ft² at 342 ft length, or 104.2416 m, at 78 knots. These conditions are quite different from a fixed, roughly one-cell-radius core drop. Johnson & Jordan, 2000

The coverage conversion is 3.785411784 L / 9.290304 m² = 0.407458 L/m² for 1 US gallon per 100 ft². A 910 L payload allocated over 180 m² averages 5.05556 L/m² before losses, over twelve times that coverage level. That comparison does not mean the game's entire footprint should be twelve times larger: the test describes coverage contours, not one uniformly wetted rectangle. Johnson & Jordan, 2000

The Canadian drop-effectiveness guide includes measured open and canopy footprints from Dryden trials in 2017-2019, using infrared estimates in the open and cups under jack pine. It combines distribution with an energetics model and distinguishes effective area from total wet area. Its assumptions about canopy and combustion-zone cooling are useful comparisons, not sources of the game's 28% and 20% losses. Relevant locations are sections 4 and 8 and Appendix A2. McFayden et al., 2023



The modeled water pathway separates allocated volume from effective load and the final cell response.

I would improve the concentrated prototype footprint before claiming realistic aircraft coverage. A measured distribution should be conditioned on platform, release height, speed, wind, material and canopy. Water conservation must be checked after resampling. A circular approximation may still be useful for teaching targeting, but its spatial scale should be identified explicitly.

27. Checking a recorded water drop

The new gameplay recording uses the updated 910 L bucket. The game stores each receiving cell's allocated liters and effective load, along with the water-demand inputs used to evaluate a burning cell. This lets the visible result be checked without guessing how the footprint was divided.

payload = escaped liters + intercepted liters + evaporated liters + effective liters

effective load per cell = effective liters / 36 m²

“Delivered” in the record means allocated inside the map before interception and evaporation. It is not added again in the conservation equation because the intercepted, evaporated, and effective terms already partition that delivered volume. Off-map allocation appears only in escaped liters.

The verification run checks that the four disjoint terms total 910 L within numerical rounding. It also verifies that affected burning cells record critical flow, required load, dose ratio, cooling fraction, and extinguishment probability. The companion run record preserves the values from the actual demonstrations.

This accounting follows the distinction between total aircraft payload and measured ground coverage in cup-and-grid testing. The game's small, weighted footprint is an analytical approximation; it does not reproduce a certified drop-pattern curve for a particular airframe. Suter, 2000 Johnson & Jordan, 2000

A successful water action therefore has several separate tests: aircraft availability, target geometry, water reaching a useful cell, sufficient effective load relative to demand, and the resulting cooling or extinguishment. The video and drop record show which of those stages occurred.

28. Water demand from published suppression equations

The updated model replaces the former resistance index with a surface-cooling demand calculation. Penney and colleagues express critical water flow as a baseline application rate plus the external heat flux divided by useful water enthalpy. Their equation gives a rate in liters per square meter per second. Penney et al., 2019

CF = 0.0129 + q_external / (0.70 × 2,640)

q_external = 0.27 × I_B / (2 L_f + D) + 0.077 × (1,090 - 588)

CF is L/m²/s, q_external is kW/m², I_B is kW/m, L_f and D are meters, and the temperatures are kelvin. The water enthalpy is 2,640 kJ/kg and water density is taken as 1 kg/L. Since kW equals kJ/s, the heat-flux term reduces to kg/m²/s and therefore L/m²/s. The source assumes transmissivity and view factor equal to one in this close-exposure calculation. Penney et al., 2019

The game estimates local flame length as its scenario flame potential in feet multiplied by 0.3048 and the square root of normalized source intensity. It converts this to a surface-fire intensity estimate using the metric Byram relation L_f = 0.0775 I_B^0.46, explained by Andrews. The inversion is I_B = (L_f / 0.0775)^(1/0.46). This calculation makes the water-demand units explicit; it does not replace the landscape spread engine with the full Rothermel model. Andrews, 2018

The selected active-flame depths are 2 m for grass, 4 m for mixed fuel, and 6 m for timber. To turn critical flow into a finite effective dose, the model uses a 30 s reference contact period and penetration multipliers 1.0, 1.5, and 2.0. The contact period and category mapping are simulation choices, while CF and its dimensional constants come from the cited paper. The paper modeled depths from 2 to 10 m; it did not fit these three game categories. Penney et al., 2019

required load = CF × 30 s × penetration multiplier

Critical flow is a theoretical limiting rate, not a guarantee of finite-time extinction. The finite-dose response is therefore defined separately in the next chapter. Penney's high-intensity forest assumptions and the surface flame-length relation also limit physical transfer to grass or crown-fire conditions.

29. A continuous water-dose response

After calculating required load, the model compares the effective water actually reaching a cell with that demand. The ratio x is dimensionless. A smooth response then determines cooling and the chance that active flame is extinguished.

x = effective load / required load

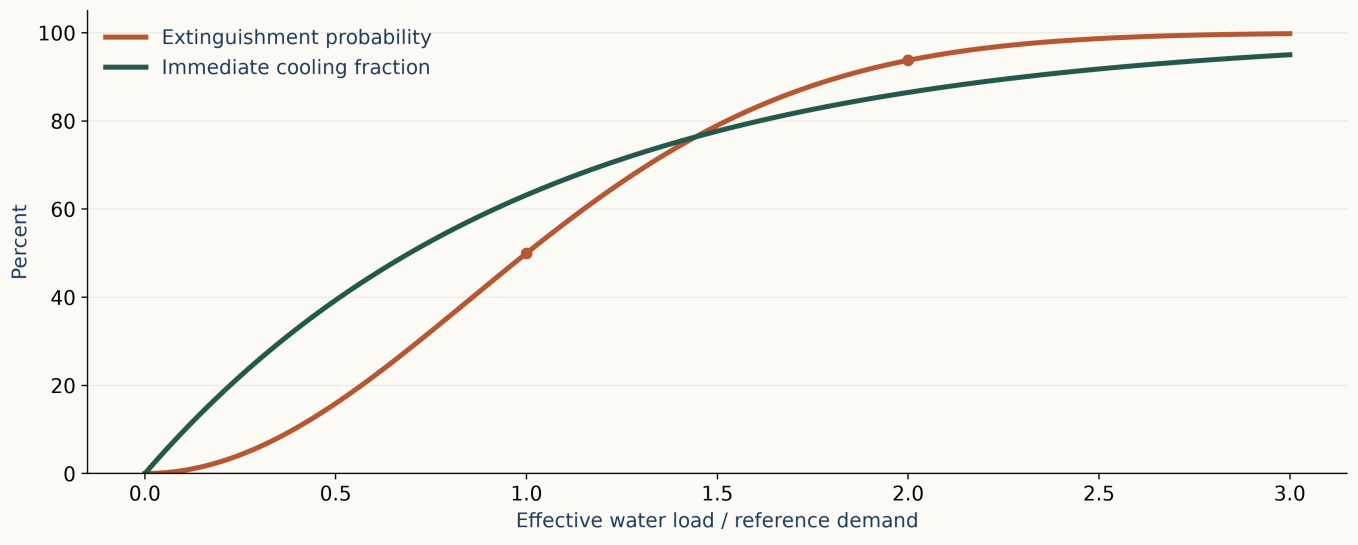
cooling fraction = 1 - exp(-x)

$P_{extinguish} = 1 - \exp[-\ln(2) \times x^2]$

Dose ratio x	Cooling fraction	Conditional extinguishment probability
0.00	0.00%	0.00%
0.50	39.35%	15.91%
1.00	63.21%	50.00%
1.50	77.69%	78.98%
2.00	86.47%	93.75%
3.00	95.02%	99.80%

The selected probability curve has its midpoint at one reference dose and increases smoothly. It removes the former jump to certain extinction at x = 2. Zero effective water produces no immediate cooling or extinction. A larger dose can still fail, while repeated treatment can further reduce active intensity and residual heat.

The exponential form is a simulation mapping. The cited water equations determine the demand term, but neither Penney's deterministic study nor the aerial mission studies measured this exact grid-cell curve. The slope, midpoint interpretation, and penetration multipliers remain explicit parameters rather than being attributed to those studies. Penney et al., 2019 Wheatley & Wotton, 2026



Updated analytical response to effective dose. Cooling and extinction are distinct outcomes; the curve has no special jump at twice the reference load.

30. Interpreting suppression research

The random draw is made once for each affected burning cell. If extinction occurs, flame intensity becomes zero and residual heat is retained. Otherwise, the model still applies cooling. Later time steps can extinguish a sufficiently weak, wet cell through the continuing combustion rule.

Penney and colleagues evaluated 540 constructed forest-fire scenarios using combinations of fuel loading, fire danger, and active flame depth. Their study shows how demand and supply must be compared with consistent units and how limited storage and refill time constrain suppression. It is a deterministic modeling study, with acknowledged limitations in validation for severe fires. I implemented its critical-flow relationship and exposed the added finite-dose choices in the parameter table. Penney et al., 2019

For a worked example, let local flame length be 4 ft = 1.2192 m in the grass mapping. Byram's inverse relation gives about 399.651 kW/m. With D = 2 m, the external heat flux is about 62.966 kW/m². Critical flow is 0.0469724 L/m²/s, and a 30 s contact period gives a reference load of 1.409173 L/m². A 36 m² cell therefore has a reference effective volume of 50.7302 L. These are calculations from the stated inputs, not an observed experimental burn. Penney et al., 2019 Andrews, 2018

If that cell receives 1.409173 L/m² after canopy and atmospheric losses, x = 1 and the selected extinction probability is 50%. If it receives twice that load, x = 2 and the probability is 93.75%. The amount released from an aircraft must be larger whenever water misses the cell or is intercepted.

Van Wagner and Taylor distinguish quenching flames, cooling hot fuel, and wetting fuel before ignition. This supports separate game outcomes for active flame, residual hotspots, and unburned wet fuel. Van Wagner & Taylor, 2022 McFayden and colleagues relate aerial drop effectiveness to distribution and fire behavior rather than payload alone. McFayden et al., 2023

Wheatley and Wotton's Ontario mission analysis relates fire-behavior reduction to initial behavior and perimeter-normalized drop stacking. It supports tracking repeated effort and defining the outcome carefully. Its aircraft, boreal fuel settings, mission scale, and ordinal fire-rank response differ from one helicopter-treated cell, so I use it to inform structure and evaluation rather than substitute a mission percentage for the cell curve. Wheatley & Wotton, 2026

31. Water, hotspots, and the recorded sequence

The updated Calm Weather Mop-up recording shows an initial dispatch interval, water directed at active cells, cooling of residual hotspots, and further observation. Captions identify the mechanic being demonstrated. The game is recorded at 2560 × 1440 pixels; the still figures retain higher-resolution captures.



Recorded updated gameplay after water attack. The interface separates active flame from residual heat and reports the remaining drop budget.

PLAY VIDEO 02 - Water demand, cooling, and follow-up observation

The drop record preserves the 910 L nominal payload and its allocation. A burning cell can be knocked down without going out. A cell that goes out can retain heat, and a later application can cool that heat. This separation follows the different suppression mechanisms discussed by Van Wagner and Taylor and avoids equating visible flame disappearance with a completed incident. Van Wagner & Taylor, 2022

The game requires a quiet observation period after both active flame and hotspots fall below its thresholds. The period is an exercise rule. Its purpose is to make the player reassess after suppression rather than end immediately after a successful drop.

The companion recorded-run file contains the final state and intermediate actions for this exact build. A demonstration establishes what happened for that seed and action sequence; statistical performance requires repeated runs with defined exposures.

32. Water drying and intensity recovery

Water load declines exponentially between applications. The hourly decay coefficient depends on wind, temperature, humidity and a fuel multiplier. The fuel multiplier is 1 for grass, 0.85 for mixed and 0.65 for timber, making the modeled surface water persist longer in timber.

$$k = [0.1 + 0.007w + 0.005 \max(0,T - 55) + 0.0015(100 - RH)] \times \text{fuelDrying}$$

$$L_{next} = L \times \exp(-k \times dt / 60)$$

For grass at 90°F, RH = 10 and wind = 30 mph, k = 0.62/h. An initial 2 L/m² becomes 1.712830 after 15 minutes. At 70°F, RH = 60 and wind = 5, k = 0.27/h and the same load becomes 1.869455. Under constant conditions the water half-lives are ln(2)/k: about 67.08 and 154.03 minutes.

Loads at or below 0.03 are removed from the wet-state list and reset to zero. This is a numerical cleanup threshold with an abrupt endpoint. It also affects which cells are eligible to be added to the wet list on another application. It has no established physical meaning as a moisture extinction threshold.

Burning intensity updates separately, with C = L/(L + 2):

$$I_{next} = \text{clamp}[I + 0.012dt(1 - C)(1 - I) - 0.018dt\,C, 0.02, 1.2]$$

For I = 0.5, L = 0 and dt = 5, intensity becomes 0.53. At L = 2, C = 0.5 and it becomes 0.47. Thus dry weak fire tends toward intensity 1 while water suppresses it. The coefficients are per modeled minute; intensity remains a normalized state rather than a temperature or kW/m.

Exponential drying is mathematically convenient and positive, but this particular coefficient is a response curve. It does not explicitly track intercepted water, infiltration, evaporation energy, rainfall or different droplet sizes after impact.

33. Residual heat and the possibility of reignition

Heat above 0.12 is retained as a hotspot. Its half-life without water is 22 minutes in grass, 40 in mixed fuel and 65 in timber. Water speeds decay by up to a factor of four.

heat_next = heat × exp[-ln(2) × dt / halfLife × (1 + min(3,0.25L))]

For an illustrative starting heat of 0.65 in mixed fuel, after 15 minutes with no water it is 0.501219. With a fixed L = 4 it becomes 0.386492. Holding water constant is an analytical simplification; the running model also dries it. The continuous time to fall from 0.65 to 0.12 is 97.496 minutes dry or 48.748 minutes with fixed L = 4. Stepwise checking rounds disappearance to a simulation checkpoint.

A drop onto an already burned hot cell now multiplies its heat by exp(-fresh load / requiredLoad). The required-load function uses the same critical-flow model, evaluated at a reference normalized intensity of 0.10. The fresh water is then added for subsequent time-dependent cooling. This reuses the revised water scale while retaining a separate hotspot state. Penney et al., 2019

The 0.12 threshold decides whether the model continues tracking heat and whether incident control is clear. It is not 0.12 degrees, a soil-temperature measurement, or a standard mop-up criterion. The response exists to make follow-up cooling meaningful. Its exact thresholds, fuel values and half-lives require calibration against appropriate thermal or reignition observations before any physical interpretation.

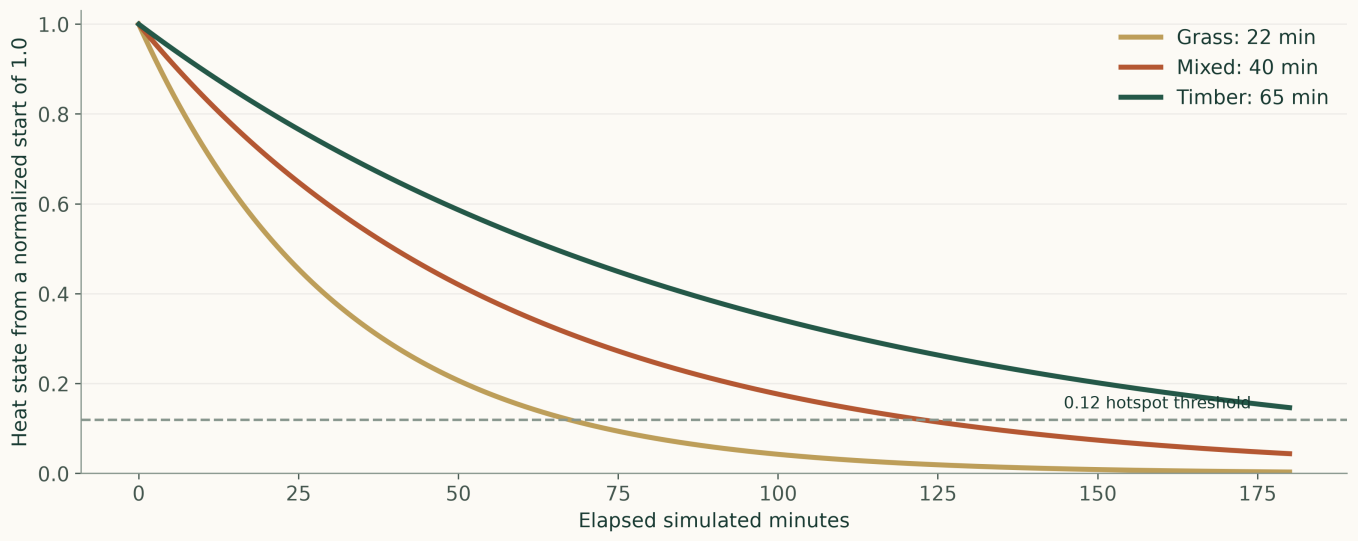
34. How a hotspot ignites adjacent fuel

For each hotspot that remains above threshold after cooling, the engine rounds the wind vector to select one nearby downwind cell. Only an unburned, fueled receiver is eligible. The probability depends on residual heat, time, fuel moisture, wind, fuel density and water in both source and receiver.

P_reignite = clamp[heat × dt/60 × 0.20 × (1 - M/45) × (1 + w/45) × F × exp(-L_source - L_receiver), 0, 0.25]

For a hypothetical post-decay heat of 0.65, dt = 5, M = 9, wind = 10 mph, full density and no water, the candidate probability is 0.65 × 5/60 × 0.20 × 0.8 × 1.222222 = 0.010593, about 1.06%. If source and receiver each carry 1 L/m², the added factor exp(-2) reduces that to about 0.1434%.

Successful reignition starts a neighboring cell at normalized intensity 0.55. The burned cell is not restored to unburned fuel. Thus the feature represents retained heat igniting adjacent receptive material, rather than regrowing and reburning all consumed fuel.



Analytical normalized heat decay. Water and the selected fuel half-life change when the model falls below its hotspot threshold.

The moisture divisor 45, wind divisor 45, multiplier 0.20, 0.25 cap, and starting intensity 0.55 are selected coefficients for the neighboring-fuel response. Their signs make wetter fuel less likely to ignite and stronger wind increase the modeled transfer. Sensitivity checks examine those directions separately from field calibration.

Verified control requires zero active cells and zero tracked hotspots plus 30 quiet minutes for the tutorial or 45 for other exercises. Earliest review is at 60 tutorial, 90 guided or 120 original-condition minutes. Regular core dispatch starts at minute 15. These periods are teaching and assessment rules.

35. Ember generation and launch opportunities

Spotting is a separate pathway from ordinary neighboring-cell spread. The weather precursor calculates a spotting index from sustained wind, a timber multiplier, and moisture bands. That index then controls whether an accumulated launch opportunity produces a flight.

spotIndex = 0.012w × timberFactor × drynessFactor

P_launch = clamp[0.60 × spotIndex × (1 - RH/100) × launchFuelFactor, 0, 0.8]

creditAdded = [3 + ceil(w/9) + √activeCells / 18] × dt/60

The precursor timber factor is 1.5 in timber and 1 otherwise. Dryness is 1.25 for M below 6, 1.10 below 10 and 1 otherwise. The additional launch fuel factor is 1.15 for timber and 1 otherwise. These separate multipliers mean timber influences both the precursor and the launch calculation.

Every whole ember credit is spent on one launch trial, whether the trial succeeds or fails. Consider mixed fuel with wind 18 mph, RH = 30%, M = 9 and 324 active cells. The spotting index is 18 × 0.012 × 1.1 = 0.2376. Launch probability is 0.2376 × 0.60 × 0.7 = 0.099792. Credits accrue at 3 + 2 + 18/18 = 6 per hour, or 0.5 per five minutes. Under constant conditions, six opportunities have an expected 0.598752 launches per hour.

For timber at wind 36 mph, RH = 20%, M = 5 and 324 active cells, the spotting index is 0.81. The launch probability is 0.44712 and credit production is 8 per hour. Expected launches are 3.57696 per hour before landing and ignition restrictions. These are calculated model expectations, not field ember counts.

Albini's spotting work separates firebrand source, lofting, transport and ignition-distance considerations. Albini, 1979 The game follows that separation at a coarse level but does not model brand size, mass, burnout, drag or plume structure. Its credit rate and multipliers are selected aggregation settings.

36. Ember loft time travel and landing

A launched flight selects an active source cell and constructs a downwind travel distance. Its variable named gust is a random multiplier, distinct from the weather gust value used elsewhere. The multiplier ranges from 0.65 up to 1.40.

G = 0.65 + 0.75U1

loftSeconds = clamp[12 + 1.4 FL + fuelLoft, 12, 70]

downwindMeters = w × 0.44704 × loftSeconds × G + 5 + 12U2

crossMeters = (U3 + U4 - 1) × 0.23 × downwindMeters

FuelLoft is 14 seconds for timber and 6 for other classes. The added 5-17 m distance prevents zero downwind travel in calm conditions. Crosswind displacement uses a triangular distribution with support approaching ±23% of downwind distance. These are assumed distributions, not measured firebrand transport parameters.

For a calculated grass example with wind 10 mph, FL = 4 ft and all uniform draws 0.5, loft time is 12 + 5.6 + 6 = 23.6 s and G = 1.025. Downwind distance is 10 × 0.44704 × 23.6 × 1.025 + 11 = 119.138976 m. Crosswind displacement is zero. Starting from cell center (50.5,78.5) with an eastward vector, grid flooring yields a center-to-center distance of 120 m.

For timber with wind 30 mph and FL = 10 ft, loft time is 40 s and distance is 560.8592 m before cell rounding; the selected receiving center is 558 m from that same source. At 90 mph in mixed fuel with FL = 14 ft, the computed 1,561.60 m flight leaves the map from that source.

The conversion from wind speed times time to distance is dimensionally sound. That does not validate the assumed loft time or travel multiplier. Real spotting depends on vertical wind, plume dynamics, brand properties and survival. The generated map boundary also censors long flights rather than representing the wider incident.

37. Ground landings and house exposure

An ember landing has two possible targets in the updated model: receptive ground and a nearby house. The ground-ignition calculation retains the version 6 fuel, moisture, wind, and water factors. A failed ground ignition can still expose a roof or vent when the landing occurs close enough to a house.

$$P_{ground} = \text{clamp}[(0.2 + 0.45 \times \text{density}/255 + w/180) \times (1 - M/40) \times \exp(-1.4L), 0, 0.86]$$

Only unburned cells with fuel are eligible for ground ignition. Roads, rivers, burned ground, and wet receiving cells reduce or remove that opportunity. The home pathway does not require the ground to ignite, reflecting firebrands entering building openings or accumulating on exterior components. NIST's experiments distinguish such exterior ignition and penetration mechanisms. Manzello et al., 2011

Each modeled flight represents a coarse ember-exposure event rather than a counted physical firebrand. For a home within five cells (30 m) of its endpoint, exposure is weighted by exp(-d/2), where d is distance in cells. The 30 m neighborhood and distance scale are selected subgrid aggregation settings; they are not measured vent-capture distances.



Recorded wind-driven spotting in the updated Marshall preset. The preset uses generated terrain and weather around historical context.

PLAY VIDEO 03 - Wind, ember transport, and house pathways

The house calculation preserves a decaying exposure store, so an earlier arrival can contribute to later ignition. This represents accumulated firebrand exposure. It does not assume that one visible particle corresponds to one physical ember or one inevitable house loss.

38. How homes can ignite

Each home now has a roof profile, vent profile, nearby-fuel condition, ember-exposure store, and separate pathway totals. A mixed neighborhood alternates vulnerable, maintained, and hardened profiles. The player can choose a uniform profile before starting an incident to compare preparation under the same seed.

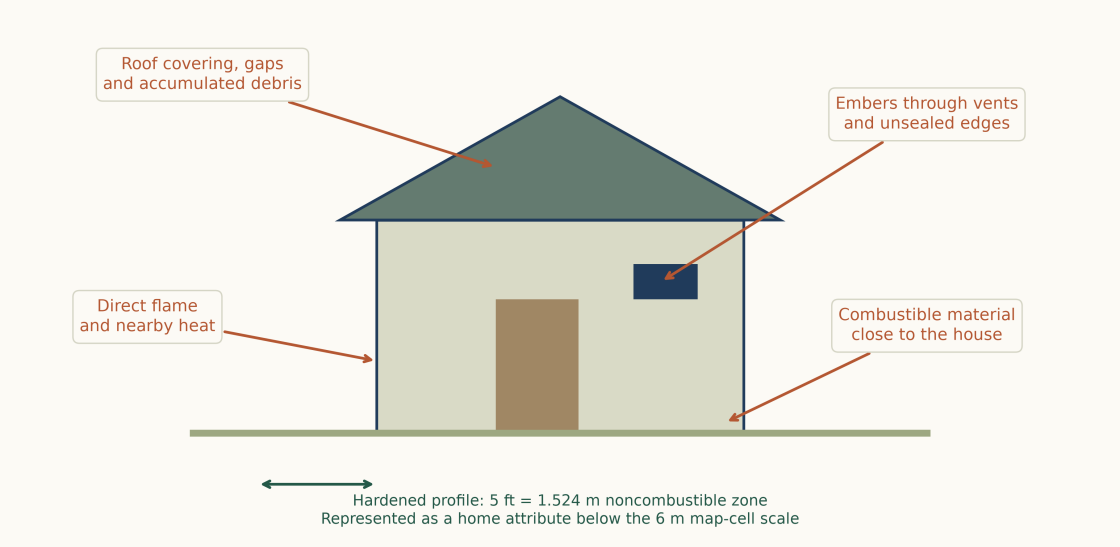
The home model combines four nonnegative hazards: nearby flame and heat, roof ember exposure, vent entry, and ember ignition of nearby combustibles. Adding hazards expresses multiple possible routes to ignition. If the random trial produces a loss, the game selects a contributing pathway in proportion to its hazard and records it in the home inspector.

$$H_{total} = H_{flame} + H_{roof} + H_{vent} + H_{nearFuel}$$

$$P_{home} = 1 - \exp(-H_{total})$$

The nearby-fire term uses the accumulated intensity-minutes within 2.5 cells of the home center, multiplied by 0.009 and (1 + gust/100). This preserves the original local heat exposure while adding independent ember routes. Embers can therefore ignite a home even when no neighboring ground cell has burned.

NIST's pathway guidance shows why that separation matters: an ember can ignite detached combustible material, then a fence, vehicle, deck, and finally a home. The game aggregates these possible near-home connections into one nearby-fuel pathway rather than constructing a detailed object network. NIST, n.d.



The updated home model distinguishes direct heat, roof exposure, vent entry, and combustible material near the building. A hardened profile reduces specific pathways.

The homes-standing score remains the fraction not lost. Evacuation is tracked separately because protecting people does not make a building resistant to ignition. A standing home and a controlled wildfire remain separate outcomes.

39. Roof assemblies, debris, and vents

NIST's full-scale firebrand experiments show that roof vulnerability depends on the assembly and openings, not just the outer covering. Firebrands can collect in gaps and interact with combustible material below tiles or with accumulated debris. The revised hardened profile therefore specifies a Class A assembly, closed edge gaps, and clean gutters rather than treating a named roof material as immunity. Manzello et al., 2011

IBHS requires a Class A roof covering or qualifying assembly, maintenance that keeps vegetative debris off the roof, and noncombustible closures at relevant barrel-tile or corrugated-covering openings. These are the features represented by the maintained and hardened roof descriptions in the game. IBHS, 2025

For vents, NIST tested openings including 5.72, 2.74, 2.00, 1.55, 1.35, and 1.04 mm at 7 m/s tunnel wind. Firebrands could continue burning until small enough to pass. Smaller mesh reduced some ignition outcomes, but did not universally stop fine-fuel ignition; for shredded paper behind 1.04 mm mesh, one of three experiments produced smoldering ignition. Those are test outcomes for that arrangement, not annual house-loss probabilities. Manzello et al., 2011

IBHS vent experiments also compare mesh alone with screened or baffled designs. Their results support treating vent design and downstream ember behavior separately from nominal opening size. The hardened game profile uses an ember-resistant vent with sealed edges, while the other profiles use 6 mm and 3 mm mesh descriptions. IBHS, n.d.

Profile	Roof response factor	Vent response factor
Vulnerable	1.00	1.00 (6 mm mesh)
Maintained	0.35	0.60 (3 mm mesh)
Hardened	0.15	0.15 (ember-resistant vent)

These factors are selected relative susceptibility coefficients. The experiments establish distinct mechanisms and comparative protection, but do not supply these exact whole-home factors. None is zero: improved construction reduces the corresponding pathway without implying a fireproof house.

40. The five-foot zone and nearby fuels

The hardened profile includes a noncombustible zone extending 5 ft, or exactly 1.524 m, from the building and attached features. IBHS identifies this as the critical immediate zone and calls for removal of vegetation, combustible groundcover, and relevant combustible items or connections. Its guidance addresses the whole area and attachments, not simply a clean patch below one window. IBHS, 2025

Because the landscape cells are 6 m wide, the game does not erase a full six-meter ring to represent a five-foot zone. It stores the zone as a home attribute. In the hardened profile, that attribute sets the ember-to-nearby-fuel susceptibility to zero. Other routes remain active: a roof or vent can still receive embers, and the direct heat term can still contribute.

The vulnerable profile has nearby-fuel factor 1; maintained homes use 0.4; hardened homes use 0. This represents the connection from ember-receptive material near a house into the building. It is not a universal measured percentage reduction from creating a five-foot zone. The zero value removes that particular modeled route because the defined zone is noncombustible.

NIST illustrates paths through fences, sheds, vehicles, decks, and other parcel components. The current game combines these into a near-fuel term. It does not simulate a separate vehicle fire, deck flame, or neighbor-to-neighbor radiation calculation. Those mechanisms explain the abstraction and define what a later, more detailed building model would need. NIST, n.d.

This distinction makes preparation meaningful without overstating it. The player can reduce one pathway completely within the simplified model while other threats continue to require suppression and observation.

41. Ember accumulation and a worked home example

Each house stores ember exposure E. Between evaluation times, existing exposure decays with a selected 20-minute half-life. A nearby flight adds exp(-d/2), where d is endpoint distance in grid cells. Only events in the elapsed interval are added.

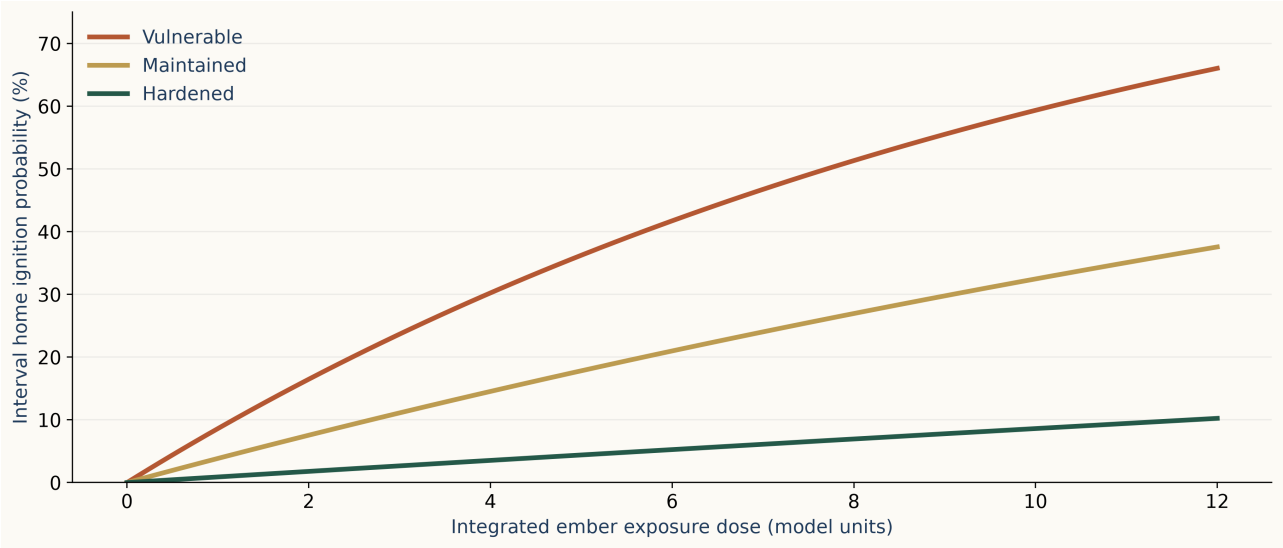
E_next = E_old × exp[-ln(2)dt/20] + arriving exposure

D_ember = {E_old × [1 - exp(-k dt)] / k + arriving exposure × dt/2} / 15

k = ln(2) / 20 per minute

The integrated dose accounts for the persistence of old exposure. New events are assigned half an interval of average exposure; this is a simple within-interval timing approximation. Roof hazard is 0.035 × D_ember × roofFactor; vent hazard is 0.025 × D_ember × ventFactor; nearby-fuel hazard is 0.030 × D_ember × nearFuelFactor.

For equal integrated dose D_ember = 4 with no nearby flame, the vulnerable home has roof hazard 0.14, vent hazard 0.10, and near-fuel hazard 0.12. Total hazard is 0.36 and ignition probability is 1 - exp(-0.36) = 30.2324%. The hardened home has roof hazard 0.021, vent hazard 0.015, and zero near-fuel hazard. Total hazard is 0.036 and probability is 3.5360%.



Equal-exposure analytical comparison from the selected coefficients. The percentages are simulation outputs, not observed household loss rates.

NIST's experiments support accumulated exposure and multiple ignition mechanisms. The coefficients above define how those mechanisms are represented in this exercise; they are not experimental whole-house risk estimates. Manzello et al., 2011

42. Ground resources and what they change

The optional structure-engine action still assigns protection to up to two homes for two simulated hours and wets adjacent terrain. In the revised house calculation, its 0.30 factor applies to nearby flame/heat and the nearby-fuel pathway. It no longer acts as one blanket multiplier on every source of house ignition. Roof and vent exposure remain governed by their building attributes.

This separation reflects the difference between reducing exterior exposure and changing a building opening. A crew's presence does not turn a mesh vent into an ember-resistant assembly. The two-home limit, duration, and 0.30 response factor remain selected resource abstractions; the NIST Waldo Fire case study supplies context for defensive actions and exposure rather than this exact factor. Maranghides et al., 2015

The optional spot-patrol action now uses the same water-response function as aircraft instead of a fixed large intensity reduction. It assigns a 400 L reference allocation over a 13-cell disk, giving $400/(13 \times 36) = 0.854701$ L/m² per potential cell. That allocation is an effective treatment budget for the exercise, not a specific engine specification. Actual reach depends on which cells are burning.

Patrol extinction can leave residual heat, so follow-up observation remains necessary. Water applied to nearby fuel continues to dry through the existing moisture rule. Hand crews and dozers create fuel discontinuities rather than applying water to active combustion.

The core controls remain line and helicopter attack. Optional resources provide additional comparisons, and their assumptions are kept explicit so a generic game resource is not confused with a field apparatus or staffing standard.

43. Scenario values and optional controls

Version 6 includes 37 cards, including the guided tutorial and fictional exercises. Each scenario configuration provides fuel, slope, temperature, wind and humidity arrays, direction, homes and incident context. The companion scenario register exports each configuration separately, with every array element retained and identified as a preset input.

Historical acres and narrative losses on a card are contextual metadata. They do not resize the 399.67-acre grid. The homes in the simulation are generated exercise targets, not a mapped sample of every historical structure. Temperature, slope and the hourly arrays are not automatically measured observations merely because the card has a real fire name.

Difficulty uses a heuristic based on peak wind, slope, fuel and minimum humidity. The coefficients are 0.75 for peak wind, 0.30 for slope and 0.45 for the humidity term; fuel offsets are 4 grass, 10 mixed and 17 timber. Scores are capped at 98 except for special tutorial and extreme cases. Tier boundaries are 20, 40, 60 and 80. These numbers rank exercises, not forecast containment probability.

Core controls provide five drops per 15-minute turn and a shared line-construction budget. The aircraft choice defaults to utility bucket, the core footprint fixes both axes to one cell and corrected aiming is selected. Expanded controls retain five drops per hour, three line methods, aircraft and pattern options, and other resources. Their retained presence does not make them part of every core demonstration.

Optional engines protect up to two homes for 120 minutes and wet nearby fuel. The updated house calculation applies protection to exterior heat and nearby-fuel hazards. Spot patrol now uses a 400 L reference allocation through the same water-response function as helicopter suppression. Evacuation protects people separately from building survival.

The register includes these values as gameplay assumptions. It also identifies the eight resource-unit hourly budget and three-orders-per-type limit. The default two-action lesson should be interpreted using its actual settings, which are recorded with each video.

44. Why the equations are structured this way

The spread model uses a product of environmental factors because several conditions act together: dry fuel, wind alignment, slope, intensity, and fuel continuity. The exponential hazard conversion makes the resulting transition bounded. Cellular-automaton studies support this local state-transition architecture, while physical fire models explain the roles of the environmental variables. Ghosh et al., 2024 Zhou et al., 2025 Rothermel, 1972

Line crossing uses survival over time, so dividing an interval into smaller equal intervals preserves the reference probability when conditions are fixed. Water uses dimensional heat and flow terms before a separate dose-to-response mapping. Home ignition uses additive hazards so improvements can affect one route while another remains active.

These choices make the mechanisms testable. Zero water should produce zero immediate water response. Doubling effective dose should not reduce cooling. A hardened vent should reduce the vent term without altering direct flame exposure. Repainting the interface should not create another ignition opportunity. A fixed seed should reproduce the same sequence on either execution path.

The random process describes uncertain outcomes within the model. Randomness does not itself make the parameters empirically calibrated. Evidence-based dimensional inputs, transparent probability mappings, and explicit exposure definitions make it possible to examine that calibration separately.

Keeping the deterministic parts inspectable is useful for teaching. The player can reason about why a line pauses, why water misses, why a wet patch later dries, or why an ember threatens a house even when the nearest ground is not burning.

45. A practical calibration dataset

I would begin calibration by choosing a specific prediction target: cell arrival time, successful line engagement, delivered water footprint, local fire reduction, or structure loss. Fitting all of them to one final burned-area total would allow errors in one mechanism to compensate for errors in another.

LANDFIRE can supply vegetation and fuel-model layers for a mapped landscape. Its products require explicit version and resolution choices. Resampling them onto six-meter cells does not create six-meter information. USGS 3DEP can supply elevation for local slope and aspect, with appropriate treatment of vertical units and coordinate systems. LANDFIRE, n.d. USGS, n.d.

Time-stamped fire perimeters, incident maps and progression observations would constrain where fire had reached at known times. Weather-station records and gridded weather would provide changing inputs, with uncertainty for distance, terrain and measurement height. Neither a final perimeter nor a single weather maximum identifies every five-minute exposure.

MTBS supplies mapped fire effects, including burn severity, and is useful for final spatial comparison. Burn severity concerns effects on biomass; it does not directly measure minute-by-minute intensity, flame length or the time a cell ignited. MTBS, n.d.

Line calibration needs completed geometry, completion time, treatment type and engagement status. Drop calibration needs aircraft type, release time, flight conditions, material, measured or defensible delivered-water distribution, and pre/post fire state. Structure calibration needs exposed surviving buildings as well as losses, construction attributes and nearby fuels.

I would keep raw observations separate from derived features and record every conversion. Missing completion times or uncertain perimeter placement should be represented as uncertainty rather than replaced with exact values. The research bibliography and evidence ledger identify which sources could supply each kind of information and which cannot.

46. Fitting testing and reporting uncertainty

Calibration should use training incidents to estimate coefficients, a separate validation set to choose among candidate model structures, and a final untouched set to assess generalization. Splitting neighboring pixels from the same fire across all sets would leak common weather and management conditions and overstate performance.

For arrival times, I would compare time error and burned-area overlap at multiple checkpoints. For conditional ignition or line crossing, reliability plots, log loss and Brier score would check whether predicted probabilities agree with observed frequencies. For structures, I would assess loss probabilities among exposed buildings and compare performance across construction and exposure groups.

Line outcome fitting must address selection: crews choose where and when to build. A held line in mild terrain is not proof that its construction type would succeed at a different intense front. Aircraft missions are similarly selected, often because a fire is difficult. This is a reason to model operating conditions and placement, rather than copying an aggregate success percentage.

Sensitivity analysis should first vary one coefficient to identify direction and saturation, then vary plausible combinations to expose interactions. Stochastic ensembles should report seeds, action policies and confidence intervals. Shared random streams can improve controlled comparisons, but different call sequences need explicit handling.

I would prioritize spatial water coverage, line engagement logging, measured weather and mapped fuels before fitting every small coefficient. A well-fitted extinguishment curve still gives misleading results if almost the whole aircraft payload lands in an unrealistically concentrated footprint.

The research revisions described in this report have been implemented and tested in the updated build. Further statistical calibration would require defined response variables, suitable field data, and evaluation on incidents not used for fitting.

47. What the updated game demonstrates

The updated game connects research to playable consequences. Crew work varies with the mapped fuel category. Water is conserved across its footprint, compared with a published critical-flow calculation, and applied through a continuous response. House preparation changes specific roof, vent, and near-fuel pathways, while the game records the contributing route when a loss occurs.

The update also improves internal consistency. Reference line probabilities compose correctly across time steps. The background worker and browser fallback agree for a matched seeded sequence. Home exposure is applied once per elapsed interval. The user interface exposes the new home attributes and links directly to this documentation.

All 37 scenario presets were started and advanced through the initial fifteen-minute interval without browser errors. Targeted checks verified production conversions, water accounting, monotonic response, the five-foot unit conversion, and the expected effect of hardening. These checks support the stated implementation behavior.

The report's research citations explain why the mechanisms and dimensional values were chosen. The selected response coefficients remain visible in the appendix because matching field data would be needed to estimate them statistically. The exercise is intended to make those choices understandable and open to revision.

The next useful comparison is between matched runs with the same weather and initial conditions: line placement and timing, water coverage, and home preparation can then be changed one at a time. The updated inspector and recorded data make these comparisons easier to audit.

48. Updated parameter register

The accompanying parameter register gives the current value and evidence class for the main revised quantities. It is tied to the updated build. The earlier register remains part of the archived version 6 audit rather than being presented as the new implementation.

Parameter	Current value	Basis
Type I indirect crew	627 / 323 / 455 ft/h	NWCG; grass / mixed / timber mapping
Utility bucket	910 L	SEI nominal BB2024 reference
Water baseline flow	0.0129 L/m²/s	Penney equation 1
Water efficiency	0.70	Penney modeling assumption, adopted
Water enthalpy	2,640 kJ/kg	Penney equation 1
Radiative fraction	0.27	Penney equation 2
Convective coefficient	0.077 kW/m²/K	Penney equation 2
Gas / fuel temperature	1,090 / 588 K	Penney modeling conditions, adopted
Dose contact time	30 s	Selected finite-dose mapping
Active depth	2 / 4 / 6 m	Selected fuel-category mapping
Penetration factor	1 / 1.5 / 2	Selected fuel-category mapping
Dose probability	$1 - \exp[-\ln(2)x^2]$	Selected smooth response
Home zone	1.524 m	IBHS 5 ft dimension
Ember-store half-life	20 min	Selected exposure memory

Source support concerns the specified quantity and conditions. For example, Penney uses a 0.70 water efficiency; adopting it does not imply every helicopter drop actually achieves that efficiency. The game separately accounts for canopy and atmospheric delivery losses before applying its effective-dose response. Penney et al., 2019

The code exposes these settings through the research panel's linked documentation and the build's research parameter object. The input register gives a compact route from a number in the report to its implemented role.

49. Verification and recorded evidence

The verification suite checks behavior that could materially contradict this report. It confirms exact production conversions, zero-dose behavior, the 50% midpoint of the selected dose curve, continuity around the former x = 2 cutoff, and increasing response as dose rises. It checks conservation of the 910 L bucket across disjoint loss and effective-delivery terms.

Home checks confirm that ember-only exposure can produce risk; the hardened profile reduces that risk; the 5 ft buffer converts to 1.524 m; and engine protection reduces the intended exterior terms without automatically reducing vent entry. Repeating a home-state refresh at the same simulation time leaves cumulative exposure unchanged.

The simulation check compares a matched seeded sequence with and without the background worker, including line construction. It also starts and advances all 37 presets. No browser errors occurred in these checks. The saved verification files give the individual pass/fail results rather than an unsupported claim of field validation.

Three new demonstrations show the updated build: crew production and line timing; water application and hotspots; and severe-wind spotting with house inspection. Each recording stores its scenario, mode, seed, build identity, and intermediate state. The credit is “Recorded gameplay demonstration.”

The PDF's images are actual frames from those recordings or analytical figures calculated from the stated equations. Posters link to public playback pages, and MP4 copies are attached inside the PDF. The downloadable game and PDF describe the same release.

50. Research bibliography 1

Ghosh et al., 2024

Ghosh, R., Adhikary, J., & Chemlal, R. (2024). Fire spread modeling using probabilistic cellular automata [Preprint]. arXiv.

Supplied PDF pp. 6 and 8; model state and probability sections.

Read source

Zhou et al., 2025

Zhou, Z., Xu, Q., Yi, J., Li, Y., Zhang, S., & Li, W. (2025). Modeling the spread of forest fires through cellular automata by leveraging deep learning to derive transition rules. Ecological Informatics, 88, 103150.

Supplied PDF pp. 6-7 and discussion p. 14.

Read source

Rothermel, 1972

Rothermel, R. C. (1972). A mathematical model for predicting fire spread in wildland fuels (Research Paper INT-115). USDA Forest Service.

Model development and equations for reaction intensity, propagating flux, wind, slope and moisture.

Read source

Andrews, 2018

Andrews, P. L. (2018). The Rothermel surface fire spread model and associated developments A comprehensive explanation (RMRS-GTR-371). USDA Forest Service.

Chapters on model inputs, wind and slope, fireline intensity and flame length.

Read source

NWCG, 2021

National Wildfire Coordinating Group. (2021, July 26). Fireline production tables [Fuels Management Committee memorandum and tables].

PDF p. 3 sustained crew rates; p. 6 initial-action caution; p. 8 dozer rates.

Read source

51. Research bibliography 2

NWCG, 2021, PMS 437-1

National Wildfire Coordinating Group. (2021, September). NWCG guide to fireline fire assessment (PMS 437-1).

Printed p. 34 (PDF p. 39), Table 1: fire behavior classes and flame-length tactical interpretation.

Read source

Gannon et al., 2020

Gannon, B. M., Thompson, M. P., Deming, K. Z., Bayham, J., Wei, Y., & O'Connor, C. D. (2020). A geospatial framework to assess fireline effectiveness for large wildfires in the western USA. Fire, 3(3), 43.

Methods 2.1-2.3; results 3.3; Table 4.

Read source

USDA Forest Service, n.d.

USDA Forest Service. (n.d.). Fireline effectiveness FLE dashboard. Rocky Mountain Research Station.

Overview and linked dashboard data description.

Read source

SEI Industries, 2021

SEI Industries. (2021). Bambi bucket [Product brochure]. DART Aerospace.

Capacity table, BB2024 row, nominal 240 US gal / 910 L.

Read source

NWCG, 2019

National Wildfire Coordinating Group. (2019). NWCG standards for helicopter operations (PMS 510).

Aircraft capability, load calculations and performance limitations.

Read source

52. Research bibliography 3

Suter, 2000

Suter, A. (2000). Drop testing airtankers A discussion of the cup-and-grid method (0057-2868-MTDC). USDA Forest Service.

History of the cup-and-grid method; 1990s testing and error discussion.

Read source

Johnson & Jordan, 2000

Johnson, G., & Jordan, C. (2000). Ground pattern performance of the SEI Industries Bambi 324-gallon helibucket (0057-2802-MTDC). USDA Forest Service.

PDF pp. 1-2; test conditions and Table 2 water tests.

Read source

McFayden et al., 2023

McFayden, C. B., Wotton, B. M., Robinson, J. W., Johnston, J. M., Cantin, A. S., Jurko, N. M., Boucher, J., Wheatley, M., Ansell, M., Boychuk, D., & Russo, B. (2023). Reference guide to the drop effectiveness of skimmer and rotary wing airtankers (GLC-X-35). Natural Resources Canada.

Printed pp. 3-5, 22-30, 63-66 and Appendix A5; PDF has 193 pages including front matter.

Read source

Van Wagner & Taylor, 2022

Van Wagner, C. E., & Taylor, S. W. (2022). Theoretical amounts of water to put out forest fires (BC-X-458). Natural Resources Canada.

Sections 2-3; equations 7b, 10b, 13b; Table 3.

Read source

Wheatley & Wotton, 2026

Wheatley, M., & Wotton, M. (2026). Empirical models of initial attack fire behaviour reduction by skimmer airtankers in Ontario's boreal forest. International Journal of Wildland Fire, 35(10), WF25311.

Methods; Results; Table 5; Figures 5-6; published 15 October 2026.

Read source

53. Research bibliography 4

Wotton, 2009

Wotton, B. M. (2009). A grass moisture model for the Canadian Forest Fire Danger Rating System [Conference paper 3A.2]. American Meteorological Society.

Introduction and moisture exchange model discussion.

Read source

Albini, 1979

Albini, F. A. (1979). Spot fire distance from burning trees A predictive model (INT-GTR-56). USDA Forest Service.

Model formulation and assumptions concerning firebrand lofting and transport.

Read source

Maranghides et al., 2015

Maranghides, A., McNamara, D., Vihnanek, R., Restaino, J., & Leland, C. (2015). A case study of a community affected by the Waldo Fire Event timeline and defensive actions (NIST TN 1910). National Institute of Standards and Technology.

Study area, event reconstruction and defensive-action analysis; Mountain Shadows destruction count.

Read source

NIST, n.d.

National Institute of Standards and Technology. (n.d.). Fire spread pathways. Hazard Mitigation Methodology.

How fire spreads in the WUI; illustrated pathways.

Read source

National Weather Service, n.d.

National Weather Service Boulder. (n.d.). Marshall Fire and high wind on December 30 2021. NOAA.

Event overview and observed wind reports.

Read source

54. Research bibliography 5

LANDFIRE, n.d.

LANDFIRE. (n.d.). Data. U.S. Department of Agriculture and U.S. Department of the Interior.

Data access and product descriptions.

Read source

USGS, n.d.

U.S. Geological Survey. (n.d.). About 3DEP products and services.

Source data products; lidar, DEMs and metadata.

Read source

MTBS, n.d.

Monitoring Trends in Burn Severity. (n.d.). Fire mapping Frequently asked questions.

What is burn severity.

Read source

Penney et al., 2019

Penney, G., Habibi, D., Cattani, M., & Carter, M. (2019). Calculation of critical water flow rates for wildfire suppression. Fire, 2(1), 3.

Section 2, equations 1-2; sections 5-7.

Read source

Manzello et al., 2011

Manzello, S. L., Suzuki, S., & Hayashi, Y. (2011). Summary of full-scale experiments to determine vulnerabilities of building components to ignition by firebrand showers. NIST SP 1126.

Vent penetration experiments and Table 1; roofing assembly experiments.

Read source

55. Research bibliography 6

IBHS, 2025

Insurance Institute for Business & Home Safety. (2025, December). Wildfire Prepared Home technical standard.

Sections 2.1.1 (roof), 2.1.3 (vents), 2.2.1 (0-5 foot noncombustible zone).

Read source

IBHS, n.d.

Insurance Institute for Business & Home Safety. (n.d.). Vulnerability of vents to wind-blown embers.

Experimental vent comparisons and discussion of mesh and baffle designs.

Read source