

Weakened downslope winds and slower simulated nighttime spread of the Camp Fire under warming

Weichen Liu

University of Wyoming

Stefan Rahimi

University of Wyoming

Maria Frediani

NSF National Center for Atmospheric Research

Benjamin Hatchett

Cooperative Institute for Research in the Atmosphere

Chad Thackeray

University of California, Los Angeles

Ertugrul Taciroglu

University of California, Los Angeles

Fanghe Zhao

fzhao@uwyo.edu

University of Wyoming

Alex Hall

University of California, Los Angeles

Karandeep Singh

University of California, Los Angeles

Timothy Juliano

AiDASH Inc.

Article

Keywords:

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11020281/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

**Weakened downslope winds and slower simulated nighttime spread of the
Camp Fire under warming**

Weichen Liu¹, Stefan Rahimi^{1,2}, Maria Frediani³, Benjamin J. Hatchett⁴, Chad Thackeray²,
Ertugrul Taciroglu², Fanghe Zhao^{1*}, Alex Hall², Karandeep Singh², Timothy W. Juliano⁵

¹University of Wyoming, Laramie, WY, United States

²University of California Los Angeles, Los Angeles, CA, United States

³NSF National Center for Atmospheric Research (NSF NCAR), Boulder, CO, United States

⁴Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins,
CO, United States

⁵AiDASH Inc., Palo Alto, CA, United States

*Corresponding author: Fanghe Zhao (fzhao@uwyo.edu)

16 **Abstract**

17 Wildfire projections often assume warming dries fuels but leaves the winds that drive fire spread
18 unchanged. This assumption may be important for downslope windstorm fires, which can cause
19 severe losses. Here we show, using coupled fire and weather simulations of the 2018 Camp Fire,
20 how warming of 1.5 to 6 K changes fire behaviour. Six simulations were run per scenario. Fire
21 spread slowed steadily with warming, especially after midnight. Warming reduced the
22 temperature contrast driving downslope winds, weakening the driving force by about half at 6 K.
23 In the flaming-front area, wind speed fell 16% and spread rate 20%, exceeding differences
24 caused by small changes in starting conditions. Prognostic fuel drying enlarged the fire, yet
25 warming shrank burned area in both fuel treatments because weaker winds outweighed the
26 drying. Our simulations exclude spotting, which drier fuels could increase. Projections holding
27 winds fixed may therefore overestimate flaming-front area spread under warming.

28

1. Introduction

At 06:33 local time on 8 November 2018, the Camp Fire ignited in the Feather River canyon of northern California during a Diablo-like offshore downslope windstorm. Winds near the ignition point had remained above 12 m s^{-1} through the preceding night, with gusts exceeding 23 m s^{-1} , while 100-h dead-fuel moisture at nearby stations was near record minima^{1,2}. Firebrands ignited spot fires in Paradise within two hours^{3,4}. By the end of the event, the fire had burned 620 km^2 , killed 85 people and destroyed 18,804 structures⁵. Much of this damage resulted from the transition of a wildland fire into an urban conflagration, which depended on wildland fuels, ember exposure, structure spacing, and community design⁶⁻⁸.

Downslope windstorms such as the one driving the Camp Fire develop when synoptic pressure gradients cross topographic barriers, producing strong, dry flow on the lee side⁹. Across the western United States from 1992 to 2020, downslope-wind-driven fires accounted for 13.4% of fires and 11.9% of burned area, yet had nearly three times the odds of causing a fatality¹⁰. This distinction matters because rapid wildland spread and dense ember exposure can occur together, linking the meteorological driver to structural loss. In Southern California, Santa Ana winds similarly produce severe fire weather and substantially increase daily burned area, especially before the first wetting rains of autumn¹¹⁻¹⁴. Recent examples include the January 2025 Eaton and Palisades fires¹⁵. Related fire-weather regimes occur during foehn-like winds east of the Australian Alps and bora winds along the Adriatic coast^{16,17}. In each setting, terrain channels and accelerates the flow, and strong winds can overwhelm fuel-age effects¹⁸⁻²⁰.

Climate change modifies both the fuels and the meteorological conditions associated with these events. Warming, earlier snowmelt, rising vapour pressure deficit (VPD), and lower fuel moisture due to both warmer and drier summers have lengthened fire-weather seasons and increased western U.S. forest fire activity. Anthropogenic warming approximately doubled cumulative forest area burned from 1984 to 2015 relative to an estimated climate without human forcing²¹⁻²⁵, and an anthropogenic signal is detectable in regional and global fire-weather indices²⁶. Future burned area remains harder to estimate because historical relationships between VPD and fire may not extrapolate reliably under continued warming and may overestimate future burned area²⁷. More importantly here, fuel drying and wind also need not change together. Projections for Southern California generally indicate fewer Santa Ana events and, in some models, weaker winds, although the remaining events may be hotter and drier²⁸⁻³⁰. Consequently, increased fuel flammability can coincide with weaker dynamical forcing, making the net fire response dependent on both processes.

Such climatological projections describe changes in event frequency and intensity over decades, but cannot isolate how a given windstorm responds under comparable synoptic conditions. Event-based thermodynamic perturbation experiments, commonly called pseudo-global-warming (PGW)³¹⁻³⁴ or storyline experiments, provide this comparison. Hydrostatic theory provides a plausible event-scale mechanism. Atmospheric-layer thickness increases with mean virtual temperature, so warming changes the pressure a column exerts at its base³⁵⁻³⁷. For a downslope windstorm the relevant column is shallow: the nocturnal force behind the flow rests largely on the cold, dense air banked over the upstream terrain, and how fast warming erodes that cold layer, relative to the cold pool in the valley below, decides how much of the force survives. Because pressure reduced to sea level over steep terrain depends on an assumed subsurface temperature³⁸, we evaluate the horizontal pressure-gradient force directly on the model's native

levels. Whether this adjustment is large enough to alter fire spread during an observed windstorm remains uncertain. Timing also matters because Diablo and Santa Ana winds often peak overnight or in the morning, placing any wind-mediated warming response within the usual period of nocturnal fire relief^{2,39,40}.

Fire suppression often benefits from cooler temperatures, humidity recovery, and reduced fire activity at night. Warming can weaken this thermodynamic barrier, and nocturnal drying can remove it abruptly in coastal California^{41–43}. During downslope windstorms, strong nighttime winds may also diminish the usual overnight relief. Determining how these competing responses interact could improve fire-behaviour forecasts, inform future public safety power-shutoff criteria, and support catastrophe-risk modelling⁴⁴.

To examine these competing responses, we simulate the Camp Fire with the coupled fire module of the Weather Research and Forecasting model^{45,46}, under baseline conditions and PGW perturbations of +1.5 to +6 K. Each scenario includes six ensemble members. Gradient-boosted regression (XGBoost) and Shapley additive explanations (SHAP)⁴⁷ quantify how the predictive contributions of meteorological and surface variables differ among warming levels. For the fuel-moisture experiments, an exact decomposition of the model's fire spread-rate function separates fuel, wind, slope and interaction contributions. These experiments are designed to compare warming scenarios. Because the model does not represent firebrand transport and spot fire ignition, the simulations are expected to underpredict the westward advance toward Paradise, as found in an earlier coupled simulation of this event⁴⁸. Radar analysis attributes the observed westward run to long-range spotting⁴⁹. This omission is common to all scenarios, although its interaction with warming cannot be quantified here. Differences among the simulated scenarios are concentrated at night. Near-surface wind averaged over the flaming front decreases by about 16%, and the corresponding spread rate decreases by about 20%, while daytime differences remain small. Dynamical diagnostics relate these changes to a weaker near-surface pressure-gradient force along the canyon, the signature of a preferentially eroded cold layer over the high terrain. Although a single event cannot establish a climatological response, it can isolate a mechanism under constrained synoptic conditions. We therefore present the result as a testable expectation for other downslope-wind regimes with comparable terrain geometry.

2. Results

2.1 Simulated fire progression and comparison with the observed perimeter

In the Base case, the fire simulation reproduces the orientation and channeling of the observed fire. Compared with the National Interagency Fire Center (NIFC) perimeter at 16:00 local standard time (LST) 8 November, the simulated fire captures the northeast-to-southwest elongation and its leading edge advances down the Feather River canyon under the downslope flow (Fig. 1a; domain configuration in Supplementary Fig. 1). Except along the eastern flank, where the simulated fire extends slightly beyond the observed perimeter, the simulated footprint lies within it. The discrepancy is concentrated at the western, downwind head, where the simulated fire reaches only 61 km² against 285 km² observed at this time. This westward deficit is due to firebrand spotting absent from these runs. Firebrands were igniting spot fires in Paradise within two hours of ignition, and radar revealed long-range spotting advancing the fire by as much as 3 km in 15 min^{3,49}, while the simulated head advances by contiguous surface spread alone (Methods). Terrain and surface roughness likely also contribute to the perimeter bias.

Model terrain is smoothed relative to the real canyon walls, and surface roughness derives from a land-cover map compiled in 2011, seven years before the event. Both act on the speed and direction of the low-level flow that drives the fire. Suppression is likewise absent from the model, but containment during the first day was minimal³, so a fire simulated without suppression should grow larger than observed. Because the same deficit recurs in every scenario at every output time, it largely cancels in scenario-to-scenario comparisons. Their simulated perimeters are similar through midday and separate primarily after dark (Supplementary Fig. 2), so differences reflect the warming response instead of the shared bias. The corridor the fire follows contains a mixture of grass, shrub and timber fuels (Fig. 1c), whose roles are quantified below. By the last output time common to all runs, 03:35 LST 9 November, the perimeters differ visibly only at the downwind head (Fig. 1b), where the warming response is expressed.




Fig. 1 | Study area, surface fuels, and observed and simulated Camp Fire perimeters. a Terrain elevation of the inner domain d02 (111.1 m atmospheric grid, 27.8 m fire subgrid), shown with hillshaded hypsometric shading and 200 m contours (thin white). The yellow star marks the ignition point (39.8134° N, 121.4347° W). Black and purple lines show the observed NIFC perimeter and simulated Base perimeter, respectively, at 16:00 LST on 8 November 2018. Corresponding burned areas are printed on the panel. The dashed rectangle identifies the area enlarged in b, and the black scale bar denotes 10 km. b Simulated perimeters near the downwind fire edge at 03:35 LST on 9 November 2018, the final output shared by all scenarios. Scenario colours are Base purple, +1.5 K blue, +3 K green, +4.5 K orange, and +6 K red throughout the figures. Pale terrain shading, 100 m (white) and 500 m (grey) contours, and road centerlines (grey) are included. The displayed area adjoins the southern boundary of d02, and the black scale bar denotes 4 km. c Surface fuels over the same extent as a, grouped by Scott and Burgan family: NB, non-burnable; GR, grass; GS, grass-shrub; SH, shrub; TU, timber-understorey; and TL, timber litter. The black line shows the Base perimeter at 03:35 LST on 9 November, and the star marks the ignition point.

2.2 Diurnal changes in spread rate, wind speed, and burned area

During the daytime, ensemble-mean spread rate and wind speed averaged over the flaming-front (fastest-spreading quartile of active burning cells) show no consistent ordering among warming levels, and their ensemble ranges overlap (Fig. 2a, b; Methods 4.3). Differences emerge near dusk, during the downslope-wind window from 21:00 to 03:00 LST (shaded in Fig. 2) and become largest after midnight. Burned area follows the same temporal progression and its

diagnostics are evaluated for the unperturbed control run of each scenario. At 16:00 LST, the five scenarios span 58 to 68 km² and are not ordered by warming level. By 22:00 LST, areas decrease monotonically from 169.5 km² at Base to 144.2 km² at +6 K. By 03:35 LST, the corresponding areas are 303 against 202 km², a difference of 101 km² (Fig. 1b, Supplementary Fig. 2). Approximately three quarters of that difference develops after 22:00 LST. Although the Base control enters the lateral relaxation zone at 03:25 LST, the last output clear of the zone, 03:20 LST, already shows a 98.9 km² difference, approximately one third of the Base burned area (Methods 4.3).




Fig. 2 | Warming dries the air but weaker winds cause slower rates of fire spread. a–c Rate of spread (a), 10 m wind speed (b), and vapour pressure deficit (c), averaged over cells at or above the 75th percentile of spread rate among all actively spreading cells (Methods 4.3). Solid lines show the mean of

six realizations per scenario. Each ensemble includes one unperturbed control and five members initialized with 0.2 K potential-temperature and 0.1 m s⁻¹ wind-component perturbations. Shading spans the minimum and maximum across members. Grey shading marks the downslope-wind window (21:00–03:00 LST). The first hour after ignition is excluded because the initial quartile sample contains few cells. Times are LST, and the date changes from 8 to 9 November 2018 as labeled.

Across the downslope-wind window, the ensemble-mean spread rate decreases monotonically from 0.52 m s⁻¹ at Base to 0.42 m s⁻¹ at +6 K, a reduction of approximately 20%. Mean sampled wind speed decreases from 4.12 to 3.46 m s⁻¹, approximately 16% (Fig. 2a, b). Both quantities are averaged first over the flaming front of all actively spreading cells and then over six runs in each scenario. They are therefore scenario-comparable averages, with values below instantaneous maxima at the fire perimeter. The mean inter-member standard deviation of sampled wind speed is 0.024 m s⁻¹, compared with a scenario range of 0.66 m s⁻¹. Calculated separately at each output time, the signal-to-noise ratio averages 28 across the window and remains above 15 throughout (Supplementary Fig. 3).

Reduced nighttime spread is also evident in burned-area growth. Area added during the downslope-wind window decreases from 133 km² at Base to 62 km² at +6 K, a 54% reduction corresponding to a decrease from 22 to 10 km² h⁻¹. Peak growth rate declines from 47 to 14 km² h⁻¹. Because the reduced area exceeds the 20% decrease in quartile-averaged spread rate, changes in the length or geometry of the actively spreading perimeter may also contribute. Maximum distance from the ignition point behaves similarly. Control runs approach approximately 18 km during the evening, then resume propagation at different times during the nighttime wind event. Final distances range from 29.2 km at Base to 22.8 km at +6 K, with the Base estimate affected by boundary contact after 03:25 LST (Supplementary Fig. 4; Methods 4.3). These distances describe perimeter geometry and are distinct from the fastest-quartile sampling used for the main spread-rate analysis.

Across warming scenarios, the pattern of the burned area shifts only modestly with warming, chiefly a larger timber-understory share in the smaller +6 K footprint (23.6% against 18.6% at Base; Supplementary Fig. 5a). The median spread rate over the fastest-spreading cells declines with warming in every burnable family, from shrub, always the fastest (530 down to 414 mm s⁻¹, a 22% drop), to timber litter, always the slowest (177 down to 153 mm s⁻¹, 13%; Supplementary Fig. 5b). Despite the increasingly favorable fire environment, the fire slows. Within the downslope-wind window, sampled VPD is 33% higher (1.48 to 1.97 kPa), while specific humidity over the same regions barely changes (Fig. 2c, Supplementary Fig. 6a, b). All 30 simulations prescribed dead fuel moisture at 5%, an already critically dry value, so the fuel state remains fixed while the atmosphere and winds respond to warming (Methods 4.2). The following analysis examines the pressure-gradient force associated with the wind decrease. Separate prognostic-moisture experiments then quantify the additional influence of fuel drying.

2.3 Elevated cold-air structure and downslope pressure-gradient forcing

Across the fire corridor, nocturnal winds weaken beyond the burned area. At 22:00 LST in the unperturbed Base control, near-surface jets follow the same canyon axis as the simulated fire (Fig. 3a). Relative to Base, corridor-mean wind speed at that time decreases by 0.83 m s⁻¹ at +3 K and 2.29 m s⁻¹ at +6 K, with the largest changes along the canyon and its lee slopes (Fig. 3b, c). Over the full downslope-wind window, the 10 m wind projected onto the corridor axis decreases from 9.89 m s⁻¹ at Base to 7.55 m s⁻¹ at +6 K, approximately 24%. Wind speed

averaged over the flaming-front decreases by 16% during the same period. These percentages differ because one diagnostic uses a fixed corridor and the other follows the changing distribution of rapidly spreading cells, but both show the same warming order and timing.

The flow weakens because of the pressure-gradient force (PGF) driving it. Evaluated from the horizontal pressure-gradient term on the lowest model levels and projected onto the ridge-to-valley axis, the force along the corridor decreases from 6.5 to $3.1 \times 10^{-3} \text{ m s}^{-2}$ across the downslope-wind window (Methods 4.4). This 52% reduction is monotonic and approximately linear with warming, corresponding to $0.57 \times 10^{-3} \text{ m s}^{-2} \text{ K}^{-1}$ (Fig. 3d). Inter-member standard deviations range from 0.02 to $0.07 \times 10^{-3} \text{ m s}^{-2}$, compared with a scenario range of $3.4 \times 10^{-3} \text{ m s}^{-2}$. The same ordered decrease occurs in an eastern control box that remains unburned in every experiment, indicating that the pressure response also occurs outside the fire-affected area. Scenario differences in forcing emerge by mid-afternoon, several hours before consistent differences appear in wind speed.




Fig. 3 | Warming erodes the elevated cold layer and weakens the force driving the downslope flow.

a 10 m wind speed in the unperturbed Base control at 22:00 LST 8 November 2018, with the ridge and valley averaging boxes (solid rectangles), the corridor box (dashed) and the corridor axis (arrow; Methods). b, c 10 m wind speed difference from the Base control at the same time for the +3 K (b) and +6 K (c) controls. Blue denotes weaker wind under warming, and the corridor-mean difference is printed in each panel. Panels a–c are single-realization snapshots at one output time; d–h are six-member ensemble statistics. d Along-corridor pressure-gradient force, evaluated on the lowest three model levels over the corridor box and projected on the ridge-to-valley axis (Methods). e Mass-weighted layer-mean virtual potential temperature difference, valley minus ridge. Each box is integrated from its own surface to the 1600 m common level, so the two columns differ in depth. In d and e, solid lines are six-member ensemble means, shading spans the member range, dotted horizontal lines mark downslope-window means, the vertical dashed line marks the time mapped in a–c, and grey shading marks the downslope-wind window (21:00–03:00 LST). f, g Window-mean force against the layer contrast (f) and window-mean along-corridor 10 m wind against the force (g). Large dots are scenario ensemble means, faint dots the 30 individual runs, and the dashed line connects the scenario means in order of warming. Each panel prints the correlation across the five scenario means and across all 30 runs. Both are descriptive, not inferential. h Hydrostatic budget of the window-mean force: bars decompose the directly evaluated force (diamonds) into the pressure-gradient force at the 1600 m common level scaled to the surface layer (top), the katabatic (slope) and cold-pool (pool) terms, all evaluated over the corridor box. The bars sum to the diamonds to within 0.1%. Circles are the independent reconstruction from the virtual-potential-temperature field alone, which matches the direct force within about 4% (Methods 4.4).

Warming removes the buoyancy contrast that the cold layer supplies. In the Base case, mass-weighted mean virtual potential temperature of the layer between each box's surface and the 1600 m common level is 3.2 K lower over the ridge than over the valley, and the contrast decreases monotonically to 1.2 K at +6 K (Fig. 3e). It decreases because warming erodes the elevated cold layer faster than the valley cold pool: the column-integrated cold-air mass anomaly of the ridge box falls from 13.5 to 9.5 kg m⁻² across the scenarios, about 30% and strictly ordered, while the valley value shows no systematic change. Warming preferentially thins the cold layer over the high terrain. Across the scenarios means, the force along the corridor follows the diminishing contrast, tracking the layer-mean difference almost exactly (Fig. 3f), and the wind along the corridor follows the force in turn (Fig. 3g). The reported correlations summarize those relationships.

The budget is closed through a hydrostatic reconstruction. Rebuilding the force from the virtual-potential-temperature field reproduces the direct evaluation within about 4% in every scenario. The reconstructed decline of 50 to 51% is insensitive to the reference level, so the weakening is driven by the thermal structure below the common level, with little contribution from pressure changes aloft (Fig. 3h). Within the budget at 1600 m, the katabatic term, defined as the product of local buoyancy and terrain slope, decreases by 28% (5.7 to $4.1 \times 10^{-3} \text{ m s}^{-2}$). However, the cold-pool term declines from $1.4 \times 10^{-3} \text{ m s}^{-2}$ to $-0.5 \times 10^{-3} \text{ m s}^{-2}$ and, at that reference, changes from supporting the flow to opposing it. How the decline divides between the two terms depends on the reference level. This is consistent with a cascading effect that begins in the cold layer.

Warming thins the cold air layer over the high terrain, relaxes the PGF, the weaker PGF slows the downslope flow, and the slower flow yields a slower rate of fire spread.

2.4 Scenario-dependent importance of spread-rate predictors

To compare the relative predictive contributions of correlated environmental variables, we fit a separate gradient-boosted model to the 5-min flaming-front records from each warming scenario and interpret its predictions with SHAP. The predictor set, its correlation structure and the exclusion of leakage-prone variables are documented in Methods 4.5 and Supplementary Fig. 7. The models quantify how the mean absolute SHAP magnitude of each candidate predictor changes across warming levels. The robustness of the ranking under resampling of records and predictor pools is characterized with a double bootstrap in Supplementary Fig. 8. Skill is high and insensitive to the validation design, with R^2 of 0.98 to 0.99 for random, contiguous-block and leave-one-member-out cross-validation alike (Supplementary Fig. 9). Leave-one-member-out values (0.981 to 0.993), from the only design that never trains and tests on the same simulation, match the rest, indicating that the surrogate generalizes across the closely related ensemble realizations. Because these realizations share event timing and boundary forcing, the validation does not establish independence across different fire events.

Wind speed ranks first in every scenario, with the largest mean absolute SHAP value in each model (19.2 to 47.5 mm s⁻¹). Higher wind speeds contribute positively to the predicted spread rate, and lower wind speeds contribute negatively (Fig. 4). The relative importance of other predictors changes with warming. VPD rises from rank 4 at Base at Base (5.5 mm s⁻¹) to rank 2 at +4.5 K (18.6 mm s⁻¹) and +6 K (14.2 mm s⁻¹). Because fuel moisture is prescribed in these simulations, this association does not measure a direct fuel-drying effect. Terrain slope moves the other way, from rank 2 at Base (9.3 mm s⁻¹) to rank 5 at +6 K (2.5 mm s⁻¹), and the wind-slope synergy term rises to rank 3 at +6 K (10.0 mm s⁻¹). Warming thus reshapes the mean absolute SHAP magnitude beneath a stable leader: the wind term keeps the largest share in every scenario while the second tier reorders. The functional forms behind these magnitudes are shown in Supplementary Fig. 10. Since each scenario has its own model, curve shapes can be compared but vertical offsets cannot.




Fig. 4 | Wind speed carries the largest attribution weight in every scenario, with VPD rising as the climate warms. a–e SHAP summaries for one gradient-boosted model per scenario, fitted to the merged six realizations with target and predictors averaged over the regions of fastest spread (Methods). Each dot is one 5-min quartile record. Its horizontal position is the SHAP value of that predictor in mm s^{-1} , the signed contribution to the predicted spread rate relative to the model mean. Dots stack vertically by local density, and colour encodes the predictor value from low (blue) to high (red), scaled between its 5th and 95th percentiles. Rows share one order, set by the mean absolute SHAP value across scenarios. The axis is clipped at $\pm 134 \text{ mm s}^{-1}$, 6% beyond the largest of the five per-scenario 99.8th percentiles of absolute SHAP values, leaving 30 of the 45,540 plotted values (0.066%) outside. Hold-out and leave-one-member-out R² are printed in each panel, and the inset table lists mean absolute SHAP values (mm s^{-1}). Surface moisture flux is given in $\text{mg m}^{-2} \text{ s}^{-1}$.

The spatial distribution of rapidly spreading cells is consistent with wind-driven propagation through the canyon. During the downslope-wind window, mean wind-slope synergy within the flaming-front ranges from -0.34 to -0.35 m s^{-1} across the five scenarios, a total range of 0.014 m s^{-1} . Negative values indicate that sampled winds are directed downslope on downwind-facing terrain (Methods 4.4). This pattern is expected when strong winds overcome the usual upslope tendency of buoyancy-driven surface spread and agrees with the observed down-canyon progression. Mean wind-slope alignment remains similar across scenarios even though the simulated fires occupy different terrain as their perimeters diverge.

2.5 Effects of warming and prognostic fuel moisture on burned area

Two additional simulations repeat the unperturbed Base and +6 K controls with prognostic dead fuel moisture enabled. Together with the original controls, they form a 2×2 experiment: A (Base, prescribed), B (Base, prognostic), C (+6 K, prescribed), and D (+6 K, prognostic). Runs A and C are the control members of the warming ensemble (Methods 4.2). With prognostic moisture enabled, mean fuel moisture within the flaming-front falls below the prescribed 5%, averaging 3.1% at Base and 2.4% at +6 K. During the downslope-wind window, these values decline to 2.3% and 1.8%, respectively (Supplementary Fig. 6c). Drying is strongest in the fine fuels. Corridor-mean 1-h moisture decreases by 3.7 percentage points at Base and 2.7 percentage points at +6 K, whereas the 100-h class changes little. The +6 K simulation starts drier and remains drier throughout the event (Supplementary Fig. 6d).

Prognostic fuel moisture-enabled fires meet the domain boundary earlier, which constrains how the four runs can be compared. Run B first places burned cells in the lateral relaxation zone at 23:20 LST 8 November and run D at 00:15 LST 9 November, whereas the prescribed moisture runs reach it only at 03:25 LST 9 November (A) or never (C). Beyond first contact a perimeter is no longer independent of the boundary forcing (Supplementary Fig. 11a). Every burned area and area difference quoted for these experiments is therefore evaluated at a single common timestamp, $t^* = 22:00$ LST 8 November, 80 min before the earliest contact, when all four perimeters are still interior (Fig. 5a, Supplementary Fig. 11b), and never at the runs' respective end times. The moisture-enabled curves in Fig. 5a end at each run's last archived output, 0 to 10 min after first relaxation-zone contact (Methods 4.1), while the fires are still spreading.




Fig. 5 | Prognostic fuel drying enlarges the fire, and warming shrinks it in both configurations. a Cumulative burned area for the four experiments: the Base and +6 K control runs with dead fuel moisture prescribed at 5% (A, C; solid) and the corresponding runs with the prognostic scheme active (B, D;

dashed), purple denotes Base and red +6 K. Triangles with dotted lines mark each run's first contact with the lateral relaxation zone, beyond which the perimeter is no longer independent of the boundary forcing. The grey band marks the boundary-affected period and the vertical black line the common comparison time $t^* = 22:00$ LST 8 November 2018, 80 min before the earliest contact. b–e Pairwise differences in burned extent at t^* on the fire subgrid, over 200 m terrain contours: grey denotes cells burned in both runs of the pair, red cells burned only in the first-named run and blue cells only in the second, for the moisture effect at Base (b, $B - A$), the warming effect with moisture prescribed (c, $C - A$), the warming effect with the scheme active (d, $D - B$) and the moisture effect at +6 K (e, $D - C$). The black bar in b is a 5 km scale. f Burned area of each run at t^* (top) and the pairwise differences with the interaction term ($D - C - (B - A)$) (bottom). In the full 2×2 design, $(D - B) - (C - A)$ is algebraically the same interaction, read as the change in the warming effect when the moisture scheme is active.

At the common comparison time (t^*), activating the moisture scheme adds 122.8 km² (+72%) at Base and 113.1 km² (+78%) at +6 K (Fig. 5), imposing +6 K removes 25.3 km² (–15%) with moisture prescribed and 34.9 km² (–12%) with it active, an interaction of –9.7 km², about 8% of the main moisture effect. Warming removes more absolute area when the scheme is on because the fire it acts on is larger, and the slightly smaller fractional loss is consistent with the additional drying partly offsetting the weaker wind. Therefore, fire slows under warming in both moisture configurations, even though the moisture-enabled pair holds the drier fuel at +6 K.

At +6 K, the burned-area increase associated with prognostic fuel drying (+113.1 km²) is approximately three times the decrease associated with warming within the prognostic-moisture treatment (–34.9 km²). Consequently, the +6 K prognostic run burns 257 km² by t^* , compared with 169.5 km² in the Base prescribed-moisture run. The simulations therefore distinguish the large effect of the moisture treatment on absolute burned area from the smaller warming effect within each treatment. Prognostic moisture was tested only at Base and +6 K, so these results cannot identify whether the relative effects change at intermediate or greater warming.




Fig. 6 | Weakening of the downslope wind under warming overrides the concurrent drying of the fuel, while the fuel-moisture scheme sets the systematic spread-rate offset and the wind sets its fluctuation. Exact decomposition of the simulated rate of spread (ROS) into fuel-state, wind, slope and interaction contributions (Methods), for the two contrasts of the fuel-moisture experiment of Fig. 5: D - B, the effect of +6 K warming with the prognostic scheme active in both runs, and D - C, the effect of the scheme itself at +6 K. All quantities are averaged over each run's own 75th percentile of spread speed and paired by output timestamp (Methods). a, b Time series of the four contributions and their sum, ΔROS , for D - B (a) and D - C (b). Light shading marks the downslope-wind window, narrow vertical bands mark times at which the interaction term exceeds half of $|\Delta ROS|$. The dashed line marks the end of each paired record (23:20 LST in a, 00:25 LST in b; Methods), and the flat-shaded interval beyond it carries no data. c Mean contribution of each term over 21:00–23:20 LST, the part of the downslope-wind window covered by both contrasts ($n = 29$ five-minute outputs per contrast). The legend gives the corresponding mean ΔROS . In c and d bar colour identifies the term, as in a, and fill pattern identifies the contrast: solid bars are D - B, hatched bars D - C. d Contribution of each term to the fluctuation of ΔROS , the covariance share $Cov(term_i, \Delta ROS)/Var(\Delta ROS)$. The four shares sum to 100% and may be negative, a negative share marking a term that varies out of phase with ΔROS and damps. Computed over each contrast's full paired record ($n = 196$ for D - B, 209 for D - C). Error bars in c and d are 95% percentile intervals from a blocked bootstrap (2,000 resamples). Those in c rest on six blocks and are indicative (Methods). $\dagger$ The wind term is not exogenous: the model is coupled, so part of the wind difference at the fastest spread area is the fire's own indraft responding to a spread-rate difference, and it marks the proximate channel through which the difference is expressed, not an independent cause.

Within the flaming-front, weaker winds offset the acceleration associated with drier fuel. We decompose the model's Rothermel spread-rate function into fuel-state, wind, slope, and interaction contributions using coefficients archived on the fire grid. Each contribution is averaged over cells at or above the 75th percentile of spread rate in its own run and paired by

output time. Their sum reproduces the run-to-run difference Δ ROS to machine precision (Fig. 6; Methods 4.6). For warming with prognostic moisture in both runs (D – B), the fuel-state contribution during 21:00–23:20 LST is $+20.9 \text{ mm s}^{-1}$ (95% CI $+15.4$ to $+27.0$). The wind contribution is -33.9 mm s^{-1} (-52.8 to -15.3), with a slope contribution of $+0.4 \text{ mm s}^{-1}$ and an interaction contribution of -5.2 mm s^{-1} , producing a net change of -17.8 mm s^{-1} . Neither the fuel nor the wind interval includes zero, although the six-block bootstrap provides only approximate interval widths (Fig. 6c). The wind term also accounts for 88% (82 to 97) of Δ ROS variability across the complete paired record (Fig. 6a, d), which ends at 23:20 LST and so samples the establishment phase in which the ensemble signal is largest.

For the effect of prognostic moisture at $+6 \text{ K}$ (D – C), the fuel-state contribution accounts for most of the mean spread-rate difference: $+53.4 \text{ mm s}^{-1}$ ($+51.4$ to $+55.4$) of a $+50.4 \text{ mm s}^{-1}$ total. A slope contribution of -6.8 mm s^{-1} and an interaction contribution of -1.1 mm s^{-1} partly offset this increase. The mean wind contribution is $+4.9 \text{ mm s}^{-1}$, with an interval extending nearly to zero ($+0.07$ to $+10.0$). Wind nevertheless accounts for 65% (54 to 90) of Δ ROS variability, and the wind-moisture interaction accounts for another 25% (20 to 38) (Fig. 6b, d). Across the two contrasts, fuel moisture primarily determines the mean spread-rate offset, while wind contributes most of its temporal variability and produces a negative net response to warming. Because fire-atmosphere coupling allows fire-induced inflow to alter the sampled wind, the wind contribution should be interpreted together with the pressure-gradient diagnostics (Fig. 3) and the controlled 2×2 experiment. Each run is evaluated on its own fastest-quartile mask, consistent with Figs. 2 and 4. Models refitted to all four experiments support the same interpretation: wind retains the largest predictive weight, while the relative importance of moisture-related predictors changes (Supplementary Fig. 12).

3. Discussion

For the Camp Fire windstorm, imposed warming affects atmospheric dryness and downslope winds in opposite directions, with the largest fire spread response occurring at night. Daytime spread changes little, whereas nighttime warming preferentially erodes cold air over elevated terrain and reduces the pressure-gradient force driving downslope flow. At $+6 \text{ K}$, that force decreases by approximately half, wind speed averaged over the flaming-front decreases by approximately 16%, spread rate decreases by approximately 20%, and overnight burned-area growth decreases by roughly half. Allowing fuel moisture to respond increases absolute burned area, with the drying-related increase at $+6 \text{ K}$ approximately three times the warming-related decrease within the prognostic-moisture treatment. Nevertheless, warming reduces burned area by 12 to 15% within both moisture treatments. The spread-rate decomposition helps reconcile these comparisons: fuel moisture controls much of the mean difference between moisture treatments, while the wind contribution counteracts drying under warming and explains most of the temporal variability in both contrasts.

The nighttime concentration of the response reflects the changing relationship between pressure-gradient forcing and near-surface wind. Although forcing differs among scenarios by mid-afternoon, daytime winds show no consistent warming order, likely because convective mixing connects the canyon flow to a deeper boundary layer. As mixing weakens near dusk, the near-surface flow settles into balance with the force, and the ordering established hours earlier re-emerges in the wind signal. This signal begins to develop at dusk, strengthens through the night,

and is largest after midnight. Traditionally, these are the hours used for more effective suppression actions (direct and indirect attack). Warming is eroding the thermodynamic side of that nighttime relief, through weaker humidity recovery and more active nocturnal burning^{42,43}, but for downslope windstorm fires the controlling nighttime variable is the wind, and here warming moves it the other way. A weaker force would act to preserve part of the relief that thermodynamic trends are removing. Whether reductions of this magnitude alter tactical operations remains a separate question, since residual winds may still preclude direct attack or make indirect firing operations and line construction unsafe, including by increasing the hazard from falling trees.

This event-scale response is consistent with, and independent of, the climatological evidence that offshore downslope events become less frequent as the cool-season land-sea contrast weakens^{28–30}. Those projections concern how often such windstorms occur. The present experiment holds one synoptic evolution fixed and finds that warming weakens the event itself. It also supplies a mechanistic reason why statistical projections that scale burned area on vapour pressure deficit may overestimate the future of this fire class²⁷. For windstorm-driven fires the wind is not a stationary boundary condition but a warming-sensitive one. The causal content of these statements comes from the experimental design. Pseudo-global warming is an intervention, so the difference between scenarios is the effect of the imposed increment within the model, and shared model biases, including the westward perimeter deficit, largely cancel in that comparison. The machine-learning component should be read accordingly: it does not identify the drivers of spread, which the coupled dynamics already contain, but quantifies how the relative weights of correlated predictors shift across warming levels and locates those shifts in the diurnal cycle and on the terrain, with wind's leading weight intact throughout.

The absence of spotting limits the interpretation of the simulated response. Whereas the model represents contiguous surface spread, the observed Camp Fire reached Paradise through firebrand transport and prolific spot fire ignition³. Warming could influence these processes in opposite directions: weaker nocturnal winds may reduce horizontal ember transport, while drier receiving fuels may ignite more readily. Plume-driven lofting depends on fire intensity rather than the ambient wind and is not constrained by these simulations. Consequently, the number of established spot fires could increase even as the modelled spread rate of flaming-front slows. Our spread-rate response applies to the fastest quartile spread. It does not represent a perimeter-only metric or the total expansion rate of a real fire that includes spotting. Because the omission is common to every scenario, the scenario differences remain interpretable for the head fire pathway, but bounding the net response of this fire type will require coupled simulations that accurately characterize firebrand generation, transport and ignition⁵⁰.

A single event under uniform warming demonstrates the mechanism and the direction of the response in the flaming-front regions but cannot establish its magnitude across events or its net effect on wildfire hazard. These broader outcomes also depend on windstorm frequency and longer-term changes in fuels. Wet years may increase grass growth and fine-fuel accumulation, whereas dry years may limit finer fuel accumulation. Fuel moisture, the vertical arrangement and horizontal continuity of fuel, and management actions such as thinning can further influence fire behaviour. Real warming will also vary across space and height. Vertically structured warming, change in land-sea temperature contrasts, and differential thermal advection could alter

atmospheric stability, the nocturnal cold layer, and the pressure gradient driving downslope winds. This implies repeating the experiment with seasonally and spatially varying deltas drawn from global model projections is the natural next step, and embedding the mechanism in ensembles of events is what would turn it into a projection. The framework also has an operational use. Our perturbation ensemble bounds only the sensitivity to grid-scale initial noise, and operational analyses differ from one another at the mesoscale and beyond. Driving the same coupled configuration with members of the National Blend of Models or other operational as well as experimental/non-operational models would quantify how that larger initial-condition uncertainty propagates into predicted fire spread, and with it the forecast confidence available for decision support focused on public safety such as proactive de-energization of utility infrastructure, wildfire prevention and preparation messaging, and evacuation decisions. Although no public safety power shutoff was implemented in advance of the Camp Fire, the Camp and Marshall fires have prompted continued reassessment of the conditions that warrant proactive de-energization. For the downslope windstorm fires that cause disproportionate harm, the wind that drives them is itself sensitive to warming, and projections that hold it fixed are susceptible to overstating the future hazard from flaming front spread.

4. Methods

4.1 Coupled fire-atmosphere model and configuration

Testing how imposed warming alters a windstorm-driven fire requires a model that resolves terrain-channeled flow and allows the fire to act back on the weather. We use WRF-Fire, the wildland fire module of the Advanced Research WRF (ARW) core, version 4.6.0^{45,46,51}. The fire component advances a surface fire perimeter with the semi-empirical Rothermel rate of spread⁵², evaluated locally from wind, slope and fuel and tracked by a level-set method with hybrid WENO5 upwinding and one reinitialization iteration per step⁵³. Coupling runs both ways: the atmosphere supplies the near-surface winds entering the spread-rate calculation, and the fire returns sensible and latent heat fluxes, injected over an exponentially decaying depth, that perturb the boundary-layer flow. The fire can therefore modify the winds that spread it, an essential capability here, as fire-induced circulations and the ambient cross-barrier flow act on the fire simultaneously. Throughout the study, the diagnostic ROS_FRONT is the Rothermel spread rate. The field represents all grid points where the fire area is between 0 and 1, indicating fire is actively advancing. All quartile-sampled quantities in this study are defined from it through the mask of Section 4.3.

The model was built from the release-v4.6.0 source (commit 0a11865f) for real-data cases with the coupled fire module compiled in. The firebrand-spotting and fire-emissions modules were not activated. The code was compiled on the NSF National Center for Atmospheric Research Derecho system with the Intel oneAPI 2024.2.1 compilers in distributed-memory (Message Passing Interface only) mode at -O3, with aggressive floating-point optimizations disabled so that the executable retains strict IEEE arithmetic, and with netCDF-4 input and output.

All simulations start at 04:00 LST (PST, UTC – 8; daylight saving ended 4 November 2018) on 8 November 2018, end 24 h later at 04:00 LST on 9 November. A run terminates early when its fire front enters the lateral relaxation zone at the domain boundary, which three runs do. A run

terminates early through WRF-Fire's boundary guard, which aborts the integration when the level-set function first becomes negative within `fire_boundary_guard` = 8 fire cells (222 m) of the d02 fire-grid edge. The last archived output of an aborted run is the last 5-min history write before the abort. Three runs terminate this way: the Base control, with a last output at 03:35 LST 9 November, and the two prognostic fuel-moisture runs, with last outputs at 23:20 LST 8 November (Base) and 00:25 LST 9 November (+6 K). The Base termination sets the last output common to the warming ensemble, 03:35 LST (section 4.3). The 2 h 33 min spin-up period allows resolved turbulence to develop in the inner domain, and even the shortest warming-ensemble run covers the nocturnal downslope-wind phase in full. Two one-way nests are used (Supplementary Fig. 1). The outer domain (d01) spans about 360 km by 270 km of northern California and western Nevada at 1 km spacing (361 by 271 cells, $\Delta t = 6$ s). The inner domain (d02) covers about 40 km by 30 km around the Feather River canyon and Paradise at 111.1 m spacing (361 by 271 cells, $\Delta t \approx 0.67$ s). The 9:1 ratio steps directly across the "terra incognita" where neither planetary-boundary-layer (PBL) parameterizations nor large-eddy closures are formally valid^{54,55}. Forty-five terrain-following levels reach a model top of 200 hPa. On d01, vertical mixing uses the YSU PBL scheme⁵⁶ with first-order Smagorinsky horizontal diffusion. On d02, the PBL scheme is off, and the domain runs as a large-eddy simulation with the 1.5-order turbulence kinetic energy (TKE) closure of Deardorff⁵⁷. Advection of moisture, scalars and TKE is positive definite. Vertical motion is damped explicitly and through implicit Rayleigh damping in the uppermost 5 km (coefficient 0.2 s^{-1}) to absorb mountain-wave energy, and the acoustic off-centering parameter is raised from 0.5 in d01 to 0.8 in d02 to suppress acoustic noise over the steep canyon sidewalls. Microphysics uses the WRF single-moment six-class scheme⁵⁸. Radiation uses RRTMG longwave⁵⁹ and Dudhia shortwave⁶⁰, called every minute. The surface layer follows the revised MM5 similarity functions⁶¹, and land surface uses the Noah model with four soil layers⁶² over the 40-category National Land Cover Database. No cumulus parameterization is applied. Inner-domain output is written every 5 min, the cadence of all fire diagnostics.

Initial and lateral boundary conditions come from the ERA5 reanalysis⁶³ at 0.25° resolution, interpolated through the standard WRF Preprocessing System (WPS) to 38 metgrid levels (37 pressure levels plus the surface) and updated 6-hourly. The 1 km outer domain then spins up its own mesoscale structure, and the inner domain regenerates the canyon-scale flow as an LES. We drive the simulations with ERA5 rather than a convection-allowing regional analysis such as High-Resolution Rapid Refresh (HRRR) for two reasons. First, the pseudo-global-warming framework requires a driving dataset whose domain encloses the synoptic-scale circulation being perturbed and whose production system is homogeneous over the full event, conditions that a limited-area, frequently re-versioned analysis does not meet. Second, the warming increments are defined as climate-change signals relative to a global base state, so adding them to a regionally assimilated analysis would mix perturbation and base state from inconsistent sources. The convection-scale detail that HRRR would otherwise supply is instead generated internally by the nested downscaling, which applies identically to the control and perturbed simulations and therefore keeps their difference attributable to the imposed increment alone.

The fire module is active on d02 only, on a fire mesh refined fourfold to 27.8 m. Topography is resolved with Shuttle Radar Topography Mission data⁶⁴, capturing the canyon relief that channels the downslope flow. Winds entering the Rothermel formula are taken at a 60 m

reference height, above the layer most perturbed by the fire's heat release and reduced along a logarithmic profile to the 1 m fire wind height. Fire-to-atmosphere feedback is fully enabled. Surface fuels come from the 2014 LANDFIRE map for the Scott and Burgan⁶⁵ 40 fuel models, the closest release predating the event (the 2020 release contains the Camp Fire burn area itself). Fuel type and load are static. In the warming ensemble the prognostic fuel-moisture scheme is off and dead fuel moisture is prescribed at a constant 5% everywhere, so among the fuel properties nothing responds to the perturbed atmosphere. The two experiments that activate the scheme are described below. The fire is ignited as a near-point source of 100 m radius at 39.8134° N, 121.4347° W, the origin near Pulga identified by the National Institute of Standards and Technology investigation³, at 14:33 UTC (06:33 LST), consistent with the reported ignition time. The patch expands at 0.15 m s⁻¹ for 800 s to seed a smooth initial perimeter.

4.2 Pseudo-global-warming design, ensemble and fuel-moisture experiments

Warming experiments follow the pseudo-global-warming (PGW) approach, in which a historical event is re-simulated under thermodynamically perturbed boundary conditions^{31,34,66}. Uniform increments of +1.5, +3, +4.5 and +6 K are added to the initial and 6-hourly lateral boundary fields of three-dimensional air temperature and skin temperature, including sea surface temperature and all soil levels. Specific humidity, the moisture variable carried by the driving files, is held fixed, so relative humidity (RH) diagnosed from the perturbed temperature decreases with warming and vapour pressure deficit rises. The fixed specific humidity adds heat without changing moisture, whereas preserving RH would inject additional water vapour and could alter latent heating and cloud cover alongside the thermal perturbation. Each scenario therefore replays the synoptic evolution of 8 to 9 November 2018 in a uniformly warmer state, and differences between scenarios are attributable to the imposed increment and the pressure and flow adjustments it induces.

Each scenario comprises one unperturbed control and five perturbed members, giving 5 scenarios × 6 realizations = 30 simulations. Member initial conditions were generated by adding zero-mean Gaussian noise to the d02 initial state only, independently at every grid point, with standard deviations of 0.2 K for potential temperature and 0.1 m s⁻¹ for each horizontal wind component. Each member used its own fixed random seed. The perturbed state was not rebalanced. Moisture, geopotential and surface pressure were left untouched, and no hydrostatic or thermal-wind adjustment followed the increments. Because the noise is added to the inner-domain initial state only, the parent domain is unchanged, so the lateral boundary forcing of the nest is identical across all six realizations. No interior nudging is applied on either domain, so the members are free to diverge, drift or re-equilibrate within the shared lateral forcing. The evolution of the resulting spread in wind speed and direction through the full integration is shown in Supplementary Fig. 3. Because the noise is uncorrelated at the 111 m grid scale, model diffusion damps it within the first minutes of integration, and the ensemble limits the sensitivity of the solution to small-scale unbalanced noise, not the mesoscale or synoptic uncertainty an operational ensemble would sample. By the downslope-wind window, approximately 17 h into the integration, the inter-member standard deviation of the wind sampled over the regions of fastest quartile spread averages 0.024 m s⁻¹ and the circular standard deviation of wind direction over the corridor box 0.13°, against a scenario range of 0.66 m s⁻¹ (Supplementary Fig. 3).

Two additional simulations repeat the unperturbed Base and +6 K controls with the prognostic dead-fuel-moisture scheme active⁶⁷, completing a 2×2 design of climate against moisture treatment (A: Base, prescribed; B: Base, prognostic; C: +6 K, prescribed; D: +6 K, prognostic). This scheme carries 1-h, 10-h and 100-h time-lag classes that equilibrate toward the simulated near-surface temperature and RH, so in these runs, and only in these runs, the fuel state responds to the perturbed atmosphere. The moisture integration begins with the first model step after initialization. Moisture of the fuel (FMC_G) sampled over the regions of fastest quartile spread, and the corridor means of the time-lag classes are reported in Supplementary Fig. 6c, d.

4.3 Fastest-quartile sampling and burned-area diagnostics

Every quartile-sampled quantity in this study, including the target and predictors of the attribution models, is averaged over the same region, defined at each 5-min output as the fire-mesh cells whose instantaneous ROS_FRONT is at or above the 75th percentile of ROS_FRONT over all cells with ROS_FRONT > 0. This 75th percentile is chosen to isolate the regions of fastest quartile spread, where the fire's heat release and the ambient downslope flow interact most strongly and where the warming response in rapid fire spread is expressed. The text refers to these averages as taken over the regions of fastest quartile spread and to the 5-min samples entering the attribution models as quartile records. The first hour after ignition (records before 07:35 LST) is excluded from the plotted windows, because the young fire spans only a few actively spreading cells and the quartile sample is correspondingly noisy. It is not excluded from the model training sample.

This choice reflects the spatial structure of ROS_FRONT. The Base perimeter at 03:35 LST 9 November covers a mixture of fuel families (grass-shrub 27%, shrub 24%, timber litter 19%, timber-understory 19%, grass 10%; Supplementary Fig. 5a) whose median spread rates differ by a factor of three, from 530 mm s⁻¹ in shrub to 177 mm s⁻¹ in timber litter (Supplementary Fig. 5b). Much of the spatial dispersion in ROS_FRONT therefore aligns with fuel boundaries. Fast segments occur where the wind-driven head run through shrub and grass-shrub, and slow segments where the front advances through timber fuels. An average over the whole front mixes these regimes and is weighted down by the slow timber tail. Two alternative samplings test this choice (Supplementary Figs. 13, 14). Averaged over all front cells (LFN = 0), spread rate and wind speed retain the nocturnal, warming-ordered decline at absolute rates several times lower than the quartile averages, because slow segments enter the average (Supplementary Fig. 13). A front-restricted quartile, in which the 75th percentile of ROS_FRONT is recomputed over the front cells alone, does not consistently maintain that ordering. Its scenario curves cross repeatedly within the downslope-wind window (Supplementary Fig. 14). This sample contains roughly a quarter of the front cells, and its averages are correspondingly noisier at each output time. The attribution is insensitive to the sampling. Refitting the per-scenario models of Fig. 4 to the whole-front records ranks wind speed first in every scenario.

For every output time of every run we compute: (1) the burned area (cell count $\times 7.716 \times 10^{-4}$ km²), (2) the maximum distance from the ignition point to any burned cell, computed with a local planar approximation (the head distance, with the advance rate a centred difference over ± 30 min; Supplementary Fig. 4), and (3) the first output at which any burned cell enters the lateral relaxation zone of the nest, spec_bdy_width = 5 atmospheric cells = 20 fire cells (0.56 km), inside which the solution is no longer independent of the boundary forcing. In the warming

ensemble, the Base control first enters the zone at 03:25 LST 9 November, 10 min before the last common output, and no warmed control ever does. No burned cell reaches the domain edge in any run, so no burned area is removed from the count. The Base areas at 03:25 and 03:35 LST are instead exposed to boundary forcing that no other run experiences. Restricting the comparison to outputs clear of the zone changes the quoted deficit by about 2%: at 03:20 LST, the last output at which all five scenarios are clear of the zone, the Base-minus-+6 K deficit is 98.9 km², against 101.0 km² at 03:35 LST. All areas and differences quoted in the paper are computed directly from the burned-cell counts defined above. The hole filling, small-component removal and Gaussian smoothing applied before contouring the perimeters shown in the figures affect the rendering only and enter no quoted number. In the fuel-moisture experiments the runs reach the relaxation zone at different times (B at 23:20 LST 8 November, D at 00:15 LST 9 November, A at 03:25 LST 9 November, C never; Supplementary Fig. 11), so all four are compared at a single common timestamp, $t^* = 22:00$ LST 8 November, 80 min before the earliest contact, and never at their respective end times.

Relaxation-zone contact and run termination are distinct events at distinct depths. Contact is diagnosed when any burned cell lies within 20 fire cells (0.56 km) of the edge, whereas the model aborts at 8 fire cells (222 m; Section 4.1). A run therefore continues while its front crosses the intervening 12 fire cells (333 m), truncated to the 5-min output grid. The crossing takes two outputs for A (03:25 to 03:35 LST) and D (00:15 to 00:25 LST), whose fronts approached the boundary at 2 to 4 fire cells per output, and none for B (23:20 LST), whose front advanced at 6 to 9 fire cells per output and crossed both thresholds within one output interval.

4.4 Pressure-gradient force and derived quantities

The pressure gradient force along the corridor is a term of the model's momentum equation, evaluated on the lowest three model levels, averaged over the corridor box and projected on the ridge-to-valley axis of Fig. 3a. Positive values point down the corridor. Every scenario shares one terrain, so the metric term of the terrain-following gradient is common to all runs to within 0.04% and the scenario differences are carried by the pressure field itself. The window-mean decline of Section 2.3 does not depend on these choices: it ranges from 41% to 59% when the force is evaluated on one to five levels and shifts by a few percentage points under horizontal smoothing of the pressure field. The ridge (121.43°–121.36° W, 39.800°–39.860° N) and valley (121.62°–121.55° W, 39.615°–39.675° N) boxes bracket the corridor (121.60°–121.42° W, 39.700°–39.820° N) that the fire descends. A fourth box east of the fire, burned in no run of any experiment, supplies the unburned control of Section 2.3. The maps of Fig. 3a–c show the unperturbed control of each scenario at a single output time. All time series and window statistics pool the six realizations.

All thermal diagnostics share one reference height, a common level at 1600 m above sea level, the lowest level that clears the highest terrain in the boxes (1528 m). Each column is integrated from its own surface to this level, so the ridge and valley columns differ in depth (about 594 and 1372 m at the box mean). The layer-mean virtual potential temperature of Fig. 3e is the mass-weighted mean over that layer, $\int \text{pref } \theta_v dz / \int \text{pref } dz$, and the contrast is the valley mean minus the ridge mean. The cold-air content of each box is the column-integrated mass anomaly $\Lambda = \int \text{pref } (\theta_{v,\text{ref}} - \theta_v) / \theta_{v,\text{ref}} dz$, in kg m⁻², where $\theta_{v,\text{ref}}$ is a scenario reference value and $\text{pref}(z)$ a

fixed reference density profile, box-averaged after integration. For Fig. 3h the force is rebuilt hydrostatically over the corridor box on the same lowest three levels and separated into three terms: the pressure-gradient force at the common level scaled to the surface layer by the density ratio (top), a katabatic term given by the pointwise product of local buoyancy and terrain slope (slope), and a cold-pool term given by the along-corridor gradient of the column-integrated cold-air anomaly (pool). The three terms sum to the directly evaluated force to within 0.1%, and the reconstruction from the θ_v field alone matches it within about 4% in every scenario. The split between the katabatic and cold-pool terms is reference-dependent by construction, and the layer diagnostics sharpen as the common level approaches the terrain: raising it from 1600 to 2200 m leaves the directly evaluated decline unchanged and the θ_v -only reconstruction at 50 to 51%, while the two-term split floats around a pinned sum and the ridge-valley contrasts dilute, because equal added depth is a larger fraction of the shallow ridge column than of the deep valley column. The 1600 m level is therefore the most discriminating admissible choice, and every level-conditional quantity in the text names it. The wind-slope synergy term is computed on the fire mesh as the product of the terrain gradient magnitude and the component of the wind entering the spread-rate calculation along the local upslope direction, so it carries the units of wind speed (m s^{-1}). The gradient and the upslope direction are taken from the model's own fire-mesh slope components rather than recomputed from the elevation field, and at each output the product is averaged without weighting over the sampled cells. The fire cells share one size, so the mean is also an area mean. Positive values mark flow directed up the local slope, which reinforces terrain-driven spread. Negative values mark downslope flow opposing it, the regime that prevails here through the downslope-wind window. Vapour pressure deficit is $e_s(T) \cdot (1 - \text{RH}/100)$ with the standard Magnus form of e_s . Wind-direction spread is the circular standard deviation in the Mardia definition⁶⁸. The signal-to-noise ratio is the range of the five scenario ensemble means divided by the mean inter-member standard deviation, computed at each output time. Window values quoted in the text are averages of this ratio over the window, not the ratio of the window-mean range to the window-mean spread.

4.5 Machine-learning attribution

To compare how the candidate predictors are weighted under each warming level, we fit gradient-boosted regression trees XGBoost⁴⁷ to the 5-min quartile records and interpret them with SHAP⁶⁹, computed as exact tree Shapley values⁷⁰. Here, the target is ROS_FRONT averaged over the regions of fastest quartile spread, and the runs entering a given model are first aligned on their common time index. Because ROS_FRONT is itself a diagnostic of the Rothermel parameterization, computed by the model from wind, slope, fuel properties and fuel moisture, the fitted surrogate redescribes a known functional relationship as sampled along the simulated fire. Its SHAP values are accordingly read as contributions to the surrogate's prediction: they quantify how predictive weight is distributed among correlated inputs within each scenario and how that distribution shifts with warming, in space and through the diurnal cycle, and they neither identify drivers of spread nor carry causal content of their own. Causal interpretation rests on the controlled PGW design and the 2×2 moisture experiments, not on the attribution.

Predictors include 10 m wind speed, 2 m temperature, VPD, the wind-slope synergy term, terrain slope and surface moisture flux, with dead fuel moisture added in the experiment models. Six

variables present in the quartile records are excluded as target leakage, being consequences of the spread rate. 2 m RH and elevation are withheld from the reported models but retained in the correlation matrix and the stability-selection pool so a reader can see what a wider pool would have done (Supplementary Figs. 7, 8). The estimator uses 400 trees of depth 4, learning rate 0.06, subsampling of 0.8 over rows and 0.9 over columns, L2 regularization of 1.0 and a fixed seed, with no per-model tuning. One model is fitted per scenario on its merged six realizations (Fig. 4), one per moisture experiment on the window common to all four runs and with an identical predictor set (Supplementary Fig. 12), and two pooled models, over the 30 warming runs and over the two prognostic-moisture runs, supply the stability and correlation diagnostics. Skill is assessed four ways: a single 85/15 hold-out split, random and contiguous-block five-fold cross-validation, and leave-one-member-out validation, in which the model is retrained six times with all records of one realization withheld. The last is the only design that never trains and tests on the same simulation and is the value to trust (Supplementary Fig. 9). Reported SHAP values are scaled to mm s^{-1} , the units of the quantity being explained. Importance is the mean absolute SHAP value, ranks are computed within a model, and where two models are compared, they are explained over a shared background sample. SHAP values are never differenced across models fitted to different fires. For the same reason, no surrogate is fitted to differences between paired runs: at any output time the paired fires occupy different terrain, so differenced quartile averages compare unlike regions. The Rothermel response is nonlinear, so the difference of the outputs is not a function of the difference of the inputs. Once the fires diverge in size and position, instantaneous differences mix intervention's effect with accumulated geometric divergence. Pairwise effects are therefore quantified where the intervention defines them, in the burned extent of the 2×2 design (Fig. 5), and at the flaming-front by exact expansion of the model's own spread-rate function, which involves no surrogate and no truncation (Fig. 6; Section 4.6). Ranking stability is probed with 240 rounds of a double bootstrap: each round resamples records within each run using a moving-block bootstrap of 12 consecutive records (1 h) to the run's original length, offers a random 80% of the candidate predictors drawn without replacement (six of eight for the pooled warming model, seven of nine for the pooled moisture model), and refits a smaller model (150 trees, depth 4, learning rate 0.08). The reported statistic is the probability that a predictor ranks in the top three given that it was offered, with chance levels of 0.50 and 0.43 respectively. Blocks rather than individual records are resampled because consecutive quartile records are strongly autocorrelated, and an independent-row bootstrap would overstate the effective sample size. SHAP dependence curves (Supplementary Fig. 10) show the functional form each surrogate has learned and are not measured dose-response relationships.

4.6 Exact decomposition of the spread-rate difference

The budget of Fig. 6 expands the model's own spread-rate function rather than fitting any statistical model to differences. WRF-Fire evaluates the Rothermel rate as a base rate times an enhancement factor, $\text{ROS} = A \cdot \Phi$ with $\Phi = 1 + \varphi_W \cos \theta_W + \varphi_S \cos \theta_S$, and archives the coefficients on the fire grid. The run-to-run difference then expands without truncation as $\Delta \text{ROS} = \Delta A \cdot \Phi_b + A_b \cdot \Delta \varphi_W + A_b \cdot \Delta \varphi_S + \Delta A \cdot \Delta \Phi$, giving a fuel-state term (which carries dead fuel moisture through A), a wind term, a slope term and an interaction term. The closure of the sum is $5 \times 10^{-14} \text{ mm s}^{-1}$. The slope term is non-zero despite identical terrain because φ_S depends on the slope resolved along the direction of spread, and the two fronts of a contrast lie on different ground and advance along different normals.

All quantities are averaged over each run's own 75th percentile of spread speed, with the two runs of a contrast paired by output timestamp. The same spread-rate percentile is applied independently to each run, so the sampled cells can differ. For D – C the two samples give opposite signs of Δ ROS and are not interchangeable. This budget requires a paired output time in both runs of a contrast, so each contrast ends at the earlier of the two runs' last archived fire output, 23:20 LST 8 November for D – B (limited by run B) and 00:25 LST 9 November for D – C (limited by run D), and the attribution covers the establishment phase of the downslope wind rather than the full nocturnal period. A few paired records therefore lie at or beyond a run's own relaxation-zone contact: one of 196 for D – B (23:20 LST, the contact of run B) and three of 209 for D – C (00:15 to 00:25 LST, the contact of run D). Dropping these records changes the variance shares of Fig. 6d by less than 0.1 percentage point in both contrasts. It leaves the D – C means of Fig. 6c unchanged, because none of the three falls in the 21:00–23:20 LST window. For D – B it moves the window-mean fuel-state, wind and net contributions from +20.9, –33.9 and –17.8 mm s^{–1} to +21.2, –35.3 and –19.1 mm s^{–1}. The full-record values are quoted throughout. They understate the wind and net responses relative to the boundary-clean values.

Confidence intervals are 95% percentile intervals from a blocked bootstrap (2,000 resamples) of contiguous time blocks drawn with replacement to the original record length. Blocking is in time only, because runs B, C and D are single realizations, so the intervals carry no initial-condition uncertainty. In Fig. 6d the block length is 3 h (36 outputs), giving six blocks per contrast. In Fig. 6c it is 25 min (five outputs), giving six blocks from the 29 available samples, because a 3 h block would leave a single block. Intervals in Fig. 6c therefore rest on six blocks and should be read as indicative rather than precise, a qualification that applies with most force to the wide D – B wind interval. The narrow interval on the D – C fuel term reflects a quasi-constant plateau in that term over the window. The propagation normal used within the model's Runge-Kutta level-set update is not archived, so the spread rate is reconstructed from the archived coefficients and the end-of-step level set. Over front cells the reconstruction reproduces the model's ROS_FRONT with a front-mean correlation of 0.96 to 0.97 and a –5 to –7% bias, a level offset that largely cancels in a run-to-run difference.

References

1. Brewer, M. J. & Clements, C. B. The 2018 Camp Fire: Meteorological Analysis Using In Situ Observations and Numerical Simulations. *Atmosphere* **11**, 47 (2020).
2. Mass, C. F. & Ovens, D. The Synoptic and Mesoscale Evolution Accompanying the 2018 Camp Fire of Northern California. *Bull. Am. Meteorol. Soc.* **102**, E168–E192 (2021).
3. Maranghides, A. et al. A Case Study of the Camp Fire: Fire Progression Timeline, Appendix C. Community WUI Fire Hazard Evaluation Framework. (US Department of Commerce, National Institute of Standards and Technology, 2021).

- 808 4. Lareau, N. P., Donohoe, A., Roberts, M. & Ebrahimian, H. Tracking Wildfires With
809 Weather Radars. *J. Geophys. Res. Atmos.* **127**, e2021JD036158 (2022).
- 810 5. California Department of Forestry and Fire Protection (CAL FIRE). Camp Fire incident
811 information. <https://www.fire.ca.gov/incidents/2018/11/8/camp-fire/> (2019).
- 812 6. Koo, E., Pagni, P. J., Weise, D. R. & Woycheese, J. P. Firebrands and spotting ignition in
813 large-scale fires. *Int. J. Wildland Fire* **19**, 818–843 (2010).
- 814 7. Caton, S. E., Hakes, R. S. P., Gorham, D. J., Zhou, A. & Gollner, M. J. Review of Pathways
815 for Building Fire Spread in the Wildland Urban Interface Part I: Exposure Conditions. *Fire*
816 *Technol.* **53**, 429–473 (2017).
- 817 8. Knapp, E. E., Valachovic, Y. S., Quarles, S. L. & Johnson, N. G. Housing arrangement and
818 vegetation factors associated with single-family home survival in the 2018 Camp Fire,
819 California. *Fire Ecol.* **17**, 25 (2021).
- 820 9. Abatzoglou, J. T., Hatchett, B. J., Fox-Hughes, P., Gershunov, A. & Nauslar, N. J. Global
821 climatology of synoptically-forced downslope winds. *Int. J. Climatol.* **41**, 31–50 (2021).
- 822 10. Abatzoglou, J. T. *et al.* Downslope Wind-Driven Fires in the Western United States. *Earths*
823 *Future* **11**, e2022EF003471 (2023).
- 824 11. Westerling, A. L., Cayan, D. R., Brown, T. J., Hall, B. L. & Riddle, L. G. Climate, Santa
825 Ana Winds and autumn wildfires in southern California. *Eos Trans. Am. Geophys. Union*
826 **85**, 289–296 (2004).
- 827 12. Billmire, M., French, N. H. F., Loboda, T., Owen, R. C. & Tyner, M. Santa Ana winds and
828 predictors of wildfire progression in southern California. *Int. J. Wildland Fire* **23**, 1119–
829 1129 (2014).

- 830 13. Jin, Y. *et al.* Contrasting controls on wildland fires in Southern California during periods
831 with and without Santa Ana winds. *J. Geophys. Res. Biogeosciences* **119**, 432–450 (2014).
- 832 14. Cayan, D. R. *et al.* Autumn precipitation: the competition with Santa Ana winds in
833 determining fire outcomes in southern California. *Int. J. Wildland Fire* **31**, 1056–1067
834 (2022).
- 835 15. Guirguis, K. *et al.* Compound atmospheric drivers of the catastrophic 2025 Los Angeles
836 urban firestorm. *npj Nat. Hazards* **2**, 103 (2025).
- 837 16. Sharples, J. J., Mills, G. A., McRae, R. H. D. & Weber, R. O. Foehn-Like Winds and
838 Elevated Fire Danger Conditions in Southeastern Australia. *J. Appl. Meteorol. Climatol.* **49**,
839 1067–1095 (2010).
- 840 17. Tomašević, I. Č., Cheung, K. K. W., Vučetić, V. & Fox-Hughes, P. Comparison of Wildfire
841 Meteorology and Climate at the Adriatic Coast and Southeast Australia. *Atmosphere* **13**,
842 755 (2022).
- 843 18. Potter, B. E. Atmospheric interactions with wildland fire behaviour – I. Basic surface
844 interactions, vertical profiles and synoptic structures. *Int. J. Wildland Fire* **21**, 779–801
845 (2012).
- 846 19. Potter, B. E. Atmospheric interactions with wildland fire behaviour – II. Plume and vortex
847 dynamics. *Int. J. Wildland Fire* **21**, 802–817 (2012).
- 848 20. Keeley, J. E. & Syphard, A. D. Twenty-first century California, USA, wildfires: fuel-
849 dominated vs. wind-dominated fires. *Fire Ecol.* **15**, 24 (2019).
- 850 21. Westerling, A. L., Hidalgo, H. G., Cayan, D. R. & Swetnam, T. W. Warming and Earlier
851 Spring Increase Western U.S. Forest Wildfire Activity. *Science* **313**, 940–943 (2006).

- 852 22. Jolly, W. M. *et al.* Climate-induced variations in global wildfire danger from 1979 to 2013.
853 *Nat. Commun.* **6**, 7537 (2015).
- 854 23. Abatzoglou, J. T. & Williams, A. P. Impact of anthropogenic climate change on wildfire
855 across western US forests. *Proc. Natl. Acad. Sci.* **113**, 11770–11775 (2016).
- 856 24. Williams, A. P. *et al.* Observed Impacts of Anthropogenic Climate Change on Wildfire in
857 California. *Earths Future* **7**, 892–910 (2019).
- 858 25. Goss, M. *et al.* Climate change is increasing the likelihood of extreme autumn wildfire
859 conditions across California. *Environ. Res. Lett.* **15**, 094016 (2020).
- 860 26. Touma, D., Stevenson, S., Lehner, F. & Coats, S. Human-driven greenhouse gas and
861 aerosol emissions cause distinct regional impacts on extreme fire weather. *Nat. Commun.*
862 **12**, 212 (2021).
- 863 27. Cheng, Y., McColl, K. A., Mickley, L. J. & Feng, X. Large Overestimation of Projected
864 Western U.S. Wildfire Burned Forest Area With Warming. *AGU Adv.* **7**, e2026AV002350
865 (2026).
- 866 28. Guzman-Morales, J. & Gershunov, A. Climate Change Suppresses Santa Ana Winds of
867 Southern California and Sharpens Their Seasonality. *Geophys. Res. Lett.* **46**, 2772–2780
868 (2019).
- 869 29. Hawkins, L. R., Abatzoglou, J. T., Li, S. & Rupp, D. E. Anthropogenic Influence on Recent
870 Severe Autumn Fire Weather in the West Coast of the United States. *Geophys. Res. Lett.*
871 **49**, e2021GL095496 (2022).
- 872 30. You, Y. & Delworth, T. L. Projected Response of Santa Ana Winds Over Southern
873 California to Global Warming by a High-Resolution Climate Model. *Geophys. Res. Lett.*
874 **52**, e2025GL114994 (2025).

- 875 31. Schär, C., Frei, C., Lüthi, D. & Davies, H. C. Surrogate climate-change scenarios for
876 regional climate models. *Geophys. Res. Lett.* **23**, 669–672 (1996).
- 877 32. Shepherd, T. G. A Common Framework for Approaches to Extreme Event Attribution.
878 *Curr. Clim. Change Rep.* **2**, 28–38 (2016).
- 879 33. Trapp, R. J., Woods, M. J., Lasher-Trapp, S. G. & Grover, M. A. Alternative
880 Implementations of the “Pseudo-Global-Warming” Methodology for Event-Based
881 Simulations. *J. Geophys. Res. Atmos.* **126**, e2021JD035017 (2021).
- 882 34. Brogli, R., Heim, C., Mensch, J., Sørland, S. L. & Schär, C. The pseudo-global-warming
883 (PGW) approach: methodology, software package PGW4ERA5 v1.1, validation, and
884 sensitivity analyses. *Geosci. Model Dev.* **16**, 907–926 (2023).
- 885 35. Geerts, B., Miao, Q. & Demko, J. C. Pressure Perturbations and Upslope Flow over a Heated,
886 Isolated Mountain. *Mon. Weather Rev.* **136**, 4272–4288 (2008).
- 887 36. Smith, R. K. & Spengler, T. The dynamics of heat lows over elevated terrain. *Q. J. R.*
888 *Meteorol. Soc.* **137**, 250–263 (2011).
- 889 37. Holton, J. R. & Hakim, G. J. An Introduction to Dynamic Meteorology. vol. 88 (Academic
890 Press, 2013).
- 891 38. Pauley, P. M. An Example of Uncertainty in Sea Level Pressure Reduction. *Weather*
892 *Forecast.* **13**, 833–850 (1998).
- 893 39. Abel, M. R., Hall, A. & Kim, J. Human-induced changes in wind, temperature and relative
894 humidity during Santa Ana events. *Clim. Change* **109**, 119–132 (2011).
- 895 40. Guzman-Morales, J., Gershunov, A., Theiss, J., Li, H. & Cayan, D. Santa Ana Winds of
896 Southern California: Their climatology, extremes, and behavior spanning six and a half
897 decades. *Geophys. Res. Lett.* **43**, 2827–2834 (2016).

- 898 41. Bagley, R. B. & Clements, C. B. Extreme Fire Weather Associated with Nocturnal Drying
899 in Elevated Coastal Terrain of California. *Mon. Weather Rev.* **149**, 2497–2511 (2021).
- 900 42. Balch, J. K. *et al.* Warming weakens the night-time barrier to global fire. *Nature* **602**, 442–
901 448 (2022).
- 902 43. Chiodi, A. M., Potter, B. E., Larkin, N. K. & Harrison, D. E. Multifactor Change in
903 Western U.S. Nighttime Fire Weather. *J. Clim.* **38**, 2305–2315 (2025).
- 904 44. Abatzoglou, J. T., Smith, C. M., Swain, D. L., Ptak, T. & Kolden, C. A. Population
905 exposure to pre-emptive de-energization aimed at averting wildfires in Northern California.
906 *Environ. Res. Lett.* **15**, 094046 (2020).
- 907 45. Mandel, J., Beezley, J. D. & Kochanski, A. K. Coupled atmosphere-wildland fire modeling
908 with WRF 3.3 and SFIRE 2011. *Geosci. Model Dev.* **4**, 591–610 (2011).
- 909 46. Coen, J. L. *et al.* WRF-Fire: Coupled Weather–Wildland Fire Modeling with the Weather
910 Research and Forecasting Model. *J. Appl. Meteorol. Climatol.* **52**, 16–38 (2013).
- 911 47. Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. in *Proceedings of the*
912 *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*
913 785–794 (Association for Computing Machinery, New York, NY, USA, 2016).
914 doi:10.1145/2939672.2939785.
- 915 48. Shamsaei, K. *et al.* Coupled fire-atmosphere simulation of the 2018 Camp Fire using WRF-
916 Fire. *Int. J. Wildland Fire* **32**, 195–221 (2023).
- 917 49. Lareau, N. P. Plume-Coupled Long-Range Spotting Drove the Explosive Spread of the
918 2018 Camp Fire. *J. Geophys. Res. Atmos.* **131**, e2025JD045798 (2026).
- 919 50. Frediani, M. *et al.* Modeling Firebrand Spotting in WRF-Fire for Coupled Fire–Weather
920 Prediction. *Fire* **8**, 374 (2025).

- 921 51. Skamarock, W. C. et al. A Description of the Advanced Research WRF Model Version 4.
922 *NCAR Tech. Note NCAR/TN-556+STR*, **145** pp. (2019) doi:10.5065/1dfh-6p97.
- 923 52. Rothermel, R. C. A Mathematical Model for Predicting Fire Spread in Wildland Fuels. Res.
924 Pap. INT-115. (US Department of Agriculture, Forest Service, Intermountain Forest and
925 Range Experiment Station, Ogden, UT, 1972).
- 926 53. Muñoz-Esparza, D., Kosović, B., Jiménez, P. A. & Coen, J. L. An Accurate Fire-Spread
927 Algorithm in the Weather Research and Forecasting Model Using the Level-Set Method. *J.*
928 *Adv. Model. Earth Syst.* **10**, 908–926 (2018).
- 929 54. Wyngaard, J. C. Toward numerical modeling in the “Terra Incognita”. *J. Atmos. Sci.* **61**,
930 1816–1826 (2004).
- 931 55. Muñoz-Esparza, D., Lundquist, J. K., Sauer, J. A., Kosović, B. & Linn, R. R. Coupled
932 mesoscale-LES modeling of a diurnal cycle during the CWEX-13 field campaign: From
933 weather to boundary-layer eddies. *J. Adv. Model. Earth Syst.* **9**, 1572–1594 (2017).
- 934 56. Hong, S.-Y., Noh, Y. & Dudhia, J. A New Vertical Diffusion Package with an Explicit
935 Treatment of Entrainment Processes. *Mon. Weather Rev.* **134**, 2318–2341 (2006).
- 936 57. Deardorff, J. W. Stratocumulus-capped mixed layers derived from a three-dimensional
937 model. *Bound.-Layer Meteorol.* **18**, 495–527 (1980).
- 938 58. Hong, S.-Y. & Lim, J.-O. J. The WRF single-moment 6-class microphysics scheme
939 (WSM6). *Asia-Pac. J. Atmos. Sci.* **42**, 129–151 (2006).
- 940 59. Iacono, M. J. et al. Radiative forcing by long-lived greenhouse gases: Calculations with the
941 AER radiative transfer models. *J. Geophys. Res. Atmos.* **113**, D13103 (2008).

942 60. Dudhia, J. Numerical Study of Convection Observed during the Winter Monsoon
943 Experiment Using a Mesoscale Two-Dimensional Model. *J. Atmos. Sci.* **46**, 3077–3107
944 (1989).

945 61. Jiménez, P. A. et al. A Revised Scheme for the WRF Surface Layer Formulation. *Mon.*
946 *Weather Rev.* **140**, 898–918 (2012).

947 62. Chen, F. & Dudhia, J. Coupling an Advanced Land Surface–Hydrology Model with the
948 Penn State–NCAR MM5 Modeling System. Part I: Model Implementation and Sensitivity.
949 *Mon. Weather Rev.* **129**, 569–585 (2001).

950 63. Hersbach, H. et al. The ERA5 global reanalysis. *Q. J. R. Meteorol. Soc.* **146**, 1999–2049
951 (2020).

952 64. Farr, T. G. et al. The Shuttle Radar Topography Mission. *Rev. Geophys.* **45**, RG2004
953 (2007).

954 65. Scott, J. H. & Burgan, R. E. Standard fire behavior fuel models: A comprehensive set for
955 use with Rothermel's surface fire spread model. Gen. Tech. Rep. RMRS-GTR-153. (US
956 Department of Agriculture, Forest Service, Rocky Mountain Research Station, Fort Collins,
957 CO, 2005).

958 66. Rasmussen, R. et al. High-Resolution Coupled Climate Runoff Simulations of Seasonal
959 Snowfall over Colorado: A Process Study of Current and Warmer Climate. *J. Clim.* **24**,
960 3015–3048 (2011).

961 67. Mandel, J. et al. Recent advances and applications of WRF-SFIRE. *Nat. Hazards Earth*
962 *Syst. Sci.* **14**, 2829–2845 (2014).

963 68. Mardia, K. V. Statistics of directional data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **37**, 349–
964 371 (1975).

69. Lundberg, S. M. & Lee, S.-I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **30**, (2017).
70. Lundberg, S. M. *et al.* From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2**, 56–67 (2020).

Code Availability

WRF-Fire is part of the open-source WRF modeling system, available at <https://github.com/wrf-model/WRF>; this study used the release-v4.6.0 source (commit 0a11865f). All analysis used Python 3.9.16 with NumPy 1.23.5, pandas 2.3.3, SciPy 1.10.0, Matplotlib 3.6.3, netCDF4 1.6.2, GeoPandas 1.0.1, scikit-learn 1.2.0, XGBoost 2.1.4 and SHAP 0.49.1. All random seeds are fixed, so every figure and quoted number is reproducible bit for bit.

Data Availability

The ERA5 reanalysis fields used to initialize and force the simulations are available from the Copernicus Climate Data Store (ERA5 hourly data on pressure levels, <https://doi.org/10.24381/cds.bd0915c6>; ERA5 hourly data on single levels, <https://doi.org/10.24381/cds.adbb2d47>). Surface fuels from the 2014 LANDFIRE release of the Scott and Burgan 40 fuel models are available from the LANDFIRE program at <https://landfire.gov>. Terrain elevation from the Shuttle Radar Topography Mission is available from the USGS at <https://earthexplorer.usgs.gov>, and land cover from the National Land Cover Database at <https://www.mrlc.gov>. The observed Camp Fire perimeter is available from the National Interagency Fire Center at <https://data-nifc.opendata.arcgis.com>. The raw WRF-Fire output is retained on NSF NCAR storage systems and is available from the corresponding author on reasonable request because of its volume.

Author contributions

WL, SR and FZ designed the experiments, conducted the simulations, and analyzed the results. MF, BH, CT, AH and ET helped with running the experiments and analyzing the data. KS and TJ provided suggestions for analyzing the data. All authors contributed to the discussion and final version of the paper.

Competing interests

The authors declare no competing interests.

Acknowledgements

We would like to acknowledge the use of the Derecho system (doi:10.5065/qx9a-pg09) supported by the NSF National Center for Atmospheric Research (NCAR) at the NSF NCAR-Wyoming Supercomputing Center, sponsored by the National Science Foundation and the State of Wyoming. WL and SR acknowledge the National Science Foundation-supported project, “Wyoming Anticipating the Climate-Water Transition (WY-ACT)” (award no. 2149105). WL, SR, and FZ acknowledge financial support from the CO-WY ASCEND grant (NSF-2315760). MF acknowledges support by the NSF National Center for Atmospheric Research, which is a major facility sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1852977. CT and AH acknowledge support from the Regional and Global Model Analysis Program for the Office of Science of the US Department of Energy through the PCMDI project. E.T. wishes to acknowledge the support he received, in part, from the NSF Grants CMMI-1953333, 2536549, 2524950; from Sandia National Laboratories (Contract No. 2169310); from the University of California Institute of Transportation Studies through its Statewide Transportation Research Program, funded by the Road Repair and Accountability Act of 2017 (SB 1), under Project 2025-34; and through a philanthropic research gift from Edison International. Any opinions, findings, and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the sponsors.