

From spatial attribution to mesh decisions in nonlinear moisture-threshold calculations

Songde You

`xsdwp_tg@163.com`

Xiamen University of Technology

Yihong Chen

Xiamen University of Technology

Jin Chen

Xiamen University of Technology

Research Article

Keywords: Nonlinear diffusion, Drying endpoint, Interface flux, Numerical verification, Solver sensitivity, Heat and mass transfer

Posted Date: October 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11259676/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

From spatial attribution to mesh decisions in nonlinear moisture-threshold calculations

Songde You*, Yihong Chen, Jin Chen

Xiamen University of Technology, Xiamen, China

*Corresponding author: Songde You

Email: xsdwp_tg@163.com

Abstract This study links local flux sensitivity to mesh selection for moisture-threshold calculations in a synthetic coupled heat and moisture model. The analysis addresses three practical questions: where a flux change affects the endpoint, how that influence varies with the diffusion law, and how the tested meshes compare in reference difference under recorded CPU budgets. Coupled tangent propagation and a face-resolved adjoint provide consistent endpoint derivatives, with a maximum relative difference of 3.7×10^{-7} in six matched checks. A kernel for a fixed constitutive perturbation changes by less than 0.17% under the final spatial refinement in three cases. Cross-law tests resolve surface concentration and signed cancellation, providing a spatial basis for comparing mesh choices. The cost comparison charges recorded pilot, construction and solve stages, including unsuccessful attempts. In the baseline exponential case, surface grading has lower recorded cost and a smaller reference difference than the tested goal strategy. The goal grid gives a smaller finest-grid reference difference for the strong positive-floor case. A derivative-guard test separates endpoint improvement from constraint satisfaction and detects time sensitivity in the difficult exponential case. A conservative synthetic control quantifies a 9.77% endpoint increase from latent cooling. The study provides an interpretable, cost-accounted procedure for assessing numerical drying endpoints within the tested constitutive laws and operating conditions.

Keywords Nonlinear diffusion; Drying endpoint; Interface flux; Numerical verification; Solver sensitivity; Heat and mass transfer

1 Introduction

A moisture limit that must hold throughout a specimen creates a demanding stopping-time calculation. The centre can remain wet after the surface has dried, and a diffusivity that falls with moisture can make a thin dry region difficult to resolve. Reliable computation requires both an account of where the endpoint is sensitive and a test of how mesh changes affect accuracy and cost.

Existing work supplies the numerical foundations. Maddix et al. (2018) analyse nonlinear diffusion artifacts for functions including $e^{-1/p}$ and develop a modified harmonic method. The factor $e^{-a/C}$ used here has the same form under $p = C/a$. The publisher’s abstract for Prakash K and Natarajan (2025) describes an α -damping correction; its full algorithm was unavailable for comparison. Takatori (2026) compares interface approximations in another moisture-transport setting. These studies and established finite-volume treatments (Patankar 1980; Crank 1975) motivate a closer examination of how interface fluxes affect a drying endpoint.

Threshold-time analysis connects a trajectory perturbation to its crossing time. Chaudhry et al. (2021) treat ordinary differential equations, and Chaudhry et al. (2023) address evolutionary partial differential equations. The present study uses this local sensitivity framework to trace the effect of individual faces,

compare constitutive laws and assess finite mesh changes. This sequence connects an interpretable spatial response to a directly tested computational decision.

The heat and moisture setting also requires compatible mass and energy definitions (Datta 2007). Moving-boundary and identification studies make these assumptions explicit (Adrover et al. 2019; Vaz Jr. and Zdanski 2024), while broccoli-stem, okra-cylinder and hazelnut studies illustrate material-specific validation (Simal et al. 1998; Nwakuba et al. 2025; Çelik et al. 2026). Here, the contest inputs define a synthetic benchmark with controlled numerical comparisons. A separate conservative closure isolates the contribution of latent cooling within a consistent water and energy balance.

This paper contributes a linked assessment of flux response, spatial influence and mesh cost. Coupled tangent and adjoint calculations give independently checked routes to the same endpoint derivative. Cross-law comparisons identify changes in surface localization and cancellation, and comparisons at fixed recorded-cost budgets account for the measured setup stages and attempted solves. A derivative-guard test adds local face checks to distinguish a completed endpoint from a satisfied refinement criterion. The added evidence connects constitutive changes in spatial attribution to case-specific mesh outcomes, including the cost of local derivative checks. This link makes the conditions for a useful mesh change directly examinable in the tested cases.

2 Benchmark and numerical formulation

2.1 Inputs, scope, and endpoint

The inputs are Problem A, *Drying of medicinal plant materials* (translated title), and two attachments from the official 2026 China Undergraduate Mathematical Contest in Modeling release (CUMCM 2026). Copies downloaded on 27 September 2026 are byte-identical to the supplied PDF and workbooks, according to SHA-256 checks. The statement does not identify the plant species, instruments, repetitions, or measurement uncertainty. The file provenance is verified; the status of the records as experimental measurements remains unknown.

The primary specimen is a fixed-radius cylinder with $R = R_0 = 0.020$ m and length $L = 0.25$ m. It starts at 28°C and dry-basis moisture $C_0 = 2.55$ kg water per kg dry solid. The specified threshold is $C_{\text{th}} = 0.15$ kg/kg, with no known pharmacopoeial basis. The endpoint is defined by

$$M(t) = \max_{0 \leq r \leq R} C(r, t), \quad t^* = \inf\{t \geq 0 : M(t) < C_{\text{th}}\}. \quad (1)$$

Every endpoint test checks the nodal maximum. The unperturbed trajectories used for centre-based differentiation are checked to have their maximum at the centre. The backward-Euler verification interpolates its two bracketing time levels; its saved passing field belongs to the later discrete state. The adaptive BDF calculations instead locate the terminal event on the solver interpolant. The verification and response studies report their respective endpoint procedures.

The environmental attachment contains 241 records at 60 s intervals over 4 h. Piecewise-linear interpolation defines the temperature and empirical ambient-moisture histories. Their final values, 50.165°C and 0.04986 kg/kg, are held constant after 4 h unless a stated scenario changes that continuation. The continuation is assumed, rather than measured over the full drying period. A separate radius attachment contains 145 records at 1800 s intervals over 72 h. The prescribed-radius comparisons are outside the present fixed-domain flux and solver study.

2.2 Effective heat and moisture equations

The model uses

$$\rho(C)c_p(C)T_t = \frac{1}{r}\partial_r(rk(C)T_r), \quad (2)$$

$$C_t = \frac{1}{r}\partial_r(rD(C, T_K)C_r), \quad T_K = T + 273.15. \quad (3)$$

The axis has zero gradients. The surface conditions are

$$-k_s T_r(R, t) = h(T_s - T_\infty), \quad -D_s C_r(R, t) = h_m(C_s - C_\infty), \quad (4)$$

where $h = 25 \text{ W/(m}^2 \text{ K)}$ and $h_m = 8 \times 10^{-7} \text{ m/s}$. The primary supplied property set B is

$$\rho_B = 650 + 128C, \quad c_{p,B} = 1450 + \frac{2736C}{1+C}, \quad k_B = 0.21 + \frac{0.38C}{1+C}, \quad (5)$$

$$D_B = 2.4 \times 10^{-3} \exp(-0.45/C - 3850/T_K). \quad (6)$$

The units of ρ, c_p, k, D are $\text{kg/m}^3, \text{J/(kg K)}, \text{W/(m K)}$, and m^2/s . The moisture parameters use the numerical kg/kg convention of C , and 3850 has units of K.

The moisture field is a mass ratio, so $-D\nabla C$ is not a water mass flux in $\text{kg/(m}^2 \text{ s)}$. The ambient and solid moisture variables may also have different mass bases. Their difference in Eq. (4) is an empirical benchmark assumption. The energy equation includes sensible heating but excludes latent heat, radiation, and reactions. End-face transport is omitted. Six constant-diffusivity finite-cylinder comparisons give centre-endpoint reductions of at most 0.0321%, but this restricted linear check does not bound axial error in the coupled nonlinear problem.

2.3 Face fluxes and production settings

The backward-Euler verification solver uses a uniform vertex-centred grid in $x = r/R$, with volume weights $V_i = (x_{i+1/2}^2 - x_{i-1/2}^2)/2$ truncated at the axis and surface. Backward Euler advances the fields. Thermal conductivity is averaged harmonically. Each interior moisture face contributes equal and opposite terms to the two adjacent volumes.

The diffusivity is written as $D = A(T_K)f(C)$, where $f(C) = e^{-a/C}$. With a common arithmetic face temperature, the harmonic and arithmetic moisture means are $2f_i f_{i+1}/(f_i + f_{i+1})$ and $(f_i + f_{i+1})/2$. The integral mean follows from

$$K_a(C) = C e^{-a/C} + a \text{Ei}(-a/C), \quad K'_a(C) = e^{-a/C}, \quad (7)$$

giving the inward transport contribution

$$F_{i+1/2} = \frac{x_{i+1/2} A(T_{K,i+1/2})}{R^2 \Delta x} [K_a(C_{i+1}) - K_a(C_i)]. \quad (8)$$

The method integrates moisture dependence at a fixed face temperature; it does not exactly integrate an arbitrary temperature gradient. H, A, and K denote the three moisture means. A further harmonic form, HF, averages the complete node-temperature diffusivities. The Picard–Newton comparisons below retain that same HF face law.

The backward-Euler verification uses K flux, 640 radial intervals, and 30 s steps as its reporting setting. The flux-propagation study uses the adaptive settings specified in Sect. 3.3 and the following subsection. The backward-Euler outer iteration alternates heat and moisture updates to a largest field-relative increment below 10^{-8} . The moisture Newton solve uses a positive-field line search and a residual scaled by the largest absolute assembled right-hand-side entry, with tolerance 10^{-10} . Temperature is stored in Celsius and converted to kelvin in the diffusivity. The tests in Sect. 3.2 impose stricter shared acceptance criteria where stated. Rejected attempts are recorded and supply no accepted endpoint.

3 Numerical attribution and comparison design

3.1 Verification and reference calculations

The verification separates spatial assembly, time integration and nonlinear acceptance. A manufactured field provides a known low-moisture solution and crossing (Salari and Knupp 2000):

$$U(x, t) = 0.04 + b(t)(1 - x^2), \quad b(t) = 0.11e^{-t}, \quad (9)$$

$$\mathcal{D}(U) = \exp(0.45/0.10 - 0.45/U), \quad (10)$$

$$s(x, t) = -b(1 - x^2) + 4b\mathcal{D} - 4b^2x^2\frac{0.45\mathcal{D}}{U^2}. \quad (11)$$

The equation is $U_t = x^{-1}\partial_x(x\mathcal{D}U_x) + s$ on the unit-radius cylinder. The axis has zero gradient, and a unit-coefficient Robin boundary uses $U_\infty = 0.04 - 2b\mathcal{D}(0.04)$. Six-point volume quadrature integrates the independently differentiated source. Spatial tests use $N = 20, 40, 80, 160$, $\Delta t = 10^{-6}$ and $t_f = 0.01$. Temporal tests use $N = 1280$, steps from 0.002 to 0.00025 and $t_f = 0.7$. The threshold 0.10 has exact crossing $\ln(11/6) = 0.6061358036$. A separate endpoint sequence uses $N = 40, 80, 160, 320$ and $\Delta t = 2 \times 10^{-5}$. Half-step, doubled-grid and tighter-residual controls distinguish the remaining errors. These tests isolate the moisture operator; they do not manufacture the full coupled thermal solution.

An independently assembled coupled reference uses 16-point Gauss–Legendre face integration and adaptive BDF, without calling the backward-Euler solver’s grid, property or flux functions. Its relative tolerance is 10^{-8} , with absolute tolerances 10^{-8} for temperature and 10^{-10} for moisture, maximum step 120 s, and an event at $M = 0.15$. The baseline and steep cases use $N = 640, 1280$. Separate steep-case controls double the quadrature order or tighten relative tolerance to 10^{-10} and absolute tolerances to $10^{-10}, 10^{-12}$. Backward-Euler K calculations use 30, 15 and 7.5 s. Both implementations share the vertex-centred spatial balance, so their agreement checks assembly and time integration without supplying an independent spatial discretization. The subsequent attribution studies use the tighter settings stated in their own protocols.

3.2 Constitutive scenarios and solver diagnostics

The exponential family retains the thermal properties and matches diffusivity at $C_{\text{ref}} = 0.15$:

$$D_a(C, T_K) = 2.4 \times 10^{-3} \exp \left[\frac{a - 0.45}{0.15} - \frac{a}{C} - \frac{3850}{T_K} \right]. \quad (12)$$

The baseline has $a = 0.45$ and the recorded ambient continuation. The steep case has $a = 0.50$ and a -0.01 kg/kg ambient change, ramped from 4 to 5 h and held afterwards. An intermediate check uses $a = 0.475$ and -0.005 kg/kg; a local-conductance control uses $a = 0.45$ and -0.01 kg/kg. Matching one moisture value does not match the complete diffusivity curves.

The initial H/A/K backward-Euler comparisons use moisture Newton tolerance 10^{-10} and outer increment tolerance 10^{-8} . Matched HF Picard and Newton controls instead impose both a scaled residual and a relative field increment:

$$\max(\hat{R}_C, \hat{R}_T) \leq \varepsilon_R, \quad \max_{U \in \{C, T\}} \frac{\|U^{q+1} - U^q\|_\infty}{\|U^{q+1}\|_\infty} < \varepsilon_I. \quad (13)$$

Each residual is divided by the largest absolute entry of its assembled right-hand side; the matched setting uses $\varepsilon_R = \varepsilon_I = 10^{-10}$. Picard freezes coefficients, while Newton uses analytic moisture derivatives and a positive-field line search. Both use alternating thermal updates, the preceding accepted state as the initial iterate, and an allowance of 5000 outer iterations. Separate controls change tolerance, relaxation, time step or the arithmetic order of an equivalent geometric factor. Failed attempts retain their checkpoints and receive no endpoint.

At fixed temperature, write $q = 2d_L d_R (C_R - C_L) / (d_L + d_R)$, where d_L, d_R are the nodal diffusivities. The exponential harmonic law has right derivative

$$q_R = \frac{2d_L d_R}{d_L + d_R} (1 - \eta), \quad \eta = \frac{a(C_L - C_R)}{C_R^2} \frac{d_L}{d_L + d_R}. \quad (14)$$

The derivative changes sign at $\eta = 1$. Analytic and finite-difference derivatives are compared at three failed steps and their preceding accepted states. The checks inspect the frozen-temperature moisture Jacobian J_C and $B = I - \Delta t J_C$, with K evaluated on the same fields. This local diagnostic does not characterize the stability of the full coupled trajectory.

A separate Cartesian comparator implements the cubic modified harmonic method (MHM) of Maddix et al. (2018). For $p_t = (p^3 p_x)_x$, the forward-Euler harmonic right-hand side receives

$$\frac{\Delta x^2}{4} [9p_i (D^- p_i)^2 D^+ D^- p_i + 5(D^- p_i)^4], \quad (15)$$

using backward first and centred second differences at the old time level. H, A, K and MHM use $N = 50, 100, 200$, $\Delta t = \Delta x^2/32$ and $t_f = 0.3$, with $p(0, t) = (3t)^{1/3}$ and an interior/right-boundary floor 10^{-3} . The comparison records the crossing of $p(0.12, t) = 0.5$. The zero-floor wave $[3 \max(t - x, 0)]^{1/3}$ is not exact for the floored problem. Controls halve the step or reduce the floor to 10^{-4} . The smooth wave $p = [3(t + 1.2 - x)]^{1/3}$ provides a separate consistency check at $N = 25, 50, 100, 200$, $t_f = 0.05$ and $\Delta t = \Delta x^2/128$. No radial extension of MHM is assumed.

The same-field harmonic-to-integral conductance ratio supplies a local transport diagnostic. Its common temperature factor cancels; a small ratio does not directly determine a drying-time error. For a differentiable crossing at a fixed maximum location, a threshold change instead gives $\Delta t_{th} \approx \Delta C_{th} / M'(t^*)$, while a field perturbation gives $\delta t \approx -\delta M(t^*) / M'(t^*)$. Dense-output crossings at 0.151 and 0.153 kg/kg test the threshold approximation on baseline, steep and intermediate K trajectories. Neither this conditioning check nor the face ratio predicts an accumulated flux error without trajectory propagation.

OpenAI Codex assisted with manuscript revision, literature screening and numerical-code development. The recorded scripts generated numerical results; no experimental observations were generated.

3.3 Propagation of a face-flux perturbation

The state $y = (T, C)$ collects the nodal fields. Let $F_K(t, y)$ be the coupled semidiscrete right-hand side and let $g_F(t, y)$ be the conservative divergence of the H-minus-K moisture face flux, with zero direct heat forcing. H and K use the same face-temperature factor. We define

$$\dot{y}_\epsilon = F_K(t, y_\epsilon) + \epsilon g_F(t, y_\epsilon), \quad \dot{z} = J_K(t)z + g_F(t, y_0), \quad z(0) = 0, \quad (16)$$

where $J_K = \partial F_K / \partial y$ is evaluated along y_0 and $z = \partial y_\epsilon / \partial \epsilon|_0$. The matrix includes the moisture dependence of thermal capacity and conductivity, and the temperature dependence of diffusivity. Holding temperature fixed in this propagation would omit part of the response.

Suppose the crossing has one active maximum node i_* , $\dot{C}_{i_*}(\tau_0) \neq 0$, and the solution depends smoothly on ϵ near zero. Differentiating $C_{i_*}(\tau_\epsilon, \epsilon) = C_{th}$ gives

$$S_F^{(s)} = \left. \frac{d\tau_\epsilon}{d\epsilon} \right|_0 = -\frac{z_{C, i_*}(\tau_0)}{\dot{C}_{i_*}(\tau_0)}, \quad \tau_\epsilon - \tau_0 = \epsilon S_F^{(s)} + O(\epsilon^2). \quad (17)$$

The integration time and τ_ϵ are measured in seconds. We report $S_F = S_F^{(s)} / 3600$ in hours per unit perturbation weight. The tangent components and adjoint face sums use this same normalization. This event-sensitivity formula computes the local endpoint difference between discrete flux laws by propagating the

specified face perturbation. The observed centre maximum at crossing is checked; a changing active maximum or tangential crossing would require a different treatment.

The forcing is split into four components using faces inside or outside $x = 0.9$, and times before or after $\tau_0/2$. Each component is propagated with the same J_K . Linearity of the tangent equation gives $S_F = \sum_{b=1}^4 S_b$. The signed contributions locate sensitivity on the unperturbed trajectory; they are not an exact partition of a finite nonlinear endpoint change. The outer region is 10% of the normalized coordinate, not 10% of specimen volume.

3.4 Transfer scenarios and resolution intervention

Table 1 specifies six scenarios before their execution. The radial operator is generalized to $r^{-m}\partial_r(r^m DC_r)$, with the same change for heat. Cell weights become $V_i = (x_{i+1/2}^{m+1} - x_{i-1/2}^{m+1})/(m+1)$. All other constitutive and boundary assumptions remain explicit. These are synthetic geometry and parameter transfers, not additional medicinal materials.

Table 1 Prespecified synthetic transfer scenarios

Case	m	R (mm)	a	δC_∞	C_{th}	h	$10^7 h_m$
S1	1	20	0.45	0	0.15	25	8
S2	1	20	0.5	-0.01	0.15	25	8
S3	0	20	0.45	0	0.15	25	8
S4	2	20	0.5	-0.01	0.15	25	8
S5	1	15	0.4	0.005	0.15	25	8
S6	2	25	0.475	-0.005	0.12	15	12

$m = 0, 1, 2$ denote a symmetric slab half-domain, cylinder, and sphere. R is half-thickness for the slab. Moisture quantities are in kg/kg; h and h_m retain the units defined in Sect. 2. All cases retain the set-B thermal laws and initial fields. The diffusivity prefactor is normalized as in Eq. (12). S1 and S2 reuse known parameter settings; S3–S6 are transfer tests, not independent material datasets.

At $N = 160$, each case uses $\epsilon = -0.1, -0.05, 0, 0.05, 0.1, 0.2, 0.5$. The reference trajectory and tangent prediction do not use the perturbed endpoints. Negative weights provide derivative checks; they are diagnostic extrapolations rather than a proposed production flux. Independent integrations supply the finite shifts. Analytic full Jacobians are checked by central directional differences at three trajectory states. The BDF relative tolerance is 10^{-9} , with absolute tolerances 10^{-9} for temperature and 10^{-11} for moisture, and maximum step 120 s. The four tangent components use relative tolerance 10^{-8} and absolute tolerances 10^{-10} and 10^{-12} . Splitting the tangent integration at $\tau_0/2$ handles the attribution discontinuity explicitly.

S1, S2, and S6 repeat $\epsilon = 0, 0.1$ at $N = 320, 640$. A separate S2 control tightens both relative tolerances tenfold and reduces the maximum step to 60 s. The analytic-primitive assembly at $N = 640$ is also checked against the earlier independent quadrature reference. The protocol specifies a 10% prediction-discrepancy level as a descriptive benchmark, not a rigorous bound or a fitted acceptance rule.

After the attribution results, an additional experiment uses $x_i = 1 - (1 - i/N)^2$ at $N = 80, 160, 320$ for S1, S2, and S6. The face geometry uses actual adjacent-node distances, and both H and K retain the same control-volume balance. The exponent two is fixed for all cases. Equal-node uniform controls test H at $N = 160$. Further HF calculations retain the complete node-temperature diffusivities of the earlier difficult solver comparison and extend the graded sequence to $N = 640$. This intervention is disclosed as a follow-up, not a prespecified independent holdout.

An independently coded manufactured check verifies the graded-grid geometry for $m = 0, 1, 2$. It retains U and $\mathcal{D}$ from Eqs. (9)–(10), with source $s_m = -b(1 - x^2) + 2b(m+1)\mathcal{D} - 4b^2x^2a\mathcal{D}/U^2$ and

$a = 0.45$. Eight-point quadrature integrates the volume source. The tests use $N = 20, 40, 80, 160$, final time 0.01, BDF relative and absolute tolerances 10^{-11} and 10^{-13} , and maximum step 10^{-4} . The exact Robin condition remains the one specified in Sect. 3.1.

3.5 Cross-law tests and a local sign diagnostic

The second transfer protocol changes the moisture function itself. The tests use

$$D(C, T_K) = D_r(T_K)f(C), \quad D_r(T_K) = 2.4 \times 10^{-3} e^{-3-3850/T_K}, \quad C_r = 0.15. \quad (18)$$

Table 2 defines the factors and their exact primitives. The normalization matches the original diffusivity at $C = C_r$ for every temperature; it does not match complete trajectories. Quadratic diffusion is an established porous-medium test family (Maddix et al. 2018). The two positive floors and the saturating factor are deliberately synthetic contrasts. They are not alternative measured constitutive fits.

Table 2 Synthetic moisture factors with a shared reference diffusivity

Label	Factor $f(C)$	Primitive $F(C)$
P2	u^2	$C^3/(3C_r^2)$
F1, F10	$(\beta + u^2)/(1 + \beta)$	$[\beta C + C^3/(3C_r^2)]/(1 + \beta)$
S	$2u/(1 + u)$	$2[C - C_r \ln(1 + C/C_r)]$
Constant	1	C

$u = C/C_r$, $C_r = 0.15$ kg/kg; F1 and F10 use $\beta = 1$ and 10. All factors satisfy $f(C_r) = 1$. The constant label refers to moisture dependence; the Arrhenius temperature factor remains active. These functions are synthetic tests, not fitted descriptions of a material.

The protocol fixes the four nonconstant factors, both cylindrical and spherical geometries, and $\epsilon = -0.05, 0, 0.05, 0.1, 0.5$ before execution. The 40 resulting trajectories retain the base radius, thermal laws, boundary coefficients, initial conditions, and threshold. Four additional constant-factor trajectories use $\epsilon = 0, 1$ and supply a zero-response control. The grid has 160 intervals. Each nonconstant cylindrical case has a 320-interval pair at $\epsilon = 0, 0.1$ and a separate tighter 160-interval pair with relative tolerance 10^{-10} and maximum step 60 s. These 16 controls distinguish small paired shifts from integration sensitivity. The spatial and temporal partitions are unchanged, and no prediction coefficient is fitted.

A Taylor expansion supplies a sign diagnostic. Let $\bar{C} = (C_L + C_R)/2$, $d = C_R - C_L$, and let f be positive and smooth near $\bar{C}$. With derivatives evaluated at $\bar{C}$, the integral and harmonic moisture factors satisfy

$$\bar{f}_H - \bar{f}_K = \frac{f f'' - 3(f')^2}{12f} d^2 + O(d^4), \quad \bar{f}_K = \frac{F(C_R) - F(C_L)}{d}. \quad (19)$$

The expansion follows by expanding the two endpoint values and the integral about their midpoint. It is a local consistency calculation, not a new general error estimate. The coefficient is negative for the quadratic, saturating, and original exponential functions on $C > 0$. For a positive-floor quadratic factor it becomes

$$\frac{f f'' - 3(f')^2}{12f} = \frac{\beta - 5u^2}{6C_r^2(1 + \beta)(\beta + u^2)}, \quad u = \bar{C}/C_r. \quad (20)$$

Its sign changes at $\bar{C} = C_r \sqrt{\beta/5}$, which is about 0.0671 for F1 and 0.2121 for F10. This coefficient describes a local conductance difference. The signed concentration gradient and subsequent coupled propagation still determine the endpoint contribution. A local sign alone cannot determine the sign of the accumulated shift.

Independent 64-point Gauss–Legendre integration checks the analytic primitives over four moisture intervals spanning 0.04–2.55 kg/kg. Directional finite differences check the full coupled Jacobian and its moisture block. Additional manufactured calculations use $U = 0.04 + b(1 - x^2)$, $b = 0.11e^{-t}$, with source $-b(1 - x^2) + 2b(m + 1)f(U) - 4b^2x^2f'(U)$ and Robin ambient $0.04 - 2bf(0.04)$. The production moisture operator is retained, while a fixed temperature and adjusted prefactor remove the thermal scale. The source is written independently of the production primitive and integrated with 16-point volume quadrature. Five factors, two geometries, and 20, 40, 80, and 160 intervals give 40 checks at dimensionless time 0.01. These tests isolate moisture assembly; they are not manufactured solutions of the coupled thermal problem.

3.6 Spatial influence kernels and cost-accounted redistribution

The four fixed partitions are replaced by a face-resolved adjoint calculation. Along the unperturbed K trajectory, the discrete adjoint solves

$$-\dot{\lambda} = J_K^T \lambda, \quad \lambda(\tau_0) = -\frac{e_{C,i_*}}{3600 \dot{C}_{i_*}(\tau_0)}. \quad (21)$$

The factor 3600 expresses the response in hours while integration time is in seconds. Each internal face has an inward K contribution q_f^K and an H-minus-K difference δq_f . Their influence is

$$\Theta_f(t) = \left(\frac{\lambda_{C,i}}{V_i} - \frac{\lambda_{C,i+1}}{V_{i+1}} \right) \delta q_f(t), \quad s_f = \int_0^{\tau_0} \Theta_f dt, \quad b_f = \int_0^{\tau_0} |\Theta_f| dt. \quad (22)$$

The identity $S_F = \sum_f s_f$ checks the adjoint against the earlier forward tangent computation. Cumulative sums over every face allow any spatial cutoff to be evaluated. The nonnegative b_f also measures activity that may cancel in the signed response. Neither quantity is a fraction of physical drying error.

A separate kernel concerns a fixed constitutive perturbation rather than a grid-dependent flux substitution. Consider $D \mapsto D[1 + \eta\beta(x, t)]$ in the interior operator, while holding the Robin exchange law fixed. For a sufficiently regular trajectory and its coupled adjoint, the formal continuum representation is

$$\left. \frac{d\tau}{d\eta} \right|_0 = \int_0^{\tau_0} \int_0^1 \Psi(x, t) \beta(x, t) dx dt, \quad \Psi = -\frac{x^m}{R^2} D C_x \partial_x p_C. \quad (23)$$

Here p_C is the moisture component of the adjoint in the $x^m dx$ pairing, with the same endpoint-time normalization as Eq. (21); the full adjoint retains thermal feedback. The boundary formulation is included in the adjoint so that the prescribed exchange law is unperturbed. In the discrete calculation, $p_{C,i}$ is represented by $\lambda_{C,i}/V_i$, and replacing δq_f by q_f^K in Eq. (22) gives a face-integrated coefficient kernel. Its piecewise spatial density is compared at 160, 320, and 640 intervals. Independent direct integrations with D multiplied by 1 ± 0.01 check the spatially uniform perturbation. Four- and eight-point quadrature on the union of primal and adjoint time intervals checks the time integral. This provides a discretized influence kernel and refinement evidence, not a proof of continuum convergence for a pointwise terminal functional.

For a smooth resolved field, Eq. (19) gives $\delta q_f = O(\Delta x_f^2)$ and $s_f = O(\Delta x_f^3)$. An absolute leading-term surrogate has the form $E_{\text{abs}} = \int_0^1 w(x) \ell(x)^2 dx$, where ℓ is local mesh spacing. Minimizing this surrogate under $\int_0^1 \ell^{-1} dx = N$ gives a monitor proportional to $w^{1/3}$. This construction uses established equidistribution (Chen 1994) and adjoint-weighting principles (Becker and Rannacher 2001). The implemented monitor uses $w_f = b_f/\Delta x_f^3$ from an 80-interval K pilot and a fixed positive floor:

$$M_f = \left(0.01 + 0.99 \frac{w_f}{\bar{w}} \right)^{1/3}, \quad \bar{w} = \sum_f w_f \Delta x_f. \quad (24)$$

The implementation interpolates M_f linearly onto 4097 equally spaced auxiliary points on $[0, 1]$, with constant continuation beyond the first and last pilot faces. A cumulative trapezoidal rule integrates that monitor.

Linear inverse interpolation maps equally spaced cumulative targets to new nodes, and the endpoints are reset to 0 and 1. The first cost protocol constructs the targets as $\text{linspace}(0, 1, N + 1)I$, where I is the total monitor integral. The guard protocol uses $\text{linspace}(0, I, N + 1)$. These equivalent formulas need not produce identical floating-point nodes; each protocol retains its own saved grids. The static redistribution uses only the pilot fields; the following comparisons test its accuracy and cost, including cases with a poorly resolved dry boundary layer.

Three cases are fixed before execution: E45 is S1, E50 is S2, and F10 is the cylindrical strong-floor case of Sect. 3.5. Four strategies use 80, 160, and 320 intervals: uniform spacing, the earlier quadratic surface grading, the goal monitor, and a gradient monitor. The last uses $[0.01 + 0.99g/\bar{g}]^{1/2}$, where g is the time-average of $|C_x|$ on the same K pilot. All final solves use H with a common face temperature. Uniform K references at 1280 and 2560 intervals supply numerical comparisons and a reference-refinement check. Eighteen manufactured-field tests check the production moisture operator on the new goal and gradient grids, using three cases and three grid levels with the independent source construction from Sect. 3.5.

The timing platform uses an AMD Ryzen 7 7840H processor. Each final candidate is timed twice in a fixed randomized serial order, with single-thread execution and a 45 CPU-second solve limit. Process CPU and wall times are retained. Recorded-cost accounting charges the pilot solve and, for the goal strategy, its adjoint once; it also charges the measured mesh construction and final attempts at 80, 160 and 320 intervals. The gradient strategy's trajectory sampling and gradient integration were performed outside these timers and are not charged. Final-attempt times include compressed result writing, while candidate-model construction occurs before the timer. Pilot-output writing, fine references and research-only kernel refinement are excluded. This account measures the recorded stages, rather than end-to-end runtime.

At recorded budgets of 30, 60 or 120 CPU seconds, selection uses the last completed candidate whose accumulated charge fits the budget. The charge includes each preceding affordable attempt, including time-outs and noncrossings; reference differences do not select candidates. Pilot timings are measured once and charged to both repetitions. Only the final attempts are repeated. The comparison is retrospective and machine-specific; the records do not establish whether unmeasured overhead would change a budget cutoff.

3.7 A bounded derivative-guard test

The difficult goal-grid result motivates a separate prospective protocol. For the inward harmonic contribution $q = \bar{f}_H(C_R - C_L)$ at a fixed positive face-temperature factor, define

$$\eta_R = \frac{f_L}{f_L + f_R} \frac{f'_R}{f_R} (C_L - C_R), \quad \eta_L = \frac{f_R}{f_L + f_R} \frac{f'_L}{f_L} (C_R - C_L). \quad (25)$$

Direct differentiation gives $q_R = \bar{f}_H(1 - \eta_R)$ and $q_L = -\bar{f}_H(1 - \eta_L)$. Testing both sides also covers a locally reversed moisture gradient. If $\max(\eta_L, \eta_R) \leq 1$ on every face, the frozen-temperature moisture Jacobian J_C has nonnegative off-diagonal entries. Internal flux cancellation gives nonpositive column sums for VJ_CV^{-1} , where V is the diagonal matrix of positive control-volume weights; the linear Robin term contributes a nonpositive boundary loss. The column Gershgorin bounds place its spectrum in the closed left half-plane. Hence $I - \Delta t J_C$ is a nonsingular M-matrix for $\Delta t > 0$. This is a local algebraic sufficient condition using established matrix properties. It does not establish the same property for the coupled heat-moisture Jacobian, all Newton trial states, or an entire time trajectory.

The protocol fixes a margin of 0.8 before execution; this margin is not optimized. The guard starts from the 80-interval pilot's goal mesh at nominal levels $B = 160, 320$. Moisture is linearly interpolated onto the proposed nodes, and both indicators are evaluated at accepted pilot times and on a 120 s lattice. Every face exceeding 0.8 is bisected simultaneously, for at most 12 interpolation passes. The proposed grid is subsequently checked by re-solving K, with at most three trajectory audits. A failed audit triggers simultaneous bisection of the violating faces if another audit remains. The last audit is not followed by an

unchecked mesh change. A proposed split above 640 intervals leaves the current grid unchanged and is reported as unresolved. No coarsening or reference-endpoint information is used. The resulting interval count N can exceed B .

E50 is the already observed development failure. Two cases are fixed before their results are examined: E55H uses cylindrical geometry, $a = 0.55$ and ambient continuation change -0.015 kg/kg; F1SH uses the F1 factor, spherical geometry, and change -0.005 kg/kg. Their other inputs retain the benchmark values. These are prospective numerical holdouts for this fixed rule, not independently measured materials or a random sample of diffusion models. The guard is compared with the original goal mesh and quadratic surface grading. All final calculations use H. Uniform K references at 1280 and 2560 intervals remain separate from mesh selection.

Each of the 18 candidates is timed twice in a fixed randomized serial order, with a 45 CPU-second final-solve limit. A K audit has the same solve limit. For the goal and guard strategies, the recorded charge includes the pilot, adjoint, pilot-trajectory sampling and storage once. Every strategy is charged for its measured construction, any K re-solves and H attempts at $B = 160, 320$. Construction timing excludes writing the prepared-grid file. Final-attempt timing excludes grid loading, candidate-model construction, post-run H diagnostics and result writing. Fine references are excluded. These boundaries differ from those of the first cost protocol, so comparisons use each protocol's own recorded-cost scale.

The fixed recorded budgets are 60, 120 and 180 CPU seconds. Selection uses the last completed affordable candidate; its reference discrepancy does not enter the rule. Setup is measured once and charged to both final-attempt repetitions. The H diagnostics never alter a mesh. Their indicator maxima concern accepted and interpolated samples, rather than bounds between sample times.

Independent two-sided face differences check Eq. (25). The source construction in Sect. 3.5 tests each guarded grid and a uniform bisection of all its intervals. This additional bisection is a verification grid, not a new timed candidate. Its paired errors test refinement on a fixed shape; no order is inferred by comparing the unrelated $B = 160$ and $B = 320$ guarded grids. Separate finer-time controls check the finest surface and guarded H candidates. All unsuccessful final solves and unresolved guards are retained.

The author-held archive `IJST_Author_Working_Package.zip` retains the code, fixed protocols, saved grids and raw attempts. Its `REPRODUCIBILITY.md` maps each displayed result to the corresponding files, and `qa/review16/reproducibility_manifest.json` records their SHA-256 hashes. Tables 7 and 9 use the separate records in `results/review11/` and `results/review12/`, respectively. The code and generated numerical results can be obtained from the corresponding author on reasonable request, as stated in the data and code availability statement.

The first cost study records Python 3.12.11 on Windows 11 (build 26100). Its historical NumPy and SciPy versions were not recorded. The guard study has no separately recorded interpreter or library versions; these versions remain unknown. The two timing protocols use order seeds 11027 and 12029, respectively. Current installed versions cannot establish the historical environment, and the unpinned dependency list does not reconstruct it.

4 Results and discussion

4.1 Verification of the endpoint calculations

The low-moisture manufactured field separates numerical consistency from agreement between drying implementations. Its maximum spatial error decreases from 3.118×10^{-5} to 5.623×10^{-7} , with orders approaching two. Halving the finest time step changes that error by about 1.9×10^{-11} . Temporal crossing errors decrease from 1.278×10^{-4} to 1.605×10^{-5} , approximately in proportion to the step. The endpoint

spatial sequence decreases from 8.184×10^{-5} to 2.534×10^{-6} ; residual time error affects its finest ratios. Three coarse attempts at residual tolerance 10^{-12} stagnate and remain recorded as failures.

The independently assembled coupled K reference supports the backward-Euler time sequence. At $N = 640$ and 7.5 s, baseline, steep and intermediate differences from BDF are 0.006535, 0.006688 and 0.006616 h. The differences approximately halve with the step. Doubling the BDF grid changes baseline and steep endpoints by 3.44×10^{-5} and 3.65×10^{-5} h. The steep quadrature-order and time-tolerance controls change the endpoint by 9.14×10^{-9} and 3.11×10^{-8} h. These are observed numerical differences, not certified continuous-solution error bounds.

The difficult harmonic calculations do not inherit that reference status. Matched baseline HF Picard and Newton endpoints agree within 2×10^{-10} h. In the steep case, requiring residual convergence changes the increment-only Picard endpoint from 200.742068 to 199.729383 h; the increment-only run has maximum scaled moisture residual 4.45×10^{-8} . Tightening both acceptance criteria to 10^{-10} gives 199.730777 h for Picard, while Newton fails at the 30 s setting. At 15 s both complete but differ by about 0.043 h. Equivalent geometric-factor arithmetic also changes a Newton endpoint from 199.294228 to 199.172856 h. Adaptive HF controls range from 199.071887 to 200.883137 h. Completion does not resolve these sensitivities, and the previously reported 261.48% harmonic–integral difference cannot be treated as an isolated spatial error.

The failure checkpoints provide a local explanation. Each has a face with $q_R < 0$ near $x = 0.99921875$. Independent derivative differences are below 1.3×10^{-9} . At rejected iterates, the largest real moisture-Jacobian eigenvalue approaches $1/30 \text{ s}^{-1}$, and $I - 30J_C$ has condition numbers 3.10×10^9 – 1.16×10^{10} . Same-field K Jacobians instead have negative spectral abscissae. The near-singular moisture block explains a local Newton difficulty without proving multiple roots or every long-time amplification mechanism.

The Cartesian MHM comparator restores propagation where H fails to cross by 0.3, but its crossing discrepancy relative to the zero-floor wave remains 0.015217–0.009062 over $N = 50$ –200. Halving the step or reducing the floor changes the $N = 100$ crossing by only 1.56×10^{-6} or 1.47×10^{-5} . The smooth exact-wave errors decrease from 1.50×10^{-4} to 2.31×10^{-6} , with orders 2.004–2.013. These checks support the implemented published correction in its stated Cartesian scope; the nonsmooth floored comparison cannot establish its radial drying accuracy.

The threshold-conditioning check has a separate local use. Increasing the threshold by 0.001 kg/kg gives prediction differences below 1.8% across the three K scenarios; a 0.003 kg/kg change raises them to about 4.8–5.2%. The same-field conductance ratio describes reduced harmonic transport near the dry surface. Neither observation quantifies the accumulated effect of changing the face law. The following response calculations address that question directly.

4.2 Flux-response prediction across geometries

The 42 prespecified trajectories complete with positive fields. The full-Jacobian directional checks have relative discrepancies below 2.7×10^{-10} . Central endpoint differences at $\epsilon = \pm 0.05$ agree with the tangent derivative within 0.20%. These checks assess the implemented derivative independently of its comparison with larger finite shifts.

Figure 1 compares the finite-perturbation discrepancy and the spatial and temporal response fractions.

At $\epsilon = 0.1$, prediction discrepancies range from 2.44% to 7.85% across the six scenarios (Table 3). At 0.05 they range from 1.21% to 3.87%. The ratio of the linearization remainders at 0.1 and 0.05 is 4.10–4.23, consistent with a quadratic remainder over these tested steps. At 0.5, discrepancies grow to 13.66–43.86%. The formula has a verified local use, and the data reject extending its small-perturbation accuracy to a large flux substitution.

The outer-coordinate region contributes 98.39–99.55% of S_F at $N = 160$. The late-half contribution varies much more, from 23.13% in S6 to 74.91% in S5. A long drying tail is not enough to identify when

Table 3 Trajectory-propagated predictions and first-order attribution at $N = 160$

Case	τ_0 (h)	S_F (h)	$\Delta\tau_{1.1}$ (h)	Difference (%)	Outer (%)	Late (%)
S1	57.170087	0.285929	0.030371	5.86	99.31	60.80
S2	55.507605	0.454686	0.049342	7.85	99.55	53.75
S3	120.266938	0.634562	0.067472	5.95	99.39	70.01
S4	35.859412	0.281059	0.030483	7.80	99.51	44.54
S5	34.600269	0.063351	0.006494	2.44	98.39	74.91
S6	109.045729	1.150561	0.124817	7.82	99.55	23.13

The predicted shift is $0.1S_F$; Difference is $|0.1S_F - \Delta\tau_{1.1}|/|\Delta\tau_{1.1}|$. Outer and Late are signed fractions of the first-order derivative, not fractions of the finite nonlinear shift. Outer means $x_{i+1/2} \geq 0.9$; Late means $t \geq \tau_0/2$. These are numerical flux-perturbation responses, not observed drying errors.




Fig. 1 Prediction discrepancy as the flux-perturbation weight increases, and spatial versus temporal fractions of the first-order endpoint response. Fractions are evaluated on the unperturbed trajectories; the two bars are overlapping classifications

a flux perturbation acquires its endpoint influence. Spatial concentration is consistent across these cases, while temporal attribution changes with geometry and operating assumptions.

The $N = 320, 640$ controls preserve the 10% descriptive level for $\epsilon = 0.1$. The outer fraction remains above 99.5% at $N = 640$ in S1, S2, and S6. The derivative magnitude itself decreases with refinement because the discrete H-minus-K forcing changes; it is not a grid-independent physical coefficient. The tighter S2 settings change the predicted derivative by less than 2.5×10^{-8} h per unit weight. The new $N = 640$ unperturbed S1 and S2 endpoints agree with the independent quadrature references within 3.1×10^{-8} h. These controls support the implementation and local transfer claim within the tested family.

4.3 Surface resolution as a tested consequence

The graded-grid results test an action suggested by the attribution (Table 4). At 160 intervals, the H-minus-K difference is 0.007212 h for S1, 0.010846 h for S2, and 0.026254 h for S6. Doubling to 320 intervals reduces these differences to 0.001804, 0.002744, and 0.006592 h. Both flux endpoints move toward the independently checked fine-grid K results as the graded mesh is refined.

The independent graded-grid manufactured tests give maximum-error orders of 1.991–2.000 across the three geometries. At $N = 160$, the maximum errors are 5.53×10^{-7} , 8.25×10^{-7} , and 1.06×10^{-6} for the

slab, cylinder, and sphere. This check supports the changed geometric assembly rather than relying only on agreement between two drying fluxes.

Table 4 Common-face-temperature endpoints on the surface-graded grid

Case	N	K (h)	H (h)	H–K (h)
S1	80	57.170221	57.198997	0.028775
S1	160	57.169589	57.176800	0.007212
S1	320	57.169430	57.171234	0.001804
S2	80	55.507678	55.560546	0.052868
S2	160	55.507077	55.517924	0.010846
S2	320	55.506927	55.509671	0.002744
S6	80	109.047042	109.150002	0.102960
S6	160	109.044262	109.070516	0.026254
S6	320	109.043566	109.050158	0.006592

The grid uses $x_i = 1 - (1 - i/N)^2$ for both fluxes, with identical coupled equations and BDF tolerances. This mesh experiment was added after the attribution results; it is not a prespecified holdout. Uniform 160-interval H gives 60.215699 h for S1; S2 and S6 do not cross by 400 h in the audited runs.

On the uniform 160-interval grid, H gives 60.215699 h for S1, while the K endpoint is 57.170087 h. The uniform S2 and S6 H integrations reach 400 h without crossing the threshold. They are right-censored calculations, not converged endpoints or failed Newton iterations. Their final states are retained in separate audits. Surface grading improves the tested result at the same node count, although this is not a comparison of equal wall-clock costs or a proof of an optimal mesh.

The node-temperature HF controls also complete on the graded grids. Their 640-interval endpoints are 57.169842 h for S1 and 55.507577 h for S2. Tightening the S2 tolerances and halving the maximum step changes its endpoint by $1.87\text{e-}09$ h. These controls show that the earlier uniform-grid HF difficulty need not persist with resolved surface transport. They support the stated graded-grid sequence, not a retroactive error estimate for the unresolved uniform-grid endpoints.

These geometric transfers support surface grading within the tested exponential family. The next comparison changes the moisture factor itself, where the spatial interpretation may differ. Moving domains remain outside this resolution experiment.

4.4 Constitutive dependence of the spatial response

All 44 cross-law trajectories and 16 control trajectories reach their thresholds with positive moisture fields. The centre remains the maximum. The constant moisture factor gives identical K and H endpoints: 97.546182 h in the cylinder and 60.747816 h in the sphere, with zero tangent response. Relative prediction discrepancies and attribution fractions are undefined for this null test.

Figure 2 relates the constitutive factors to their signed spatial and temporal contributions.

Across the eight nonconstant cases, the largest prediction discrepancy at $\epsilon = 0.1$ is approximately 0.113% (Table 5). The observed shifts are only 0.196–6.410 s. Tightening the cylindrical integration settings changes the paired shift by at most 0.0021%. Doubling the grid reduces the paired shifts by factors of 3.978–4.000, consistent with a difference between converging discrete fluxes. The largest discrepancy between central endpoint differences and the tangent derivative is approximately 0.0046%. At weight 0.5, the largest discrepancy is 0.57%; this better performance than in the exponential cases is observed here, not guaranteed for other constitutive functions.

Table 5 Cross-law endpoint shifts and signed first-order attribution at $N = 160$

Case	τ_0 (h)	$\Delta\tau_{.1}$ (s)	Difference (%)	Outer (%)	Late (%)
P2-1	51.940257	6.4096	0.112	91.38	37.40
P2-2	33.828991	4.2086	0.113	91.12	30.94
F1-1	53.465031	0.7539	0.006	66.10	10.87
F1-2	34.708505	0.5341	0.007	67.27	10.38
F10-1	67.579419	0.2757	0.006	58.01	-3.79
F10-2	42.928149	0.1961	0.003	61.04	-3.44
S-1	82.101144	1.1571	0.006	78.84	41.82
S-2	51.486677	0.7673	0.005	77.59	38.72

The final digit of each case denotes geometry $m = 1$ (cylinder) or 2 (sphere). Difference compares $0.1S_F$ with the directly integrated shift. Fractions use the same signed derivative and partitions as Table 3. A negative Late fraction denotes cancellation of the earlier response. The constant-factor controls give identical K and H endpoints, so their relative discrepancies and fractions are undefined.




Fig. 2 Normalized moisture factors and cross-law signed attribution. P2 denotes a quadratic factor; F1 and F10 add positive floors; S denotes saturation. Geometry digits are defined in Table 5. Outer and Late are overlapping classifications, and negative fractions indicate cancellation

The spatial attribution does not transfer unchanged. The outer-region share falls to 91.1–91.4% for P2, 66.1–67.3% for F1, 58.0–61.0% for F10, and 77.6–78.8% for S. The F10 late-half shares are -3.79% and -3.44% : late contributions oppose part of the early response. The positive-floor sign change in Eq. (20) permits this contrast, while the trajectory calculation measures its net effect. These results reject treating the exponential family’s greater-than-98% surface localization as a constitutively general rule. These perturbation tests alone do not compare mesh costs; Sect. 3.6 defines that separate experiment for selected cases.

The maximum relative discrepancy between analytic face integrals and independent quadrature is approximately $8.8\text{e-}14$ on the tested intervals. Maximum full-Jacobian and moisture-block directional discrepancies are approximately $2.2\text{e-}10$ and $1.2\text{e-}10$, respectively. The final manufactured-solution refinement orders are 1.994–1.999. These controls verify the face integrals, Jacobians and spatial implementation used in the cross-law comparison.

4.5 Spatial influence and mesh cost

The forward and adjoint calculations evaluate the same derivative by different propagation directions. Their largest relative difference in the six matched 160- and 320-interval checks is approximately $3.7\text{e-}07$. The face-resolved sums also recover the earlier four partitions. The coefficient-kernel density changes by 0.158–0.166% from 320 to 640 intervals, while the largest discrepancy against the central 1 ± 0.01 diffusivity perturbation is approximately 0.0103%. These checks support the numerical influence representation independently of a chosen spatial cutoff.

Figure 3 distinguishes the cumulative coefficient response from the activity of the discrete flux defect.

Table 6 Spatial kernel checks and unpartitioned flux-defect attribution

Case	S_D (h)	Kernel change (%)	x_{10}	$\sum b_f/ S_F $
E45	-50.4797	0.159	0.99305	1.0000
E50	-48.7562	0.158	0.99583	1.0000
F10	-59.3536	0.165	0.74634	1.2190

S_D is the derivative for uniform multiplicative diffusivity perturbation at $N = 640$. Kernel change is the relative L^1 change of its piecewise density from 320 to 640 intervals. For the separate H-minus-K defect at $N = 320$, x_{10} is the coordinate below which 10% of time-absolute activity accumulates; the final column measures cancellation. These are different perturbation questions.




Fig. 3 Cumulative spatial influence for a fixed diffusivity perturbation and activity of the H-minus-K defect. In (a), dashed and solid lines use 320 and 640 intervals. In (b), they use 80 and 320 intervals. Normalization concerns the respective perturbation; the two panels must not be interpreted as the same error distribution

The constitutive kernel and the flux-defect distribution answer different questions. At 320 intervals, 90% of the time-absolute H-minus-K activity lies beyond $x = 0.99305$ in E45, beyond 0.99583 in E50, and beyond 0.74634 in F10. The exponential locations still move with grid refinement. Their scaled H-minus-K response has not reached a grid-independent leading coefficient over 80–320 intervals. The fixed coefficient kernel has a much smaller refinement change. The two representations distinguish constitutive sensitivity from the spatial activity of a particular discrete flux change. F10 has $\sum b_f/|S_F| \approx 1.219$, which exposes cancellation even without the earlier inner–outer or early–late split.

Figure 4 compares reference differences against the cumulative recorded charges for each strategy.

Of the 72 timed attempts, 62 complete, 8 reach 400 h without crossing, and 2 reach the 45-second CPU limit. There are 0 other recorded failures. The E50 goal grid at 320 intervals reaches the CPU limit in both

Table 7 Results within 60 recorded CPU seconds including measured setup and preceding attempts

Case	Strategy	N (runs 1/2)	Charged cost (s)	$ \tau_H - \tau_{K,2560} $ (h)
E45	Uniform	320/320	20.8/19.3	0.395426/0.395426
E45	Surface	320/320	21.4/24.1	0.001854/0.001854
E45	Goal	320/320	49.8/49.7	0.002686/0.002686
E45	Gradient	320/320	27.2/27.2	0.019971/0.019971
E50	Uniform	—/—	37.8/39.1	—/—
E50	Surface	320/320	39.5/23.3	0.002791/0.002791
E50	Goal	160/160	52.0/56.1	2.097475/2.097475
E50	Gradient	320/320	39.2/38.7	9.178915/9.178915
F10	Uniform	320/320	21.3/20.3	0.000389/0.000389
F10	Surface	320/320	23.1/21.9	0.000170/0.000170
F10	Goal	320/320	54.2/53.8	0.000032/0.000032
F10	Gradient	320/320	27.4/28.3	0.000116/0.000116

The slash separates the two final-solve repetitions, which are not confidence limits. A dash means no completed endpoint fits the budget. Precomputation is charged once per strategy; failed and noncrossing preceding attempts also consume cost. Values compare H with a numerical K reference, not an exact or experimental endpoint. Both 30- and 120-second accounts are retained in the working records.

Charges follow Sect. 3.6: gradient-pilot sampling and integration, pilot-output writing and candidate-model construction are excluded; final-result writing is included. These are recorded-stage charges, not end-to-end runtimes.




Fig. 4 Work–precision comparison using recorded CPU charges for setup and preceding attempts, with the timing boundaries in Sect. 3.6. Curves show candidates completed in both repetitions; missing endpoints are not assigned artificial error values. Dotted lines show the 1280-to-2560 K reference change, not certified error bounds. Budget outcomes, including unavailable results, are reported in Table 7

repetitions. Its 160-interval endpoint is about 57.604354 h, still far from the fine K reference near 55.506880 h. The surface-graded 320-interval run gives 55.509671 h. The comparison identifies surface grading as the more reliable tested choice for this E50 flux substitution.

The preferred mesh depends on the constitutive case and budget. At 320 intervals, F10's goal-grid endpoint is 67.578652 h, closer to the fine reference 67.578619 h than the surface-graded value 67.578790 h. The goal strategy also pays for an adjoint and has no completed result within the 30-second budget. The surface strategy has a smaller reference difference and lower recorded charge than the tested goal strategy for E45. The setup cost explains why the smallest finest-grid discrepancy need not give the preferred result at a smaller budget. The two timing repetitions can select different affordable grid levels; their separate outcomes are retained in Table 7.

The K references change by $8.7\text{e-}06$ – $9.4\text{e-}06$ h from 1280 to 2560 intervals. Their uncertainty matters when interpreting the smallest F10 differences. The last manufactured-field refinement orders on the new meshes range from 1.991 to 2.013. These tests check the mesh assembly and resolution of the numerical references.

A follow-up diagnostic evaluates the harmonic face derivative on K trajectories on the goal and surface meshes. For a general increasing moisture factor, the right-face derivative has the sign of $1 - \eta_f$, where $\eta_f = [f_L/(f_L + f_R)](f'_R/f_R)(C_L - C_R)$; Eq. (14) is its exponential specialization. In E50, the largest η_f is 3.149 on the goal grid and 0.509 on the surface grid. This diagnostic was added after the timed experiment and did not retune its meshes. It checks the unperturbed K trajectory, not the rejected H states, and supplies a local consistency warning rather than proof of the full coupled failure mechanism.

4.6 Derivative diagnostics and guarded refinement

The guard protocol was fixed after the earlier E50 failure and before the E55H and F1SH outcomes were available. Its distinction between development and prospective cases is retained. Of its 36 timed H attempts, 30 complete, 2 reach the CPU limit, 4 reach 400 h without crossing, and 0 have other recorded failures. Only 2 of the six constructed guard candidates satisfy the sampled K limit of 0.8.

Table 8 Finest nominal-level results of the fixed derivative-guard protocol

Case	Mesh	N	Ends	$ \tau_H - \tau_{K,2560} $ (h)	η_K	η_H
E50	Surface	320	2/2	0.002791	–	0.546
E50	Goal	320	2/2	0.127726	–	9.254
E50	Guard attempt	322	2/2	0.005245	1.494	2.262
E55H	Surface	320	2/2	0.004152	–	6.078
E55H	Goal	320	0/2	–	–	71.483
E55H	Guard attempt	322	2/2	12.107774 ^a	5.310	41.457
F1SH	Surface	320	2/2	0.000204	–	0.005
F1SH	Goal	320	2/2	0.000145	–	0.007
F1SH	Guard attempt	320	2/2	0.000145	0.007	0.007

The nominal level is $B = 320$; N includes inserted nodes. Ends counts completed repetitions. Indicator columns give sampled maxima of both face indicators; η_K is the final guard audit, and η_H is the post-run H audit. A dash denotes an unavailable or inapplicable value. An endpoint on an unresolved guard grid is not certified as satisfying the constraint. Uncompleted attempts have no endpoint difference.

^a The E55H guard difference uses the default time settings and remains sensitive to time integration. The stricter control changes its endpoint by 1.436622 h (Sect. 4.6); this entry is not a time-converged error estimate.

The coarse pilot can substantially understate the indicator on its redistributed mesh. In E50, the interpolated pilot gives a finest-level maximum near 0.133, while the first fresh K trajectory gives about 3.149. The two allowed subsequent bisections reduce that audit maximum but do not satisfy the margin. E55H is more difficult. F1SH already satisfies the limit, so its guarded and original goal grids coincide; the extra K audit adds cost without changing the endpoint. No post-result adjustment was made to the margin, audit count, or holdout cases.

The finest E50 guard attempt has 322 intervals and a final sampled K indicator of 1.494. Its completed H endpoint is 55.512125 h, with reference difference 0.005245 h and sampled H indicator 2.262. The finest E55H guard attempt has 322 intervals and a final sampled K indicator of 5.310. Its completed H endpoint is 66.325618 h, with reference difference 12.107774 h and sampled H indicator 41.457. These outcomes distinguish improvement of an endpoint from enforcement of the sufficient condition. A derivative violation also does not by itself prove that a completed endpoint is inaccurate.

The two cost protocols retain their separately generated mesh files and timing records. A completed E50 goal run in Table 8 does not supersede the timed-out candidate in the first protocol. A post-result comparison identifies sensitivity in the unguarded E50 goal calculation. The old and recomputed pilot activities are identical, while mathematically equivalent cumulative-grid sampling changes the 160-level nodes by at most $1.7\text{e-}16$. The associated H endpoints are 57.604354 and 57.642945 h. This observed sensitivity cautions against treating completion or identical repeated runs as convergence. It is a historical implementation comparison, not an additional prespecified holdout or evidence of a unique failure mechanism.

Table 9 Outcomes within 120 recorded CPU seconds for the prospective protocol

Case	Mesh	N (runs 1/2)	Charged CPU (s)	Reference difference (h)
E50	Surface	320/320	15.6/15.1	0.002791
E50	Goal	320/320	102.2/101.6	0.127726
E50	Guard attempt	322/322	102.2/102.6	0.005245
E55H	Surface	320/320	53.5/53.5	0.004152
E55H	Goal	—/—	69.6/69.2	—
E55H	Guard attempt	322/322	110.5/112.6	12.107774 ^a
F1SH	Surface	320/320	15.3/14.6	0.000204
F1SH	Goal	320/320	39.5/39.7	0.000145
F1SH	Guard attempt	320/320	53.2/53.2	0.000145

Pilot, adjoint, pilot sampling and storage costs are charged once for the goal and guard strategies; construction, K re-solves and final H attempts are charged at both nominal levels in order. A completed candidate may still fail the derivative constraint, as shown in Table 8. The two final-solve repetitions are separate timings, not confidence bounds. Records also retain the 60- and 180-second accounts. Prepared-grid writing, final grid loading, candidate-model construction, H diagnostics and final-result writing are excluded; Sect. 3.7 defines this protocol's timing boundary.

^a The E55H guard difference uses the default time settings and remains sensitive to time integration. The stricter control changes its endpoint by 1.436622 h (Sect. 4.6); this entry is not a time-converged error estimate.

The guard must pay for every K re-solve. An apparently small insertion of two nodes can cost much more than generating a simple graded mesh. Budget availability, completion and constraint satisfaction are reported separately. In these tests, the additional audits improve some endpoints but leave the exponential guard candidates above the prescribed margin.

The finer-time control exposes an unresolved E55H guard endpoint: tightening relative tolerance from 10^{-9} to 10^{-10} and reducing the maximum step from 120 to 60 s changes it from 66.325618 to 67.762240 h, a difference of 1.436622 h. The other five controls change by at most $1.7\text{e-}09$ h. The E55H guard discrepancy

in Table 8 is a value at the stated settings, not a converged estimate. The repeated default runs agree with each other but do not resolve this integration sensitivity.

The K reference changes by $5.7\text{e-}06$ – $9.5\text{e-}06$ h between 1280 and 2560 intervals. Independent finite differences check both face derivatives with a largest scaled discrepancy of $2.4\text{e-}09$. The twelve manufactured-field calculations cover all six guarded grids and their complete interval bisections; the paired orders range from 1.996 to 2.000. These checks support the grid implementation; the preceding time controls identify E55H as the unresolved endpoint.

5 Conservative closure and latent-heat response

A separate synthetic extension gives the water inventory an explicit dry-solid basis and includes the enthalpy carried by moisture. The domain is rigid, dry-solid density ρ_d is constant, and water is transported as condensed moisture until it evaporates at the surface. With $T_0^{\text{ref}} = 0^\circ\text{C}$, define $\theta = T - T_0^{\text{ref}}$ and

$$\rho_w = \rho_d C, \quad \mathbf{j}_w = -\rho_d D \nabla C, \quad \partial_t(\rho_d C) + \nabla \cdot \mathbf{j}_w = 0, \quad (26)$$

$$H = \rho_d(c_d + c_w C)\theta, \quad H_t + \nabla \cdot [-k \nabla T + c_w \theta \mathbf{j}_w] = 0. \quad (27)$$

Here ρ_w is water mass per total specimen volume. The energy flux includes the sensible enthalpy carried by moisture, giving a consistent balance with surface evaporation (Datta 2007).

At the surface, $j_s = \mathbf{j}_w \cdot \mathbf{n} = \rho_d h_m(C_s - C_\infty)$ and

$$-k \partial_n T = h(T_s - T_\infty) + L_v j_s, \quad \frac{d}{dt} \int_\Omega H dV = - \int_{\partial\Omega} [h(T_s - T_\infty) + (L_v + c_w \theta_s) j_s] dA. \quad (28)$$

The first condition balances conduction, convective heating, and surface evaporation. The second is the total sensible-energy inventory balance, including the enthalpy exported with water. Dry-solid mass is fixed by construction, wet density is $\rho_d(1 + C)$, and cumulative boundary loss must equal the decrease in water inventory.

The extension assumes $\rho_d = 500 \text{ kg/m}^3$, $c_d = 1450$ and $c_w = 4180 \text{ J/(kg K)}$. It retains the base radius, initial fields, conductivity function, exponential diffusivity, and empirical exchange history. The pair $L_v = 0$ and $2.4 \times 10^6 \text{ J/kg}$ isolates latent cooling within this new closure. These constants are specified synthetic inputs, not recovered specimen properties. The $L_v = 0$ run is a numerical control rather than a physically complete evaporation model. Its different heat capacity and enthalpy transport mean that comparison with the old benchmark cannot isolate latent heat alone.

The solver advances dry-basis moisture and scaled sensible-energy density, with opposite contributions from each internal face. Integrated boundary losses check whole-domain inventories. Runs at 80, 160 and 320 intervals use relative tolerance 10^{-9} and maximum step 120 s; the finest pair is repeated at 10^{-10} and 60 s. Initial finite-difference-Jacobian runs emitted perturbation-factor warnings for accumulator columns. Eight analytic-Jacobian repeats retain the original evidence and agree in endpoint within 5.7×10^{-9} h; directional differences check the new derivatives.

At 320 intervals, the paired endpoints are 57.235596 h without latent heat and 62.825585 h with it, a change of 5.589989 h or 9.767%. The latent endpoint changes by 0.000652 h from 160 to 320 intervals and by at most 1.9×10^{-9} h under the tighter time settings. The minimum nodal temperature is about 2.994°C during the initial transient, while the final centre and surface temperatures are both about 50.062°C . A small final temperature difference does not make early latent cooling negligible. Maximum relative water- and energy-inventory residuals are about 6.6×10^{-15} and 7.2×10^{-15} . The residuals verify the discrete water and energy balances.

The physical interpretation concerns a rigid synthetic specimen with surface evaporation and empirical moisture exchange. Material identity and independent observations are unavailable, so the results have not

been physically calibrated or validated. Shrinkage, internal evaporation and sorption-based exchange remain outside the solved model; humidity conversion would require a defined mass basis and pressure (ASHRAE 2021). The spruce-wood dataset described by Richter and Staněk (2024) concerns different conditions, and its measurement files were not obtained.

The numerical comparisons cover three selected cases per mesh protocol and two final-solve repetitions on one machine. The K reference shares the radial spatial balance, and the uniform coefficient perturbation checks the integrated kernel response rather than its full spatial shape. The MHM comparison concerns the Cartesian cubic equation. Mesh rankings apply to the reported laws, grids and budgets; the E55H guard endpoint remains time-sensitive.

6 Conclusions

This study connects interface-flux changes to moisture-threshold times and uses that connection to assess mesh choices. Tangent–adjoint agreement and direct perturbation checks support the endpoint derivatives. The face-resolved analysis identifies surface concentration in the exponential cases and signed cancellation in other constitutive laws, giving a spatial explanation for differences that a single endpoint value cannot show.

The recorded-cost comparisons turn this attribution into a testable mesh decision. Surface grading has lower recorded charges and smaller reference differences than the tested goal strategy for E45, while the F10 goal grid gives the smaller finest-grid difference. The derivative checks reveal when a coarse pilot understates local face sensitivity and separate endpoint improvement from constraint satisfaction. These results make accuracy, completion and computational cost visible within the same comparison; the E55H guard endpoint still requires tighter time-resolution assessment.

The conservative synthetic control makes water and energy inventories explicit and quantifies a 9.77% increase in the paired endpoint from latent cooling. Together, the calculations provide an interpretable procedure for locating influential flux changes, checking their constitutive dependence and evaluating the cost of a numerical response. The findings concern the tested synthetic models and support numerical assessment rather than material-specific prediction.

Statements and Declarations

Funding. This research received no funding.

Competing interests. The authors declare no financial or non-financial competing interests related to this work.

Data and code availability. The original benchmark files are available from the official contest release cited in CUMCM (2026). The code and generated numerical results supporting this study, including saved grids and unsuccessful-attempt records, are available from the corresponding author on reasonable request. The study generated no experimental observations.

Author contributions. SY developed the algorithms, conducted the numerical experiments, produced the core results, and wrote the initial manuscript. YC assisted with the simulations, conducted comparative numerical experiments, and revised the manuscript. JC reviewed the literature, proofread the manuscript, and contributed to discussions of the study. All authors reviewed and approved the final manuscript.

References

- Adrover A, Brasiello A, Ponso G (2019) A moving boundary model for food isothermal drying and shrinkage: General setting. *Journal of Food Engineering* 244:178–191. <https://doi.org/10.1016/j.jfoodeng.2018.09.018>
- ASHRAE (2021) Psychrometrics. In: *ASHRAE Handbook—Fundamentals*, SI edn, Chapter 1. ASHRAE, Atlanta. https://handbook.ashrae.org/Handbooks/F21/SI/F21_Ch01/F21_Ch01_si.aspx. Accessed 27 September 2026
- Becker R, Rannacher R (2001) An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica* 10:1–102. <https://doi.org/10.1017/S0962492901000010>
- Çelik E, Parlak N, Meral ZK, Çay Y (2026) Experimental investigation of convective drying behavior of in-shell hazelnuts: moisture diffusivity, transport phenomena, and pressure drop. *Heat and Mass Transfer* 62:99. <https://doi.org/10.1007/s00231-026-03715-3>
- Chaudhry JH, Estep D, Stevens Z, Tavener SJ (2021) Error estimation and uncertainty quantification for first time to a threshold value. *BIT Numerical Mathematics* 61:275–307. <https://doi.org/10.1007/s10543-020-00825-0>
- Chaudhry JH, Estep D, Giannini T, Stevens Z, Tavener SJ (2023) Error estimation for the time to a threshold value in evolutionary partial differential equations. *BIT Numerical Mathematics* 63:12. <https://doi.org/10.1007/s10543-023-00947-1>
- Chen K (1994) Error equidistribution and mesh adaptation. *SIAM Journal on Scientific Computing* 15:798–818. <https://doi.org/10.1137/0915050>
- China Undergraduate Mathematical Contest in Modeling (2026) Problem A: Drying of medicinal plant materials and accompanying data files [in Chinese; translated title]. Official 2026 contest problem release, 10 September 2026. https://www.mcm.edu.cn/html_cn/node/27b6e148f8113f09b0269f64a02629fb.html. Accessed 27 September 2026
- Crank J (1975) *The Mathematics of Diffusion*, 2nd edn. Clarendon Press, Oxford
- Datta AK (2007) Porous media approaches to studying simultaneous heat and mass transfer in food processes. I: Problem formulations. *Journal of Food Engineering* 80:80–95. <https://doi.org/10.1016/j.jfoodeng.2006.05.013>
- Maddix DC, Sampaio L, Gerritsen M (2018) Numerical artifacts in the Generalized Porous Medium Equation: Why harmonic averaging itself is not to blame. *Journal of Computational Physics* 361:280–298. <https://doi.org/10.1016/j.jcp.2018.02.010>
- Nwakuba N, Ezeanya N, Horsfall IT, Okafor V, Ononogbo C, Ndukwu M, Simo-Tagne M, Asoegwu S (2025) Heat and moisture transport in okra cylinders with shrinkage effects under solar drying: a multiphysics-based simulation approach. *Sustainable Food Technology* 3:520–536. <https://doi.org/10.1039/D4FB00343H>
- Patankar SV (1980) *Numerical Heat Transfer and Fluid Flow*. Hemisphere Publishing Corporation, New York
- Prakash K V, Natarajan G (2025) Revisiting numerical artifacts in the generalized porous medium equation with continuous coefficients: Does averaging really matter? *Journal of Computational Physics* 537:114101. <https://doi.org/10.1016/j.jcp.2025.114101>
- Richter J, Staněk K (2024) Experimental data_2021_Dynamic moisture transport in wood. Mendeley Data, version 1. <https://doi.org/10.17632/dbfdm6sy6x.1>
- Salari K, Knupp P (2000) Code verification by the method of manufactured solutions. Report SAND2000-1444, Sandia National Laboratories, Albuquerque. <https://doi.org/10.2172/759450>
- Simal S, Rosselló C, Berna A, Mulet A (1998) Drying of shrinking cylinder-shaped bodies. *Journal of Food Engineering* 37:423–435. [https://doi.org/10.1016/S0260-8774\(98\)00095-8](https://doi.org/10.1016/S0260-8774(98)00095-8)
- Takatori N (2026) Comparison of interface moisture conductivity averaging methods for moisture-induced deterioration prediction of outdoor cultural properties. *Journal of Building Physics*. Advance online publication. <https://doi.org/10.1177/17442591261447478>
- Vaz M Jr, Zdanski PSB (2024) Optimisation-based strategies to identification of material parameters of hygro-thermo-mechanical problems. *Discover Mechanical Engineering* 3:8. <https://doi.org/10.1007/s44245-024-00038-7>