Perfecting Steak Searing with Computation

Heat, Moisture, Flipping, and the Maillard Reaction

Bob Zhao
The Webb Schools

Abstract

Getting a good crust without turning the middle gray sounds like a cooking problem. It is also a heat-and-mass-transfer problem with a stopwatch. I built a one-dimensional finite-volume model that follows heat conduction, moisture movement, surface evaporation, and a relative Maillard exposure index while a steak is flipped. The model predicts temperature and vapor loss, but the browning number is only a comparison tool; it is not an exact flavor, color, or chemical measurement. Against published tests for a 19 mm steak on a 215 °C pan, the predicted center temperatures were 30.2 °C, 43.4 °C, and 55.8 °C at 180, 300, and 420 s. The RMSE was 0.81 °C, and every prediction landed inside the reported experimental standard deviation. Numerical checks also showed second-order combined space–time refinement and very small discrete balance errors. The sweeps gave a few useful patterns. Hotter pans and a drier starting surface raised modeled exposure. A perfectly even flip shared the browning more fairly between the two faces, but it did not maximize the average. Thicker steaks collected more surface exposure before reaching the same center target because their centers needed longer to warm. Contact conductance made the biggest difference among the parameter changes I tested. In short, the model turns familiar cooking advice into testable ideas—without pretending that one equation can taste the steak.

Keywords: computational food science; Maillard reaction; steak; heat transfer; finite-volume model

1 Background

1.1 A kitchen question becomes a transport problem

The usual advice for getting a good steak crust is simple: use a hot pan and keep the surface dry. That advice works, but it skips most of the physics. The meat does not instantly jump to the pan temperature. Heat has to cross a messy contact surface filled with tiny gaps, fat, and sometimes water. Meanwhile, heat travels into the steak much faster than liquid water moves through it. Evaporation also steals energy from the surface. A wet surface can therefore stay much cooler than the pan while the center keeps warming. In short, the pan temperature matters, but it is not the whole story.

The cooking goals also compete with one another. Browning happens at the surface, while doneness depends on the temperature inside. Food safety depends on the full time–temperature

history at the relevant cold point, along with the condition of the meat and where contamination could be located. Juiciness involves evaporation as well as liquid pushed out when proteins contract. Evenness depends on when the steak is flipped. A change that improves the crust may overcook the center, so slogans such as “higher heat is better” or “flip once” need some conditions attached. At minimum, a fair comparison should account for center temperature, cooking time, thickness, and starting surface moisture.

This is where computational science is useful. A numerical model can keep track of surface water activity, temperature gradients, evaporation, and accumulated heating all at once. Those quantities would be awkward to measure together over a frying pan. The model also lets us change one condition at a time, which is harder to do with real steaks unless we have a very large grocery budget. Pan searing is still a tricky case because every flip suddenly swaps the boundary conditions on both sides.

1.2 What browning represents

The Maillard reaction is not a single neat reaction. It is a whole network that begins with reducing sugars and amino compounds, then branches through rearrangements, fragmentation, Strecker chemistry, and polymerization [1]. Some products help create roasted aromas. Others add brown color, and some would never show up in a simple color measurement. One number therefore cannot capture every chemical and sensory change in a crust.

Temperature speeds up this network and can be approximated with Arrhenius behavior over a limited range [1]. The approximation is useful, but it is not universal. A fitted activation energy changes with the food, reactants, pH, measured endpoint, and temperature range. Moisture complicates the picture again. Too much available water dilutes reactants and increases evaporative cooling. Too little water limits molecular movement. Browning is therefore often fastest somewhere in the middle, although the best water activity depends on the food system [1].

Color needs its own warning label. Inside a steak, losing redness mainly reflects changes in myoglobin and proteins, not surface Maillard conversion. A cooking model for chicken handled interior lightening and high-temperature surface browning as separate processes [2]. We follow that basic distinction by giving the surface its own index, but we do not borrow the fitted chicken color constants for beef chemistry.

Real steak experiments show why that distinction matters. Yoo and coauthors compared an oven-only steak with one finished by a high-temperature pan sear, while bringing both to the same core endpoint [3]. Searing increased measurements connected with roasted flavor and Maillard intermediates. It did not support the familiar claim that a crust “seals in” the water. The crust is doing flavor chemistry, not building a waterproof jacket.

1.3 Existing cooking models and the remaining gap

Cooking models range from simple heat-conduction equations to large three-dimensional simulations that include deformation and moving liquid. The contact with the pan deserves special attention because its heat flux changes with time and is not automatically set by the pan temperature [4]. More detailed meat models can also include protein denaturation, shrinkage, evaporation, and dripping [5]. Cooking is clearly more than heat diffusing through a motionless block.

The closest benchmark for this project is the three-dimensional steak model by Moya and coauthors [5]. Their model combined heat transfer, moisture transport, shrinkage, water-holding capac-

ity, dripping, and evaporation. They then compared center temperature and total mass loss with experiments on double-sided pan cooking. For a 19 mm steak starting at 20 °C on a 215 °C pan, they fitted an effective contact conductance of 120 W/(m^2 K) and a moisture diffusivity of 1×10^{-9} m^2/s. That makes their work a strong thermal benchmark. However, it does not convert the temperature and moisture fields into a direct comparison of Maillard exposure.

Our model takes a deliberately smaller bite out of the problem. A one-dimensional slab runs quickly enough for grid tests, parameter sweeps, matched-endpoint comparisons, and structural ablations. It still includes the main competition among contact heating, internal conduction, evaporation, and water-activity-dependent reaction exposure. It does not simulate shrinkage, pressure-driven liquid flow, or the full Maillard network. The goal is not a perfect virtual steak. It is a clear computational experiment that shows which conclusions remain when the assumptions change.

1.4 Research questions and contributions

We organized the project around four questions:

1. Can a discrete-balance-closing one-dimensional model reproduce published steak center-temperature data using physically interpretable parameters?
2. How do pan temperature and initial surface moisture interact to control a relative Maillard exposure index?
3. How do single-flip timing and repeated flipping redistribute exposure between the two surfaces?
4. At a matched center threshold, how do steak thickness and pan temperature change the tradeoff between modeled surface exposure, time, and vapor loss?

The main contribution is how these pieces are connected and compared. Heat, moisture, and browning use the same spatial grid. A flip changes the physics at both surfaces, and protocols are compared both after the same amount of time and at the same core-temperature threshold. We also report numerical verification, a published experimental benchmark, grid refinement, one-at-a-time sensitivity, and structural ablation. These checks keep three different claims from getting mixed together: the thermal prediction is benchmarked, the vapor-loss prediction is only partial, and the Maillard index is phenomenological.

2 Model Concept

2.1 One steak, two surfaces with different jobs

We represent the steak as a flat slab of thickness L, with coordinate $z \in [0, L]$. The other two dimensions are treated as large compared with L, so heat and moisture move only through the thickness. This leaves out edge heating but keeps the two broad faces separate. At any moment, one face touches the pan while the other exchanges heat and vapor with the air. Flipping simply swaps those boundary conditions. The mesh does not turn over, and the steak does not forget its previous temperature or moisture.

Figure 1 shows the setup. The two main state variables are temperature $T(z, t)$ and moisture ratio $X(z, t)$, measured as kilograms of water per kilogram of dry solid. A sorption isotherm converts

X into water activity $a_w(z,t)$. We calculate a local kinetic rate from T and a_w, but the reported surface index uses the average state across an outer 0.5 mm layer. A real crust has some thickness, so treating one infinitely thin point as the whole crust would be hard to defend.

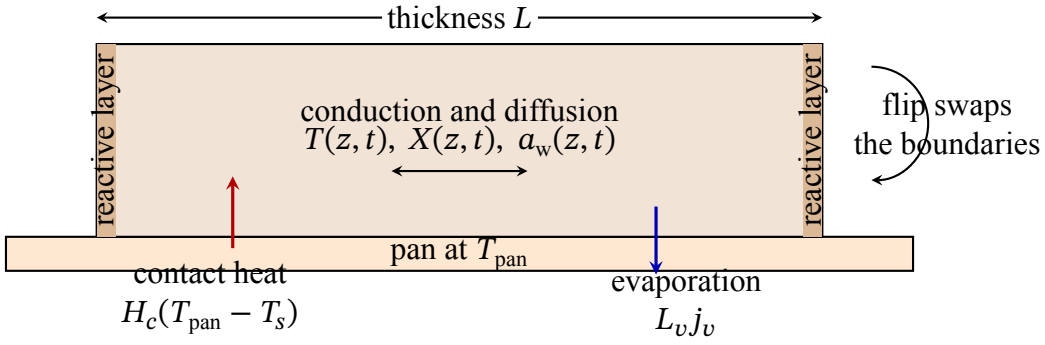

Figure 1: Conceptual model. The full-thickness one-dimensional grid keeps a separate history for each surface. Contact heating, conduction, moisture diffusion, evaporation, and a water-activity-dependent reaction rate are coupled at each time step.

2.2 Heat and moisture run on different clocks

Using representative wet-beef properties, the thermal diffusivity is

$$\alpha = \frac{k}{\rho c_p} = \frac{0.49}{(1070)(3810)} \approx 1.20 \times 10^{-7}\ \mathrm{m^2/s}. \tag{1}$$

For the 19 mm baseline steak, the half-thickness is $L_c = 9.5$ mm. At $t = 300$ s, the thermal Fourier number is $Fo_T = \alpha t / L_c^2 = 0.400$. The contact Biot number is $Bi = H_c L_c / k = 2.33$. Since Bi is much larger than 0.1, the steak cannot reasonably be treated as one uniform temperature. We should expect a strong surface–core gradient.

Moisture follows a much slower clock. With $D = 1 \times 10^{-9}$ m²/s, the mass Fourier number is only $Fo_m = 0.00332$. The ratio $\alpha/D \approx 120$ means heat diffuses about two orders of magnitude faster than the effective Fickian moisture field. This difference explains how a thin surface layer can dry and brown while most of the steak remains water-rich. It also shows why a Fickian model misses part of the total cooking loss: a fuller steak model needs effects such as protein denaturation, shrinkage, and dripping [5].

2.3 A useful, but intentionally modest, Maillard index

Our chemical output measures accumulated exposure, not the concentration of one compound. For each surface s,

$$\Omega_s(t) = \int_0^t r_M\left(\overline{T}_s(\tau), \overline{a}_{w,s}(\tau)\right) d\tau, \qquad B_s(t) = 1 - \exp[-\Omega_s(t)], \tag{2}$$

where the bars represent averages through the outer 0.5 mm. B_s stays between zero and one. It helps rank cases that use the same kinetic assumptions, but it does not convert directly into CIELAB color,

absorbance, flavor intensity, melanoidin concentration, or the amount of one volatile compound. Whenever this paper says "Maillard index," it means Equation (2) and nothing more ambitious.

Two combined metrics summarize what happens on the two faces:

$$\overline{B} = \frac{B_A + B_B}{2}, \qquad U_B = 1 - \frac{|B_A - B_B|}{B_A + B_B + \epsilon}, \tag{3}$$

where ϵ is a small numerical safeguard. $U_B = 1$ means the exposures are equal, while a value near zero means one surface did most of the browning work. Maximizing $\overline{B}$ and maximizing U_B are different goals: the darkest average crust is not always the most even crust.

2.4 Where the model has to stop

The results follow an evidence ladder. First, global discrete-balance closure and an exact heat-equation test check the code. Next, published center-temperature measurements benchmark the thermal part of the model. Parameter sweeps then test how the proposed mechanisms behave. Only the first two levels support quantitative thermal claims. We did not fit r_M to measured crust chemistry, so the browning results only support comparisons such as "case A produces more modeled exposure than case B." They do not predict an exact color or flavor.

The same boundary matters for food safety. The center-temperature targets used later are mathematical stopping points, not serving advice. A nice brown crust does not certify pathogen reduction. The U.S. Department of Agriculture recommendation for whole beef steak is 62.8 °C (145 °F) followed by at least a three-minute rest [6]. None of our quality indices replace a validated time–temperature safety calculation.

3 Methodology

3.1 Heat: from pan to steak

I used constant effective properties to represent wet beef. This is a simplification, since real meat changes as it cooks, but it keeps the first version of the model understandable. With no internal heat generation, the energy equation is

$$\rho c_p \frac{\partial T}{\partial t} = \frac{\partial}{\partial z}\left(k \frac{\partial T}{\partial z}\right). \tag{4}$$

The steak surface is not simply assigned the pan temperature. That would make the math easy, but it would also make the steak heat unrealistically fast. Instead, the contacting face uses an imperfect-contact condition that balances incoming heat, conduction into the next control volume, and evaporative cooling:

$$q''_{\text{in}} = H_{\text{c}}(T_{\text{pan}} - T_{\text{s}}) - L_v j_v. \tag{5}$$

The face exposed to air uses

$$q''_{\text{in}} = h_a(T_\infty - T_{\text{s}}) + \varepsilon\sigma\left(T_{\infty,K}^4 - T_{s,K}^4\right) - L_v j_v. \tag{6}$$

Temperatures with subscript K are absolute. Positive flux points into the steak. These boundary conditions let the surface stay much cooler than the pan when contact is imperfect or evaporation is strong—which is much closer to what happens in a real skillet.

The effective contact coefficient is one of the shakiest inputs. In continuous heat-flux measurements on minced meat, the value dropped from about 500 toward 100 W/(m^2 K) as cooking continued [4]. I therefore treated Moya's value of 120 W/(m^2 K) as a fitted effective baseline, not a magical constant for every steak [5].

3.2 Moisture: where the water goes

The moisture ratio X follows a constant-diffusivity balance,

$$\rho_d \frac{\partial X}{\partial t} = \frac{\partial}{\partial z}\left(\rho_d D \frac{\partial X}{\partial z}\right), \tag{7}$$

where ρ_d is the dry-solid mass per original steak volume. At either surface, water supplied by diffusion is matched to water removed as vapor:

$$j_v = h_m \max(\rho_{v,s} - \rho_{v,\infty}, 0), \qquad \rho_{v,s} = \frac{\min[a_{w,s} p_{sat}(T_s),\ 0.99 P_{atm}]}{R_v T_{s,K}}. \tag{8}$$

Ambient vapor density is calculated for 25 °C air at 50% relative humidity. I used the following saturation-pressure closure so the calculation can be reproduced exactly:

$$p_{sat}(T_C) = \min\left\{101325,\ 610.94 \exp\left[\frac{17.625 T_C^*}{T_C^* + 243.04}\right]\right\} \text{ Pa}, \qquad T_C^* = \min[100, \max(-20, T_C)]. \tag{9}$$

The Magnus-style coefficients are intended for water temperatures up to approximately 50 °C. Using them above that range is an explicit extrapolation, followed by an atmospheric-pressure cap because the model does not resolve an internal steam phase. The same j_v appears in Equations (5), (6), and (8). In plain language, losing water also costs heat, so the energy and moisture calculations cannot ignore each other.

Equation (8) is applied to both the exposed face and the face touching the pan. This is an effective mixed-contact closure: vapor can escape through gaps and edges even while the surface receives contact heat. It does not pretend there is a neat, uniform air layer between the pan and the meat. A real contact surface is messy—solid contact spots, oil, meat juice, and possibly boiling liquid all share the same small area. Separating them would require measurements or a multidimensional contact model.

This formulation counts evaporated water only. It does not count liquid squeezed out by contracting protein, pooled in the pan, or drained from an edge. Those processes can be a large part of total cooking loss. Therefore, the reported "evaporative loss" is not a complete prediction of measured weight loss.

3.3 Turning moisture into water activity

Moisture ratio is not the same thing as water activity, so the model needs a bridge between them. I used the Guggenheim–Anderson–de Boer (GAB) isotherm:

$$X = \frac{X_m C K a_w}{(1 - K a_w)(1 - K a_w + C K a_w)}. \tag{10}$$

The baseline uses $X_m = 0.036$, $C = 97.21$, and $K = 0.99$. Ishikawa and coauthors called the first parameter q_m and fitted the set to room-temperature beef-slice data in a long-duration drying study

[7]. I renamed it X_m to match the moisture-ratio notation. Equation (10) is inverted analytically in every control volume. This fit is a reasonable starting point, but a rapidly heating and denaturing steak surface is not the same as a room-temperature drying experiment. To see how much that choice matters, a structural ablation replaces GAB with the simpler closure $a_{\mathrm{w}} = \max(0, 1 - 0.08/X)$.

3.4 A useful, but not literal, browning score

The local rate in Equation (2) is

$$r_M(T, a_{\mathrm{w}}) = k_{\mathrm{ref}} \exp\left[-\frac{E_a}{R}\left(\frac{1}{T_K} - \frac{1}{T_{\mathrm{ref},K}}\right)\right] \exp\left[-\left(\frac{a_{\mathrm{w}} - a_{\mathrm{w},*}}{s_a}\right)^2\right] \times \left[1 + \exp\left(-\frac{T - T_g}{s_T}\right)\right]^{-1}. \tag{11}$$

The Arrhenius term makes the rate strongly temperature-sensitive. The Gaussian term creates a broad middle-moisture sweet spot, and the logistic term prevents the warm interior from quietly collecting fake "browning" points. Baseline values are $k_{\mathrm{ref}} = 0.010$/s at $T_{\mathrm{ref}} = 150\,°\mathrm{C}$, $E_a = 90\,\mathrm{kJ/mol}$, $a_{\mathrm{w},*} = 0.68$, $s_a = 0.22$, $T_g = 88\,°\mathrm{C}$, and $s_T = 3\,°\mathrm{C}$. The gate is based on the fitted onset of surface browning in the chicken-roasting model of Rabeler and coauthors [2], not on a universal beef threshold. I chose the activation energy as an order-of-magnitude food-browning sensitivity within the broad ranges discussed in Maillard-reaction research [1]. Neither it nor k_{ref} is a measured beef-crust constant. The Damköhler-like reference exposure $k_{\mathrm{ref}}t$ equals 3 at 300 s; the actual exposure is much smaller because most of the steak is too cool or too wet for the model's preferred region.

Table 1: Baseline model parameters. "Effective" values are modeling choices, not universal steak properties.

Parameter	Baseline value	Basis or interpretation
Thickness, L	19 mm	Published benchmark geometry
Initial meat / pan / air temperature	20 °C, 215 °C and 25 °C	Moya benchmark [5]
Density, ρ	1070 kg/m^3	Representative wet beef
Heat capacity, c_p	3810 J/(kg K)	Representative wet beef
Thermal conductivity, k	0.49 W/(m K)	Representative wet beef
Contact conductance, H_c	120 W/(m^2 K)	Fitted effective value [5]
Air coefficient / relative humidity	8 W/(m^2 K) / 50%	Ambient-boundary assumptions
Emissivity, ε	0.90	Surface assumption
Initial moisture ratio, X_0	2.60 kg/kg dry solid	Lean-meat assumption
Dry-solid density, ρ_d	290 kg/m^3	Separate mass-bookkeeping assumption
Moisture diffusivity, D	1×10^{-9} m^2/s	Fitted effective value [5]
Mass-transfer coefficient, h_m	0.008 m/s	Published effective value [5]
Latent heat, L_v	2.26×10^6 J/kg	Water vaporization
Reactive-layer depth	0.50 mm	Resolved reporting depth

The remaining constants are $R = 8.314\,462\,618\,\mathrm{J/(mol\,K)}$, $R_v = 461.52\,\mathrm{J/(kg\,K)}$, $\sigma = 5.670\,374\,419 \times 10^{-8}\,\mathrm{W/(m^2\,K^4)}$, and $P_{\mathrm{atm}} = 101\,325\,\mathrm{Pa}$. The regularizer in Equation (3) is $\epsilon = 10^{-12}$.

3.5 How the code steps through the steak

For the baseline, I split the steak into 152 cell-centered control volumes, so $\Delta z = 0.125\,\mathrm{mm}$. Diffusive face fluxes are conservative: whatever leaves one cell enters the next. Time integration is explicit, and the step is chosen from thermal-diffusion and boundary-flux stability limits; the baseline step is 0.0286 s. At each outer face, a half-cell reconstruction solves for T_s and X_s. Ten bisection iterations match Fickian water supply to evaporation, and four fixed-point iterations feed evaporation back into the surface energy balance. These are fixed iteration counts, not residual-based tolerances. Therefore, the global closure test below does not prove that the interface iteration itself has fully converged. The reconstruction does at least avoid pretending that a cell-center value is the exact surface value.

The baseline cook lasts 300 s. Surface A touches the pan for the first 200 s, and Surface B touches it for the remaining 100 s. The internal cells stay attached to their physical positions during the flip; the steak changes orientation, but its insides do not teleport. Heat and moisture are accumulated independently so I can calculate the final global balance residuals.

3.6 Checks, benchmark, and computational experiments

Before trusting the steak results, I checked the heat solver against a problem with a known answer. The exact solution was $u(x,t) = \sin(\pi x)\exp(-\pi^2 t)$ for the one-dimensional heat equation on $x \in [0,1]$, with homogeneous Dirichlet boundaries through nondimensional time $t = 0.05$. Four grids from 20 to 160 intervals test combined space–time refinement of the interior heat operator, using $\Delta t \approx 0.4\Delta x^2$. I then repeated the coupled steak problem with 76, 152, and 304 cells and refined its stable time step with the grid. I compared global discrete balance closure, center temperature, evaporative loss, and Maillard index. The exact test checks the interior heat calculation, but not the nonlinear outer-surface iteration.

For the thermal benchmark, I copied the published 19 mm thickness, 215 °C pan condition, endpoint times of 180, 300, and 420 s, and two-thirds flip fraction. The values of H_c, D, and h_m are also transferred effective assumptions from Moya and coauthors, while D appears here in a simpler Fickian closure. For that reason, this is a reproduction benchmark, not independent validation. I compare the model with their reported center-temperature means and standard deviations [5].

I ran four groups of computational experiments. First, pan temperature varies from 160 to 230 °C, while the initial moisture ratio in each outer 0.5 mm layer is multiplied by 1.00, 0.75, or 0.50. These multipliers are numerical starting conditions, not calibrated versions of paper towels, air drying, or any other kitchen trick.

Second, I moved one flip from 35% to 75% of the 300 s cook and also tested regular alternation every 15, 30, or 50 s. Each interval divides the full 300 s into an even number of contacts, so both sides spend exactly 150 s on the pan.

Third, I compared thicknesses of 12, 19, 26, and 34 mm and pan temperatures of 180, 200, 215, and 230 °C at nominal center thresholds of 50, 55, and 60 °C. These cases use a 0.75 surface-moisture multiplier, a 0.125 mm grid, and idealized rapid alternation every 0.5 s. I recorded the endpoint after the first completed 1 s two-face cycle at or above the threshold. Stored output times

occur just after the requested times, so contact is balanced to within one numerical time step, and the rule also limits threshold overshoot. This is a balanced-contact computational limit, not a serious suggestion that someone should flip a steak every half second. The longest simulated cook is 1200 s.

Finally, I varied one sensitivity input at a time. The tested contact conductances were 80 W/(m^2 K), 120 W/(m^2 K) and 240 W/(m^2 K); moisture diffusivities were 0.3×10^{-9} m^2/s, 1.0×10^{-9} m^2/s and 3.0×10^{-9} m^2/s; activation energies were 75 kJ/mol, 90 kJ/mol and 105 kJ/mol; and thicknesses were 15 mm, 19 mm and 26 mm. Structural ablations change the water-activity closure, change the contact-face vapor-transfer factor from 1 to 0, or remove latent heat and/or evaporation.

4 Results, Verification, and Benchmarking

4.1 First question: does the solver behave?

The independent heat-equation test gave the expected second-order combined space–time refinement. Doubling the grid produced observed orders of 1.99, 1.99, and 2.00, while the L_2 error dropped from 6.01×10^{-4} on 20 intervals to 9.55×10^{-6} on 160 intervals. That was a reassuring first check: the interior solver gets closer to an exact answer as the grid improves. It still does not verify the nonlinear surface closure.

Table 2 gives the coupled grid and time-step refinement. Between 152 and 304 cells, center temperature changes by only 0.009 °C, and evaporative loss changes by 0.0002 percentage points. The surface index is less settled: the 152-cell mean is 7.2% below the 304-cell value. This makes sense because the Arrhenius rate magnifies small errors in surface temperature. I kept the 0.125 mm grid for the sweeps as a practical compromise. Differences in the index smaller than roughly ten percent are therefore not worth arguing over. Even this grid sensitivity is smaller than the structural kinetic uncertainty shown later. The relative global balance-closure residuals are 6.6×10^{-16} for energy and 8.0×10^{-14} for water. Those tiny numbers show that the bookkeeping closes, not that the fixed-count interface iteration is perfect.

Table 2: Grid refinement for the coupled 300 s baseline model.

Cells	Δz (mm)	Center T (°C)	$\overline{B}$	Evaporative loss (%)
76	0.2500	43.390	0.00903	3.1948
152	0.1250	43.375	0.01003	3.2118
304	0.0625	43.384	0.01080	3.2116

4.2 The temperature benchmark works better than the weight loss

The reduced model stays close to the published center-temperature data (Figure 2). Moya and coauthors tested six pieces for each thickness–endpoint combination. For the 19 mm steaks, their three means were 30 ± 3, 44 ± 3, and 57 ± 2 °C at 180, 300, and 420 s [5]. My model predicts 30.15 °C, 43.37 °C, and 55.75 °C. Across only those three published means, the RMSE is 0.81 °C, and every prediction falls within one reported standard deviation. The more detailed published finite-element model predicted 33, 47, and 58 °C at the same endpoints.

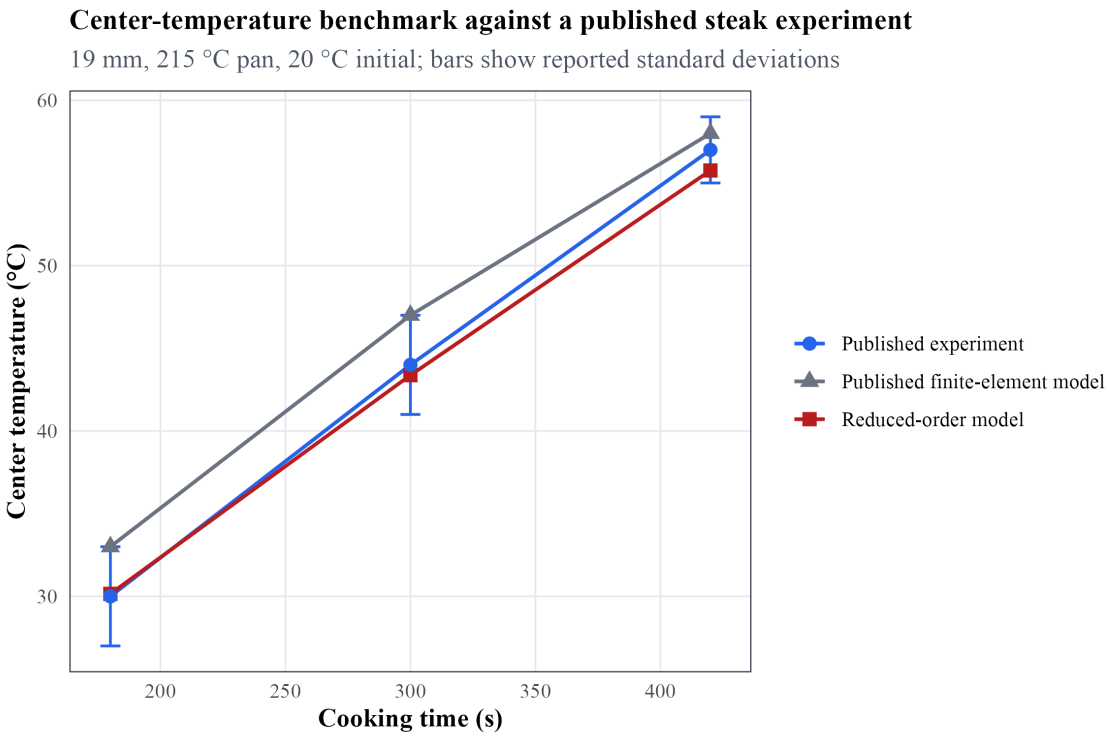

Figure 2: Thermal reproduction benchmark. Error bars are the experimental standard deviations reported by Moya and coauthors [5]. Because this reduced model borrows effective coefficients from the same study, the agreement is a reproduction check rather than independent validation.

The weight-loss result is a different story. Modeled vapor losses are 2.11%, 3.21%, and 4.00% of initial mass, while the measured total losses were $4 \pm 2\%$, $8 \pm 1\%$, and $11 \pm 2\%$ [5]. The model has no protein-contraction-driven squeezing, deformation, or edge drainage. I could force a match by turning up evaporation, but that would fix the number for the wrong reason and would distort the surface energy balance. The gap is consistent with those missing liquid-loss mechanisms, so I keep the label "evaporative loss" on purpose.

4.3 The surface races ahead of the center

Figure 3 shows the baseline history. Surface A heats quickly during its 200 s contact, but the center reaches only 43.4 °C by the end. After the flip, Surface A cools while Surface B heats. Neither one comes close to copying the 215 °C pan because contact resistance and evaporative cooling get in the way. The large surface–core gap agrees with $Bi = 2.33$. In other words, "the steak temperature" is not one temperature during this cook.

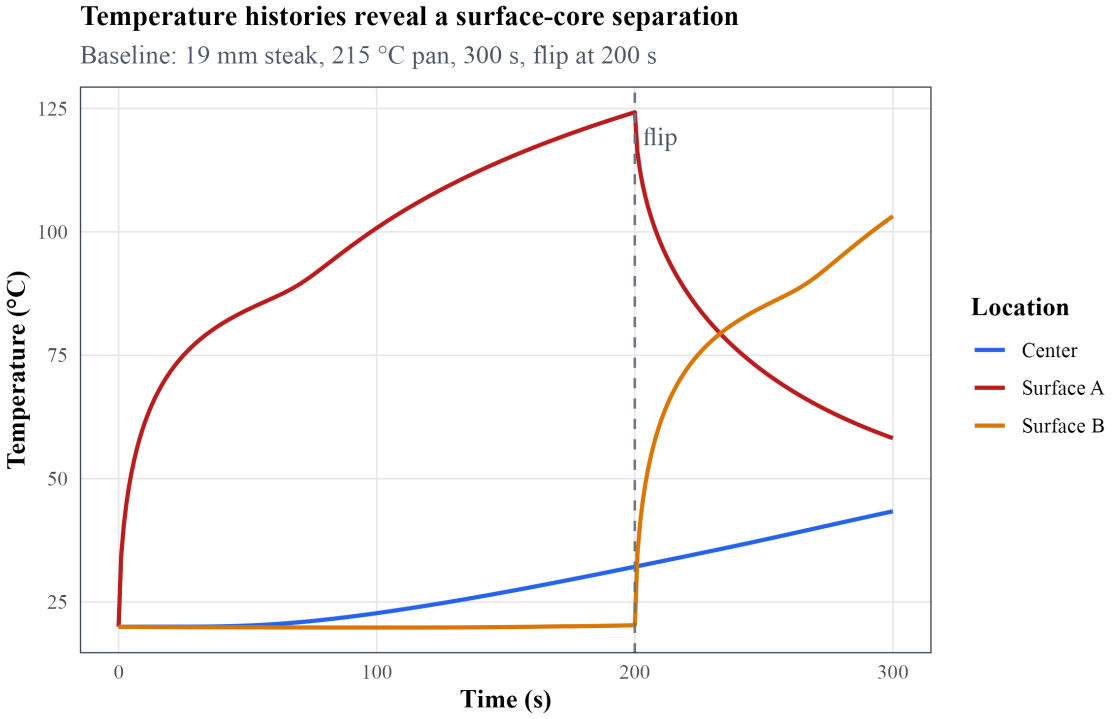

Figure 3: Outermost-control-volume temperatures in the baseline single-flip case. The dashed line marks the flip at 200 s. Face labels stay attached to the meat, so their boundary conditions swap.

The reaction history is even more uneven (Figure 4). Both surfaces start with water activity near one, where the model's moisture factor slows the reaction. Heating and evaporation then move the contacting side toward the intermediate-moisture window. Surface A reaches $B_A = 0.0191$, while the shorter second contact gives only $B_B = 0.000960$. That produces $\overline{B} = 0.0100$ but $U_B = 0.0957$. The poor balance comes directly from the two-thirds/one-third schedule. It does not prove that a real side would look "burned" or "pale."

The final profiles show that the modeled reaction stays close to the surfaces (Figure 5). Temperature changes through the full steak, but meaningful index values occupy only thin outer regions.

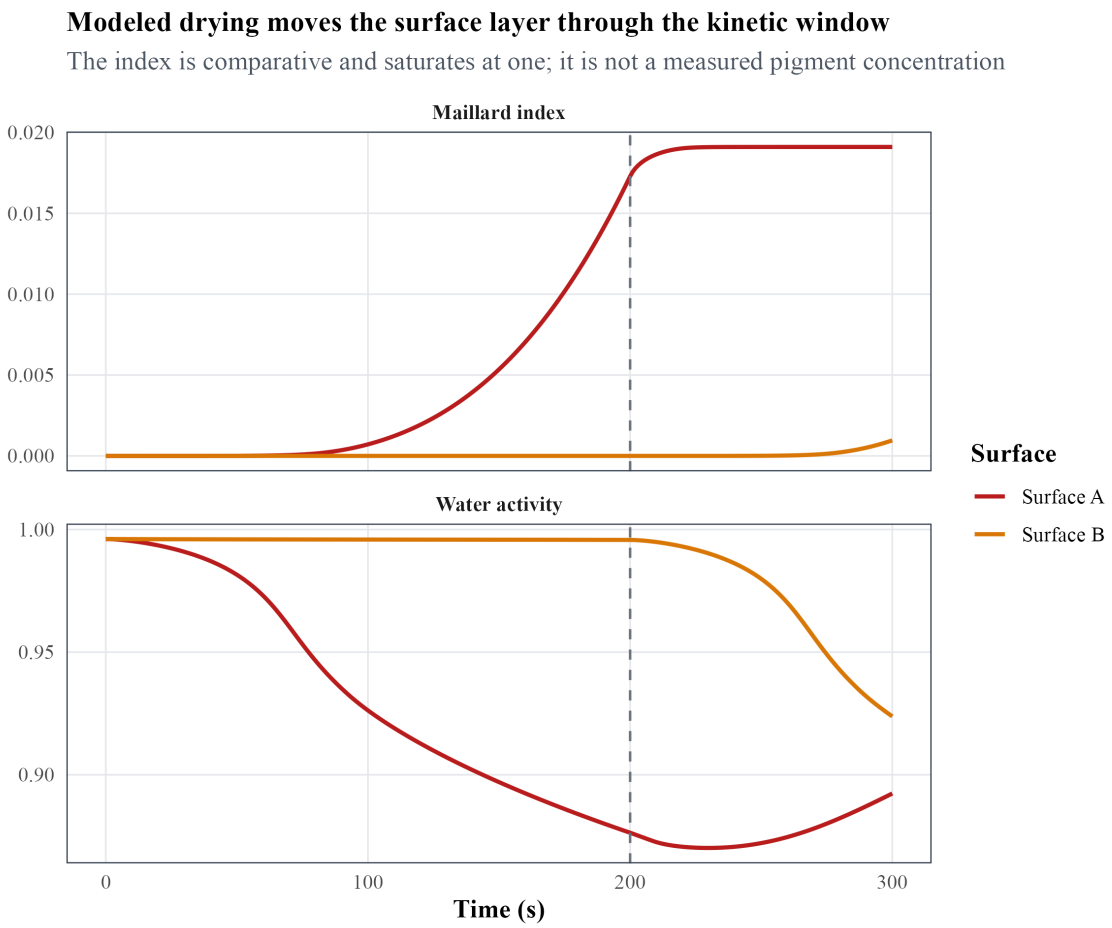

Figure 4: Outermost-control-volume water activity and accumulated reactive-layer Maillard index. Evaporation changes water activity while using heat, and the index rises only in a favorable temperature–moisture region. The curves show cell-center states, not reconstructed interface values, and the index is comparative.

That is why I treated center doneness and surface exposure as separate outputs. A model of interior redness cannot simply stand in for a model of crust formation.

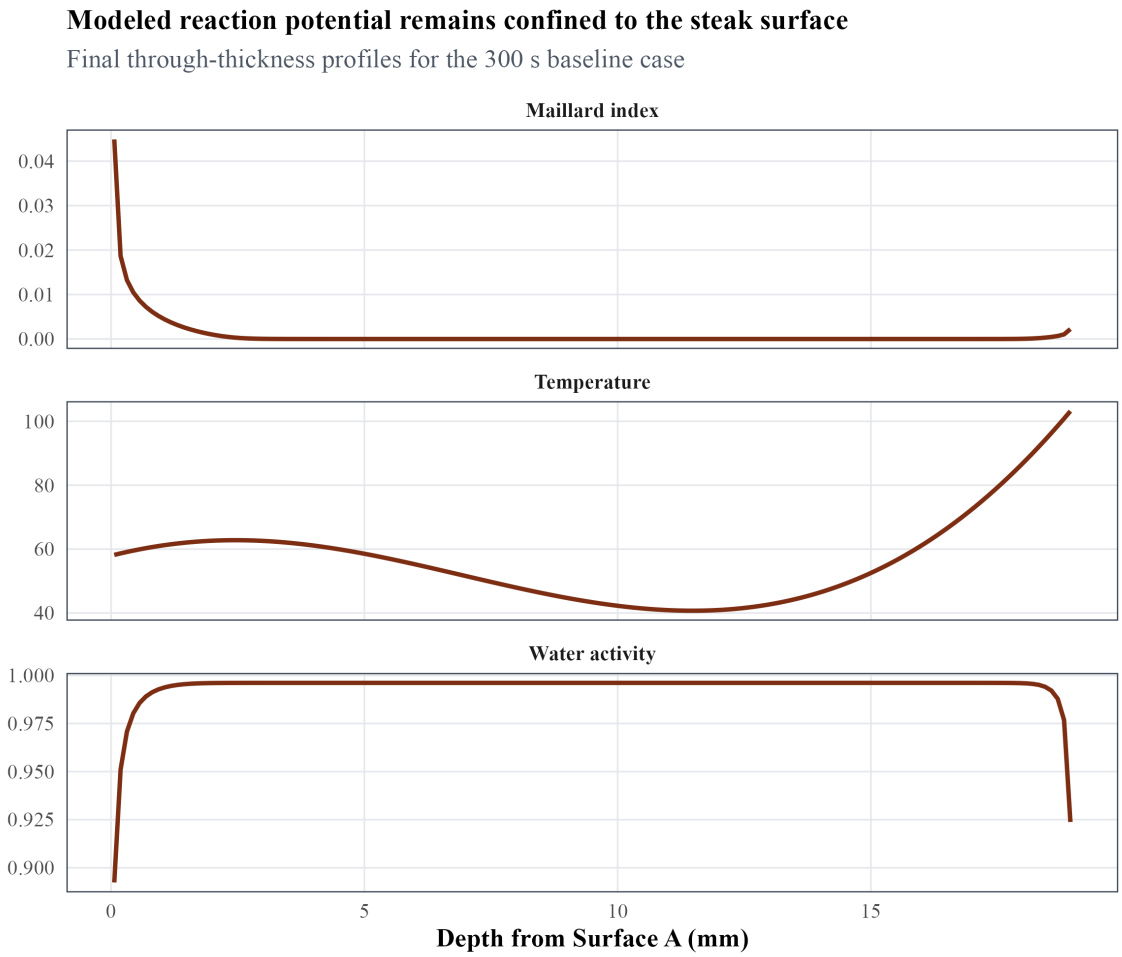

Figure 5: Final baseline profiles. Heat reaches much farther into the slab than moisture redistribution, while the modeled reaction stays near the contacting surfaces.

4.4 Heat and starting moisture work together

The fixed-time sweep is curved instead of linear (Figure 6). At 215 °C, lowering the initial surface-layer moisture multiplier from 1.00 to 0.50 raises $\overline{B}$ from 0.0100 to 0.0177, a 76.6% increase. At 230 °C, it rises from 0.0191 to 0.0320, or 67.6%. At 160 °C, every index stays near zero because the surface spends too little time in the model's high-temperature region. A slightly drier surface helps, but it cannot turn a barely hot pan into a searing machine.

The basic pattern makes physical sense, but the exact percentages are not universal. A drier starting layer needs less water removal before it reaches the assumed activity window and uses less incoming energy on evaporation. A hotter pan also gives a larger contact driving force. Since both changes act on the same surface history, they build on each other. The result still depends on the sorption closure, contact conductance, surface-layer definition, and kinetic gate. It is a reason to run an experiment, not a promise that dinner will improve by exactly 76.6%.

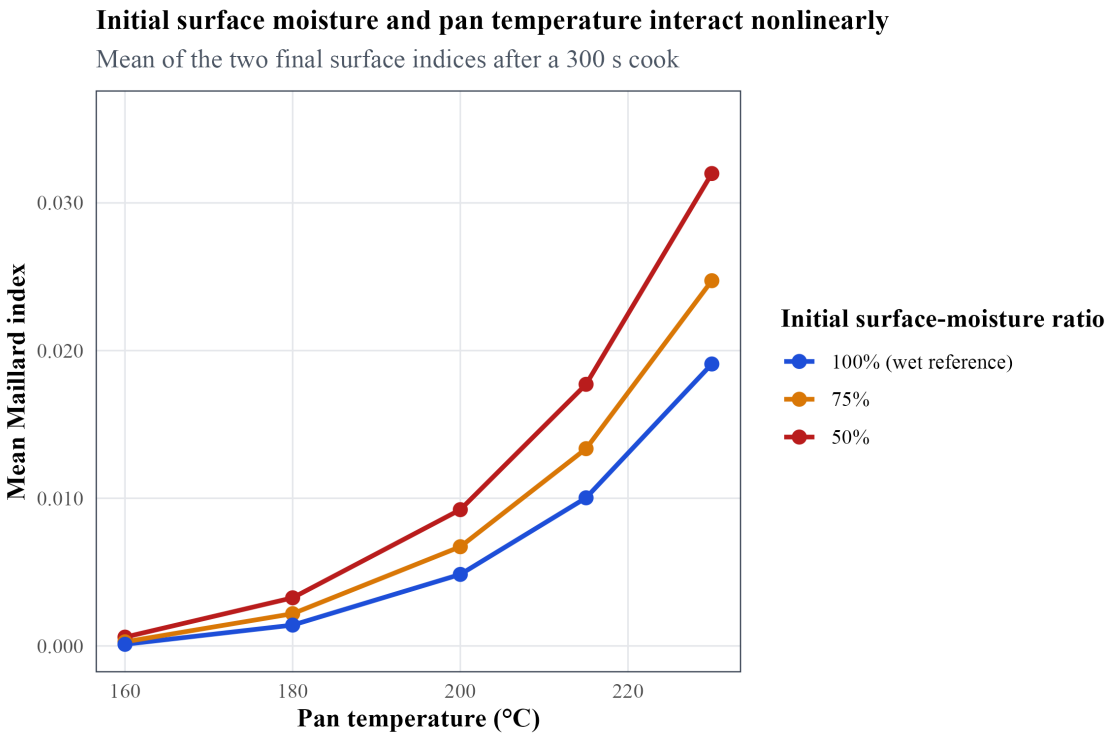

Figure 6: Fixed-time interaction between pan temperature and the numerical initial surface-moisture multiplier. Lower values shorten the evaporation and drying delay, but they are not calibrated to a specific preparation method.

4.5 Flipping is a tradeoff, not a magic trick

With one flip, an even 50/50 schedule gives the best uniformity, $U_B = 0.974$, but also the smallest mean index in the tested range, $\overline{B} = 0.00606$ (Figure 7). Waiting longer to flip increases Surface A's exposure much faster than it reduces Surface B's, so the mean rises while the balance gets worse. At a 75% first-side fraction, $\overline{B}$ reaches 0.0150, but most of it belongs to the first face. Under these assumptions, no single flip time wins at everything. The answer depends on whether the goal is more total exposure, two similar sides, or a compromise.

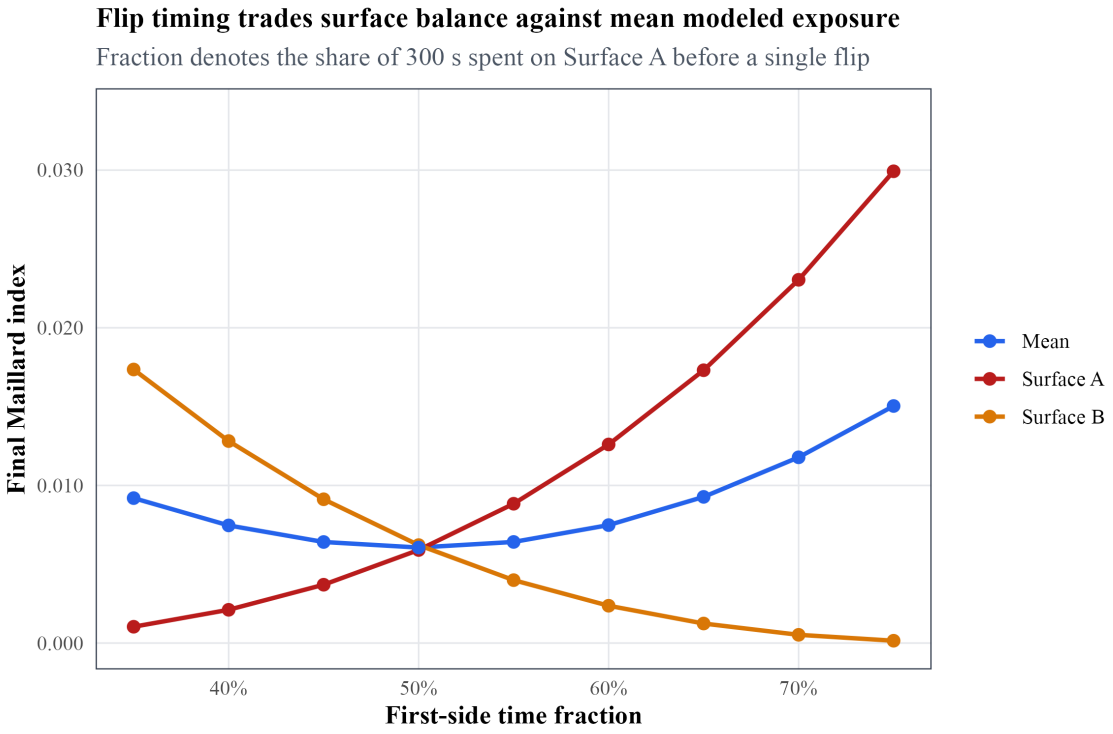

Figure 7: Single-flip timing at fixed total time. Equal contact nearly balances the two indices; uneven contact raises the mean mainly by concentrating exposure on one face.

Repeated flipping has a similar tradeoff even when total contact time is exactly equal. Alternating every 15 s gives $U_B = 0.956$ and a center temperature of 45.2 °C, but $\overline{B} = 0.000793$. Alternating every 30 seconds gives $U_B = 0.978$ and $\overline{B} = 0.00134$; every 50 s gives $U_B = 0.992$ and $\overline{B} = 0.00206$. In this boundary model, faster swapping slightly raises the center temperature, but it keeps interrupting each surface's hot period and lowers the fixed-time index. This only compares equal cooking times. A matched center temperature or real experimental endpoint could change the order.

4.6 Comparing steaks at the same center temperature

Thickness produces the clearest pattern in the design space (Figure 8). With a 215 °C pan and a 55 °C center target, the 12 mm slab gets there in 190 s with $\overline{B} = 1.45 \times 10^{-5}$. The 19, 26, and 34 mm slabs need 367, 579, and 870 s, reaching indices of 0.00420, 0.0314, and 0.132. The chemistry is not somehow faster just because the steak is thicker. Its surface simply keeps heating for much longer

while the center catches up. Across all 48 matched-threshold cases, the largest center-temperature overshoot above the target is only 0.171 °C.

For a fixed 19 mm thickness and the same 55 °C target, raising the pan from 180 to 230 °C cuts the modeled time from 442 to 344 s, while $\overline{B}$ rises from 0.000423 to 0.00694. The hotter surface more than makes up for the shorter cook. Modeled vapor loss stays in the narrower range of about 3.83–4.23%, but the missing liquid-loss physics means this is not a complete answer about juiciness.

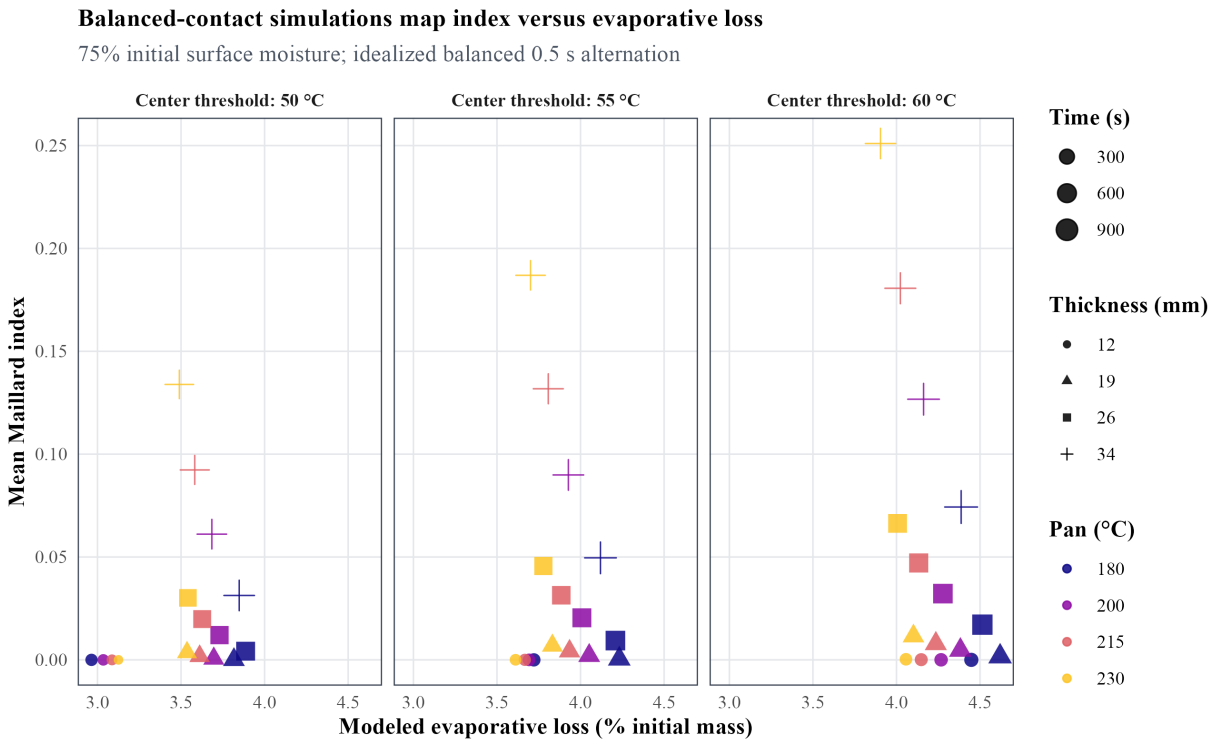

Figure 8: Balanced-contact matched-threshold design space. Color shows pan temperature, shape shows slab thickness, and point size shows time to threshold. Thicker slabs get more surface-heating time while reaching almost the same center target. These targets are research constraints, not safety recommendations.

4.7 Which assumptions move the result the most?

The one-at-a-time study shows how much the outputs move over the chosen ranges (Figure 9). Lowering H_c from 120 to 80 W/(m^2 K) reduces the baseline index to 0.063 times its reference value and drops the center temperature to 37.2 °C. Raising it to 240 W/(m^2 K) multiplies the index by 13.8 and raises the center to 54.7 °C. Contact conductance causes the biggest response among the perturbations I tested. Since the ranges are not normalized or probabilistic, this is not a global ranking of every uncertainty in steak cooking.

Changing D from 1.0×10^{-9} to 0.3×10^{-9} m^2 s^{-1} raises the index by 49%. Increasing it to 3.0×10^{-9} m^2 s^{-1} lowers the index by 61%. Faster moisture resupply keeps the boundary wetter and more strongly cooled. Changing E_a from 90 to 75 kJ/mol raises the index by 54%, while changing

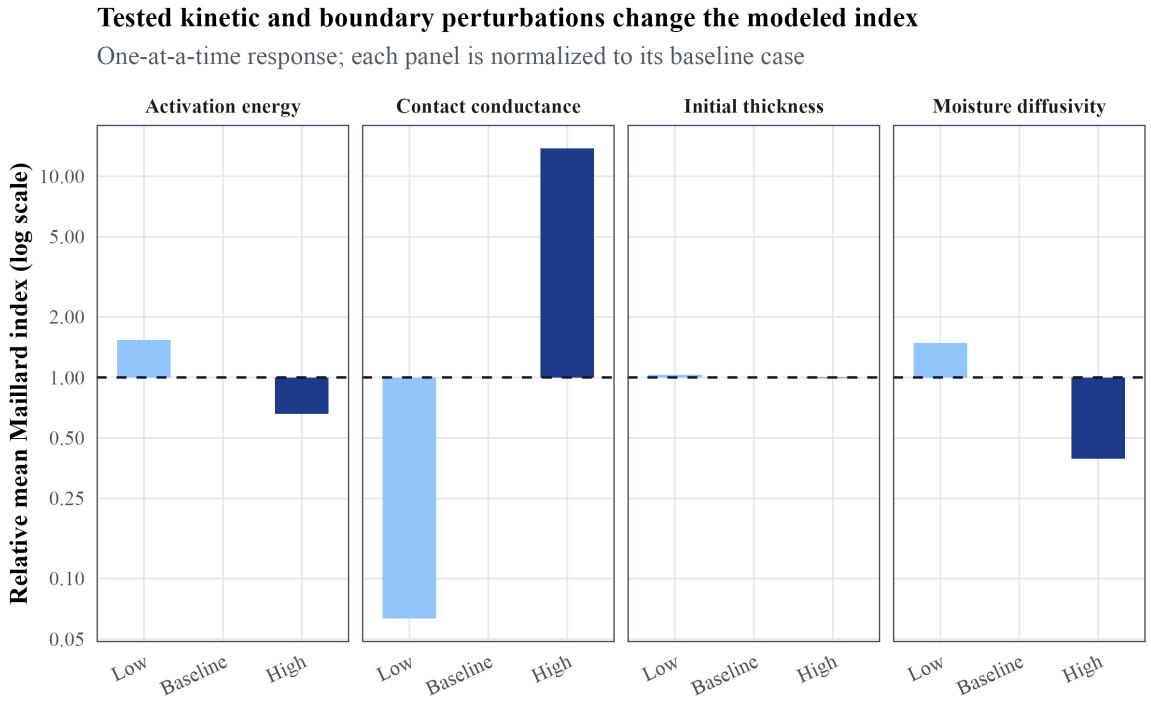

Figure 9: One-at-a-time changes around the 300 s baseline. Response is relative to baseline on a logarithmic scale. Because the parameter ranges differ, bar height is not a normalized global sensitivity score.

it to 105 kJ/mol lowers it by 34%. Thickness barely changes fixed-time surface exposure, but it strongly changes center temperature. That is why thickness becomes so important in the matched-threshold experiment.

The structural ablation is even more revealing (Table 3). Replacing the GAB closure nearly doubles the index. Blocking vapor escape only at the pan-contact face raises center temperature by 3.87 °C, more than doubles the index, and cuts predicted evaporative loss from 3.21% to 0.94%. So the mixed-contact evaporation assumption is not just decorative math. Removing the latent-heat sink raises center temperature by 5.19 °C and multiplies the index by 3.60. Turning evaporation off completely causes the same temperature rise but gives an index 2.27 times baseline, because the surface stays too wet for the Gaussian moisture optimum. Evaporation is therefore more than a weight-loss output: it cools the surface while also changing water activity. Both effects matter, and the contact-face closure still needs experimental calibration.

Table 3: Structural ablations for the 300 s baseline.

Model structure	Center T (°C)	$\overline{B}$	Vapor loss (%)
Baseline GAB	43.37	0.01003	3.212
No vapor escape at contact face	47.24	0.02105	0.936
Simple a_w closure	43.52	0.01915	3.104
No latent-heat sink	48.57	0.03615	3.380
No evaporation	48.57	0.02273	0.000

5 Discussion and Future Work

5.1 The main result is a regime map, not a recipe

The clearest result is that there is no magic combination of pan temperature and flip time. This is slightly disappointing for anyone hoping the computer would produce the One True Steak Recipe, but it is also more interesting. The center and surface live through very different cooking experiences. The center warms slowly, while a thin outer layer heats, loses water, cools by evaporation, and sometimes reaches conditions that favor Maillard chemistry. Pan temperature matters, but it is not the same thing as surface temperature.

This helps explain why normal cooking advice can sound contradictory. Frequent flipping improved the temperature balance between the two faces, but it did not maximize the fixed-time browning index. A thicker slab built up more modeled surface exposure when cases were stopped at the same center temperature, even though thickness had little effect in a fixed-time comparison. Starting with a drier surface reduced the evaporative delay; it did not create a magical waterproof crust. A hotter pan shortened the cook while still increasing the index because reaction rates rose sharply with temperature. In other words, the answer changes with the question being asked.

For that reason, I think the matched-threshold plots are more useful than the fixed-time plots. If two steaks cook for the same number of seconds but finish at different center temperatures, the comparison is not completely fair. Stopping cases near the same center threshold removes much of that problem and shows the tradeoff between heating, water loss, and modeled surface exposure. The map is not ready to control a real pan, but it gives a sensible way to design an experiment

instead of changing several kitchen variables and hoping for enlightenment.

5.2 What is robust and what is conditional

Some conclusions survived every test I tried within this model. The steak cannot be treated as one uniform temperature. Contact heat transfer has a large effect, evaporation removes both mass and energy, and heat moves inward much faster than the simplified moisture field. Comparing cases at a shared center-temperature target also makes more sense than comparing them only at a shared clock time. These points come from the balances and overall trends, not from one lucky choice of reaction constant.

The exact value of B is much less secure. Changing only the water-activity equation almost doubled the index, while changing contact conductance moved it by more than an order of magnitude. The activation energy, reference rate, reactive-layer thickness, and temperature gate were not fitted to measured steak-crust chemistry. Therefore, $B = 0.02$ does not mean "twice as delicious" as $B = 0.01$, tempting though that interpretation may be. The index is useful for comparing simulations under the same assumptions, not for predicting a particular color, smell, or taste.

The thermal benchmark needs the same honesty. An RMSE of 0.81 °C is encouraging, but several effective parameters came from the same published steak study used for comparison [5]. The result shows that this simpler slab model can reproduce those center-temperature points. It does not prove that the model will work unchanged for a different cut, pan, oil, pressure, or starting temperature. Also, because the browning index changed more slowly with mesh refinement than center temperature did, future studies should always report their surface resolution.

5.3 Model limitations

The geometry is a flat one-dimensional slab, so edges, corners, uneven contact, and sideways heat flow disappear. Real meat is not nearly that cooperative. Marbling, connective tissue, surface fat, salt, pH, oil, pan texture, and the direction of muscle fibers can all matter. The model also holds material properties constant while actual meat denatures, shrinks, and releases liquid. Contact conductance is fixed even though direct measurements show that heat flux can change during cooking [4]. Ambient humidity is fixed, and the thin vapor-and-liquid layer under a steak is not resolved.

The moisture model has an even harder job. It treats moisture as diffusion plus surface evaporation, so it misses liquid squeezed out by shrinking proteins and liquid that runs into the pan. It also allows simplified evaporation from a face labeled as contacting the pan. A real contact surface is a patchwork of contact points, gaps, oil, steam, and sizzling liquid—basically a tiny mess. These missing processes explain why predicted vapor loss was lower than measured total mass loss in the benchmark. The model should therefore not be used to predict juiciness.

Finally, Equation (11) compresses a large chemical network into one relative exposure rate. The Maillard reaction really contains many pathways and products [1]. This model does not track reactant depletion, pH, lipid oxidation, aroma compounds, or harmful compounds that can form during high-temperature cooking. The index is not a health-risk measure, and a dark simulated crust is not automatically a safe or tasty real steak.

5.4 A validation experiment suggested by the model

A follow-up experiment could use steaks from the same cut and thickness and test two pan temperatures, two measured starting surface-moisture levels, and two flip schedules. Small thermocouples

and calibrated infrared images could record center and surface temperatures. A pan heat-flux measurement would be especially valuable because the sensitivity study pointed to contact conductance as one of the biggest uncertainties. Cases should stop at matched center temperatures rather than just matched times.

Surface browning should also be measured separately from interior color. Standardized photographs could provide CIELAB L^* values, and a water extract could be tested for absorbance at 420 nm. Yoo and coauthors used physicochemical and sensory measurements to distinguish seared steaks cooked to a matched core endpoint [3]. Total sample loss, liquid left in the pan, and estimated evaporation should be recorded separately. The kinetic constants could then be fitted to part of the dataset and tested on the remaining cases. That would be a much stronger challenge than fitting every point and congratulating the model for remembering them.

After that, uncertain inputs such as thickness, surface moisture, contact conductance, and probe placement could be sampled in repeated simulations. Prediction intervals would show whether two cooking methods are meaningfully different or merely look different because one input was guessed. A two- or three-dimensional model should be added only if the experimental errors show that edges or shrinkage are important enough to justify the extra complexity.

6 Conclusion

This project built a one-dimensional heat-and-moisture model to study when a steak surface enters conditions that favor Maillard browning. The solver matched three published center-temperature means with an RMSE of 0.81 °C, but it underpredicted total cooking loss because it does not model shrinking meat squeezing out liquid. In the simulations, browning depended on the whole temperature–moisture path, not just the number printed on the pan. Surface moisture and contact conductance had especially large effects, while flip timing created a tradeoff between average exposure and equal browning on both faces.

The main lesson is simple: center temperature, water loss, surface browning, and food safety are different outputs. One cannot stand in for all the others. The temperature model has a useful benchmark, while the Maillard index is still a comparison tool that needs experimental calibration. Even with that limitation, the model turns vague cooking claims into questions that can actually be tested. It did not discover the perfect steak, but it did show what should be measured next—which is a respectable outcome for a computer that never gets dinner.

Reproducibility Statement

The source package includes this LaTeX manuscript, its bibliography, the R simulation and plotting scripts, tabulated outputs, and rendered figures. The scripts reproduce the parameter sweeps, tables, and figures reported here.

References

[1] Sara I. F. S. Martins, Wim M. F. Jongen, and Martinus A. J. S. van Boekel. A review of Maillard reaction in food and implications to kinetic modelling. *Trends in Food Science & Technology*, 11(9–10):364–373, 2000. doi: 10.1016/S0924-2244(01)00022-X.

[2] Felix Rabeler, Jacob Lercke Skytte, and Aberham Hailu Feyissa. Prediction of thermal

induced color changes of chicken breast meat during convective roasting: A combined mechanistic and kinetic modelling approach. *Food Control*, 104:42–49, 2019. doi: 10.1016/j.foodcont.2019.04.018.

[3] Ji Hyun Yoo, Ji Won Kim, Hae In Yong, Ki Ho Baek, Hyun Jung Lee, and Cheorun Jo. Effects of searing cooking on sensory and physicochemical properties of beef steak. *Food Science of Animal Resources*, 40(1):44–54, 2020. doi: 10.5851/kosfa.2019.e80.

[4] Richard Rocca-Poliméni, Nuria Zárate Vilet, Stéphanie Roux, Jean-Luc Bailleul, and Bertrand Broyart. Continuous measurement of contact heat flux during minced meat grilling. *Journal of Food Engineering*, 242:163–171, 2019. doi: 10.1016/j.jfoodeng.2018.08.032.

[5] J. Moya, S. Lorente-Bailo, M. L. Salvador, A. Ferrer-Mairal, M. A. Martínez, B. Calvo, and J. Grasa. Development and validation of a computational model for steak double-sided pan cooking. *Journal of Food Engineering*, 298:110498, 2021. doi: 10.1016/j.jfoodeng.2021.110498.

[6] U.S. Department of Agriculture, Food Safety and Inspection Service. Safe minimum internal temperature chart. https://www.fsis.usda.gov/food-safety/safe-food-handling-and-preparation/food-safety-basics/safe-temperature-chart, n.d. Accessed 26 August 2026.

[7] Daitaro Ishikawa, Genjiro Ueno, and Tomoyuki Fujii. Estimation method of moisture content at the meat surface during drying process by NIR spectroscopy and its application for monitoring of water activity. *Japan Journal of Food Engineering*, 18(3):135–143, 2017. doi: 10.11301/jsfe.17493.

AI Assistance Disclosure

Artificial intelligence tools were used solely to guide understanding, improve clarity, organization, and grammatical precision of the written manuscript. All research design, mathematical modeling, data analysis, code implementation, experimental validation, and scientific conclusions were developed and executed independently by the author.