Hierarchical Bayesian inference of β-lactoglobulin denaturation and lactic acidification kinetics for a stirred-yoghurt line: a meta-analysis of published kinetics and its consequences for process design

A.Benahmed

Conservatoire National des Arts et Métiers (CNAM), France

abenahmed1@gmail.com

November 2026

Abstract

In an earlier preprint I optimised a stirred-yoghurt line with a deterministic model, and the optimiser pushed pasteurisation to the upper bounds of temperature and holding time. That result rests on single values of the kinetic constants. Published work on β-lactoglobulin (β-Lg) denaturation in milk does not agree on those constants: reported orders run from 1 to 2, and the Arrhenius break sits anywhere between 83 and 95 °C. Here I treat the disagreement as uncertainty and follow it through to the design. The model part gives a normalised two-regime Arrhenius law with a smooth break, a rational lifting for symbolic identifiability analysis, the gradient of the stopping time at pH 4.60, and a chance-constrained version of the optimisation. The data part is a hierarchical Bayesian meta-analysis of four studies, seven entries in all. The data are of three kinds: 107 points of the Dannenberg and Kessler curves, digitised from a published re-plot; nine isothermal rate constants from Claeys et al.; and the Arrhenius regressions of Manji and Kakuda and of Oldfield. The population activation energy is 280 kJ mol^{-1} (95 % interval 261–299) below the break and 51 kJ mol^{-1} (44–62) above it. The break lies at 89.6 °C (86.6–92.4), and the mean order is 1.56 (1.21–1.90). At 90 °C the rate constants of the seven entries span a factor of four. For a milk that has not been characterised, the predicted denaturation after 300 s at 90 °C has a 95 % interval from 0.54 to 1.00. The value used in the earlier preprint, 0.80, falls at the 18th percentile. StructuralIdentifiability.jl shows that one non-isothermal curve identifies E_a/R globally. From pH data alone, the growth rate, the lag variable and the product of the buffering peak with its shape factor are locally identifiable. Uncertainty does not move the optimum away from 95 °C and 600 s. If the plant also sets a ceiling on denaturation, no operating point meets both chance constraints until the plant's own milk has been measured. With two (simulated) measurements on that milk, a ceiling of 0.92 is met at 92.5 °C and 420 s. Batch variability means that four tanks sized deterministically are too small for about 62 % of simulated batches.

Keywords: β-lactoglobulin; denaturation kinetics; lactic acidification; hierarchical Bayesian inference; identifiability; event sensitivity; chance constraints; stirred yoghurt.

1 Introduction

A stirred-yoghurt line chains three operations: heat treatment of the milk, which denatures whey proteins and above all β-Lg; lactic fermentation by *Streptococcus thermophilus* and *Lactobacillus delbrueckii* subsp. *bulgaricus* down to about pH 4.6; and breaking, pumping and cooling of the gel. In an earlier preprint I built a deterministic model of such a line from public data and optimised

it [1]. The optimiser drove pasteurisation to the upper bounds of the search box, 95 °C and 600 s. That result rests on point estimates of the kinetic constants, and two of those estimates are less certain than a single number suggests.

The first is the denaturation of β-Lg. Dannenberg and Kessler reported an order of 1.5 in skim milk [2], as did Anema and McKenna in reconstituted whole milk, together with an abrupt change in the Arrhenius slope [5]. Oldfield found orders from 1.0 to 1.6 in skim milk heated by direct steam injection and a change in activation energy at 90 °C from 51.18 to 301.73 kJ mol^{-1} [6, 7]. Claeys et al. found first-order kinetics in raw milk [10], and Galani and Owusu Apenten found first order between 75 and 85 °C and second order between 90 and 120 °C for a whey protein mixture in buffer [8]. Matrix, heating equipment and assay all differ between these studies. Treating one set of constants as exact hides this spread.

The second is the starter. Lag, growth rate and the buffering capacity of the milk vary from batch to batch, and with them the time needed to reach pH 4.6, which in turn sets the number and size of fermentation tanks.

Hierarchical Bayesian models are the standard way to pool studies or batches while keeping their differences visible. Luong et al. used one to describe the pH of meat products across production batches [11]. I have not found a hierarchical treatment of the dispersion between β-Lg denaturation studies, nor one that carries both sources of uncertainty through to the design of a yoghurt line. This paper does that for the denaturation block and sets up the same machinery for fermentation. It contains:

1. a normalised, smooth form of the two-regime Arrhenius law and an explicit bound on the error introduced by smoothing;
2. a rational lifting of the model, which is what symbolic identifiability software requires, and an analytic statement of what an isothermal design can identify;
3. the sensitivity of the stopping time at pH 4.60, which lets a gradient-based sampler use endpoint-only data;
4. a hierarchical likelihood that separates tabulated from digitised data, and a chance-constrained form of the optimisation;
5. a hierarchical meta-analysis of the denaturation kinetics reported in four studies, its propagation through the line model, and the effect of a plant-specific calibration;
6. a symbolic identifiability analysis of the lifted models and numerical verification on synthetic data.

The meta-analysis uses whatever form of data each study makes available: raw curves for one study, rate constants with standard errors for another, regression lines or fitted parameters for the rest. The likelihood of Section 5 handles each form. All numbers in Sections 5 and 7 come from scripts that are part of the supplementary material.

2 Process model

The model is the one of [1], with one addition: a lag term in the growth rate, in the form proposed by Baranyi and Roberts [14]. Table 1 lists the symbols. Units are SI throughout.

2.1 Denaturation of β-lactoglobulin

The rate law $\mathrm{d}C_\beta/\mathrm{d}t = -k(T)\,C_\beta^{n}$ has a rate constant whose unit depends on n. When n varies between studies that used different initial concentrations $C_{\beta,0}$, $\ln k$ becomes correlated with

Table 1: Main symbols.

Symbol	Meaning	Unit
T_{past}, t_{hold}	pasteurisation temperature, holding time	K, s
T_{ferm}, t_{ferm}	fermentation temperature, time to pH 4.60	K, s
y_β, $D_\beta = 1 - y_\beta$	native fraction and denatured fraction of β-Lg	–
$\tilde{k}(T)$, n, T_b	normalised rate constant, reaction order, break temperature	s^{-1}, –, K
$E_{a,1}, E_{a,2}$	activation energies below and above T_b	$\mathrm{J\,mol}^{-1}$
X, S, P	biomass, lactose, lactic acid	$\mathrm{kg\,m}^{-3}$
q	Baranyi physiological state	–
β_{buf}, β_{CCP}	buffering capacity and its peak amplitude	$\mathrm{mol\,m}^{-3}\,\mathrm{pH}^{-1}$
$\dot{\gamma}$	shear rate during pumping	s^{-1}
η_{50}	apparent viscosity at 50 s^{-1}	Pa s

$(1-n)\ln C_{\beta,0}$ for reasons that have nothing to do with kinetics. Working with the native fraction $y_\beta = C_\beta / C_{\beta,0}$ and $\tilde{k} = k\, C_{\beta,0}^{n-1}$ removes this:

$$\frac{\mathrm{d}y_\beta}{\mathrm{d}t} = -\tilde{k}(T)\, y_\beta^n, \qquad y_\beta(0) = 1, \tag{1}$$

which at constant temperature gives $y_\beta = \bigl(1 + (n-1)\tilde{k}t\bigr)^{-1/(n-1)}$ for $n \neq 1$ and $y_\beta = e^{-\tilde{k}t}$ for $n = 1$. Near $n = 1$ I evaluate $\ln(1+u)/u$ with $u = (n-1)\tilde{k}t$ through its Taylor series, which keeps the expression smooth.

For the temperature dependence, write $\xi(T) = 1/T - 1/T_b$ and $\tilde{k}_b = \tilde{k}(T_b)$. A continuous two-regime Arrhenius law has four free parameters, $(\ln \tilde{k}_b, E_{a,1}, E_{a,2}, T_b)$:

$$\ln \tilde{k}_0(T) = \ln \tilde{k}_b - \frac{E_{a,2}}{R}\,\xi - \frac{E_{a,1} - E_{a,2}}{R}\,\max(0, \xi). \tag{2}$$

Below T_b ($\xi > 0$) the slope with respect to $1/T$ is $-E_{a,1}/R$; above it is $-E_{a,2}/R$. The kink at T_b is a problem for Hamiltonian Monte Carlo when T_b is itself a parameter. Replacing $\max(0,\xi)$ by the softplus $s_\tau(\xi) = \tau \ln(1 + e^{\xi/\tau})$ gives

$$\ln \tilde{k}_\tau(T) = \ln \tilde{k}_b - \frac{E_{a,2}}{R}\,\xi - \frac{E_{a,1} - E_{a,2}}{R}\, s_\tau(\xi). \tag{3}$$

Proposition 1. *$\ln \tilde{k}_\tau$ is infinitely differentiable, has the two Arrhenius slopes as limits far from T_b, and satisfies $\sup_T |\ln \tilde{k}_\tau(T) - \ln \tilde{k}_0(T)| \le |E_{a,1} - E_{a,2}|\,\tau \ln 2 / R$, with the maximum reached at $T = T_b$.*

Proof. For any real z, $s_\tau(z) - \max(0, z) = \tau \ln(1 + e^{-|z|/\tau})$, which lies in $(0, \tau \ln 2]$ and equals $\tau \ln 2$ at $z = 0$. Multiplying by $|E_{a,1} - E_{a,2}|/R$ gives the bound; smoothness and the limits follow from those of s_τ. □

The bound matters in practice. With $E_{a,1} - E_{a,2} = 211$ $\mathrm{kJ\,mol}^{-1}$ and $\tau = 5 \times 10^{-6}$ K^{-1} (a transition about 0.65 K wide), the error at T_b is 0.088 in $\ln \tilde{k}$, and D_β after 300 s at 90 °C drops from 0.800 to 0.780. With $\tau = 10^{-6}$ K^{-1} the error is 0.018 and $D_\beta = 0.796$. I recommend $\tau = 10^{-6}$ K^{-1} for the sampler and use it in the identifiability analysis below; forward propagation uses the sharp form.

The nominal values are those of [1], after Dannenberg and Kessler [2]: $n = 1.5$, $E_{a,1} = 265$, $E_{a,2} = 54$ $\mathrm{kJ\,mol}^{-1}$, $T_b = 363.15$ K and $C_{\beta,0} = 3.2$ $\mathrm{kg\,m}^{-3}$. They give $D_\beta = 0.462$, 0.800 and 0.849 after 300 s at 85, 90 and 95 °C.

2.2 Fermentation

Growth follows the cardinal temperature model with inflection (CTMI) [12], Monod limitation by lactose and product inhibition. Acid production follows Luedeking and Piret [13], and the pH follows from the buffering capacity of the milk:

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \mu_{\max} q, \qquad \alpha(q) = \frac{q}{1+q}, \tag{4}$$

$$\mu = \mu_{\max}\,\alpha(q)\,\gamma_{\mathrm{CTMI}}(T_{\mathrm{ferm}})\,\frac{S}{K_S + S}\Big(1 - \frac{P}{P_{\max}}\Big)^{m_p}, \tag{5}$$

$$\frac{\mathrm{d}X}{\mathrm{d}t} = (\mu - k_{\mathrm{dec}})X, \qquad \frac{\mathrm{d}P}{\mathrm{d}t} = \alpha_{\mathrm{LP}} \max\Big(0, \frac{\mathrm{d}X}{\mathrm{d}t}\Big) + \beta_{\mathrm{LP}} X, \qquad \frac{\mathrm{d}S}{\mathrm{d}t} = -\frac{\mu X}{Y_{X/S}} - m_S X, \tag{6}$$

$$\frac{\mathrm{dpH}}{\mathrm{d}t} = -\frac{1}{M_{\mathrm{LA}}\,\beta_{\mathrm{buf}}(\mathrm{pH})}\frac{\mathrm{d}P}{\mathrm{d}t}, \qquad \beta_{\mathrm{buf}}(\mathrm{pH}) = \beta_0 + \beta_{\mathrm{CCP}} \exp\Big(-\frac{(\mathrm{pH} - \mathrm{pH}_{\mathrm{CCP}})^2}{2\sigma_{\mathrm{buf}}^2}\Big). \tag{7}$$

The Gaussian term stands for the buffering peak that comes from the solubilisation of colloidal calcium phosphate; for casein micelle suspensions Salaün et al. found the maximum near pH 5.2 [16]. The parameters are those of [1]: $\mu_{\max} = 3.2 \times 10^{-4}$ s^{-1}, $K_S = 1.5$, $P_{\max} = 16$ kg m^{-3}, $m_p = 1.2$, $k_{\mathrm{dec}} = 1.2 \times 10^{-5}$ s^{-1}, $\alpha_{\mathrm{LP}} = 3.8$, $\beta_{\mathrm{LP}} = 8.5 \times 10^{-5}$ s^{-1}, $Y_{X/S} = 0.14$, $m_S = 2.5 \times 10^{-5}$ s^{-1}, $\beta_0 = 22$, $\beta_{\mathrm{CCP}} = 28$ mol m^{-3} pH^{-1}, $\mathrm{pH}_{\mathrm{CCP}} = 5.15$, $\sigma_{\mathrm{buf}} = 0.45$, cardinal temperatures 293.15, 315.15 and 325.15 K. The batch stops at t_{ferm}, where $\mathrm{pH}(t_{\mathrm{ferm}}) = 4.60$.

The lag term $\alpha(q)$ is new. It corresponds to a lag of $\mu_{\max}^{-1}\ln(1 + 1/q_0)$ and I take $q_0 = 10$ as the nominal value (about 5 min). This lengthens the batch at 42 °C from 14 726 s without the term to 15 033 s (4.18 h).

2.3 Texture, energy and tank size

Pumping at shear rate $\dot\gamma$ leaves a structural integrity $\lambda = k_{\mathrm{rec}}/(k_{\mathrm{rec}} + k_{\mathrm{rup}}\dot\gamma)$ with $k_{\mathrm{rec}} = 0.015$ s^{-1} and $k_{\mathrm{rup}} = 3.5 \times 10^{-4}$. The Herschel–Bulkley [15] parameters at 4 °C and the syneresis index are the phenomenological, unfitted relations of [1]. Since fermentation always stops at pH 4.60, the pH factors of the original equations equal 1:

$$\tau_0 = (2.5 + 14 D_\beta^{1.5})(0.25 + 0.75\lambda), \qquad K = (1.8 + 6.4 D_\beta^{1.3})(0.35 + 0.65\lambda), \qquad n_{\mathrm{HB}} = 0.42, \tag{8}$$

$$\eta_{50} = \tau_0/50 + K\,50^{\,n_{\mathrm{HB}}-1}, \qquad S_{\mathrm{syn}} = 0.28\,(1 - 0.72 D_\beta)(1.15 - 0.45\lambda). \tag{9}$$

Specific energy comprises heating with regeneration (efficiency 0.85, capped by T_{ferm}), losses from a DN80 holding tube (inner diameter 0.081 m, $U = 1.2$ W m^{-2} K^{-1}) and chilling to 4 °C with a coefficient of performance of 3.2:

$$E = c_p\big[\Delta T - \min(0.85\,\Delta T,\; T_{\mathrm{past}} - T_{\mathrm{ferm}})\big] + \frac{4U(T_{\mathrm{past}} - T_{\mathrm{amb}})\,t_{\mathrm{hold}}}{\rho D} + \frac{c_p(T_{\mathrm{ferm}} - 277.15)}{3.2}, \tag{10}$$

with $\Delta T = T_{\mathrm{past}} - 277.15$ K, $c_p = 3930$ J kg^{-1} K^{-1}, $\rho = 1030$ kg m^{-3} and $T_{\mathrm{amb}} = 298.15$ K. For a continuous feed of 2.5 kg s^{-1} served by $N = 4$ tanks in rotation, with 5400 s for emptying and cleaning and a fill ratio of 0.82, each tank needs

$$V_{\mathrm{tank}} = \frac{\dot V\,(t_{\mathrm{ferm}} + 5400)}{(N-1) \times 0.82}. \tag{11}$$

3 Identifiability and the gradient of the stopping time

3.1 What an isothermal design can identify

From (1), $y'_\beta(0) = -\tilde k$ and $y''_\beta(0) = n\tilde k^2$. A single noise-free isothermal curve with known $y_\beta(0) = 1$ therefore fixes both $\tilde k(T)$ and n. In a single Arrhenius regime, two temperatures fix $\tilde k_{\mathrm{ref}}$ and E_a.

With two regimes and a continuity constraint, $\ln \tilde{k}$ is a pair of lines in $1/T$ that meet at T_b: at least two temperatures strictly on each side of T_b are needed to determine both slopes and the point where they meet. With temperatures on one side only, the slope on the other side and T_b cannot be determined. Section 7.2 checks this numerically.

3.2 Rational lifting

Symbolic tools for structural identifiability, such as StructuralIdentifiability.jl [17], require rational right-hand sides. The model is not rational as written, but it can be lifted exactly:

- under a known temperature profile $T(t)$ within one regime, $\tilde{k}(t)$ obeys $\mathrm{d}\tilde{k}/\mathrm{d}t = E_a \dot{T}\, \tilde{k}/(RT^2)$;
- for $n = 3/2$, $w = y_\beta^{1/2}$ obeys $\mathrm{d}w/\mathrm{d}t = -\frac{1}{2}\tilde{k}\, w^2$;
- the buffering term $z = \exp(-(\mathrm{pH} - \mathrm{pH_{CCP}})^2/2\sigma_{\mathrm{buf}}^2)$ obeys $\mathrm{d}z/\mathrm{d}t = -(\mathrm{pH} - \mathrm{pH_{CCP}})\, z\, \dot{\mathrm{pH}}/\sigma_{\mathrm{buf}}^2$, so that $\beta_{\mathrm{buf}} = \beta_0 + \beta_{\mathrm{CCP}} z$;
- the inhibition term with non-integer $m_p = 1.2$ is handled by $v = (1 - P/P_{\max})^{m_p}$, with $\mathrm{d}v/\mathrm{d}t = -m_p\, v\, \dot{P}/(P_{\max} - P)$.

The softplus break is not rational, so a symbolic analysis has to be run regime by regime. The $\max(0, \dot{X})$ term reduces to $\dot{X}$ wherever biomass is growing. At the nominal point, $\dot{X}$ stays positive over the whole batch: its minimum is 1.07×10^{-5} kg m^{-3} s^{-1}.

3.3 Symbolic results

I ran StructuralIdentifiability.jl [17] (version 0.6.0, Julia 1.13.4 from conda-forge) on the lifted models. Its answers are probabilistic, with a correctness probability of at least 0.99 by default.

Denaturation, one regime, known temperature profile. With $u(t) = \dot{T}/T^2$ as a known input, $a = E_a/R$ as parameter and the native fraction as output, a is globally identifiable for $n = 1$, $n = 3/2$ and $n = 2$. The rate $\tilde{k}$ and the native fraction are globally identifiable for $n = 1$ and $n = 2$. For $n = 3/2$ they are locally identifiable, because the lifted variable $w = y_\beta^{1/2}$ is only fixed up to its sign, which the physics settles ($w > 0$). Under isothermal conditions a disappears from the model and only $\tilde{k}$ remains identifiable, as expected. One non-isothermal curve therefore carries information on the activation energy of the regime it covers.

Fermentation. I analysed the lifted seven-state model (biomass, lactose, acid, pH, buffering shape z, inhibition variable v, lag variable q) at T_{opt}, with every constant except $\mu_{\max}$ and β_{CCP} fixed to its nominal value, all initial conditions unknown and pH as the only output. $\mu_{\max}$ and $q(t)$ are locally identifiable, which includes q_0. β_{CCP} is not, but the product $\beta_{\mathrm{CCP}}\, z(t)$ is. This is a scaling symmetry of the lifting, $(\beta_{\mathrm{CCP}}, z) \mapsto (c\beta_{\mathrm{CCP}}, z/c)$, and it disappears once $z(0) = \exp(-(\mathrm{pH}_0 - \mathrm{pH_{CCP}})^2/2\sigma_{\mathrm{buf}}^2)$ is known, as it is when the initial pH is measured. The version of the package I used accepts known initial conditions only in the global test. That test, run with all initial conditions unknown, had not finished within the 50-minute limit I set for the seven-state model, so I report the local results only.

3.4 Gradient of the stopping time

Fermentation data often give only the time to pH 4.6. To use such data with a gradient-based sampler, one needs the derivative of t_{ferm} with respect to the parameters.

Proposition 2. *Let* $\Phi(t,\theta) = \mathrm{pH}(t;\theta) - 4.60$ *and assume* $\beta_{\mathrm{LP}} > 0$, $X(0) > 0$, *and* $\dot{X} > 0$ *on a neighbourhood of* t_{ferm}. *Then* $t_{\mathrm{ferm}}(\theta)$ *is locally unique and continuously differentiable, with*

$$\nabla_\theta t_{\mathrm{ferm}} = -\frac{\nabla_\theta \mathrm{pH}(t_{\mathrm{ferm}};\theta)}{\dot{\mathrm{pH}}(t_{\mathrm{ferm}})}. \tag{12}$$

Proof. Where $\dot{X} > 0$ the right-hand side is smooth, so Φ is C^1 in (t,θ). Since $\dot{X} \geq -k_{\mathrm{dec}} X$, we have $X(t) \geq X(0)e^{-k_{\mathrm{dec}}t} > 0$, hence $\dot{P} \geq \beta_{\mathrm{LP}} X > 0$ and $\dot{\mathrm{pH}} < 0$: the crossing is transversal. The implicit function theorem applied to $\Phi(t_{\mathrm{ferm}}(\theta),\theta) = 0$ gives the result. □

This is the standard event-sensitivity result for hybrid systems [18], stated here for the pH stopping condition. It is what makes it possible to mix full pH curves and endpoint-only records in one likelihood.

4 Hierarchical model

Both blocks use non-centred parameterisations, to avoid the funnel geometry that slows NUTS [19, 20].

Block A: denaturation across studies. For study i, $\psi_i = (\ln \tilde{k}_{b,i}, E_{a,1,i}, E_{a,2,i}, \eta_{n,i}, T_{b,i})$ with $n_i = 0.5 + 2\,\mathrm{logistic}(\eta_{n,i})$, and

$$\psi_i = \mu_\psi + X_i\gamma + \mathrm{diag}(\tau_\psi)\, z_i, \qquad z_i \sim \mathcal{N}(0, I_5), \tag{13}$$

where X_i codes the matrix (skim milk, whole milk, whey or buffer) and the assay. The interval $(0.5, 2.5)$ contains every order reported in Table 2, including exact first order. Section 5 fits this block without covariates, because seven entries cannot support them. When raw curves are available, the observation variance has a digitisation term:

$$y_{ij}^{\mathrm{obs}} \sim \mathcal{N}\big(y_\beta(t_{ij}; T_{ij}, \psi_i),\ \sigma_{\mathrm{tab}}^2 + \mathbb{1}_{\mathrm{fig},ij}\,\sigma_{\mathrm{dig}}^2\big), \tag{14}$$

In principle σ_{dig} can be estimated by digitising a subset of curves twice. The only curves used in Section 5 were digitised twice in sequence, first by Loveday from the original figures and then by me from his re-plot, so the two error sources cannot be separated and I estimate a single noise term.

Block B: batches. For batch b, $\ln \mu_{\max,b}$, $\ln \beta_{\mathrm{CCP},b}$ and $\ln q_{0,b}$ are normal around population means. A batch contributes either a pH curve, $y_{bj} \sim \mathcal{N}(\mathrm{pH}(t_{bj};\theta_b), \sigma_{\mathrm{pH}}^2)$, or only its endpoint, $t_b^{\mathrm{obs}} \sim \mathcal{N}(t_{\mathrm{ferm}}(\theta_b), \sigma_t^2)$, with the gradient from (12).

Diagnostics. I plan four chains, with split-$\hat{R} < 1.01$, bulk and tail effective sample sizes above 400 [21], no divergences, simulation-based calibration [22], leave-one-study-out validation and prior sensitivity on T_b and τ_ψ.

5 Meta-analysis of published denaturation kinetics

5.1 Data

Table 2 lists the seven entries. They come from four studies and are of three kinds.

Curves. Dannenberg and Kessler printed their native fractions as figures. Loveday digitised them and re-plotted the 12 isothermal curves (70 to 150 °C) on a single untransformed graph, Figure 1 of [4]. I digitised that figure by colour segmentation of the markers, after calibrating the axes on the tick marks. Six colours are each shared by two temperatures, and I assigned those markers by their position. All 107 markers were recovered, and I checked each one on an overlay of the source image. One pixel is 0.0017 in native fraction and 1.5 % in time. As a check, I fitted an n-th-order curve to each temperature separately. The orders ranged from 1.58 ± 0.24 to 1.97 ± 0.55, against 1.57 ± 0.17 to 1.96 ± 0.63 in Loveday's own re-analysis. The caption does not name the variant. Loveday compares his fit with the original activation energies of 279.96 and 47.75 kJ mol^{-1}, which are the values that Oldfield's review lists for β-Lg B (Table 2.5 of [6]), so I treat the curves as variant B. I also follow Loveday in treating each holding as isothermal. For variant A, only the order and activation energies from Oldfield's Table 2.5 are available to me.

Rate constants. Claeys et al. report nine isothermal decimal reduction times, each with its standard error, between 70 and 95 °C (Table 3 of [10]). I converted them to first-order rates, $k = \ln 10 / D$, with a standard error on $\ln k$ of SE_D / D. The z values in the same table are derived from these D values, so I do not use them a second time.

Regressions and fitted parameters. Manji and Kakuda give, for each variant and each range (70–90 and 90–140 °C), a straight line $\log_{10}(C_0 k_2) = a - b\,(10^3/T)$ with its root-mean-square error (Table I of [3]). The product $C_0 k_2$ is exactly the normalised rate $\tilde{k}$ of (1) for $n = 2$. A straight line is summarised without loss by its value at the centroid of $1/T$ and its slope, which are uncorrelated. I computed both, and their standard errors, from the reported RMSE. For that I assumed temperatures of 70, 80, …, 140 °C with three rate constants each; the paper states eight temperatures over that range and three rate constants per temperature. Their slope for variant B above the break gives 47.5 kJ mol^{-1}, not the 46.0 that Oldfield's review lists. Oldfield's thesis reports, for each variant and each range, the order, the activation energy, $\ln k_0$, $\ln k_{\mathrm{ref}}$ and the fitted initial concentration, all with 95 % intervals (Table 4.4 of [6]). The reference temperatures follow from $T_{\mathrm{ref}} = E_a / [R(\ln k_0 - \ln k_{\mathrm{ref}})]$. They lie between 80 and 82.5 °C for the lower range and close to 112.5 °C for the upper one. I converted the rates to the normalised form using the fitted initial concentrations.

Where a source gives no uncertainty, I had to assume one. This applies to the reported orders, which were chosen by comparing integer orders or by visual linearity, and to the values taken from Oldfield's review. The table marks these.

Table 2: The seven entries. $E_{a,1}$, $E_{a,2}$: activation energies below and above the break (kJ mol^{-1}), with 95 % intervals as reported or derived. Asterisks mark assumed uncertainties or values, daggers values I derived from the published regression.

Entry	Milk, heating	Data used	n	$E_{a,1}$	$E_{a,2}$	Break [°C]
Oldfield 1996, A	skim, direct steam injection	order, E_a and $\ln k_{\mathrm{ref}}$ for each range	1.0±0.3; 1.2±0.2	285.5±30.8	58.8±5.7	90–95
Oldfield 1996, B	same	same	1.4±0.4; 1.6±0.2	296.7±42.2	44.0±5.8	90–95
Claeys 2001	raw, vials in water bath	nine D values ± SE, 70–95 °C	1 (SE 0.1*)	from rates	from rates	80–83
Dannenberg & Kessler 1988, B	skim, pilot plant	107 native fractions, 12 temperatures, 70–150 °C	from curves	from curves	from curves	90–95
Dannenberg & Kessler 1988, A	same	order and E_a, via [6]	1.5 (SE 0.1*)	265.2 (SE 10 %*)	54.1 (SE 10 %*)	90–95
Manji & Kakuda 1986, A	skim, capillary tubes	Arrhenius lines with RMSE, two ranges	2 (SE 0.1*)	276.6±10.1†	46.1±5.4†	85–95*
Manji & Kakuda 1986, B	same	same	2 (SE 0.1*)	258.7±11.3†	47.5±4.5†	85–95*

5.2 Model

Each entry j has parameters $(\ln \tilde{k}_j(363.15\,\mathrm{K}), E_{a,1,j}, E_{a,2,j}, n_j, T_{b,j})$. They are drawn from normal population distributions in non-centred form, with $n_j = 0.5 + 2\,\mathrm{logistic}(\eta_j)$, so that orders between 0.5 and 2.5 are allowed. The likelihood combines whatever each entry provides.

- A native fraction at time t and temperature T is a normal observation of the n-th-order solution of (1), with a common standard deviation σ_c.
- A rate constant is a normal observation of $\ln \tilde{k}_j(T)$ from (3). Its variance is the reported one plus a discrepancy term $\sigma^2_{\ln k}$, which absorbs what a sharp two-line Arrhenius law cannot represent; Oldfield, for instance, fits a separate order in each range.
- Reported orders and activation energies are normal observations of n_j, $E_{a,1,j}$ and $E_{a,2,j}$.
- The break temperature is interval-censored, because the sources only say between which temperatures the slope changes. I use a soft interval with edges of width 0.5 K.

The hyperpriors are weakly informative: normal means centred on 270 and 60 kJ mol^{-1}, 361 K and -4, with standard deviations of 60, 40, 6 K and 2, and half-normal scales. I wrote the hyper-means themselves in non-centred form and started the chains at prior medians. Without this, chains started from the default initialisation got stuck far from the data. Variants A and B of one study are treated as exchangeable entries, which ignores their shared equipment and assay. I ran four NUTS chains with 1500 warm-up and 2000 retained draws each and a target acceptance of 0.99. There were no divergences, $\hat{R} \le 1.002$ for all hyperparameters, and every effective sample size was above 1400.

5.3 Results

Table 3 gives the population parameters. Figure 1 shows the rate data with the fitted entries, and the Dannenberg and Kessler curves with their fit. Figure 2 compares reported and posterior activation energies.

Table 3: Population parameters (median and 95 % interval).

	Mean	Between-entry s.d.
$\ln \tilde{k}$ at 90 °C [s^{-1}]	-3.94 [-4.60, -3.28]	0.70 [0.40, 1.44]
$E_{a,1}$ [kJ mol^{-1}]	280.0 [261.1, 299.5]	17.6 [8.0, 41.3]
$E_{a,2}$ [kJ mol^{-1}]	50.6 [44.3, 61.7]	6.2 [1.5, 24.8]
n	1.56 [1.21, 1.90]	–
T_b [°C]	89.6 [86.6, 92.4]	3.4 [1.9, 6.3] (K)
$\sigma_{\ln k}$ (discrepancy)	0.24 [0.11, 0.42]	
σ_c (curves)	0.034 [0.030, 0.040]	

The curve noise is 0.034 in native fraction; it covers the assay and both digitisations together. The fit in Figure 1b holds from 70 to 150 °C with one order and one Arrhenius law per regime. That order is 1.71 (1.60–1.82), higher than the 1.5 that Dannenberg and Kessler obtained from linearised plots, and in line with Loveday's non-linear re-analysis. For this entry the break lies at 89.2 °C (88.6–90.1).

Below the break the studies agree closely. With rate data in the model, Claeys et al. give 292 kJ mol^{-1}, and Manji and Kakuda's two variants differ by less than 20 kJ mol^{-1}. Above the break the agreement is also good once the Claeys entry is fitted on its rates rather than on its z value: its upper activation energy is 54 kJ mol^{-1}, with a wide interval (42–100), because few of

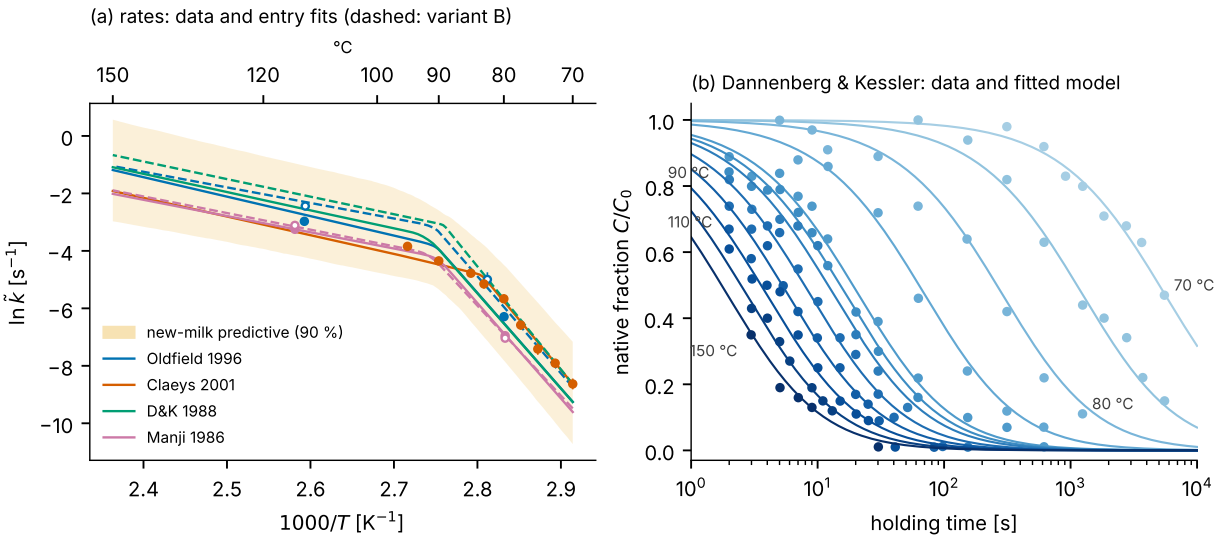


Figure 1: (a) Normalised rate constants against $1000/T$: Claeys et al. (nine points), the centroid values of the Manji and Kakuda regressions and the reference rates of Oldfield (open symbols: variant B; bars: ±1.96 SE), with the posterior median of each entry (dashed: variant B) and the 90 % predictive band for a new milk. (b) The Dannenberg and Kessler curves as digitised from [4], from 70 °C (light) to 150 °C (dark), with the posterior median of the model.

its rates lie above the break. The real outlier is the Claeys break itself, at 83.2 °C (81.8–84.2), against 89 to 92 °C for the other entries. The order also varies, from about 1 (Claeys, Oldfield A) to 2 (Manji and Kakuda).

The level of the rate is what matters most for design, and it varies most. At 90 °C the entry medians of $\ln \tilde{k}$ run from −4.47 (Claeys) to −3.05 (Dannenberg and Kessler B), a factor of four in rate. The capillary tubes of Manji and Kakuda and the vials of Claeys et al. give the slowest rates, and the pilot plants of Oldfield and of Dannenberg and Kessler the fastest. With rate data from six entries, the between-entry standard deviation of $\ln \tilde{k}$ is now reasonably well determined (0.40–1.44).

For a milk that has not been characterised, denaturation after 300 s at 90 °C has a median of 0.93 and a 95 % interval from 0.54 to 1.00. The calibration used in [1], $D_\beta = 0.80$, falls at the 18th percentile. I also refitted the model six times, leaving out in turn each entry that carries rates or curves. Each of the 17 held-out rate constants fell inside its 90 % predictive interval. With Dannenberg and Kessler left out, 96 % of their 107 points fell inside their pointwise 90 % intervals. The rate intervals are 2.8 to 3.7 units of $\ln k$ wide, so the check mainly shows that the model is not overconfident.

6 Chance-constrained optimisation

With decision vector $x = (T_{\text{past}}, t_{\text{hold}}, T_{\text{ferm}}, \dot{\gamma})$ and uncertain parameters Θ, the expected objective keeps the weights and references of [1]:

$$\min_x \ \mathbb{E}_\Theta\Big[0.4\,\frac{t_{\text{ferm}}}{t_{\text{ref}}} + 0.4\,\frac{S_{\text{syn}}}{S_{\text{ref}}} + 0.2\,\frac{E}{E_{\text{ref}}}\Big] \quad \text{s.t.} \quad \mathbb{P}_\Theta(\eta_{50} \geq 0.45\ \text{Pa s}) \geq 0.95. \tag{15}$$

The references are the values at the nominal point (90 °C, 300 s, 42 °C, 140 s^{-1}) with nominal parameters. I also consider an optional constraint $\mathbb{P}(D_\beta \leq D_{\max}) \geq 0.95$. It is meant as a plant-imposed ceiling standing in for effects the model does not represent, such as fouling of the heater or cooked flavour. I have not established a link between a given value of D_β at the end of

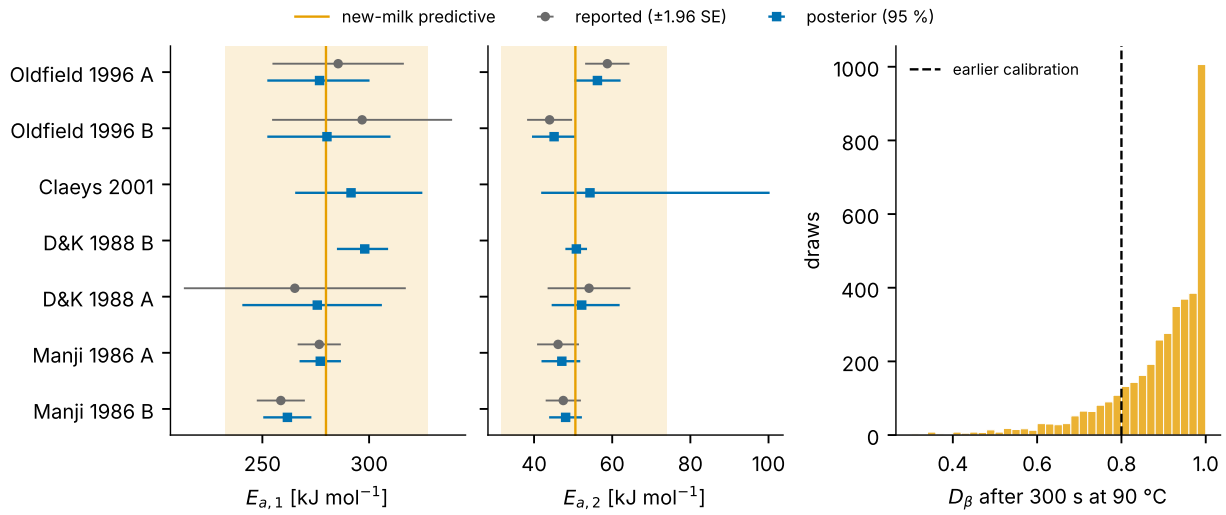


Figure 2: Activation energies as reported or derived (±1.96 standard errors, grey; none for entries fitted on rates or curves), entry posteriors (95 %, blue) and the predictive for a new milk (median and 95 % band, orange). Right: predictive distribution of D_β after 300 s at 90 °C for a new milk, with the calibration of [1].

holding and fouling, so $D_{\max}$ is a scenario parameter. Tank volume is not constrained: I report its 95th percentile as a sizing result.

7 Numerical verification

The scripts use SciPy (DOP853 with event location) for forward runs, JAX with Diffrax for sensitivities, and NumPyro for NUTS [23, 24].

7.1 Gradient of the stopping time

At 42 °C with $\theta = (\mu_{\max}, \beta_{\mathrm{CCP}}, q_0) = (3.2 \times 10^{-4}, 28, 10)$, the batch reaches pH 4.60 at $t_{\mathrm{ferm}} = 15\,033.46$ s, where $\dot{\mathrm{pH}} = -3.262 \times 10^{-4}$ s^{-1}. I obtained $\nabla_\theta \mathrm{pH}(t_{\mathrm{ferm}})$ by forward-mode differentiation through the solver and applied (12). The reference is a central finite difference of the event time (relative step 10^{-5}, tolerances 10^{-12}). Table 4 shows the agreement.

Table 4: Sensitivity of the stopping time: implicit formula (12) against central finite differences.

Parameter	Eq. (12)	Finite difference	Relative error
$\mu_{\max}$ [s^2]	-4.80926×10^7	-4.80926×10^7	1.4×10^{-9}
β_{CCP} [s per mol m^{-3} pH^{-1}]	87.1438	87.1438	1.5×10^{-10}
q_0 [s]	−29.3196	−29.3196	1.7×10^{-9}

A 1 % drop in $\mu_{\max}$ thus lengthens the batch by about 154 s, while a 1 % rise in β_{CCP} lengthens it by about 24 s.

7.2 Local identifiability of the denaturation block

I computed the Fisher information of $(\ln \tilde{k}_b, E_{a,1}, E_{a,2}, n, T_b)$ for synthetic isothermal designs: seven sampling times from 30 to 1800 s, $\sigma = 0.02$ on y_β, and the smooth model with $\tau = 10^{-6}$ K^{-1}.

- **Bracketing design** (75, 80, 85, 88, 92, 95 and 100 °C): the information matrix is regular (condition number 2.8×10^3). Standard errors are 0.061 on $\ln \tilde{k}_b$, 5.7 and 7.7 kJ mol^{-1} on $E_{a,1}$

and $E_{a,2}$, 0.045 on n and 0.31 K on T_b. The largest correlations are -0.78 ($\ln \tilde{k}_b$, $E_{a,2}$) and 0.78 ($\ln \tilde{k}_b$, T_b).

- **One-sided designs** (only above or only below T_b): the matrix is singular to machine precision (condition number about 10^{19}), with $\ln \tilde{k}_b$, T_b and the opposite activation energy confounded.

This agrees with the analytic argument of Section 3. It also matters for the meta-analysis: a study that only covers one side of the break says nothing about the other side, and in a hierarchical model it can only borrow that information from the other studies.

7.3 Recovery on a synthetic batch

I generated 15 pH values between 1008 s and 4.2 h at 42 °C from $\theta_{\mathrm{true}} = (2.95 \times 10^{-4}, 30.5, 10)$, with Gaussian noise $\sigma_{\mathrm{pH}} = 0.015$. The priors were log-normal around the nominal values, with log standard deviations of 0.2, 0.3 and 0.7. I ran four NUTS chains (500 warm-up and 1000 retained draws each, target acceptance 0.9). There were no divergences, $\hat{R} \le 1.004$ and effective sample sizes between 735 and 841. Table 5 and Figure 3 summarise the posterior.

Table 5: Posterior on a synthetic batch (median and 95 % interval).

	True value	Posterior median	95 % interval
$\mu_{\max}$ [10^{-4} s^{-1}]	2.95	2.933	[2.859, 3.079]
β_{CCP} [mol m^{-3} pH^{-1}]	30.5	30.05	[26.29, 35.79]
q_0 [–]	10	11.5	[4.5, 37.6]
t_{ferm} [h]	4.605	4.605	[4.569, 4.646]

The three parameters are far from independent. The posterior correlations of the logarithms are +0.86 between $\mu_{\max}$ and β_{CCP}, -0.90 between $\mu_{\max}$ and q_0, and -0.64 between β_{CCP} and q_0. A faster starter can be offset by a stronger buffer or a longer lag. The data constrain the time to pH 4.6 much better than any single parameter: the 95 % predictive interval on t_{ferm} stays within 1 % of the median on either side, while q_0 is barely learnt beyond its prior. In practice an independent titration curve, which informs β_{CCP} directly, would be the most useful extra measurement.

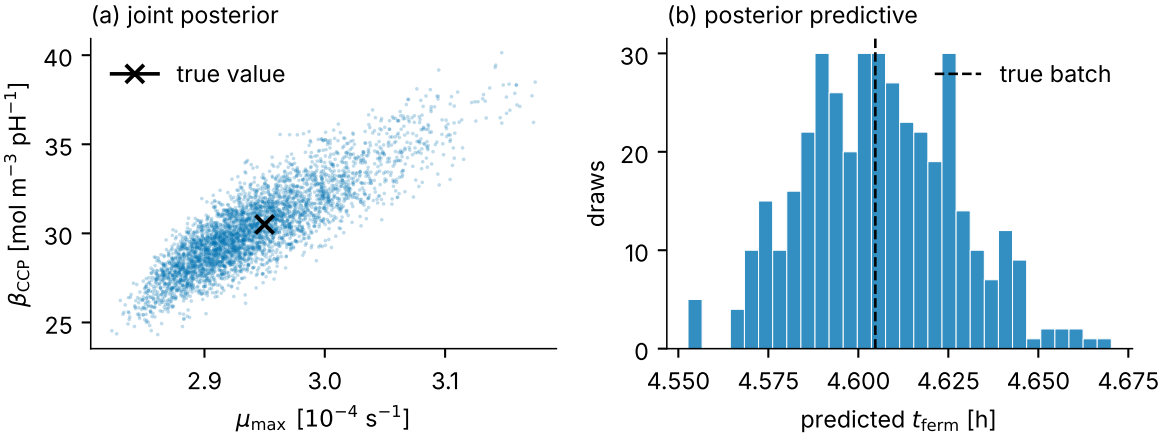


Figure 3: Synthetic batch. (a) Joint posterior draws of $\mu_{\max}$ and β_{CCP}; the cross marks the true values. (b) Posterior predictive distribution of the time to pH 4.60 (400 draws); the dashed line is the true batch.

7.4 Propagation and the robust optimum

Set-up. For the denaturation block I drew 4000 parameter sets from the predictive distribution of a new entry (Section 5), which represents a milk whose kinetics have not been measured. About 11 % of these draws have $n < 1$. For them the native fraction reaches zero in finite time, and I clip it there. For batch variability I have no data, so I kept an illustrative prior: log-normal $\mu_{\max}$, β_{CCP} and q_0 with log standard deviations of 0.08, 0.10 and 0.7, and 1000 draws. The search grid covered 85 to 95 °C in steps of 0.5 K, 180 to 600 s in steps of 20 s, 39 to 44 °C in steps of 0.25 K, and six shear rates between 60 and 250 s^{-1}.

Plant calibration. A plant can measure the denaturation of its own milk. I simulated two such measurements, of D_β after 300 s at 90 and 95 °C, each with a standard error of 0.03. Having no plant data, I set them to the predictions of the earlier model (0.800 and 0.849). This amounts to assuming that the plant's milk behaves like the one in the earlier calibration. I then reweighted the new-milk draws by the likelihood of the two values (importance sampling, effective sample size 379 out of 4000). After reweighting, D_β after 300 s at 90 °C is 0.807 (95 % interval 0.754–0.852). For this step I fixed $T_{\mathrm{ferm}} = 41$ °C and $\dot{\gamma} = 60$ s^{-1}, the values found at every optimum without calibration, and searched over temperature and holding time only. Table 6 gives the results.

Table 6: Operating points. P5 and P95 are the 5th and 95th percentiles. "New milk": predictive of the meta-analysis. "Calibrated": the same draws reweighted by two plant measurements. "Det." is the deterministic optimum with nominal parameters and $\eta_{50} \geq 0.45$ Pa s; CC is the solution of (15), without or with a ceiling $D_{\max}$. Expected objectives can only be compared within the same distribution.

	New milk		Calibrated		
	Nominal	Det. = CC (no ceiling)	CC (no ceiling)	CC, $D_{\max} = 0.95$	CC, $D_{\max} = 0.92$
T_{past} [°C]	90.0	95.0	95.0	91.5	92.5
t_{hold} [s]	300	600	600	580	420
T_{ferm} [°C]	42.0	41.0	41.0	41.0	41.0
$\dot{\gamma}$ [s^{-1}]	140	60	60	60	60
Expected objective	0.944	0.869	0.891	0.901	0.919
D_β median [P5, P95]	0.926 [0.637, 1.000]	0.988 [0.852, 1.000]	0.927 [0.896, 0.967]	0.911 [0.878, 0.949]	0.883 [0.847, 0.919]
$\mathbb{P}(\eta_{50} \geq 0.45)$	0.816	0.994	1.000	1.000	1.000
$\mathbb{P}(D_\beta \leq 0.92)$	0.473	0.139	0.347	0.684	0.954
t_{ferm} median / P95 [h]	4.20 / 4.84	4.23 / 4.88	4.23 / 4.88	4.23 / 4.88	4.23 / 4.88
V_{tank} median / P95 [m^3]	20.3 / 22.5	20.4 / 22.6	20.4 / 22.6	20.4 / 22.6	20.4 / 22.6

Uncertainty on its own does not move the optimum. With the new-milk predictive, the deterministic grid optimum (95 °C, 600 s, 41.0 °C, 60 s^{-1}) also solves (15). Syneresis falls as D_β rises, and that gain outweighs the small energy cost of longer holding. At that point the viscosity constraint holds in more than 99 % of draws.

A ceiling on denaturation is a different matter. Without plant data it cannot be met at all. The viscosity chance constraint alone is satisfied on 265 of the 462 grid points, all at 86.5 °C or above. None of them also satisfies $\mathbb{P}(D_\beta \leq 0.95) \geq 0.95$, let alone the 0.92 ceiling. The spread between milks is simply too wide: a treatment gentle enough for the fast milks leaves the slow ones short of viscosity.

The two calibration measurements change this. They narrow the P5–P95 range of D_β from about 0.36 at the nominal point to about 0.07 at the operating points of Table 6, and both ceilings become feasible. With $D_{\max} = 0.92$ the optimum moves to 92.5 °C and 420 s. The expected objective is then 3.1 % higher than at the calibrated optimum without a ceiling. Whether a ceiling is justified at all needs a physical argument that this work does not supply.

Batch variability shows up most directly in tank sizing. Figure 4b gives the volume needed per tank at 41 °C. The earlier deterministic sizing, 19.95 m^3, was computed without the lag term. It is too small for 62 % of simulated batches, and covering 95 % of batches takes 22.6 m^3 per tank

with four tanks. This figure depends on the illustrative batch prior.

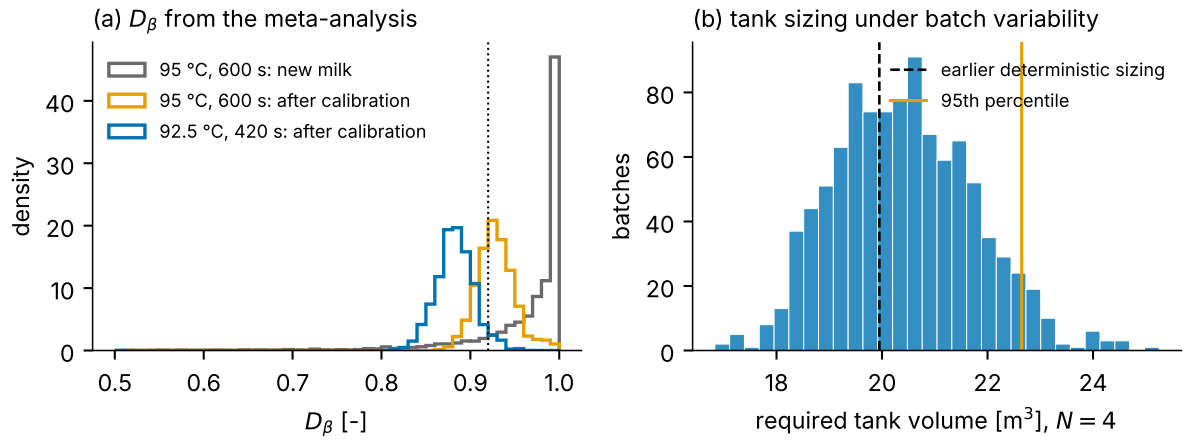


Figure 4: Propagation. (a) Distribution of D_β at 95 °C and 600 s for a new milk (grey) and after calibration on two plant measurements (orange, weighted), and at the calibrated chance-constrained point with $D_{\max} = 0.92$ (92.5 °C, 420 s, blue); the dotted line marks 0.92. (b) Required volume per tank at $T_{\text{ferm}} = 41$ °C for 1000 simulated batches ($N = 4$ tanks), with the earlier deterministic sizing (dashed, computed without lag) and the 95th percentile.

8 Data inventory

Table 7 lists the sources and what I used from each. Of the denaturation studies, I have read Manji and Kakuda and Claeys et al. in full, the Oldfield thesis, and Loveday's author post-print, which is distributed under a CC BY-NC-ND licence. I could not reach the full text of Anema and McKenna, nor the original paper of Dannenberg and Kessler. The fermentation files on Mendeley could not be downloaded from the environment used for this work.

The digitised Dannenberg and Kessler points are in `dk_points.csv`, produced by `digitise_dk.py`; the other study-level data are in `meta_data.py`.

9 Discussion and limitations

The meta-analysis puts numbers on the familiar remark that published kinetics disagree, and it shows where the disagreement lies. The activation energies agree well on both sides of the break. The order and the break temperature disagree more. The level of the rate disagrees most, by a factor of four between the slowest and the fastest set-up. Part of that spread probably comes from the heating method (capillary tubes or vials against pilot plants) and part from the assay, but seven entries are too few to separate the two. Propagated through the line model, the spread has a concrete consequence. For a milk nobody has measured, the denaturation reached at a given heat treatment is uncertain enough that no operating point can guarantee both the viscosity target and a denaturation ceiling. Two measurements on the plant's own milk remove most of that uncertainty. If I had to give one design recommendation from this work, it would be to calibrate the denaturation kinetics on the milk actually processed before trusting a model-based heat treatment.

The symbolic analysis points the same way. A non-isothermal curve identifies the activation energy of the regime it covers, and an isothermal curve does not. So a plant calibration that logs the temperature profile through the heater and holding tube tells more than a laboratory isothermal test of the same cost. On the fermentation side, pH data identify the growth rate and

Table 7: Data sources for the two blocks.

Source	System	What I checked	Still to check
Dannenberg & Kessler 1988 [2]	β-Lg A, B in skim milk	used: curves for one variant (B, inferred), digitised from Loveday's re-plot [4]; variant A via Oldfield's Table 2.5	variant of the curves; heating profile of the pilot plant
Manji & Kakuda 1986 [3]	β-Lg A, B in skim milk, capillary tubes	used: Table I (Arrhenius lines with RMSE, 70–90 and 90–140 °C)	exact temperature grid; individual rate constants (only plotted)
Anema & McKenna 1996 [5]	β-Lg A, B in reconstituted whole milk, pilot plant	order 1.5, abrupt change of Arrhenius slope (abstract)	full text not accessible; not used
Oldfield 1996 [6]; Oldfield et al. 1998 [7]	β-Lg in skim milk, direct steam injection	used: Table 4.4 of the thesis (order, E_a, $\ln k_0$, $\ln k_{\mathrm{ref}}$, C_0 with 95 % intervals for 70–90 and 95–130 °C)	raw data in the thesis appendices
Galani & Owusu Apenten 1997 [8]	whey protein mixture in Tris-HCl buffer	75–120 °C; first order 75–85, second order 90–120 °C	data format
Galani & Apenten 1999 [9]	whey protein mixture, pH 6.7	GP-FPLC and SDS-PAGE; biphasic Arrhenius plots with different breaks	data format
Claeys et al. 2001 [10]	β-Lg in raw milk	used: Table 3, nine isothermal D values with standard errors, 70–95 °C	raw retention data
Ge & Liu 2023 [25]	*S. thermophilus* CICC 6038 + *L. bulgaricus* CICC 6047	Mendeley Data v2, CC BY 4.0; time to pH 4.6 for six inoculum ratios; post-acidification	whether full pH curves are included
Zanatta & Basso 1992 [26]	37 *S. thermophilus* strains	acidification parameters from pH curves (max. rate and its time)	whether curves or only parameters are given
Béal & Corrieu 1991 [27]	*S. thermophilus* 404 + *L. bulgaricus* 398	effects of pH, temperature and inoculum composition on mixed cultures	data format
ComBase	microbial growth and survival curves	open database of over 60 000 records	whether lactic starter pH curves exist

the lag. The buffering peak is identified only through its product with the buffering shape, and a measured initial pH resolves that.

I had expected, when planning this work, that uncertainty alone would pull the optimum away from the bounds. In this model it does not, because nothing in the objective penalises high denaturation.

Several limitations remain. Only one study contributes raw curves, and those were digitised twice in sequence; their variant is inferred, and the holding is treated as isothermal although the pilot plant needed a few seconds to come up to temperature, which matters most for the shortest times at 130 to 150 °C. For Manji and Kakuda I had to assume the temperature grid to recover standard errors. Dannenberg and Kessler A still rests on a secondary source. The reported orders carry uncertainties I assumed. Variants A and B of one study are treated as exchangeable, which probably understates the variation between studies. I cannot tell from the paper whether the low break of Claeys et al. reflects raw milk, the assay or the heating method. On the fermentation side, batch variability and the lag term are still illustrative, so the tank-sizing result demonstrates the method rather than giving a design figure. The plant calibration is simulated, and its importance sampling rests on 379 effective draws. The texture and syneresis relations are unfitted, as in [1]. The global symbolic test of the fermentation model did not finish, so those results are local. Finally, the grid search resolves 0.5 K and 20 s.

10 Conclusion

I have fitted a hierarchical Bayesian model to four studies of β-Lg denaturation in milk, using curves, rate constants and regression summaries in the same likelihood. I have also run a symbolic identifiability analysis of the lifted line model and carried the result through a chance-constrained

optimisation. Without measurements on the plant's milk, a denaturation ceiling cannot be combined with the viscosity target. With two such measurements it can. The next steps are to add the Anema and McKenna data and the curves of Oldfield's thesis, to model the heating profile of the pilot-plant studies, and to find fermentation data with full pH curves for block B.

Data and code availability

The scripts that produce every number and figure (`model.py`, `jaxmodel.py`, `ift_check.py`, `fim.py`, `recover.py`, `digitise_dk.py`, `dk_points.csv`, `meta_data.py`, `meta.py`, `mc2.py`, `mc3.py`, `figs.py`, `figs_meta.py`, and the Julia scripts `siA.jl`, `si2.jl`) and their outputs are provided as supplementary material, and will be deposited on Zenodo.

Conflict of interest

The author declares no conflict of interest and received no funding for this work.

References

[1] A. Benahmed, A coupled mechanistic and reservoir-computing model of a stirred-yoghurt line built from public data, preprint, 2026.

[2] F. Dannenberg, H. G. Kessler, Reaction kinetics of the denaturation of whey proteins in milk, *Journal of Food Science* 53 (1988) 258–263. doi:10.1111/j.1365-2621.1988.tb10223.x

[3] B. Manji, Y. Kakuda, Thermal denaturation of whey proteins in skim milk, *Canadian Institute of Food Science and Technology Journal* 19 (1986) 163–166. doi:10.1016/S0315-5463(86)71624-6

[4] S. M. Loveday, β-Lactoglobulin heat denaturation: a critical assessment of kinetic modelling, *International Dairy Journal* 52 (2016) 92–100. doi:10.1016/j.idairyj.2015.08.001

[5] S. G. Anema, A. B. McKenna, Reaction kinetics of thermal denaturation of whey proteins in heated reconstituted whole milk, *Journal of Agricultural and Food Chemistry* 44 (1996) 422–428. doi:10.1021/jf950217q

[6] D. J. Oldfield, Heat-induced whey protein reactions in milk: kinetics of denaturation and aggregation as related to milk powder manufacture, PhD thesis, Massey University, Palmerston North, 1996.

[7] D. J. Oldfield, H. Singh, M. W. Taylor, K. N. Pearce, Kinetics of denaturation and aggregation of whey proteins in skim milk heated in an ultra-high temperature (UHT) pilot plant, *International Dairy Journal* 8 (1998) 311–318. doi:10.1016/S0958-6946(98)00089-2

[8] D. Galani, R. K. Owusu Apenten, The comparative heat stability of bovine β-lactoglobulin in buffer and complex media, *Journal of the Science of Food and Agriculture* 74 (1997) 89–98.

[9] D. Galani, R. K. O. Apenten, Heat-induced denaturation and aggregation of β-lactoglobulin: kinetics of formation of hydrophobic and disulphide-linked aggregates, *International Journal of Food Science and Technology* 34 (1999) 467–476.

[10] W. L. Claeys, L. R. Ludikhuyze, A. M. Van Loey, M. E. Hendrickx, Inactivation kinetics of alkaline phosphatase and lactoperoxidase, and denaturation kinetics of β-lactoglobulin in raw milk under isothermal and dynamic temperature conditions, *Journal of Dairy Research* 68 (2001) 95–107. doi:10.1017/S002202990000460X

[11] N.-D. M. Luong, L. Coroller, M. Zagorec, N. Moriceau, V. Anthoine, S. Guillou, J.-M. Membré, A Bayesian approach to describe and simulate the pH evolution of fresh meat products depending on the preservation conditions, *Foods* 11 (2022) 1114. doi:10.3390/foods11081114

[12] L. Rosso, J. R. Lobry, J. P. Flandrois, An unexpected correlation between cardinal temperatures of microbial growth highlighted by a new model, *Journal of Theoretical Biology* 162 (1993) 447–463.

[13] R. Luedeking, E. L. Piret, A kinetic study of the lactic acid fermentation. Batch process at controlled pH, *Journal of Biochemical and Microbiological Technology and Engineering* 1 (1959) 393–412.

[14] J. Baranyi, T. A. Roberts, A dynamic approach to predicting bacterial growth in food, *International Journal of Food Microbiology* 23 (1994) 277–294.

[15] W. H. Herschel, R. Bulkley, Konsistenzmessungen von Gummi-Benzollösungen, *Kolloid-Zeitschrift* 39 (1926) 291–300.

[16] F. Salaün, B. Mietton, F. Gaucheron, Influence of mineral environment on the buffering capacity of casein micelles, *Milchwissenschaft* 62 (2007) 20–23.

[17] R. Dong, C. Goodbrake, H. A. Harrington, G. Pogudin, Differential elimination for dynamical models via projections with applications to structural identifiability, *SIAM Journal on Applied Algebra and Geometry* 7 (2023) 194–235. doi:10.1137/22M1469067

[18] S. Galán, W. F. Feehery, P. I. Barton, Parametric sensitivity functions for hybrid discrete/continuous systems, *Applied Numerical Mathematics* 31 (1999) 17–47. doi:10.1016/S0168-9274(98)00125-1

[19] M. D. Hoffman, A. Gelman, The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo, *Journal of Machine Learning Research* 15 (2014) 1593–1623.

[20] M. Betancourt, M. Girolami, Hamiltonian Monte Carlo for hierarchical models, in: *Current Trends in Bayesian Methodology with Applications*, CRC Press, 2015, pp. 79–101.

[21] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, P.-C. Bürkner, Rank-normalization, folding, and localization: an improved $\widehat{R}$ for assessing convergence of MCMC, *Bayesian Analysis* 16 (2021) 667–718.

[22] S. Talts, M. Betancourt, D. Simpson, A. Vehtari, A. Gelman, Validating Bayesian inference algorithms with simulation-based calibration, arXiv:1804.06788, 2018.

[23] P. Kidger, On Neural Differential Equations, PhD thesis, University of Oxford, 2021. arXiv:2202.02435

[24] D. Phan, N. Pradhan, M. Jankowiak, Composable effects for flexible and accelerated probabilistic programming in NumPyro, arXiv:1912.11554, 2019.

[25] Y. Ge, C. Liu, Fermentation characteristics and post-acidification of yogurt by *Streptococcus thermophilus* CICC 6038 and *Lactobacillus delbrueckii* subsp. *bulgaricus* CICC 6047 at optimal inoculum ratio, Mendeley Data, version 2, 2023. doi:10.17632/w59pnm88xx.2

[26] P. Zanatta, A. Basso, A new approach to the characterization of *Streptococcus salivarius* subsp. *thermophilus* based on acidification rates, *Le Lait* 72 (1992) 285–295. doi:10.1051/lait:1992321

[27] C. Béal, G. Corrieu, Influence of pH, temperature, and inoculum composition on mixed cultures of *Streptococcus thermophilus* 404 and *Lactobacillus bulgaricus* 398, *Biotechnology and Bioengineering* 38 (1991) 90–98.