The Cleaning Equation

A calculable model for industrial parts cleaning:
from Sinner's circle to a determining equation

Tomasz Benert

Foundation paper, Edition 1.3, September 2026

Developed in collaboration with the AI Claude (Anthropic)

Summary. Since the 1950s, Sinner's circle has described cleaning through four mutually exchangeable factors (temperature, time, chemistry, mechanics), but it does not say at what exchange rate. This paper supplies that exchange rate. The factors do not add up, they reinforce one another; we therefore calculate multiplicatively. Soil removal is approximated as first-order kinetics, temperature follows Arrhenius, chemistry and mechanics follow power laws with a base effect that prevents a missing factor from bringing the process to a numerical standstill. The product gives the cleaning rate k; the task is met when $k \cdot t \geq A$, where $A = \ln(m_0/m_{\text{perm}})$ combines soil level and cleanliness requirement in a single number. Taking logarithms turns the product into a sum: Sinner's circle as an equation, exchange rates included. The material side enters at three separate points: the soil level sets A, the type of soil determines the exponents, material and section thickness bound the permitted window. The calculation is not made against an absolute value but against a known working point, in practice the chemical supplier's specification. The unknown base rate k_0 thereby cancels out, and the calculation needs no test series of one's own.

This equation is an approximation with a limited window of validity, not a law. The paper states throughout where it breaks down, gives a test plan of 29 samples that checks the parameters *and* the soundness of the assumptions, and prescribes a safety factor. Without one, roughly half the parts miss the target at realistic scatter.

Contents

1 Purpose, scope and honest limits **3**

1.1 What kind of cleanliness is meant 3

1.2 What the paper delivers and what it does not 4

2 Starting point: Sinner's circle **4**

2.1 What it says 4

2.2 What it lacks 4

2.3 What the literature supports and what it disputes 5

3 Building block 1: the decay curve **6**

3.1 Testability with caution 7

4 Building block 2: the cleaning task A **7**

5 Building block 3: the four factors as laws **8**
5.1 Temperature: Arrhenius . . . 8
5.2 Chemistry: diminishing returns . . . 9
5.3 Mechanics: wall shear stress . . . 9
5.4 Mechanics in production: the ratio of two process types . . . 10
5.5 Time: the only linear factor . . . 11

6 The cleaning equation **12**
6.1 Base effects: why no factor goes to zero . . . 12
6.2 The complete equation . . . 12
6.3 The logarithmic form: Sinner as a sum . . . 12
6.4 What the product form presupposes and how far it holds . . . 13
6.5 Exchange rule, sensitivity and headroom . . . 13

7 The reference point: calculating without a test series of your own **14**
7.1 The device: not calculating against zero, but against a working point . . . 14
7.2 From the degree of fulfilment to the residual soil . . . 15
7.3 What the device costs: three honest restrictions . . . 15

8 The material and soil side **16**
8.1 Three separate roles . . . 16
8.2 Type of soil: the parameter set . . . 17
8.3 Material: the permitted window . . . 18
8.4 Section thickness: heat-up lag and bath cooling . . . 18

9 Three extensions, without which the model fails in practice **19**
9.1 Redeposition: the bath load . . . 19
9.2 Mixed soil: the kink . . . 22
9.3 Thresholds: where the model jumps . . . 22

10 The special case of aluminium: etching as a second condition **22**
10.1 Why etching is not a fifth factor . . . 22
10.2 The etching equation . . . 23
10.3 The three numbers and how certain they are . . . 24
10.4 The reversal of the levers . . . 24
10.5 Etching improves the degreasing . . . 25
10.6 The common range and when there is none . . . 25

11 The design problem **27**

11.1 Scatter and the safety factor 27
11.2 The design rule 28

12 A worked example **29**

13 Test plan: determining the parameters in your own plant **30**

14 What the model cannot do **32**
14.1 Where it breaks inside its own window 32

15 Next steps **33**

A List of symbols **34**

B Derivation of the heat-up loss **34**

Sources and further reading **35**

1 Purpose, scope and honest limits

This paper answers a question asked daily in every plant and almost always settled by feel: *How do I set temperature, time, concentration and mechanics so that a given part carrying a given soil comes out reliably clean, and as cheaply as possible?*

The **scope** is aqueous industrial parts cleaning ahead of a surface treatment: degreasing, inter-stage cleaning, pre-treatment before coating, electroplating or bonding. Immersion and spray processes are covered; for ultrasound the framework applies but the mechanics term does not. Cavitation is not a wall shear stress (Section 5.4). Solvent cleaning follows different laws and is not treated here.

1.1 What kind of cleanliness is meant

The model works with **film-type residual soil as a mass per unit area in** mg/m^2: oil, grease, emulsion residues, water-soluble salts.

This paper does not apply to particulate technical cleanliness to VDA 19.1 / ISO 16232. There the target quantity is not a mass but the number of particles per size class, as a rule together with a largest-particle limit. The two quantities cannot be converted into one another: a component can meet every gravimetric limit and still fail on a single oversized chip. The reason is a different removal mechanism: the adhesive force of a particle grows roughly with its diameter, the flow force acting on it with the square of the diameter, which is why large particles leave first. Residual mass therefore falls faster than residual count, and the largest remaining particle follows a curve of its own again. For particulate requirements only the *directions* carry over from this paper: mechanics and shadowing govern the result, chemistry and temperature are almost without effect.

1.2 What the paper delivers and what it does not

It delivers: a closed equation that weights the factors correctly, states the exchange rates and builds the material limits in as constraints.

It does not deliver: dependable absolute minutes without measurements of your own. The equation contains six parameters (k_0, a, b, E_A, κ_0, μ_0) that are *not constants of nature*. They depend on the combination of soil, detergent and surface and have to be measured. The figures in Section 8.2 are *estimated starting values for planning trials*, not design data.

How this paper came about. It was developed in collaboration with the AI Claude (Anthropic): the derivation, the literature research and the wording were worked out in dialogue, while the technical choices and the responsibility for the result rest with the author. Every figure has been recomputed: the worked examples against the source code of the accompanying calculator, the material and standards data against the works listed in the sources. Where a plausible-sounding statement could not be substantiated, it is no longer in this paper.

New in Edition 1.3. The calculation for shadowed spots (blind hole, $\tau = 0$) has been removed; it was the model's only extrapolation and is no longer in the calculator (Section 6.1). The residual ratio has been extended by the bath load (Section 9.1). And the supplier margin the calculator assumes for its verdict is now stated (Section 7.3).

In one sentence: the model reliably says *which lever gives how much and where headroom is left*. What it delivers in absolute terms is only as good as the test series it was calibrated against, and holds only within a limited window around that calibration point.

That calibration point, however, need not be measured by you. Take the chemical supplier's specification as the reference and k_0 cancels out, which makes the calculation usable without a laboratory, with a limitation that Section 7 states openly.

2 Starting point: Sinner's circle

2.1 What it says

Herbert Sinner (1900–1988), surfactant chemist and head of detergent applications at Henkel, condensed cleaning into four influencing quantities: **chemistry** (detergent and concentration), **mechanics** (flow, spray pressure, ultrasound), **temperature** and **time**. This is shown as a pie chart with four sectors. The core statement: the factors are interdependent but can be varied against one another in magnitude. Reduce one and another must be raised. That is correct, and experience confirms it daily.

2.2 What it lacks

There is no equation. The circle says *that* you can trade, but not *at what rate*. Whether 10 °C more replaces five minutes or fifteen, whether doubling the concentration saves half the time or only a fifth of it, the picture gives no answer. That answer is exactly what is needed when designing a plant.

And the usual additive reading is misleading. A pie chart of constant area suggests that the four shares always add up to the same total. As an image of exchangeability that is serviceable; as a calculation rule it is wrong. Doubling the concentration does not double the effect, and the effects do not add, they **reinforce one another**. Raise temperature and mechanics together

and you gain the product of the individual factors, not their sum. That is why we calculate multiplicatively.

What expressly does *not* follow from this: that a missing factor brings the process to a halt. Hot water without detergent cleans: water-soluble salts completely, thin oil films in part. A still bath without circulation also cleans, only slowly. A pure product of power laws would predict a cleaning rate of zero at zero concentration or zero flow, that is, *never clean.* This is wrong, and Section 6.1 corrects it.

2.3 What the literature supports and what it disputes

The building blocks have been studied for decades, but the evidence is less uniform than the model would like. That belongs at the front, not in the small print.

In favour of exponential decay is Kissa (1978): for particulate soil on textiles he measures a reaction order of 1.1, so close to first order. The JAOCS series on detergency kinetics likewise finds formal first order for tristearin.

Against simple exponential decay stand several findings that have to be known:

- Kissa himself records that the rate coefficient *falls* as washing proceeds, so decay departs from first order. Schott (1976) gives the reason: binding strengths are distributed, the loosely bound goes first, the remainder becomes ever harder to remove. A single k is therefore always a mean over the interval considered.
- For closed, soluble deposits, more recent modelling assumes **zero order**: the layer thickness falls linearly because the attackable interface stays constant. Diffusion-driven rinsing follows a $\sqrt{t}$ law.
- For aged oil soil on hard surfaces a **two-rate model** is needed, and the overall rate depends very strongly on **pH**.
- The JAOCS work finds that around **90 % of the removal** does not occur through the flow but through the air–liquid interface travelling across the surface as the part enters and leaves the bath (Dupré mechanism), a contribution *independent* of dwell time and not contained in this model. *On the scope of this finding see below;* it applies to liquid oil films and does not carry over to machining greases.
- In CIP practice a **concentration optimum** is documented, above which cleaning gets worse again, as is a **critical wall shear stress** below which adhering deposits are practically not removed at all.

The Dupré finding, were it general, would be an attack on the design equation itself. A removal that happens as the part passes through the bath surface and does not depend on dwell time could not be described by $k \cdot t \geq A$. One would then be designing the wrong quantity. Its scope therefore needs stating:

It applies to liquid oil films. A travelling meniscus can roll up a thin, mobile oil film and carry it away; that is the interfacial mechanics the work describes. A grease that is solid at room temperature (machining, drawing or preservation grease) it cannot strip off, because there is nothing there that flows and could be carried along. Plant experience in metalworking is unambiguous: greases cannot be removed by dipping in and out alone.

And when spraying, the mechanism does not exist at all. There the part never passes through a bath surface. A finding from an immersion trial does not carry over to spray degreasing.

For the scope of this paper (degreasing, predominantly oils and greases, immersion as well as spraying) $k \cdot t \geq A$ therefore holds up. Where thin, liquid oil films are cleaned in an immersion bath, the calculation is conservative instead: part of the work happens on withdrawal, which the model attributes to dwell time. That is the direction in which a model error is forgivable.

An honest appraisal. We are not inventing physics, but neither are we simply assembling settled results. We choose a *simple form that is serviceable over a limited window* and state where it breaks. Integrated cleaning models of this kind already exist in dishwasher and CIP engineering; what is new here is the adaptation to industrial parts cleaning with the material limits as constraints.

3 Building block 1: the decay curve

We assume that the removal rate is proportional to the amount of soil still present:

$$\frac{\mathrm{d}m}{\mathrm{d}t} = -\,k \cdot m \qquad \Longrightarrow \qquad \boxed{m(t) = m_0 \cdot e^{-k\,t}} \tag{1}$$

m_0 is the soil at the start (in mg/m^2), $m(t)$ the remainder after time t, and k the **cleaning rate** in min^{-1}. Everything that temperature, chemistry and mechanics do sits inside that single k. That is the trick that keeps the matter manageable.

Range of validity of this assumption. It holds as long as the soil is present as mutually independent individual amounts: broken-up oil islands, single particles, adsorbed residual layers. It does *not* hold while a closed film covers the area completely. Then the attackable interface is constant and so is the rate, and decay is initially **linear** (zero order), not exponential. Where removal is diffusion-controlled, a $\sqrt{t}$ curve results.

A curve spanning a whole decade of soil removal therefore usually passes through several regimes: 500 mg/m^2 of drawing oil corresponds to a film about 0.5 μm thick, 5 mg/m^2 to only a few molecular layers. **The parameter set holds per regime, not across the whole curve.** When in doubt, calibrate in the last decade before the target value. That is where the cycle time is decided.

3.1 Testability with caution

Plot $\ln m$ against t and a straight line must result. If it does not, the base model is provisionally **refuted** for that case. The extensions in Section 9 may rescue it, but only on one condition:

An additional parameter must be measured independently or be independently plausible, for instance the soil carried by the bath for the bath load m_∞, or a known second soil component for mixed soil. If a parameter is introduced only because the curve would not fit otherwise, the model is not confirmed but merely adjusted. **Five measured points support at most two free parameters.** The family $m(t) = m_\infty + \sum_i A_i e^{-k_i t}$ has five of them with two terms, and will fit any monotonically falling curve. That would not be a test but interpolation.

And the shape of a curve alone permits *no* unambiguous diagnosis: a plateau may be redeposition, an irreducible residual layer, or the detection limit of the measurement. A kink may be mixed soil, a change of removal regime, or a distribution of binding strengths. The diagnosis always needs a second, independent argument.

4 Building block 2: the cleaning task A

Rearranging equation (1) for the condition that at most the permitted residual soil m_{perm} may remain at the end:

$$m_0\, e^{-kt} \le m_{\text{perm}} \qquad \Longrightarrow \qquad \boxed{k \cdot t \;\ge\; \underbrace{\ln\!\left(\frac{m_0}{m_{\text{perm}}}\right)}_{A}} \tag{2}$$

We call A the **cleaning task**. It is dimensionless and combines in one number two things usually discussed separately: how dirty the part arrives and how clean it has to leave. The product $k \cdot t$ is the **cleaning work delivered**.

Table 1 The cleaning task A for typical cases

Case	m_0 [mg/m²]	m_{perm} [mg/m²]	A
Light oil film, simple requirement	200	20	2.30
Ordinary interstage cleaning	500	50	2.30
Pre-treatment before coating	500	5	4.61
Heavily oiled, high requirement	1000	5	5.30

From this follows immediately a statement that contradicts gut feeling, *if* first-order kinetics apply: what counts is not the absolute amount of soil but the ratio of initial to final state. Twice the soil raises the task by only $\ln 2 = 0.69$, at $A = 4.6$ therefore by some 15 %, not by 100 %. The *requirement*, conversely, weighs heavily: lowering the limit from 50 to 5 mg/m² doubles the task from 2.30 to 4.61.

Careful: this is a consequence of the assumption, not a measured finding. For thick, closed films, where the rate is governed by the area and not by the amount, $t \propto (m_0 - m_{\text{perm}})$ applies instead: twice the oil, twice the time. The widespread shop-floor experience that heavily oiled parts take noticeably longer than the logarithmic calculation says is therefore to be taken seriously and not dismissed as an error. Check in series 0 which case applies to you: if m against t gives a straight line on a *linear* plot rather than $\ln m$ against t, work with the difference, not with the ratio.

5 Building block 3: the four factors as laws

5.1 Temperature: Arrhenius

$$k \;\propto\; \exp\!\left(-\frac{E_A}{R\,T}\right) \quad \text{or, referred to } T_0: \quad \frac{k}{k_0} = \exp\!\left[-\frac{E_A}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right] \tag{3}$$

$R = 8.314\,\text{J/(mol K)}$ is the gas constant, T the *absolute* temperature in kelvin.

On notation: absolute temperatures and the gas constant are given in kelvin, because the equation of Arrhenius requires it. **Temperature *differences*, by contrast, are given throughout in °C.** A difference of 10 °C and one of 10 K are the same number, and °C is the unit read off at the bath.

E_A is an *apparent* activation energy, that is, a fitted number summarising the combined temperature dependence of soil viscosity, diffusion, surfactant action and interfacial tension. It is *not* a molecular energy barrier and must not be read as one. In particular it may itself be temperature dependent. The Arrhenius plot is then curved, and the measured slope depends on the range in which you measure. Where *transport away* becomes rate determining, E_A falls to the value for the viscosity of water, that is, around 15 to 20 kJ/mol.

In practice this is known as the **rule of thumb** that the rate roughly doubles per 10 °C. That is the same statement, only approximated, with $Q_{10} = \exp[E_A \cdot 10/(R\,T\,(T+10))]$:

Table 2 Converting between activation energy and the familiar doubling rule (at 50 °C)

E_A [kJ/mol]	Q_{10} per +10 K	Rating	typical of
30	1.40	weak	thin oils, transport controlled
45	1.65	medium	emulsions, cutting fluids
60	1.95	**the classic doubling rule**	greases, general degreasing
80	2.44	strong	viscous, polymerised residues

Worth remembering: $E_A \approx 62\,\text{kJ/mol}$ **corresponds at** 50 °C **exactly to the rule of thumb "twice the rate per** 10 °C**".** With no measured activation energy to hand, use that. **But:** Q_{10} is itself temperature dependent. At a fixed $E_A = 62\,\text{kJ/mol}$ it is 2.09 at 40 °C and only 1.76 at 85 °C. Over that span of 45 °C, using a fixed Q_{10} taken from the bottom of the range is therefore 19 % too optimistic at the top end, and considerably more than that once compounded over the whole stretch. This paper therefore uses Arrhenius throughout, not a constant Q_{10}.

Temperature is the factor with the *greatest practical leverage*, not because its functional form is special, but because raising it by 10 °C is usually cheap in this application and usually still has

headroom. Across the usable window from 40 to 90 °C the Arrhenius term acts like a very steep power law with an exponent of the order $E_A/(RT)$, that is, around 20.

5.2 Chemistry: diminishing returns

$$\frac{k}{k_0} = \left(\frac{c}{c_0}\right)^a, \qquad a \approx 0.2 \ldots 0.6 \tag{4}$$

The exponent a lies well below one: diminishing returns. At $a = 0.5$, doubling the concentration gives only a factor of $\sqrt{2} = 1.41$, that is, 41 % more speed, not 100 %.

This is a fit, not a derivation, and the reason is *not* the CMC. The critical micelle concentration of technical surfactants lies between 0.01 and 1 g/l. A bath at 2 % detergent contains around 6 g/l of surfactant and is therefore *one to nearly three orders of magnitude above* the CMC. If the CMC were the cause of the diminishing returns, a would have to be zero in every real bath, which the recommended values would immediately refute.

The actual causes are: limited **emulsifying and holding capacity** of the bath, **alkalinity consumption** and, with strong chemistry, the transition to **transport-controlled** removal, where more detergent no longer speeds anything up. A power law represents this serviceably over roughly a doubling to a quadrupling of the dosage, and not beyond.

There is an optimum, not merely a saturation. CIP practice documents well that the cleaning rate *falls again* above a certain caustic concentration. A power law with $a \geq 0$ fundamentally cannot represent that. Anyone going beyond the customary dosage must measure rather than calculate. Overdosing additionally increases drag-out, rinse water demand, foam and attack on the material: costs the equation does not contain.

5.3 Mechanics: wall shear stress

$$\frac{k}{k_0} = \left(\frac{\tau}{\tau_0}\right)^b, \qquad b \approx 0.3 \ldots 1.0 \tag{5}$$

τ is the wall shear stress, that is, the force with which the flow drags at the surface (in $\mathrm{N/m^2}$). It can rarely be measured directly in production; the substitute quantities are: for immersion with circulation the flow velocity, with $\tau \propto v^{1.8}$; for spraying, approximately the nozzle pressure; for ultrasound, the power density in W/l above the cavitation threshold.

Two restrictions and one warning about the substitute quantities.

First, firmly adhering deposits have a **critical shear stress** τ_{crit}: below it practically nothing is removed, above it removal sets in. In hygienic and CIP engineering this is an established test criterion; the power law does not represent it. *If your test series yields a b well above 1, that is as a rule not a genuine exponent but a sign that you are measuring near* τ_{crit}. Then move two steps higher.

Second, for soluble soils the removal is often transport controlled; the rate then depends on the concentration profile in the boundary layer and not in any simple way on the wall shear stress. A b below 0.3 is an indication of this.

The substitute quantities are assumptions, not conversions. Anyone determining b from a series over flow velocity is measuring the slope against $\ln v$ and has to assume $\tau \propto v^{1.8}$. Every error in that goes straight into b. Always state, therefore, which quantity was measured. Ultrasound is a special case: cavitation is not a wall shear stress but a threshold- and frequency-dependent process, distributed very unevenly in space within the tank by the standing wave. A b obtained from ultrasonic trials cannot be transferred to any other plant.

5.4 Mechanics in production: the ratio of two process types

The equation contains τ/τ_0, a ratio. That is why mechanics remains usable despite the restrictions just listed: τ **is never needed as a number, only ever as the factor between two plant states.** A plant that measures neither wall shear stress nor flow velocity can still say whether it immerses or sprays, and that is exactly enough.

In practice this means that a ladder of process types takes the place of a measured quantity, its rungs given as multiples of a chosen reference point. The reference point of the ladder is ordinary spraying.

Table 3 Mechanics steps as multiples of τ_0 (reference: ordinary spraying)

Process type	τ/τ_0	typical
Immersion, still bath	0.05	no circulation, convection only
Immersion with circulation	0.20	pump, nozzle manifold in the tank
Immersion with part movement	0.41	raising and lowering, rotating, shaking
Spraying	**1.00**	**reference point**
Spraying, high pressure	2.00	more pressure or more nozzles
High pressure, directed	4.00	lance, point jet

Where the three immersion rungs come from. They are not estimates but were back-calculated from plant experience. For grease and drawing compound (the soil type with $b = 0.60$ and $\mu_0 = 0.20$) the following holds relative to spraying: immersion with part movement costs 1.5 times the time, with circulation 2.0 times, in a still bath 3.0 times. Solved for τ this gives

$$\frac{\tau}{\tau_0} = \left(\frac{1/\zeta - \mu_0}{1 - \mu_0}\right)^{1/b}, \qquad \zeta = \text{time factor},$$

and hence exactly the three values in the table. The three spray rungs remain estimates; what is dependable there is the ranking, not the absolute value.

Those time factors apply to grease, not to everything. Since b and μ_0 depend on the soil type, the same step works out differently for other soils. In a still bath relative to spraying: light oil 2.3 times, grease and wax 3.0 times, dried-on emulsion 3.9 times, polishing paste 5.0 times, baked-on polymer 6.2 times. This is intended: a tenacious polymer suffers more from the absence of movement than a thin oil film does.

The reference point of the ladder is not the reference point of the design. The ladder is tied to spraying because otherwise it would have no fixed zero. The *chemical supplier's specification*, by contrast, is as a rule written for the still immersion bath: if the data sheet says „ten minutes at 60 °C“, immersion without part movement is what is meant. For the degree of fulfilment the correct reference is therefore $\tau_{\text{spec}} = 0.05$ and not 1.00. Anyone who puts spraying here calculates roughly a third too favourably.

Two steps, one ratio. And the order must not matter. Because μ_0 is an *absolute* floor (what still comes off without any flow at all), it belongs at $\tau = 0$ and not at whichever reference point happens to be chosen. The mechanics term is therefore formed as the quotient of two evaluations of the same curve:

$$f_M \;=\; \frac{g(\tau_{\text{actual}})}{g(\tau_{\text{spec}})}, \qquad g(\tau) \;=\; \mu_0 + (1-\mu_0)\left(\frac{\tau}{\tau_0}\right)^{b}. \tag{6}$$

Writing instead $f_M = \mu_0 + (1-\mu_0)\,(\tau_{\text{actual}}/\tau_{\text{spec}})^b$ (that is, placing the floor at the reference point rather than at zero) makes the same physical step differ depending on which way it is taken: from spraying into the still bath the factor would be 2.49, but back again 0.30 instead of the reciprocal 0.40. With equation (6), $f_M(\text{there}) \cdot f_M(\text{back}) = 1$ holds for every pair of steps, and $f_M = 1$ at the reference point is preserved.

The figures on the ladder are not the effect. The rungs are values for τ/τ_0; how much of that reaches the result is decided by the mechanics term with b and μ_0. The difference is large and easily missed:

Step	rung (τ/τ_0)	factor on k
still bath → circulation	×3.9	1.4 … 1.5
circulation → part movement	×2.1	1.2 … 1.4
part movement → spraying	×2.5	1.3 … 1.7
spraying → high pressure	×4.0	1.6 … 2.6

(Range for $b = 0.4 \ldots 0.8$ at $\mu_0 = 0.2$.) Going from a still bath to spraying triples the cleaning rate. It does not multiply it sixteenfold, even though the rungs are a factor of 20 apart.

And they are not real wall shear stresses. Between a still bath and a spray jet there are in fact several decades in τ. The ladder is compressed so that the *effects* match experience; it is therefore useless for estimating shear stresses. Anyone with figures of their own should use them. And anyone comparing two plants of the *same* process type bypasses the ladder entirely: there, the ratio of the flow velocities raised to the power 1.8 is the better number. **The ladder is at its weakest across process boundaries.** Ultrasound deliberately does not appear in it (see the warning above), and the jump from immersion to spraying changes not only τ but also the entry-and-exit event, which the model does not know at all.

5.5 Time: the only linear factor

Time sits outside k in equation (2) and therefore acts linearly: twice the time, twice the work. **Measured by relative adjustment** (doubling against doubling) it is thus more effective than chemistry (exponent a) and mechanics (exponent b).

With temperature it is *not* comparable in that way: “twice the temperature” is not a meaningful quantity, and how temperature and time stand against each other depends solely on how many degrees you are allowed to move. That is precisely why the right question is not “which factor is stronger” but “which factor still has headroom within the permitted window” (Section 6.5).

6 The cleaning equation

6.1 Base effects: why no factor goes to zero

A pure power law predicts a cleaning rate of zero for $c \to 0$ or $\tau \to 0$, that is, *never clean.* This is wrong and moreover contradicts our own soil table, in which $a \approx 0$ is given for water-soluble salts. We therefore give each of the two factors a **base effect**: the share of the effect that remains even without that factor.

$$f_C(c) = \kappa_0 + (1-\kappa_0)\left(\frac{c}{c_0}\right)^a, \qquad f_M(\tau) = \mu_0 + (1-\mu_0)\left(\frac{\tau}{\tau_0}\right)^b \tag{7}$$

κ_0 is the effect of plain hot water (typically $0.05 \ldots 0.3$, close to 1 for water-soluble soil), μ_0 the effect in a still bath (typically $0.1 \ldots 0.3$). At the reference point $c = c_0$, $\tau = \tau_0$ both are exactly 1, so the calibration is unchanged.

Shadowed spots lie outside the model. In blind holes, under overlaps and at contact points in the basket the flow is practically zero. Arithmetically the effect curve would fall back to μ_0 there; but that figure cannot be measured, because a genuinely stationary volume does not occur in a running plant. It would be the only extrapolation in the whole model: everywhere else the comparison is against the supplier's specification, and everything unknown cancels out in the process (Section 7). At the dead end nothing cancels. This paper therefore does not calculate such spots, and neither does the calculator. They are solved through jigging and racking, part orientation, rotation or targeted nozzles, not through time.

6.2 The complete equation

$$\boxed{k \;=\; k_0 \cdot \underbrace{\left[\kappa_0 + (1-\kappa_0)\left(\tfrac{c}{c_0}\right)^a\right]}_{\text{chemistry}} \cdot \underbrace{\exp\left[-\frac{E_A}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right]}_{\text{temperature}} \cdot \underbrace{\left[\mu_0 + (1-\mu_0)\left(\tfrac{\tau}{\tau_0}\right)^b\right]}_{\text{mechanics}}} \tag{8}$$

and the condition for a successful cleaning step reads

$$\boxed{k(T,c,\tau)\cdot t \;\geq\; A \;=\; \ln\left(\frac{m_0}{m_{\text{perm}}}\right)} \tag{9}$$

That is **the cleaning equation**. Everything that follows is either its evaluation or a correction to it.

6.3 The logarithmic form: Sinner as a sum

In the range where the base effects are small compared with the power terms (that is, at and above the reference point), taking logarithms turns the product into a sum:

$$\ln k_0 + \underbrace{a \ln \tfrac{c}{c_0}}_{\text{chemistry}} - \underbrace{\tfrac{E_A}{R}\left(\tfrac{1}{T} - \tfrac{1}{T_0}\right)}_{\text{temperature}} + \underbrace{b \ln \tfrac{\tau}{\tau_0}}_{\text{mechanics}} + \underbrace{\ln t}_{\text{time}} \;\geq\; \ln\left(\frac{m_0}{m_{\text{perm}}}\right) \tag{10}$$

This is Sinner's circle as an equation. Every factor is a summand, each contributes with its own weight, and the sum has to reach a minimum value. So the circle picture was right after all, but only in the logarithm, and only with the weights a, b and E_A/R that it never named itself.

6.4 What the product form presupposes and how far it holds

The product form claims that the effect of one factor does not depend on where the others stand. In general that is **false**. Equation (10) is the first-order expansion of $\ln k$ about the reference point; a, b and E_A are the three slopes at that point. Close to the reference point this is a good and very serviceable approximation. Over large adjustments it breaks down.

The most important known couplings:

- a **depends on** T**.** Non-ionic surfactants have their optimum close to the cloud point and lose effect above it. The temperature effect is *not monotonic* there, and Arrhenius fundamentally cannot represent a maximum.
- b **depends on** T**.** Above the softening or melting point of the soil the removal mechanism changes from cohesive failure to viscous shearing: b jumps.
- τ **depends on** T**.** At a fixed pump operating point, bath viscosity changes the wall shear stress along with it. *Change T and you change τ unintentionally too.* A test plan that holds "everything else fixed" does not hold τ fixed.
- E_A **depends on** τ**.** Where transport away becomes rate determining, the apparent activation energy falls to 15 to 20 kJ/mol.
- c **and pH are coupled.** With alkaline detergents the pH effect sits unstated inside the exponent a.

A practical rule for the window of validity. One parameter set holds for roughly ±15 °C in temperature, a factor of 2 in concentration and a factor of 3 in mechanics around the calibration point. Outside that, recalibrate. The test plan in Section 13 checks expressly whether the couplings are negligible within the chosen window.

6.5 Exchange rule, sensitivity and headroom

Table 4 answers the question from Section 1 in numbers: what does it cost to halve the cleaning time? Table 5 sets alongside it what a doubling of each adjustment buys.

Table 4 What *halving the cleaning time* costs ($a = 0.5$, $b = 0.8$, $\kappa_0 = 0.15$, $\mu_0 = 0.2$, $E_A = 62$ kJ/mol)

Adjustment	change required	assessment
Temperature	+10.0 °C from 50 °C	usually the cheapest, often barred by the material
	+11.3 °C from 70 °C	*the exchange rate gets worse the higher you go*
Mechanics	shear stress × 2.8	rebuilding the plant, but permanent
Chemistry	concentration × 4.7	usually impossible or uneconomic

Table 5 Effect of a doubling on the cleaning work delivered, $k \cdot t$

Adjustment	change	factor on $k \cdot t$
Temperature	+10 °C from 50 °C	2.00
Time	× 2	2.00
Mechanics	× 2	1.59
Chemistry	× 2	1.35

The temperature row is not directly comparable with the others: it shows an absolute change of 10 °C*, the others a doubling. At* +20 °C *the entry would be* 3.84*, at* +5 °C 1.42*.*

The mechanics row applies to a plant whose reference point is spraying ($\tau_0 = 1.00$). It is the only row that depends on the starting point: out of a still immersion bath the same doubling gives only 1.20 *instead of* 1.59*, and halving the time would require* × 7.0 *there instead of* × 2.8*. The closer a plant runs to standstill, the more of the cleaning comes from chemistry and heat anyway, and the less more flow achieves.*

Sensitivity is not headroom. A large exponent says that k changes strongly *if* you change the factor. It does not say that you can change it. What a factor can actually contribute is

$$\text{headroom} \;=\; \ln \frac{f(\text{highest value in the permitted window})}{f(\text{present value})}$$

that is, sensitivity times the adjustment available.

Example. A plant may go from 60 to 75 °C, may dose from 2 to 8 %, and the pump still has 20 % in reserve:

temperature 0.97 chemistry 0.62 mechanics 0.12

Among the three factors inside k, mechanics has the second-highest exponent and by far the smallest headroom. Work out the headroom for all four factors before deciding. The column "most sensitive factor" in Table 9 is only the starting point of that calculation.

7 The reference point: calculating without a test series of your own

Section 1 conceded what the paper cannot deliver: dependable absolute minutes without measurements of your own. The reason is a single number: the base rate k_0. It depends on the combination of soil, detergent and surface, it is not a constant of nature, and no plant knows it.

This section shows the way out. It is not a new piece of physics but a device of engineering practice, and it is the reason the equation could become a usable tool at all.

7.1 The device: not calculating against zero, but against a working point

Every plant already has a point it knows to be sound: the **chemical supplier's specification**. It is in the technical data sheet, it names temperature, concentration and time, and it is a commitment: run it this way and you get clean parts.

Take that point as the reference (T_0, c_0, τ_0, t_0) and the commitment reads, in the language of this paper, simply

$$k_0 \cdot t_0 \;=\; A \tag{11}$$

Substitute this into equation (8) and divide the actual condition by the specified condition, and k_0 disappears completely:

$$E = \frac{k \cdot t}{k_0 \cdot t_0} = \underbrace{\left[\kappa_0 + (1-\kappa_0)\left(\tfrac{c}{c_0}\right)^a\right]}_{\text{chemistry}} \cdot \underbrace{\exp\left[-\frac{E_A}{R}\left(\frac{1}{T}-\frac{1}{T_0}\right)\right]}_{\text{temperature}} \cdot \underbrace{\left[\mu_0 + (1-\mu_0)\left(\tfrac{\tau}{\tau_0}\right)^b\right]}_{\text{mechanics}} \cdot \underbrace{\frac{t}{t_0}}_{\text{time}} \tag{12}$$

We call E the **degree of fulfilment**. It states what share of the promised cleaning work the plant actually delivers today. $E = 1$ means: exactly on specification. $E = 0.41$ means: 41 % of the work required. $E = 1.8$ means: 80 % headroom.

Equation (12) **contains not a single quantity that would have to be measured.** Six numbers from the plant (three from the data sheet, three from the line) plus the process type are enough. That is the whole trick, and it is the same one sterilisation engineering has used since the 1920s with the F_0 value: do not calculate the absolute kill, but set the delivered against the required.

7.2 From the degree of fulfilment to the residual soil

E is a ratio. It can be translated into the quantity the shop floor actually asks about (*by what factor is the residual soil above the target?*), and exactly so, without approximation. With $m(t) = m_0 e^{-kt}$, $k\,t = E \cdot A$ and $m_{\text{perm}} = m_0 e^{-A}$ it follows that

$$\frac{m(t)}{m_{\text{perm}}} = \frac{e^{-EA}}{e^{-A}} = e^{(1-E)\,A} \tag{13}$$

The formula shows at once why small deviations bite so hard: the degree of fulfilment enters linearly, but in the *exponent*. For a demanding task ($m_0 = 500$, $m_{\text{perm}} = 1\,\text{mg/m}^2$, hence $A = 6.21$) a degree of fulfilment of 0.414 already produces the factor $e^{0.586 \cdot 6.21} = 38.2$. *To check this:* the 0.414 belongs to thin oil ($E_A = 38\,\text{kJ/mol}$, $a = 0.5$, $\kappa_0 = 0.15$) at the deviation named above. With the reference value $E_A = 62\,\text{kJ/mol}$ from Section 5.1 it would read 0.317 and a factor of 70. So the choice of soil type carries a great deal of weight.

7.3 What the device costs: three honest restrictions

First: the specification does not sit on the limit. Equation (12) assumes the supplier's specification is only just sufficient. That is certainly false: every supplier builds in a margin, otherwise they would face constant complaints. The calculation is thereby *systematically pessimistic*, by exactly that undisclosed margin:

For a warning tool this is the right direction: a plant is never wrongly cleared. But it also means that the multiple is a **warning level and not a measurement**. The accompanying calculator draws a consequence from this: its figures stay raw, but its verdict (green, amber, red) assumes a margin of $R = 1.5$. Green means $E \geq 1$, amber means $E \cdot R \geq 1$, red below that. Without this assumption, three degrees too cold would already read *target missed*, and a tool that shows red in normal operation stops being looked at. R is a setting; anyone who knows their own supplier's commitment enters it.

Table 6 Effect of an undisclosed supplier margin R. Specification 60 °C / 3.0 % / 120 s, actually run at 50 °C / 2.0 % / 90 s, mechanics unchanged; thin oil with $E_A = 38\,\mathrm{kJ/mol}$, $a = 0.5$, $\kappa_0 = 0.15$, giving $A = 6.21$ and $E = 0.414$. The true value is then residual/target $= e^{(1-E\,R)\,A}$.

Supplier margin R	the calculation says	actually
none (1.0×)	38.2× above target	38.2×
1.2×	38.2× above target	22.8×
1.5×	38.2× above target	10.5×
2.0×	38.2× above target	2.9×
3.0×	38.2× above target	**below the target**

Second: A remains an estimate, but E does not depend on it. Nobody knows the initial soil m_0 precisely, and equation (13) depends on it exponentially. The degree of fulfilment itself, by contrast, contains no A at all:

Table 7 The same plant, four assumptions about the amount of oil

assumed m_0 [mg/m²]	A	residual/target	degree of fulfilment E
200 (light oil film)	5.30	22.3×	41.4 %
500 (ordinary)	6.21	38.2×	41.4 %
1000 (drawn part)	6.91	57.3×	41.4 %
2000 (preserved)	7.60	86.0×	41.4 %

The percentage is the dependable number, the multiple the understandable one. A factor of 10 in an estimated input gives a factor of 3.9 in the multiple, and does not change a single digit of the degree of fulfilment. If you want to quote one number, quote the degree of fulfilment; the multiple only says what order of magnitude you are in.

Third: only k_0 cancels, the exponents do not. a, b and E_A stand in equation (12) unchanged. They still have to be chosen (via the type of soil from Table 9), and every error there goes straight into the result. The device removes the *calibration*, not the *choice of model*.

The reference point has to match the task. Equation (12) presupposes that specification and actual condition concern *the same* cleaning task: the same part, the same soil, the same requirement. Set the specification for drawing oil against a process with a baked-on residue and you are comparing two different parameter sets and get a number without meaning. The reference point must also be *valid*: after a change of detergent, a bath change or a modification to the flow, the old specification is no longer a reference point.

8 The material and soil side

8.1 Three separate roles

The part to be cleaned is not a fifth adjustment. It enters at three structurally different points, and keeping these three roles apart is the most important idea in this section:

Property	acts on	meaning
Soil level	m_0, hence on A	Sets the size of the task: under the first-order assumption only logarithmically, so weakly.
Type of soil	k_0, a, b, E_A, κ_0, μ_0	Determines *which lever bites at all.* The whole parameter set changes with the type of soil.
Material and section thickness	constraints	Set the *limits* of the permitted window ($T_{\max}$, pH, $\tau_{\max}$) and, through thermal inertia, a loss of time. They do not affect the rate; they forbid settings.

Table 8 The three roles of the material side

Type of soil	E_A [kJ/mol]	a chemistry	b mechanics	most sensitive factor
Thin oil, cutting fluid	30–45	0.4–0.6	0.3–0.5	chemistry and temperature
Grease, wax, solid drawing compound	50–80	0.3–0.5	0.5–0.7	temperature (melting threshold!)
Polymerised, baked on	60–90	0.2–0.4	0.8–1.0	mechanics and time
Polishing paste (grease + solids)	40–60	0.3–0.5	0.7–0.9	mechanics and temperature
Chips, dust, loose particles	0–10	0.0–0.2	0.9–1.2	mechanics alone
Salts, water-soluble	10–20	≈ 0	0.3–0.5	time and volume of water
Scale, corrosion products	50–70	0.7–1.0	0.2–0.4	chemistry (acid)
Biofilm	40–60	0.4–0.7	0.6–0.9	chemistry and mechanics

Table 9 Starting values for planning trials

8.2 Type of soil: the parameter set

Where these figures come from. Please read. The values are **estimated, not measured**. To our knowledge no published collection of activation energies and power-law exponents broken down by type of soil exists for aqueous parts cleaning; the figures available come from dairy CIP and from textile washing and are not transferable. The table therefore gives the *ranking of sensitivities* as it follows plausibly from the mechanism. **Where your measurement differs from the table, your measurement wins.**

Only two rows can be made plausible independently: for water-soluble salts the activation energy follows essentially from the temperature dependence of the viscosity of water and therefore lies at 15 to 20 kJ/mol; for particles to be removed purely mechanically it is close to zero. *A measured value well below* 15 kJ/mol *is an indication that the process is governed not by detachment but by transport away.* A b above 1 indicates a critical shear stress within the measuring range, not a genuine exponent.

How to read the table: for loose soil $a \approx 0$: no detergent in the world helps against chips, only flow. For scale $b \approx 0.3$: brushing hardly helps, the acid does the work. Pull the wrong lever and you spend money without effect.

8.3 Material: the permitted window

Table 10 names the limits within which the settings may be moved at all. They are not a recommendation but a prohibited zone: exceed them and you clean faster but damage the part.

Material	pH	$T_{\max}$	mechanics	critical
Mild steel	2–14	90°C	high	flash rust after rinsing
Stainless steel	2–14	90°C	high	avoid chlorides (pitting)
Galvanised steel	8–11	65°C	medium	attack on the zinc coating above pH 11
Aluminium, non-attacking	8.5–10.5	60°C	medium	etching attack, matting
Aluminium, etch cleaning	*its own window, see Section 10*			
Copper alloys (Cu, brass)	8–11	60°C	medium	tarnishing; avoid ammonia and amines
Magnesium	9–11	50°C	low	very narrow window
Plastics, composites	6–9	50°C	low	stress cracking from surfactants
Thin sheet $< 0.5\,\mathrm{mm}$	–	–	low	deformation, tangling in the basket

Table 10 Material limits: guide values for *non-attacking* cleaning. In case of doubt the chemical supplier's approval applies, not this table. In the etch cleaning of aluminium the attack is the purpose and not the damage; there, pH 12 to 13.5 and up to 70 °C apply (alkaline), or pH 1 to 3 (acid), see Section 10.

This turns the cleaning equation into an *optimisation under constraints*.

8.4 Section thickness: heat-up lag and bath cooling

A thick part is cold when it enters and cleans more slowly in the first minutes than the bath temperature promises:

$$T(t) = T_{\mathrm{bath}} - (T_{\mathrm{bath}} - T_{\mathrm{start}})\, e^{-t/\theta}, \qquad \theta = \frac{\rho\, c_p\, s}{2\,\alpha} \tag{14}$$

with section thickness s, density ρ, specific heat capacity c_p and heat transfer coefficient α (agitated aqueous bath: $300 \ldots 1000\,\mathrm{W/(m^2K)}$). Appendix B derives the loss of time:

The heat-up lag costs between half and just under two time constants of pure cleaning time, depending on the temperature difference, at 40 °C exactly $1.61 \cdot \theta$ (Appendix B). For 10 mm steel at $\alpha = 300\,\mathrm{W/(m^2K)}$, $\theta \approx 60\,\mathrm{s}$, so the loss is about **one and a half minutes**. For 1 mm sheet it is ten seconds, so nothing.

Table 11 Time constant θ [s] for steel ($\rho c_p = 3.61\,\mathrm{MJ/(m^3K)}$)

Section thickness	1 mm	3 mm	10 mm	20 mm
$\alpha = 300\,\mathrm{W/(m^2K)}$ (still bath)	6	18	60	120
$\alpha = 1000\,\mathrm{W/(m^2K)}$ (strongly agitated)	2	5	18	36

Limit of validity: the table does not hold everywhere. The formula is the lumped-capacitance approximation. It presupposes that the component is approximately at one temperature throughout, that is, $\mathrm{Bi} = \alpha\, s/(2\lambda) \leq 0.1$ with the thermal conductivity λ of the material. Worked out:

Material (λ in W/(m K))	valid up to s at $\alpha = 300$	at $\alpha = 1000$
Aluminium (200)	133 mm	40 mm
Mild steel (50)	33 mm	10 mm
Stainless steel (15)	10 mm	3 mm
Plastics (0.2)	0.1 mm	**not at all**

For plastics the formula is useless: there the surface heats up quickly, the core stays cold for a long time, and the governing time is the conduction time $s^2/(4\,\mathfrak{a})$ with the thermal diffusivity $\mathfrak{a}$. The factor 2 in the denominator applies to parts wetted on both sides; for solid parts and parts lying flat it drops out.

The second effect: **the batch cools the bath.**

$$\Delta T_{\mathrm{bath}} = \frac{m_{\mathrm{part}}\, c_{p,\mathrm{part}}\, (T_{\mathrm{bath}} - T_{\mathrm{part}})}{m_{\mathrm{bath}}\, c_{p,\mathrm{bath}} + m_{\mathrm{part}}\, c_{p,\mathrm{part}}} \tag{15}$$

800 kg of steel at 20 °C entering a 1000 l bath at 60 °C gives $\Delta T = -3.2$ °C, and via the doubling rule around 20 % loss of rate. For 200 kg in 2000 l it is only 0.4 °C and thus 3 %, a question of the ratio, not of the absolute quantity.

The loss is temporary, if the heating makes it up. What matters is not the drop but how long it lasts. What has to be replaced is the heat the cold part takes up: $m_{\mathrm{part}}\, c_{p,\mathrm{part}}\, (T_{\mathrm{bath}} - T_{\mathrm{part}})$, not the bath drop. In the example with 800 kg of steel that is 4.1 kWh: a 40 kW heater supplies that in 6 minutes, a 10 kW heater only in 25. Heating power therefore belongs among the design quantities, even though it does not appear in the cleaning equation. With marginal heating the second batch is colder than the first and the third colder still.

9 Three extensions, without which the model fails in practice

Three observations bring the base model down in practice. All three leave a trace in the decay curve, and Figure 1 shows what those traces look like, and why they are not unambiguous.

9.1 Redeposition: the bath load

A bath gets dirty in service. Above a certain soil load, what has just been removed is deposited again. The curve then runs towards a bath load m_∞:

$$m(t) = m_\infty + (m_0 - m_\infty)\, e^{-kt} \qquad \Longrightarrow \qquad A_{\mathrm{req}} = \ln \frac{m_0 - m_\infty}{m_{\mathrm{perm}} - m_\infty} \tag{16}$$

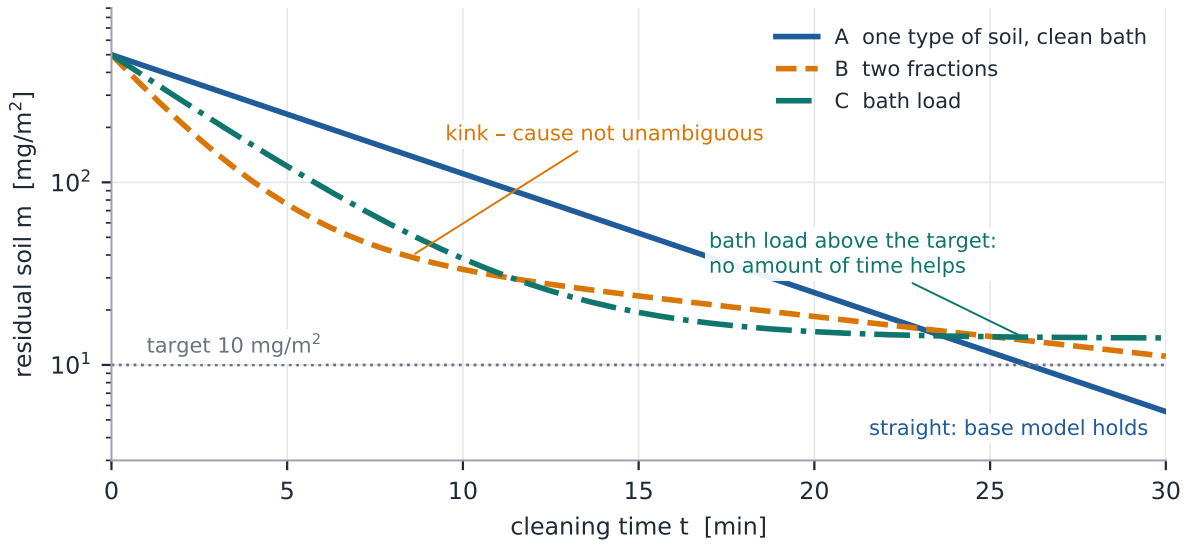

Figure 1 The semi-logarithmic plot of residual soil against time is the diagnostic tool, but not an unambiguous one: a straight line confirms the base model; a kink *may* be mixed soil or a change of regime; a plateau *may* be redeposition or the detection limit of the measurement. The diagnosis always needs a second, independent argument.

One of three cases in which more time is of no help in principle: the other two are the critical wall shear stress and an irreducible residual layer (Section 9). As m_∞ approaches the target value the required task explodes; once m_∞ reaches the target value it cannot be met at all, at any temperature, concentration and duration. Only bath maintenance helps then: oil separator, filtration, partial exchange, bath change.

Table 12 Effect of the bath load ($m_0 = 500$, $m_{\text{perm}} = 10\,\text{mg/m}^2$, $k = 0.3\,\text{min}^{-1}$)

Bath load m_∞ [mg/m²]	A_{req}	time required [min]
0 (fresh bath)	3.91	13.0
5	4.60	15.3
8	5.51	18.4
9.5	6.89	23.0
10	unattainable, maintain the bath	

m_∞ **is not a constant.** The bath load is the equilibrium between removal and redeposition and depends on the soil load of the bath, which cleaning itself increases. At high loading (much part area, little bath volume) m_∞ is higher at the end of a batch than at the start; calculating with a fixed m_∞ then underestimates the time required. m_∞ also depends on concentration and temperature, because both govern the emulsifying capacity. So it is not an independent additional parameter.

A rule of thumb instead of a formula: as long as m_{perm} stays above three times m_∞, the bath load costs about a tenth more time, annoying, but affordable. Below one and a half times it gets steep, and at parity the task is not solvable by time at all.

And the residual ratio with a floor. Equation (13) holds only for a clean bath. With

$s = m_\infty / m_{\text{perm}}$ and the extended task A_{req} from equation (16) the exact relation is

$$\frac{m(t)}{m_{\text{perm}}} \;=\; s + (1 - s)\, e^{(1-E)\, A_{\text{req}}} \tag{17}$$

The share below the floor is never removed, however long the bath runs. This cuts both ways, calculated for $m_0 = 500$, $m_{\text{perm}} = 1\,\text{mg/m}^2$, hence $A_{\text{req}} = 8.5$: at $s = 0.9$ and $E = 0.8$ the residue sits at 1.45 times the target, not at the 5.5 times the formula without the floor would give; and at $E = 1.2$ it stays at 0.92 times instead of 0.18 times, because nothing goes below 0.9. Anyone who puts the floor into A but not into the ratio is too pessimistic below the specification and too optimistic above it.

What the literature offers and what it does not

The mechanism is undisputed and consistently described in the technical literature: the more oil is held in the bath, the more surfactant is bound up with it and unavailable to the part; re-oiling on withdrawal through the bath surface comes on top. For emulsifying immersion cleaners a holding capacity of the order of 8 to 10 per cent by volume of oil is quoted, until the emulsion is saturated and the bath has to be discarded. Demulsifying, that is splitting, systems are run by design at 0.05 to 1 g/l and therefore some two orders of magnitude below that, because the oil is continuously separated off.

The figure this model needs does not exist in the literature. What would be wanted is the relation *oil load in the bath → treatment time required*. It is published nowhere. A controlled comparison between a fresh and an aged bath at *fixed* time does exist (there the yield visibly drops), but no study of how much the time would have to be extended. The plant limits for oil content that can be found scatter by more than a factor of 10 (3000 ppm to 6 %) and are unfit as a design basis; the technical literature itself states that no universally valid limit exists and that the tolerable load is to be determined per process.

The industry therefore does not steer by time at all, but by trend quantities (the ratio of total to free alkalinity, concentrate consumption per day, the acid-split test) and then changes the bath. That is precisely the gap m_∞ closes: it translates the condition of the bath into the only quantity that counts in cycle planning.

Why topping up has an economic limit

In practice a rising bath load is offset by topping up with a degreasing booster, a surfactant concentrate dosed separately from the alkalinity builder. This acts in two places at once: it raises the rate through f_C, and it raises the emulsifying capacity, thus lowering m_∞. Both effects fall off while the requirement grows, and from that the limit follows without any additional parameter being needed.

For grease and wax ($\kappa_0 = 0.10$, $a = 0.40$), starting from $m_\infty = 8$ at $m_{\text{perm}} = 10$ and assuming that twice the dose creates twice the capacity:

Dose	m_∞	Time saved	Added benefit per further dose unit
×1.5	5.3	27 %	54 percentage points
×2	4.0	38 %	22 percentage points
×3	2.7	49 %	9 percentage points
×5	1.6	59 %	4 percentage points
×8	1.0	66 %	2 percentage points

The marginal benefit collapses between twice and three times the dose. This is not a rule of thumb but follows from the structure: f_C has an exponent below 1 (for grease, twice the dose gives only +29 %), while A_{req} runs to infinity as soon as m_∞ approaches the target value. Two diminishing levers against an exploding requirement. Exactly where the cut falls is decided by the price comparison; the only published decision rule on this is coarse and reads: change the bath once the sum of the top-ups equals the original make-up quantity.

9.2 Mixed soil: the kink

$$m(t) = m_A\, e^{-k_A t} + m_B\, e^{-k_B t} \tag{18}$$

Example: 450 mg/m² of a fast fraction ($k_A = 0.5$) plus 50 mg/m² of a slow one ($k_B = 0.05$), both in min^{-1}. Removing 90 % of the soil takes 6.9 minutes. Removing 99 % takes 46 minutes.

The last nine per cent take almost six times as long as the first ninety. Extrapolate a cleaning time from the behaviour at the start and you will dramatically underestimate it for mixed soil. And extending the time attacks the slow fraction only. Often a different detergent aimed at exactly that fraction would be far the cheaper route.

9.3 Thresholds: where the model jumps

- **Melting point of the soil.** Below it almost nothing happens, above it a great deal. For waxes and drawing compounds, reaching the melting point matters more than any further optimisation: 5 °C can be the difference between scrap and success.
- **Concentration optimum.** For some detergent–soil pairs the rate passes through a maximum; above it more chemistry acts *negatively*. That is not a saturation but a reversal, and no formula in this paper captures it.
- **Critical wall shear stress.** Below it an adhering deposit is practically not removed, above it removal sets in.
- **Cavitation threshold in ultrasound.** Below it only the weak flow acts, above it the actual cleaning effect starts abruptly.

10 The special case of aluminium: etching as a second condition

For aluminium ahead of powder coating, degreasing is almost never just degreasing. The bath attacks the substrate and dissolves metal from it: **etch cleaning**. That removal is not a side effect but the purpose: the microcrystalline **deformation layer** from extrusion or rolling has to go so that the conversion coating grows evenly.

Two processes therefore run in one tank, they share one cycle time, and they want different things.

10.1 Why etching is not a fifth factor

The whole paper so far rests on a sentence that is true for cleaning:

More is better. Longer, hotter, more concentrated never harms cleanliness. It only costs money and throughput. That is why equation (9) has only *one* direction: a $\geq$.

For etching this sentence does not hold. Too little removal leaves the deformation layer in place, adhesion suffers, and filiform corrosion is encouraged. Too much removal costs aluminium, dimensional accuracy and appearance, fills the bath with sludge, and every gram of it has been paid for. Etching has a **window**, not a direction.

A window cannot be multiplied into k. Etch removal therefore does not belong in equation (8) but **alongside it, as a second, independent condition**:

$$\underbrace{k(T, c, \tau) \cdot t \;\geq\; A}_{\text{cleanliness}} \qquad \text{and} \qquad \underbrace{G_{\min} \;\leq\; G(T, c, \mathrm{pH}) \cdot t \;\leq\; G_{\max}}_{\text{etch removal}} \tag{19}$$

Both over *the same* time t. From this arises the really interesting question, which nobody can answer by checking them separately: *is there any cycle time at all at which both are met?*

10.2 The etching equation

Unlike the soil, the metal does not decay exponentially: it is inexhaustible. The attackable interface stays constant, and within the technically relevant window removal is **linear with time**:

$$\boxed{G\,[\mathrm{g/m^2}] \;=\; r_0 \cdot \underbrace{\exp\!\left[-\frac{E_A}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right]}_{\text{temperature}} \cdot \underbrace{\left(\frac{c}{c_0}\right)^{n}}_{\text{concentration}} \cdot \underbrace{f_{\mathrm{pH}}}_{\text{acid route only}} \cdot t} \tag{20}$$

r_0 is the removal rate at the reference point, in $\mathrm{g/(m^2\,min)}$.

Here (uniquely in the whole model) there is a genuine measurement from production. QUALICOAT prescribes the etch removal and requires it to be demonstrated by weighing a panel before and after. The plant therefore already knows its value from the last check. *One* measured point at known conditions gives r_0, and with r_0 any other setting can be scaled, the same reference-point device as in Section 7, only with a measurement that genuinely exists.

The QUALICOAT requirement applies across *all* etching steps together:

Table 13 Required minimum removal $G_{\min}$, determined as weight loss on a test panel

Case	removal across all etching steps
Standard	$\geq 1.0\ \mathrm{g/m^2}$
SEASIDE (marine climate)	$\geq 2.0\ \mathrm{g/m^2}$
SEASIDE, final etching step alone	$\geq 0.5\ \mathrm{g/m^2}$

There is no standard for the upper limit $G_{\max}$. Neither QUALICOAT nor GSB names one; GSB expressly refrains from doing so. The only published window figure that could be found is $0.5 \ldots 4.0\ \mathrm{g/m^2}$. $G_{\max}$ is therefore an *in-house limit* and must remain labelled as such: it follows from the dimensional accuracy, appearance and material cost of your own range of parts, not from any code of practice.

10.3 The three numbers and how certain they are

$E_A = 50$ kJ/mol**, and for both routes.** For aluminium in dilute caustic soda between 15 and 45 °C, 45 … 51 kJ/mol is measured. Across all alloys, concentrations and temperature ranges, however, the literature scatters widely, from around 40 to over 90 kJ/mol. The value is a mean, not a constant of nature, and therefore belongs in any evaluation as an *input field*. The widespread claim that acid etching responds less strongly to temperature could not be substantiated; the only direct comparison found on the same alloy in fact shows the opposite trend. Quoting a figure that asserts an unsupported direction would be worse than making no distinction at all.

$n = 1.0$**: concentration enters linearly.** Half the caustic concentration, half the etching rate; the practical literature is unambiguous on this. It is the most striking difference from cleaning, where the exponent a lies well below one.

f_{pH}**: the only term in the whole paper that is not an estimate.** Aluminium is also etched in acid, below pH 4, with fluoride. What attacks is the undissociated hydrogen fluoride, and how much of it is present in free form follows solely from the pH and the pK_a of hydrofluoric acid (3.17):

$$\text{HF fraction} \;=\; \frac{1}{1 + K_a \cdot 10^{\mathrm{pH}}}, \qquad f_{\mathrm{pH}} \;=\; \frac{\text{HF fraction(pH)}}{\text{HF fraction(pH}_0\text{)}} \tag{21}$$

Table 14 Fraction of attacking free hydrogen fluoride

pH	2.0	3.0	4.0	5.0
HF fraction	94 %	60 %	13 %	1.5 %

This is the justification for the shop rule "the lower the better", and it supplies the limit of that rule at the same time. Below pH 2, 94 % is already active. A pH of 1 instead of 2 buys a further six per cent of effect, for twice the acid consumption and the corresponding wear on the plant. That figure appears in no data sheet and is not to be had without the calculation.

10.4 The reversal of the levers

From $n = 1.0$ against $a \approx 0.5$ follows a statement that directly contradicts the rules of thumb of degreasing, and which is therefore regularly overlooked in production:

Table 15 Effect of the same adjustment on both conditions (from 60 °C; cleaning with $E_A = 60$ kJ/mol, $a = 0.5$, $\kappa_0 = 0.15$, $b = 0.8$, $\mu_0 = 0.2$)

Adjustment	factor on **etch removal**	factor on **cleaning**
+10 °C	1.69	**1.88**
Concentration ×2	**2.00**	1.35
Time ×2	2.00	2.00
Mechanics ×2	–	1.59*

** Reference point spraying ($\tau_0 = 1.00$); out of a still immersion bath the entry would be* 1.20.

In cleaning, chemistry is the weakest lever; in etching it is the strongest. Halve the dosage and you lose a quarter of the cleaning effect. You can hardly see that on the parts. But you lose *half* the etch removal, and that goes into the test report. The same plant, the same bath, two opposing rules.

10.5 Etching improves the degreasing

When the metal dissolves beneath the soil, the soil loses its grip, regardless of what the detergent is doing. This effect is **mechanical**, not chemical, and therefore belongs in the mechanics term. Two decisions here, both with reasons:

Additive, not multiplicative. Undercutting works in a still bath too, where $\tau = 0$. A factor applied to τ would give zero there, and that would be wrong. With the relative removal rate $\rho = G/G_0$ and the share p that undercutting has in the mechanical effect at the reference point,

$$\left(\frac{\tau}{\tau_0}\right)_{\text{effective}} = \frac{\dfrac{\tau}{\tau_0} + \gamma\,\rho}{1+\gamma}, \qquad \gamma = \frac{p}{1-p} \tag{22}$$

γ follows from a share, not from a measurement. With the plant on specification, $\tau/\tau_0 = 1$ and $\rho = 1$, and the expression gives exactly 1. The building block cancels out completely at the reference point. Like the rest of the model, it needs no calibration.

An etching bath run too cold loses twice over. The detergent slows down and the etching attack falls off as well: the mechanics goes with it. This coupling is why an etch cleaning bath reacts more sensitively to a drop in temperature than the cleaning equation alone predicts.

10.6 The common range and when there is none

Put both conditions from equation (19) on the same time axis and three statements emerge: from when the cleaning is right, from when to when the removal is right, and the overlap.

The cleaning condition gives a *lower* bound, the etching condition a window bounded on *both* sides. The common range is

$$t \;\in\; \left[\; \max\!\left(t_{\text{clean}},\; \frac{G_{\min}}{G/t}\right),\;\; \frac{G_{\max}}{G/t} \;\right] \tag{23}$$

and this range may be **empty**. The shortest time needed for cleanliness then already exceeds the longest that the removal still permits. Adjusting the cycle no longer helps; one of the other two settings has to move. Which one and in which direction is less obvious than it looks.

Checked separately, each of the two conditions looks solvable on its own. So the time gets extended and nobody understands why it does not get better, and that happens exactly when an answer is needed most. Only side by side on one time axis does the situation become visible at all.

Figure 2 shows both bounds against the two settings that move them. The comparison of the two panels is the real message of this section, and on one point it contradicts what one would expect.

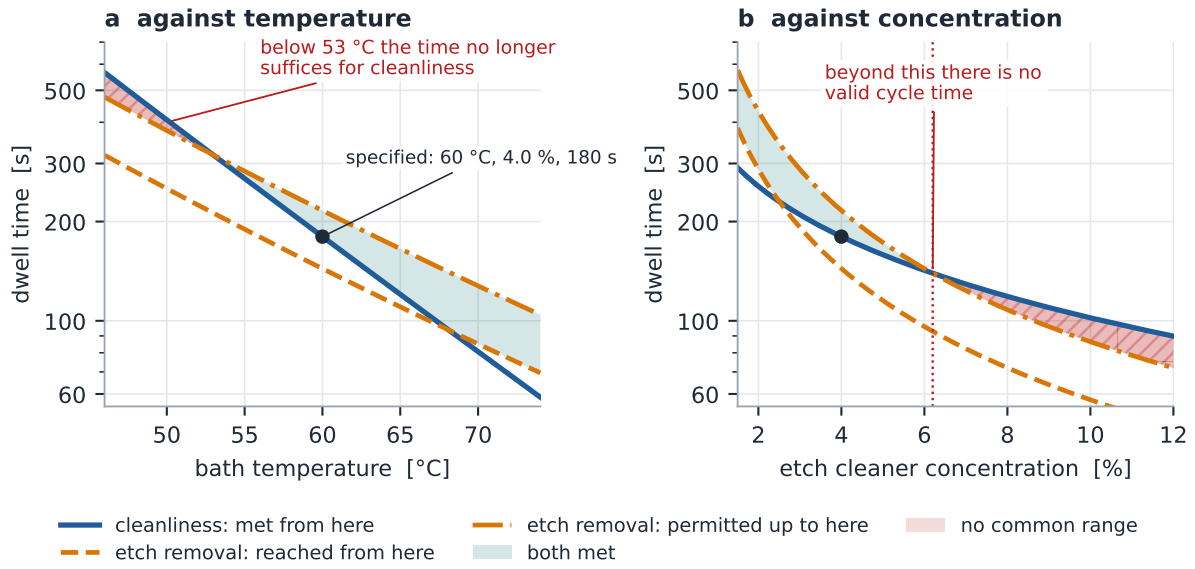

Figure 2 Both conditions plotted against the two quantities that move them. Green: dwell times that satisfy both. Red hatching: there are none. Calculated for an etch cleaning bath removing 2.5 g/m^2 per pass, the SEASIDE requirement ($G_{\min} = 2.0$), an in-house upper limit of $G_{\max} = 3.0$ and a solid grease soil ($E_A = 65\,\text{kJ/mol}$, $a = 0.4$, $b = 0.6$). The figures come from the calculator itself, not from a re-implementation.

Concentration decides whether a common range exists at all. It enters the etch removal linearly and the cleaning only with the exponent $a < 1$ (Section 10.4). More chemistry therefore shortens the *permitted* time faster than it shortens the *required* one. Beyond a certain point no cycle time satisfies both, in the case calculated at 6.2 %. This direction holds for every soil that occurs on aluminium, because $n = 1.0$ exceeds its exponent a. The limiting case in Table 9 is scale, with a up to 1.0: there both bounds move at the same rate and concentration loses its special role.

For temperature the direction depends on the soil. It moves both bounds, but not equally: what governs is which of the two activation energies is larger. For greases and baked-on residues (E_A above 50 kJ/mol) a cold bath closes the window, in the figure below 53 °C. For thin oil (E_A around 38) it is the other way round: there, cooling *opens* the window and heating closes it.

This makes a widespread expectation wrong. Saving on chemistry does *not* make the problem worse: it pushes the lower bound out, so the required gram arrives later, but the common range becomes *larger* in the process. What closes it is overdosing, and along with it everything that lengthens the cleaning alone: dirtier parts, lost circulation, shadowed spots. Those act on the cleanliness bound and leave the etching bounds untouched.

What is estimated in this building block. First E_A (see above): inhibited etch cleaners, whose additives limit the removal, also respond less strongly to temperature than equation (20) assumes; the calculation then overestimates the gain from more heat. Second, the share p (taken as 40 % alkaline, 25 % acid): there is no dependable series of measurements for it; it only changes how strongly a deviation additionally feeds through to the cleaning. Third, $G_{\max}$.

Not covered are dissolved aluminium in the bath (it demonstrably slows the process, so after a bath change the reference value has to be measured again), foreign ions in the acid bath, and the question of whether the connection between too little removal and filiform corrosion is more than plausible. It has been debated in the industry for years; the paper therefore says "encourages" and not "causes".

11 The design problem

$$\begin{aligned} &\text{minimise} && K_E(T) + K_C(c) + K_M(\tau) + K_t \cdot t && \text{(cost per batch)} \\ &\text{subject to} && k(T,c,\tau)\cdot t \;\geq\; S \cdot A && \text{(cleanliness with margin)} \\ & && T \leq T_{\max}, \quad \mathrm{pH} \in [\mathrm{pH}_{\min}, \mathrm{pH}_{\max}] && \text{(material)} \\ & && \tau \leq \tau_{\max} && \text{(part geometry)} \\ & && t \leq t_{\text{cycle}} && \text{(throughput)} \end{aligned} \tag{24}$$

As a rule the cleanliness condition is met with equality, unless the dwell time is bound from below by the conveyor anyway. Temperature and concentration are then the free quantities and are wound down until the condition binds.

11.1 Scatter and the safety factor

The equation works with *one* number for m_0 and *one* for k. Both scatter: m_0 depends on tool life, part position and the waiting time before cleaning, k on bath age, batch size and position in the basket.

Figure 3 plots the reject rate against the safety factor.

Design to $k \cdot t = A$ and you design so that roughly half the parts miss the target. That is not an exaggeration but a property of the mean: the calculated point is by definition the middle of the distribution. Simulated with 300 000 parts and log-normally distributed m_0 and k:

Scatter (m_0 / k)	$S = 1.0$	$S = 1.2$	$S = 1.4$	$S = 1.6$
realistic (40 % / 15 %)	**49.3 %**	13.9 %	2.3 %	0.3 %
unfavourable case (60 % / 25 %)	**49.2 %**	24.7 %	10.5 %	4.0 %

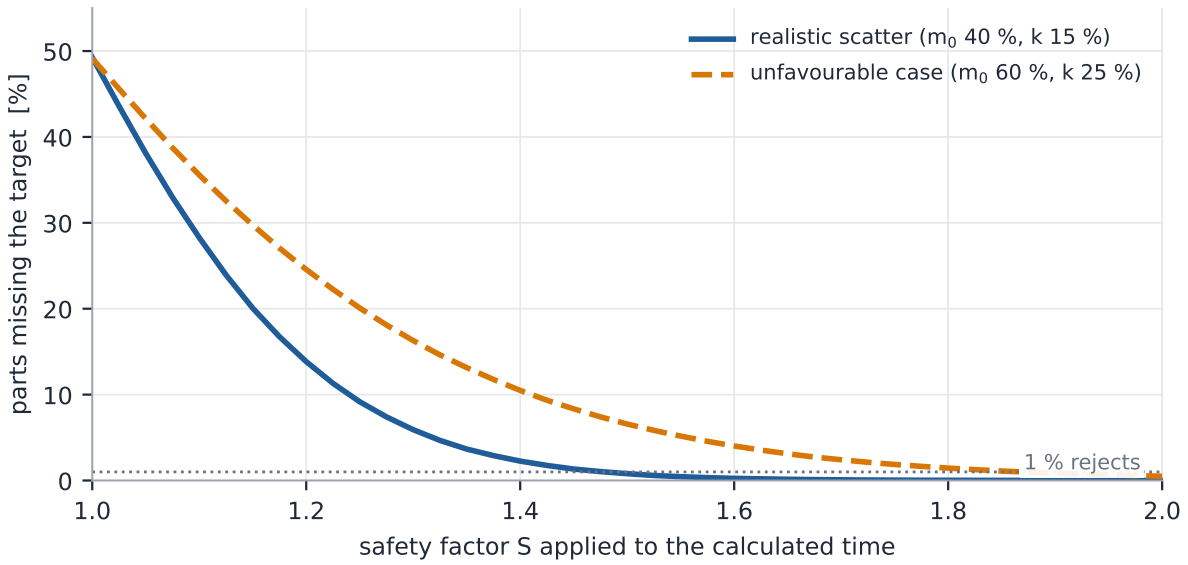

Figure 3 Share of parts missing the target value, against the safety factor. Without a margin it sits at around 50 % irrespective of the scatter: the scatter only determines how quickly the curve falls.

Two corrections are mandatory.

1. **Form A from the upper end of the m_0 distribution**, not from the mean. Measure m_0 on at least ten parts and use the 90 % quantile.
2. **Apply a safety factor to the time:** $t_{\text{design}} = S \cdot A/k$ with $S = 1.4$ where k is well known and the scatter moderate, $S = 1.6 \ldots 1.8$ where k is uncertain or the consequence of one bad part is severe.

That typically amounts to a margin of 40 to 80 % on the calculated minimum time. It sounds like a lot, and it follows from the equation describing means while production delivers individual parts.

11.2 The design rule

Figure 4 shows what follows from this for one concrete case. What is striking there is less the family of curves than the vertical line at 60 °C: for aluminium the time does not suffice even with twice the chemistry and twice the mechanics: here the material is the limit, not the chemistry. That is the normal case as soon as temperature is capped.

Run the temperature as high as the material allows. Then raise the mechanics to the limit set by the part geometry. Set the chemistry to the value recommended by the supplier, and beyond it only where measurements show a benefit. Time comes out as the remainder, multiplied by the safety factor.

But check the headroom first (Section 6.5): a factor already at its limit contributes nothing, however large its exponent. *Exception to the rule:* where the cycle time is free (batch operation, overnight running) and energy is expensive, everything reverses: time is then the cheapest factor of all.

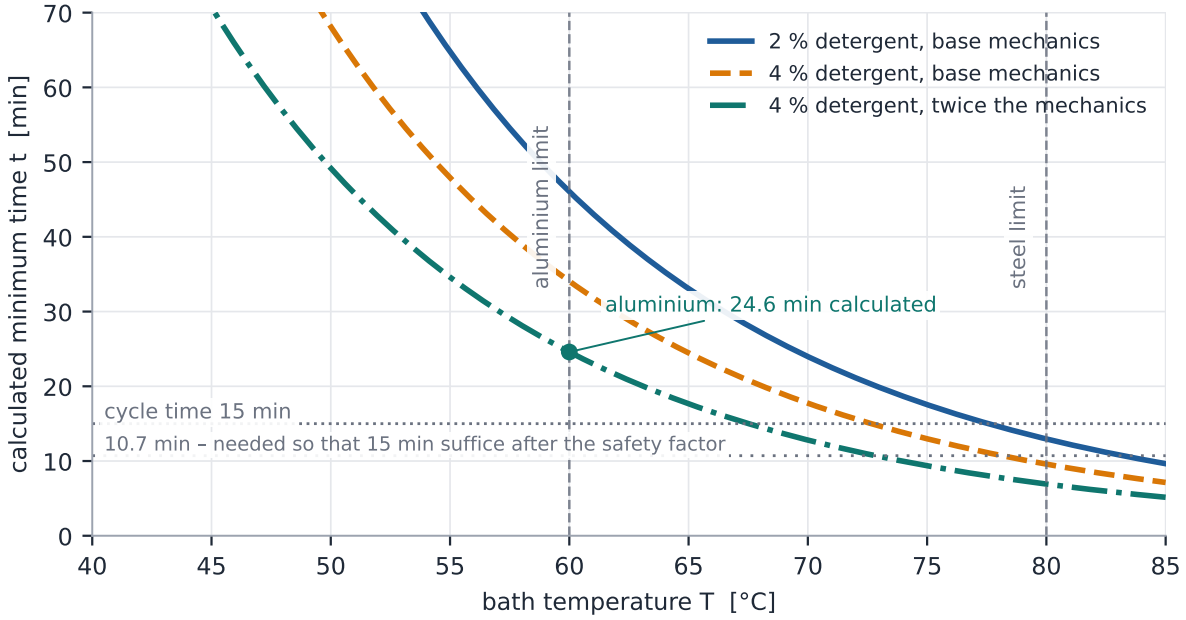

Figure 4 Calculated minimum time against bath temperature for $A = 4.61$, with true Arrhenius and base effects. The lower dotted line shows what is required by calculation so that, after the safety factor ($S = 1.4$), the cycle time of 15 min is met. For aluminium, 60 °C is not enough even with twice the chemistry and twice the mechanics: here the material, not the chemistry, is the limiting factor.

12 A worked example

Task: a component with 500 mg/m^2 of drawing oil, requirement 5 mg/m^2 before powder coating, hence $A = \ln 100 = 4.61$. **Starting point:** 50 °C, 2 % detergent, moderate circulation, measured in trials as $k_0 = 0.05\,\mathrm{min}^{-1}$. **Parameters:** $E_A = 62\,\mathrm{kJ/mol}$, $a = 0.5$, $b = 0.8$, $\kappa_0 = 0.15$, $\mu_0 = 0.2$. **Safety factor** $S = 1.4$.

The calibration here sits at *moderate circulation*, that is at $\tau_0 = 0.20$ on the ladder of Table 3. That statement used to be dispensable and, since equation (6), no longer is: what doubling the mechanics achieves depends on *where* the doubling starts from, and in the direction one does not at first expect. Out of an almost stationary bath it achieves *less* than out of one that is already strongly flushed. The reason is the base effect. In a still bath 73 % of the cleaning happens without any flow at all, from chemistry and heat; what gets doubled there is only the remaining good quarter. When spraying it is the other way round: 80 % hangs on the flow. The figures for $b = 0.8$ and $\mu_0 = 0.2$:

Doubling the impingement starting from …	factor on k	of which flow-dependent
Immersion, still bath	1.20	27 %
Immersion with circulation	1.39	52 %
Immersion with part movement	1.49	66 %
Spraying	1.59	80 %
High pressure	1.69	92 %

In a quiet immersion bath, more circulation is therefore the weakest lever of all: there it is heat and chemistry that help. This is not an artefact of the arithmetic but the very statement this term exists to make.

Setting	k [min^{-1}]	calculated	designed
Starting point 50 °C, 2 %, base mechanics	0.050	92.1 min	128.9 min
+20 °C to 70 °C	0.192	24.0 min	33.6 min
plus 2 → 4 % detergent	0.260	17.7 min	24.8 min
plus twice the mechanics	0.359	**12.8 min**	**17.9 min**
The same part, but in aluminium, only 60 °C *permitted (Table 10):*			
+10 °C to 60 °C, 4 %, twice the mechanics	0.187	**24.6 min**	**34.4 min**

Two statements from one example. *First:* changing the material from steel to aluminium costs just under 12 minutes of cycle time, even though chemistry and mechanics are already doubled and the temperature is run right up to the material limit. *Second:* the safety factor eats up the apparently solved task again: 12.8 minutes calculated becomes 17.9 minutes designed, and a cycle time of 15 minutes is therefore *missed*, not met. The formula says both *before* the plant is built.

13 Test plan: determining the parameters in your own plant

Two blocks, 29 samples, one to two working days. What is determined is the parameter set for *one* combination of soil, detergent and material.

Measured quantity. Residual soil in mg/m^2. Suitable methods: gravimetry after extraction (accurate, slow), fluorescence measurement of oil residues (fast, ideal for series). The water-break test is only qualitative and *not* sufficient. **State the blank and the detection limit of the method**: with m_{perm} in the range of a few mg/m^2 and a sample area of 100 cm^2 that is 0.05 mg in absolute terms, at the limit of ordinary analytical balances. Without a blank the lower part of the curve cannot be evaluated, and that is where the design is decided.

Block 1: testing the model (7 samples). Fixed conditions, residual soil after 1, 2, 4, 8, 16 minutes, plus **two repeats at 4 minutes**. Plot $\ln m$ against t. The repeat points give the measurement scatter; *only a deviation larger than that scatter counts as refuting the straight line.* Determine m_0 on at least three parallel samples and report its scatter. Do this test first: it can refute the base model, and it is meant to be able to.

Block 2: the exponents (22 samples, full factorial with centre points). Do **not** vary each factor on its own. Run a 2^3 design in T, c and τ (two levels each, all eight combinations) plus **three repeats at the centre point**. **Two time points per condition**, for the reason given two boxes below. Hence $8 \times 2 + 3 \times 2 = 22$ samples, not 11. Set the levels as far apart as the material window allows, say $T = 45/65$ °C, $c = 1.5/6$ %, τ in the ratio 1:4. Evaluation: linear regression of $\ln k$ against $\ln(c/c_0)$, $(1/T - 1/T_0)$ and $\ln(\tau/\tau_0)$.

The slopes are *not* a and b: this is where the calculation goes wrong most often. The base effects from Section 6.1 damp the slope at the reference point:

$$\left.\frac{\mathrm{d}\ln f_C}{\mathrm{d}\ln(c/c_0)}\right|_{c_0} = a\,(1-\kappa_0), \qquad \left.\frac{\mathrm{d}\ln f_M}{\mathrm{d}\ln(\tau/\tau_0)}\right|_{\tau_0} = b\,(1-\mu_0)$$

The regression therefore yields a' and b', from which $a = a'/(1-\kappa_0)$ and $b = b'/(1-\mu_0)$.

Skip this step and you feed too small an exponent into an equation that already contains κ_0 and μ_0: the base effect is then counted twice. Over the range proposed above that is some 15 % too little for chemistry and 13 % for mechanics. Only E_A is unaffected: the temperature term is purely multiplicative and has no base effect. Determine κ_0 and μ_0 *first*, then work back.

Why not a one-factor-at-a-time plan. A plan that varies only one factor while holding the others fixed measures main effects by construction and *cannot in principle detect interactions*. It would thereby presuppose exactly the separability that the model claims, a circular argument. The 2^3 design costs the same number of samples and additionally yields all three two-factor interactions plus a curvature test.

If an interaction is clearly larger than the scatter of the centre points, the product form does not hold in this window, and the parameters are usable only at the reference point. If the mean of the centre points differs clearly from the mean of the eight corners, there is a threshold within the range examined.

Never determine k from a single measurement after a fixed time. If the decay is not exactly exponential (and by Section 3 it usually is not), such a one-point k is systematically biased, and more so for the fast conditions than for the slow ones. Worked through for a mixed soil with a true ratio $k(70\,^\circ\mathrm{C})/k(50\,^\circ\mathrm{C}) = 4.00$: measuring after a fixed time of 8 minutes yields a ratio of 1.57, E_A **comes out 67 % too small**, and a and b are dragged down with it. That would falsify exactly the ranking of the levers that the paper offers as its core statement.

Remedy: two time points per condition, placed so that both lie in the same range of residual soil (say 30 % and 10 % remaining).

Randomise the order, monitor the bath. Over a working day the soil load rises and with it the bath load m_∞. Work through the points in sequence and you will confuse every factor effect with the ageing of the bath. Work in random order and repeat the centre point at the start, in the middle and at the end. If those three values drift, the ageing of the bath is larger than the effect you are looking for: change the bath, split the day.

And give a confidence interval for every parameter, not just the point estimate. A coefficient of determination without an error estimate says nothing with so few points. Residual soil measurements scatter by 10 to 30 % depending on the method.

The base effects κ_0 and μ_0 are determined with one extra sample each: once with plain hot water and no detergent, once in a still bath with no circulation. Where that cannot be arranged in production, use the conservative values $\kappa_0 = 0.1$ and $\mu_0 = 0.15$.

14 What the model cannot do

- It does not correctly describe the **initial range of thick, closed films**. There the rate is governed by area, not by amount.
- It does not capture the **entry and exit event**. In laboratory trials around 90 % of the removal was attributable to the air–liquid interface travelling across the surface as the part passes through the bath surface. In practice that means: for *liquid oil films in an immersion bath*, repeated dipping in and out can be more effective than a longer dwell. The model does not know this lever, and it calculates conservatively there. *The finding is narrowly bounded:* a travelling meniscus rolls up a mobile oil film but strips off no grease that is solid at room temperature: machining, drawing and preservation greases cannot be removed by dipping in and out. When spraying the mechanism does not exist at all, because the part passes through no bath surface.
- It does not know the **age of the soil**. Dried, partly oxidised or polymerised oil has a different kinetics from freshly applied oil: k can fall by an order of magnitude. In practice the time between machining and cleaning is often the single largest cause of fluctuating results. Keep it constant during calibration and document it.
- It does not contain **pH as a factor in the cleaning equation**, only as a material limit. With alkaline detergents the pH effect sits unstated inside the exponent a. *Exception:* in the etch removal on the acid route (Section 10), pH appears explicitly.
- It does not know **water hardness or bath conductivity**.
- It does not know **foam**, with many detergents the real upper limit for dosage and mechanics in spray plants.
- It says **nothing about the choice of detergent** and does not cover **rinsing and drying**. Of the **chemical reactions with the base material**, exactly one is covered: the etch removal on aluminium (Section 10), and that as a condition of its own, not as part of k. Passivation, hydrogen uptake, the pickling of other materials and the effect of dissolved aluminium remain outside.
- It has **no spatial resolution**: it treats the whole part as one place. Real parts have local differences that can only be represented by dividing them into zones, at least the visible surface and the worst spot. The worst spot itself the model deliberately does not calculate: $\tau = 0$ would lie outside the supported range of the effect curve (Section 6.1).
- It does **not apply to particulate cleanliness** to VDA 19.1 / ISO 16232.

14.1 Where it breaks inside its own window

The list above says what the model does not *contain*. The following points are less comfortable: there it is *wrong* within its own scope. Each is already conceded at the place named; here they stand together so that nobody reads past them.

- **The window of validity** is narrow: roughly ±15 °C in temperature, a factor of 2 in concentration, a factor of 3 in mechanics around the calibration point (Section 6.4). Outside it the parameter set no longer holds.
- **The factors are coupled**, so the product form is false in general: five known couplings are listed in Section 6.4. The equation is the first-order expansion about the reference point, no more.

- **The concentration optimum** cannot be represented by a power law with $a \geq 0$ (Section 5.2). Above the customary dosage, measure rather than calculate.
- **The critical wall shear stress** is not represented by the mechanics term. Below it an adhering deposit is practically not removed, not even with more time.
- **Ultrasound** has neither a rung on the mechanics ladder nor a transferable exponent; the cavitation threshold is a jump, not a power law.
- **A single k is always a mean** over the interval considered, because binding strengths are distributed (Schott 1976). Across a whole decade of soil removal it is exact at neither end.
- **In the etching part** the upper limit $G_{\max}$ has no standard behind it, and the share p of undercutting is assumed without a series of measurements (Section 10).

15 Next steps

1. **Measure block 1 on a real case.** Seven samples are enough to see whether the base model holds. The only step requiring no groundwork, and the only one that could refute the model. That the model can be calculated with does not replace this test; it raises its value.
2. **Check the etching part against your own test reports.** Unlike residual grease, a figure the plant genuinely measures exists here. Anyone recording temperature, concentration and time across several QUALICOAT checks can test equation (20) against their own figures, without a single additional trial. That is the cheapest test of the model this paper has to offer.
3. **Bring in bath maintenance.** The bath load m_∞ remains the economically most important parameter of all and at the same time the one least often determined. As long as it is estimated rather than measured, any design close to the target value is uncertain.

A List of symbols

Symbol	Unit	Meaning
m_0	mg/m^2	amount of soil at the start
m_{perm}	mg/m^2	permitted residual soil
m_∞	mg/m^2	bath load from redeposition
A	–	cleaning task, $\ln(m_0/m_{\text{perm}})$
E	–	degree of fulfilment, $(k\,t)/(k_0\,t_0)$; share of the promised work
S	–	safety factor applied to the time
k	min^{-1}	cleaning rate
k_0	min^{-1}	cleaning rate at the reference point
t	min	cleaning time
T	K	absolute bath temperature
E_A	J/mol	*apparent* activation energy
R	8.314 J/(mol K)	universal gas constant
Q_{10}	–	rate factor per +10 °C (itself temperature dependent)
c	mass-% or g/l	detergent concentration; the same unit throughout in the ratio c/c_0
a	–	exponent of chemistry
κ_0	–	base effect of chemistry (effect without detergent)
τ	N/m^2	wall shear stress
b	–	exponent of mechanics
μ_0	–	base effect of mechanics (effect without flow)
θ	s	time constant of component heat-up
s	m	section thickness
α	W/(m^{2}K)	heat transfer coefficient
λ	W/(m K)	thermal conductivity of the material
Bi	–	Biot number, $\alpha s/(2\lambda)$; validity measure of the lumped model
etching part only, Section 10		
G	g/m^2	etch removal, across all etching steps together
$G_{\min}$, $G_{\max}$	g/m^2	lower (standard) and upper (in-house) window limit
r_0	g/(m^2 min)	removal rate at the reference point, from the audit
n	–	exponent of concentration in etch removal; $n = 1.0$
K_a	–	acid constant of hydrofluoric acid, p$K_a = 3.17$
ρ	–	removal rate relative to the reference point, G/G_0
p	–	share of undercutting in the mechanical effect
γ	–	$p/(1-p)$; weight of undercutting in the mechanics term

B Derivation of the heat-up loss

With the approximation $k(T)/k(T_{\text{bath}}) = Q_{10}^{(T-T_{\text{bath}})/10}$ and the heat-up curve $T(t) = T_{\text{bath}} - \Delta T_0\, e^{-t/\theta}$, writing $\beta = \frac{\ln Q_{10}}{10}\,\Delta T_0$:

$$\Delta t_{\text{loss}} = \int_0^\infty \left[1 - \frac{k(T(t))}{k(T_{\text{bath}})}\right] \mathrm{d}t = \int_0^\infty \left[1 - e^{-\beta e^{-t/\theta}}\right] \mathrm{d}t = \theta \int_0^\beta \frac{1 - e^{-u}}{u}\, \mathrm{d}u = \theta \cdot \mathrm{Ein}(\beta) \qquad (25)$$

For large β, $\mathrm{Ein}(\beta) \approx \ln\beta + \gamma_{\text{E}}$ with the Euler–Mascheroni constant $\gamma_{\text{E}} = 0.5772$.

The loss grows only logarithmically with the temperature difference. So entering cold is less

Table 16 Time lost, in multiples of the time constant θ (at $Q_{10} = 2$)

ΔT_0 [°C]	β	$\Delta t_{\text{loss}}/\theta$
10	0.69	0.59
20	1.39	1.02
30	2.08	1.35
40	2.77	1.61
50	3.47	1.83

damaging than one might expect, as long as the time constant stays small. What matters is θ, and θ depends linearly on the section thickness.

A note on accuracy: for this estimate Q_{10} is taken as constant, although it is itself temperature dependent. Over a heat-up span of 40 °C this *under*estimates the loss, by 0.3 % when checked against true Arrhenius, so immaterial for an estimate. The closed form itself was checked against a direct simulation of the same temperature curve and agrees with it exactly; that is a check of the algebra, not of the model.

Sources and further reading

- Sinner's circle, background and history: `https://en.wikipedia.org/wiki/Sinner%27s_circle`
- E. Kissa: *Kinetics and Mechanisms of Detergency*, Textile Research Journal, 1978: `https://journals.sagepub.com/doi/10.1177/004051757804800705`
- H. Schott: *Limitations in Studying the Kinetics of Fabric Detergency*, Textile Research Journal, 1976, the key critical paper on the applicability of simple rate laws: `https://journals.sagepub.com/doi/10.1177/004051757604600614`
- *Kinetic studies of detergency I: Analysis of cleaning curves*, JAOCS: `https://link.springer.com/article/10.1007/BF02822456`
- *Kinetic studies of detergency II: Effect of age, temperature, and cleaning time on rates of soil removal*, JAOCS, source for the Dupré mechanism with its share of around 90 % of the removal: `https://link.springer.com/article/10.1007/BF02822457`
- T. Sandle: *Making cleaning better: Sinner's Circle and beyond* (extension by soil and surface factors after Wilson and Fryer): `https://www.rssl.com/insights/life-science-pharmaceuticals/making-cleaning-better-sinners-circle-and-beyond/`
- CIP modelling with removal and redeposition, Food and Bioproducts Processing: `https://www.sciencedirect.com/science/article/abs/pii/S0960308502703386`
- *On the nature of the optimum cleaning concentration for dairy fouling*, evidence for the concentration optimum: `https://www.sciencedirect.com/science/article/abs/pii/S0023643818310016`
- *Critical wall shear stress for the EHEDG test method*, critical wall shear stress: `https://www.sciencedirect.com/science/article/abs/pii/S0255270103001016`
- *The Fluid Mechanics of Cleaning and Decontamination of Surfaces*, Annual Review of Fluid Mechanics: `https://www.annualreviews.org/content/journals/10.1146/annurev-fluid-022820-113739`
- The doubling rule and its relation to Arrhenius: `https://en.wikipedia.org/wiki/Q10_(temperature_coefficient)`

- Technical cleanliness, test standards VDA 19.1 and ISO 16232: https://www.rjl-microanalytic.de/en/restschmutzanalyse-technische-sauberkeit-vda-19-iso-16232/pruefnormen-pruefvorschriften-restschmutz-technische-sauberkeit/
- Component Cleanliness Code CCC to VDA 19 / ISO 16232, why particle count and not mass: https://www.crownindservices.com/pdfs/Understanding%20the%20Component%20Cleanliness%20Code%20or%20CCC%20of%20VDA%2019%20_%20ISO%2016232%20-%20Crown%20March%202018.pdf

Additional sources for the etching part (Section 10)

- QUALICOAT, Specifications 2024 edition: required minimum removal, permitted etching sequences and the determination of removal by weighing before and after.
- GSB AL 631: likewise requires a minimum removal and expressly refrains from setting an upper limit.
- M. Miadoková, V. Molnárová-Plchová: kinetics of the dissolution of aluminium in caustic soda, *Chemical Papers* 39 (1985): linearity of removal with time and activation energies in the range 45 . . . 51 kJ/mol.
- POLYSURFACES / Polymedia, one- and two-component etching processes: free fluoride, bath temperature and foreign ions as the three governing quantities of acid etching.
- pK_a of hydrofluoric acid: the common reference works give 3.17 to 3.20; the difference shifts the HF fractions by at most two percentage points.
- JOT / Springer, Hüllweck et al.: filiform corrosion and etch removal; at the same time evidence that this connection is disputed within the industry.