

WORKING PAPER

A Continuum Theory of Channeling in Espresso and the Grind Paradox

Growth Law, Channel Spacing, and the Defects That Make a Finer Grind Extract Less

Supervised by Dr. T. Zamrik

16-OCT-2026

Gradient flow of the Peaks function: $\dot{\mathbf{x}} = \nabla f(\mathbf{x})$, $f(x, y) = |3(1-x)^2 e^{-x^2-(y+1)^2} - 10(x/5 - x^3 - y^5) e^{-x^2-y^2} - \frac{1}{3} e^{-(x+1)^2-y^2}|$



WP-2026-15315794
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2026 T. Zamrik. All rights reserved. Published by Zamrik Quant Research. No part of this document may be reproduced, stored, distributed or adapted, in whole or in part, by any means, without the prior written permission of the publisher, except for brief quotation with full citation. Requests for permission: info@zamrik.com.

Suggested Citation

Applied Mathematics Desk, Zamrik Quant Research (2026). ‘A Continuum Theory of Channeling in Espresso and the Grind Paradox: Growth Law, Channel Spacing, and the Defects That Make a Finer Grind Extract Less’. *Zamrik Quant Research Working Paper WP-2026-15315794*. Available at <https://zamrik.com/research-items/wp-2026-15315794/>.

About Zamrik Quant Research

Zamrik Quant Research is the research arm of Zamrik Quant Lab. It works through five research desks, from market microstructure and derivatives to optimal control, stochastic analysis and applied mathematics, and publishes everything it finds in the open: research papers, notebooks that run on Kaggle, free tools that answer one practical question, and essays.

The five desks are Market Microstructure & Trading; Derivatives & Volatility; Optimal Control & Decision; Stochastic Analysis; and Applied Mathematics. Each paper, notebook, tool and essay belongs to exactly one desk. Zamrik Quant Research is run by a small, independent team, headed by Dr. T. Zamrik. We are self-funded: no sponsors, no advertisers, no outside capital.

Every publication passes through the same five functions before anyone outside sees it: research, data, software, verification and production. We never retrofit the mathematics to a conclusion chosen in advance. Every number is checked two independent ways. Every figure is re-run from its code. When we find an error, or replace a piece with a better one, we record it in the public corrections log at <https://zamrik.com/corrections/>.

Legal Notice

This paper is research and education. Nothing in it is investment, financial, legal or engineering advice; nothing is a recommendation to buy or sell shares of stock, derivative contracts, bonds, or any other instrument or property; and nothing is a forecast. You are responsible for your own decisions. Figures and code are provided as they are. Care and independent checks reduce errors but cannot rule them out, so no warranty is given. The views expressed are those of the Applied Mathematics Desk.

WORKING PAPER



A Continuum Theory of Channeling in Espresso and the Grind Paradox

Growth Law, Channel Spacing, and the Defects That Make a Finer Grind Extract Less

Supervised by Dr. T. Zamrik¹

16-OCT-2026

WP-2026-15315794

Abstract. Finer coffee can brew weaker espresso, and the only mechanism proposed so far uses two lumped flow paths. We treat channeling as an instability of a continuous dissolving bed at fixed pressure: Darcy flow through a Kozeny–Carman bed that opens as its solubles dissolve, solute transport and first-order extraction, in a two-dimensional section of the puck with one calibrated constant. The extraction front moves at $\gamma_f q$, and a thin front grows at a closed-form rate, confirmed against a direct simulation within 0.11 percent and a moving-frame eigenvalue within 5.0 percent. The exact linear problem along the moving base puts the fastest channel spacing at 3.6 mm at every grind and shows unevenness decaying over a shot coarser than 1.58 times the reference grind. In 1264 simulated shots, millimeter-scale porosity noise costs up to 5.35 percent of the extraction at the finest grind and flattens the extraction-grind curve without reversing it; streaks running the full height of the puck do not reverse it while they are 1 mm wide. The curve turns over only when the unevenness runs the full height and is at least about one channel spacing wide, and then shallowly, by 0.39 points or 1.7 percent, far less than the 2020 experiment’s drop. Only better puck preparation reduces channeling, recovering 54 percent of the loss.

WP-2026-15315794 · Applied Mathematics Desk · Working Paper · Published

16-OCT-2026 · Reading time: about 57 min

1. Introduction

A finer grind exposes more surface and slows the water, so it should extract more coffee. Cameron et al. [1] measured the opposite below a point. With 20 g of coffee, 40 g of beverage and 6 bar, the extraction yield of their espresso rose as the grind was made finer down to setting 1.7 of their EK43 grinder and fell below it, so that settings 1.3 and 2.0 gave the same 22 percent. Their own one-dimensional homogenized model rises monotonically as the grind gets finer; to recover the drop they imposed a fraction of the coffee that the water cannot reach, and attributed it to fines clogging part of the bed.

Lee, Smith and Arshad [2] proposed a mechanism instead of a fraction. In their model the bed is two parallel paths sharing a fixed total flow. Each path dissolves and opens as it extracts, the more porous path takes more of the flow and extracts faster, and once it has given up all its solubles the water through it carries nothing more: the mean extraction

¹Correspondence: Applied Mathematics Desk, Zamrik Quant Research, info@zamrik.com.



falls. The model reproduces a peak in extraction against grind, but only after fitting a transfer coefficient, a porosity contrast and the largest extractable fraction, and after doubling the density of the grains, a value its authors call unphysical. They name a model with more paths as the next step.

Two questions follow: what happens when the paths form a continuum and water can cross into slower neighbors, and whether the fast path is an accident of the packing or is selected by the flow.

Channeling is the barista's word for water finding a fast way through the puck. The physics behind it is older than espresso. When a fluid dissolves the porous matrix it flows through, a region that carries more flow dissolves faster, its permeability rises, and it draws still more flow. This feedback is the reactive-infiltration instability. Chadam et al. [7] analyzed it for a reaction front moving through a porous medium, Hinch and Bhatt [8] for an acid front moving through porous rock, and Szymczak and Ladd [9] for the dissolution of a porous matrix.

An espresso puck has every ingredient. Water is driven at a fixed pressure through a packed bed; up to 30 percent of the bed's mass is soluble [1], so the permeability of a fully extracted region exceeds that of fresh coffee many times over; and a front between extracted and fresh coffee moves down the bed within a single shot. Three features of the puck matter below. The shot is short, so the instability has only a few growth times to act, and it can switch off again as the solubles run out. A unit volume of saturated water holds more solute than a unit volume of bed has to give, so the front between extracted and fresh coffee moves at a sizable fraction of the speed of the pore water. And at the grinds baristas use, the front is as thick as the bed, so the thin-front picture of the classical analyses is a guide rather than an answer.

The questions are therefore quantitative: how fast a channel grows at fixed pressure, how far apart the fastest channels are, at which grind the growth over a shot vanishes, and what channels cost in the cup.

Models of the espresso puck so far are lumped, one-dimensional, or coupled in one direction only. Moroney et al. built a multiscale model of extraction from a coffee bed, resolved in depth and validated against experiment [4], and analyzed the dominant mechanisms of the extraction process asymptotically [5]; the homogenized model of Cameron et al. [1] belongs to the same line and is one-dimensional. Moroney et al. [6] computed the flow through coffee beds and asked how uniform the extraction is, with the flow coupled to the extraction in one direction only: the bed shapes the flow, but the extraction does not change the bed's resistance. Lee, Smith and Arshad [2] closed the loop in both directions, but in two lumped paths, with no space, no pressure field and no stability analysis. Waszkiewicz et al. [3] showed that the puck behaves as a poroelastic bed whose flow saturates as the pressure rises, and that dissolution opens the bed; their model, like the others, resolves depth only.

None of these lets water cross from one path to its neighbor, which is what decides whether an instability selects a spacing, whether a small unevenness grows or decays over a shot, and whether two lumped paths are a fair stand-in for a bed in which every path borders



another.

We treat channeling as an instability of a continuous dissolving bed at fixed pressure and follow it from the linear law to the cup, in six steps.

First, the extraction front moves at a speed fixed by a mass balance (Proposition 1), and a thin front at fixed pressure grows at a closed-form rate (Proposition 2), with a parallel-column limit and a growth proportional to the wavenumber at short waves (Corollary 1). A direct simulation, the formula and a co-moving eigenvalue agree on the growth of a seeded mode. Second, the exact linear problem along the moving base gives the fastest channel spacing, 3.6 mm at every grind, and a threshold, 1.58 times the reference grind, coarser than which unevenness decays over a shot. Third, in two-dimensional shots with exact solute bookkeeping (Lemmas 1 and 2), millimeter-scale unevenness costs extraction at fine grinds and flattens the extraction-grind curve, but does not reverse it. Fourth, the curve reverses only when the unevenness runs the full height of the puck and is at least about one channel spacing wide, and then shallowly. Fifth, run with the 2020 recipe, the model matches the measured ratio of shot times but neither the shape nor the size of the measured drop. Sixth, of three levers a barista has, only better puck preparation reduces channeling.

One constant is calibrated, the permeability scale, on a published shot from a different machine and coffee from the experiment we compare with. Every other parameter comes from the literature, and every number reported below is computed by the scripts that accompany the paper.

2. The Model

Definition 1 (The puck and the shot). A puck is a vertical section through a coffee bed: $x \in [0, W]$ across, periodic in x , and $z \in [0, H]$ down, from the shower screen at $z = 0$ to the filter at $z = H$. Its state at time t is the porosity $\phi(x, z, t)$, the undissolved solubles $s(x, z, t)$ per unit bed volume, and the concentration $c(x, z, t)$ of solute in the pore water. The grind enters through one effective grain size d , measured against a reference grind d_{ref} . Water enters at $z = 0$ with $c = 0$ under a pressure drop ΔP held across the puck; the bed starts full of fresh water; and the shot ends when the volume of beverage that has left through $z = H$, per unit basket area, reaches V_{out} , the beverage mass divided by the density of water and by the basket area. The extraction yield EY is the mass of solute that reached the cup divided by the dose.

The main recipe is an 18 g dose, 36 g of beverage and $\Delta P = 9$ bar. The section spans the $W = 58$ mm basket of Waszkiewicz et al. [3]; their 18.50 g filled it to 14 mm, a dry bulk density of $\rho_b = 500.15$ kg/m³, so 18 g fill it to $H = 13.62$ mm. The reference grind d_{ref} is setting 1.7 of the grinder of Cameron et al. [1], where their extraction peaked, and the bed starts full of fresh water as in their model. The shot is ended by volume, as a barista ends it by weight, with the beverage at the density of water, 997 kg/m³, as in [1]. The grind grid used throughout is $d/d_{\text{ref}} = 2^{j/6}$, $j = -6, \dots, 6$: thirteen grinds from 0.5 to 2.

Definition 2 (Permeability and the opening of the bed). The initial porosity $\phi_i(x, z)$ has



mean ϕ_0 , and the initial solubles are $s_i = \text{EY}_{\max} \rho_s (1 - \phi_i)$, with ρ_s the grain density and $\text{EY}_{\max}$ the extractable fraction. Dissolved solid keeps its volume in solution, so the porosity rises by the volume dissolved, $\phi = \phi_i + (s_i - s)/\rho_s$, and a fully extracted bed has gained $\Delta\phi_{\max} = s_0/\rho_s$, with $s_0 = \text{EY}_{\max} \rho_b$. The permeability is of Kozeny–Carman form,

$$k(\phi) = k_{\text{ref}} \left(\frac{d}{d_{\text{ref}}} \right)^2 \left(\frac{\phi}{\phi_0} \right)^3 \left(\frac{1 - \phi_0}{1 - \phi} \right)^2,$$

where k_{ref} is the permeability of fresh coffee at the reference grind.

Waszkiewicz et al. [3] use the same form. The values are $\phi_0 = 0.1728$, the volume fraction of grounds 0.8272 quoted by Cameron et al. [1]; $\text{EY}_{\max} = 0.30$, their upper limit of extraction [1]; $\rho_s = \rho_b/(1 - \phi_0) = 604.63 \text{ kg/m}^3$ and $s_0 = 150.04 \text{ kg/m}^3$, so $\Delta\phi_{\max} = 0.2482$. A fully extracted bed therefore reaches porosity 0.421 and 29.5 times the permeability of fresh coffee: the opening is the feedback that drives channeling. Because ϕ_i has mean ϕ_0 exactly, the dose is the same for every unevenness. The grind enters the permeability only through the factor $(d/d_{\text{ref}})^2$; k_{ref} is the one calibrated constant of the paper (Example 1).

Definition 3 (The extraction rate). Solubles pass from the grains into the pore water at the rate

$$r = k_m a(\phi) (c_{\text{sat}} - c) \frac{s}{s_i}, \quad a(\phi) = a_{\text{ref}} \frac{d_{\text{ref}}}{d} \frac{1 - \phi}{1 - \phi_0},$$

per unit bed volume, where k_m is a mass-transfer coefficient, a the grain surface per unit bed volume, a_{ref} its value for fresh coffee at the reference grind, and c_{sat} the concentration at which the pore water is saturated. Write $K_0 = k_m a_{\text{ref}} d_{\text{ref}}/d$ for the rate constant of fresh coffee at grind d .

The rate is linear in the undersaturation, as in the transfer term of Lee, Smith and Arshad [2], which goes back to the multiscale models of Moroney et al. [5]. It vanishes as the grain empties, as the rate of Cameron et al. [1] does; we keep that dependence to first order in the solubles left, where [2] switch the rate off with a step at the largest extractable fraction. The values are $k_m = 8.0 \times 10^{-6} \text{ m/s}$, the best fit of [2]; $c_{\text{sat}} = 212.4 \text{ kg/m}^3$, as listed in [2] from [1]; and $a_{\text{ref}} = 8817.6 \text{ m}^{-1}$, from the linear fit in [2] of grain surface area to grind setting for the grinder of [1], evaluated at setting 1.7. The rate constant at the reference grind is $K_0 = 0.0705 \text{ s}^{-1}$, a dissolution time of 14.2 s. The grind enters the rate only through $a \propto 1/d$, so one factor d/d_{ref} carries the whole grind dependence of the model: $k \propto d^2$ and $a \propto d^{-1}$.

Assumption 1 (Standing assumptions). The puck is a two-dimensional section with periodic sides. Darcy’s law holds with a constant viscosity μ . Dissolved solid keeps its volume in solution. The bed is rigid and saturated from the first instant: it neither compacts nor swells, and there is no wetting stage. The shot is isothermal, fines do not migrate, and there is no gas. Solute spreads by molecular diffusion with diffusivity D and by advection, with no mechanical dispersion. One effective grain size sets both the permeability and the grain surface.



Each item is a limit, not a finding. Remark 5 shows what the missing wetting stage does to pre-infusion, and Remark 6 lists what the rest leave out. The viscosity, $\mu = 3.0 \times 10^{-4}$ Pa s for water near 93 °C, cancels from every result, because the calibration fixes k_{ref}/μ ; the diffusivity is $D = 2.0 \times 10^{-9}$ m²/s.

Start. Under Assumption 1 the water moves by Darcy's law through the bed of Definitions 1 and 2, carries solute, and gains it at the rate of Definition 3. Write $\mathbf{u} = (u, w)$ for the Darcy flux, w its downward component, and p for the pressure.

Steps. The liquid gains the volume of the solid it dissolves, r/ρ_s per unit time, and so does the pore space, since $\partial_t \phi = r/\rho_s$ by Definition 2. The two cancel in the volume balance of the liquid, $\partial_t \phi + \nabla \cdot \mathbf{u} = r/\rho_s$, which leaves the flow divergence-free at every instant:

$$\mathbf{u} = -\frac{k(\phi)}{\mu} \nabla p, \quad \nabla \cdot \mathbf{u} = 0, \quad p = \Delta P \text{ at } z = 0, \quad p = 0 \text{ at } z = H. \quad (2.1)$$

Solute in the pore water is advected, diffuses and is supplied by dissolution; the grains lose what the water gains:

$$\partial_t(\phi c) + \nabla \cdot (\mathbf{u} c) - \nabla \cdot (\phi D \nabla c) = r, \quad \partial_t s = -r, \quad (2.2)$$

with $c = 0$ in the water entering at $z = 0$, no diffusive flux through the top or the bottom, and periodic sides. Adding the two, the total solute $\phi c + s$ is conserved,

$$\partial_t(\phi c + s) + \nabla \cdot (\mathbf{u} c - \phi D \nabla c) = 0, \quad (2.3)$$

so the only way out is through the filter. With $\langle \cdot \rangle$ the mean over x , the beverage volume per unit basket area is $V(t) = \int_0^t \langle w \rangle_{z=H} dt'$, the shot time T is the first time with $V(T) = V_{\text{out}}$, and

$$\text{EY} = \frac{A_b}{M_{\text{in}}} \int_0^T \langle w c \rangle_{z=H} dt, \quad (2.4)$$

with A_b the basket area and M_{in} the dose.

End. If the bed does not vary across, the flux does not either, and $\nabla \cdot \mathbf{u} = 0$ makes it the same at every depth: the column is a series of resistances, and

$$q(t) = \frac{\Delta P}{\mu \int_0^H dz/k(\phi)}. \quad (2.5)$$

As the bed dissolves, k rises and so does q : at fixed pressure the shot speeds up as it goes. This one-dimensional identity is checked against the two-dimensional solver at every step of a uniform shot, in the checks of the full puck below.

Start. Measure depth in units of H , the Darcy flux in units of a scale U , and concentrations in units of c_{sat} . The equations of the model then depend on five groups,

$$\text{Da} = \frac{K_0 H}{U}, \quad \text{Pe} = \frac{UH}{\phi_0 D}, \quad \gamma = \frac{c_{\text{sat}}}{s_0}, \quad \alpha = \frac{c_{\text{sat}}}{\rho_s}, \quad \Delta \phi_{\text{max}} = \frac{s_0}{\rho_s}, \quad (2.6)$$



and on the shape of the permeability law. The Damköhler number Da compares the puck height with the length over which fresh water saturates, the Péclet number Pe compares advection with diffusion, γ is the capacity of the water relative to the solubles, α the volume the solute occupies, and $\Delta\phi_{\max}$ how far the bed can open. Only Da and Pe depend on the grind. The question is which flux to use for U .

Steps. At the first instant the bed is fresh and uniform, so the flux is the same everywhere and equals $q_0 = \Delta P k_{\text{ref}} (d/d_{\text{ref}})^2 / (\mu H)$ (Definition 2 with $\phi = \phi_0$). With $K_0 = k_m a_{\text{ref}} d_{\text{ref}} / d$ from Definition 3,

$$Da_0 = \frac{K_0 H}{q_0} = \frac{k_m a_{\text{ref}} \mu H^2}{k_{\text{ref}} \Delta P} \left(\frac{d}{d_{\text{ref}}} \right)^{-3}, \quad Pe_0 = \frac{q_0 H}{\phi_0 D} \propto \left(\frac{d}{d_{\text{ref}}} \right)^2. \quad (2.7)$$

At fixed pressure the initial Damköhler number goes exactly as the inverse cube of the grind: one power from the grain surface, $a \propto 1/d$, and two from the permeability, $k \propto d^2$. Halving the grind multiplies it by eight.

End. The flux does not stay at q_0 . As the solubles dissolve the bed opens and, at fixed pressure, the flux rises with it. Over the shot the natural scale is the mean flux V_{out}/T , which gives

$$Da = \frac{K_0 H T}{V_{\text{out}}} \propto \left(\frac{d}{d_{\text{ref}}} \right)^{-1} T(d). \quad (2.8)$$

The shot time would scale as d^{-2} only if the bed kept its initial permeability. It does not: a fine bed spends longer in the basket and opens further before the beverage is out, so the mean-flux Damköhler number falls more slowly than d^{-3} and obeys no power law.

Method. The uniform puck is solved in one dimension at each of the thirteen grinds by the independent solver (300 cells, adaptive Runge–Kutta, the shot ended exactly at V_{out} by event location). The groups of the scaling derivation are evaluated with the initial flux q_0 and with the mean flux V_{out}/T ; the ratio of the two fluxes measures how far the bed opens during the shot.

Parameters. The main recipe of Definition 1; $d/d_{\text{ref}} = 2^{j/6}$, $j = -6, \dots, 6$; the values of Definitions 2 and 3; k_{ref} from Example 1.

Result. Table 1. On the initial flux, Da_0 runs from 119 at $d = 0.5 d_{\text{ref}}$ to 1.85 at $d = 2 d_{\text{ref}}$, a ratio of 64.0. On the mean flux, Da runs from 11.4 to 0.581, a ratio of 19.6, not 64; a least-squares slope of $\log Da$ against $\log d$ over the grid is -2.12 , not -3 . The reason is the opening of the bed: the mean flux is 10.4 times the initial flux at $d = 0.5 d_{\text{ref}}$ and 3.19 times at $2 d_{\text{ref}}$, so the shots last 81.1 and 16.5 s, a ratio of 4.90 rather than the 16 of a bed that never opened. The Péclet number is 638 or more, so advection dominates at every grind. The groups that do not depend on the grind are $\gamma = 1.42$, $\alpha = 0.351$ and $\Delta\phi_{\max} = 0.248$, and the beverage is 5.81 pore volumes.



Table 1: The dimensionless groups of the uniform puck across the grind grid, on the initial and on the mean flux, with the shot time and the factor by which the flux grows over the shot.

d/d_{ref}	Da, initial flux	Da, mean flux	Pe, initial flux	Pe, mean flux	shot (s)	mean / initial flux
0.50	119	11.4	638	6642	81.1	10.40
0.63	59.3	6.15	1013	9776	55.1	9.65
0.71	42.0	4.64	1277	11554	46.6	9.05
1.00	14.8	2.17	2553	17476	30.8	6.84
1.41	5.24	1.10	5107	24283	22.2	4.75
1.59	3.71	0.889	6434	26825	20.1	4.17
2.00	1.85	0.581	10214	32568	16.5	3.19

Error. The inverse-cube law of the initial Damköhler number holds to round-off: $\text{Da}_0 (d/d_{\text{ref}})^3$ is the same at all thirteen grinds to a relative 2.2×10^{-16} , the error expected of an identity. The mean-flux number departs from the inverse cube by a factor of $64.0/19.6 = 3.26$ over the grid; that departure is the measured effect of the opening, not an error.

Remark 1 (Where espresso sits: a front as thick as the puck). In the reactive-infiltration framing of Szymczak and Ladd [9], a dissolving porous matrix is placed by a Damköhler number, a Péclet number and the capacity of the fluid. Espresso is strongly advective, $\text{Pe} \geq 638$ at every grind, and its Damköhler number at the reference grind is 2.17 on the mean flux and 14.8 on the initial one: fresh water saturates over a length between $H/14.8$ and $H/2.17$, so the zone where grains are being extracted is as thick as the puck. Its capacity is high, $\gamma = 1.42$: a front with fresh coffee ahead of it moves at 1.14 times the Darcy flux, about a fifth (0.197) of the speed of the pore water. And the bed opens during the shot, so the mean flux is 6.84 times the initial flux at the reference grind.

A thin front, the setting of the classical analyses [7], [8], is reached only at grinds finer than a barista uses: at $d = 0.25 d_{\text{ref}}$, where the initial Damköhler number is 949, the front's 10–90 percent width is 0.80 mm in a puck 13.62 mm high. There the thin-front law is tested; at practical grinds the growth over a shot comes from the exact linear problem along the moving base.

Example 1 (The reference shot). Take 18 g in, 36 g out, 9 bar and a uniform bed at the reference grind. The one calibrated constant, $k_{\text{ref}} = 2.9416 \times 10^{-16} \text{ m}^2$, is fixed by root-finding so that the uniform puck pours the recipe of Waszkiewicz et al. [3], 37 g from 18.50 g in 32 s at 9 bar. With it the reference shot takes 30.82 s and extracts $\text{EY} = 22.42$ percent, inside the 20–23 percent range that Cameron et al. [1] report for espresso, although nothing was tuned to extraction.

The calibration recipe itself extracts 22.66 percent. The two-dimensional solver, run on the same uniform puck, pours in 30.88 s, 0.20 percent longer than the one-dimensional



solver, and extracts 22.44 percent. Over the shot the mean flux is 6.84 times the initial flux, and the beverage is 5.81 pore volumes.

The calibrated permeability corresponds, through the Kozeny–Carman law $k = d_h^2 \phi_0^3 / (180(1 - \phi_0)^2)$, to a grain diameter $d_h = 2.65 \mu\text{m}$, far below the size of real espresso grounds. The constant also absorbs what the model leaves out of the calibration shot, such as the wetting delay and the compaction that [3] measure, so only relative changes, the factor $(d/d_{\text{ref}})^2$ and the porosity factor of Definition 2, are used.

3. The Growth Law

Start. Before any channel grows, the puck is uniform across and its fields depend on depth and time only. This base state is what a disturbance grows on, and it is not steady.

Steps. With the flux uniform in depth (the end of the flow derivation), the equations of the model reduce to

$$\partial_t(\phi c) + q(t) \partial_z c = \partial_z(\phi D \partial_z c) + r, \quad \partial_t s = -r, \quad q(t) = \frac{\Delta P}{\mu \int_0^H dz/k(\phi)}, \quad (3.1)$$

with $c = 0$ entering at $z = 0$ and the bed full of fresh water at $t = 0$. The grains near the top are washed by water that is never saturated and are emptied; the grains below receive water close to saturation and give up little. Between them lies an extraction front: above it the bed is exhausted, $s \approx 0$, and its pore water is fresh, $c \approx 0$; below it the bed holds solubles s_d at porosity ϕ_d and its water is close to saturation, $c \approx c_{\text{sat}}$. Behind the front the bed has opened, so its permeability $k_1 = k(\phi)$ exceeds the permeability k_2 below the front; at the contrast of a fully extracted bed the ratio is 29.5.

End. The front is thin when fresh water saturates over a length short against the puck, a large Damköhler number, as at $d = 0.25 d_{\text{ref}}$; at practical grinds it is as thick as the puck (Remark 1). Either way it moves at a speed fixed by a mass balance (Proposition 1), and as it moves the opened region behind it lengthens and the flux at fixed pressure rises,

$$\frac{1}{q} = \frac{\mu}{\Delta P} \left(\frac{\zeta}{k_1} + \frac{H - \zeta}{k_2} \right) \quad (3.2)$$

for a sharp front at depth ζ . A region whose front runs ahead of its neighbors' has less resistance and draws more flow, which is the seed of the instability.

Proposition 1 (The speed of the extraction front). Under Assumption 1, let a front at depth $\zeta(t)$ separate an upstream state with no solubles and fresh water, $s = 0$ and $c = 0$, from a downstream state with solubles s_d , porosity ϕ_d and saturated water, $c = c_{\text{sat}}$, under a Darcy flux q that is uniform in depth. Suppose the diffusive flux is negligible on either side of the front ($\text{Pe} \gg 1$) and the front's profile translates without change of shape over the time it takes to cross its own width. Then it moves down the puck at

$$v_f = \gamma_f q, \quad \gamma_f = \frac{c_{\text{sat}}}{s_d + \phi_d c_{\text{sat}}}.$$



Proof. Write $e = \phi c + s$ for the total solute per unit bed volume and $F = qc - \phi D \partial_z c$ for its flux. Adding the two equations of the base state gives $\partial_t e + \partial_z F = 0$. Take a window $[a(t), b(t)]$ that contains the front and moves with it at v_f , with a in the upstream state and b in the downstream state. Then

$$\frac{d}{dt} \int_a^b e \, dz = \int_a^b \partial_t e \, dz + v_f (e(b) - e(a)) = F(a) - F(b) + v_f (e(b) - e(a)).$$

A profile that translates without change of shape keeps $\int_a^b e \, dz$ fixed, so the left side vanishes. Upstream $e(a) = 0$ and $F(a) = 0$; downstream $e(b) = s_d + \phi_d c_{\text{sat}}$ and, with the diffusive flux negligible, $F(b) = q c_{\text{sat}}$. Hence $v_f (s_d + \phi_d c_{\text{sat}}) = q c_{\text{sat}}$. $\square$

The kinetics set the width of the front, not its speed. With fresh coffee ahead, $s_d = s_0$ and $\phi_d = \phi_0$, it is $\gamma_f = 1.14$, so the front moves at about a fifth of the speed q/ϕ_0 of the pore water. A front whose downstream bed has already given up part of its solubles, as after the first pore volume, moves faster.

Method. Proposition 1 is checked where its hypotheses hold, on a thin front. The uniform puck at $d = 0.25 d_{\text{ref}}$ is solved by the independent one-dimensional solver (300 cells, dense output). The front position $\zeta(t)$ is where the solubles fall to half their downstream value, $s = s_d/2$, interpolated between cells; the measured speed v_f is its centered difference over $\pm 0.005 T$. The formula $\gamma_f q$ takes s_d and ϕ_d from the bottom cell and q from the series resistance, all at the same instant.

Parameters. $d = 0.25 d_{\text{ref}}$, a shot of $T = 320.6$ s; the instant $t = 112.2$ s, which is $0.35 T$, with the front at $\zeta = 3.44$ mm and a 10–90 percent width of 0.78 mm.

Result. The measured front speed is $v_f = 3.4450 \times 10^{-5}$ m/s, and the formula gives $\gamma_f q = 3.4456 \times 10^{-5}$ m/s, so the front moves 0.0345 mm per second, crossing its own width in 22.7 s.

Error. The measured speed is below the formula by a relative 1.7×10^{-4} , within the 10^{-2} tolerance set for this check. A second, indirect check sits inside the growth test below: Proposition 2, which moves every point of the front at γ_f times its local flux, reproduces the growth measured in the direct simulation within 0.11 percent.

Proposition 2 (The growth rate of a thin front at fixed pressure). Under Assumption 1, let a thin front at depth ζ separate bed of permeability k_1 above from bed of permeability k_2 below, with a uniform Darcy flux q , the pressure held at $z = 0$ and $z = H$, and periodic sides, and let each point of the front move down at γ_f times the local normal flux, as in Proposition 1. Displace the front to $z = \zeta + \eta(x, t)$, $\eta = \hat{\eta} e^{\sigma t} \cos \kappa x$, small against ζ , $H - \zeta$ and $1/\kappa$, with the front thin against $1/\kappa$. Then

$$\sigma(\kappa) = \gamma_f q \kappa \frac{k_1 - k_2}{k_2 \tanh(\kappa \zeta) + k_1 \tanh(\kappa(H - \zeta))}.$$

Proof. On each side of the front the permeability is uniform, so $\nabla^2 p = 0$, and the base pressure is linear, $\partial_z \bar{p} = -\mu q/k_1$ above and $-\mu q/k_2$ below. Write $p = \bar{p} + \hat{p}(z) \cos \kappa x$. The



faces hold the pressure, so $\hat{p}(0) = \hat{p}(H) = 0$, and $\partial_{zz}\hat{p} = \kappa^2\hat{p}$ gives

$$\hat{p} = A \sinh(\kappa z) \quad (z < \zeta), \quad \hat{p} = B \sinh(\kappa(H - z)) \quad (z > \zeta).$$

Pressure and normal flux are continuous across the displaced front. To first order in η , with $\bar{p}$ continuous at ζ ,

$$-\frac{\mu q}{k_1} \eta + A \sinh(\kappa \zeta) = -\frac{\mu q}{k_2} \eta + B \sinh(\kappa(H - \zeta)), \quad -k_1 A \cosh(\kappa \zeta) = k_2 B \cosh(\kappa(H - \zeta)),$$

the second being continuity of the perturbation of $w = -(k/\mu) \partial_z p$; the tilt of the front multiplies the base horizontal flux, which is zero. Eliminating B ,

$$A \frac{k_2 \sinh(\kappa \zeta) + k_1 \cosh(\kappa \zeta) \tanh(\kappa(H - \zeta))}{k_2} = \mu q \eta \frac{k_2 - k_1}{k_1 k_2},$$

so the flux perturbation at the front is

$$\hat{w} = -\frac{k_1}{\mu} \kappa A \cosh(\kappa \zeta) = q \kappa \eta \frac{k_1 - k_2}{k_2 \tanh(\kappa \zeta) + k_1 \tanh(\kappa(H - \zeta))}.$$

Each point of the front moves at $\gamma_f(q + \hat{w})$, the base part moving ζ , so $\partial_t \eta = \gamma_f \hat{w}$, which is $\sigma \eta$ with σ as stated. $\square$

This is the reactive-infiltration mechanism [7], [8] in a finite puck at fixed pressure: where the front runs ahead, low-permeability bed has become high-permeability bed, the local resistance has fallen, and more water arrives. We checked the formula independently by solving the linearized Darcy problem for a smoothed front, the logarithm of the permeability rising over a width of $H/20000$ on 400 000 cells. At $k_1/k_2 = 29.5$ and the front depth of the test case below, the two agree to a relative 5.8×10^{-5} at $\kappa H = 0.5$, 1.7×10^{-4} at 2, 5.6×10^{-4} at 6 and 1.5×10^{-3} at 20, the error growing with κ as the smoothing predicts.

Corollary 1 (The column limit and short waves). Under the hypotheses of Proposition 2: (i) as $\kappa \rightarrow 0$,

$$\sigma \rightarrow \sigma_{\text{col}} = \gamma_f q \frac{k_1 - k_2}{k_2 \zeta + k_1(H - \zeta)},$$

the growth rate of two independent columns held at the same pressure; (ii) when $\kappa \zeta$ and $\kappa(H - \zeta)$ are both large,

$$\sigma = \gamma_f q \kappa \frac{k_1 - k_2}{k_1 + k_2} (1 + o(1)),$$

so the growth rate rises in proportion to the wavenumber. The law holds only while $\kappa w \ll 1$, with w the width of the front, so short waves stop growing where the front's own thickness, not the formula, sets the scale.

Proof. $\tanh y = y + O(y^3)$ as $y \rightarrow 0$, so the denominator of Proposition 2 is $\kappa(k_2 \zeta + k_1(H - \zeta)) + O(\kappa^3)$. Independently, a column held at pressure ΔP carries $q = \Delta P / (\mu(\zeta/k_1 + (H - \zeta)/k_2))$; moving its front by η changes the flux by $\eta \partial_\zeta q = \eta q (k_1 - k_2) / (k_2 \zeta + k_1(H - \zeta))$, and the front moves at γ_f times the flux, which is (i) again. (ii) $\tanh y = 1 - O(e^{-2y})$ as



$y \rightarrow \infty$. The restriction to $\kappa w \ll 1$ is the hypothesis of Proposition 2 that the front be thin against $1/\kappa$. $\square$

At $k_1/k_2 = 29.5$, the contrast of a fully extracted bed, the short-wave coefficient $(k_1 - k_2)/(k_1 + k_2)$ is 0.934. In `derive.py` the column limit is checked both ways, to a relative 2×10^{-15} and 1.8×10^{-11} .

The cut-off is measured, not derived. At $d = 0.25 d_{\text{ref}}$ and $t = 112.2$ s the moving-frame eigenvalue peaks at 2.26 mm, with $\sigma = 0.0315 \text{ s}^{-1}$ and $\kappa w = 2.18$, a wave 2.89 front widths long, and does not grow below 0.92 mm. Molecular diffusion cannot be what stops it: the diffusion length $\phi_0 D/U = H/\text{Pe}$ is at most 85 μm there, even on the initial flux, the smallest of the shot, and 2.05 μm or less at practical grinds on the mean flux. The front thickness sets the cut-off.

Remark 2 (Why a frozen analysis in the fixed frame overstates growth). Freeze the base state at an instant, linearize the equations about it in fixed cells, and take the leading eigenvalue as the growth rate. Integrated over the window of the test case, 64.2 to 176.9 s at $d = 0.25 d_{\text{ref}}$ and $\lambda = 2H$, this gives a growth of 2.535 against 0.7202 in the direct simulation: 3.52 times too much. A frozen analysis assumes that the base changes slowly compared with the disturbance. In fixed cells it does not. The front sweeps through a cell in the time it takes to cross its own width, 22.7 s, while the disturbance takes 156 s to grow by a factor e . Each cell's base state therefore changes completely while the disturbance barely moves, and the frozen operator mistakes the passage of the front through the cell for growth. In a frame that moves with the front, the base is nearly steady, the frozen assumption holds, and the leading eigenvalue agrees with Proposition 2 within 5.0 percent.

The overstatement is not confined to one wavelength. At $t = 112.2$ s the fixed-frame eigenvalue peaks at $\sigma = 0.0513 \text{ s}^{-1}$ at 2.82 mm, against 0.0315 s^{-1} at 2.26 mm in the moving frame, and at the longest wavelength computed, 200 mm, it is 4.25 times the moving-frame value. The fixed-frame eigenvalues were the method declared before any computation and were replaced, with no parameter changed, when they failed the check against the direct simulation.

Start. At practical grinds the front is as thick as the puck and the base changes throughout the shot, so neither Proposition 2 nor a frozen eigenvalue gives the growth over a shot. We linearize about the time-dependent base itself and integrate through the shot, freezing nothing and choosing no frame.

Steps. Write every field as its base value plus a transverse mode, $f = \bar{f}(z, t) + \hat{f}(z, t) \cos \kappa x$, and let the disturbance be a streak of initial porosity $\hat{\phi}_i$, uniform in depth, as a packing defect would be. Then $\hat{s}_i = -E Y_{\text{max}} \rho_s \hat{\phi}_i$, the porosity perturbation is $\hat{\phi} = (1 - E Y_{\text{max}}) \hat{\phi}_i - \hat{s}/\rho_s$, and with $m = \phi c$ the concentration perturbation is $\hat{c} = (\hat{m} - \bar{c} \hat{\phi})/\bar{\phi}$. The permeability perturbation is $\bar{k} \bar{\beta} \hat{\phi}$, with $\bar{\beta} = 3/\bar{\phi} + 2/(1 - \bar{\phi})$ the logarithmic derivative of the Kozeny–Carman law. The pressure perturbation solves

$$\partial_z (\bar{k} \partial_z \hat{p}) - \kappa^2 \bar{k} \hat{p} = -\partial_z (\bar{k} \bar{\beta} \hat{\phi} \partial_z \bar{p}), \quad \hat{p}(0) = \hat{p}(H) = 0, \quad (3.3)$$

and gives the flux perturbations $\hat{w} = -(\bar{k} \partial_z \hat{p} + \bar{k} \bar{\beta} \hat{\phi} \partial_z \bar{p})/\mu$ downward and $\kappa \bar{k} \hat{p}/\mu$ across.



Solute balance and dissolution become

$$\partial_t \hat{m} + \partial_z (q\hat{c} + \hat{w}\bar{c}) + \frac{\kappa^2 \bar{k} \bar{c}}{\mu} \hat{p} - \partial_z (D\bar{\phi} \partial_z \hat{c} + D\hat{\phi} \partial_z \bar{c}) + \kappa^2 D\bar{\phi} \hat{c} = \hat{r}, \quad \partial_t \hat{s} = -\hat{r}, \quad (3.4)$$

with $\hat{r}$ the linearization of the rate of Definition 3 in ϕ , c , s and s_i . Discretized in depth like the one-dimensional solver (150 cells, first-order upwind, harmonic face permeability), the state $v = (\hat{m}, \hat{s}, \hat{\phi}_i)$ obeys $\partial_t v = B(t) v$, with $B(t)$ built on the solver's base at time t . We integrate it by exponential midpoint steps,

$$v_{j+1} = \exp(B(t_{j+1/2}) \Delta t) v_j, \quad (3.5)$$

with 100 steps per shot, and record the flux unevenness at the outlet, $\hat{w}(H, t)/q(t)$.

End. The amplification of a wavelength $\lambda = 2\pi/\kappa$ over the shot is

$$A = \frac{|\hat{w}(H, T)/q(T)|}{|\hat{w}(H, 0)/q(0)|}, \quad (3.6)$$

above one, the puck ends the shot more uneven than it began. The problem is solved for the column limit $\kappa = 0$ and 15 wavelengths from 58 to 0.45 mm at each of the 13 grinds; the time derivative of $\log |\hat{w}(H, t)/q(t)|$ is the instantaneous growth rate.

Method. Proposition 2 is tested where a front is thin, three independent ways, on the growth of one transverse mode over the same window. (a) A direct simulation: the two-dimensional solver of Algorithm 1 on a strip one wavelength wide, started from the uniform shot's state at $0.15T$ with every column shifted down by $\eta_0 \cos(2\pi x/\lambda)$, $\eta_0 = 1 \mu\text{m}$; the growth is that of the first Fourier mode of the extracted depth of each column. (b) Proposition 2, evaluated at every recorded instant with the front depth, the two permeabilities, the flux and the downstream state, taken once from the direct simulation's own column-mean base and once from the one-dimensional solver's base, and integrated in time. (c) The leading eigenvalue of the linearized operator in the frame moving with the front, at 11 instants across the window, integrated by the trapezoid rule.

Parameters. $d = 0.25 d_{\text{ref}}$, a shot of 320.6 s; $\lambda = 2H = 27.24 \text{ mm}$; the window 64.2 to 176.9 s, $0.20T$ to $0.55T$, over which $\kappa w = 0.185$ on average; grid 0.2 mm.

Result. The integrated growth, the logarithm of the amplitude ratio across the window, is 0.7202 in the direct simulation, 0.7194 from the formula on the simulation's base and 0.7229 on the one-dimensional base, and 0.7609 from the moving-frame eigenvalue.

Error. The formula differs from the direct simulation by a relative 0.11 percent and from the moving-frame eigenvalue by 5.0 percent, and the eigenvalue from the simulation by 5.4 percent, against tolerances of 15, 10 and 15 percent declared beforehand; the two bases change the formula by 0.49 percent. At $t = 112.2 \text{ s}$, formula and eigenvalue agree within 4.9 percent at $\lambda = 2H$ and within 27 percent wherever $\kappa w \leq 1$, that is $\lambda \geq 5.04 \text{ mm}$, the gap opening as the thin-front hypothesis fails.

Method. The exact linear problem along the moving base, for a depth-uniform streak of



initial porosity, at 16 wavelengths and 13 grinds. At each grind λ^* is the wavelength of largest amplification A over the shot, and the threshold d_c is the grind where $\max_{\lambda} A = 1$, by linear interpolation of $\log \max_{\lambda} A$ against $\log d$ between the two grid grinds that bracket it.

Parameters. The main recipe; wavelengths from 58 to 0.45 mm in steps of about $\sqrt{2}$ and the column limit; 150 cells, 100 exponential steps per shot. For the check, a streak of amplitude 10^{-5} on a strip one wavelength wide, run through the whole shot by Algorithm 1.

Result. The fastest wavelength is $\lambda^* = 3.6$ mm at every one of the 13 grinds, about a quarter (0.264) of the puck height, bounded by its grid neighbors 2.6 and 5.1 mm; waves of 0.64 mm and shorter decay at every grind. The amplification at λ^* falls from 4.62 at $d = 0.5 d_{\text{ref}}$ through 1.38 at the reference grind to 0.91 at $2 d_{\text{ref}}$ and crosses one at $d_c = 1.58 d_{\text{ref}}$, where the mean-flux Damköhler number is 0.894; coarser than d_c every wavelength computed leaves the puck less uneven than it entered. Long waves barely grow: at most 1.27 at 58 mm and 1.09 in the column limit, both at $d = 0.63 d_{\text{ref}}$. Through the reference shot the largest instantaneous growth rate is -0.129 s^{-1} at the start, while the fresh water filling the bed is pushed out; it turns positive by $0.2 T$, peaks at 0.0823 s^{-1} at $0.5 T$, and is negative again from $0.8 T$ as the solubles run out. Meanwhile the fastest wavelength shrinks from the column limit through 29 mm at $0.4 T$ and 5.1 mm at $0.5 T$ to 2.6 mm at 0.6 and $0.7 T$. The unevenness peaks mid-shot, at 12.5 times its initial value at $d = 0.5 d_{\text{ref}}$ and 1.66 times at the reference grind (Table 2).

Error. Against the direct simulation at the reference grind, the logarithmic growth from $0.1 T$ to T is 0.4743 against 0.4623 at $\lambda = 3.6$ mm, a relative 2.6 percent, and 0.1675 against 0.1610 at 58 mm, 4.1 percent; the whole curves agree within 2.3 and 0.7 percent. At $d = 0.5 d_{\text{ref}}$ on a 0.1 mm grid the whole curves agree within 14.5 and 4.0 percent, inside the declared 20 percent. The end point of the long wave at that grind, 0.260 in the simulation against 0.191, is not claimed: the last percent of the shot is a steep swing, and the two codes' shot times differ by 0.7 percent.

Source. The exact linear problem along the moving base for a depth-uniform streak of initial porosity, at seven of the 13 grinds (`model.py` stage `ivp`, `results.npz`), printed by `derive.py`.

Columns. Grind; mean-flux Damköhler number; amplification A of the outlet flux unevenness over the shot at $\lambda^* = 3.6$ mm; the largest unevenness reached during the shot at λ^* , relative to its start; A at 58 mm; A in the column limit; and the end-of-shot contrast in extraction yield between a streak and its neighbors, per unit amplitude of the initial porosity streak, at λ^* .



Table 2: Amplification of flow unevenness over a shot, by grind, from the exact linear problem; above one the puck ends the shot more uneven than it began.

d/d_{ref}	Da	$A(\lambda^*)$	peak	A , 58 mm	A , columns	contrast
0.50	11.4	4.62	12.54	1.17	0.94	2.01
0.63	6.15	2.64	4.41	1.27	1.09	1.48
0.71	4.64	2.14	3.14	1.22	1.09	1.31
1.00	2.17	1.38	1.66	1.02	0.96	0.89
1.41	1.10	1.06	1.18	0.87	0.84	0.56
1.59	0.889	1.00	1.10	0.84	0.81	0.47
2.00	0.581	0.91	1.00	0.80	0.78	0.32

The amplification at λ^* , its mid-shot peak and the extraction contrast fall steadily as the grind coarsens, while at 58 mm and in the column limit the amplification never exceeds 1.27 and 1.09. Finer coffee wins twice: a higher Damköhler number and a longer shot. For a streak of amplitude 0.02, the last column gives, in the linear theory, an end-of-shot contrast of 4.0 points of extraction yield at $d = 0.5 d_{\text{ref}}$, 1.8 at the reference grind and 0.6 at $d = 2 d_{\text{ref}}$.

Script. `fig_growth.py`: left, the growth law at $d = 0.25 d_{\text{ref}}$ and $t = 112.2$ s on 30 wavelengths from 0.3 to 200 mm, the closed form drawn only where $\kappa w < 1$; right, the amplification over a whole shot from the exact linear problem at five grinds, with the column limit drawn as squares. Both panels read `results.npz`.

Shows. The thin-front law against the moving-frame eigenvalue and the fixed-frame eigenvalue, and the selection of a channel spacing over a real shot.




Figure 1: Left: growth rate against channel wavelength at one instant of a fine-grind shot; the closed-form thin-front law, the eigenvalue in the frame of the moving front, and the fixed-frame eigenvalue that overstates growth. Right: amplification of outlet flow unevenness over a whole shot against wavelength for five grinds, with the column limit as squares.

On the left the closed form follows the moving-frame eigenvalue at long waves and leaves it as κw nears one; the eigenvalue peaks at 2.26 mm and vanishes below 0.92 mm, while the fixed-frame curve lies above it wherever it grows, its peak 1.63 times higher. On the right every grind peaks at 3.6 mm, with amplifications from 4.62 at $d = 0.5 d_{\text{ref}}$ to 0.91 at $2 d_{\text{ref}}$; the column limits lie between 0.78 and 1.09.

4. Channels in a Full Puck

Lemma 1 (The finite-volume scheme conserves solubles exactly). Let Algorithm 1 hold cell averages of $m = \phi c$ and s on an $N_z \times N_x$ grid of cells $\Delta x \times \Delta z$, periodic in x ; let each stage of its Heun step change a cell by minus the divergence of its face fluxes plus r for m , and by $-r$ for s , with face fluxes that vanish on the top face; and let the cup content C , per unit basket area, gain at each stage the mean flux through the bottom face, with the same Heun weights. Then for every step Δt , every set of face fluxes and every rate, in exact arithmetic

$$\frac{\Delta x \Delta z}{W} \sum_{\text{cells}} (m + s)^{n+1} + C^{n+1} = \frac{\Delta x \Delta z}{W} \sum_{\text{cells}} (m + s)^n + C^n.$$

Proof. The rate enters m with $+r$ and s with $-r$ in the same cell, so it cancels from $m + s$ cell by cell. The divergence of cell (i, j) is $(F_{i,j+1/2} - F_{i,j-1/2})/\Delta x + (G_{i+1/2,j} - G_{i-1/2,j})/\Delta z$, with F and G the horizontal and vertical face fluxes. Summed over a row, the horizontal differences telescope to zero, because the faces are periodic. Summed over a column, the vertical differences telescope to the flux through the bottom face minus the flux through the top face, which is zero: the entering water carries $c = 0$, and there is no diffusive



flux through the top. So one stage changes the weighted sum on the left by $-\Delta t$ times $(\Delta x/W) \sum_j G_{N_z+1/2,j}$, the mean flux through the bottom face, and adds exactly that to C . The Heun step is the average of two states: the starting state, and the state after two such stages taken in succession, whose cup has gained both stage fluxes. Both carry the same total of grains, pores and cup, and an average with weights summing to one preserves a linear quantity, so the identity holds for the step. Nothing in the argument uses the size of Δt , the velocity, or the form of r . $\square$

What is conserved is the solute in grains, pores and cup together. Freezing the velocity within a step and shortening the last step to land on V_{out} leave the identity untouched; positivity of c and s , and the accuracy of the fluxes, are not given by the lemma. Round-off enters only through the sums over the 290×68 cells and the accumulation over the steps, and the checks after Algorithm 1 find the balance error below 3.3×10^{-15} of the initial solubles in every shot.

Lemma 2 (A uniform puck stays uniform). In Algorithm 1, suppose that at step n the cell values of ϕ_i , m and s do not depend on the column index, so every column of the grid holds the same numbers, and that the pressure drop is the same for every column. Then, in exact arithmetic, the same holds at step $n + 1$: a puck that is uniform across a periodic strip stays uniform across it.

Proof. Each cell's permeability depends only on its own porosity, so the cell, face and boundary permeabilities are the same in every column. Let S shift the unknowns by one column, cyclically. The pressure system $Ap = b$ is then invariant, $SA = AS$ and $Sb = b$. The matrix A is symmetric and irreducibly diagonally dominant, strictly so in the rows that touch the faces held at fixed pressure, so it is nonsingular. Since Sp solves the same system, $Sp = p$: the pressure is the same in every column. The horizontal face fluxes, proportional to differences of pressure between neighboring columns, vanish; so do the concentration differences across vertical faces, hence the van Leer slopes across, and the horizontal diffusive fluxes. What remains in each column is the same one-dimensional update applied to the same numbers, with one time step, a single number, shared by all. $\square$

In floating point every step keeps the columns bitwise identical except the sparse factorization of the pressure system, whose fill-reducing ordering treats the columns differently; it leaves differences at the rounding level, which the instability then amplifies. The measured departure from uniformity at the end of a shot is 9.3×10^{-15} in porosity over all 13 grinds, so any unevenness in the results below comes from the initial porosity field, not from the solver.

Algorithm 1 — A two-dimensional shot by the method of lines

```

1 Input:  grind d/d_ref; initial porosity phi_i on an Nz x Nx grid of
         spacing h,
2         periodic across; pressure program Delta P(t, V/V_out); beverage
3         volume V_out per unit basket area; CFL number 0.4
4 Output: EY, the shot time T, and the fields at chosen fractions of V_out

```



```

5
6 1. s_i = EYmax * rho_s * (1 - phi_i); s = s_i; m = phi*c = 0 (fresh
   water);
7   t = 0; V = 0; cup = 0.
8 2. In every cell: phi = phi_i + (s_i - s)/rho_s; c = m/phi;
9   k = k_ref (d/d_ref)^2 (phi/phi0)^3 ((1 - phi0)/(1 - phi))^2.
10 3. Pressure: solve the 5-point finite-volume problem div(k grad P) = 0
    with
11   harmonic face permeabilities, P = 1 on the top face, P = 0 on the
    bottom
12   face, periodic across, by sparse LU; scale the face Darcy fluxes by
13   Delta P(t, V/V_out)/mu; q = mean flux through the bottom face.
14 4. Time step: dt = 0.4 * min over cells of phi / (outgoing face flux per
15   unit cell volume), and dt <= 0.02/K0; if V + q*dt >= V_out, shorten dt
16   so that V lands exactly on V_out.
17 5. With the face fluxes frozen, take one Heun (SSP-RK2) step of
18   dm/dt = -div(F) + r, ds/dt = -r,
19   where F is upwind MUSCL advection with the van Leer limiter plus
    central
20   diffusion, c = 0 entering at the top, no diffusive flux at top and
21   bottom; add dt * (mean of the two stage fluxes through the bottom face)
22   to cup.
23 6. V = V + q*dt; t = t + dt; if V < V_out, go to 2.
24 7. EY = cup * A_basket / dose; T = t.

```

Cost. One sparse factorization of the pressure system per step, $N_x N_z = 19720$ unknowns on the 0.2 mm production grid of 290×68 cells. The 976 shots of the coherence sweep, including its checks on the 0.1 mm grid, took 52.3 CPU-hours, 193 s per shot on average, on 32 workers.

Every flux of step 5 leaves one cell and enters its neighbor or the cup, which is Lemma 1; step 3 is the only place where the columns of a uniform puck can differ, which is Lemma 2's remark on round-off. The last step lands on V_{out} exactly, so the shot ends by volume, as Definition 1 requires.

Method. Every two-dimensional shot reports at its end the relative solute balance error, the initial solubles against the solute in grains, pores and cup; every uniform shot reports the largest variation of porosity across the puck. A uniform shot at the reference grind compares, at every step, its flux with the series-resistance flux of the flow derivation, and its shot time with the independent one-dimensional solver. The grid study repeats shots on $h = 0.4, 0.2$ and 0.1 mm and reports the observed order $p = \log_2(|EY_{0.4} - EY_{0.2}|/|EY_{0.2} - EY_{0.1}|)$.

Parameters. Production grid $h = 0.2$ mm; tolerances declared before any run: 10^{-6} for the balance, 10^{-10} for uniformity, 10^{-9} for the flux identity, 1 percent for the shot time and for the change of extraction from 0.2 to 0.1 mm. The grid study uses $d = 0.63 d_{\text{ref}}$ and d_{ref} , porosity noise $\varepsilon = 0.02$ with 1 mm blobs, seeds 0 to 2; the coherence sweep adds 0.1 mm checks on depth-uniform streaks 1 and 3.6 mm wide, five seeds.

Result. The largest balance error is 2.6×10^{-15} over the shots of the grind sweep and



the grid study, 2.2×10^{-15} over the lever and 2020-recipe shots, 3.3×10^{-15} over the 976 shots of the coherence sweep, and 8.7×10^{-16} in the one-dimensional solver. The largest variation of porosity across a uniform puck at the end of a shot is 9.3×10^{-15} over the 13 grinds. The two-dimensional flux matches the series-resistance flux to 2.9×10^{-13} at every step, and the two solvers' shot times differ by 0.20 percent. Table 3 gives the grid study: from 0.2 to 0.1 mm the mean extraction changes by 0.099 percent at $d = 0.63 d_{\text{ref}}$ and 0.027 percent at the reference grind, with observed orders 1.62 and 1.43; the coherence checks change by 0.140 percent or less.

Error. Every value is within its declared tolerance. The balance and uniformity errors sit at the rounding level, as Lemmas 1 and 2 predict for a scheme that is exact in exact arithmetic. The production grid meets the declared rule, a change of less than 1 percent from 0.2 to 0.1 mm, by a factor of ten or more, and the verdicts of the coherence sweep are the same on both grids.

Source. Algorithm 1 on three grids: the first two rows are means of seeds 0 to 2 (`model.py` stage `conv`), the last four means of seeds 0 to 4 from the coherence sweep (`coherence.npz`), printed by `derive.py`.

Columns. Case; mean cup extraction yield, in percent, on grids of 0.4, 0.2 and 0.1 mm; relative change from one grid to the next, in percent; observed order where all three grids were run.

Table 3: Grid convergence of the mean cup extraction yield, with observed orders, on the grind sweep and at the onset of the coherence sweep.

case	0.4 mm	0.2 mm	0.1 mm	$\Delta_{0.4 \rightarrow 0.2}$ (%)	$\Delta_{0.2 \rightarrow 0.1}$ (%)	order
$d = 0.63 d_{\text{ref}}$, $\varepsilon = 0.02$, 1 mm blobs	26.78	26.70	26.67	0.305	0.099	1.62
$d = d_{\text{ref}}$, $\varepsilon = 0.02$, 1 mm blobs	22.29	22.27	22.27	0.074	0.027	1.43
z-uniform streaks, $\ell_x = 1$ mm, $d = 0.50 d_{\text{ref}}$	—	24.24	24.21	—	0.140	—
z-uniform streaks, $\ell_x = 1$ mm, $d = 0.71 d_{\text{ref}}$	—	23.69	23.68	—	0.059	—
z-uniform streaks, $\ell_x = 3.6$ mm, $d = 0.50 d_{\text{ref}}$	—	23.19	23.18	—	0.053	—
z-uniform streaks, $\ell_x = 3.6$ mm, $d = 0.71 d_{\text{ref}}$	—	23.39	23.39	—	0.017	—



The changes shrink by 3.1 at $d = 0.63 d_{\text{ref}}$ and 2.7 at the reference grind per refinement, an observed order between one and two, consistent with a second-order scheme whose limiter lowers the order at steep gradients. The 0.2 mm grid is within 0.14 percent of the 0.1 mm grid in every case.

Script. `fig_puck.py`: the porosity of the full section, 58 mm across and 13.62 mm deep, at one third, two thirds and the whole of the beverage, for seed 0 of the porosity noise $\varepsilon = 0.02$ with 1 mm blobs, at $d = 0.5 d_{\text{ref}}$ (left) and the reference grind (right), on one color scale. It reads the snapshots saved in `results.npz`.

Shows. Where the bed opens first, and how unevenly, at the grind where channeling costs most and at the reference grind.




Figure 2: Porosity in a vertical section of the puck at one third, two thirds and the end of the shot, for a fine grind on the left and the reference grind on the right, same initial porosity noise; at the fine grind the opened bed advances in tongues and narrow columns of little-extracted coffee survive to the end, while the reference puck opens nearly evenly.

At $d = 0.5 d_{\text{ref}}$ (snapshots at 48, 68 and 80 s) the opened bed advances from the shower screen in tongues, and narrow vertical columns of little-extracted coffee are bypassed and survive to the end; this seed extracts 27.44 percent. At the reference grind (18, 26 and 31 s) the extraction zone is as deep as the puck from the start and the bed opens almost evenly; the seed extracts 22.27 percent. Porosity runs from 0.178 to 0.465.

Method. Algorithm 1 at the 13 grinds for a uniform puck and for Gaussian porosity noise of standard deviation $\varepsilon = 0.01, 0.02$ and 0.04 with 1 mm correlation length, five seeds each: 208 shots. The noise is made on the 0.1 mm grid, block-averaged, clipped to $[0.05, 0.40]$ and shifted to mean ϕ_0 , so the dose is exact. The paradox rule was declared before the first shot: the seed-mean extraction must peak at an interior grind, and its drop from the peak to the finest grind must exceed twice the combined standard error of the two means. The rule is also applied to each seed alone.

Parameters. The main recipe; $h = 0.2$ mm; seeds 0 to 4 at every grind, reused by the



levers and the 2020 recipe.

Result. Table 4. No amplitude makes a paradox. At $\varepsilon = 0.01$ and 0.02 the extraction rises to the finest grind in every seed alone. At $\varepsilon = 0.02$ the finest grind extracts 27.51 ± 0.25 percent against 29.07 for a uniform puck, a loss of 1.55 points or 5.35 percent, against 0.65 percent at the reference grind and 0.06 percent at $2d_{\text{ref}}$. Doubling the noise raises these losses to 11.75 , 2.91 and 0.40 percent and flattens the fine end: $d = 0.5d_{\text{ref}}$ extracts only 0.10 points more than $0.56d_{\text{ref}}$, and one seed of five peaks at 0.56 . The flattening is channeling: at $\varepsilon = 0.02$ the spread of local extraction yield grows from 2.50 points at the reference grind to 3.24 at $0.5d_{\text{ref}}$, against 0.90 for a uniform puck there, and the tenth of the outlet that carries most water carries 14.3 and 16.2 percent of it, against 10 percent for an even flow.

Error. The seed-to-seed standard deviation at the finest grind is 0.11 , 0.25 and 0.57 points at the three amplitudes, so the loss at $\varepsilon = 0.02$ is six standard deviations of the scatter; the grid checks put the discretization error at 0.14 percent of the yield or less. The losses follow the linear theory: they grow with fineness as the amplification at λ^* does, and coarser than $d_c = 1.58d_{\text{ref}}$ they stay within 0.20 percent at $\varepsilon = 0.02$.

Source. Algorithm 1, a uniform puck and five seeds at each noise amplitude, $h = 0.2$ mm, the main recipe (`model.py` stage `f3`, `numbers.json`), at seven of the 13 grinds, printed by `derive.py`.

Columns. Grind; shot time of the uniform puck; cup extraction yield, in percent, of the uniform puck and, as mean $\pm$ standard deviation over five seeds, with porosity noise of standard deviation ε in 1 mm blobs; and the loss at $\varepsilon = 0.02$ relative to the uniform puck, in percent.

Table 4: Mean cup extraction against grind and heterogeneity amplitude for millimeter-scale porosity noise, five seeds: the curve rises to the finest grind at every amplitude.

d/d_{ref}	shot, uniform (s)	uniform puck	$\varepsilon = 0.01$	$\varepsilon = 0.02$	$\varepsilon = 0.04$	loss at $\varepsilon =$ 0.02 (%)
0.50	81.5	29.07	28.35 ± 0.11	27.51 ± 0.25	25.65 ± 0.57	5.35
0.56	66.5	28.41	28.01 ± 0.07	27.29 ± 0.18	25.55 ± 0.51	3.95
0.63	55.3	27.54	27.32 ± 0.04	26.79 ± 0.14	25.28 ± 0.44	2.72
0.71	46.8	26.48	26.35 ± 0.02	26.00 ± 0.10	24.78 ± 0.36	1.83
1.00	30.9	22.44	22.41 ± 0.01	22.29 ± 0.04	21.79 ± 0.17	0.65
1.41	22.2	17.73	17.72 ± 0.01	17.68 ± 0.03	17.49 ± 0.09	0.28
2.00	16.6	12.93	12.93 ± 0.01	12.93 ± 0.02	12.88 ± 0.06	0.06



The noise subtracts more at finer grinds, so the top of each curve flattens, but none turns over.

Script. `fig_paradox.py`: left, the mean cup extraction yield against grind for the uniform puck and the three noise amplitudes of Table 4, error bars one standard deviation over five seeds; right, the 2020 recipe in the model at eight settings of the EK43 grinder, with the facts Cameron et al. [1] state in their text. It reads `results.npz`.

Shows. Flattening without reversal in the main recipe, and the shape the model gives the 2020 recipe.




Figure 3: Left: cup extraction yield against grind size for a uniform puck and three amplitudes of millimeter-scale porosity noise, each rising to the finest grind. Right: the 2020 recipe in the model against grinder setting, falling monotonically with coarser grinding, with the reported equal extraction at settings 1.3 and 2.0 and the reported peak setting marked.

On the left the noisy curves fall below the uniform one only at fine grinds, by 1.55 points at $0.5d_{\text{ref}}$ and $\varepsilon = 0.02$, and all rise to the finest grind. On the right the model falls monotonically from 26.33 percent at setting 1.1 to 22.24 at 2.3, where Cameron et al. measured a peak near 1.7 and 22 percent at both 1.3 and 2.0; the model gives 25.88 and 23.60 there.

5. When the Paradox Appears

Method. The model reduced to two side-by-side columns of equal area that share the pressure drop and exchange no water: each is a series of resistances with its own flux, extracting by Definition 3, and the beverage is the sum of their outflows. This is the continuum model’s version of the two paths of Lee, Smith and Arshad [2], at fixed pressure in place of their fixed total flow. The columns start at porosity $\phi_0 \pm \delta$, uniform in depth, and are solved by the one-dimensional solver (300 cells each) at the 13 grinds, with the paradox rule of Table 4.



Parameters. The main recipe; $\delta = 0.02$, the standard deviation of the full puck's noise, and $\delta = 0.035$, the contrast fitted in [2].

Result. Two columns at fixed pressure reverse the curve. At $\delta = 0.02$ the extraction peaks at $d = 0.71 d_{\text{ref}}$ with 22.94 percent and falls to 22.14 percent at $0.5 d_{\text{ref}}$, a drop of 0.80 points and a loss of 23.8 percent against the uniform puck. At $\delta = 0.035$ the peak is 18.62 percent at the same grind, the finest grind extracts 18.39, and the loss there is 36.7 percent. The mechanism is exhaustion, as in [2]: at the finest grind and $\delta = 0.02$ the grains of the more porous column have given up 29.91 percent of their mass against an extractable 30, while those of the other have given up 20.06; water keeps pouring through the spent column and reaches the cup diluted.

Error. Corollary 1 identifies the column limit with two independent columns, and the exact linear problem puts its amplification over a shot between 0.94 and 1.09 at the reference grind and every finer one (Table 2): close to neutral. The reversal is therefore not an instability amplifying the contrast but a contrast imposed over the full height and half the width of the puck exhausting the faster column. Its loss at the finest grind, 23.8 percent, is 4.5 times the 5.35 percent that the same standard deviation costs the full puck in millimeter blobs.

Method. Table 4 shows no paradox and two puck-high columns do; this sweep asks where in between it switches on. The porosity noise keeps its standard deviation, $\varepsilon = 0.02$ exactly on the computational grid, and changes shape: a Gaussian covariance $\exp(-x^2/\ell_x^2 - z^2/\ell_z^2)$, periodic across and not in depth. Sweep A keeps blobs $\ell_x = 1$ mm wide and stretches them in depth, $\ell_z = 1, 2, 4, 7$ and 13.62 mm, up to streaks uniform in depth. Sweep B keeps the streaks depth-uniform and widens them, $\ell_x = 1, 3.6, 7, 14.5$ and 29 mm; at 29 mm a field is close to a single sinusoid across the section, the continuum cousin of two columns. The paradox rule of Table 4 was written down, unchanged, before the first shot, together with two conditional stages: $\varepsilon = 0.04$ at the two settings that bracket the onset, and the 0.1 mm grid at the onset.

Parameters. The main recipe, 13 grinds, seeds 0 to 4, $h = 0.2$ mm: 976 shots and 30 declared checks, all passed. The sweep's field generator reproduces the extraction of Table 4 to a relative 1.9×10^{-15} .

Result. Table 5 and Figure 4. Depth coherence alone never makes the paradox: as the blobs stretch from 1 mm to the full depth, the loss at $d = 0.5 d_{\text{ref}}$ grows from 5.2 through 7.5, 9.9, 12.9 and 14.2 to 16.6 percent and the fine end flattens, depth-uniform 1 mm streaks extracting 24.24 percent at 0.5 against 24.22 at $0.56 d_{\text{ref}}$, but the seed mean never turns over. Width switches it on: streaks 3.6 mm wide peak at $d^* = 0.63 d_{\text{ref}}$, 0.39 points or 1.7 percent above the finest grind, every seed peaking alone, and every wider setting peaks in every seed, with seed-by-seed drops of 0.39 ± 0.17 , 0.44 ± 0.18 , 0.46 ± 0.31 and 0.35 ± 0.06 points at $\ell_x = 3.6, 7, 14.5$ and 29 mm. The rule compares the drop with the seed-to-seed scatter of the two levels, which grows with the width, so it calls the paradox present at 3.6, 14.5 and 29 mm and absent at 7 mm, where the drop, 0.44 points, is 8.5 percent short of its threshold of 0.48. The reversal is shallow: 0.35 to 0.46 points against 0.80 for two columns, though the loss at the finest grind, 21.5 percent at 29 mm, nearly



matches their 23.8. Doubling the amplitude does not deepen it: at $\varepsilon = 0.04$ the loss at the finest grind grows to 29.0 percent for 1 mm streaks and 34.8 percent for 3.6 mm streaks, while the 3.6 mm drop stays at 0.37 points, 0.37 ± 0.15 seed by seed, which the rule calls absent because the larger scatter raises its threshold to 0.93.

Error. On the 0.1 mm grid the onset settings change by 0.14 percent or less at $d = 0.5$ and $0.71 d_{\text{ref}}$ with the same verdicts; the onset drop is 0.400 points against 0.393. The solute balance holds to 3.3×10^{-15} and every field has standard deviation 0.02 to 6.9×10^{-18} . The seed-to-seed standard deviation at the finest grind is 0.26 to 0.57 points, and 0.11 at 29 mm, where the seeds are nearly one realization. The onset is bracketed only between $\ell_x = 1$ and 3.6 mm; that the second equals λ^* is an observation on a coarse grid, not a fit.

Source. Algorithm 1, five seeds per setting, $\varepsilon = 0.02$ unless marked, $h = 0.2$ mm, the main recipe (`sweep_coherence.py`, `coherence.npz`, `coherence.json`); uniform puck and two columns from `numbers.json`; printed by `derive.py`.

Columns. Shape of the unevenness, by its correlation lengths across (ℓ_x) and in depth (ℓ_z); mean cup extraction yield, in percent, at seven grinds d/d_{ref} ; and the verdict of the declared rule: the peak grind d^* , the drop from it to the finest grind against the rule's threshold, in points, and the number of seeds that peak on their own.



Table 5: Mean cup extraction against grind by the shape of the porosity unevenness, with the verdict of the paradox rule declared before the sweep.

unevenness	0.50	0.56	0.63	0.71	1.00	1.41	2.00	paradox; d^* ; drop/thresh- old; seeds
uniform puck	29.07	28.41	27.54	26.48	22.44	17.73	12.93	—
$\ell_x = 1, \ell_z = 1$ mm	27.56	27.32	26.82	26.02	22.30	17.68	12.93	no —; —/0.37; 0/5
$\ell_x = 1, \ell_z = 2$ mm	26.89	26.73	26.32	25.63	22.13	17.57	12.84	no —; —/0.42; 0/5
$\ell_x = 1, \ell_z = 4$ mm	26.18	26.05	25.72	25.13	21.88	17.44	12.75	no —; —/0.50; 1/5
$\ell_x = 1, \ell_z = 7$ mm	25.32	25.28	25.06	24.56	21.59	17.29	12.66	no —; —/0.32; 1/5
$\ell_x = 1, \ell_z = 13.62$ mm	24.95	24.86	24.63	24.18	21.38	17.17	12.58	no —; —/0.53; 0/5
$\ell_x = 1$ mm, z-uniform	24.24	24.22	24.06	23.69	21.14	17.05	12.52	no —; —/0.72; 3/5
$\ell_x = 3.6$ mm, z-uniform	23.19	23.49	23.58	23.39	21.08	17.05	12.52	yes 0.63; 0.39/0.37; 5/5
$\ell_x = 7$ mm, z-uniform	22.97	23.27	23.41	23.29	21.10	17.06	12.52	no 0.63; 0.44/0.48; 5/5
$\ell_x = 14.5$ mm, z-uniform	22.77	23.06	23.23	23.16	21.06	17.04	12.52	yes 0.63; 0.46/0.40; 5/5
$\ell_x = 29$ mm, z-uniform	22.82	23.06	23.17	23.06	21.02	17.04	12.52	yes 0.63; 0.35/0.11; 5/5
two columns $\phi_0 \pm 0.02$	22.14	22.47	22.80	22.94	21.05	17.04	12.51	yes 0.71; 0.80
$\ell_x = 1$ mm, z-uniform, $\varepsilon = 0.04$	20.64	20.56	20.39	20.11	18.30	15.13	11.21	no —; —/0.95; 2/5
$\ell_x = 3.6$ mm, z-uniform, $\varepsilon = 0.04$	18.94	19.18	19.32	19.30	17.99	15.02	11.19	no 0.63; 0.37/0.93; 5/5

Down the table the fine end falls further below the uniform puck as the unevenness grows more coherent; a row turns over only from the depth-uniform streaks 3.6 mm wide onward.

Script. `fig_coherence.py`: the mean cup extraction yield against grind, error bars one standard deviation over five seeds, for Sweep A on the left and Sweep B on the right, with the uniform puck dashed, the two columns $\phi_0 \pm 0.02$ dotted on the right, and a red dot at the peak d^* wherever the declared rule finds the paradox. It reads `coherence.npz` and, for



the two columns, `results.npz`.

Shows. Where the paradox switches on as the unevenness becomes coherent in depth and then wide.




Figure 4: Cup extraction yield against grind size for porosity unevenness of growing coherence at a fixed standard deviation. Left: blobs one millimeter wide stretched in depth up to streaks uniform in depth, every curve rising to the finest grind. Right: depth-uniform streaks of growing width, which peak at an interior grind from 3.6 mm wide onward, with two puck-high columns for comparison.

On the left every curve rises to the finest grind; stretching the blobs lowers the fine end from 27.56 to 24.24 percent but turns no curve over. On the right the streaks 3.6 mm wide and wider peak at $d = 0.63 d_{\text{ref}}$, a grid step finer than the two columns, and drop by 0.35 to 0.46 points against 0.80. The 7 mm streaks carry no red dot: every seed peaks, but the drop misses the declared threshold.

Method. The two-pathway model of Lee, Smith and Arshad [2], their Eqs. (1) to (9), rerun from the arXiv version of their paper: two beds of equal volume share a fixed total flow in proportion to their Kozeny–Carman permeabilities, each dissolves at their transfer rate, and dissolution in a bed stops when its extraction reaches the largest extractable fraction, handled exactly by event location. The bed volume follows from the dose through their Eq. (4), and the parameter $\alpha = c_{\text{sat}}/\rho_c$ is taken from their Eq. (8). The extraction is counted, as in their Eq. (3), from each path’s starting porosity, and for comparison from the mean porosity ε_0 .

Parameters. Their Table I: a 40 g shot from 20 g, $\varepsilon_0 = 0.173$, $c_{\text{sat}} = 212.4 \text{ kg/m}^3$, $\lambda/D = 0.125 \times 10^6 \text{ s/m}$, $\delta = 0.035$, a largest extraction of 33.8 percent, grain density $\rho_c = 798 \text{ kg/m}^3$ (their imposed value) or 399 kg/m^3 (the measured one); their fits of shot time and grain surface to the grind setting g ; g from 1.1 to 2.3 in steps of 0.05.

Result. With $\rho_c = 798 \text{ kg/m}^3$ the rerun has a shallow interior peak, 24.31 percent at $g = 2.0$ against 24.10 percent at $g = 1.1$ and 24.25 percent at 2.3; counted from ε_0 , the



peak is 23.09 percent at $g = 1.45$. At $g = 1.1$ the more porous bed is exhausted at dimensionless time 0.620 (0.564 counted from ε_0), within a shot of 1.12. With $\rho_c = 399$ kg/m³ there is no peak: the extraction rises monotonically as the grind gets finer, from 28.61 percent at $g = 2.3$ to its ceiling of 33.8 percent at $g = 1.1$, and does not decrease, as [2] state for this value.

Error. The rerun reproduces the model of [2] qualitatively, not to their figures. Their Figure 7 caption places the exhaustion of the first bed at $\tau \approx 0.8$ for $g = 1.1$, against our 0.620 or 0.564. One entry of their Table I we could not use as printed, and we give the reason only as our reading. Their Eq. (8) defines " $\alpha = c_{\text{sat}}/\rho_c$ ", which is 0.266 at the imposed density, while their Table I lists " α 3.76 Calculated", which equals ρ_c/c_{sat} at that density to the precision printed (3.757). Read literally, the tabulated value exhausts both beds and gives 33.80 percent at $g = 1.1$, 1.7 and 2.3 alike, with no dependence on grind, so we use the value of Eq. (8).

Script. `fig_twopath.py`: left, the two-pathway model of Lee, Smith and Arshad [2] rerun at fixed total flow against the grind setting, for their imposed grain density with the extraction counted two ways and for the measured density; right, this model at fixed pressure against grind, as two columns at $\phi_0 \pm 0.02$ and $\phi_0 \pm 0.035$, as the full puck with millimeter noise $\varepsilon = 0.02$, and as a uniform puck. It reads `results.npz`.

Shows. What two lumped paths do and what a continuum does with the same contrast.




Figure 5: Left: extraction yield against grinder setting in the rerun two-pathway model, a shallow peak at the imposed grain density and a monotone rise at the measured one. Right: this model at fixed pressure, where two columns peak at an interior grind and fall at the finest grinds, while the full puck with the same standard deviation keeps rising.

On the left the imposed density peaks at 24.31 percent at $g = 2.0$ and the measured density does not peak. On the right both pairs of columns peak at $d = 0.71 d_{\text{ref}}$ and lose 23.8 and 36.7 percent at $0.5 d_{\text{ref}}$, where the full puck with the narrower pair's standard deviation loses 5.35 percent and still rises.



Remark 3 (Why only coherent, wide defects reverse the curve). Two lumped paths and a continuum differ in what makes the fast path. In the column limit of Corollary 1 the paths exchange no water, and over a shot the exact linear problem finds that limit nearly neutral, an amplification between 0.94 and 1.09 at the reference grind and finer. The reversal of two columns therefore owes little to an instability: it comes from a contrast imposed over the whole height and half the width of the puck, which exhausts the faster column while the slower one still holds solubles, the mechanism of Lee, Smith and Arshad [2]. In a continuum with millimeter blobs, water crosses between neighbors. A fast blob cannot become a fast path unless the instability makes it one, and the instability grows fastest near its own spacing, $\lambda^* = 3.6$ mm, and only over the part of the shot when its growth rate is positive. The cross-flow caps the cost: the same standard deviation that costs two columns 23.8 percent of the extraction at the finest grind costs the full puck 5.35 percent, and the curve does not turn over.

The coherence sweep fills in between. Coherence in depth gives a fast path its first ingredient, a defect that no fresh bed blocks: the loss at the finest grind rises from 5.2 to 16.6 percent as millimeter blobs stretch into depth-uniform streaks, but the curve still rises. Width gives the second, and streaks at least about one fastest spacing wide turn the curve over, though shallowly, 0.35 to 0.46 points against 0.80 for two columns, consistent with water still crossing from a spent streak into neighbors that keep extracting. What survives from the two-path picture is its mechanism, exhaustion of a coherent fast path; what does not is the expectation that any small unevenness of the packing will produce it.

6. Against the Experiment

Method. The recipe of Cameron et al. [1] in the model: 20 g of coffee, 40 g of beverage, 6 bar, at eight settings of their grinder, mapped to grinds through the grain surface, $d/d_{\text{ref}} = S(1.7)/S(g)$, with $S(g)$ the linear fit of Lee, Smith and Arshad [2] to the surface areas measured in [1]. Each setting is run as a uniform puck and with noise $\varepsilon = 0.02$ in 1 mm blobs, five seeds; nothing is refitted. The comparison uses only what [1] state in their text, three facts: the extraction falls at settings finer than 1.7; it is 22 percent at both 1.3 and 2.0; and their one-dimensional model exceeds their measurement by 13.1, 6.1 and 2.6 percent at 1.1, 1.3 and 1.5. Their figures were not digitized. For shot times, the comparison is the linear fit of [2] to the shot times measured in [1].

Parameters. Settings 1.1, 1.3, 1.5, 1.7, 1.9, 2.0, 2.1 and 2.3, which map to d/d_{ref} from 0.840 to 1.235; $h = 0.2$ mm; seeds 0 to 4.

Result. Table 6. The model extracts 26.33 percent at setting 1.1 and falls monotonically to 22.24 percent at 2.3. Its curve peaks at the finest setting, not near 1.7, and gives 25.88 percent at 1.3 against 23.60 percent at 2.0, not equal yields. The loss to channeling is 1.8 percent of the yield at 1.1 and 0.5 percent at 2.3: the range of settings in the experiment lies between 0.84 and $1.24 d_{\text{ref}}$, where the full puck loses little. The ratio of shot times between the finest and the coarsest setting is 1.59 in the model and 1.57 in the fit to the measurements.

Error. The shot-time ratio agrees within 1.2 percent, which supports the mapping from



setting to grind. The absolute shot times are 1.59 times the measured ones at 1.1 and 1.57 times at 2.3, a nearly constant factor, as expected when k_{ref} is calibrated on a different machine and coffee; the yields, 22.24 to 26.33 percent, sit above the 20 to 23 percent that [1] report. The shape and the size of the drop do not match: the losses in the model are 7.3, 4.1 and 2.1 times smaller than the excess of the one-dimensional model of [1] over their measurement at 1.1, 1.3 and 1.5.

Source. Algorithm 1 with the recipe of Cameron et al. [1], a uniform puck and five seeds of $\varepsilon = 0.02$ per setting, $h = 0.2$ mm (`model.py` stage `f6`, `numbers.json`); the last column is the fit of Lee, Smith and Arshad [2] to the shot times measured in [1]; printed by `derive.py`.

Columns. Grinder setting; the grind it maps to; cup extraction yield, in percent, with noise (mean $\pm$ standard deviation) and uniform; loss to channeling, in percent; shot time, seed mean; the published linear fit to the measured shot times.

Table 6: The 2020 recipe in the model: extraction, loss to channeling and shot time by grinder setting, with the fit to the measured shot times.

g	d/d_{ref}	EY, noise	EY, uniform	loss (%)	shot (s)	fit (s)
1.1	0.840	26.33 ± 0.08	26.80	1.8	60.2	37.9
1.3	0.887	25.88 ± 0.06	26.26	1.5	55.9	35.5
1.5	0.940	25.34 ± 0.05	25.65	1.2	51.8	33.2
1.7	1.000	24.72 ± 0.04	24.96	1.0	48.0	30.9
1.9	1.068	24.00 ± 0.04	24.19	0.8	44.4	28.7
2.0	1.105	23.60 ± 0.04	23.77	0.7	42.7	27.5
2.1	1.146	23.17 ± 0.03	23.33	0.7	41.0	26.3
2.3	1.235	22.24 ± 0.03	22.36	0.5	37.8	24.1

The model’s yield falls at every step from setting 1.1 to 2.3, while the measured one peaks near 1.7.

Remark 4 (What else can make a finer grind extract less). The model leaves three candidates open. First, fines: Cameron et al. [1] attribute the drop to fines clogging part of the bed, which makes a fraction of the coffee inaccessible, and the model has no fines and no migration of particles. Second, coherent defects: Remark 3 shows that unevenness running the full height of the puck and at least about one channel spacing wide does reverse the curve, but shallowly, by 0.35 to 0.46 points at a standard deviation of 0.02, and doubling the amplitude does not deepen it. A clogged region that runs through the puck would be such a defect, so the two explanations are not exclusive. Third, the mechanics of the bed: Waszkiewicz et al. [3] show that the puck compacts under pressure, so that its flow saturates as the pressure rises, and the model holds the bed rigid.

The comparison rules out something narrow: millimeter-scale unevenness, acting through



the dissolution instability alone, does not account for the drop that [1] measured. It leaves a sharper question for experiment: whether finer grinds carry defects coherent over the height of the puck, and how wide they are.

7. What a Barista Can Do

Method. Three things a barista can change, tried at the grind where channeling costs most, the finest grind of the sweep, against the shots of Table 4 with the same five noise seeds. (a) Pre-infusion: 3 bar for the first 8 s, then 9 bar. (b) A declining profile: the pressure falls linearly with the beverage poured, from 9 bar at the start to 6 bar at the end. (c) Better puck preparation: the porosity noise halved, from $\varepsilon = 0.02$ to 0.01, same seeds. For (a) and (b) a uniform puck is run with the same profile; its change in extraction is the part of the gain that comes from the pressure history alone, and the remainder is the part that comes from less channeling. All three levers and their sizes were declared before any shot.

Parameters. $d = 0.5 d_{\text{ref}}$, the main recipe; the baseline extracts 27.51 percent in a shot of 79.7 s, against 29.07 percent for a uniform puck, a loss of 1.55 points.

Result. Pre-infusion changes the extraction by $+0.01 \pm 0.01$ points and the uniform puck by $+0.02$: nothing. The declining profile gains $+0.25 \pm 0.03$ points, but the uniform puck under the same profile gains $+0.36$, because the lower pressure lengthens the shot, to 89.4 s for the five seeds against 79.7 s; the part due to channeling is -0.11 points, so the profile wins only by brewing longer. Halving the noise gains $+0.84 \pm 0.16$ points, all of it from less channeling, and narrows the spread of local extraction by 1.03 ± 0.17 points.

Error. The seed-to-seed standard deviations are the uncertainties quoted with each change. Better puck preparation recovers 54 percent of the 1.55-point loss, five standard deviations of its scatter above zero. The declining profile recovers 16 percent of it, less than the slower shot alone gives a uniform puck, and pre-infusion 0.6 percent. Only puck preparation reduces channeling.

Script. `fig_levers.py`: at $d = 0.5 d_{\text{ref}}$, the change in cup extraction yield (left) and in the spread of local extraction yield (right) for each lever against the 9 bar baseline with the same five seeds, error bars one standard deviation over the seeds; grey bars on the left give the same pressure profile on a uniform puck. It reads `numbers.json`.

Shows. Which lever buys extraction by reducing channeling, and which only by changing the shot.




Figure 6: Change in cup extraction yield and in the spread of local extraction for pre-infusion, a declining pressure profile and better puck preparation at the finest grind; only better puck preparation gains extraction and narrows the spread, while the declining profile’s gain is matched by the same profile on a uniform puck.

Pre-infusion is indistinguishable from zero on both panels. The declining profile’s +0.25 points stand below its uniform control, +0.36. Better puck preparation gains +0.84 points and narrows the spread by 1.03; it has no control, since a uniform puck has no noise to halve.

Remark 5 (Why pre-infusion does nothing here). Under Assumption 1 the bed is saturated from the first instant, so the model has no wetting stage. Pre-infusion at 3 bar for 8 s then only lowers the flux for a few seconds: the same water passes later, the extraction changes by $+0.01 \pm 0.01$ points, and the spread by nothing measurable. This is a limit of the model, not a verdict on pre-infusion.

Pre-infusion is meant to wet the puck before full pressure arrives, so that dry pockets do not become preferential paths. That is the stage the model leaves out, along with the swelling of the grains and the escape of gas. A model with a wetting front would start each shot from a puck whose unevenness depends on how it was wetted; the coherence sweep suggests where to look, since wetting defects that run through the puck and are a few millimeters wide are the kind that reverse the extraction curve. The declining profile, by contrast, is inside the model’s scope, and its verdict stands: it gains extraction by lengthening the shot, not by reducing channeling.

For practice, the results read as a diagnostic. With a sound puck the model never lets a finer grind extract less: fine-scale unevenness lowers the extraction at fine grinds but does not reverse the extraction-grind curve in the grind sweep. In the coherence sweep a reversal needs a defect that runs through the full height of the puck and is at least about one channel spacing, 3.6 mm, wide. So if grinding finer makes a shot weaker, the first suspect is the puck, not the grinder. The case the model covers directly is a crack: a streak through the bed that is more open than the coffee around it. Two further suspects



come from practice and lie outside the model, which is two-dimensional, periodic and has no basket wall: a gap where the puck meets the wall of the basket, and a tilted tamp that leaves one side of the bed thinner than the other. We name them as candidates for experiment, not as results.

8. Limits

Remark 6 (What the model leaves out). The bed is saturated from the start, so there is no wetting stage, which is what pre-infusion acts on. Fines neither migrate nor clog, though Cameron et al. [1] attribute the measured drop to clogging. The bed is rigid, while Waszkiewicz et al. [3] show that it compacts under pressure and that its flow saturates as the pressure rises. There is no gas, no swelling and no change of temperature. One effective grain size sets both the permeability, $k \propto d^2$, and the surface, $a \propto 1/d$, while real grinds are bimodal. Solute spreads by molecular diffusion only; mechanical dispersion would act like a larger diffusivity across the flow and could widen λ^* and weaken the growth. The porosity opens by the full volume dissolved, an upper bound on the feedback. The section is two-dimensional and periodic, not a three-dimensional puck bounded by the basket wall. The beverage is taken at the density of water. And the one calibrated constant comes from a different machine and coffee than the experiment we compare with.

The coherence sweep brackets its onset only between streaks 1 and 3.6 mm wide; at 29 mm the seeds are nearly one realization; wide but shallow defects were not swept; and the declared rule is unpaired, so its verdicts at 3.6 and 14.5 mm are marginal and at 7 mm negative although every seed peaks. Two results were computed and not claimed: a closed-form long-wave threshold for the fresh puck, which disagreed with its eigenvalue check by 4.6 percent and had no direct simulation behind it, and the end-of-shot amplification of long waves at the finest grind, where the linear problem and the simulation diverge in the last percent of the shot.

9. Conclusion

Channeling in an espresso puck is a reactive-infiltration instability: where the front between extracted and fresh coffee runs ahead, the bed has opened, its resistance has fallen, and more water arrives. At fixed pressure the front moves at $\gamma_f q$, and a thin front grows at a closed-form rate whose long-wave limit is the growth of two independent columns and which rises with the wavenumber until the wave is as short as the front is thick. A direct simulation confirms the law within 0.11 percent and a moving-frame eigenvalue within 5.0 percent, while a frozen analysis in fixed cells overstates the growth 3.52 times. At practical grinds the exact linear problem along the moving base puts the fastest channel spacing at 3.6 mm at every grind, and unevenness decays over a shot coarser than 1.58 times the reference grind.

In two-dimensional shots with exact solute bookkeeping, millimeter-scale porosity noise costs 5.35 percent of the extraction at half the reference grind and flattens the extraction-grind curve without reversing it. The curve turns over only when the unevenness runs the



full height of the puck and is at least about one channel spacing wide, and then by 0.39 points, 1.7 percent, at onset, against 0.80 points for two lumped columns. With the 2020 recipe the model matches the measured ratio of shot times, 1.59 against 1.57, but not the shape or the size of the measured drop. Only better puck preparation reduces channeling, recovering 54 percent of the loss. On this evidence the finer-grind paradox needs defects coherent through the puck, such as clogged or poorly wetted regions, more than it needs the instability alone.



References

- [1] Cameron, M. I., Morisco, D., Hofstetter, D., Uman, E., Wilkinson, J., Kennedy, Z. C., Fontenot, S. A., Lee, W. T., Hendon, C. H., & Foster, J. M. (2020). Systematically improving espresso: insights from mathematical modeling and experiment. *Matter*. <https://doi.org/10.1016/j.matt.2019.12.019>
- [2] Lee, W. T., Smith, A., & Arshad, A. (2023). Uneven extraction in coffee brewing. *Physics of Fluids*. <https://doi.org/10.1063/5.0138998>
- [3] Waszkiewicz, R., Myck, F., Białas, Ł., Puciata-Mroczynska, M., Dzikowski, M., Szymczak, P., & Lisicki, M. (2026). Under pressure: poroelastic regulation of flow in espresso brewing. *Physics of Fluids*. <https://doi.org/10.1063/5.0319611>
- [4] Moroney, K. M., Lee, W. T., O'Brien, S. B. G., Suijver, F., & Marra, J. (2015). Modelling of coffee extraction during brewing using multiscale methods: an experimentally validated model. *Chemical Engineering Science*. <https://doi.org/10.1016/j.ces.2015.06.003>
- [5] Moroney, K. M., Lee, W. T., O'Brien, S. B. G., Suijver, F., & Marra, J. (2016). Asymptotic analysis of the dominant mechanisms in the coffee extraction process. *SIAM Journal on Applied Mathematics*. <https://doi.org/10.1137/15M1036658>
- [6] Moroney, K. M., O'Connell, K., Meikle-Janney, P., O'Brien, S. B. G., Walker, G. M., & Lee, W. T. (2019). Analysing extraction uniformity from porous coffee beds using mathematical modelling and computational fluid dynamics approaches. *PLOS ONE*. <https://doi.org/10.1371/journal.pone.0219906>
- [7] Chadam, J., Hoff, D., Merino, E., Ortoleva, P., & Sen, A. (1986). Reactive infiltration instabilities. *IMA Journal of Applied Mathematics*. <https://doi.org/10.1093/imamat/36.3.207>
- [8] Hinch, E. J., & Bhatt, B. S. (1990). Stability of an acid front moving through porous rock. *Journal of Fluid Mechanics*. <https://doi.org/10.1017/S0022112090001963>
- [9] Szymczak, P., & Ladd, A. J. C. (2014). Reactive-infiltration instabilities in rocks. Part 2. Dissolution of a porous matrix. *Journal of Fluid Mechanics*. <https://doi.org/10.1017/jfm.2013.586>



Zamrik Quant Research

The research arm of Zamrik Quant Lab

About the Desk

Applied Mathematics Desk

Zamrik Quant Research


Mathematics in the wild — floods, fires, cellars, patterns and maps — and everything else worth solving.


The Applied Mathematics desk solves problems numerically, whatever field they come from. Its subject is the method: how a class of equations is discretised, how fast the approximation converges, and how to know that the number on the screen is the solution and not an artefact of the grid.

Its work spans the main classes of partial differential equations. Parabolic problems — heat in a sphere, the seasonal lag that leaves a cellar warmest in winter, axial dispersion in a tubular reactor — are solved with finite elements and finite differences. Hyperbolic problems cover waves solved by spectral methods, shocks and tidal bores. Reaction–diffusion systems produce the patterns on an animal’s coat, and the induction equation of magnetohydrodynamics carries magnetic diffusion through the Sun.


Alongside the equations sit the general tools every modeller depends on: quadrature for rapidly oscillating integrands by the Filon and Levin methods, Monte Carlo and the quasi-Monte Carlo methods that beat its $1/\sqrt{N}$ rate, matrix factorisations for sampling correlated variables, and stochastic simulation of systems of many interacting parts, from flocks of birds to molecular motors. Each method is checked against an exact solution or a second, independent method before its result is published.

The desk is supervised by **Dr. T. Zamrik**, Head of Zamrik Quant Research, PhD, Imperial College London (2011).





Hamiltonian flow of the Ricker wavelet: $\dot{\mathbf{x}} = J \nabla f(\mathbf{x})$, $J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, $f(r) = \left| (1 - r^2/\sigma^2) e^{-r^2/(2\sigma^2)} \right|$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



WP-2026-15315794
zamrik.com