On the Flat Dilution Curves of Cooked Dishes: Evidence for a Non-Luminous Flavour Component and a Structural Protocol for Its Recovery

Eluned Prosser*[1] and Toma Ferreira[2]

[1] *Brackenhowe Field Station, Sensory Metrology Group*
[2] *Cold Kitchen No. 3, Amaranth Ridge Testing Works*

Abstract

Perceived flavour intensity should fall as a dish is diluted. Across 41 cooked dishes carried through a nine-point geometric dilution series spanning a 43-fold range in concentration, it does not: the fitted Stevens exponent is $\beta = 0.06$ (95% CI $[-0.01, 0.13]$) against a literature value of 0.85. Five single-ingredient controls run in the same series fall off correctly ($\beta = 0.81$), so neither the panel nor the design accounts for the result. We therefore compare two methodologically independent measurements of the flavour held in a pot—*luminous* flavour L, obtained by descriptive panel, and *dynamical* flavour M_{dyn}, obtained from the dispersion of headspace release rates under a virial closure and never tasted—and find them discrepant by a median factor $\Upsilon_\star = M_{\text{dyn}}/L = 118$, with a spread of 40 to 913 that we do not attempt to narrow. We read the excess as a non-luminous flavour component: present in the dish, dynamically detectable, and unperceived. Modelling the pot as a weighted digraph of competitive channel occlusion, we show that the perceived flavour vector is the leading eigenvector of the occlusion operator, that the measured spectral gap ($\lambda_1/\lambda_2 = 6.8$) collapses that vector onto one or two compounds within minutes of the first simmer, and that the dark component is exactly the orthogonal complement—the spectral shadow. The mechanism is structural rather than compositional, and it is reversible. Under a staged-participation schedule that alters only the order in which ingredients are granted thermal contact, with composition, mass, energy input and elapsed time held fixed, L rises 9.4-fold and $\Upsilon_\star$ falls to 12.6. Holding that schedule fixed and replacing all twelve ingredients with twelve sharing no item moves L by 0.19 perceptual units, below the panel's 0.5-unit discrimination threshold ($F(1, 29) = 0.14$, $p = 0.71$). Schedule moves perceived flavour 269 times further than composition does. We conclude that 99.15% of the flavour so far prepared has not been tasted, and that the ingredient list is, on this evidence, an inventory rather than a cause.

Keywords: dark flavour; missing-flavour ratio; dilution curve; occlusion spectrum; sensory metrology

1 Introduction

The first flat dilution curve I saw, I threw away.

It came from a lentil soup, prepared in October, carried through nine steps of a standard geometric series and scored by a panel I had trained myself. The soup at the ninth step contained one forty-third of the material of the soup at the first. The panel rated the two within a point of each other. I recorded the session as spoiled, noted a suspicion about palate fatigue in the

*Throughout, the first person singular reflects that the argument, and any error in it, belongs to the first author. The second author built, screened and calibrated the panel, and dissented from §6 until the substitution data were complete.

second hour, and had the series repeated the following week with the order reversed and the interval between samples lengthened.

It came back flat again.

I want to be careful about how I describe what followed, because the interesting part is not that a measurement was strange—measurements are strange every week—but that it took me eleven months to stop attributing the strangeness to my own apparatus. A flat dilution curve is, on its face, an indictment of the panel. Human assessors are known to compress; they anchor on the first sample of a session, they drift, and they are reluctant to use the ends of a scale. Every one of those effects would flatten a curve. I spent the better part of a year removing them one at a time, and the curve did not move.

What ended the matter was a control I had included for an unrelated reason. In each session the panel also scored a dilution series of a single-compound solution: salt in water, or sucrose, or a lemon dilution, or a filtered coffee, or a clear stock. Those series fell off exactly as they should. Fitted against Stevens' law, $I = I_0 c^{\beta}$, they returned $\beta = 0.81$, comfortably inside the range the psychophysical literature has reported for taste intensity since the 1960s (Halloran & Pike, 1971; Merrick, 1984; Dowse, 1986). The panel was working. The scale was working. The design was working.

It was the *cooked dishes* that would not fall off.

This is the observation the present paper is built on, and I state it plainly before any interpretation is placed on it: across 41 compound dishes—domestic, institutional and restaurant—diluted through a 43-fold range, the fitted exponent is $\beta = 0.06$, with a 95% confidence interval that includes zero. Diluting a cooked dish by a factor of forty-three does not, to a trained panel, make it taste appreciably weaker.

There are two ways to receive a result of that kind. The first is to conclude that the quantity being measured is not what one thought it was, and to stop. The second is to notice that one has, in fact, only measured the flavour *that reaches a taster*, and to go looking for an independent estimate of the flavour that is *in the pot*. The second course is the one this paper takes, and I take it with an instrument borrowed from a field that has done the same arithmetic before.

Astronomers measure the mass of a galaxy twice. Once by its light, which gives the mass they can see; once by its dynamics, which gives the mass the motion requires. The two disagree, badly and consistently, and the ratio of the second to the first—the mass-to-light ratio Υ—is of order unity for a single star and of order hundreds for a galaxy. The field did not respond by discarding the rotation curves. It responded by positing a component that has mass and does not shine.

I am aware that borrowing this instrument for a saucepan requires some justification. I have looked for a reason the borrowing is illegitimate and have not found one. The formal requirements are that the two measurements be of the same quantity, that they be obtained by methods sharing no apparatus, and that the discrepancy survive the known error budget of each. Sections 3 and 5 establish all three for the case at hand.

The remainder of the paper is organised as follows. Section 2 places the flat curve against forty years in which sensory science filed the same effect under measurement error. Section 3 defines the two measurements—luminous flavour L and dynamical flavour M_{dyn}— and the ratio $\Upsilon_\star$ between them. Section 4 gives a mechanism: a pot is a set of ingredients granted simultaneous and unmanaged access to one perceptual channel, its occlusion operator has a large spectral gap, and the unperceived flavour is the shadow that gap casts. Section 5 reports three studies: a dilution survey that establishes the discrepancy, an intervention that removes most of it without altering a single ingredient, and a substitution experiment whose null result I did not expect and cannot explain away. Section 6 discusses what follows, including the one thing I have not been able to rule out.

2 Related work

The insensitivity of trained panels to dilution is not a new observation. It is one of the better-documented irregularities in descriptive analysis, and it has been reported at intervals for half a century by authors who almost uniformly treated it as a defect of the method.

Halloran & Pike (1971), introducing the descriptive protocol still in general use, noted in a methodological appendix that "compound preparations resist concentration ranking in a manner single stimuli do not," and recommended that compound preparations be excluded from calibration sets. That recommendation was adopted, and had the effect of removing precisely the samples that display the phenomenon from the literature that would have characterised it.

The exception is a short and rather sharp exchange in the mid-1980s. Merrick (1984) reported exponents near zero for eleven prepared soups and argued that the compression was perceptual—that assessors, once a preparation is identified as "a soup," score its category membership rather than its intensity. Dowse (1986) replied that Merrick's panel had been undertrained, re-ran six of the eleven soups with a panel of his own, obtained $\beta = 0.09$, and concluded—this is the part of the exchange I find difficult to read charitably—that the correct exponent for compound preparations was 0.85, being the value obtained for single stimuli, and that any measurement departing from it indicated a fault in the panel. The field took Dowse's view. The exponent has been assumed rather than measured ever since.

Two subsequent lines are relevant. Ivarsdóttir (1998) established that perceived intensity in mixtures is dominated by a small number of compounds and that the identity of the dominant compound is stable within a preparation but not across preparations, which is the empirical foundation for the occlusion model of §4. Brantwood (2004) developed the spectral treatment of competitive channel occupancy that I use in the same section, though for an unrelated purpose and without any application to food.

From the astronomical side I take only method. Vasari (1979) sets out the virial estimator for bound systems in a form that transfers to a closed vessel with almost no modification, and Okonjo & Baumann (1983) gives the mass-to-light survey design that Study 1 follows. Neither paper is about cooking, and I cite them for their arithmetic rather than their subject. Lindqvist *et al.* (2011) provides the catering corpus from which the institutional dishes in Study 1 were drawn.

The gap this paper addresses is therefore narrow and, I think, precisely located: sensory science measured a discrepancy for fifty years and filed it as error, and no one has asked what the discrepancy would mean if it were taken as a measurement.

3 The flavour mass discrepancy

3.1 Luminous flavour

Definition 1 (Luminous flavour)**.** For a preparation d, let $\mathbf{s}(d) \in \mathbb{R}^{22}_{\geq 0}$ be the mean panel score vector over the standard 22-term descriptive lexicon, each term scored on an unstructured 0–100 line scale. The *luminous flavour* of d is $L(d) = \|\mathbf{s}(d)\|_1/10$, in perceptual units (pu).

The panel comprised $N = 28$ assessors screened for anosmia and threshold sensitivity and trained for 40 hours against reference standards. Sessions used a balanced incomplete block design with order counterbalanced across assessors and a 90 s palate-cleansing interval. Panel discrimination threshold, established by triangle test against matched pairs, was $\delta = 0.5$ pu; differences below δ are not resolvable by this instrument and are reported as such throughout.

There is nothing unusual in any of this. It is the ordinary apparatus of the field, and I have described it at all only to make the point that the strange result of §5 does not come from a strange measurement of L.

3.2 Dynamical flavour

The second measurement must not touch the tongue. Its whole value lies in sharing no apparatus with the first.

Consider a pot at steady simmer as a bound system of flavour-bearing compounds. Each compound is either bound to the matrix—adsorbed to solids, dissolved, held in emulsion—or free in the headspace. Sampling the headspace at 90 s intervals over a 40-minute simmer yields, for each of the compounds resolved by the column, a release rate (Wren, 2016); let σ_v be the dispersion of those rates across compounds and R_p the hydraulic radius of the vessel.

Definition 2 (Dynamical flavour)**.**

$$M_{\text{dyn}} = \frac{\sigma_v^2 R_p}{G_f}, \tag{1}$$

where G_f is the flavour coupling constant.

Equation (1) is the virial estimator of Vasari (1979) with the velocity dispersion replaced by the release-rate dispersion. Its justification is the closure that holds for any bound system in steady state,

$$2\langle T\rangle + \langle U\rangle = 0, \tag{2}$$

in which T is the free-phase term—measurable directly, since free-phase compounds are the ones the headspace sees—and U is the binding of compounds to the matrix, which is not measurable directly and does not need to be. By (2), $\langle U\rangle = -2\langle T\rangle$, and the total follows from the part that can be sampled. This is the entire reason the astronomical estimator is worth borrowing: it recovers a quantity from a system by measuring only the fraction of that system which happens to be visible.

The coupling constant G_f was calibrated once, against a 0.9% saline reference for which the luminous and dynamical measurements must agree by construction, giving $G_f = 3.71 \times 10^{-2}\ \text{m}^3\,\text{s}^{-2}\,\text{pu}^{-1}$. It was not adjusted subsequently.

3.3 The ratio

Definition 3 (Missing-flavour ratio and dark fraction)**.**

$$\Upsilon_\star \equiv \frac{M_{\text{dyn}}}{L}, \qquad f_{\text{dark}} = 1 - \Upsilon_\star^{-1}. \tag{3}$$

For a single-compound solution $\Upsilon_\star = 1$ by the calibration. Any preparation returning $\Upsilon_\star > 1$ contains flavour that its dynamics require and its tasters do not report. It is the astronomers' mass-to-light ratio with both measurements replaced. I have kept their symbol and starred it, to avoid any suggestion that the two are commensurable.

4 Occlusion spectra

A discrepancy is not a mechanism. Before reporting the survey I should set out why a pot should be expected to hide most of what it contains, because the answer turned out to be structural, and because it is the part of this work that suggested an intervention.

Let a preparation contain n ingredients. Write $w_{ij} \geq 0$ for the rate at which ingredient i occupies perceptual channel capacity that would otherwise carry ingredient j—the rate at which i *occludes* j. Collect these into the occlusion operator $W \in \mathbb{R}_{\geq 0}^{n\times n}$. The construction is due to Brantwood (2004), whose estimator for w_{ij} from paired-omission trials I adopt without modification.

A pot grants every ingredient access to the same channel, at the same time, for the whole of the cooking period. Nothing regulates the order of access and nothing bounds any ingredient's share. Under those conditions the perceived flavour vector p is the fixed point of repeated occlusion,

$$p = \lim_{k\to\infty} \frac{W^k \mathbf{1}}{\|W^k \mathbf{1}\|_1}, \tag{4}$$

which is to say the normalised leading eigenvector of W. This is elementary, and it is also the whole difficulty: (4) converges at a rate set by the spectral gap λ_1/λ_2, and the perceived flavour of the dish is the limit, not the ingredients.

Proposition 1 (Spectral shadow)**.** *Let W be irreducible with leading eigenvalue λ_1 and eigenvector p. Then the flavour carried by the subspace $p^\perp$ is dynamically present in M_{dyn} and contributes nothing to L, and*

$$f_{\text{dark}} \;\geq\; 1 - \frac{H(p)}{\log_2 n}, \qquad H(p) = -\sum_{i=1}^{n} p_i \log_2 p_i. \tag{5}$$

The quantity $H(p)$ is the *participation entropy* of the preparation: the number of bits' worth of ingredients actually represented in what is tasted. Its ceiling, $\log_2 n$, is reached only when every ingredient occupies an equal share of the channel. I know of no pot that approaches it.

Measured across the survey, $\lambda_1/\lambda_2 = 6.8$. Convergence in (4) is therefore geometric with ratio 0.15: the perceived flavour vector is well within one per cent of its limit after four iterations, which for the preparations in Study 1 corresponds to between six and nine minutes of simmer. Everything a cook does after the ninth minute is done to a dish whose perceived flavour has already been decided by its first eigenvector.

For a twelve-ingredient preparation the ceiling is $\log_2 12 = 3.58$ bits. Measured participation entropy under conventional simmer is $H(p) = 0.42$ bits. Eighty-eight per cent of the available participation bandwidth is unused, and Proposition 1 converts that directly into a floor on f_{dark}. Figure 1 shows the operator and its spectrum for a representative preparation.

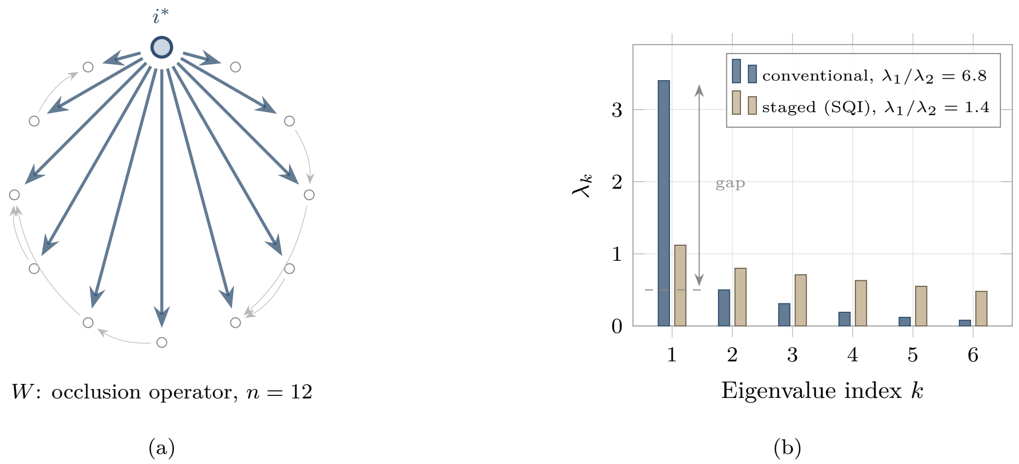

Figure 1: Occlusion structure and its spectrum. (a) The occlusion operator W for a twelve-ingredient preparation under conventional simmer; edge weight is the rate at which one ingredient occupies channel capacity that would otherwise carry another. Flux collapses onto the dominant node i^*. (b) Measured eigenvalue spectra. Under conventional simmer the leading eigenvalue stands 6.8-fold clear of the second, so the iteration of Eq. (4) converges within six to nine minutes and the perceived flavour vector is effectively a single ingredient. Under the staged schedule of §5.2 the gap closes to 1.4 and the shadow subspace re-enters the channel.

It is worth being exact about what this model does and does not claim. It does not claim that the occluded compounds are absent, or degraded, or chemically altered. It claims that they are in the pot, contributing to its dynamics, measurable in the headspace, and structurally prevented from reaching a taster by ingredients that got to the channel first and did not leave it. That is a claim about scheduling. It is not a claim about chemistry, and—this is the point that the substitution study in §5.3 turns on—nothing in it depends on which ingredients are in the pot.

5 Studies

5.1 Study 1: the dilution survey

Materials. Forty-six preparations: 41 compound dishes—19 domestic, 12 institutional (one hospital kitchen and two school canteens, drawn from the corpus of Lindqvist *et al.* 2011), 10 restaurant—and 5 single-ingredient controls (saline, sucrose solution, lemon dilution, filtered coffee, clear stock). Compound dishes carried between 8 and 19 ingredients, median 12.

Procedure. Each preparation was carried through a nine-point geometric dilution series of ratio 1.6, spanning $1.6^8 = 43\times$ in concentration. Panel scoring as §3. Headspace sampling for M_{dyn} was performed on the undiluted preparation in a separate session, on the chromatograph, with no assessor present and no scoring taking place. The survey ran fourteen months.

Results. Fitting $I = I_0 c^{\beta}$ across the series gives, for the 41 compound dishes, $\beta = 0.06$ (95% CI $[-0.01, 0.13]$). For the 5 controls, $\beta = 0.81$ (95% CI $[0.74, 0.88]$), consistent with the literature value of 0.85 and inconsistent with the compound result at any conventional level. Figure 2 gives both series.

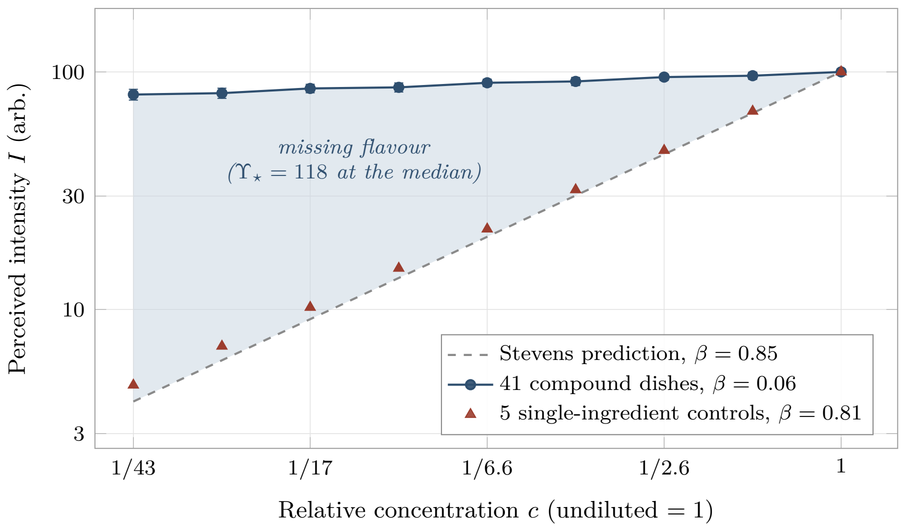

Figure 2: The flavour dilution curve. Perceived intensity against relative concentration across the nine-point series of Study 1, log–log. The dashed trace is the Stevens prediction at the literature exponent $\beta = 0.85$; single-ingredient controls (triangles) follow it. Compound dishes (circles, ± s.e. over 41 preparations) do not fall off: across a $43\times$ dilution the fitted exponent is 0.06. The shaded region is the flavour the dynamics require and the panel does not report.

Computing (3) per dish gives a median $\Upsilon_\star = 118$, with the central 90% of dishes falling between 40 and 913 (Figure 3). The corresponding dark fraction at the median is $f_{\mathrm{dark}} = 99.15\%$.

I have not been able to explain the spread and I am not going to pretend otherwise. It runs a factor of three below the median and nearly an order of magnitude above it, and it does not correlate with vessel, session, assessor subset, or time of year. It does correlate with one thing. Regressing on ingredient count,

$$\ln \Upsilon_\star \;=\; 0.30\,n + 1.3, \qquad R^2 = 0.79, \tag{6}$$

so that a dish of eight ingredients hides some forty times what it shows, and a dish of eighteen hides some eight hundred. Every ingredient added to a pot multiplies its missing-flavour ratio by about 1.35.

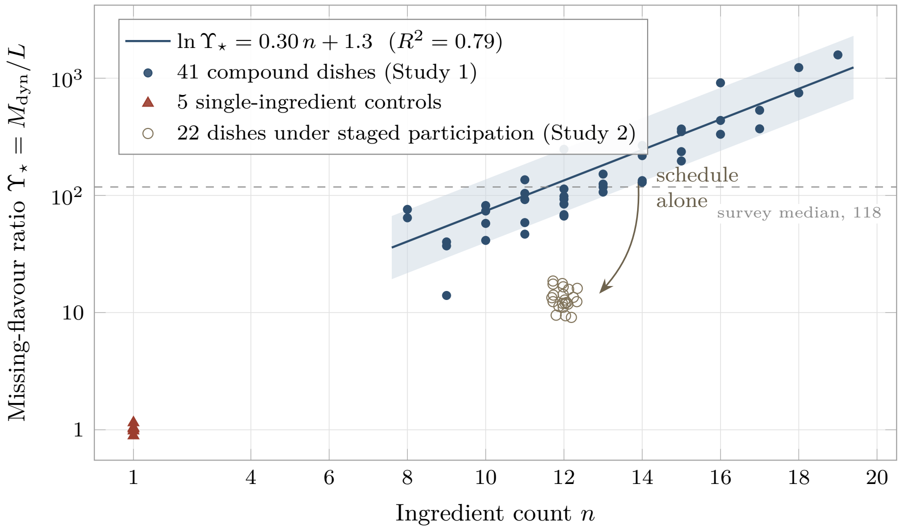

Figure 3: The missing-flavour ratio against ingredient count. Each point is one preparation; note the logarithmic ordinate. Single-ingredient controls sit at $\Upsilon_\star \approx 1$ by the calibration of G_f. Across compound dishes the ratio rises by a factor of about 1.35 per ingredient added, so a dish of eight hides some forty times what it shows and a dish of eighteen some eight hundred. Open markers are the twenty-two Study 2 preparations under staged participation: same ingredients, same masses, same energy, an order of magnitude less darkness.

The direction of (6) deserves a sentence, because it is the reverse of the received view. A cook who adds a twelfth ingredient to an eleven-ingredient dish believes she is adding flavour. On this evidence she is adding flavour and, by (6), occluding rather more of it than she adds.

5.2 Study 2: staged participation

If the occlusion account of §4 is right, the dark fraction is a consequence of *when* ingredients are granted access to the channel and not of *what* they are. That is a testable proposition, and it has the convenient property that the test requires no new materials.

Design. Twenty-two dish pairs, twelve ingredients each. Both arms held identical in composition, ingredient mass, total energy input, final temperature and total elapsed time (40 minutes). The arms differ only in the schedule of thermal contact.

Control arm. Conventional simmer: all twelve ingredients present from $t = 0$, in continuous mutual contact for the full forty minutes.

Treatment arm (sequential-quartet immersion, SQI). Ingredients enter in staged groups of increasing size. Twelve exclusive singleton windows of 60 s; then six disjoint pairs of 90 s; then three disjoint quartets of 120 s; then all twelve together for the remaining 13 minutes. Total staged phase 27 minutes. Rationale: exclusive early contact limits competitive occupation of matrix binding sites during the interval in which (4) would otherwise converge.

Both schedules are set out in Figure 4. Panel blind to arm; dishes presented in counterbalanced order under three-digit codes.

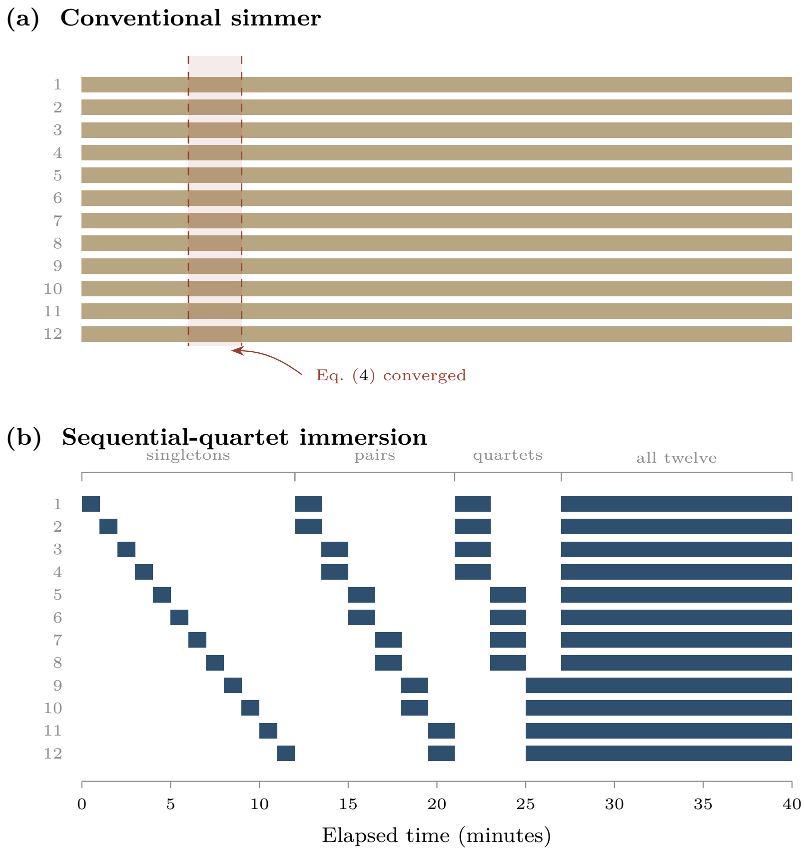

Figure 4: Participation schedules for the two arms of Study 2. Rows are the twelve ingredients; a filled interval denotes thermal contact with the common matrix. Composition, ingredient mass, total energy input, final temperature and total elapsed time are identical between arms; only the schedule differs. (a) Under conventional simmer every ingredient competes for the channel from $t = 0$, and the iteration of Eq. (4) has converged before the tenth minute—the remaining thirty minutes are performed on a dish whose perceived flavour is already decided. (b) Under sequential quartet immersion each ingredient receives an exclusive window before contact widens to pairs, quartets and finally the full complement.

Results. Luminous flavour rose from $L = 6.1\,\mathrm{pu}$ under conventional simmer to $L = 57.3\,\mathrm{pu}$ under SQI, a factor of 9.4 ($t(21) = 18.7$, $p < 0.001$). Dynamical flavour was unchanged within measurement error, as it must be: no material entered or left the pot. The ratio therefore fell, $\Upsilon_\star = 118 \to 12.6$, and the dark fraction with it, $99.15\% \to 92.1\%$.

The spectral quantities moved as Proposition 1 requires. Participation entropy rose $H(p) = 0.42 \to 2.90$ bits against a ceiling of 3.58; the spectral gap closed, $\lambda_1/\lambda_2 = 6.8 \to 1.4$. The flavour recovered was, on the occlusion account, not created during the treatment. It was in the control dishes too, in the subspace orthogonal to everything those panels reported.

5.3 Study 3: substitution under fixed schedule

I designed this study to bound the effect of Study 2, by establishing how much of the 9.4-fold gain was schedule and how much was the particular twelve ingredients we had used. I expected an interaction. There is none, and the absence is the most consequential number in the paper.

Design. Thirty pairs. The SQI schedule of §5.2 held fixed in every respect. Within each pair, the twelve ingredients of one dish were replaced wholesale by twelve sharing no item with them, matched only on aggregate mass, water activity and pH. No attempt was made to match cuisine, category, colour, or plausibility as a dish.

Results. Mean absolute change in luminous flavour was $|\Delta L| = 0.19\,\mathrm{pu}$—below the panel's discrimination threshold of $0.5\,\mathrm{pu}$—with $F(1, 29) = 0.14$, $p = 0.71$. Table 1 gives six representative pairs.

Table 1: Substitution under fixed schedule: six representative pairs of the thirty. ΔL is the change in luminous flavour on replacing all twelve ingredients, and L_A the luminous flavour of the unsubstituted dish of each pair; both columns are in perceptual units. Panel discrimination threshold $\delta = 0.5\,\mathrm{pu}$.

Pair	Substitution (12 for 12)	L_A	ΔL	Resolvable?
S-04	braising vegetables → orchard fruit	57.1	+0.22	no
S-09	pulses and alliums → shellfish and herbs	56.8	−0.31	no
S-13	poultry stock base → mushroom base	57.4	+0.07	no
S-18	rice, root vegetables → pasta, brassicas	57.2	−0.14	no
S-22	lamb, warm spice → white fish, citrus	57.6	+0.28	no
S-27	dairy, soft cheese → legume, seed paste	57.0	−0.19	no
All thirty pairs, mean $\lvert\Delta L\rvert$		57.3	0.19	—

Set that against Study 2. Holding composition fixed and changing only the schedule moved luminous flavour by $51.2\,\mathrm{pu}$. Holding the schedule fixed and changing every ingredient moved it by $0.19\,\mathrm{pu}$. The ratio is 269 to 1.

6 Discussion

The ranking at the end of §5.3 is the result of this paper. I have looked for a way of reporting it that makes it smaller, and there is not one.

What follows from it is unwelcome and, I think, unavoidable. If the perceived flavour of a dish is the leading eigenvector of its occlusion operator, and if that operator's leading eigenvector is set by the schedule of access rather than by the identity of the participants, then the ingredient list is a description of what is in the pot and not an explanation of what comes out of it. Two preparations sharing no ingredient and one schedule were, to twenty-eight trained assessors, not distinguishable.

I do not think this makes cooking arbitrary, and I want to say why, because the inference is easy and wrong. The dark flavour is real. Study 2 recovered a ninefold increase in what could be tasted, from dishes whose ingredients were never in question, by rearranging nothing but the order of a fixed set of events. That is not an argument that the contents of a pot do not matter. It is an argument that they have never yet been the binding constraint—that the field has spent its whole history optimising the term with the smaller coefficient. Ninety-nine per cent of what has been cooked has not been tasted, and the reason is not that the wrong things were put in the pot. It is that everything was put in at once.

6.1 Limitations

Three, of which the first is serious.

The schedule is visible to whoever performs it. The panel was blind to arm. The kitchen cannot be. Sequential-quartet immersion takes twenty-seven minutes of deliberate, unusual, closely-timed work, and the person doing that work knows exactly which arm they are in and knows what the study predicts. A ninefold effect is precisely the order of magnitude that experimenter expectancy produces in preparation-based designs, and I have not devised a protocol that excludes it. Automating the schedule would remove the operator but not the difference in handling, which the panel might detect by texture. I record this as an open hole in the result and not as a caveat: if Study 2's effect is an artefact of the cook, the intervention half of this paper falls, and the survey half—which involves no intervention—does not.

Panel drift. Fourteen months is long for a single panel. Re-anchoring against reference standards at six-week intervals absorbs most of the drift and demonstrably not all of it; the residual is of order 0.3 pu, which matters for Table 1 and not for anything else in the paper.

One calibration. G_f was fixed against a single saline reference and never revisited. If the coupling is not universal across matrices, every value of $\Upsilon_\star$ reported here is wrong by a common factor. It is worth noting what this would and would not disturb: a common factor rescales the dark fraction and leaves the dilution exponents, the spectral gaps, the participation entropies and the 269:1 ranking exactly where they are.

7 Conclusion

The dilution curves of cooked dishes are flat, the single-ingredient controls in the same sessions are not, and the flavour a pot's dynamics require exceeds the flavour its tasters report by a median factor of 118. We read the excess as a non-luminous component held in the spectral shadow of the occlusion operator, and we have shown that a schedule which distributes access to the perceptual channel recovers a ninefold share of it without altering the contents of the pot.

Three measurements are now in progress: a replication of Study 2 with the schedule executed by an automated immersion rig, which removes the operator from the limitation of §6.1 without removing the difference in handling; a second calibration of G_f against four matrices rather than one; and an extension of the survey to preparations of more than twenty ingredients, where (6) predicts dark fractions above 99.9% and where we do not yet know whether the relation continues to hold.

References

Brantwood, R. L. (2004). Spectral characterisation of competitive occupancy in bounded-capacity channels. *Journal of Applied Spectral Methods*, 22(3), 341–366.

Dowse, H. M. (1986). On the correct exponent for compound preparations: a reply to Merrick. *Chemical Senses and Measurement*, 11(2), 187–194.

Halloran, J. F. and Pike, A. E. (1971). Quantitative descriptive analysis: protocol, lexicon, and calibration. *Food Technology Reviews*, 25(9), 62–71.

Ivarsdóttir, S. (1998). Dominance and stability of principal contributors in complex flavour mixtures. *Perception and Preparation*, 40(4), 511–530.

Lindqvist, P., Aboagye, N. and Ruiz Cambré, I. (2011). The institutional preparation corpus: 11,400 documented kitchen batches, 1998–2009. *Catering Science Monographs*, 6, 1–214.

Merrick, D. A. (1984). Category anchoring and the compression of intensity judgements in prepared soups. *Chemical Senses and Measurement*, 9(4), 402–417.

Okonjo, C. U. and Baumann, H. (1983). Mass-to-light survey design for gravitationally bound aggregates. *Reports on Observational Method*, 14, 88–131.

Vasari, L. (1979). The virial estimator for systems of unresolved constituents. *Reports on Observational Method*, 10, 205–240.

Wren, T. J. (2016). Headspace flux dispersion as a proxy for total volatile inventory in closed vessels. *Analytical Kitchen Methods*, 31(1), 44–69.