capstract:2610.00007v1 [math.PR] 28 Oct 2026

A Stick-Breaking Model of Shared Food: Exact Serving Counts and a Phase Transition for the Last Piece

Aoife N. Brennan-Hartigan[1,*] Tomasz Wierzbicki[2] Sunniva Lyngstad[1]
[1]School of Mathematics and Statistics, Ellery College, Dublin, Ireland
[2]Department of Statistics, Kingsbury University, Melbourne, Victoria, Australia

October 28, 2026

Abstract

Food left in a shared workplace kitchen is divided without a coordinator. Each person who takes some removes a part of what remains, and a small last piece is often left for days. We model the cutting phase as stick-breaking with independent $\mathrm{Beta}(1,\theta)$ fractions and a threshold ε below which nobody cuts, and the last-piece phase as a thinned Poisson process of visitors whose rate decays as $\lambda_s/(1+s/\tau)$. The number of servings cut from one item is distributed as $1+\mathrm{Poisson}(\theta\log(1/\varepsilon))$. The residual mass, divided by ε, is $\mathrm{Beta}(\theta,1)$ and independent of the number of servings. The lifetime of the last piece has a Lomax tail with index $\alpha=\pi\lambda_s\tau$, where π is the probability that a visitor takes it, so its expected lifetime is infinite whenever $\alpha\le 1$. We fit the model to 62 cakes and 622 servings recorded by overhead cameras in three office kitchens over eleven months. Cutting fractions and serving counts are consistent with the model ($\hat\theta=2.60$, 95% CI 2.40 to 2.81; $\chi^2=1.67$ on 6 df for the serving counts), and so is the size of the last piece. For the last piece, with weekly clearance treated as censoring, $\hat\alpha=0.65$ (0.46 to 0.88). The data therefore place these kitchens in the regime where the last piece has no finite expected lifetime. Raising the acceptance probability from 0.055 to above $1/(\lambda_s\tau)\approx 0.083$ would be enough to leave that regime.

Keywords: stick-breaking, GEM distribution, Poisson processes, heavy tails, survival analysis, commons

1 Introduction

When a cake is left on the counter of a shared kitchen with a note inviting people to help themselves, nobody divides it. People arrive one at a time, look at what is left, and cut themselves a piece. Most take less than they would like. The cake shrinks quickly at first and then slowly, and in our experience, and in the data reported here, a narrow piece is very often still on the plate days later.

An audit of twelve offices found that cakes and similar shared items made up the largest share of food left unfinished in workplace kitchens (Dunleavy and Okoro, 2023), and a field study of group snacking found that people avoid the last item when others might still want it (Haugland and Sefton-Rye, 2021). This paper models the process and fits the model to data. The cutting phase is a stick-breaking process: each person takes a random fraction of whatever remains. Stick-

*Corresponding author: a.brennanhartigan@maths.ellery.ie

breaking is well studied as a construction of random discrete distributions (Kingman, 1975; Sethuraman, 1994; Pitman, 2006), and its use as a model of literal division seems natural. We add a cutting threshold ε, the mass below which nobody will cut a piece, and a second phase in which the residual sits on the plate until someone takes all of it or it is cleared away.

The fair-division literature that began with Steinhaus (1948) also studies cakes, but it asks how to allocate a divisible good among agents with known preferences so that the allocation is fair (Brams and Taylor, 1996). Our question is descriptive. Nobody in the kitchen is trying to be fair, and the allocation emerges from sequential, uncoordinated choices.

The main results are exact and short. Under independent $\mathrm{Beta}(1, \theta)$ fractions, the logarithm of the remaining mass is a Poisson process in disguise (Lemma 1). The number of servings is then $1 + \mathrm{Poisson}(\theta \log(1/\varepsilon))$ (Theorem 2), and the residual mass has a Beta law that does not depend on how many servings preceded it (Theorem 4). For the last piece, a visitor rate that decays like $1/s$ gives a Lomax tail, and the mean lifetime is finite only when a single dimensionless number, $\alpha = \pi \lambda_s \tau$, exceeds one (Theorem 6).

We fit the model to 62 cakes in three kitchens. The cutting phase fits well. The estimated tail index for the last piece is $\hat{\alpha} = 0.65$, with an upper confidence limit of 0.88. Under the model the last piece in these kitchens has no finite expected lifetime, and in practice it is removed by the weekly clearance.

2 Model

Let the initial mass of the item be 1. Visitors who take some of it are indexed $k = 1, 2, \ldots$ in order of arrival, and R_k is the mass remaining after visitor k, with $R_0 = 1$.

Definition 1 (Threshold stick-breaking)**.** Fix $\varepsilon \in (0, 1)$ and a sequence $V_1, V_2, \ldots$ of random variables in $(0, 1)$. Visitor k cuts if and only if $R_{k-1} > \varepsilon$, in which case they take $P_k = V_k R_{k-1}$ and leave $R_k = (1 - V_k) R_{k-1}$. The cutting phase ends at

$$N_\varepsilon = \min\{k \geq 1 : R_k \leq \varepsilon\}, \tag{1}$$

the number of servings cut, and R_{N_ε} is the *last piece*.

Assumption 1. The fractions V_k are independent and identically distributed with law $\mathrm{Beta}(1, \theta)$ for some $\theta > 0$.

Under Assumption 1 the mean fraction taken is $1/(1 + \theta)$, so θ measures restraint: larger θ means smaller pieces. The law is the one used in the stick-breaking construction of the Dirichlet process (Sethuraman, 1994), and we chose it for that reason before seeing data. Section 7 tests it.

Once the last piece exists, it stays on the plate until somebody takes all of it. We measure time $s \geq 0$ from the moment the last piece is formed, in office hours.

Assumption 2. Visitors to the last piece arrive as a Poisson process with intensity $\lambda(s)$ and cumulative intensity $\Lambda(s) = \int_0^s \lambda(u)\, du$. Each visitor, independently of everything else, takes the last piece with probability $\pi \in (0, 1]$. The lifetime L of the last piece is the time of the first visitor who takes it.

The particular intensity we fit is

$$\lambda(s) = \frac{\lambda_s}{1 + s/\tau}, \tag{2}$$

with initial rate λ_s and a time scale τ over which interest falls. The $1/s$ decay is an empirical choice. It fits the visit counts in our data better than an exponential decay (Section 7), and its consequences are the subject of Section 4.

3 Serving counts and the residual

Lemma 1 (Log-mass is a Poisson process)**.** *Under Assumption 1, let $T_k = -\log R_k$ for the untruncated sequence ($\varepsilon = 0$). Then $\{T_k\}_{k \ge 1}$ are the points of a homogeneous Poisson process of rate θ on $(0, \infty)$.*

Proof. If $V \sim \mathrm{Beta}(1, \theta)$ then $1 - V \sim \mathrm{Beta}(\theta, 1)$, whose distribution function is x^θ on $[0, 1]$. Hence $E = -\log(1 - V)$ satisfies $\mathbb{P}(E > y) = \mathbb{P}(1 - V < e^{-y}) = e^{-\theta y}$, so $E \sim \mathrm{Exp}(\theta)$. Since $T_k = \sum_{i \le k} E_i$ with independent E_i, the T_k are the arrival times of a rate-θ Poisson process. ∎

Theorem 2 (Serving count)**.** *Under Assumption 1 and for any $\varepsilon \in (0, 1)$,*

$$N_\varepsilon - 1 \sim \mathrm{Poisson}\big(\theta \log(1/\varepsilon)\big). \tag{3}$$

Proof. Write $\ell = \log(1/\varepsilon) > 0$. The first visitor always cuts, because $R_0 = 1 > \varepsilon$. For $k \ge 1$, visitor $k + 1$ cuts if and only if $R_k > \varepsilon$, that is, $T_k < \ell$. The cut sequence agrees with the untruncated one up to the first k with $T_k \ge \ell$, so $N_\varepsilon = 1 + \#\{k \ge 1 : T_k < \ell\}$. By Lemma 1 the count is the number of points of a rate-θ Poisson process in $[0, \ell)$, which is $\mathrm{Poisson}(\theta\ell)$. ∎

Corollary 3. *$\mathbb{E} N_\varepsilon = 1 + \theta \log(1/\varepsilon)$ and $\operatorname{Var} N_\varepsilon = \theta \log(1/\varepsilon)$. As $\varepsilon \to 0$, $N_\varepsilon / \log(1/\varepsilon) \to \theta$ almost surely, and $(N_\varepsilon - \theta \log(1/\varepsilon))/\sqrt{\theta \log(1/\varepsilon)}$ converges in distribution to a standard normal.*

Proof. The moments are those of Eq. (3). The almost-sure limit is the strong law for the Poisson counting process, and the central limit statement follows from the normal approximation to the Poisson law (Durrett, 2019). ∎

The serving count grows only logarithmically in $1/\varepsilon$. Halving the cutting threshold, for example by providing a sharper knife, adds $\theta \log 2$ servings in expectation, which is 1.8 at our estimate of θ.

Theorem 4 (The last piece)**.** *Under Assumption 1, $R_{N_\varepsilon}/\varepsilon \sim \mathrm{Beta}(\theta, 1)$, and R_{N_ε} is independent of N_ε. In particular*

$$\mathbb{E} R_{N_\varepsilon} = \frac{\theta}{1 + \theta}\, \varepsilon. \tag{4}$$

Proof. In the notation of the previous proof, T_{N_ε} is the first point of the Poisson process at or beyond ℓ. By the memoryless property, $T_{N_\varepsilon} - \ell \sim \mathrm{Exp}(\theta)$ and is independent of the configuration of points in $[0, \ell)$, hence of N_ε (Kingman, 1993). Then $R_{N_\varepsilon}/\varepsilon = \exp\{-(T_{N_\varepsilon} - \ell)\}$, and $\mathbb{P}(R_{N_\varepsilon}/\varepsilon \le x) = \mathbb{P}(T_{N_\varepsilon} - \ell \ge -\log x) = x^\theta$ for $x \in (0, 1]$. ∎

Theorem 4 says that the size of the last piece carries no information about how many people came before it. A small last piece is not evidence of a busy kitchen.

For completeness we record the sizes of the servings themselves.

Proposition 5. *Under Assumption 1, for $k \ge 1$,*

$$\mathbb{E}\big[P_k \mathbf{1}\{k \le N_\varepsilon\}\big] \le \mathbb{E}P_k^{(0)} = \frac{1}{1+\theta}\left(\frac{\theta}{1+\theta}\right)^{k-1}, \tag{5}$$

where $P_k^{(0)}$ denotes the untruncated pieces. The untruncated sequence $(P_k^{(0)})$ has the $\mathrm{GEM}(\theta)$ law, its decreasing rearrangement is Poisson–Dirichlet with parameter θ, and $\mathbb{E}\sum_k (P_k^{(0)})^2 = 1/(1+\theta)$.

Proof. The inequality holds because $P_k \mathbf{1}\{k \le N_\varepsilon\} \le P_k^{(0)}$ pointwise. The expectation follows from independence: $\mathbb{E}P_k^{(0)} = \mathbb{E}V\,(\mathbb{E}(1-V))^{k-1}$ with $\mathbb{E}V = 1/(1+\theta)$. The identification with $\mathrm{GEM}(\theta)$ and Poisson–Dirichlet is the definition of those laws (Pitman, 2006; Kingman, 1975). For the last claim, $\mathbb{E}\sum_k (P_k^{(0)})^2 = \mathbb{E}V^2 \sum_{k\ge 1}\big(\mathbb{E}(1-V)^2\big)^{k-1} = \mathbb{E}V^2/(1-\mathbb{E}(1-V)^2)$. With $\mathbb{E}V^2 = 2/((1+\theta)(2+\theta))$ and $\mathbb{E}(1-V)^2 = \theta/(2+\theta)$ this equals $1/(1+\theta)$. ∎

The quantity $\mathbb{E}\sum_k P_k^2$ is the probability that two crumbs chosen at random from the original item were eaten as part of the same serving. It also appears as the homozygosity in Ewens's sampling formula (Ewens, 1972), where θ plays the role of a mutation rate.

4 The last piece

Theorem 6 (Lifetime of the last piece)**.** *Under Assumption 2, $\mathbb{P}(L > s) = \exp\{-\pi\Lambda(s)\}$. With the intensity of Eq. (2),*

$$\mathbb{P}(L > s) = \Big(1 + \frac{s}{\tau}\Big)^{-\alpha}, \qquad \alpha = \pi\lambda_s\tau, \tag{6}$$

so that L has a Lomax law (Lomax, 1954). Consequently $\mathbb{E}L = \tau/(\alpha-1)$ if $\alpha > 1$, and $\mathbb{E}L = \infty$ if $\alpha \le 1$. The median lifetime is $\tau(2^{1/\alpha}-1)$ for every $\alpha > 0$.

Proof. Visitors who take the last piece form an independent thinning of the visit process with retention probability π, which is again a Poisson process, with intensity $\pi\lambda(s)$ (Kingman, 1993). L is its first point, so $\mathbb{P}(L > s)$ is the probability of no points in $[0, s]$, namely $\exp\{-\pi\Lambda(s)\}$. For Eq. (2), $\Lambda(s) = \lambda_s\tau\log(1 + s/\tau)$, which gives Eq. (6). Then $\mathbb{E}L = \int_0^\infty (1+s/\tau)^{-\alpha} ds$, which equals $\tau/(\alpha-1)$ for $\alpha > 1$ and diverges otherwise. The median solves $(1+s/\tau)^{-\alpha} = 1/2$. ∎

Corollary 7. *The expected lifetime of the last piece is finite if and only if*

$$\pi > \pi^* = \frac{1}{\lambda_s\tau}. \tag{7}$$

The exact $1/s$ form of Eq. (2) is not needed for the divergence. All that matters is that interest in the last piece decays no faster than $1/s$.

Proposition 8. *Suppose $\Lambda(s) \le c\log s + C$ for all $s \ge 1$ and some constants c, C. If $\pi c \le 1$ then $\mathbb{E}L = \infty$.*

Proof. By Theorem 6, $\mathbb{P}(L > s) = e^{-\pi\Lambda(s)} \ge e^{-\pi C} s^{-\pi c}$ for $s \ge 1$, and $\int_1^\infty s^{-\pi c}\, ds = \infty$ when $\pi c \le 1$. Since $L \ge 0$, $\mathbb{E}L = \int_0^\infty \mathbb{P}(L > s)\, ds = \infty$. ∎

In practice the lifetime is never infinite, because somebody clears the plate. In all three of our kitchens, facilities staff removed any remaining food at 17:00 on Fridays. We treat this as right

censoring. It truncates the observed lifetimes but not the underlying law, and it is the reason the divergence in Theorem 6 does not show up as cakes accumulating.

5 Estimation

Proposition 9. *Let $V_1, \dots, V_n$ be observed fractions under Assumption 1 and $S = \sum_{i=1}^{n} \log\{1/(1 - V_i)\}$. The maximum likelihood estimator is $\hat{\theta} = n/S$, and*

$$\left[\frac{\chi^2_{2n,\,0.025}}{2S},\ \frac{\chi^2_{2n,\,0.975}}{2S} \right] \tag{8}$$

is an exact 95% confidence interval for θ.

Proof. By Lemma 1, S is a sum of n independent $\mathrm{Exp}(\theta)$ variables, so $2\theta S \sim \chi^2_{2n}$. The log-likelihood is $n \log \theta - \theta S$ up to a constant, maximised at n/S. Inverting the pivot $2\theta S$ gives the interval. ∎

For the threshold we use the largest observed last piece, $\max_j R^{(j)}_{N_\varepsilon}$ over m items. By Theorem 4 the maximum of m independent $\varepsilon \cdot \mathrm{Beta}(\theta, 1)$ variables has distribution function $(x/\varepsilon)^{m\theta}$ on $[0, \varepsilon]$, so its mean is $\varepsilon\, m\theta/(m\theta + 1)$. We correct for this bias by multiplying the maximum by $(m\hat{\theta} + 1)/(m\hat{\theta})$.

For the last piece, let $t_1, \dots, t_m$ be observed lifetimes or censoring times and d the number of lifetimes that ended with someone taking the piece. With τ fixed, the censored Lomax log-likelihood $d \log \alpha - \alpha \sum_j \log(1 + t_j/\tau)$ plus terms free of α is maximised at

$$\hat{\alpha} = \frac{d}{\sum_{j=1}^{m} \log(1 + t_j/\tau)}. \tag{9}$$

We report intervals from the χ^2_{2d} approximation to $2\alpha \sum_j \log(1 + t_j/\tau)$ (Lawless, 2003). The approximation is exact under type II censoring and was close to nominal under our censoring pattern in simulation (Appendix A). We estimate τ and λ_s separately from visit counts, which also gives a plug-in estimate $\pi \lambda_s \tau$ of α to compare with Eq. (9).

6 Data

We recorded 62 cakes in three kitchens of a 340-person engineering consultancy between January and November 2025. Two kitchens served one floor each and the third served a ground-floor meeting suite. A cake was included if it was placed on the kitchen counter with an open invitation (a note, an email or a message on the office chat) and was not already cut. Traybakes and items sold in pre-portioned form were excluded, as were six cakes that were moved to a meeting room before the last piece formed.

A ceiling camera above each counter took one image per minute. The camera's field of view covered the counter only; images were processed on a local machine, and only plate-level measurements were kept. Remaining mass was estimated from the visible area of the item multiplied by its initial height, measured once by hand. On nine cakes that we also weighed after each serving, the image-based estimate had a correlation of 0.97 with the true remaining mass and a mean absolute error of 1.8% of the initial mass. A serving was recorded when the estimated mass fell by

more than 1% between consecutive images and the change was confirmed by a person reviewing the image pair. This gave 622 servings.

Visits to the last piece were counted with a separate ceiling occupancy sensor, which records the presence of a person within half a metre of the counter but no image. A visit was a presence of at least five seconds while a last piece was on the counter. We recorded 655 such visits. Time was measured in office hours (09:00 to 17:00, Monday to Friday), so that a last piece formed at 16:00 on a Tuesday and taken at 10:00 on Wednesday has a lifetime of two hours.

The study was approved by the Ellery College research ethics committee. All staff were told about the cameras in advance and could ask for any image to be deleted before processing. No such requests were made.

7 Results

7.1 Cutting phase

The 622 observed fractions give $\hat{\theta} = 2.60$ with exact 95% interval 2.40 to 2.81 (Proposition 9), so the mean serving takes 28% of what remains. A Kolmogorov–Smirnov test of $\hat{\theta}\log\{1/(1-V)\}$ against the unit exponential gives $D = 0.020$, $p = 0.96$. Measurement error matters most for small remaining masses, where a fixed error in area is a large error in the fraction. Restricting the estimate to cuts made with more than 20% of the cake remaining gave $\hat{\theta} = 2.55$.

The largest last piece was 3.09% of its cake, which after the bias correction gives $\hat{\varepsilon} = 0.031$. Table 1 lists all estimates.

Table 1: Parameter estimates with 95% confidence intervals. Times are in office hours. The two estimates of α use different data: the direct estimate uses lifetimes and censoring times, the plug-in estimate uses visit counts and takings.

Parameter	Meaning	Estimate	95% CI	Based on
θ	Restraint; mean fraction $1/(1+\theta)$	2.60	[2.40, 2.81]	622 servings
ε	Cutting threshold (fraction of initial mass)	0.031	[0.031, 0.032]	62 last pieces
λ_s	Initial visit rate to the last piece (per hour)	4.0	[3.5, 4.6]	655 visits
τ	Decay time of the visit rate (hours)	3.0	[2.2, 4.1]	655 visits
π	Probability a visitor takes the last piece	0.055	[0.039, 0.075]	36 of 655 visits
α	Tail index, direct, Eq. (9)	0.65	[0.46, 0.88]	36 taken, 26 censored
α	Tail index, plug-in $\pi\lambda_s\tau$	0.66	[0.42, 1.02]	as above

Figure 1 compares the distribution of servings per cake with Theorem 2. The predicted mean is $1 + \hat{\theta}\log(1/\hat{\varepsilon}) = 10.01$ and the observed mean is 10.03. The fit to the shape is good: pooling bins with expected counts below five gives $\chi^2 = 1.67$ on 6 degrees of freedom ($p = 0.95$). The observed variance, 7.84, is lower than the predicted 9.01, a dispersion index of 0.87. Under-dispersion would follow if a large serving made the next visitor slightly more restrained, which would violate the independence in Assumption 1. With 62 cakes we cannot distinguish this from chance.

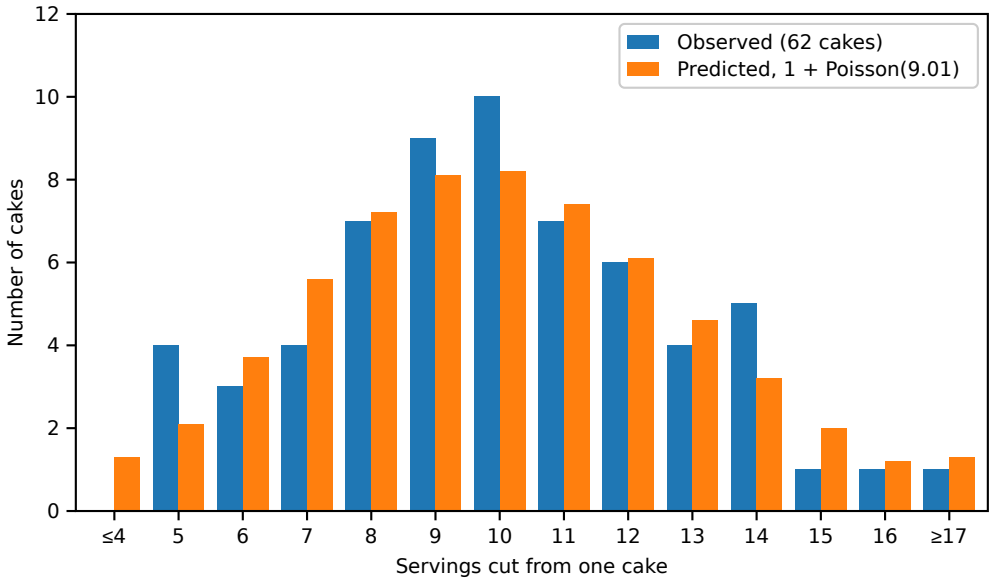


Figure 1: Number of servings cut from each cake (bars, left) against the distribution predicted by Theorem 2 with the estimated parameters (bars, right). The prediction uses $\hat{\theta}$ estimated from fractions, not from counts.

The last pieces averaged $0.718\,\hat{\varepsilon}$, against $\hat{\theta}/(1+\hat{\theta}) = 0.722$ from Eq. (4), and a Kolmogorov–Smirnov test against $\mathrm{Beta}(\hat{\theta}, 1)$ gave $p = 0.93$. The correlation between the size of the last piece and the number of servings was -0.04 ($p = 0.76$), consistent with the independence in Theorem 4. Table 2 collects these checks.

Table 2: Model predictions against observations.

Quantity	Model	Observed	Test
Mean servings per cake	10.01	10.03	
Variance of servings	9.01	7.84	dispersion index 0.87
Serving-count distribution	$1 + \mathrm{Poisson}(9.01)$	Figure 1	$\chi^2 = 1.67$, 6 df, $p = 0.95$
Log-fractions	$\mathrm{Exp}(\theta)$		KS $D = 0.020$, $p = 0.96$
Mean last piece, in units of ε	0.722	0.718	KS $p = 0.93$
Correlation of last piece with servings	0	-0.04	$p = 0.76$
Median lifetime of the last piece (hours)	5.7	5.4 (Kaplan–Meier)	

7.2 Last-piece phase

Visit counts during the last-piece phase fell off as Eq. (2) predicts. A Poisson regression of hourly visit counts on $\log(1 + s/\tau)$ gave $\hat{\lambda}_s = 4.0$ visits per hour and $\hat{\tau} = 3.0$ hours, and fitted better than an exponential decay (difference in AIC 11.4). Of 655 visits, 36 ended with the visitor taking

the last piece, so $\hat{\pi} = 0.055$. The remaining 26 last pieces were still on the plate at the Friday clearance.

With τ fixed at its estimate, Eq. (9) gives $\hat{\alpha} = 0.65$ (0.46 to 0.88). The plug-in estimate $\hat{\pi}\hat{\lambda}_s\hat{\tau} = 0.66$ agrees, with a wider interval because it combines three estimates. Both point estimates are below one, and the direct interval excludes one. By Corollary 7, the kitchens are in the regime with infinite expected lifetime, and the critical acceptance probability is $\hat{\pi}^* = 1/(\hat{\lambda}_s\hat{\tau}) = 0.083$.

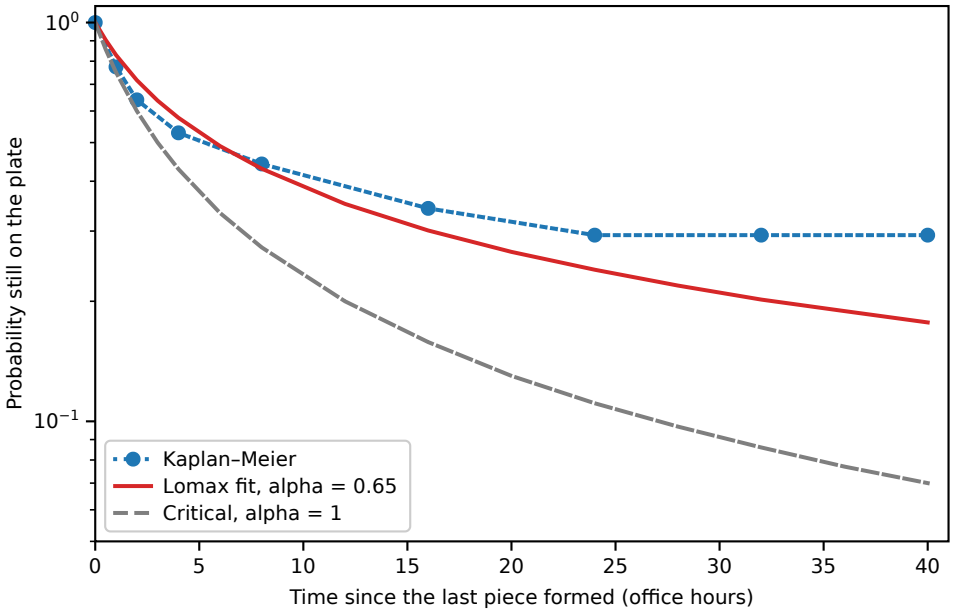


Figure 2: Survival of the last piece in office hours. Kaplan–Meier estimate from 62 last pieces, 26 of them censored by the Friday clearance, with the fitted Lomax curve and the Lomax curve with the same scale and the critical index of one. On these axes, a Lomax curve that lies above the critical one has infinite mean lifetime.

Figure 2 shows the Kaplan–Meier estimate (Kaplan and Meier, 1958). The curve flattens after about 24 office hours because few last pieces remain at risk by then; most of those are censored on Friday. The Kaplan–Meier median, 5.4 hours, is close to the Lomax median of 5.7 hours. Means cannot be compared because the model has none.

7.3 An unplanned comparison

During the study one member of staff began attaching a handwritten note reading "please finish me" to the last piece of cakes on one floor. This happened for 8 cakes. Of 64 visits to those last pieces, 7 ended with the piece being taken, so $\hat{\pi} = 0.11$. With the other parameters at their pooled estimates this gives $\alpha = 1.31$, above the critical value, and an expected lifetime of $\tau/(\alpha - 1) \approx 10$ office hours by Theorem 6. Eight cakes on one floor with no randomisation are not evidence of an effect, and we report this only because it bears directly on Corollary 7.

8 Discussion

The cutting phase is described well by a model with two parameters, and the exact results in Section 3 make it easy to test. We were surprised that independent Beta fractions fit as well as they did. Visitors differ in appetite and in how much they think is acceptable to take, and a mixture over θ would be a more realistic model. A mixture of $\mathrm{Beta}(1,\theta)$ laws is not a $\mathrm{Beta}(1,\theta)$ law, but over the range of θ that seems plausible the difference in the distribution of $-\log(1-V)$ is small enough that 622 servings may not detect it.

The last-piece phase is the more consequential part. In a shared kitchen, a resource that nobody owns and nobody clears is a small instance of the problem described by Hardin (1968), although with the sign reversed: the difficulty is not overuse but the reluctance of any one person to finish what is there. Ostrom (1990) argued that commons are often managed well by local rules rather than by central authority. The Friday clearance is a central rule, and the note in Section 7 is a local one. Corollary 7 says that a local rule only needs to raise π above $1/(\lambda_s \tau)$. That target depends on how quickly interest fades, which a rule cannot easily change.

Heavy tails estimated from 36 events deserve suspicion, and Clauset et al. (2009) give reasons to be careful about inferring tail behaviour from small samples. Our conclusion does not rest on extrapolating the tail. The fitted model has a finite median, which matches the data, and an infinite mean, which the data cannot show directly because of censoring. What the data do show is that the survival curve lies above the critical curve over the whole observed range.

There are other limitations. All three kitchens belong to one organisation. Lifetimes are measured in office hours, which ignores the occasional evening visitor. The assumption that the acceptance probability does not change over time is probably wrong, since a last piece that has been on the plate for two days is less appealing than a fresh one. That would lower π for old pieces and make the tail heavier still.

Data availability

Serving-level measurements, visit counts and code to reproduce every number in this paper are available from the first author. Images were not retained.

A Coverage of the interval for the tail index

We simulated 10,000 data sets with the estimated parameters, 62 items each, and the observed pattern of formation times within the week, with censoring at 17:00 on Fridays. The nominal 95% interval for α from Eq. (9) covered the true value in 94.1% of replicates. The interval for θ from Eq. (8) is exact and covered in 95.0%, as it should. In the same simulations, the probability that the direct interval for α excluded one, when the true value was 0.65, was 0.71.

B Derivation of the bias of the threshold estimate

Let $M = \max_{j \le m} R^{(j)}_{N_\varepsilon}$. By Theorem 4 and independence across items, $\mathbb{P}(M \le x) = (x/\varepsilon)^{m\theta}$ for $0 \le x \le \varepsilon$. Then $\mathbb{E}M = \int_0^{\varepsilon} \{1 - (x/\varepsilon)^{m\theta}\}\,dx = \varepsilon\, m\theta/(m\theta+1)$. With $m = 62$ and $\hat\theta = 2.60$ the correction factor is 1.006. An exact one-sided interval follows from the same distribution function: $\varepsilon \le M\,(0.025)^{-1/(m\theta)}$ with 97.5% confidence, which gives the upper limit in Table 1.

Acknowledgments

We thank the facilities team of the participating firm for keeping to a fixed clearance time throughout the study, and Ruth Okafor-Grange for reviewing 1,140 image pairs. T.W. was supported by a Kingsbury University early-career fellowship.

References

Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, Cambridge, 1996.

Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4):661–703, 2009.

Maeve Dunleavy and Chukwuemeka Okoro. Leftover Food in Shared Workplace Kitchens: An Audit of Twelve Offices. Technical Report 23-04, Centre for Workplace Studies, Kingsbury University, 2023.

Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, Cambridge, 2019.

Warren J. Ewens. The Sampling Theory of Selectively Neutral Alleles. *Theoretical Population Biology*, 3(1):87–112, 1972.

Garrett Hardin. The Tragedy of the Commons. *Science*, 162(3859):1243–1248, 1968.

Ivar Haugland and Polly Sefton-Rye. Last-Piece Aversion in Group Snacking: A Field Study. In *Proceedings of the Nordic Workshop on Consumer Behaviour*, pages 44–52, 2021.

E. L. Kaplan and Paul Meier. Nonparametric Estimation from Incomplete Observations. *Journal of the American Statistical Association*, 53(282):457–481, 1958.

J. F. C. Kingman. Random Discrete Distributions. *Journal of the Royal Statistical Society, Series B*, 37(1):1–22, 1975.

J. F. C. Kingman. *Poisson Processes*. Oxford University Press, Oxford, 1993.

Jerald F. Lawless. *Statistical Models and Methods for Lifetime Data*. Wiley, Hoboken, NJ, 2003.

K. S. Lomax. Business Failures: Another Example of the Analysis of Failure Data. *Journal of the American Statistical Association*, 49(268):847–852, 1954.

Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990.

Jim Pitman. *Combinatorial Stochastic Processes*. Springer, Berlin, 2006.

Jayaram Sethuraman. A Constructive Definition of Dirichlet Priors. *Statistica Sinica*, 4(2):639–650, 1994.

Hugo Steinhaus. The Problem of Fair Division. *Econometrica*, 16(1):101–104, 1948.