capstract:2611.00063v1 [cond-mat.stat-mech] 20 Nov 2026

Room allocation with fragmentable access credentials

Iona V. Selnar
Institute for Transport Theory, Sintermere College, United Kingdom
(Dated: November 20, 2026)

We study a room scheduler in which each unclaimed portion of catering food carries a reservation credential. A guest who subdivides a portion leaves the credential with the remainder. At scheduled departure, any credential detected by mass prevents a new booking until a room inspection. We derive the first-passage law for a single credential with mass-dependent division and retirement rates, and use it to compute the frequency and duration of holds on vacant rooms. When the retirement rate vanishes faster than the division rate, the credential can survive at vanishing mass. With a finite cutting floor and four 120 g starting portions, changing the sensing threshold from 1 g to 0.5 g raises the probability of holding a room at departure from 0.093 to 0.988 in the faster of two cutting protocols. The choice of threshold reverses the effect of cutting rate on room availability. Simulated first-passage and room-renewal histories agree with the exact model.

I. INTRODUCTION

We consider a room-allocation system in which the credentials are issued to pizza. Each plate arrives with one portion and one credential. A guest who takes the entire portion retires the credential. A guest who takes part leaves the credential with the remainder. At the end of a booking the sensor checks for outstanding credentials. If it detects one, the room remains reserved for the departing party until a facilities inspection.

There are two relevant clocks. The first runs while guests use the room and can take portions. The second starts at the scheduled end of the booking, when the guests depart. The food remaining then determines whether the second clock has to run at all. This arrangement assigns a cost in vacant room time to a process that can remove food almost completely without retiring its credential.

Group-dining studies distinguish the last unclaimed piece from the pieces that preceded it [1]. We use that distinction to specify separate rates for taking part and taking all. In binary fragmentation both daughters ordinarily remain in the system [2, 3]. Here the daughter served leaves immediately. The other retains the entire credential. Shattering can bring mass to zero size in finite time through increasingly rapid breakups [4]. In the present process mass approaches zero during an indefinitely long sequence of cuts, with a nonzero probability that the credential survives. Aggregating several such credentials raises the probability of a hold even if their combined mass is a small fraction of the delivery.

II. FRAGMENTATION AND CREDENTIAL RETIREMENT

Time starts with one slice of mass m_0 on each plate. Let n count partial servings at a plate. For equal cuts the remainder has mass $m_n = m_0 2^{-n}$. A partial serving removes one half and advances n by one, leaving one live credential. Taking the whole remainder sends the process to the absorbing state E. The number of credentials is integer-valued throughout a cut.

While the room is in use, two independent exponential clocks control the events on each plate. Their rates are

$$c_n = c\,2^{-\gamma n}, \qquad u_n = u\,2^{-\beta n}, \tag{1}$$

where c and u have units of inverse time, $u > 0$, and $\beta, \gamma \geq 0$. The exponents describe the separate responses to shrinking mass. A whole-claim exponent of $\beta = 1$ gives a 1 g whole claim one hundredth of the rate of a 100 g whole claim. A positive γ models the increasing difficulty of dividing a small piece. The rates are model inputs.

A plate permits at most K cuts. Thus $c_K = 0$, while the whole-claim clock keeps running at u_K. With guests present indefinitely and K finite, the only stationary distribution is concentrated at E. We also study an unbounded chain, with positive c_n at every finite n. Its total event rate is at most $c + u$. A trajectory with infinitely many cuts takes infinite time.

The booking ends at time D. All guests leave at D and both serving clocks stop. A tared scale detects a live credential if the residual mass on its plate is at least ϵ. The next booking can begin at once if every plate passes the release check. Otherwise a facilities attendant removes the residual portions and releases the room at the first inspection after departure. The post-departure inspection time is exponential, with rate b.

III. CLEARANCE IN THE UNBOUNDED CHAIN

Write $\Delta = \beta - \gamma$, $r = u/c$ and $q = 2^{-\Delta}$ for $c > 0$. The probability that the cut clock rings first at state n is $(1 + r q^n)^{-1}$. The probability S_n that the process makes at least n cuts is therefore

$$S_n = \prod_{j=0}^{n-1} (1 + r q^j)^{-1}, \qquad S_0 = 1. \tag{2}$$

At a finite floor, S_K is the probability of reaching it. For the unbounded chain the limit S_∞ is the probability of never clearing the plate. For $\Delta \le 0$, $S_n \le (1+r)^{-n}$, and a whole claim occurs almost surely. For $\Delta > 0$, the sum of $\log(1 + r q^j)$ converges, yielding $S_\infty > 0$. The nonzero never-cleared fraction appears only on the side where whole claims become rare faster than partial servings.

For any infinitely cut trajectory m_n tends to zero while the occupied-state indicator stays at one. Consequently the expected residual mass tends to zero in both regimes, including the regime where the occupancy probability tends to S_∞. The probability distribution on the occupied branch moves to arbitrarily high n. There is no normalizable stationary distribution on its finite states. The transition shown in Fig. 1 concerns a first-passage probability.

The never-cleared fraction has an essential singularity at the boundary. Setting $h = \Delta \ln 2$ and applying a Riemann sum to Eq. (2) gives

$$\lim_{\Delta \downarrow 0} h \log S_\infty = -\int_0^\infty \log(1 + r e^{-x})\, dx. \tag{3}$$

Thus a positive Δ can give a very small S_∞. The absence of indefinitely occupied plates in a small ensemble would give weak evidence about the sign of Δ.

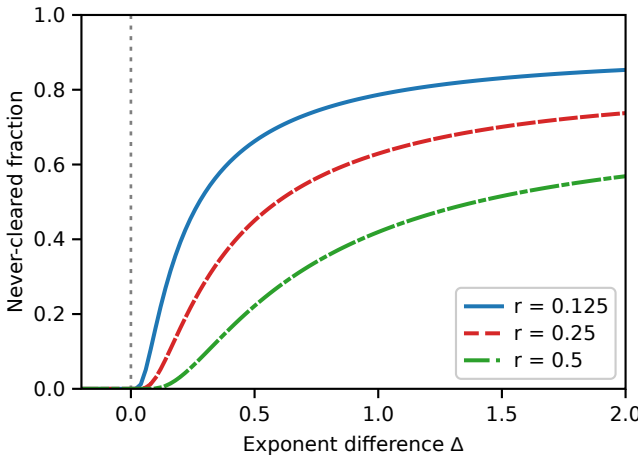


FIG. 1. Probability of never reaching the empty state when cutting is possible at every positive mass. The dotted line marks the equal-exponent boundary.

IV. CUTTING FLOOR AND THE MEAN CLEARANCE TIME

A trajectory reaching n waits an average time $(c_n + u_n)^{-1}$ there. Summing the stage residence times weighted by their probabilities of being visited gives the finite-floor mean

$$\tau_K = \sum_{n=0}^{K-1} \frac{S_n}{c_n + u_n} + \frac{S_K}{u}\, 2^{\beta K}. \tag{4}$$

These contributions separate the time before the smallest permitted remnant is formed from the time spent waiting for someone to take it. When $\Delta > 0$, $S_K \to S_\infty$ and the pre-floor sum grows at most as $2^{\gamma K}$ for $\gamma > 0$, or linearly in K for $\gamma = 0$. The terminal term dominates, and $\tau_K \sim S_\infty 2^{\beta K}/u$. The clearance probability is one for every finite K. Taking the long-time limit first therefore gives zero occupied fraction. Removing the floor first gives S_∞.

There is a second divergence on the line $\beta = \gamma$. Write $z = (1+r)^{-1}$ and $a = z 2^\gamma$. On this line $S_n = z^n$ and the pre-floor sum in Eq. (4) is a geometric series in a with prefactor $(c+u)^{-1}$. The floor contribution is a^K/u. The unbounded mean is finite for $a < 1$. With a floor, the mean grows linearly in K at $a = 1$ and exponentially for $a > 1$. The probability of eventual clearance is one on both sides of this mean-lifetime threshold. For $\Delta < 0$, $\log S_n = -|\Delta|(\ln 2) n^2/2 + O(n)$ and the mean remains finite. The equal-exponent curve in Fig. 2 uses $a = 0.8\sqrt{2} > 1$.

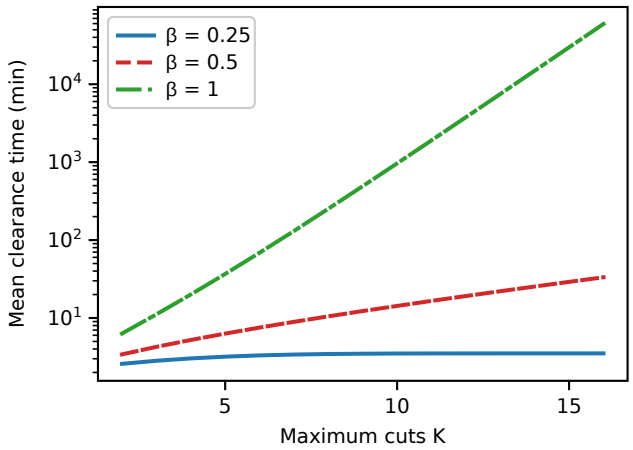


FIG. 2. Mean first-passage time with a floor after K cuts, with $c = 2\ \mathrm{min}^{-1}$, $u = 0.5\ \mathrm{min}^{-1}$, and $\gamma = 0.5$. For these rates the equal-exponent mean diverges as the floor is removed.

V. TRANSIENT CALCULATION AND SIMULATION

Let $p_n(t)$ be the probability of being occupied at stage n at time t, with $p_0(0) = 1$. The transient master equation is

$$\dot{p}_n = c_{n-1} p_{n-1} - (c_n + u_n) p_n, \tag{5}$$

where the incoming term vanishes for $n = 0$. The fraction of plates with a live credential is $P(t) = \sum_{n=0}^{K} p_n(t)$. The mean mass on the plate, including cleared plates in the denominator, is $m_0 \sum_{n=0}^{K} 2^{-n} p_n(t)$. For detection at threshold ϵ, we count only the terms with $m_n \geq \epsilon$. In the example below the coarser threshold excludes stage K alone.

We integrated Eq. (5) using Poisson uniformization [5]. The uniformizing rate was $c + u$. Time steps were chosen so their Poisson means were at most 12 and the sum included terms through event count 64. The omitted weight per step is below 1.3×10^{-26}; the transient probability vector was propagated without renormalization. As a check we simulated both exponential clocks directly at each occupied state [6], using 100,000 independent trajectories per row of Table I. The seeds were the consecutive integers 64271 through 64275 in table order, using the standard Python Mersenne Twister. All trajectories ran through to clearance. Sampling errors refer to the simulation alone.

TABLE I. Monte Carlo estimates of the fraction of plates reaching the floor, S_7, and the mean clearance time, τ_7. Errors are one standard error. Parameters are $K = 7$, $u = 0.5\ \mathrm{min}^{-1}$, and $\gamma = 0.5$.

c (min^{-1})	β	S_7	τ_7 (min)
2	0.25	0.0772 ± 0.0008	3.37 ± 0.01
2	0.5	0.2093 ± 0.0013	9.00 ± 0.05
2	1	0.4855 ± 0.0016	132.32 ± 0.71
1	1	0.2569 ± 0.0014	73.46 ± 0.56
4	1	0.6889 ± 0.0015	181.68 ± 0.77

For $c = 2\ \mathrm{min}^{-1}$ and $\beta = 1$, the exact S_7 is 0.4853 and the exact τ_7 is 130.90 min. Across all runs the simulated S_7 and τ_7 differ from their exact values by at most 2.1 standard errors. The occupancy at 12 min agreed within 2.9 standard errors in each case. The latter check uses elapsed time and can detect errors in the holding times that leave Eq. (2) unchanged.

VI. DEPARTURE CHECK

Consider a 120 g initial portion and a floor of $K = 7$ cuts. The smallest remnant is 0.9375 g. A 1 g scale threshold ceases to record a credential after the seventh cut. A 0.5 g threshold records every residual portion, including the one at the floor. We use $\beta = 1$, $\gamma = 1/2$, and $u = 0.5\ \mathrm{min}^{-1}$. The initial whole-claim time scale is 2 min and the corresponding time scale at the floor is 256 min. The service interval $D = 12$ min is long compared with one cut, and short compared with whole claims at the floor.

For $D = 12$ min, the mean residual mass per plate falls from 1.32 to 0.66 g as the initial cut rate rises from 1 to $4\ \mathrm{min}^{-1}$ (Fig. 3). The probability of a live credential rises from 0.296 to 0.672. The probability of registering one at a threshold of 1 g falls from 0.259 to 0.024 (Fig. 4). Roughly 65% of fast-cut plates carry a credential on a remnant below the 1 g threshold at departure. The mass of a detected credential enters the release check only through the threshold crossing.

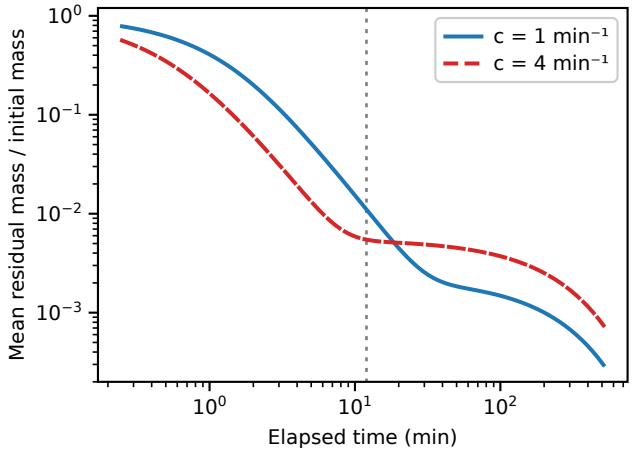


FIG. 3. Mean mass at a cutting floor of $K = 7$, with cleared plates included in the average. Both protocols use $u = 0.5\ \mathrm{min}^{-1}$, $m_0 = 120$ g, $\gamma = 0.5$, and $\beta = 1$. The dotted line marks departure at 12 min.

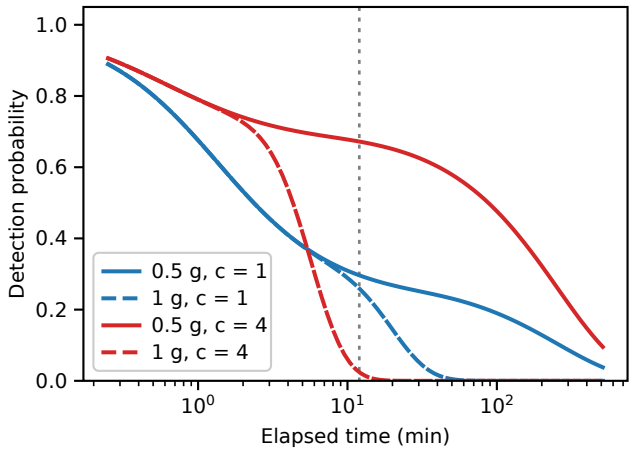


FIG. 4. Probability of a detected credential with thresholds of 0.5 g (solid) and 1 g (dashed), for the conditions in Fig. 3. The dotted line marks departure. Curves beyond that line show continued service with guests present.

The ordering at the finer threshold holds at every departure time in this example. Couple two processes so a slow-protocol cut also occurs in the fast protocol whenever both are at the same stage. Once the stages differ, a slow cut can at most close a gap of one stage. At a common whole-claim proposal, drawn at rate u, each process accepts with probability $2^{-\beta n}$. With the same

uniform random variate, a proposal accepted at the later stage is also accepted at the earlier one. Hence the fast-cut process retains a credential at least as long. This argument requires the finer threshold to include the floor. The coarse detector can release the fast process as soon as it reaches that floor.

VII. HOLDS ON VACANT ROOMS

Let $P_\epsilon(D)$ be the single-plate detection probability at departure. For a room with B independent plates the incidence Q_ϵ of a hold on the vacant room is

$$Q_\epsilon = 1 - [1 - P_\epsilon(D)]^B. \tag{6}$$

For small P_ϵ, the crossover from few holds to nearly universal holds occurs at BP_ϵ of order one. In the limit of increasing B at fixed $x = BP_\epsilon$, Eq. (6) tends to $1 - e^{-x}$. More generally, independent plates with probabilities $P_{\epsilon,i}$ give a no-hold probability $\prod_i (1 - P_{\epsilon,i})$. Residual mass grows additively with plate number, while the chance of releasing the room decreases geometrically at fixed P_ϵ.

The sensitivity to plate number depends on the threshold (Fig. 5). With the fast cut rate and a 1 g threshold, even twelve portions yield a hold in only 25% of vacant rooms. With the 0.5 g threshold, four portions give a hold in 98.8% of rooms. At the finer setting the fast protocol makes the room harder to release than the slow protocol, although its total residual mass at departure is about half as large.

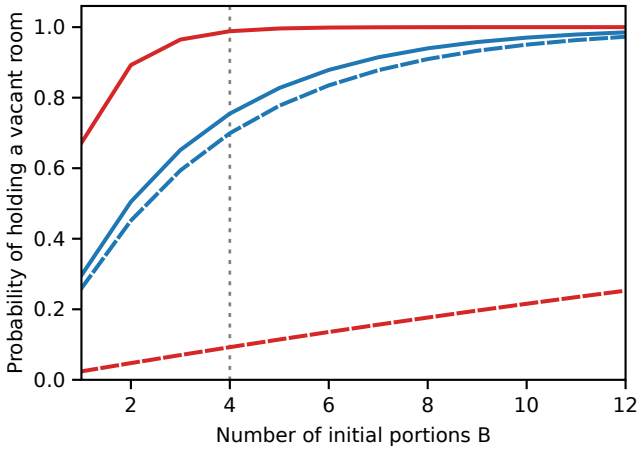


FIG. 5. Probability of holding a vacant room at departure under the two cut protocols. Solid and dashed curves use thresholds of 0.5 and 1 g, respectively. Blue denotes $c = 1\,\mathrm{min}^{-1}$ and red denotes $c = 4\,\mathrm{min}^{-1}$. The vertical line marks four starting portions.

VIII. INSPECTION AND ROOM CAPACITY

With requests waiting, each release starts a new booking. The kitchen replaces the plates at turnover. One cycle consists of D minutes of actual use, followed with probability Q_ϵ by an inspection wait of mean $1/b$. Let V_ϵ be the long-run fraction of room time for which the room is vacant but the credential check prevents its use. The renewal-reward theorem [7] yields

$$V_\epsilon = \frac{Q_\epsilon}{bD + Q_\epsilon}. \tag{7}$$

At fixed D and b, the ordering of V_ϵ follows that of Q_ϵ. For $b \to 0$ any positive Q_ϵ drives $V_\epsilon \to 1$. At $b = 0$ the first hold locks the room indefinitely. With independent cycles the number of completed releases before that hold has a geometric distribution with mean $(1 - Q_\epsilon)/Q_\epsilon$. This failure requires neither infinite divisibility nor an asymptotically small remnant. Serving stops at D, and custody of the residual credential passes to an inspection process with rate zero.

Table II gives the results for $B = 4$ and a mean inspection wait of 30 min. At the low cut rate, switching from the coarse threshold to the fine one raises vacant reserved time from 63.6% to 65.4%. At the high cut rate the same switch raises it from 18.8% to 71.2%. The two thresholds would therefore lead a room allocator to opposite decisions about which cut protocol to require.

TABLE II. Departure and turnover for four plates. Q_1 and $Q_{0.5}$ use thresholds of 1 and 0.5 g, respectively. $V_{0.5}$ is the corresponding fraction of vacant reserved time at the finer threshold. Mass is the mean per plate at departure.

c (min^{-1})	Mass (g)	Q_1 (%)	$Q_{0.5}$ (%)	$V_{0.5}$ (%)
1	1.32	69.9	75.5	65.4
2	0.82	60.0	93.3	70.0
4	0.66	9.3	98.8	71.2

For each $\beta = 1$ condition in Table I we grouped the 100,000 independent plates into 25,000 rooms of four and drew inspection waits using a second random stream with the original seed plus 10000. A shared inspection wait for each pair of threshold readings let both readings be made on the same room history. The means of the vacant reserved intervals, including zeros, were within 1.1 simulation standard errors of Q_ϵ/b across the six conditions. For the fast cut protocol, the fine and coarse estimates were 29.86 ± 0.19 min and 2.70 ± 0.08 min, respectively. The exact values are 29.65 min and 2.78 min.

The floor and threshold need to be specified independently. Changing K by one changes the final mass by a factor of two. In the illustrated setup the change from seven cuts to six would raise the final mass above the 1

g threshold. The discrepancy between sensors would then vanish, although the rate of retiring credentials before departure would also change. The example uses fixed departure times and independent plate histories. Common changes in attendance or serving frequency can correlate outcomes within a room. Equation 7 still applies to identically distributed renewal cycles if Q_ϵ is obtained from the joint room distribution instead of Eq. (6).

IX. UNEQUAL CUTS AND TERMINAL SERVICE

Equal halves give an exact mass ladder. The arrest boundary also holds for independent retained fractions F_j with support in $[f_-, f_+] \subset (0,1)$. Use the same rate laws in the actual mass ratio $x_n = \prod_{j=1}^{n} F_j$, so that the ratio of whole-claim to cut hazards is $r x_n^{\Delta}$. One can draw the sequence of retained fractions before running the clocks. For $\Delta > 0$, $x_n^{\Delta} \le f_+^{\Delta n}$ on every sequence, which gives a positive conditional probability of no clearance. For $\Delta \le 0$, each conditional factor for making another cut is at most $(1+r)^{-1}$ and the product vanishes. At $\Delta = 0$ the whole-claim probability is independent of the split fraction, and the expectation of $x_n^{-\gamma}$ is $\langle F^{-\gamma} \rangle^n$. The equal-exponent mean therefore diverges when $z\langle F^{-\gamma} \rangle \ge 1$. In this regime, occasional small daughters can have a substantial effect on the mean without changing the eventual probability of clearance.

An unconditional release at departure gives the scheduler a direct clock for the end of occupancy. A mass-independent rate of retiring credentials during the booking also reduces their survival. For an independent clearance rate h on each plate, the departure detection probability is multiplied by e^{-hD} and the mean lifetime of a credential is bounded by h^{-1} even without a cutting floor. This rule can be implemented by assigning the remaining portion to a guest at random and transferring it out of the room. An instruction to consider taking all of the food, represented by a finite increase in u, preserves the small-mass arrest boundary.

A room trial would require separate records of service events, guest departure, and facility release. For a one-gram scale, a reading below threshold may follow complete consumption or seven successive cuts. The room log alone would assign the same release to both histories, even though the second would still carry a credential at the finer threshold.

DATA AVAILABILITY

The numerical results use the rate laws, cutoff rules, and seeds specified above. All plate and inspection histories were generated computationally. The model parameters were fixed before simulation.

[1] V. Filimonau, J. Matute, H. Sezerel, and T. Özoğul Balyalı, J. Sustain. Tour. **34**, 60 (2026).

[2] P. L. Krapivsky and E. Ben-Naim, Phys. Rev. E **50**, 3502 (1994).

[3] P. L. Krapivsky, S. Redner, and E. Ben-Naim, *A Kinetic View of Statistical Physics* (Cambridge University Press, Cambridge, 2010).

[4] E. D. McGrady and R. M. Ziff, Phys. Rev. Lett. **58**, 892 (1987).

[5] S. Asmussen and P. W. Glynn, *Stochastic Simulation: Algorithms and Analysis* (Springer, New York, 2007).

[6] D. T. Gillespie, J. Phys. Chem. **81**, 2340 (1977).

[7] R. G. Gallager, *Stochastic Processes: Theory for Applications* (Cambridge University Press, Cambridge, 2013).