Surfacing of a Heavy Disc in a Vibrated Granular Medium

Censored survival analysis of waiting time, density dependence and critical behaviour in a two-dimensional soft-disc simulation

The Brazil nut effect · exploratory research report

Prepared by Baturhan Bıktım **with the assistance of** Claude Sonnet 5.5 (Anthropic)

28 November 2026

Scope note
This document reports an exploratory study that has not been peer reviewed. All results belong to the two-dimensional, non-rotating-particle, air-free simulation model defined in Section 2 and cannot be generalised directly to real granular systems. Raw data, formulas and charts are in the accompanying Excel workbook (Appendix A).

Abstract

The rise of a large object to the surface of a vertically vibrated granular bed (the Brazil nut effect) has been studied for a long time, but the effect of the object's density on this process has rarely been treated as a measurable *waiting time*. Here, the first time at which a single large disc of radius 4 reaches the 75th-percentile height of a bed of 230 small discs was measured with a two-dimensional soft-disc (DEM) simulation. A total of 168 runs (at least 4.0×10^5 vibration cycles) were analysed with right-censored survival methods (Kaplan–Meier, censored exponential and Weibull models); 112 further runs in exploratory extensions (extended density range, drive strength, binary mixture) are reported separately.

Findings. (1) As the large disc's density ρ rises from 2 to 2.75, the mean waiting time grows from about 143 to 2,600 vibration cycles. At higher densities surfacing becomes very rare but not zero: for ρ≥2.9 only 3 of 38 runs surfaced (in 322,515 cycles in total; 1/6 at ρ=2.95, 2/14 at ρ=3); the pooled mean wait is ≈107,505 cycles (95% CI 36,786–521,302). (2) A log-linear (exponential) increase of the waiting time with ρ is rejected (p=0.0003). A critical form diverging near ρ_c≈3.11 (95% interval 3.03–3.41) is preferred by AIC (ΔAIC=16.3), but there are only three events at ρ≥2.9, so a genuine divergence has not been demonstrated. The earlier estimate ρ_c≈2.9 and the result ‘no surfacing at ρ=3’, obtained from shorter runs, were invalidated by the long runs. (3) Surfacing is not a memoryless process (Weibull shape parameter 1.9–3.6). (4) At ρ=2.5 the waiting time of the large disc is indistinguishable from that of a neutral tracer of the same size as the small discs (≈514 cycles); density removes the size advantage at about this point. (5) A disc that does not surface does not sink to the bottom; it fluctuates in the lower-middle region of the bed, and its relative depth decreases smoothly with ρ (≈0.61 → 0.37). (6) Wall friction lengthens the time appreciably only for μ<0.5; no significant difference was observed over the radius range R=3–5. (7) Additional exploratory experiments: the waiting time is flat for ρ≤1.5 and no surfacing was seen for ρ=3.5–8 (the disc still does not sink); the drive strength Γ matters little for ρ=2 once Γ≥3; and in a 2D vibrated binary mixture a size ratio of 2–3 drives the minority component upward significantly, whereas a density difference at equal size does not in 500 cycles.

Limitations. The results belong to a two-dimensional, non-rotating, air-free model with small samples (6–8 runs per condition); the location of the critical density is weakly constrained and the proposed mechanisms are hypotheses not tested in this study. Possible relevance to pharmaceutical blends and asteroid regoliths is discussed, but not tested.

Keywords: Brazil nut effect, granular segregation, vertical vibration, discrete element method, censored survival analysis, critical slowing down.

Contents

Abstract **2**

1 Introduction **5**

1.1 The Brazil nut effect and the role of density . . . 5

1.2 Position of this study . . . 5

1.3 Research questions . . . 5

2 Methods **6**

2.1 Simulation model . . . 6

2.2 Exit criterion and measured quantities . . . 6

2.3 Experimental design . . . 6

2.4 Statistical methods . . . 7

2.5 Reproducibility . . . 7

2.6 Additional experiments (E1–E3) . . . 8

3 Results **8**

3.1 An illusion: the 500-cycle ‘threshold’ . . . 8

3.2 Density dependence . . . 10

3.3 Distribution of the waiting time: not memoryless . . . 11

3.4 The non-floating state: the disc does not sink, it stays suspended . . . 12

3.5 Effect of the initial layout . . . 13

3.6 Wall friction . . . 14

3.7 Radius and the neutral control . . . 15

3.8 Extended density range: ρ = 0.25 to 8 . . . 15

3.9 Dependence on the drive strength Γ . . . 17

3.10 A pharmaceutical-style test: does a minority component segregate? . . . 18

3.11 Summary of the additional experiments . . . 20

4 Discussion **20**

4.1 What density does: two baselines . . . 20

4.2 Relation to the literature . . . 20

4.3 Possible mechanisms (hypotheses; not tested in this study) . . . 21

4.4 Meaning of the critical behaviour . . . 21

4.5 A methodological lesson . . . 22

4.6 What this does and does not say about pharmaceutical blends and asteroid regoliths . . . 22

5 Limitations **23**

6 Conclusions and future work **23**

6.1 Summary of the conclusions . . . 23

6.2 What we concluded, in plain words . . . 24

6.3 The mathematics behind the conclusions 25
6.4 Future work 27

References **28**

Appendix A Excel workbook and sheet guide **28**

Appendix B Raw exit times (group G1) **29**

Appendix C The earlier stage (S1–S4) and this report **30**

1 Introduction

1.1 The Brazil nut effect and the role of density

When a box of mixed nuts is shaken, the largest nuts rise to the top. This is known as the Brazil nut effect (BNE). Rosato and colleagues (1987) showed with Monte Carlo simulations that the effect can be produced by a local, geometric mechanism, and that a large particle can rise even when it is denser [1]. Knight and colleagues (1993) reported that the rise of a single large glass bead in a vertically vibrated cylinder is linked to a convective circulation in the bed: particles move up in the middle and down through a narrow region near the wall, and particles larger than the width of this region are kept at the top [2]. In a two-dimensional molecular dynamics study, Pöschel and Herrmann (1995) reported that the rise occurs together with convection cells and that the probability of rising by this route depends strongly on the depth of the large particle [3].

The density effect is less clear. Shinbrot and Muzzio (1998) reported that in deep beds under large-amplitude vibration, large and heavy particles rise while light particles of the same size sink (‘reverse buoyancy’) [4]. Hong, Quinn and Luding (2001) proposed a theory that explains the transition from the BNE to the reverse BNE as a competition between percolation and condensation [5]; Breu and colleagues (2003) studied this transition experimentally [6]. Möbius and colleagues (2001) showed that in a three-dimensional bed the rise time does not depend monotonically on density [7]; in a detailed 2005 study they reported that this density dependence is caused by air drag [8]. The review by Kudrolli (2004) notes that segregation is a complex phenomenon sensitive to small changes in physical conditions and that substantial work is still needed for an accurate description [9]; Shinbrot (2004) published a commentary discussing the reverse effect [10].

1.2 Position of this study

Most of the literature reports the *speed* of the rise or its *direction* (rising or sinking). This study uses a different measure: the time that elapses before the centre of the large disc **first** reaches the 75th-percentile height of the bed (the first-passage time). The reason is statistical: for dense discs the exit may not occur at all within a given observation time, and leaving non-surfacing runs out, or counting the run limit as if it were a time, makes the waiting time look systematically *shorter* (Section 3.1). The data were therefore treated as right-censored survival data.

The study grew out of a search for ‘easy-looking unsolved physics problems’ and is exploratory; it makes no claim to contribute to the peer-reviewed literature. Granular convection is tagged with an unsolved-problem template on Wikipedia [11]; this is only an indication that the question is considered open and is not a primary source.

1.3 Research questions

- **Q1.** How does the waiting time change as the large disc's density ρ increases? Is the change exponential, or a critical form diverging at a finite ρ_c?
- **Q2.** Is surfacing a memoryless process (constant hazard rate)?
- **Q3.** If the disc does not surface at high density, what does it do: sink to the bottom, or stay suspended?

- **Q4.** How much do the initial layout, the wall friction and the radius of the large disc affect the result?
- **Q5.** What is the right comparison (control) for measuring the density effect?

2 Methods

2.1 Simulation model

The simulation is a discrete element (DEM) soft-disc model in the style of Cundall and Strack (1979) [12]. The particles are circles that interact through a linear spring and dashpot in the normal direction and through Coulomb-limited viscous friction in the tangential direction. There is no rotational degree of freedom (non-rotating particles). Units: g=1 and small-disc radius ≈1. Table 1 gives the model parameters.

Table 1. Simulation parameters.

Parameter	Value	Description
Container width W	32	open-topped container; side-wall height H=44 (bed surface ≈25)
Number of small discs	230	radius uniformly distributed in 0.9–1.1 (to prevent crystallisation)
Large disc	R=4 (default)	mass $m=\rho R^2$; ρ is the ratio to the small-disc density (small disc: $m=r^2$)
Gravity g	1	dimensionless
Normal stiffness K	1500	$f_n = K\delta - c\cdot v_n$ (δ: overlap)
Damping ratio ζ	0.1	$c = 2\zeta\sqrt{(K\cdot m_{eff})}$
Inter-particle friction	μ=0.5	Coulomb limit $\mu\cdot f_n$, tangential coefficient 0.5c
Wall/floor friction μ_w	0.8 (default)	scanned between 0.1 and 1.2 in group G2
Vibration	ω=2, Γ=4	floor $y=A\cdot\sin(\omega t)$, $A=\Gamma/\omega^2=1$ ($\Gamma=A\omega^2/g$); one cycle ≈3.14 time units = 1571 steps
Time step	Δt=0.002	semi-implicit Euler
Initial state	—	hexagonal-like arrangement + ±0.1 noise; large disc at the bottom centre; 2500 damped settling steps
Seed	101–108	determines the initial noise and layout

2.2 Exit criterion and measured quantities

Measurements are taken ten times per vibration cycle. A run is recorded as an **exit** the first time the centre height y of the large disc is greater than or equal to the 75th percentile p_{75} of the small-disc heights (excluding the first cycle); the exit time T is the (fractional) number of vibration cycles. If no exit occurs by the run limit, the run is right-censored (T > limit). In most runs the simulation was continued for 100 more cycles after the exit, but only the first passage is analysed. As an auxiliary measure, the ratio of the disc's mean height before the exit (cycle ≥50) to the mean p_{75} is computed as the ‘relative depth’.

2.3 Experimental design

Table 2. Experimental groups (168 runs in total).

Group	Purpose	ρ	μ_w	R	Seeds	Runs	Run limit (cycles)
G1	density scan	2; 2.25; 2.5; 2.75; 2.875; 2.9; 2.95; 3; 3.05; 3.1	0.8	4	101–108 (ρ=2.875: 101–106; ρ=2.9–3.1: 201–206)	76	1500 (ρ≤2.25); 3000 (2.5; 2.75); 4000 (2.875); 3000/4000/10,000 (ρ=3)[a]; 10,000 (2.9–3.1)
G2	wall friction × layout	2; 2.25	0.1; 0.2; 0.5; 0.8; 1.2	4	101–106	60	1500
G3	radius	1	0.8	3; 4; 5	101–106	18	1500
G4	control (short)	1	0.8	1	101–106	6	600
G4b	control (long)	1	0.8	1	101–108	8	3000

[a] There are 14 runs at ρ=3: four from the previous stage (seeds 101–104, limit 4000), four from the first campaign (limit 3000) and six from the long runs (seeds 201–206, limit 10,000); the total exposure is 86,269 cycles. The ρ=2.9–3.1 runs (G5) were added after the first version of this report and were merged into group G1.

2.4 Statistical methods

- **Right censoring and Kaplan–Meier.** Runs that do not surface are treated as censored at the limit; the probability of not yet having surfaced is computed with the Kaplan–Meier product-limit estimator [13].
- **Censored exponential model.** Under a constant hazard rate, the exit rate is $\hat{\lambda}$ = (number of exits) / (total exposure time); the mean wait is $\tau = 1/\hat{\lambda}$. The confidence interval is the exact Poisson (chi-square) interval; if there are no exits, only a lower bound for τ can be given.
- **Density models.** (i) $\ln \lambda = a + b(\rho - 2.5)$ (exponential); (ii) $\lambda = A(\rho_c - \rho)^{\beta}$, $\rho < \rho_c$ (critical). Parameters were found by censored maximum likelihood, confidence intervals were derived from the profile likelihood, and the models were compared with the Akaike criterion (AIC) [14]. The fit of the exponential model was tested with a likelihood-ratio (LR) test against a saturated model with a separate λ for each ρ (df=8).
- **Weibull fit.** A censored Weibull model was fitted for each ρ [15]; the hypothesis that the shape parameter k=1 (exponential) was tested with an LR test.
- **Paired and rank tests.** Wilcoxon signed-rank test across identical seeds [16], Spearman rank correlation, Kendall's coefficient of concordance W, and the Kruskal–Wallis and Mann–Whitney U tests.
- **Caution.** There are 6–8 runs per condition, so the tests have low power. No multiple-comparison correction was made; p-values should be read as exploratory.

2.5 Reproducibility

All raw runs (group, ρ, μ_w, R, seed, exit, heights) are on the *Data_Campaign* sheet of the Excel workbook; the summary tables are computed from this sheet with live formulas. The same seed produces the same initial layout; the simulation core and run scripts are kept in the accompanying working folder.

2.6 Additional experiments (E1–E3)

After the main campaign three further experiments were run to probe the reach of the results. They use the same simulation core and exit criterion, and the same censored-survival estimates (Section 2.4); they are exploratory and have only 4–5 runs per condition.

- **E2, extended density.** ρ=0.25, 0.5, 1.5, 3.5, 4, 5, 6 and 8 (R=4, μ_w=0.8, Γ=4, seeds 101–104; run limit 600–800 cycles for $\rho \le 1.5$ and 1,500 for $\rho \ge 3.5$), combined with the existing ρ=1–3 data (Section 3.8).
- **E3, drive strength.** Γ=2, 3, 5, 6 and 8 at ρ=1 and ρ=2 (ω=2, seeds 101–104, limit 1,200), compared with the existing Γ=4 data (Section 3.9).
- **E1, binary mixture.** A pharmaceutical-style test (Section 3.10). In place of one intruder, a minority component (called 'API-like') makes up 20% of the total *mass* of a bed of 230 'excipient' discs of radius ≈1. The API-like discs have radius $s_A \in \{1, 2, 3\}$ times the excipient radius and density $d_A \in \{0.5, 1, 2\}$ times the excipient density, so their number is fixed by the mass fraction (e.g. 29 discs for s_A=2, d_A=0.5; 3 discs for s_A=3, d_A=2). The bed starts as a staggered lattice in which the API-like discs occupy randomly chosen sites (random sequential addition) and is briefly settled, which gives an approximately well-mixed start; it is then shaken for 500 cycles (Γ=4, ω=2, μ_w=0.8; five seeds per case, 40 runs in total). Two quantities are recorded: the segregation index $S=(\bar{y}_{API}-\bar{y}_{exc})/H$, where $\bar{y}$ is the mean height of the discs of each type and H the 90th-percentile height of the bed (S=0 well mixed, S>0 the API-like discs have moved up, S<0 down); and the content of four layers of equal total mass, expressed relative to the overall API fraction (1 means on target; layers outside 0.85–1.15 are counted as failing, in analogy with the 85–115% of label claim used in content-uniformity testing). The case (s_A=1, d_A=1) is the neutral control: discs of identical size and density.

3 Results

3.1 An illusion: the 500-cycle 'threshold'

In the first scans the run limit was 500 vibration cycles, and there appeared to be a sharp transition between ρ=2.5 and 2.75: the fractions of runs surfacing within 500 cycles were 8/8, 7/8, 4/8 and 0/8 for ρ=2, 2.25, 2.5 and 2.75. When the limit was extended with the same seeds, these fractions rose to 8/8, 8/8, 8/8 and 4/8 at 1500 cycles, and to 6/8 at 3000 cycles for ρ=2.75 (Table 3). The 'threshold' moved as the observation window widened: it was a censoring artefact. The subsequent analysis therefore uses methods that model censoring explicitly (Figure 2).

Table 3. The ‘threshold’ moves as the observation window grows: number of runs surfacing by run limit (same seeds; for ρ≥2.875 the final limits are 4000–10,000 cycles and the number of runs differs between groups).

ρ	Runs	Surfaced by 500 cycles	Surfaced by 1500 cycles	Surfaced by final limit
2	8	8/8	8/8	8/8
2.25	8	7/8	8/8	8/8
2.5	8	4/8	8/8	8/8
2.75	8	0/8	4/8	6/8
2.875	6	0/6	1/6	1/6
2.9	6	0/6	0/6	0/6
2.95	6	0/6	0/6	1/6
3	14	0/14	0/14	2/14
3.05	6	0/6	0/6	0/6
3.1	6	0/6	0/6	0/6

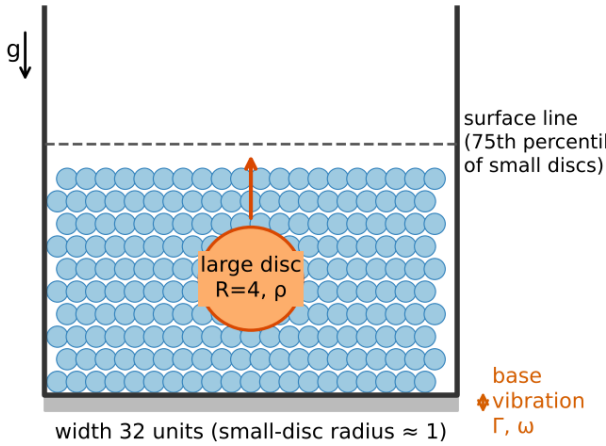

Figure 1. Schematic of the experimental set-up: small discs (blue) form a bed, the large disc (orange) starts at the bottom centre, and the floor is vibrated vertically. The dashed line is the 75th-percentile surface line used in the exit criterion.

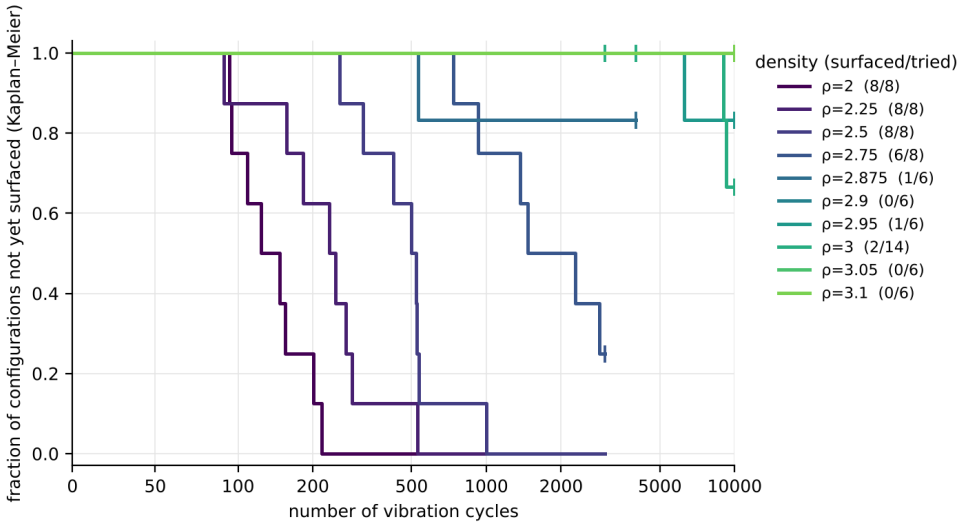

Figure 2. Kaplan–Meier estimate: probability that the large disc has not yet surfaced (6–8 runs per ρ; the horizontal axis is symmetric-logarithmic). Vertical ticks are censored runs. The curve shifts to the right as ρ increases. The horizontal axis has been extended to 10,000 cycles. Curves for ρ≥2.9 stay in the 0.8–1.0 range because of censoring; at ρ=3 two exits are seen at about cycle 9,000.

3.2 Density dependence

Table 4 and Figure 3 show that the mean waiting time grows rapidly from ρ=2 to 2.75: τ is about 143 (ρ=2), 250 (2.25), 510 (2.5) and 2,600 (2.75) cycles. At ρ=2.875 only one of six runs (seed 106, at cycle 530) surfaced. We also scanned this region with long runs (6 new seeds per ρ, limit 10,000 cycles): 0/6 exits were observed at ρ=2.9, 1/6 at ρ=2.95 (cycle 6,246), 2/6 at ρ=3 (cycles 8,999.7 and 9,268.9), and 0/6 at ρ=3.05 and 3.1. The earlier result for ρ=3, 'no exit in 28,000 cycles', was therefore not a firm prohibition but a censoring result showing that the exit rate is very small. Because each ρ at ρ≥2.9 has only 0–2 events, separate estimates by ρ are highly uncertain; the pooled estimate is more reliable: 3 events / 322,515 cycles, $\lambda \approx 9.3\times10^{-6}$ /cycle (95% CI $1.9–27.2\times10^{-6}$), τ≈107,505 cycles (95% CI 36,786–521,302).

Table 4. Exit rate and waiting time by density (censored exponential estimate). For ρ values with no exits only a lower bound for τ can be given; rows with 1–2 events are very uncertain.

ρ	Runs	Exits	Total exposure (cycles)	λ (10^{-3} /cycle)	τ=1/λ (cycles)	95% CI (τ)	KM median
2	8	8	1,144	6.992	143	73 – 331	124
2.25	8	8	1,996	4.007	250	127 – 578	233
2.5	8	8	4,080	1.961	510	259 – 1,181	500
2.75	8	6	15,627	0.384	2,604	1,197 – 7,097	1,470
2.875	6	1	20,530	0.049	20,530	3,685 – 810,884	—
2.9	6	0	60,000	0	—	> 16,265	—
2.95	6	1	56,246	0.018	56,246	10,095 – 2,221,614	—
3	14	2	86,269	0.023	43,134	11,941 – 356,174	—
3.05	6	0	60,000	0	—	> 16,265	—
3.1	6	0	60,000	0	—	> 16,265	—

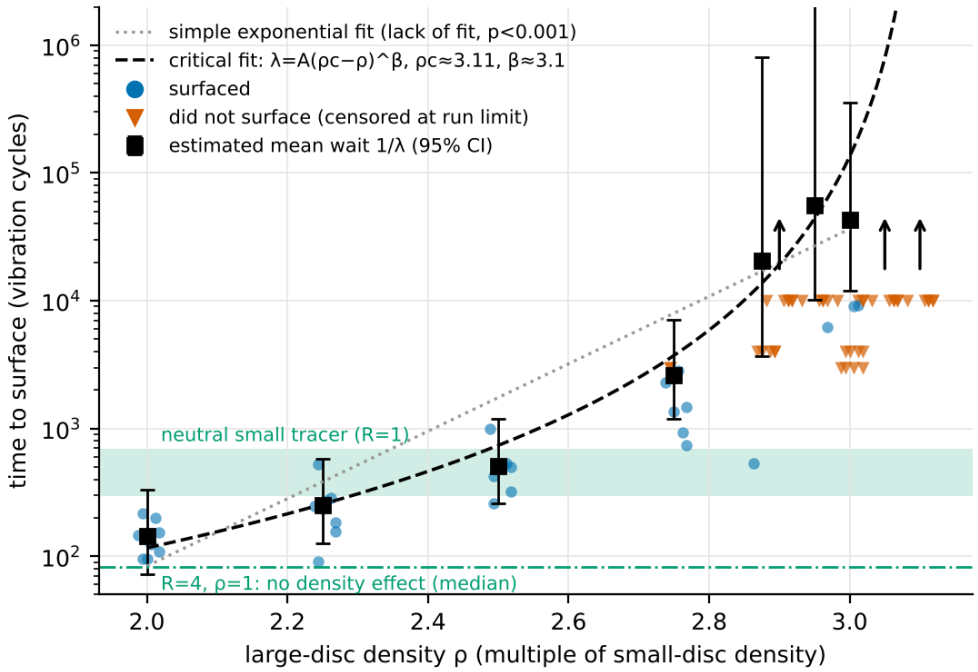

Figure 3. Exit time (log scale) and estimated mean wait 1/λ (black squares, 95% CI). Blue dots are runs that surfaced, orange triangles are runs that did not (censored). Upward arrows show the 95% lower bound for τ at ρ values with no observed exit. The dotted grey line is the simple exponential fit and the dashed black curve is the critical fit. The green band is the range of exit times of a neutral tracer of the same size as the small discs (R=1, ρ=1; 8 runs); the green dash-dotted line is the median for R=4 and ρ=1 (geometric rise in the absence of a density effect).

The exponential model is rejected. In the simple exponential model each 0.25 increase in ρ multiplies the exit rate by ×0.22 (95% CI 0.18–0.27); that is, the mean wait lengthens by a factor of about 4.6 at each step. However, the fit is deficient against the saturated model that gives each ρ its own λ (LR=28.8, df=8, p=0.0003): the observed τ(ρ) is not a straight line on a log scale: the exponential fit over-predicts the waiting time by about 1.5–3.4× at ρ=2.25–2.75 and under-predicts it by about 1.7× at ρ=2 (Table 5).

The critical form fits better. The model $\lambda = A(\rho_c - \rho)^{\beta}$ with ρ_c=3.115 and β=3.11 raises the log-likelihood by 9.2 units; AIC=530.8 (547.1 for the exponential, ΔAIC=16.3). Table 5 compares the observed waiting time with the predictions of the two models. The 95% profile interval for ρ_c is wide, 3.03–3.41; its lower end is already implied by events having been observed at ρ=3 (ρ_c>3). ρ_c≈2.9 was found in the previous stage; once events were observed at ρ=3 the estimate shifted upwards. This shift shows that the estimate is sensitive to the data grid and the run length.

Table 5. Observed mean waiting time (cycles; in parentheses: surfaced/tried) and the predictions of the two models. The critical model's prediction at ρ=3.1 is meaninglessly large because it is very close to ρ_c.

ρ	Observed τ	Exponential fit	Critical fit
2	143 (8/8)	84	116
2.25	250 (8/8)	382	255
2.5	510 (8/8)	1,746	737
2.75	2,604 (6/8)	7,980	3,729
2.875	20,530 (1/6)	17,062	13,723
2.9	no exit (τ > 16,265)	19,862	19,316
2.95	56,246 (1/6)	26,916	43,968
3	43,134 (2/14)	36,477	135,005
3.05	no exit (τ > 16,265)	49,432	794,949
3.1	no exit (τ > 16,265)	66,990	75,729,224

Important caveat. The advantage of the critical model means that it ‘describes the data better’; it is not evidence of a genuine divergence (infinite waiting time). This three-parameter model was fitted to ten ρ values, and there are only three events at ρ≥2.9. Observing no exit in 12 runs and 120,000 cycles in total at ρ=3.05 and 3.1 is consistent both with the pooled rate for ρ≥2.9 (≈1.1 expected events, P(0)≈0.33) and with the critical model. The exponential model, in contrast, predicts a total of ≈9.6 events over ρ=2.9–3.1, whereas 3 were observed (Poisson P(≤3)≈0.014). So in the high-ρ region surfacing is rarer than a log-linear extrapolation predicts; but whether this arises from a threshold or from a slowly decreasing continuous rate cannot be distinguished with these data.

3.3 Distribution of the waiting time: not memoryless

When a censored Weibull model is fitted for each ρ, the shape parameter k is 3.6, 2.2 and 2.5 for ρ=2.0–2.5; the hypothesis k=1 (constant hazard rate, exponential) is rejected at these three densities (p=0.0004; 0.014; 0.004). At ρ=2.75, k=1.9, but power is low because only six of eight runs surfaced and because of censoring (p=0.12). k>1 means that the exit hazard grows with time: the exit behaves like the outcome of a process that needs preparation. The 1/λ in Table 4 should therefore be read as the ‘mean wait under the exponential assumption’; the true distribution is narrower than an exponential. The shape parameter rests on only 6–8 observations at each ρ; confidence intervals were not computed and the k values should be read as approximate.

Table 6. Weibull shape parameter (k=1 corresponds to the exponential model).

ρ	Runs	Exits	Weibull shape k	Scale	LR test (k=1): p
2	8	8	3.55	159	0.0004
2.25	8	8	2.19	283	0.014
2.5	8	8	2.55	575	0.0042
2.75	8	6	1.85	2,474	0.116
≥ 2.875	44	4	—	—	cannot be estimated (too few exits)

3.4 The non-floating state: the disc does not sink, it stays suspended

Before the exit (or, for runs that do not surface, over the whole duration, cycle ≥50), the mean height of the large disc relative to the bed surface line is 0.61–0.65 for ρ≤2.5, and falls to 0.55 at ρ=2.75, 0.44 at ρ=2.875, 0.39 at ρ=3 and 0.37 at ρ=3.1 (Table 7, Figure 4); the decrease continues smoothly over ρ=2.9–3.1 and the spread between runs is very narrow. At ρ=3 all 14 runs lie in a narrow range (SD=0.02): a systematic suspended/floating state. The disc does not sink to the bottom (resting on the floor its centre height would be ≈4; the observed mean is 9.8 and the surface is ≈24.9); however, it mostly does not reach the surface either: in the 42 runs that did not surface, the highest position reached stayed on average 1.4 units below the mean surface line. The example trajectories in Figure 5 show that at ρ=2.75–2.875 the exit follows a long excursion in the lower-middle region, and at ρ=3 it occurs at about cycle 9,000; before the exit the disc fluctuates in the lower-middle region.

Table 7. Relative depth before the exit (disc-centre height / surface height).

ρ	2	2.25	2.5	2.75	2.875	2.9	2.95	3	3.05	3.1
Relative depth (mean)	0.61	0.65	0.64	0.55	0.44	0.43	0.41	0.39	0.38	0.37
Standard deviation	0.09	0.05	0.08	0.08	0.03	0.02	0.01	0.02	0.01	0.01
Mean disc-centre height	15.2	16.3	15.8	13.5	10.9	10.9	10.3	9.8	9.4	9.3

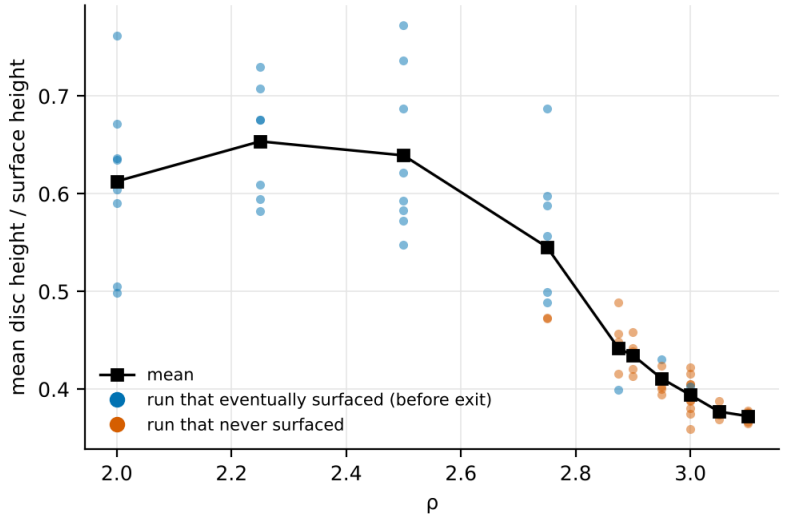

Figure 4. Relative depth (disc-centre height / surface-line height) decreases with ρ. Blue: runs that eventually surfaced (period before the exit); orange: runs that did not surface; black squares: mean.

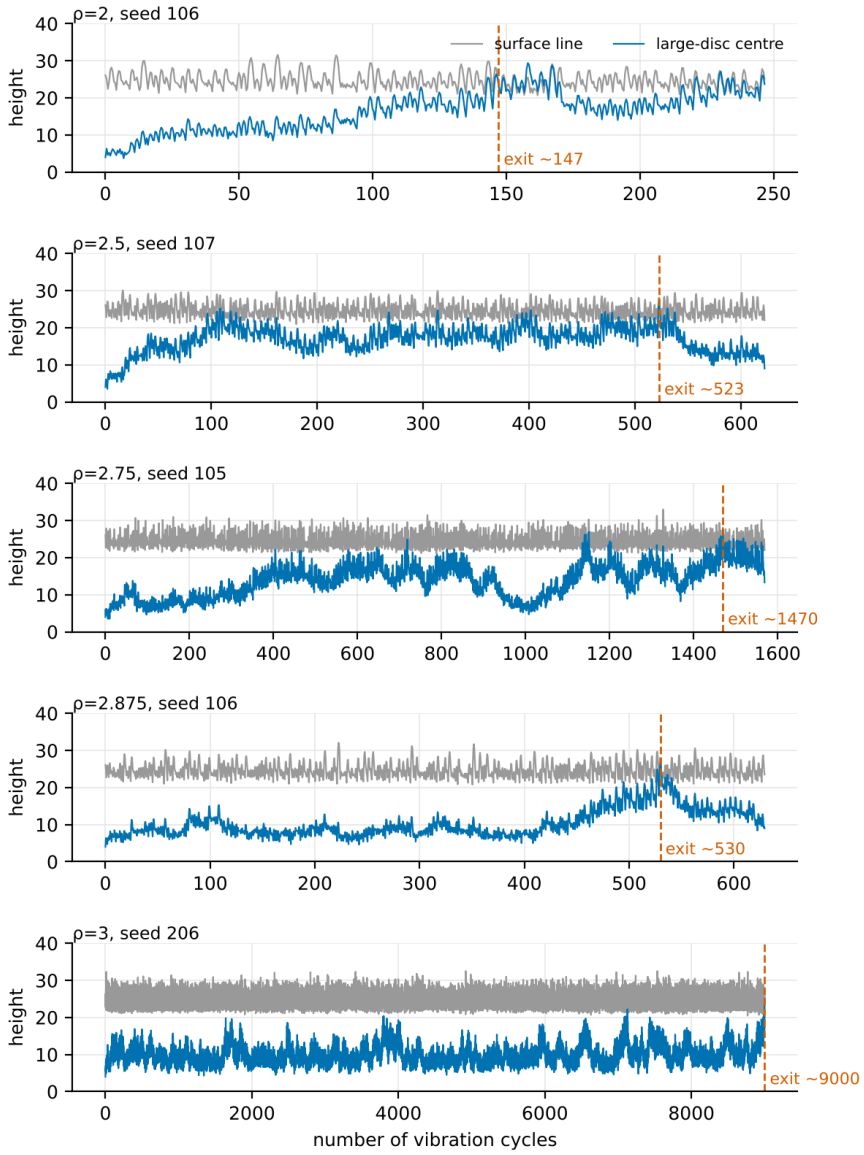

Figure 5. One example trajectory for each ρ (the run whose exit is closest to the median; the first run if there is no exit). Grey: surface line (p_{75}), blue: large-disc centre, orange dashed line: first exit. The horizontal-axis scale differs in each panel.

3.5 Effect of the initial layout

The same seed represents the same initial layout. If a layout exits quickly at low ρ, does it also exit quickly at higher ρ? Spearman correlations between pairs of ρ for eight seeds: r_s=-0.36 (p=0.39) for ρ=2.0–2.25; 0.52 (p=0.18) for 2.25–2.5; -0.21 (p=0.61) for 2.0–2.5. No consistent layout effect was seen. However, with n=8 an association cannot reach significance unless $|r_s|$ exceeds about 0.74, so a medium-sized layout effect cannot be ruled out.

3.6 Wall friction

At ρ=2 the median exit time is 701, 289, 159, 151 and 125 cycles for μ_w=0.1, 0.2, 0.5, 0.8 and 1.2 respectively. At ρ=2.25 only one of six runs surfaced at μ_w=0.1 (median at the limit, ≥1500); the others are 614, 415, 208 and 167 (Table 8, Figure 6). In the paired Wilcoxon test against μ_w=0.8, only μ_w=0.1 is significant at ρ=2 (ratio 5.4; p=0.031), and μ_w=0.1, 0.2 and 0.5 at ρ=2.25 (ratios 6.8, 2.9, 1.7; each p=0.031); there is no difference at μ_w=1.2. That is, low friction (roughly ≤0.5) slows the exit, and above ≈0.8 the effect saturates. For n=6 the smallest p-value a two-sided exact Wilcoxon test can take is 0.031; no correction was made for the eight tests. The ratio for μ_w=0.1 at ρ=2.25 is a lower bound because of censoring (the time was taken at the limit value).

Table 8. Exit time by wall friction (group G2, R=4, 6 seeds).

ρ	μ_w	Runs	Exits	Median exit (cycles)	Mean (non-exits at the limit)	Median ratio (μ_w/0.8)	Wilcoxon p
2	0.1	6	5	701.1	817	5.37	0.031
2	0.2	6	6	289.0	250	2.14	0.156
2	0.5	6	6	158.5	164	1.07	0.562
2	0.8	6	6	150.8	154	1 (reference)	—
2	1.2	6	6	125.3	121	0.86	0.312
2.25	0.1	6	1	1,500.0	1,466	6.76	0.031
2.25	0.2	6	6	613.5	660	2.86	0.031
2.25	0.5	6	6	414.6	492	1.73	0.031
2.25	0.8	6	6	208.1	200	1 (reference)	—
2.25	1.2	6	6	167.1	193	1.07	1.000

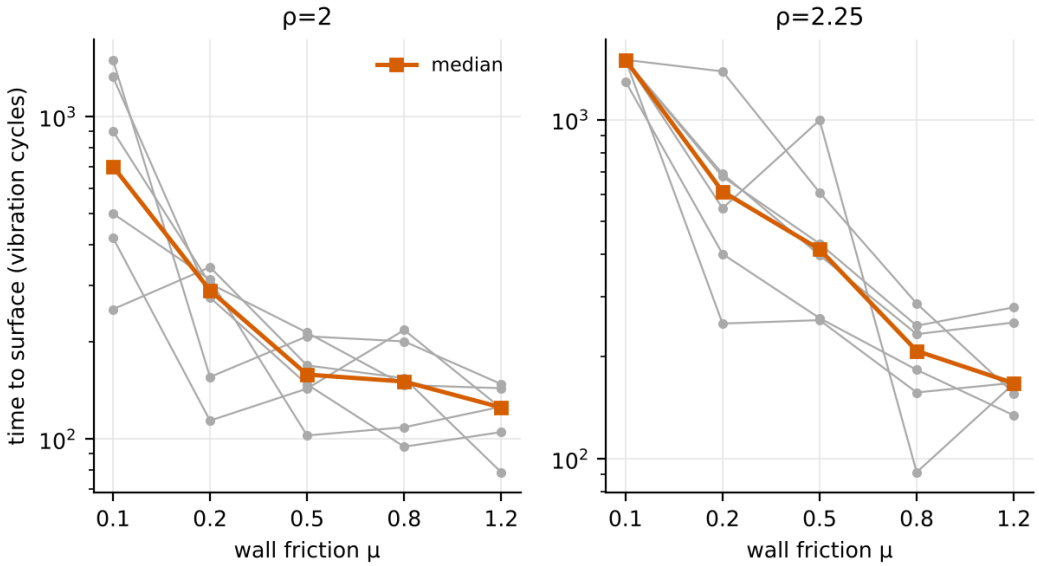

Figure 6. Exit time falls as wall friction μ_w increases (log scale). Thin grey lines connect runs with the same seed; the orange square is the median. At ρ=2.25 and μ_w=0.1, five runs are censored at the limit.

The consistency of the seed rankings across friction levels was measured with Kendall's W: W=0.22 (p=0.36) for ρ=2 and W=0.40 (p=0.074) for ρ=2.25. This agrees with the result in Section 3.5: the effect of the initial layout is weak and not significant. A mechanism explaining the effect of wall friction was not tested in this study; it is known in the literature that wall friction feeds convection [2, 8] (see Section 4.3).

3.7 Radius and the neutral control

At ρ=1, large discs with R=3, 4 and 5 surface in similar times (medians 92, 82 and 104 cycles, all surfaced); no difference is seen between groups (Kruskal–Wallis p=0.37). A neutral tracer of the same size as the small discs (R=1), on the other hand, is much slower at ρ=1: in the long control run (G4b, 8 runs) the median is 517.5 and the mean 514.1 cycles; in the short control (G4, limit 600) three of six runs surfaced. The roughly six-fold difference between the large disc and the tracer is statistically clear (R=4 vs R=1, Mann–Whitney p=0.0007). This shows that size alone, even without a density difference, speeds up the rise (Table 9, Figure 7).

Table 9. Radius and control runs (ρ=1, μ_w=0.8). The number of exits excludes runs that did not surface at the limit.

Group	R	ρ	Runs	Exits	Median	Mean	Shortest	Longest	Limit
G3	3	1	6	6	92.3	104.1	64.7	171.6	1500
G3	4	1	6	6	82.2	81.2	65.8	101.3	1500
G3	5	1	6	6	104.0	95.7	57.5	125.7	1500
G4	1	1	6	3	579.3	531.9	391.1	600.0	600
G4b	1	1	8	8	517.5	514.1	297.5	691.9	3000

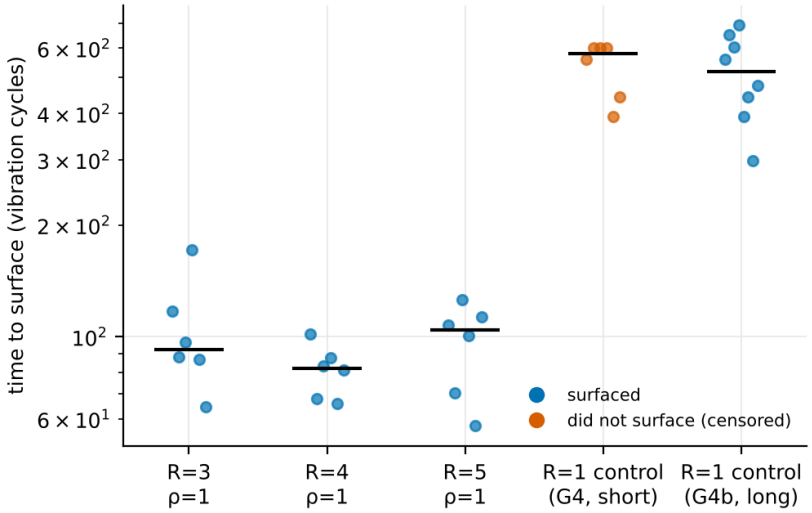

Figure 7. Exit time at ρ=1: R=3, 4, 5 (G3) and the R=1 neutral control (G4 short limit, G4b long limit). Black lines are medians; orange dots are runs that did not surface (censored) at the run limit.

The density effect is also visible when ρ=1 and ρ=2 are compared for R=4 with the same seeds (101–106): the median exit time rises from 82.2 to 150.8 cycles; the per-seed ratio is 1.1–3.2 (median 1.7) and the paired Wilcoxon test gives p=0.031 (ρ=2 is slower for all six seeds). At ρ=2.5 the exit time of the R=4 disc (median 511; 1/λ=510) cannot be distinguished from the R=1 neutral control (median 518; mean 514) (Mann–Whitney p=0.65, n=8+8); of course this is not proof of equivalence, only a failure to find a difference.

3.8 Extended density range: ρ = 0.25 to 8

The main scan covered ρ=1–3.1. To see what happens on either side, 32 new runs (R=4, μ_w=0.8, Γ=4, seeds 101–104) were made at ρ=0.25, 0.5 and 1.5 (limits 600–800 cycles) and at ρ=3.5, 4, 5, 6 and 8 (limit 1,500 cycles). For comparison, Table 10 and Figure 8 also include ρ=1 (G3) and ρ=2–3 (the first six seeds of G1), which are therefore not identical to Table 4, where all seeds are used.

Table 10. Extended density scan. Where there are no exits only the one-sided 95% lower limit (E/3.0) is given (>). Relative depth = mean height of the disc centre / bed surface line, before the exit or over the whole run (cycle ≥50).

ρ	Runs	Exits	Mean wait 1/λ̂ (95% interval)	Median exit time	Relative depth (mean ± SD)
0.25	4	4	68 (27–249)	73	0.74 ± 0.06
0.5	4	4	81 (32–296)	77	0.75 ± 0.17
1	6	6	81 (37–221)	82	0.74 ± 0.04
1.5	4	4	92 (36–338)	96	0.72 ± 0.13
2	6	6	154 (71–420)	151	0.61 ± 0.10
2.25	6	6	200 (92–544)	208	0.66 ± 0.05
2.5	6	6	522 (240–1,424)	511	0.63 ± 0.08
2.75	6	5	1,954 (837–6,019)	1,366	0.57 ± 0.08
2.875	6	1	20,530 (3,685–810,884)	530	0.44 ± 0.03
3	6	0	>7,344	–	0.39 ± 0.02
3.5	4	0	>2,003	–	0.34 ± 0.01
4	4	0	>2,003	–	0.32 ± 0.01
5	4	0	>2,003	–	0.30 ± 0.00
6	4	0	>2,003	–	0.28 ± 0.00
8	4	0	>2,003	–	0.26 ± 0.00

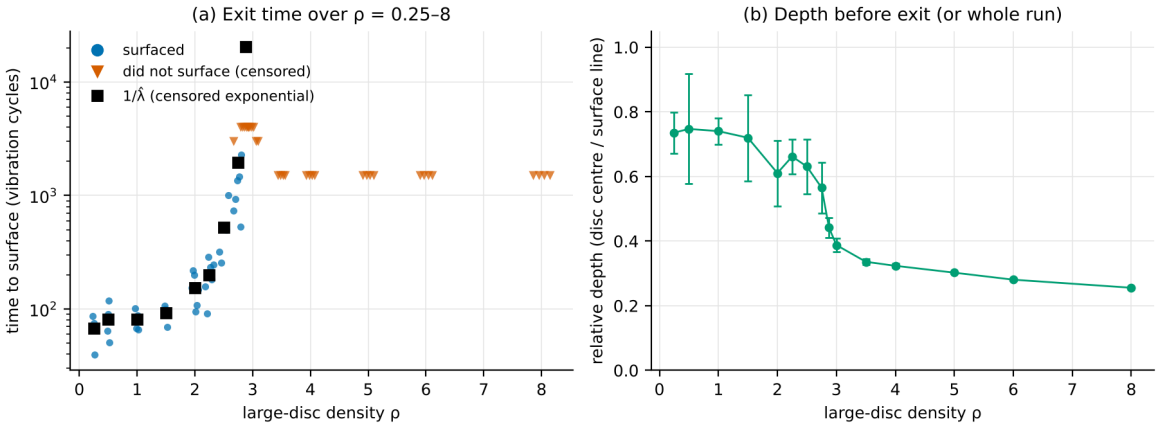

Figure 8. (a) Exit time against ρ for all runs (blue dots: surfaced; orange triangles: not surfaced by the run limit) with the censored-exponential mean wait (black squares; the point at ρ=2.875 rests on a single event). (b) Relative depth of the disc against ρ.

Light side. Below ρ=2 the waiting time is almost flat: the mean wait is 68, 81, 81 and 92 cycles at ρ=0.25, 0.5, 1 and 1.5, with heavily overlapping intervals (Kruskal–Wallis over the four groups, p=0.42; Spearman correlation of exit time with ρ, +0.38, p=0.12). A weak rise cannot be excluded, but making the large disc lighter than the small discs does not make it surface faster than at ρ=1. We therefore see no sign of a non-monotonic dependence on the light side within this resolution (four runs per point); in this model the size advantage is already fully expressed at ρ≈1, and below that density becomes irrelevant.

Heavy side. Between ρ=3.5 and 8 none of the 20 runs surfaced (6,000 cycles per density; 30,000 cycles in total, one-sided lower limit of the mean wait about 2,000 cycles at each density). This is not strong evidence on its own, because these lower limits are only comparable to the waiting times already seen at ρ=2.75 (about 2,000–2,600 cycles); but together with the main scan it shows that the exit rate does not recover at high density. The disc does *not* sink to the floor either: its relative depth decreases slowly and smoothly, 0.34

(ρ=3.5), 0.32, 0.30, 0.28 and 0.26 (ρ=8). At ρ=8 that is a mean height of about 6.4 length units, whereas a disc lying on the floor would have its centre at about 4 (one radius above the floor, which itself oscillates by ±1); its highest excursions still reach 13.5 on average. A very heavy disc therefore stays in the lower part of the bed but keeps moving in it.

In summary, over ρ=0.25–8 the dependence is: flat (≈70–90 cycles) up to ρ≈1.5, a rapid rise between ρ=2 and 2.75, and then a regime in which exits are rare or absent and the disc rides at a depth that decreases only slowly with ρ.

3.9 Dependence on the drive strength Γ

All earlier runs used Γ=4 (the peak acceleration of the floor in units of g, at ω=2). To test whether the picture depends on the drive, 40 further runs (R=4, μ_w=0.8, 4 seeds each) were made at Γ=2, 3, 5, 6 and 8 for ρ=1 and ρ=2 (limit 1,200 cycles). The Γ=4 values come from the existing data (G3 for ρ=1, G1 for ρ=2).

Table 11. Waiting time by drive strength Γ (ω=2, R=4, μ_w=0.8). At Γ=4 the number of runs is six.

ρ	Γ	Runs	Exits	Mean wait 1/λ̂ (95% interval)	Relative depth
1	2	4	1	4,442 (797–175,450)	0.37
1	3	4	4	142 (56–523)	0.61
1	4	6	6	81 (37–221)	0.74
1	5	4	4	72 (28–265)	0.73
1	6	4	4	62 (24–227)	0.85
1	8	4	4	39 (15–143)	0.87
2	2	4	0	>1,602	0.24
2	3	4	4	250 (98–919)	0.60
2	4	6	6	154 (71–420)	0.61
2	5	4	4	163 (64–598)	0.67
2	6	4	4	137 (54–504)	0.69
2	8	4	4	160 (63–588)	0.67

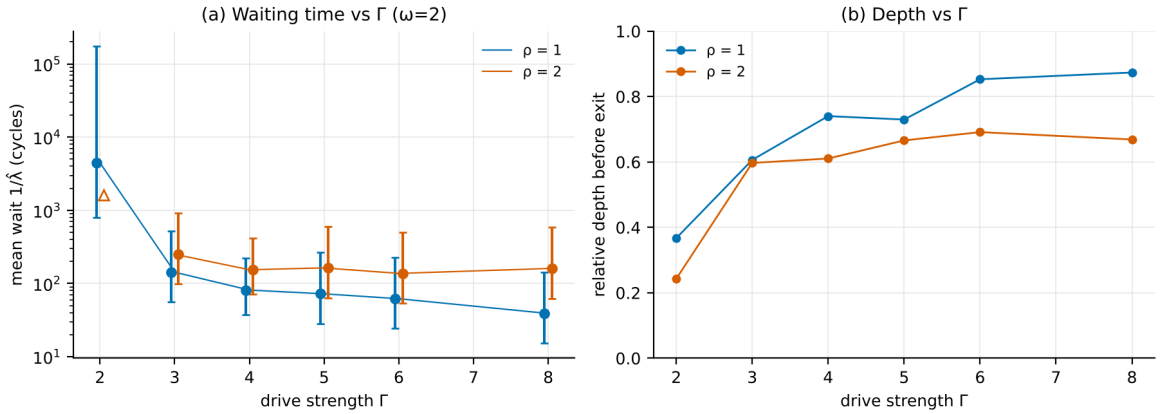

Figure 9. (a) Mean wait against Γ (error bars: exact 95% interval; the open triangle at Γ=2, ρ=2 is the lower limit because no run surfaced). (b) Relative depth before the exit against Γ.

Weak drive (Γ=2). At Γ=2 the bed hardly moves: only 1 of 4 runs surfaced at ρ=1 (after 842 cycles) and none of 4 at ρ=2, with a relative depth of 0.37 and 0.24 respectively, so the disc stays near the bottom. Surfacing needs a minimum drive, and a drive of Γ=2 is below it for the 1,200-cycle window used. **Above the minimum (Γ=3–8).** At ρ=1 the mean wait falls from about 140 (Γ=3) to 39 cycles (Γ=8). Fitting ln λ = a + α ln(Γ/4) to the 22 runs with Γ=3–8 gives α=1.21 (95% profile interval -0.1 to 2.5; likelihood-ratio p=0.065), so a trend is

suggested but not established. At ρ=2 the wait stays between about 140 and 250 cycles for every Γ from 3 to 8 (α=0.41, interval -0.8 to 1.6, p=0.50); within the available precision it is independent of the drive.

A tentative reading is that for the light disc (ρ=1) a stronger drive speeds up the rise, while for the heavier disc (ρ=2) something other than the drive limits it, so that the density effect becomes *stronger* as the drive increases: the ratio of the mean waits at ρ=2 and ρ=1 is about 1.8 at Γ=3 and 1.9 at Γ=4, but 4.1 at Γ=8 (with intervals that overlap). With four runs per point this is only a suggestion. What can be said firmly is that the large drive does not rescue the heavy disc: at Γ=8 the ρ=2 disc takes about 160 cycles, as at Γ=4. The location of the critical density ρ_c may therefore depend on Γ, which could not be explored here and is listed under future work.

3.10 A pharmaceutical-style test: does a minority component segregate?

The single-intruder experiments say little about a mixture in which a minority component is spread through the bed. Experiment E1 (Section 2.6) therefore shakes a bed of 230 small ‘excipient’ discs containing 20% by mass of ‘API-like’ discs that are larger and/or denser, and follows how far the API-like discs move relative to the excipient (segregation index S) and how the API content of four equal-mass layers changes. Table 12 and Figures 10–11 give the result after 500 cycles (S and layer contents averaged over cycles 400–500; five runs per case).

Table 12. Binary-mixture test (5 runs per case, 500 cycles). S = segregation index; ΔS = S late − S start per run, with the p-value of a one-sample t-test against 0 (not corrected for eight comparisons). Bottom and top layer = relative API content (1 = on target) at cycles 400–500.

Size ratio	Density ratio	API discs	S start	S late (± SE)	ΔS (± SE)	p	Bottom layer	Top layer
1	0.5	115	0.02	0.04 ± 0.01	0.02 ± 0.02	0.496	0.92	1.08
1	1	58	0.03	-0.00 ± 0.02	-0.03 ± 0.03	0.336	0.96	0.92
1	2	29	0.05	0.03 ± 0.02	-0.01 ± 0.04	0.789	0.97	1.12
2	0.5	29	0.05	0.19 ± 0.03	0.14 ± 0.02	0.003	0.51	1.56
2	1	14	0.01	0.13 ± 0.03	0.11 ± 0.03	0.031	0.41	1.40
2	2	7	0.01	0.18 ± 0.03	0.18 ± 0.08	0.106	0.17	1.48
3	1	6	0.05	0.28 ± 0.06	0.22 ± 0.09	0.069	0.37	1.98
3	2	3	0.04	0.19 ± 0.06	0.15 ± 0.11	0.260	0.15	1.25

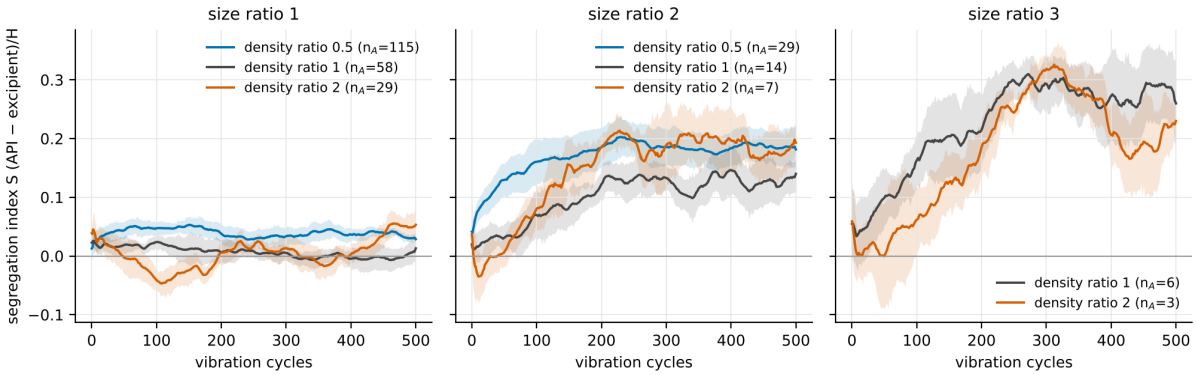

Figure 10. Segregation index S against vibration cycle for each size ratio (panels) and density ratio (lines); mean of five runs, light bands ±1 SE, 5-point smoothing. S=0 is well mixed; S>0 means the API-like discs have moved upward.

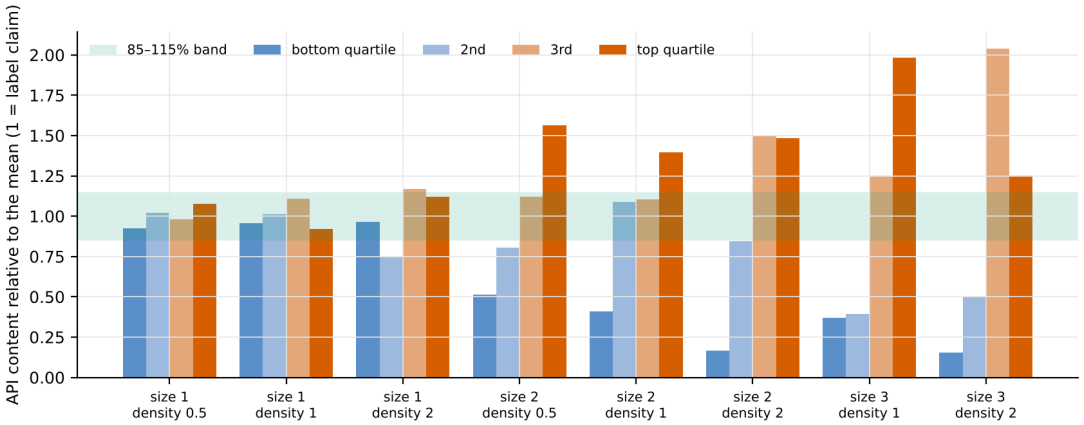

Figure 11. Relative API content of the four equal-mass layers at cycles 400–500. The shaded band is 0.85–1.15 of the target content.

Density alone did nothing detectable. With the API-like discs the same size as the excipient (size ratio 1), neither a lighter (0.5) nor a heavier (2) density produced segregation in 500 cycles: over the 15 runs ΔS=-0.008 ± 0.017 (p=0.67), and S ends within about ±0.05 of zero in all three cases (density ratio 1 is the neutral control). The layer contents stay close to 1 (bottom and top 0.92–1.12), and the deviations seen are of the size expected from random sampling of only 29–115 discs. The three densities are not distinguishable (Kruskal–Wallis on S late, p=0.25).

Size produced clear upward segregation. With a size ratio of 2 the API-like discs rise: ΔS=0.14 ± 0.03 over the 15 runs (p=0.0003), significantly more than at size ratio 1 (Mann–Whitney p=0.0003); S reaches 0.13–0.20 by about 200 cycles and then stays level. By cycles 400–500 the top layer holds 1.4–1.6 times the target API content and the bottom layer only 0.2–0.5 times (Figure 11). At size ratio 3 the effect is larger (ΔS=0.19 ± 0.07, p=0.026 against 0; against size ratio 1, p=0.007) but rests on very few API discs (3–6 per bed), so the numbers fluctuate strongly from run to run (the top layer holds up to 2 times the target content). Within size ratio 2 the three densities give similar S (late 0.19, 0.13 and 0.18 for density ratio 0.5, 1 and 2; Kruskal–Wallis p=0.40): at this resolution a lighter or heavier API of the same size rises about as far as one of equal density. The bottom layer is emptied most for the heavy API (0.17 against 0.51 for the light one), but that case has only seven API discs, so we do not read this as a density effect.

Content uniformity. Comparing layer contents with the 0.85–1.15 band is a much cruder measure here. The share of layers outside the band is 65–100% at the end for size ratios 2 and 3, but it is already 45–100% at the start, where the small number of API discs (3–29) alone gives sampling scatter; at size ratio 1 it is 20–55%, with no systematic change. The layer-failure fraction is therefore dominated by sampling noise and is not a reliable indicator on its own; the segregation index S and the top-to-bottom ratio, which use all discs, are more informative. These are properties of the mixture in this 2D model after a fixed number of cycles; they are not a prediction of the content uniformity of any real blend.

3.11 Summary of the additional experiments

- **E2.** The waiting time is flat (≈70–90 cycles) for ρ≤1.5, rises steeply for ρ=2–2.75, and no surfacing was seen at ρ=3.5–8 (4 runs × 1,500 cycles each); the disc does not sink to the floor (relative depth 0.34 → 0.26).
- **E3.** At Γ=2 the disc almost never surfaces; for Γ=3–8 the wait at ρ=2 is independent of Γ within the precision available (≈140–250 cycles) while at ρ=1 it tends to fall with Γ (≈140 → 39 cycles; not significant).
- **E1.** In a vibrated binary mixture, a size ratio of 2–3 moves the minority component upward significantly; a density ratio of 0.5 or 2 at equal size does not, in 500 cycles.

4 Discussion

4.1 What density does: two baselines

There are two reference points for interpreting the density effect. The first is the exit time of a disc of the same size with no density difference (ρ=1) (≈82 cycles): this is the pure *size* effect. The second is the exit time of a neutral tracer of the same size as the small discs (R=1, ρ=1) (≈514 cycles): this is the time for 'a particle wandering randomly inside the bed' to reach the surface. Within this framework the ordering is as follows:

Table 13. Ordering of waiting times together with the baselines.

Condition	Mean/median wait (cycles)	Comment
R=4, ρ=1	≈82 (median)	size effect only: fast rise
R=4, ρ=2	≈143	density slows the rise by ≈1.7×
R=4, ρ=2.25	≈250	
R=4, ρ=2.5	≈510	indistinguishable from the R=1 neutral tracer (≈514)
R=4, ρ=2.75	≈2,600	≈5× slower than the neutral tracer
R=4, ρ=2.875	≈20,500 (single event)	very uncertain
R=4, ρ=2.9–3.1 (pooled)	≈107,505 (3 events / 38 runs)	exit very rare; cannot be separated by ρ

Density steadily reduces the geometric advantage provided by size; at ρ≈2.5 this advantage effectively disappears, and at higher ρ the large disc exits even *later* than a small tracer. This is not a full 'reverse BNE' (the disc does not sink to the bottom, see Section 3.4); it is rather a regime in which floating becomes progressively harder.

4.2 Relation to the literature

- **Size effect.** That an R=4 disc surfaces quickly even at ρ=1 is consistent with the density-independent geometric rise mechanism of Rosato and colleagues [1]. That exits continue up to ρ=2.75 is also consistent with the observation that a large particle can rise even though it is denser.
- **Depth dependence.** Pöschel and Herrmann reported that in 2D simulation the probability of rising by convection depends strongly on the depth of the large particle [3]. In this study the disc staying for a long time at a relative depth of ≈0.4 for ρ≥2.9 is qualitatively compatible with the idea of a 'depth threshold'; however, this study did not test the link directly.
- **Density dependence and air.** Möbius and colleagues (3D, with air) reported that the density-dependent rise time is not monotonic and that this dependence arises from air drag [7, 8]. Our model has no air, yet it shows a strong and monotonic density

dependence. The two results cannot be compared directly (2D/3D, container geometry, Γ, non-rotation); the difference is striking and deserves separate investigation.

- **Reverse buoyancy.** Shinbrot and Muzzio reported that in deep beds and at large amplitude, large heavy particles rise while light particles sink [4]. In this study the heavy disc rises *more slowly*; because the regimes (depth, amplitude, dimensionality, air) differ, this should not be read as a contradiction.
- **Percolation–condensation theory.** The theory of Hong and colleagues is for mixtures with many large particles [5]; this study contains a single large disc and does not directly test the competition between percolation and condensation.

4.3 Possible mechanisms (hypotheses; not tested in this study)

The explanations below are hypotheses that *may be compatible* with the data; none has been distinguished by the measurements of this study.

- **H1 – Void-filling ratchet.** When the large disc rises and falls with the bed, the voids that open beneath it are filled by small discs and the disc climbs step by step [1]. As density increases, the disc's lift relative to the bed decreases; the ratchet effect weakens, and together with the observation of Weibull k>1 (hazard increasing with time) one may think that the disc must first reach a sufficient ‘height’.
- **H2 – Convection capture.** Bed convection arises through the effect of wall friction [2, 8]; a disc that is high enough is carried to the surface by the rising stream, whereas a disc lower down cannot be carried [3]. The slowing of the exit at low μ_w (Section 3.6) and the long lower-middle state at $\rho \geq 2.9$ are compatible with this hypothesis.
- **H3 – Effective density (granular Archimedes).** If the vibrated bed behaves on average with an ‘effective density’ lower than the small-disc density, then at a certain ρ the disc exceeds the bed's effective density and stays suspended rather than floating. Whether the transition near $\rho \approx 3$ is related to such a criterion is unknown.

Proposed tests. Measuring the local packing fraction and the velocity field (convection) of the bed; scanning Γ and ω; scanning container width and particle number; repeating the exit criterion with thresholds other than p_{75}.

4.4 Meaning of the critical behaviour

The censored exponential model assumes that ln λ decreases linearly with ρ (a ‘barrier’ growing linearly with ρ) and is rejected. The critical form is $\ln \lambda = \ln A + \beta \ln(\rho_c - \rho)$: the effective barrier diverges logarithmically as ρ approaches ρ_c. The data distinguish these two forms at $\rho \leq 2.75$. The long runs ($\rho = 2.9$–3.1; 10,000 cycles) refuted the earlier estimate $\rho_c \approx 2.9$: exits occur at $\rho = 3$. However, at $\rho \geq 2.9$ the data still cannot distinguish ‘a genuine ρ_c exists’ from ‘the exit rate is very small but non-zero’: there are no events in 12 runs and 1.2×10^5 cycles at $\rho = 3.05$ and 3.1; the pooled rate ($\approx 9 \times 10^{-6}$ /cycle) predicts ≈1.1 events for these two ρ, so zero events is not excluded ($P \approx 0.33$). Telling them apart needs $\geq 10^6$ cycles and ≥20 seeds over ρ=3–3.4.

4.5 A methodological lesson

The most instructive point in this study is that the sharp ‘threshold’ seen in the 500-cycle window was a censoring artefact (Section 3.1). The same lesson repeated at ρ=3: the ‘no exit’ result at 28,000 cycles was broken at 86,000 cycles (Section 3.2). Similarly, computing the mean only from runs that surfaced, counting the run limit as a time, or choosing the wrong baseline as a control (only the ρ=1 large disc instead of the R=1 tracer) could have led to a misinterpretation of the effect of ρ. Censoring must be modelled explicitly, and it must be stated against which baseline each new result is evaluated.

4.6 What this does and does not say about pharmaceutical blends and asteroid regoliths

The Brazil nut effect is often cited in two applied settings: the segregation of components in vibrated or handled powder blends in pharmaceutical manufacturing, and the surface structure of small rubble-pile asteroids. We tested the first with a toy mixture (Section 3.10) and can comment on the second only by analogy. The following points separate what this study shows from what it does not.

- **Pharmaceutical blending.** In the mixture test, API-like particles with a size ratio of 2–3 migrated upward and were depleted at the bottom, whereas a density difference at equal size produced no detectable segregation within 500 cycles. The direction agrees with the pharmaceutical literature, in which a size difference (a ratio of about 1.3 or more in free-flowing powders) is regarded as the main driver of sifting and vibration-induced segregation, and acceptance is judged against 85–115% of the label claim. It is only an agreement of direction. Real blends are three-dimensional, often cohesive and fine, can fluidise in interstitial air and are handled by flow, tumbling and filling rather than by a sinusoidal floor; a density effect that is absent here (no air, no cohesion, one drive) can well appear there. The model should be used to generate hypotheses, not to set blending limits.
- **Heavy intruder in a bed.** The main result, that a single dense intruder surfaces after a waiting time that grows steeply with density and that it hovers part-way down instead of sinking, concerns one large object, not a dilute population. It does not translate directly into a rule for the fate of a dense API particle in a blend, which is surrounded by similar particles and can also be carried by bulk convection.
- **Asteroids and low gravity.** Rubble-pile asteroids have low gravity and very low cohesion, and their surfaces have been shaken by impacts and tidal effects. Our parameter Γ is the ratio of the peak floor acceleration to g, so reducing g at a fixed shaking amplitude raises Γ; in that sense Γ=3–8 corresponds to weaker gravity. In this range the waiting time at ρ=2 did not depend on Γ and at ρ=1 it fell only slowly (Section 3.9), so these runs give no support for a strong gravity dependence of the rise in this model; they do not test cohesion, which dominates at the smallest scales, nor the geometry, rotation and varied sizes of real regolith. Observed surface patterns on bodies such as Itokawa or Bennu have several competing explanations (granular segregation, re-accumulation after disruption, the erasure of craters, low cohesion with loose packing), and the present simulations neither favour nor exclude any of them.
- **What the study does support.** In a simple controlled model, (i) the time for a large disc to rise is a random variable with a long, density-dependent tail that must be measured with censoring-aware statistics; (ii) size acts as a strong segregation driver, whereas density at fixed size does not in the tested window; (iii) a ‘no result’ after a fixed number

of cycles can be an observation-window artefact. Each of these is a candidate to check in experiments or in more realistic simulations, not a conclusion about a real material.

5 Limitations

- **Model.** Two-dimensional with non-rotating particles; no air. It cannot be generalised directly to real (3D, rotating, air-filled) systems.
- **Sample.** There are 6–8 runs per condition; small effects and a medium-sized layout effect cannot be ruled out. No multiple-comparison correction was made.
- **Parameter space.** Γ=4, ω=2, W=32, N=230, K=1500, ζ=0.1 is a single set; ρ_c may depend on these parameters. Numerical sensitivity (Δt, K) was not scanned systematically in this study, and the simulation was not cross-validated with an independent DEM package.
- **Exit criterion.** The threshold is the 75th percentile; no other threshold (e.g. 50th or 90th) was tested.
- **Censoring.** At $\rho \geq 2.875$ the data are mostly censored (4 events in 44 runs); τ and ρ_c are weakly constrained in this region. Data from earlier stages were collected with different run limits; the analysis handles censoring correctly, but the protocol is not entirely homogeneous.
- **Mechanism.** The explanations in Section 4.3 are unmeasured hypotheses.
- **Additional experiments (E1–E3).** These have only 4–5 runs per condition and short or single windows (500–1,500 cycles); the heavy side ($\rho \geq 3.5$), the weak-drive case (Γ=2) and the mixtures with size ratio 3 (3–6 API discs) are especially uncertain. The mixture test is a two-dimensional toy: no air, no cohesion, one drive, one mass fraction, and a start that is only approximately well mixed. No correction was made for multiple comparisons.
- **Applications.** Section 4.6 discusses possible relevance to pharmaceutical blends and asteroid regoliths; none of this study's runs tests those systems.

6 Conclusions and future work

6.1 Summary of the conclusions

The main conclusions of this study are:

- The first-passage time of the large disc increases rapidly with density: ≈143 (ρ=2), 250, 510 and 2,600 (ρ=2.75) cycles; at $\rho \geq 2.9$ surfacing is very rare but not zero (3/38 runs, pooled $\tau \approx 107{,}505$ cycles).
- The time–density relation is not log-linear (p=0.0003); a form diverging near $\rho_c \approx 3.11$ (95% interval 3.03–3.41) is better by AIC (ΔAIC=16.3); a genuine divergence has not been shown. The earlier estimate $\rho_c \approx 2.9$ is invalid.
- Surfacing is not memoryless (Weibull $k \approx 2$–3.6).
- The waiting time of the R=4 disc at ρ=2.5 is indistinguishable from that of a neutral tracer of the same size as the small discs; the pure size effect (ρ=1) is ≈82 cycles.
- Discs that do not surface do not sink to the bottom; the relative depth decreases smoothly with ρ (≈0.61 → 0.37) and at $\rho \geq 2.9$ the disc stays suspended at ≈0.4 for a long time.
- Wall friction slows the exit only for $\mu_w \leq$ about 0.5; no significant effect was found for R=3–5 or for the initial layout.

- Extended range (E2): the waiting time is flat (≈70–90 cycles) for $\rho \leq 1.5$; no run surfaced at $\rho = 3.5$–8 (4 runs each, 1,500 cycles), and the disc stays suspended at a relative depth falling slowly from 0.34 to 0.26 rather than sinking to the floor.
- Drive strength (E3): at $\Gamma = 2$ the disc almost never surfaces; for $\Gamma = 3$–8 the wait at $\rho = 2$ is independent of Γ within the available precision, while at $\rho = 1$ it tends to fall with Γ (not significant).
- Binary mixture (E1): in a 2D vibrated mixture, a size ratio of 2–3 moves the minority component upward ($\Delta S \approx +0.14$ to $+0.19$, $p < 0.03$), whereas a density ratio of 0.5 or 2 at equal size produced no detectable segregation in 500 cycles.

6.2 What we concluded, in plain words

The central finding is about time, not about a threshold. The heavier the large disc is, the longer it takes to reach the surface of the vibrated bed. The mean waiting time grows from about 143 vibration cycles at $\rho = 2$ to about 2,600 at $\rho = 2.75$, a factor of roughly 18 for a density change of only 0.75; on the simple exponential description every extra 0.25 of density lengthens the wait about 4.6-fold. At first sight the experiments seemed to show a sharp boundary: with a 500-cycle limit, 6 of 6 layouts surfaced at $\rho = 2.25$, 3 of 6 at $\rho = 2.5$ and none at $\rho = 2.75$ or 3. That boundary was produced by the observation window, not by the physics. When the same layouts were followed for longer, the runs that had ‘failed’ at $\rho = 2.75$ surfaced after 700–1,400 cycles. A run that has not finished is not a run that cannot finish.

Surfacing behaves like a random event, but a fairly regular one. If the disc had a fixed small chance of surfacing in every cycle (a memoryless process), the spread of the waiting times would be as large as their mean (coefficient of variation 1; Weibull shape $k = 1$). We find $k \approx 2$–3.6, a coefficient of variation of roughly 0.3–0.5. The waiting times are therefore more tightly clustered than pure chance would give: the disc seems to have to ‘build up’ to the surface, for example by climbing in stages as voids open beneath it, rather than being drawn from a lottery with constant odds. A practical consequence is that, where we can measure it, a disc that has waited a long time is *more* likely, not less, to surface soon.

Density cancels the advantage of size. A large disc of the same density as its surroundings ($\rho = 1$) surfaces in about 82 cycles, much faster than a small tracer disc, which needs about 514. Adding density eats into this advantage step by step: at $\rho = 2.5$ the large disc is no faster than the tiny tracer, and at $\rho = 2.75$ it is about five times slower. In this sense density and size act against each other, and the net result depends on which wins.

A heavy disc that does not surface does not sink either. We expected a heavy disc to fall to the bottom (the ‘reverse’ Brazil nut effect). It does not. Its average height, as a fraction of the height of the bed surface, declines smoothly from about 0.61 at $\rho = 2$ to about 0.4 at $\rho \geq 2.9$ (0.37 at $\rho = 3.1$) and varies little between runs: the disc hovers part-way down, makes short upward excursions (to about 22–23 against a surface line near 24.6 in the 4,000-cycle runs at $\rho = 3$) and falls back. In our model heavier therefore means ‘rides lower and surfaces more rarely’, not ‘sinks’.

What happens at the highest densities is not settled. Over $\rho = 2.9$–3.1 we followed 38 runs for 322,515 cycles in total and saw 3 surfacings. This proves that surfacing at $\rho = 3$ is possible, and so retracts our earlier estimate of a limiting density near 2.9, which was a censoring artefact. It does not tell us whether the rate keeps falling smoothly or vanishes at a critical density: a form that diverges near $\rho_c \approx 3.11$ describes the data better than a

straight line in log-time (ΔAIC=16.3), but a three-parameter curve fitted to ten densities with only three events in the decisive region cannot distinguish a real divergence from a very small but finite rate. We state the result as ‘better description’, not as ‘critical point found’.

The other factors matter less, or only in a corner of the parameter space. Wall friction slows surfacing only when it is low (μ_w ≲ 0.5), which fits the idea that wall-driven convection helps to carry the disc up (a hypothesis, not tested here). Changing the disc radius from 3 to 5, or changing the initial arrangement of the small discs, produced no statistically significant effect with 6 seeds per condition; a modest effect cannot be excluded.

What the additional experiments add. Making the disc lighter than ρ=1 does not speed it up further; making it much heavier (ρ=3.5–8) did not make it surface in 1,500 cycles, but did not make it sink either. A weak drive (Γ=2) stops the process, a stronger drive speeds up the light disc but not the heavy one. In a mixture, a size difference of 2–3 drives the minority component upward within a few hundred cycles, while a density difference at equal size did not. Taken together, in this model size matters more robustly than density, and density matters mainly by slowing a large intruder that would otherwise rise quickly.

How far these conclusions go. They describe a two-dimensional, air-free model of non-rotating discs with one set of vibration parameters and 6–8 layouts per condition. They are exploratory results about *this* model. They do not by themselves say how a real three-dimensional granular mixture, with air, cohesion and irregular grains, would behave.

6.3 The mathematics behind the conclusions

This section sets out the statistical machinery used above, so that each number in Section 6.1 can be traced to a formula. Notation: T is the first-passage (surfacing) time in vibration cycles; for each run i we observe $x_i=\min(T_i, c_i)$, where c_i is the run limit, and $\delta_i=1$ if the disc surfaced ($\delta_i=0$ if the run was censored).

M1 Waiting time as a first-passage problem

The survival function (probability of not having surfaced by cycle t) and the hazard rate h(t) (instantaneous chance of surfacing, given that it has not yet) are related by

$$S(t) = P(T > t) = \exp\left[-\int_0^t h(s)\, ds \right] \qquad (1)$$

A constant hazard $h(t)=\lambda$ gives the exponential law, whose mean is $\tau=1/\lambda$ and which is *memoryless*: the time still to wait does not depend on the time already waited.

$$S(t) = e^{-\lambda t}, \quad E[T] = \tau = 1/\lambda, \quad P(T > t+u \mid T > t) = P(T > u) \qquad (2)$$

M2 Likelihood with censoring; the exit rate and its exact interval

A surfaced run contributes the density $f(x_i)$ to the likelihood and a censored run contributes $S(x_i)$, the probability of not having surfaced by the limit:

$$L = \prod_i f(x_i)^{\delta_i} \cdot S(x_i)^{1-\delta_i} \qquad (3)$$

For the exponential law this reduces to $\ln L = d \ln \lambda - \lambda E$, where $d=\Sigma\delta_i$ is the number of events and $E=\Sigma x_i$ the total *exposure* (cycles watched, whether or not the disc surfaced). Maximising gives

$$\hat{\lambda} = d / E, \quad \mathrm{Var}(\hat{\lambda}) \approx d / E^2 \qquad (4)$$

Because $2\lambda E$ follows a χ^2 distribution with 2d degrees of freedom (a Poisson count of events), an exact $(1-\alpha)$ interval is

$$\chi^2_{\alpha/2}(2d) / 2E \;\le \lambda \le\; \chi^2_{1-\alpha/2}(2d+2) / 2E \qquad (5)$$

With no events the lower limit is zero and only the one-sided bound $\lambda < -\ln \alpha / E$ exists; for $\alpha=0.05$ this means $\tau > E/3.0$, which is the ‘τ at least’ column in Tables 4 and 5. **Worked example (pooled ρ≥2.9).** d=3, E=322,515 cycles, so $\hat{\lambda}=9.30\times10^{-6}$ per cycle and $\hat{\tau}=107{,}505$ cycles; the interval (5) gives λ from 1.92 to 27.2×10^{-6}, i.e. τ from 36,786 to 521,302 cycles, a range of more than a factor of 10 because d is only 3.

Why censoring must be modelled. Averaging only the runs that surfaced estimates $E[T \mid T\le c]$, not τ. For the exponential law

$$E[T \mid T \le c] = \tau - c / (e^{c/\tau} - 1) \qquad (6)$$

With τ=2,600 (ρ=2.75) and a limit c=500, only $1-e^{-500/2600}\approx17\%$ of runs surface, and the mean of those that do is about 240 cycles, one tenth of the true mean. The same numbers explain the ‘threshold’ of Section 3.1: the chance that all six runs at ρ=2.75 fail to surface within 500 cycles is $e^{-6\cdot500/2600}\approx0.32$, which is not unusual.

M3 The Weibull form and what k tells us

The Weibull law generalises the exponential by letting the hazard change with time:

$$h(t) = (k/s)(t/s)^{k-1}, \quad S(t) = \exp[-(t/s)^k], \quad E[T] = s\,\Gamma(1+1/k) \qquad (7)$$

The coefficient of variation depends on the shape k alone:

$$CV = \sqrt{[\,\Gamma(1+2/k) / \Gamma(1+1/k)^2 - 1\,]} \qquad (8)$$

Table 14. Weibull shape, hazard and coefficient of variation (CV).

Shape k	Hazard h(t)	CV	Meaning
1	constant	1.00	memoryless; waiting time as variable as its mean
2	grows linearly with t	0.52	the longer the wait, the more likely the exit
3	grows as t^2	0.36	tightly clustered exit times
3.6	grows as $t^{2.6}$	0.31	largest fitted value (ρ=2)

The fitted shapes (k≈2–3.6; Section 3.3) therefore mean that surfacing is much more regular than a constant-odds lottery, and the hypothesis k=1 is rejected by a likelihood-ratio test at ρ=2, 2.25 and 2.5 (p=0.0004, 0.014, 0.004; at ρ=2.75, k=1.9 but p=0.12, with low power).

M4 Density laws and model comparison

The simple law lets the logarithm of the exit rate fall linearly with density; the critical law lets it fall as a logarithm of the distance to a limiting density ρ_c:

$$\text{(i) } \ln\lambda(\rho) = a + b(\rho - 2.5); \quad \text{(ii) } \ln\lambda(\rho) = \ln A + \beta \ln(\rho_c - \rho) \qquad (9)$$

Under (i) the wait is $\tau(\rho)=e^{-a}\,e^{-b(\rho-2.5)}$; the fit gives b=-6.08, so each +0.25 in ρ multiplies τ by $e^{0.25|b|}\approx4.6$. Under (ii) $\tau \propto (\rho_c-\rho)^{-\beta}$, a power-law divergence. The two laws differ in how steeply ln τ rises with ρ. In (i) the slope is the constant |b|; in (ii) it is

$$s(\rho) = -d\ln\lambda / d\rho = \beta / (\rho_c - \rho) \qquad (10)$$

Table 15. Local slope of ln τ against ρ for the two laws (ρ_c=3.115, β=3.11). The data are better described by a slope that steepens with ρ.

ρ	2	2.5	2.75	3
slope s(ρ), critical law (ii)	2.8	5.1	8.5	27.0
slope |b|, exponential law (i)	6.1	6.1	6.1	6.1

Lack-of-fit test. Compare law (i) with a saturated model that gives every density its own rate $\hat{\lambda}_\rho = d_\rho / E_\rho$. With $\ln L(\lambda) = \Sigma_\rho (d_\rho \ln \lambda_\rho - E_\rho \lambda_\rho)$,

$$\Lambda = 2 \left[\ln L(\hat{\lambda}_{sat}) - \ln L(\hat{\lambda}_{fit}) \right] \sim \chi^2_{m-2}, \quad m = 10 \qquad (11)$$

we find Λ=28.83 on 8 degrees of freedom, p=0.0003: the straight-line law is rejected.

Comparing (i) and (ii). The Akaike criterion penalises the number of parameters k_p:

$$AIC = 2k_p - 2 \ln \hat{L} \qquad (12)$$

For (i) k_p=2 and AIC=547.1; for (ii) k_p=3 and AIC=530.8. The difference ΔAIC=16.3 corresponds to a relative likelihood $e^{\Delta AIC/2}$≈3,530 in favour of (ii). Two cautions apply. First, ρ_c sits just beyond the densest data (95% profile interval 3.03–3.41), where the usual asymptotic χ^2 theory is only approximate. Second, models (i) and (ii) differ chiefly at ρ≥2.75, where there are only a handful of events, so the preference for (ii) is a statement about *description*, not proof of a divergence.

M5 Probability statements used in the text

When a pooled rate $\hat{\lambda}$ is applied to a new exposure E′, the number of events is Poisson distributed with mean $\hat{\lambda}E'$, so the probability of seeing none is

$$P(0 \text{ events}) = e^{-\hat{\lambda}E'} \qquad (13)$$

At ρ=3.05 and 3.1 together E′=120,000 cycles, and with $\hat{\lambda}=9.30\times10^{-6}$ this is $e^{-1.12}$≈0.33, so seeing no events there does not conflict with the pooled rate. For the log-linear law (i), the expected number of events over ρ=2.9–3.1 is 9.6, and observing 3 or fewer has Poisson probability P(≤3)=0.014; surfacing at the highest densities is therefore rarer than the straight-line extrapolation predicts.

Reading the p-values. The wall-friction, layout and radius comparisons (Sections 3.5–3.7) rest on 6 runs per condition; the Wilcoxon, Kruskal–Wallis, Mann–Whitney and Kendall tests have low power and were not corrected for multiple comparisons, so they are exploratory.

6.4 Future work

- Runs at ρ=3–3.4 that are much longer (≥10^6 cycles) and with many more seeds (≥20), to test whether the exit rate at ρ≥2.9 declines with ρ, that is, whether a genuine ρ_c exists.
- Scans of Γ, ω, N and W; a $\rho_c(\Gamma)$ map; different exit thresholds.
- Measuring the convection field and local packing fraction to distinguish hypotheses H1–H3.
- Adding rotating particles, three dimensions and (if relevant) air to the model; cross-validation with an independent DEM package.
- Longer and more numerous mixture runs, with realistic size distributions, more API discs at large size ratios, cohesion and 3D geometry, before any comparison with pharmaceutical content-uniformity data.
- A Γ scan at ρ=2.5–3.1 to see whether the critical density depends on the drive, and runs with reduced gravity or added cohesion to address the asteroid-regolith questions of

Section 4.6.

References

[1] Rosato, A., Strandburg, K. J., Prinz, F., Swendsen, R. H. (1987). Why the Brazil nuts are on top: Size segregation of particulate matter by shaking. *Physical Review Letters* 58, 1038.

[2] Knight, J. B., Jaeger, H. M., Nagel, S. R. (1993). Vibration-induced size separation in granular media: The convection connection. *Physical Review Letters* 70, 3728.

[3] Pöschel, T., Herrmann, H. J. (1995). Size segregation and convection. *Europhysics Letters* 29, 123.

[4] Shinbrot, T., Muzzio, F. J. (1998). Reverse buoyancy in shaken granular beds. *Physical Review Letters* 81, 4365.

[5] Hong, D. C., Quinn, P. V., Luding, S. (2001). Reverse Brazil nut problem: Competition between percolation and condensation. *Physical Review Letters* 86, 3423.

[6] Breu, A. P. J., Ensner, H.-M., Kruelle, C. A., Rehberg, I. (2003). Reversing the Brazil-nut effect: Competition between percolation and condensation. *Physical Review Letters* 90, 014302.

[7] Möbius, M. E., Lauderdale, B. E., Nagel, S. R., Jaeger, H. M. (2001). Brazil-nut effect: Size separation of granular particles. *Nature* 414, 270.

[8] Möbius, M. E., Cheng, X., Eshuis, P., Karczmar, G. S., Nagel, S. R., Jaeger, H. M. (2005). Effect of air on granular size separation in a vibrated granular bed. *Physical Review E* 72, 011304.

[9] Kudrolli, A. (2004). Size separation in vibrated granular matter. *Reports on Progress in Physics* 67, 209.

[10] Shinbrot, T. (2004). The Brazil nut effect — in reverse. *Nature* 429, 352.

[11] Wikipedia, “Granular convection” article (unsolved-problem tag; for context only, not a primary source).

[12] Cundall, P. A., Strack, O. D. L. (1979). A discrete numerical model for granular assemblies. *Géotechnique* 29, 47.

[13] Kaplan, E. L., Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* 53, 457.

[14] Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19, 716.

[15] Weibull, W. (1951). A statistical distribution function of wide applicability. *Journal of Applied Mechanics* 18, 293.

[16] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin* 1, 80.

Appendix A Excel workbook and sheet guide

The accompanying file *brazil_nut_all_findings_en.xlsx* contains 30 sheets; the first sheet (*Contents*) is a clickable index. All calculations from raw data to summaries (live formulas) have been recalculated without errors.

Table A1. Structure of the Excel workbook.

Sheet	Content
Contents	Linked index of all sheets
Methods	Model, exit criterion, censoring, group definitions, limitations
Data_Campaign	All raw runs (group, ρ, μ_w, R, seed, exit, heights); event/exposure formulas
Summary_Density	Event count, exit rate, τ and 95% CI by ρ (COUNTIFS/SUMIFS/CHIINV); scatter chart
Exit_Matrix	Exit times, seed × ρ
Friction_Layout	Wall friction × layout (ρ=2 and 2.25)
Radius_Control	R=3/4/5 and the R=1 control (G4, G4b)
Exponential_Fit	Parameters of the exponential and critical models, Weibull table
Earlier_Runs	Runs from the previous stage collected in one table
Ext_Density	Extended density range ρ=0.25–8: runs and summary by ρ (formulas)
Gamma_Scan	Drive-strength scan (Γ=2–8) at ρ=1 and 2: runs and summary (formulas)
Mixture_Test	Binary-mixture test: segregation index and layer content by size and density ratio (formulas)
Charts	The 11 figures of this report
S1 …, S2 …, S3 …, S4 … (17 sheets)	All tabs of the four earlier working sheets (records, fine scans, waiting-time test, long run at ρ=3, comments); cross formulas re-linked to the new sheet names

Appendix B Raw exit times (group G1)

Table B1 gives the exit time (cycles) of every run in the density scan. The “>” sign marks a run that did not surface at its run limit (right-censored); an empty cell means that no run was made with that seed.

Table B1. Group G1: exit time (vibration cycles), seeds 101–108.

ρ	101	102	103	104	105	106	107	108
2	154.5	217.9	200.9	94.7	108.8	147.1	124.2	96.1
2.25	156.9	91.3	287.0	233.1	183.0	246.9	271.7	526.4
2.5	318.7	257.1	535.1	523.4	499.5	1,001.2	522.9	421.9
2.75	735.7	>3,000	925.5	1,366.5	1,469.6	2,274.0	2,855.6	>3,000
2.875	>4,000	>4,000	>4,000	>4,000	>4,000	529.8		
3	>4,000	>4,000	>4,000	>4,000	>3,000	>3,000	>3,000	>3,000

Table B2. Long runs (limit 10,000 cycles): exit time (vibration cycles), seeds 201–206.

ρ	201	202	203	204	205	206
2.9	>10,000	>10,000	>10,000	>10,000	>10,000	>10,000
2.95	>10,000	>10,000	>10,000	6,246.4	>10,000	>10,000
3	>10,000	>10,000	9,268.9	>10,000	>10,000	8,999.7
3.05	>10,000	>10,000	>10,000	>10,000	>10,000	>10,000
3.1	>10,000	>10,000	>10,000	>10,000	>10,000	>10,000

The raw values of the other groups (G2, G3, G4, G4b) are on the Data_Campaign sheet of the Excel workbook.

Appendix C The earlier stage (S1–S4) and this report

Before this report, four working sheets (Brazil Nut Experiments, Fine Density Scan, Waiting-Time Test, Fine ρ Scan and Page Update) carried out the ρ scan, the μ_w validation, paired runs of 500–1500 cycles and long runs of 4000 cycles at ρ=3. All tabs of those sheets were added to the Excel workbook with the prefixes S1–S4. The 500-cycle results of that stage gave the impression of a ‘threshold’, which turned out to be a censoring artefact, as shown in Section 3.1. The ρ=3 runs from the earlier stage (seeds 101–104, 4000 cycles) were included in group G1 of this report; the other earlier runs are listed separately on the *Earlier_Runs* sheet. **Update (28 November 2026).** The first version reported ‘no exit at ρ=3’ and $\rho_c \approx 2.9$. To test this, 30 runs of 10,000 cycles were made at ρ=2.9, 2.95, 3, 3.05 and 3.1 with six new seeds each (201–206); events were seen at ρ=2.95 and ρ=3, so these two findings of the first version are invalid and have been corrected in this version.

List of figures: Figure 1 set-up; 2 Kaplan–Meier; 3 waiting time and models; 4 relative depth; 5 example trajectories; 6 friction; 7 radius and control.