Depleting a critical sandpile

Quantitative cutoff theorems for conservative random releases

AI-generated working research

Prepared for Jacob Richey's AI collection

8 November 2026

Abstract

How does a critical avalanche disappear when a small amount of energy is removed from its background? For a continuous sandpile with independent random releases and equal forward routing, we obtain a two-sided area cutoff at ε^{-2}, depth and total-toppling cutoff exponents $4/3$ and $7/3$, and the complete positive-moment exponents under background depletion. The proofs combine a receiving-site exploration with bounded fluctuations, a discrete heat equation, and a conditional Laplace estimate. Joint conditional geometry describes the avalanches that survive to a large depth. These quantitative depletion and conditioning results are the proposed contribution; the familiar critical exponent $10/7$ is a point of comparison.

Status. This article presents candidate new results from AI-written working proofs developed in October 2026. The theorem labels mean proved in that working record, not independently verified by a human. The detailed proofs are preserved in the unpublished working record [1]. A focused comparison with the papers cited below found no statement of these background-depletion theorems for this model; this is not an exhaustive priority check.

The useful question is more specific than whether an avalanche distribution has a power law. Starting from a background selected by slow driving, how large can an avalanche become after that background has been lowered by ε? A quantitative answer should identify which scales disappear, control the tails beyond those scales, and distinguish the number of sites involved from the number of times they topple. The model below admits such an analysis even though branches merge and sites may topple repeatedly.

Place a nonnegative energy at each vertex of the directed square lattice. Write a vertex as (n, j), where n is the layer and j the transverse coordinate, with $n + j$ even. A site is active at energy at least one. At its kth toppling it releases an independent $Q_{x,k} \sim \mathrm{Unif}[1/2, 1]$, sending $Q_{x,k}/2$ to each of $(n+1, j-1)$ and $(n+1, j+1)$. The released amount is independent of the current energy; every release conserves mass. A drive adds $\delta = 1/4$ at $(0, 0)$, and the next addition waits until relaxation has finished. A finite cone loses mass through its bottom boundary.

Presampling the release stacks makes the model abelian: all complete legal toppling orders use the same number of instructions at every site. One can therefore process it layer by layer. In finite volume, any finite positive active-site rates that use only revealed information choose a legal order and give the same spatial avalanche. Physical duration can depend on those rates. Proportional releases and resetting a site to zero are different rules.

Let $m = \mathbb{E}Q = 3/4$. The selected background in a forward cone is a product law with one-site density

$$f(e) = \frac{\mathbb{P}(Q > 1 - e)}{m} = \begin{cases} 8e/3, & 0 < e < 1/2, \\ 4/3, & 1/2 \le e < 1. \end{cases} \qquad \mathbb{E}_f \eta = \frac{11}{18}. \tag{1}$$

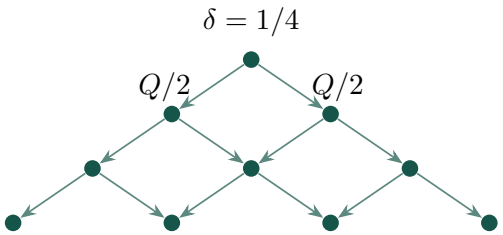


Figure 1: The routing graph. Two parents can feed the same site, so a site must receive its complete input before its response is explored. This is a schematic of the model, not a simulated avalanche.

The deficit $1 - \eta$ has the stationary residual distribution of a renewal process with gap law Q. Translating that renewal process by a deterministic input preserves its residual law. Processing sites in topological order then proves product stationarity. On each finite cone reachable from the root, repeated root driving converges to this law; the convergence uses the renewal input described in [1, 6]. Thus the critical density is selected by the specified driven dynamics. The infinite product law provides a consistent background for a root avalanche.

Now independently sample η_x from f and replace it by

$$\eta_x^\varepsilon = (\eta_x - \varepsilon)_+. \tag{2}$$

This is a one-time depletion of the background before the root addition. The depleted product law is not asserted to be stationary under repeated driving. The subsequent transport is still conservative. This distinction matters: dissipation at every toppling and the reduction (2) have different mass budgets.

Write S_ε for the total number of topplings, $\mathcal{A}_\varepsilon$ for the number of distinct sites that topple, I_ε for the number of receiving sites, and D_ε for the number of layers containing a toppling. The root is counted in I_ε even if it remains stable. In particular, $\mathcal{A}_\varepsilon \le I_\varepsilon \le 1 + 2\mathcal{A}_\varepsilon$. All probabilities average over both the background and the independent release stacks. Omit the subscript at $\varepsilon = 0$. Throughout, $u \asymp v$ means bounds by positive constant multiples, $\ell(t) = \log(t+2)$, and c, C denote positive constants that can change between statements.

Theorem 1 (Background-depletion cutoffs)**.** *For the directed model above, there are $c, C, \varepsilon_0 > 0$ such that for all integers $a \ge 1$ and $0 < \varepsilon \le \varepsilon_0$,*

$$ca^{-1/2}e^{-C\varepsilon^2 a} \le \mathbb{P}(I_\varepsilon \ge a) \le Ca^{-1/2}e^{-c\varepsilon^2 a}. \tag{3}$$

The same bounds hold for $\mathcal{A}_\varepsilon$, with possibly different constants. For each fixed $q > 0$,

$$\lim_{\varepsilon \downarrow 0} \frac{\log \mathbb{E} S_\varepsilon^q}{\log(1/\varepsilon)} = \max\{0, 7q/3 - 1\}, \tag{4}$$

$$\lim_{\varepsilon \downarrow 0} \frac{\log \mathbb{E} D_\varepsilon^q}{\log(1/\varepsilon)} = \max\{0, 4q/3 - 1\}. \tag{5}$$

All positive moments are finite when $\varepsilon > 0$. For every fixed $\alpha, M > 0$, as $\varepsilon \downarrow 0$,

$$\mathbb{P}(D_\varepsilon \ge \varepsilon^{-4/3-\alpha}) + \mathbb{P}(S_\varepsilon \ge \varepsilon^{-7/3-\alpha}) = O_{\alpha,M}(\varepsilon^M). \tag{6}$$

Conversely, writing $h_\varepsilon = \log(1/\varepsilon)$, for all sufficiently small $\varepsilon > 0$,

$$\mathbb{P}\Big(D_\varepsilon \ge c\varepsilon^{-4/3}h_\varepsilon^{-1/3}, \quad S_\varepsilon \ge c\varepsilon^{-7/3}h_\varepsilon^{-13/12}\Big) \ge c\varepsilon. \tag{7}$$

The three cutoff exponents are therefore 2 for area, 4/3 for depth, and 7/3 for total topplings. The area estimate has a two-sided exponential form. For depth and size, (6)–(7) locate the cutoff powers with logarithmic uncertainty; they do not give a limiting profile or a constant-factor cutoff. The lower witness has probability of order at least ε, so these are scales of rare large avalanches, not the size of a typical unconditioned avalanche.

For example, the susceptibility, meaning the expected number of topplings, satisfies

$$c\varepsilon^{-4/3}h_\varepsilon^{-1/3} \le \mathbb{E}S_\varepsilon \le C\varepsilon^{-4/3}h_\varepsilon^{2/3}. \tag{8}$$

The area moments have the stronger constant-factor description

$$\mathbb{E}\mathcal{A}_\varepsilon^q \asymp \begin{cases} \varepsilon^{-(2q-1)}, & q > 1/2, \\ \log(1/\varepsilon), & q = 1/2. \end{cases} \tag{9}$$

Below $q = 1/2$ they converge to finite critical moments. The corresponding thresholds for size and depth are 3/7 and 3/4; at those endpoints the depleted moments diverge but have logarithmic exponent zero. None of these statements assumes a finite-size scaling function.

The proof starts with an exact response calculation at one fresh site. If $M(u)$ is its total emitted mass after input $u \ge 0$, and η' its final energy, renewal stationarity gives

$$M(u) = u + \eta - \eta', \qquad \eta' \sim f, \qquad \mathbb{E}M(u) = u, \qquad |M(u) - u| \le 1. \tag{10}$$

Initial and final energies need not be independent. Renewal mixing implies $\operatorname{Var} M(u) \to 2\operatorname{Var}_f(\eta) > 0$ as $u \to \infty$, and continuity and positivity give a uniform variance floor for $u \ge 1/4$. Every positive input in this model is at least 1/4.

For the depleted response $M_\varepsilon(u)$, set $b_\varepsilon(u) = u - \mathbb{E}M_\varepsilon(u)$. Uniformly for $u \ge 1/4$ and sufficiently small positive ε,

$$d_\varepsilon := \varepsilon - \frac{4}{3}\varepsilon^2 \le b_\varepsilon(u) \le \varepsilon, \qquad |M_\varepsilon(u) - u| \le 1, \qquad \operatorname{Var} M_\varepsilon(u) \ge c_0 > 0. \tag{11}$$

Without clipping at zero, depletion simply replaces input u by $u - \varepsilon$, giving mean loss exactly ε. Clipping matters only when $\eta < \varepsilon$, an event of probability $4\varepsilon^2/3$, and changes the emitted mass by at most one. The same coupling controls the variance. The effective loss is thus approximately ε per receiving site, regardless of the number of topplings there.

Explore receiving sites in layer order, processing each only once its total input has arrived. Let W_k be the total pending input after k sites have been processed, with $W_0 = \delta$. Its extinction time is $\tau = I_\varepsilon$. Given the exploration history $\mathcal{G}_k$, the next site's energy and stack are fresh. Before extinction,

$$|\Delta W_k| \le 1, \qquad \mathbb{E}(\Delta W_k \mid \mathcal{G}_k) = -b_k, \qquad d_\varepsilon \le b_k \le \varepsilon, \qquad c_0 \le \operatorname{Var}(\Delta W_k \mid \mathcal{G}_k) \le 1. \tag{12}$$

This is a nonnegative process with a uniformly controlled small negative drift. It need not be Markov when only W_k is observed. All arguments below use the full exploration history.

Here is the argument for the two-sided area cutoff. At criticality the same frontier is a martingale. Stopping on reaching zero or height H, bounded optional stopping gives a probability of order H^{-1} to hit H, and the square-martingale budget bounds the expected stopping time by CH. After hitting H, a maximal inequality gives a fixed positive probability of surviving another cH^2 steps. Taking H of order $\sqrt{a}$ therefore gives $\mathbb{P}(I \ge a) \asymp a^{-1/2}$. The coupling with the same backgrounds and stacks gives $I_\varepsilon \le I$.

For the depleted upper bound, take $\lambda = d_\varepsilon/2$. Taylor's inequality in (12) shows that $e^{\lambda W_k} - 1$ is a nonnegative supermartingale and hence

$$\mathbb{P}\Big(\sup_k W_k \geq H\Big) \leq \frac{e^{\lambda\delta} - 1}{e^{\lambda H} - 1}. \tag{13}$$

At step $\lfloor a/2 \rfloor$, either the surviving frontier exceeds a constant multiple of $d_\varepsilon a$, or the compensated martingale must make an upward deviation of that order to survive the remaining steps. The first alternative is controlled by (13); the second by a bounded-increment exponential inequality, multiplied by the critical survival bound $Ca^{-1/2}$. For $\varepsilon^2 a \geq 1$ this gives

$$\mathbb{P}(\tau \geq a) \leq C(a^{-1/2} + \varepsilon)e^{-c\varepsilon^2 a} \leq C' a^{-1/2} e^{-c'\varepsilon^2 a}.$$

For $\varepsilon^2 a < 1$, critical domination suffices.

For the lower bound, start after any history with frontier mass $0 < w < H$ and stop on reaching zero or $[H, \infty)$. If $\varepsilon(H+1) \leq c_0/4$, the expected increment of W^2 before this stop is at least $c_0/2$. Since the stopped mass is at most $H+1$, the expected stopping time is at most $2(H+1)w/c_0$. The drift budget then leaves expected terminal mass at least $w/2$, yielding a hitting probability at least $w/[2(H+1)]$. Once mass H is reached, a maximal inequality gives a fixed positive probability of surviving $\lfloor H^2/64 \rfloor$ steps and retaining at least $H/2$.

Choose $H = h_0/\varepsilon$ with $h_0 > 0$ sufficiently small. Reaching H from δ costs probability of order ε. Each subsequent block of order ε^{-2} steps, including a return from $H/2$ to H if necessary, has a uniformly positive conditional success probability. Return periods only increase the lifetime. Repeating these blocks gives $\mathbb{P}(\tau \geq a) \geq c\varepsilon e^{-C\varepsilon^2 a}$ above the cutoff scale, which implies the required lower bound there. Below that scale, the same hitting argument with H of order $\sqrt{a}$ gives $ca^{-1/2}$. This proves (3); tail integration gives (9). The proof uses a lattice exploration with merging, without identifying its genealogy with an independent branching process.

Depth and total size require additional spatial information. Let $a_n(j)$ be the input to site (n, j), $Z_n = \sum_j a_n(j)$ the layer flux, $J_n = \#\{j : a_n(j) > 0\}$, and $U_n = \max_j a_n(j)$. Let P average the two neighboring transverse coordinates: $(Pv)(j) = [v(j-1) + v(j+1)]/2$. At criticality

$$a_{n+1} = Pa_n + P\xi_n, \qquad \xi_n(j) = M(a_n(j)) - a_n(j). \tag{14}$$

Conditional on earlier layers, the noises in one layer are independent, centered, and bounded by one. Equation (14) is an exact lattice identity. Under depletion it acquires a nonpositive drift term, and its centered noises remain uniformly bounded.

For the simple symmetric walk kernel p_r, $\sum_j p_r(j)^2 \leq C(r+1)^{-1/2}$. Conditional Hoeffding estimates in the convolution of (14) therefore give, for every fixed $K > 0$,

$$\mathbb{P}\Big(\max_{n \leq N} U_n > B_N\Big) \leq C_K N^{-K}, \qquad B_N = C_K (N+1)^{1/4}\sqrt{\ell(N)}. \tag{15}$$

The same cap holds under depletion. An integrated version also bounds the expected flux on cap-failure events by an arbitrarily high inverse power.

To see how the cap controls survival, kill the critical flux permanently when it first fails and call the killed process X_n. For $0 \leq \theta \leq 1$, the variance floor and $J_n \geq Z_n/B_N$ imply

$$\mathbb{E}(e^{-\theta X_{n+1}} \mid \mathcal{F}_n) \geq \exp\Big[-\Big(\theta - c\theta^2/B_N\Big) X_n\Big]. \tag{16}$$

Here $\mathcal{F}_n$ contains the preceding layers. Iterating $\theta \mapsto \theta - c\theta^2/B_N$ gives $\mathbb{P}(X_N > 0) \leq CB_N/N$. Killing is essential: conditioning on future cap survival would destroy the centering of the fresh

noise. Combining this upper bound with the spatial and second-moment estimates in the working record yields

$$cn^{-3/4}\ell(n)^{-1/4} \le \mathbb{P}(D \ge n) \le Cn^{-3/4}\ell(n)^{1/2}, \tag{17}$$

$$cs^{-3/7}\ell(s)^{-13/28} \le \mathbb{P}(S \ge s) \le Cs^{-3/7}\ell(s)^{2/7}. \tag{18}$$

The logarithmic factors here are bounds, not identified correction exponents. In particular avalanches are finite almost surely, but $\mathbb{E}S = \infty$.

Let S_L count topplings in the first L layers of this infinite root avalanche. Conservation gives the exact mean $\mathbb{E}S_L = L/3$. The working proof also gives, for every fixed $q > 0$,

$$\lim_{L\to\infty} \frac{\log \mathbb{E}S_L^q}{\log L} = \max\{0, (7q-3)/4\}. \tag{19}$$

These critical estimates are inputs to the depletion analysis. The value $10/7 = 1 + 3/7$ is the usual size-density exponent suggested by (18). It is already familiar from other stochastic directed sandpiles [3, 4]; (18) itself does not assert $\mathbb{P}(S = s) \sim cs^{-10/7}$.

Write $y_n = \mathbb{E}Z_n^\varepsilon$, and choose caps B_n so that the exceptional expected mass $e_n = \mathbb{E}[Z_n^\varepsilon; U_n^\varepsilon > B_n]$ satisfies $e_n \le C_K(n+1)^{-K}$. The loss per receiving site in (11) gives the scalar recursion

$$y_{n+1} \le (1 - d_\varepsilon/B_n)y_n + (d_\varepsilon/B_n)e_n. \tag{20}$$

Its product contracts on the scale determined by $\sum_{k<n} d_\varepsilon/B_k$, which is of order $\varepsilon n^{3/4}/\sqrt{\ell(n)}$. Splitting the error sum at $n/2$ makes this quantitative:

$$\mathbb{E}Z_n^\varepsilon \le C_K\left[\exp\left\{-\frac{c_K\varepsilon n^{3/4}}{\sqrt{\ell(n)}}\right\} + (n+1)^{-K}\right]. \tag{21}$$

Summing gives the upper susceptibility bound in (8). Because a nonzero flux is at least $1/4$, this also gives the far-depth bound in (6).

The matching lower witness comes from a late block of layers. The loss budget keeps expected flux bounded below through $2L$, for L of order $\varepsilon^{-4/3}h_\varepsilon^{-1/3}$ with a small constant. The mass V_L emitted in the second block of L layers satisfies

$$\mathbb{E}V_L \ge cL, \qquad \mathbb{E}V_L^2 \le CL^{11/4}\ell(L)^{1/4}, \qquad \mathbb{P}(V_L > 0) \le CL^{-3/4}\ell(L)^{1/2}.$$

Removing the contribution below $cL^{7/4}\ell(L)^{-1/2}$ leaves a fixed proportion of the mean. Cauchy–Schwarz then gives a probability at least $cL^{-3/4}\ell(L)^{-1/4}$ of exceeding that mass. Substitution of L proves (7).

For the moment upper bounds, truncate at $R_\varepsilon = \lceil \varepsilon^{-4/3-\alpha} \rceil$ and use critical domination and (19). Beyond that depth, the deterministic bound $S_L \le CL^3$ and the arbitrary-power depth tail control every positive moment. The deterministic bound follows by weighting energy in layer $n < L$ by $L - n$: each toppling decreases weighted energy by at least $1/2$. Letting $\alpha \downarrow 0$ gives the upper exponents in (4)–(5); the joint lower witness gives the reverse inequalities above the moment thresholds. The critical avalanche uses only finitely many continuous threshold tests, so under the monotone coupling its counts are eventually unchanged as $\varepsilon \downarrow 0$. Monotone convergence settles the subthreshold moments and the endpoints. Finally high moments and Markov's inequality give the far-size assertion of (6).

There is also joint information about the avalanches that reach a large depth. Define $\mathbb{Q}_N = \mathbb{P}(\cdot \mid D \ge N)$, and let $R_\perp$ be the largest absolute transverse coordinate of any receiving site in the entire avalanche. The notation $Y_N = N^{\gamma + o_{\mathbb{Q}_N}(1)}$ means that for every $\alpha > 0$, $\mathbb{Q}_N(N^{\gamma-\alpha} \le Y_N \le N^{\gamma+\alpha}) \to 1$.

Theorem 2 (Joint geometry under survival)**.** *As $N \to \infty$, jointly under $\mathbb{Q}_N$,*

$$S = N^{7/4+o_{\mathbb{Q}_N}(1)}, \quad \mathcal{A}, I = N^{3/2+o_{\mathbb{Q}_N}(1)}, \quad R_\perp = N^{1/2+o_{\mathbb{Q}_N}(1)}, \quad D = N^{1+o_{\mathbb{Q}_N}(1)}. \tag{22}$$

The size and area statements also hold for their first-N-layer versions. For every fixed $t \in (0,1)$,

$$Z_{\lfloor tN \rfloor} = N^{3/4+o_{\mathbb{Q}_N}(1)}, \quad J_{\lfloor tN \rfloor} = N^{1/2+o_{\mathbb{Q}_N}(1)}, \quad U_{\lfloor tN \rfloor} = N^{1/4+o_{\mathbb{Q}_N}(1)}.$$

For each fixed exponent tolerance, the fraction of layers violating these bounds tends to zero in conditional probability in every fixed interval $[aN, bN]$, $0 < a < b < 1$.

This statement is stronger than matching marginal tail exponents. Restart (16) at an intermediate layer k, before conditioning on survival to N. For every event $E \in \mathcal{F}_k$, it gives

$$\mathbb{P}(X_N > 0, E) \le \frac{CB_N}{N-k}\mathbb{E}[X_k; E].$$

Taking $E = \{Z_k \le u\}$, and using (17) before dividing by $\mathbb{P}(D \ge N)$, proves that a surviving avalanche rarely has intermediate flux much below $N^{3/4}$. The cap (15) and heat-kernel confinement then force the receiving count and maximum-input exponents. A bounded-increment maximal inequality for W forces the area exponent: reaching flux $N^{3/4}$ in much fewer than $N^{3/2}$ exploration steps has negligible probability even after conditioning. Summing flux over most intermediate layers forces the size exponent. Critical tails supply the matching conditional upper bounds.

The same geometric statements hold uniformly under $\mathbb{P}(\,\cdot \mid D_\varepsilon \ge N)$ whenever $N \to \infty$ and

$$\varepsilon N^{3/4} \ell(N)^{1/4} \le c_1 \tag{23}$$

for a sufficiently small fixed $c_1 > 0$. This window does not cover every sequence near the depletion cutoff. The conclusions give logarithmic exponents; tightness under normalization by the exact powers, a positive limiting occupied fraction, and limiting avalanche shapes remain open.

A one-way chain gives a useful exact comparison. Send the entire release to the next site, retaining the same background, root addition, and depletion. The working record proves

$$\lim_{\varepsilon \downarrow 0} \varepsilon^2 \mathbb{E} S_\varepsilon = \frac{\delta \operatorname{Var}_f(\eta)}{m} = \frac{37}{1944}, \qquad \operatorname{Var}_f(\eta) = \frac{37}{648}. \tag{24}$$

Here the frontier consists of a single pending input. The identity follows from its square-martingale budget, renewal mixing at large inputs, and a uniform bound on time spent at bounded frontier masses [1, Section 10.3]. The prefactor is a model-specific calculation; its numerical value is not a universality claim.

The area-cutoff, moment-exponent, and conditional-geometry statements extend to every fixed release law with a continuous strictly positive density on $[a, 1]$, $0 < a < 1$, using its renewal background and equal forward routing. Constants depend on the law. The one-site variance floor, bounded response noise, and positive minimum input are the decisive hypotheses. Arbitrary atomic laws, proportional releases, resetting, and physical-time exponents are outside this statement.

The literature comparison places the proposed contribution precisely. Kim and Wang's SSP already uses random release amounts with fixed neighbor allocation and develops an abelian stationary-state theory in a discrete setting [2]. Thus random releases and abelianness are precedents, not claims of priority here. Kloster, Maslov, and Tang obtain the exponent values 7/4 and 10/7 for a different rule with random routing [3]. Paczuski and Bassler derive a discrete avalanche equation and use a continuum analysis to obtain critical exponents and geometric scaling [4]. The heat-equation viewpoint and those exponent values also have clear precedents.

Pastor-Satorras and Vespignani study directed models with bulk loss at topplings, including cutoff scaling and the balance $\varepsilon \mathbb{E}S = 1$ in their normalization [5]. Here ε instead measures the clipped reduction of the initial background in (2); the mean-size exponent is 4/3. The candidate new content is the quantitative theorem package for this perturbation: the two-sided area bound, the depth and size cutoffs with explicit probability quantifiers, the full moment spectra, and the conditional geometry with a controlled depleted extension. These are derived within the random-release model, rather than transferred from a proposed equivalence of universality classes.

Several questions remain substantive. The cutoff functions and exact logarithmic corrections are unidentified. Slow driving is imposed, so no claim is made that a finite-rate drive spontaneously creates separated timescales. On undirected lattices, returning mass destroys the fresh-site exploration used here. Finite-volume stationary sampling and infinite-mean bulk avalanches are available in the working record, but undirected spatial extinction and uniqueness of the bulk law remain open. Even in dimensions at least three, the proved local fixation does not imply that an avalanche visits only finitely many sites.

The full source of Theorem 1 is Sections 9.3–9.5 and 9.7–9.10 of [1]; Theorem 2 is Section 9.11. The present article supplies the model, precise statements, and proof mechanisms, with the longer second-moment and renewal arguments retained there. No new simulations are used. The mathematical claims remain working proofs, and their independence from the literature still warrants specialist review before any claim of established novelty.

References

[1] *A conservative random-release sandpile: problem, proofs, obstructions, and continuation instructions.* AI-written research record prepared for Jacob Richey, 8 October 2026. Unpublished working record; not included in this reader.

[2] S. Kim and Y. Wang, *A stochastic variant of the abelian sandpile model* (2019). arXiv:1906.01107. doi:10.1007/s10955-019-02453-7.

[3] M. Kloster, S. Maslov, and C. Tang, *Exact solution of stochastic directed sandpile model*, Phys. Rev. E **63**, 026111 (2001). arXiv:cond-mat/0005528.

[4] M. Paczuski and K. E. Bassler, *Theoretical results for sandpile models of self-organized criticality with multiple topplings*, Phys. Rev. E **62**, 5347 (2000). arXiv:cond-mat/0005340.

[5] R. Pastor-Satorras and A. Vespignani, *Critical behavior and conservation in directed sandpiles*, Phys. Rev. E **62**, 6195–6205 (2000). arXiv:cond-mat/0007131.

[6] J.-B. Bardet, A. Christen, and J. Fontbona, *Quantitative exponential bounds for the renewal theorem with spread-out distributions.* arXiv:1504.06184, version 3 (2016).