capstract:2610.00006v1 [physics.soc-ph] 27 Oct 2026

Collective switching and return-point memory in directory capitalization

Mira L. Halvessen,[1, *] Nadim E. Vassour,[2] and Mina K. Teshari[3, 1]

[1]*Laboratory for Driven Systems, Norbjerg Institute, Aarhus, Denmark*
[2]*Department of Statistical Mechanics, Centre Valcienne, Lausanne, Switzerland*
[3]*School of Information Practice, Sefidar Institute, Shiraz, Iran*

(Dated: October 27, 2026)

We investigate the response of ordered directories to cyclic changes in an audit score favoring initial capitals. Two different case configurations persist at the same score. In an experiment with 48 document editors, calibration against fixed neighboring titles gives a nearest-neighbor coupling of 0.180 audit points and a coupling-to-disorder ratio of 1.16. A driven random-field model fitted to the calibration experiment predicts the principal loop and the response to nested reversals in separate editors. Visible flanking titles increase the mean area of the principal loop from 0.129 to 0.324 audit points. When a previous maximum is exceeded, coupling predicts the number of titles that will fail to recover more closely than independent switching does. Whole-directory recovery at a return point is less reliable than recovery measured by the fraction of titles in capitals.

I. INTRODUCTION

A change in a filing standard is commonly issued as a list of examples. The new rule can be checked against those examples. The previous standard may still be evident in directories to which both rules have applied. This occurs even when the file names are visible to a single editor and no approval is required to change them. We isolate the effect for capitalization of the first letter of a folder title. At an intermediate audit score, the same editor presented with the same titles can leave most of the capitals in place or remove most of them, according to the preceding audit.

Threshold models describe how a fixed distribution of preferences can give a discontinuous response to social information [1, 2]. The number of choices already made can serve as a signal to the next chooser, an effect isolated experimentally in other settings [3, 4]. In our setting the signal is a state of the directory. There is one chooser at a time. A change to a title at one location alters the evidence at its two neighbors even when the audit score is held fixed. The change can propagate down a list, stopping at a title whose preferred case is sufficiently strong.

The driven random-field Ising model has two properties relevant to this problem [5, 6]. It distinguishes the effect of local heterogeneity from a uniform field. It also predicts more specific properties of the history than the presence of a loop. After a field reversal and a subsequent return to the reversal value, the configuration can be recovered site by site. Once the field moves beyond that value the return-point condition no longer applies. In a directory the predicted stored variable is the location of each capital letter. The count of capitalized titles alone provides a weaker test.

II. LOCAL FIELDS AND HISTORY

The ith folder in a list of fixed order has state $s_i = +1$ if its initial letter is uppercase and $s_i = -1$ otherwise. The initial alphabetic character is the only part of the displayed title that changes. The ordering key is independent of this character. The model assigns each title a fixed bias h_i in audit points. Larger h_i means the audit must favor uppercase more strongly to convert the title. Each visible, immediately adjacent title adds a local contribution of the sign of its state. With symmetric graph A and coupling $J \geq 0$, the effective field is

$$f_i(H, \mathbf{s}) = H - h_i + J \sum_j A_{ij} s_j. \tag{1}$$

Here H is the score difference in favor of uppercase. The coupled experiment uses an open chain on 24 titles. A missing neighbor contributes zero at either end. In the masked experiment the graph has no edges. The distinction is operational. An editor in the masked trial sees one title at a time, including its current case, and can still recall something about a title on another screen.

The local form of the score has the quasistatic energy representation

$$\mathcal{E}(\mathbf{s}; H) = -\sum_i (H - h_i) s_i - J \sum_{i<j} A_{ij} s_i s_j. \tag{2}$$

Energy here is stated in audit-point units. We use it to keep track of the sign of a change, and assign an additional empirical threshold $c \geq 0$. Either flip decreases Eq. (2) by at least $2c$ if accepted. At fixed H an upward change is possible if $f_i > c$ and a downward change if $f_i < -c$. Between these bounds a folder remains in the case in which it was read. This is the usual zero-temperature single-site dynamics when $c = 0$. A

∗ m.halvessen@norbjerg.example

positive c gives an isolated folder a local hysteresis width as in a rectangular hysteron [7, 8]. Paired sweeps with fixed neighboring titles identify c.

A stable list at one score can become unstable after a score increment. During upward relaxation all accepted flips are upwards. Every such flip raises the field of its neighbors and therefore cannot invalidate an eligible upward flip. The terminal state is the same under all update orders that visit every unstable site until stability. The analogous statement holds in the opposite direction. This is the finite-chain version of the order argument used in driven disordered media [6, 9]. The proof of recovery for the histories used here is given in Appendix A.

The random field was centered by the calibration fit, with its standard deviation denoted by σ_h. We define the Halvessen number

$$\mathrm{Ha} = \frac{2J}{\sigma_h}. \tag{3}$$

This is the shift in a switching threshold caused by reversing a single neighbor, measured in units of the spread of bias across titles for the same editor. At an interior site, changing both neighboring cases shifts a threshold by $4J$. On a d-neighbor graph the available range would be $2dJ$. The definition in Eq. (3) remains a one-neighbor contrast. In particular, it can be estimated even when only one flank is controlled experimentally.

For an order-of-magnitude estimate, suppose the thresholds of the remaining susceptible titles have the unconditioned Gaussian density near its maximum. A neighboring flip moves a threshold past the current score if it lies in an interval of length $2J$. The largest unconditional mass of such an interval is $2\Phi(J/\sigma_h) - 1$. With J near 0.18 and σ_h near 0.31, this is about 0.44. A displacement of one digit after the decimal point in the score is smaller than the interaction shift, so a group of flips at one score can be separated experimentally from the change in the audit. The 0.44 estimate ignores conditioning on a title still being lowercase. It gives the size of the interval effect and is not a branching ratio for the observed chain. After an interior seed fires, a propagating front on either side has at most one fresh neighbor to test.

III. EXPERIMENT

A. Instrument and participants

A local editor presented a directory of empty experimental folders. Titles were generated from a screened list of two-word noun phrases used for work files. The controlled sets used generic headings such as “supplier copies” and “annual returns.” We excluded acronyms and titles that were proper names. The list appeared in a fixed code-based order unrelated to words or case. We also disabled auto-capitalization and intercepted system shortcuts that could change several rows together. Holding the title text fixed across a reversal avoids the changes in sort position that confound older audit logs [10, 11].

The screen gave an audit score per title, with displayed difference H in points between the uppercase and lowercase options. The score was the difference of two stated audit credits. Participant compensation was fixed. Editors were told to use the audit as the current preference of the filing office and to mark each title in the case they would leave at that score. They could leave a title as it stood. To begin a new run the titles were set to lowercase and the score to -1.50. An editor could see every review opportunity, including ones at which the score did not cause a change. No accepted edit was planted by the software during a run.

The instrument was used with 48 adult document editors. Participants gave written consent for anonymous edit records, and the Norbjerg Institute review board approved the protocol. The calibration cohort comprised 24 editors. The other 24 provided full-directory outcomes. We requested experience with a shared directory at work but did not select for allegiance to a house style. The practice trial used different stems and presented the same toggles at fixed score. Following practice, adjacent titles remained visible only in the visible-list task. In every task the active row had the same font, spacing and position.

B. Calibration of local thresholds

Calibration used 12 titles per editor. Each title was reviewed against two pinned flanks in each of three states of the neighbor sum m: -2, 0 and 2. The two mixed-flank orientations were allocated evenly across titles and mirrored across editors. The three conditions were alternated by a balanced order schedule. At each m, a forward sweep from lowercase and a backward sweep from uppercase bracketed the two switch scores. Each sweep changed score in steps of 0.05 points over the common range. The resulting 1,728 calibration intervals preserve the last score at which the old state was retained and the first score at which it was changed. The calibration used 288 distinct editor-title blocks. Pinned neighbors could be seen but could not themselves be renamed.

The threshold expressions obtained from Eq. (1) are

$$\tau_i^{\uparrow}(m) = h_i + c - Jm, \qquad \tau_i^{\downarrow}(m) = h_i - c - Jm \tag{4}$$

We fit a Gaussian random-intercept model with editor and title-within-editor effects, integrating the likelihood over each interval of width 0.05 [12]. The residual scale describes within-title threshold variability conditional on these effects. A participant-specific shift moves both thresholds together. The estimand σ_h is the between-title spread within an editor, with the editor intercept removed. The paired pin contrast estimates J independently of that intercept. A practice score with the row shown alone supplies the origin of the scale. For the validation lists, the means of the two isolated-row threshold estimates were used as the h_i inputs to the predictions. We retained their midpoint grid precision. These measurements preceded all trials with visible neighbors.

Fig. 1 shows the isolated dependence on m. The means on the upper and lower branches move together to the available precision. The fit and the order estimate from Sec. II are in agreement about the relative sizes of the interaction shift and the distribution width.

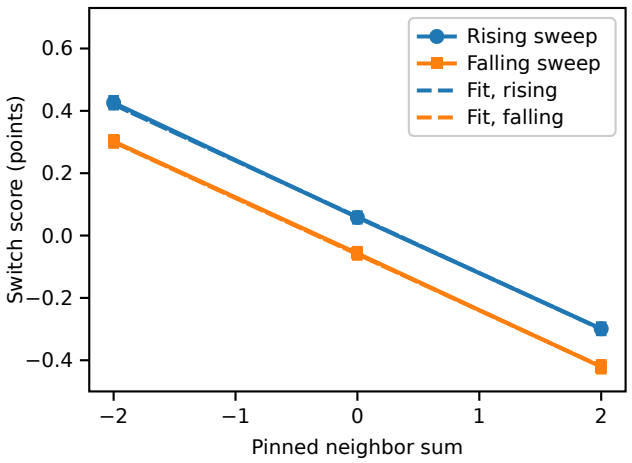


FIG. 1. Mean switching scores with pinned flanks. Bars are editor-block standard errors. Fitted lines have common slope equal to the negative of the neighbor coupling.

C. Validation histories

Each validation editor reviewed the same 24-title set in both visibility conditions, in counterbalanced order. Sets remained fixed within an editor and were taken from screened directories unrelated to the calibration titles. Each condition began at lower saturation and ran the score up to $+1.50$, then down to -1.50, by increments of 0.05. There were 121 score holds per principal loop. On each hold the instrument prompted a sweep through the directory and offered subsequent sweeps until an entire pass produced no edits. In the visible list the active row was highlighted and the cases of its immediate neighbors were visible. The masked interface showed only the active row. The sweep protocol fixes the reading order and makes changes with the score held constant observable. It also makes that order an interpretive constraint on reported avalanche locations.

A third run, at full visibility and following a fresh lowercase initialization, used the nested history in Table III. The field moved upward from -1.50 to $+0.35$, down to -0.40, up to $+0.15$, and back down to -0.40. It then passed through $+0.35$ on its way to $+0.65$, before returning to $+0.35$ from above. One snapshot came from the quiescent sweep at each comparison score. The interface displayed neither past scores nor previous file names. Each history had 102 score holds. We included the last return from above to distinguish a true return to an unexceeded reversal from approaching the same score by a new path.

For uncertainty we resampled editors as blocks, keeping paired runs together [13]. The protocol produces a large number of title observations, but retaining each as an independent editor would shrink confidence intervals for between-editor effects. Pinned trials and validation runs were resampled separately. No parameters in the prediction were fitted to the principal or nested loops. We used the same c for both visibility conditions and set $J = 0$ for masking. The independent-hysteron comparison uses the same h_i and c, and always sets $J = 0$. It is an informative comparison because it retains the possibility of return-point memory. The effects estimated by a model with independent equilibrium thresholds would confound the isolated barrier with the title-to-title coupling.

IV. CALIBRATION AND PRINCIPAL LOOPS

The threshold fit yields the values in Table I. The estimated switch shift from swapping two lowercase flanks for uppercase ones is $4J = 0.720$ audit points. That contrast is larger than twice the within-editor standard deviation among titles. The estimated mean offset of individually assessed titles was below 0.01 points in absolute value. We used the centered calibration in Fig. 1 and preserved the validation intercepts in the predictions.

TABLE I. Calibration fit from 24 editors. Intervals are the 2.5th and 97.5th percentiles of 2,000 editor-block bootstrap replicates. The global field and thresholds have the same unit.

Parameter	Meaning	Estimate	95% interval	Units
J	Visible neighbor coupling	0.180	0.164 to 0.196	points
c	Isolated half-width	0.060	0.051 to 0.069	points
σ_h	Bias spread within editor	0.310	0.282 to 0.339	points
σ_e	Editor offset spread	0.051	0.032 to 0.070	points
σ_η	Residual threshold spread	0.047	0.041 to 0.054	points
Ha	One-neighbor shift divided by bias spread	1.16	1.02 to 1.34	dimensionless

The fraction of uppercase titles is $p = (1 + L^{-1} \sum_i s_i)/2$ for $L = 24$. The principal loops in Fig. 2 show the mean of that fraction across 24 independent editors. At zero field, a mean fraction of 0.2361 was uppercase after upward driving and 0.8212 after downward driving with the flanks visible. These fractions correspond to 136 and 473 of 576 titles, respectively. In the masked runs they were 250 and 358 of 576. Differences of that size remain visible even if an editor converts the same number of titles as another editor at a different place on the list.

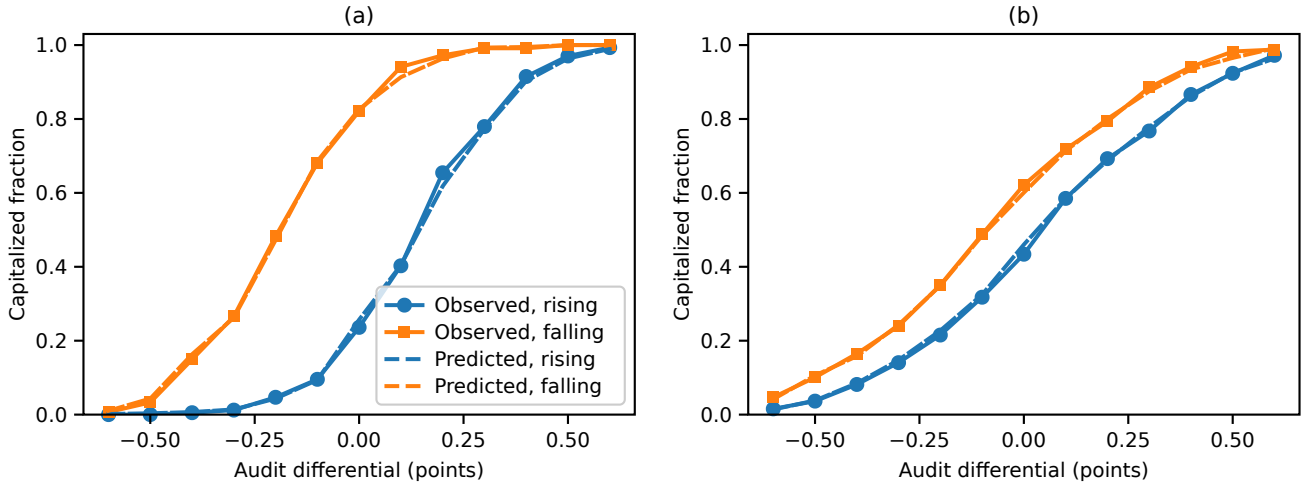


FIG. 2. Principal loops in the held-out editors. Panel (a) shows visible flanks and panel (b) shows masking. Color distinguishes the upward and downward branches. Solid curves join observed mean capitalization fractions and dashed curves join the forward predictions. Axes show the central part of the field range.

We integrate the gap between the branches over the full score range,

$$\mathcal{A} = \int_{-1.50}^{1.50} [p_\downarrow(H) - p_\uparrow(H)]\, dH. \tag{5}$$

The paired difference in area is 0.195 points, with a 95% editor-block bootstrap interval from 0.172 to 0.218 points. The prediction is 0.207 points. The root mean square error in Table II is evaluated by squaring residuals at every held score for each editor before averaging them. Agreement of the two population-average curves can conceal offsetting errors in individual trajectories.

TABLE II. Principal loop area and trajectory error in validation. Brackets give 95% editor-block bootstrap intervals for the observed means. The models share the same graph in the masked condition.

Flanks	Observed area (points)	Predicted area (points)	Coupled RMSE	Independent RMSE
Visible	0.324 [0.300, 0.349]	0.327	0.049	0.103
Masked	0.129 [0.123, 0.136]	0.121	0.025	0.025

The difference in loop area has contributions from the early and late sections of each sweep. With lowercase flanks an upward flip of an interior title needs a score higher by **0.36** than when it is displayed alone. When both flanks are already capitalized the difference is **−0.36**. Thus the model can delay the first change and then hasten changes farther down the list at a common score. At the descending edge it yields the converse effect. A screen-wide imitation effect would have a different prediction for the first and last titles, where a change in one neighboring row supplies the entire modeled shift. In the residuals, the endpoint titles account for more of the early rising flips than their fraction of list length. The open chain includes this effect before fitting any model to those flips.

A. Rearrangements at fixed score

We counted the number S of distinct titles changed during each nonempty score hold. Both conditions give 1,152 total transitions in the principal loop, since the observed initialization and both saturation endpoints were complete and changes on each branch had the same direction. Their allocation to holds differs markedly. There were 382 nonempty holds in the visible condition and 724 in the masked condition. Their mean sizes were 3.02 and 1.59. In Fig. 3 we normalize by the respective numbers of nonempty holds. Zero-event holds were omitted from the denominator.

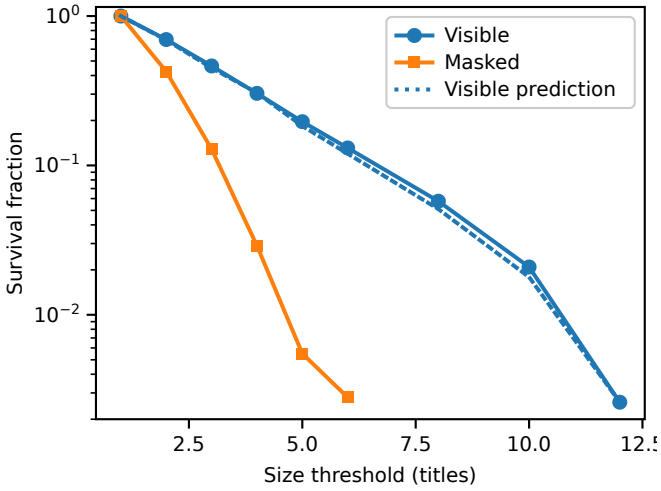


FIG. 3. Fraction of nonempty holds with at least the indicated number of changes. The dotted curve is the result for visible flanks from the fitted chain. The masked curve ends at six, the largest observed burst under masking.

Eight visible-list holds changed ten or more titles. The change of one title could alter the recommendation applicable to a later row in the same hold. In contrast, two isolated titles with thresholds inside the same 0.05 interval would register as a burst under masking. Our score grid therefore measures the entire response to an increment, including simultaneous seeding events. Avalanche definitions in the zero-step model distinguish those events from successive flips caused by an initial seed [14, 15]. We used the experimental step size in the prediction. Applying a continuous-field code directly to the comparison would systematically downweight multisite events in the masked condition.

The graph and the distribution of h_i should both affect the upper tail [16]. In 24 rows one data point supplies the last survival count. A fitted straight line in a log-log plot would spend most of its range on fewer than eight holds. Models of crackling noise predict other quantities such as the avalanche shape and the spectrum of resolved event times [17, 18]. Our timestamp is partly a measurement of the time the editor takes to reach the next row. We therefore report survival fractions at selected sizes without estimating a critical exponent [19].

V. NESTED REVERSAL AND ERASURE

Let $\mathbf{s}^{(a)}$ and $\mathbf{s}^{(b)}$ be the vectors observed at two visits to the same score. We use the positional mismatch

$$d(a,b) = \frac{1}{2}\sum_{i=1}^{24} |s_i^{(a)} - s_i^{(b)}|. \tag{6}$$

Exact recovery is $d = 0$. This test disallows any redistribution of capitals among titles. Table III records whole-directory recovery and pooled flips by comparison, and Fig. 4 shows the mean fraction over time. Between the two measured visits to -0.40, the observed count rose by 199 capitalized positions in aggregate at $+0.15$, from 86 to 285, before falling to 97. The fit recovers the 94 incoming capitals at the second visit. Eleven lists retain a single extra capital, which accounts for all failures of exact recovery in the inner cycle.

TABLE III. Positional memory in the visible validation lists. Each row compares 24 pairs of snapshots, with 576 title positions per comparison. Capitalized counts refer to observed snapshots.

Score and path	Capitals, first and second	Exact directories	Mismatched positions	Coupled prediction	Independent prediction
−0.40, inner return	86 and 97	13/24	11/576	0/576	0/576
+0.35, outer return	488 and 477	13/24	11/576	0/576	0/576
+0.35, after +0.65	488 and 571	7/24	83/576	80/576	49/576

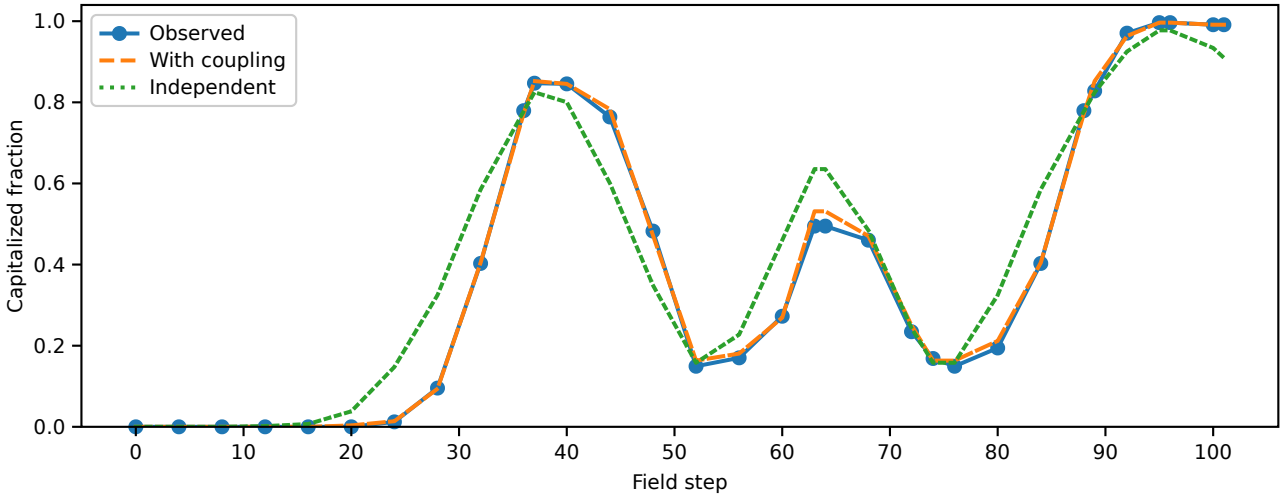


FIG. 4. Response to the nested field history. The displayed score-hold indices start at zero, with increasing indices corresponding to later observations. The comparison holds are 37 and 89 for the outer return, 52 and 74 for the inner return, and 37 and 101 for the return after overshoot.

The after-overshoot comparison retains 571 capitals, out of 574 at the immediately preceding maximum. At the first visit to $+0.35$ there were 488. The state acquired at $+0.65$ still supplies the neighbors of each folder as the score passes down through $+0.35$. In the model, a folder that was hard to convert while both flanks were lowercase can stay uppercase at this score if both have since changed. That mechanism does not enter the independent prediction, which retains 524 capitals at the last visit. This is a difference in the amount

of lost positional memory. Both models have exact return-point memory before the previous outer extremum is exceeded.

Our primary closure prediction is stricter than the observed one. In each of the two closed comparisons there were 11 one-title discrepancies. The pinned trials estimate the barrier for a single assessed title. Biases can also change over the course of reviewing a list. A change of bias by a few hundredths of a point between traversals is enough to prevent a marginal folder from returning. In the inner comparison all the discrepancies are retained capitals. In the outer comparison all are missing capitals. This distribution in the signs is consistent with incomplete reversal near the level of a return point. A symmetric independent bit-flip model would yield some discrepancies of the opposite sign. The current experiment supplies just one loop nesting per editor. It does not estimate a reliable distribution of history-dependent shifts in the local field.

There is also a direct error budget for the coupled forecast. At the calibration estimate, rounding score holds to the nearest interval midpoint gives up to 0.025 points of threshold uncertainty. The residual standard deviation of a calibrated threshold is 0.047 points. Evaluated at one sensitive interior row, one standard error in J shifts its two-flank threshold by approximately 0.016 points. These terms are smaller than the one-neighbor shift. They can still matter for the start of a burst. A substitution of the lower and upper calibration bounds for J moves the predicted gap on the principal loop in opposite directions, while the discrete exact-recovery prediction on the closed paths survives each substitution. An interaction uncertainty alone cannot explain those failures, because the closure identity holds for any fixed nonnegative J in Eq. (1).

VI. DISCUSSION

The pin experiment measures a change in the local threshold at a fixed title. That is the evidence for a coupling. The nested history probes a different property. The list itself carries state sufficient to recover the exact earlier assignment under the deterministic update rule. The independent model carries state at each site too. Its failure on the overshoot comparison follows from its weaker ability to retain capitals that were introduced above the original return level.

An experiment with identity and adjacency varied separately could distinguish wider context effects from the nearest-neighbor interaction estimated here. A future interface might display a stationary reference title at a distance of two rows while masking the intermediate one. The single-edge threshold shift in Eq. (3) would then give a measured unit with which to compare every added interaction. We do not know the appropriate graph for two different panes that show the same directory. The open chain was chosen because both intervening visibility and saturation can be controlled. A pane sorted by case-sensitive title would have a graph that changes as sites switch, and the no-passing inequality applies to a fixed graph. A bulk rename operation also makes previously measured local thresholds irrelevant for that operation.

The timescale matters. Our field changes are separated by at least one no-edit pass. In a workplace the order can be revised while the audit itself is being revised. The run then includes edits that precede local relaxation, and stable return points may never be reached. Earlier instrument runs with an automatic field advance showed delays concentrated at the end of the reading order [20]. Related field studies of capitalization distinguish differences across file types, especially names that must remain stable under software lookup [21]. We held identifiers fixed. The predicted order of revising a directory with fixed identifiers can be evaluated by comparison to a log, but the model gives no estimate for an entire migrating filesystem.

One remaining test needs no further estimation of the bulk loop. After an inner reversal, insert a neutral row between an already converted title and its first susceptible neighbor. Keep the score fixed. Equation Eq. (1) predicts exactly which change in local field the new boundary removes. The case of the far side of the directory should then depend on whether its first conversion preceded the insertion. This comparison would constrain how far a single boundary moves during an audit hold, even if the same total number of capitals is eventually reached.

DATA AVAILABILITY

The aggregate switch counts and calibrated coefficients are given in the figures and tables. The individual release file contains data for titles identified by run-local numbers. Requests for that file are handled by the Norbjerg instrument archive under the terms of the participant consent.

APPENDIX A: ORDER AND RETURN ON THE OPEN CHAIN

Use the partial order $\mathbf{s} \le \mathbf{t}$ when every component in the first vector is at most the corresponding component in the second. Let U_a denote relaxation through upward flips only at score a. If $\mathbf{s} \le \mathbf{t}$ and $a \le b$, then a site that becomes eligible along the lower path can be flipped in the upper path unless it is already uppercase there. To prove this, take two current states in order, and consider a site at which both are lowercase and the lower

field exceeds c. The upper field is at least as large by Eq. (1). At full relaxation it is impossible to have an extra upward flip on the initially lower path. The argument for D_a, relaxation through downward flips, reverses the roles of upper and lower. Finite monotone iteration gives unique terminal states of both maps at fixed input.

The order argument gives $U_a U_b = U_a$ for $b \le a$. The two-step result is bounded below by U_a of the initial state, by monotonicity in the input of U_a. It is bounded above by U_a of its own fixed point, because U_b can add no sites beyond that fixed point. Repetition yields the same property along any monotone sequence up to a. The analogous identity is $D_b D_a = D_b$ for $b \le a$. These identities account for the effect of intermediate score holds in the return-point calculation.

In the experimental path, write $S_A = U_a S_-$ and $S_B = D_b S_A$, with $a = +0.35$, $b = -0.40$ and S_- the lowercase configuration. For the inner rise to $v = +0.15$ the state $S_V = U_v S_B$ satisfies $S_B \le S_V \le S_A$. The first bound follows because U_v adds states. The other follows from $v \le a$ and stability of S_A at a. Consequently

$$S_B = D_b S_B \le D_b S_V \le D_b S_A = S_B. \tag{7}$$

For the outer return, $S_- \le S_B \le S_A$. Apply U_a to both bounds. The lower bound maps to S_A by definition, and the upper bound maps to itself. After the higher field $a' = +0.65$, the first argument used to establish $U_v S_B \le S_A$ fails for $U_{a'}$. A subsequent D_a can then stop at a different state. With both flanks uppercase an interior site in uppercase survives at score a if $h_i \le a + 2J + c$. On the rising branch with both flanks lowercase that title would require $h_i < a - 2J - c$ to convert at the same score. The interval between the two bounds admits different cases at the same score after an excursion.

APPENDIX B: LIKELIHOOD AND FINITE-STEP PREDICTION

For each bracket $(l, u]$ in score units the conditional interval probability is $\Phi((u-\mu)/\sigma_\eta) - \Phi((l-\mu)/\sigma_\eta)$, with μ from Eq. (4). Endpoint choice is immaterial under a continuous residual distribution. We multiply these probabilities conditional on an editor intercept and the title effect, then integrate those two effects. The editor and title effects are shared by the up and down brackets from a given title. With uncensored switches, $(\tau^{\uparrow} + \tau^{\downarrow})/2$ estimates the zero-barrier threshold and $(\tau^{\uparrow} - \tau^{\downarrow})/2$ estimates c. The same two directions are present in the interval likelihood. We use the isolated midpoint of the two brackets as the input field in validation. The population fit supplies the spread of h_i and the uncertainty in the calibration. The validation inputs vary by title.

The field indices are exact multiples of the displayed step. Each prediction starts at the observed initialized state. We visit positions in order, update every eligible site, and repeat until the terminal scan changes nothing. To reproduce the reported area, store the capitalization fraction after every scan with no accepted edits. Use both branches at every score, with the upper endpoint shared by the two visits. In the nested history the first 38 holds ascend, including the initial score. Successive legs contribute 15, 11, 11, 21 and 6 holds. The outer recovery at hold 89 is inside the 21-hold ascending leg. The coupling is evaluated on the current list at every visit, including when an earlier site changes during the same scan [22].

ACKNOWLEDGMENTS

We thank Bitte Norevik for verifying the collation and Ana Darves for auditing the score histories against the saved edit records.

[1] T. C. Schelling, Journal of Mathematical Sociology **1**, 143 (1971).
[2] M. Granovetter, American Journal of Sociology **83**, 1420 (1978).
[3] M. J. Salganik, P. S. Dodds, and D. J. Watts, Science **311**, 854 (2006).
[4] D. Centola, Science **329**, 1194 (2010).
[5] E. Ising, Zeitschrift für Physik **31**, 253 (1925).
[6] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, Physical Review Letters **70**, 3347 (1993).
[7] F. Preisach, Zeitschrift für Physik **94**, 277 (1935).
[8] I. D. Mayergoyz, Physical Review Letters **56**, 1518 (1986).
[9] A. A. Middleton, Physical Review Letters **68**, 670 (1992).
[10] N. E. Vassour and F. Sortil, *Row Identity in Repeated Directory Audits*, Centre Valcienne (2022).
[11] B. Norevik and S. Mehnal, *An Experimental Editor with Stable Collation Keys*, Norbjerg Institute (2023).
[12] B. W. Turnbull, Journal of the Royal Statistical Society: Series B **38**, 290 (1976).
[13] B. Efron, Annals of Statistics **7**, 1 (1979).
[14] O. Perković, K. Dahmen, and J. P. Sethna, Physical Review Letters **75**, 4528 (1995).
[15] K. Dahmen and J. P. Sethna, Physical Review B **53**, 14872 (1996).
[16] D. Dhar, P. Shukla, and J. P. Sethna, Journal of Physics A: Mathematical and General **30**, 5259 (1997).
[17] J. P. Sethna, K. A. Dahmen, and C. R. Myers, Nature **410**, 242 (2001).
[18] M. C. Kuntz and J. P. Sethna, Physical Review B **62**, 11699 (2000).
[19] A. Clauset, C. R. Shalizi, and M. E. J. Newman, SIAM Review **51**, 661 (2009).
[20] M. K. Teshari and E. Salvar, *Time Limits and Row-Order Bias in List Review*, Sefidar Institute (2023).

[21] T. Keldaran and E. M. Vestri, *Changes of Case in Paths with Stable External References*, Norbjerg Institute (2024).

[22] M. L. Halvessen and M. K. Teshari, *Sweep Logs for the Norbjerg Case-Review Instrument*, Norbjerg Institute (2024).