Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Where turnover contagion tips: Coworker network structure, job embeddedness, and retention targeting in an agent-based model

Md Monam Uddin
Lincoln University College (LUC)
Muhammad Abdul Mannan
mamannan4975@gmail.com

EXIM Bank Agricultural University Bangladesh
Norazman Saharum
Lincoln University College (LUC)

Research Article

Keywords: Employee turnover, Turnover contagion, Job embeddedness, Agent-based simulation, Social network, Tipping point

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11572304/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Where turnover contagion tips: Coworker network structure, job embeddedness, and retention targeting in an agent-based model

Abstract

Employee turnover spreads among coworkers, but the organizational consequences of that contagion are poorly understood. We build an agent-based model in which employees embedded in a coworker network start searching for a new job when shocks and the share of searching or vacant coworkers exceed an embeddedness-dependent threshold. Across about 63,000 simulation runs, contagion strength produced a sharp tipping point between modest amplification of turnover and a self-sustaining regime. Clustered networks tipped earlier than random ones, whereas hubs alone did not lower the tipping point. Higher mean job embeddedness delayed tipping strongly. Under a fixed budget, retention practices that reached hubs or brokers and were attached to positions rather than individuals raised the tipping point most. The model is stylized and uncalibrated, so the findings are hypotheses for empirical testing.

Keywords Employee turnover; Turnover contagion; Job embeddedness; Agent-based simulation; Social network; Tipping point

1. Introduction

Employee turnover has been studied for more than a century (Hom et al. 2017), and its aggregate form, collective turnover at the group, unit or organizational level, is now recognized as a phenomenon with its own causes and consequences (Hausknecht and Trevor 2011). A meta-analysis of 82 studies and 694 effect sizes found that collective turnover is negatively related to a range of performance outcomes, more strongly for proximal than for distal outcomes (Heavey et al. 2013), whereas an earlier meta-analysis on firm performance reported a mean corrected correlation of −0.03 that nevertheless varied systematically across contexts (Hancock et al. 2013). Because identical turnover rates can carry very different meanings, and therefore consequences (Hausknecht and Holwerda 2013), the temporal and structural pattern of departures is a managerial concern in its own right, not only their average rate.

At the individual level, the unfolding model describes voluntary exit as the outcome of distinct decision paths, most of which are set in motion by shocks (Lee and Mitchell 1994). Job embeddedness, comprising links to people and groups, fit with job, organization and community, and the sacrifice that leaving would entail, predicts intent to leave and voluntary turnover over and above satisfaction, commitment, perceived alternatives and job search (Mitchell et al. 2001); a later study separated on-the-job from off-the-job embeddedness and found that the two predict different outcomes (Lee et al. 2004).

Individuals, however, do not decide alone. Felps et al. (2009) developed and tested a model of turnover contagion in which coworkers' job embeddedness and job search behaviors influence an employee's decision to quit. In a sample of 45 branches of a regional bank and 1,038 departments of a national hospitality firm, these coworker variables explained variance in individual voluntary turnover over and above individual-level and group-level predictors. That evidence establishes that a coworker effect exists and estimates its average strength on individual decisions. It does not show what such an effect implies when it is transmitted repeatedly through a network of ties over time. Whether contagion merely nudges turnover upward or can make departures self-sustaining, how the pattern of ties among coworkers alters the answer, and whether a limited retention budget is better spent on some employees than on others are questions about system dynamics, and a regression on a single cross-section cannot answer them. Earlier network work on turnover points the same way: a survey of 170 employees in one organization supported structural-equivalence, social-influence and erosion models of turnover (Feeley and Barnett 1997), but it, too, describes a snapshot rather than the dynamics that such influence sets in motion.

Simulation is well suited to such questions. Management scholars have argued that simulation offers internal validity and facility with longitudinal, nonlinear and process phenomena, and that its primary value lies in creative experimentation that produces novel theory (Davis et al. 2007); it has

nonetheless been embraced more slowly in management than in some neighboring social sciences (Harrison et al. 2007). Formal precursors exist outside the turnover literature: threshold rules on random networks can generate global cascades (Watts 2002), and comparisons of agent-based and aggregate models show that heterogeneity and network structure change diffusion dynamics (Rahmandad and Sterman 2008). This paper brings that logic to turnover. It builds and analyzes an agent-based model (ABM) of turnover contagion on a coworker network and asks four research questions:

- **RQ1.** Does social contagion of search and quitting produce a tipping point between a regime of modest amplification and a regime of self-sustaining turnover, and where does it lie?
- **RQ2.** How do the structure of coworker ties (clustering, degree heterogeneity, team structure) shift that tipping point?
- **RQ3.** How do the average level and the heterogeneity of job embeddedness shift it?
- **RQ4.** Under a fixed budget, which rule for targeting embeddedness-building retention practices raises the tipping point most?

The study is computational throughout: no survey, interview or archival organizational data are collected or used. Its contribution is theoretical and methodological. It converts a verified empirical mechanism, coworker contagion (Felps et al. 2009), into a dynamic system whose regimes, sensitivities and policy levers can be mapped exhaustively, and it documents the model so that every result can be regenerated from code. Section 2 develops the theoretical background and expectations, Section 3 describes the model, experiments and verification procedures, Section 4 reports the results, and Section 5 discusses implications, limitations and an empirical-calibration agenda.

The business relevance is direct. Retention budgets are finite, and managers must decide which employees and which units to invest in; the model shows how that decision interacts with the structure of the workforce's ties, and why a unit with an unremarkable turnover rate can nevertheless be close to a self-sustaining wave of departures.

2. Theoretical background and expectations

2.1 Embeddedness as a threshold on withdrawal pressure

Threshold models of collective behavior define an actor's threshold as the proportion of others who must make a decision before that actor does so, and they stress that the full distribution of thresholds, not its mean, determines the collective outcome: groups with similar average preferences can generate very different results (Granovetter 1978). We import this logic into turnover. In the unfolding model, shocks push employees toward withdrawal (Lee and Mitchell 1994); job embeddedness raises the cost of leaving (Mitchell et al. 2001). We therefore treat embeddedness as raising the threshold that accumulated withdrawal pressure must cross before an employee starts searching, so that heterogeneity in embeddedness is heterogeneity in thresholds.

2.2 Contagion as social signaling

Coworkers' job search behaviors and embeddedness influence an employee's decision to quit (Felps et al. 2009). We operationalize the channel as an observable signal: an employee whose coworkers are searching, or whose coworkers have left and left their positions vacant, experiences additional pressure proportional to the share of ties that are signaling. Pressure from signals is transient, disappearing when the signals do, whereas pressure from shocks decays gradually.

2.3 Network structure and expectations

Four canonical structures bracket the range of workplace tie patterns. Random graphs (Erdős and Rényi 1959) have no structure beyond density. Small-world networks combine strong local clustering with short path lengths, and contagion such as infectious disease is known to spread more easily on them than on regular lattices (Watts and Strogatz 1998). Scale-free networks, generated by growth and preferential attachment (Barabási and Albert 1999), contain a few very highly connected individuals. Stochastic block models (Holland et al. 1983) represent organizations divided into dense teams with sparser ties between them. Our expectations are:

Expectation 1 (nonlinearity). Because every quit or search can push coworkers who themselves have thresholds over those thresholds, aggregate turnover should respond nonlinearly to contagion strength, with a tipping point separating a regime of modest amplification from a regime of self-sustaining turnover.

Expectation 2 (clustering). Overlapping neighborhoods concentrate signals within cohesive groups, so that several signaling coworkers reach the same employees at the same time. Holding mean degree fixed, greater clustering should lower the contagion strength at which the tipping point occurs, whereas degree heterogeneity alone should matter less because the signal is a share of ties rather than a count.

Embeddedness (RQ3) and targeting (RQ4) are treated as exploratory because competing arguments exist. Higher mean embeddedness should raise thresholds and therefore the tipping point, but at a fixed mean, greater heterogeneity creates both fragile, low-threshold employees who ignite cascades and highly embedded employees who may act as firebreaks. For targeting, meta-analytic evidence ranks human resource inducements and investments, and job embeddedness signals, among the stronger antecedents of collective turnover (Heavey et al. 2013); which employees should receive such practices when the budget covers only a fraction is an open question that depends on network position and on baseline embeddedness.

3. Model and methods

3.1 Modelling approach and reporting standards

The model is an agent-based model, in which aggregate phenomena are investigated by simulating the behavior of individual agents (Rand and Rust 2011). We follow the verification and validation guidelines of Rand and Rust (2011) and describe the model with the Overview, Design concepts and Details (ODD) protocol (Grimm et al. 2006; Grimm et al. 2010); the complete ODD description is provided in Appendix A. Time is measured in abstract periods, and no parameter is calibrated to a specific organization: the model is a theory-building instrument intended to map regimes and sensitivities, not to forecast turnover in any firm.

3.2 Model description

An organization has N positions, each occupied by one employee or temporarily vacant, connected by an undirected coworker network A. Each employee i has job embeddedness $E_i \in [0, 1]$ drawn from a Beta distribution with mean μ_E and concentration κ_E, and a search threshold that rises with it:

$$E_i \sim \text{Beta}(\mu_E\kappa_E, (1-\mu_E)\kappa_E), \quad \theta_i = \theta_0 + \lambda E_i \qquad (1)$$

Withdrawal pressure P_i accumulates from exogenous shocks, following the unfolding model's emphasis on shocks (Lee and Mitchell 1994), and decays geometrically:

$$P_i(t) = \rho P_i(t-1) + x_i(t)y_i(t), \quad x_i \sim \text{Bernoulli}(p_{\text{shock}}),\ y_i \sim \text{Exponential(mean } m_{\text{shock}}) \qquad (2)$$

A coworker j emits a signal $\sigma_j(t) = 1$ if j is searching or its position is vacant at the start of period t, and 0 otherwise. The social input to employee i is the share of its ties that are signaling, and effective pressure adds this input weighted by the contagion strength β:

$$I_i(t) = \Sigma_j A_{ij}\, \sigma_j(t) / \max(d_i, 1), \quad Q_i(t) = P_i(t) + \beta I_i(t) \qquad (3)$$

Employed agents start searching when effective pressure exceeds their threshold ($Q_i > \theta_i$). A searching agent whose effective pressure has fallen below a fraction h of the threshold reverts to employment with probability δ; otherwise it quits with hazard α. A quit creates a vacancy that signals to coworkers for τ periods and is then backfilled by a new hire with a fresh embeddedness draw and zero pressure. All agents update synchronously. The model is therefore a Granovetter-type threshold process (Granovetter 1978) embedded in a birth-death system that keeps the workforce constant and allows turnover to reach a steady state without contagion. A robustness variant replaces the share by a count of signaling ties scaled by mean degree. Fig. 1 summarizes the model.

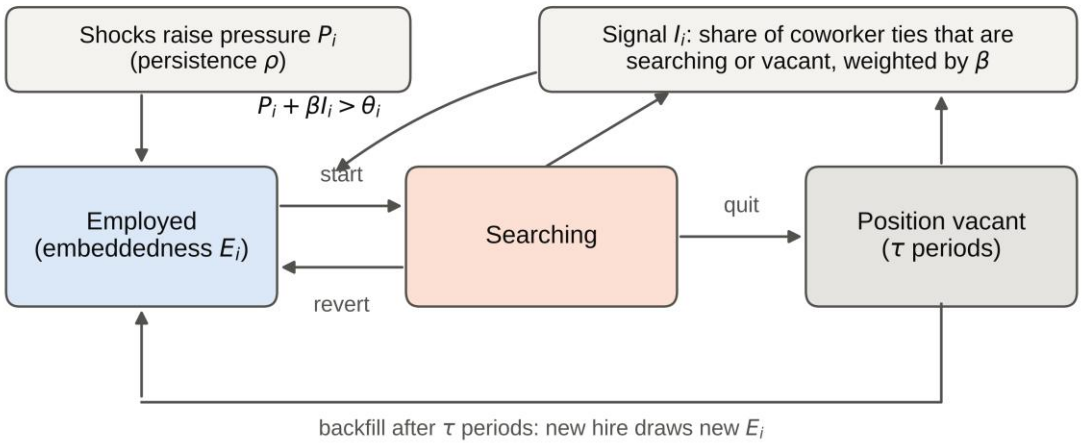


Fig. 1 Structure of the turnover-contagion model: states, transitions and the social signal

3.3 Coworker network structures

All networks have $N = 500$ nodes and a target mean degree of 8: Erdős–Rényi random graphs (ER; Erdős and Rényi 1959); Watts–Strogatz small-world graphs (WS; rewiring probability 0.1 unless stated; Watts and Strogatz 1998); Barabási–Albert scale-free graphs (BA; Barabási and Albert 1999); and stochastic block model graphs (SBM; Holland et al. 1983) with 25 teams of 20, in which 75% of each employee's expected ties lie inside the team. A fresh network is generated for every replicate with the NetworkX library (Hagberg et al. 2008). With contagion off ($\beta = 0$) the network is irrelevant, which we use as a verification test.

3.4 Retention interventions

A retention practice is represented as an embeddedness boost of $\Delta E = 0.3$ (capped at 1) for a share $b \in \{0.10, 0.20\}$ of positions, so that a fixed budget treats a fixed number of employees. Five targeting rules choose the treated positions: random; highest degree (hubs); lowest initial embeddedness; highest betweenness centrality (brokers, approximated from 100 sampled sources); and whole teams, defined by Louvain communities (Blondel et al. 2008), treated until the budget is exhausted. In the main analysis the boost is attached to the position and persists across hires, as with a structural practice such as team design or mentoring; in a robustness analysis it is attached to the person and lost when that person quits.

3.5 Outcomes

Each run discards the first 100 periods as burn-in and measures 300 periods. The **turnover rate** is quits per position per period; **prevalence** is the share of positions that are searching or vacant; and the **overdispersion index** is the variance-to-mean ratio of quits per period (about 1 for independent departures, above 1 when departures cluster in time). A run shows a **sustained turnover collapse** if

prevalence reaches 0.5 in any period and averages at least 0.5 over the final 50 periods. The central estimand is the **tipping point** β_{50}, the contagion strength at which half of runs collapse. It is estimated by isotonic (monotone) smoothing of collapse frequency over a grid of β followed by linear interpolation to 0.5, with 95% percentile bootstrap intervals from resampling runs within each β cell. These intervals express sampling uncertainty conditional on the grid and exclude discretization error, which is bounded by half the local grid step. Time to collapse is compared across structures with log-rank tests.

3.6 Experimental design

Table 1. Simulation experiments.

Experiment	Purpose	Design	Runs
E1	Topology × β	4 structures × 12 β values × 100 replicates	4,800
E1b	WS rewiring sweep	6 rewiring probabilities × 10 β values × 60 replicates	3,600
E2	Embeddedness mean × heterogeneity (WS)	5 × 5 cells; coarse β grid, then local grid of step 0.125 × 40 replicates	20,500
E3	Retention targeting	4 structures × (1 + 5 rules × 2 budgets) × 10 β values × 40 replicates; plus person-attached robustness	18,800
E4	Global sensitivity (Sobol')	11 parameters, base sample 256, 3 replicates per design point	9,984
E5	Robustness	influence specification, organization size, vacancy duration, hysteresis, collapse definition	5,200

Every run uses its own seed, so any cell can be regenerated exactly. Parameter values and the ranges explored are listed in Table 2; baseline values were chosen so that, without contagion, the organization sits at a low, stable turnover rate (0.48% of positions per period) with 5.4% of positions searching or vacant on average.

Table 2. Model parameters.

Symbol	Meaning	Baseline	Range explored
N	Positions (employees) in the organization	500	250–2,000 (robustness)
k	Target mean degree (coworker ties per position)	8	4–16 (sensitivity)
μ_E	Mean job embeddedness	0.5	0.3–0.7
κ_E	Beta concentration of embeddedness (larger = more homogeneous)	8	2–64
θ_0, λ	Threshold intercept and embeddedness slope	1, 2	0.5–1.5; 1–3 (sensitivity)
ρ	Persistence of withdrawal pressure	0.8	0.6–0.9 (sensitivity)
p_{shock}, m_{shock}	Shock probability per period; mean shock size	0.05, 1	0.02–0.10; 0.5–1.5 (sensitivity)
β	Contagion strength (weight on the share of signaling ties)	varied	0–8
α	Quit hazard while searching	0.15	0.05–0.30 (sensitivity)
δ, h	Revert probability; reversion cut-off as share of threshold	0.2, 0.6	fixed
τ	Vacancy duration before backfilling (periods)	4	2–8
T, burn-in	Simulated periods; discarded initial periods	400, 100	fixed
ΔE, b	Embeddedness boost; share of positions treated	0.3, —	b = 0.10, 0.20

3.7 Verification and validation

Verification asks whether the code implements the conceptual model. We (i) wrote a second, deliberately naive implementation in plain Python with one agent object at a time and compared the

distributions of its outputs with those of the vectorized implementation (Kolmogorov–Smirnov tests for continuous outputs, Fisher's exact test for collapse frequency); (ii) tested that with $\beta = 0$ outcomes do not depend on the network; (iii) tested that on a network without ties β has no effect; (iv) tested reproducibility from the seed and accounting identities (bounds on prevalence, quits and searching shares); (v) tested that turnover rises monotonically with β below the tipping point; and (vi) checked convergence in the number of replicates and stationarity of the measurement window (Section 4.1).

Validation asks whether the model is an adequate representation of its target. No empirical calibration was attempted, so we claim only structural validity and an internal consistency check of emergent individual-level behavior. Using panel data recorded from simulation runs, we estimate logistic models of search onset on own embeddedness, on the coworker signal share, and on coworkers' mean embeddedness, and test whether coworker variables add explanatory power over individual-level predictors, the logic of the test of Felps et al. (2009). This checks that the model generates, at the individual level, the type of pattern that the empirical contagion literature reports; it does not validate the magnitudes.

3.8 Sensitivity and software

Global sensitivity analysis used Sobol' variance decomposition with first- and total-order indices on a WS network, using the SALib library (Herman and Usher 2017) and the sampling design and estimators of Saltelli et al. (2010) without second-order terms. Simulations used Python 3.13.16, NumPy 2.5.3, SciPy 1.18.1, NetworkX 3.6.1, pandas 3.0.5, statsmodels 0.15.0, scikit-learn 1.9.1 and SALib 1.6.0. Experiments were parallelized across two processes with independent seeds.

4. Results

4.1 Verification and internal validation

The two independent implementations were statistically indistinguishable: across 12 comparisons (4 outputs × 3 contagion strengths, 150 replicates each) the smallest Kolmogorov–Smirnov p-value was 0.29, and at $\beta = 4$ both implementations produced a sustained collapse in 150 and 150 of 150 runs, respectively (Fisher exact $p = 1.00$). With $\beta = 0$ the four network structures returned identical mean turnover (0.476% per period; the runs use common random numbers, so equality to numerical precision is the expected outcome and is a stricter test than a non-significant ANOVA). On a network without ties $\beta = 0$ and $\beta = 8$ gave identical results; identical seeds reproduced identical output and different seeds did not; all accounting identities held; and turnover rose monotonically with β below the tipping point (Spearman $\rho = 1.0$ over $\beta = 0$–3). Estimated mean turnover stabilized quickly with replicates: the 95% half-width was 0.8% of the mean with 100 replicates. One check revealed a small blemish: turnover in the second half of the measurement window was slightly lower than in the first half (0.481% vs 0.469% at $\beta = 0$, Wilcoxon $p = 0.0003$; 0.656% vs 0.635% at $\beta = 3$, $p = 0.0027$), a

decline of about 2–3% that indicates a slowly decaying transient after the 100-period burn-in (the initial workforce is drawn without regard to hiring history). It is too small to affect the tipping-point comparisons reported below, which are all differences between conditions measured over the same window.

In the internal consistency check on 1,670,160 employee-period observations (16,244 search onsets) from 12 WS runs at $\beta = 3$, the standardized coefficient of own embeddedness was negative (-0.32) and that of the coworker signal share positive (0.35, 95% CI 0.34 to 0.37; odds ratio 1.42 per standard deviation). Adding the signal share raised McFadden's pseudo-R^2 from 0.009 to 0.024 (likelihood-ratio $\chi^2 = 2{,}877$, df = 1, $p < 0.001$), whereas coworkers' mean embeddedness added nothing once the signal was included ($\chi^2 = 0.47$, $p = 0.49$). The model therefore reproduces the qualitative pattern that coworker behavior predicts individual exit above individual-level predictors (Felps et al. 2009), but because the signal is built into the model this is a consistency check, not evidence of validity. The finding that coworkers' embeddedness is redundant is a feature of the model's mechanism: embeddedness acts only through the individual threshold, and it is not a claim about the empirical effect that Felps et al. (2009) report for coworkers' embeddedness.

4.2 Contagion produces a tipping point, and clustering lowers it

Without contagion the baseline organization has a turnover rate of 0.48% per period, 5.4% of positions searching or vacant, and an overdispersion index of 1.00, consistent with independent departures. As β rises, turnover at first increases modestly: at $\beta = 3$ it is 1.34 to 1.37 times the baseline in all four structures and no run collapses. The system then tips abruptly. In collapsed runs 91% of positions are searching or vacant on average and the turnover rate is 8.4% per period, about 17 times the baseline. Expectation 1 is therefore supported: the response to contagion is nonlinear, with a narrow band separating modest amplification from self-sustaining turnover. We stress that the collapse state is the model's stylized runaway regime, in which management never responds, and is not a forecast of what an organization would look like.

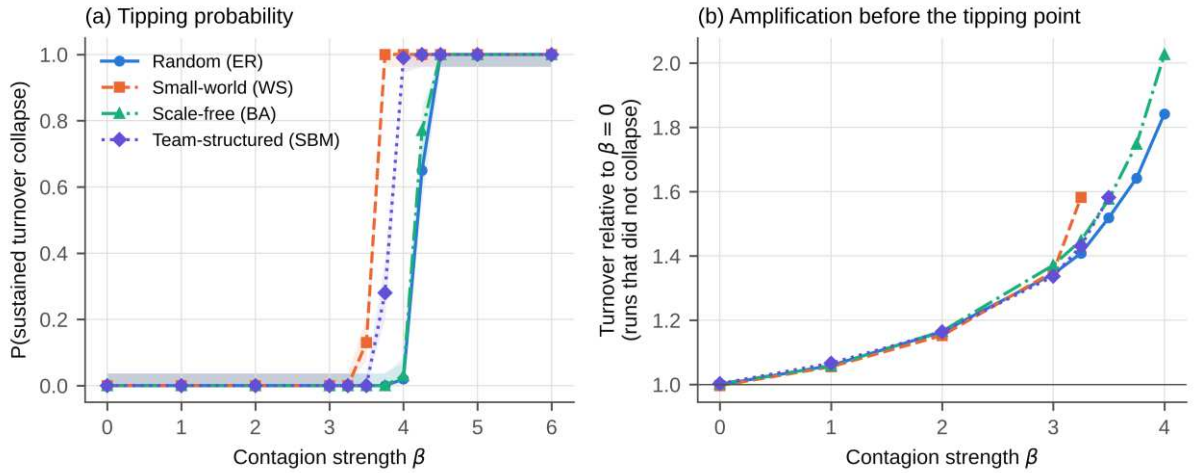


Fig. 2 Contagion strength and sustained turnover collapse by network structure (500 positions, 100 replicates per point). Panel (a) shows the share of runs in sustained collapse as a function of β (the tipping point β_{50} is where it crosses one half; bands are 95% intervals) and panel (b) turnover relative to β = 0 in runs that did not collapse. Structures are distinguished by marker and line style as well as color

The tipping point depended on structure (Fig. 2; 95% bootstrap interval in brackets, conditional on the β grid): WS 3.61 [3.60, 3.62], SBM 3.83 [3.80, 3.85], BA 4.16 [4.14, 4.18], ER 4.19 [4.16, 4.23]. The clustered structures tip at lower contagion strength than the random graph, by 0.58 for WS and 0.36 for SBM. Time to collapse showed the same ordering: at β = 4.0 the median time to collapse was 53 periods for WS, 83 for SBM, 60 for BA and 77 for ER, and log-rank tests against ER were highly significant for WS and SBM ($p < 0.001$ at β = 3.75, 4.0 and 4.5). Degree heterogeneity alone did not lower the tipping point: BA, whose largest hub had on average 78 ties against 18 in ER, was indistinguishable from ER at β = 4.0 (log-rank $p = 0.65$) and only marginally different at 4.5 ($p = 0.03$). Expectation 2 is supported, with the qualification that clustering is not the only property that differs across structures (mean clustering coefficients were 0.47 for WS, 0.18 for SBM, 0.06 for BA and 0.016 for ER).

The rewiring sweep isolates this property within one family (Fig. 3). Raising the WS rewiring probability from 0 to 1 lifted the tipping point from 3.40 to 4.17 while clustering fell from 0.64 to 0.01; across the six settings the rank correlation between clustering and the tipping point was ρ = -0.94 ($p = 0.005$; 6 points, so this is descriptive). Rewiring also shortens average path length, so the contributions of clustering and path length are not separated here. The overdispersion index at β = 3 stayed close to 1 (1.02–1.07), so the clustering of departures in time gave no early warning of the approaching tipping point.

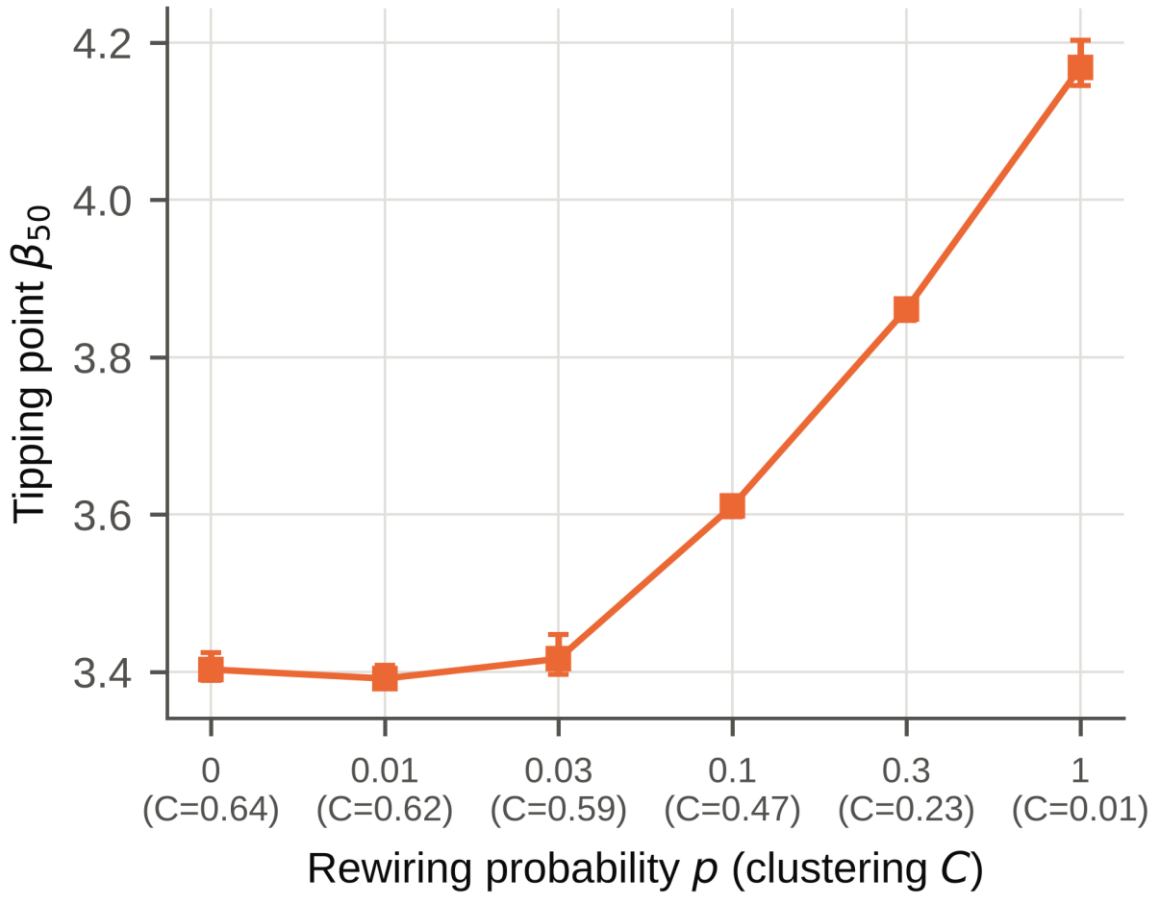


Fig. 3 Tipping point β_{50} as a function of the rewiring probability in Watts–Strogatz graphs (mean clustering coefficient *C* in parentheses; six rewiring probabilities; 60 replicates per β value; 95% bootstrap intervals conditional on the β grid)

4.3 Embeddedness: level raises the tipping point, heterogeneity raises it too

Mean embeddedness had the largest effect of any organizational lever examined (Fig. 4a). Averaged over heterogeneity levels, the tipping point rose from 2.77 at $\mu_E = 0.3$ to 4.44 at $\mu_E = 0.7$. An ordinary least squares fit over the 25 cells ($R^2 = 0.992$) gave +4.17 per unit of mean embeddedness (95% CI 4.00 to 4.34), or about 0.42 for each 0.1 increase. Heterogeneity mattered in the direction of a firebreak effect: the more concentrated (homogeneous) the distribution, the lower the tipping point, from 3.81 at $\kappa_E = 2$ to 3.44 at $\kappa_E = 64$ (coefficient on $\ln \kappa_E = -0.069$, 95% CI -0.083 to -0.055). The magnitude is small next to that of the mean (the full range of κ_E shifts β_{50} by about 0.37, against 1.67 for the range of μ_E), but it is notable that heterogeneity left baseline turnover virtually unchanged (Fig. 4b): two organizations with the same turnover rate can have different distances to the tipping point. Because the cells were estimated on a local β grid with step 0.125 and 40 replicates, differences below about 0.1 should not be interpreted.

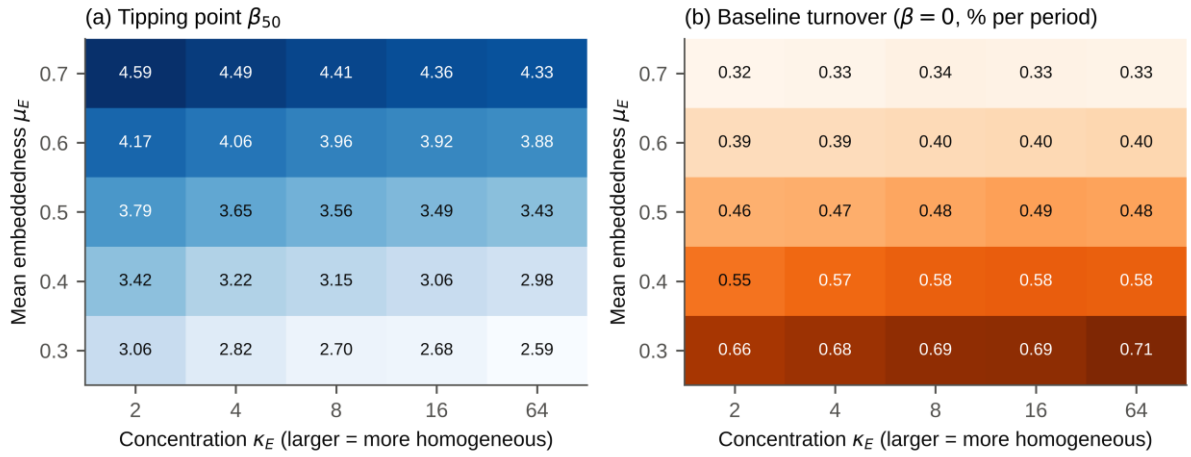


Fig. 4 Embeddedness and the tipping point in the WS network. (a) Tipping point β_{50} for each combination of mean embeddedness μ_E and concentration κ_E. (b) Baseline turnover without contagion (% of positions per period). Cell values are printed so that the figure does not rely on color

4.4 Retention targeting

Table 3. Gain in the tipping point $\Delta\beta_{50}$ from an embeddedness-building practice that treats 20% of positions, by targeting rule and network structure.

Targeting rule (20% of positions)	WS	ER	BA	SBM
No intervention (β_{50})	3.59	4.19	4.15	3.80
Random	+0.12	+0.30	+0.27	+0.23
Hubs	+0.27	+0.43	+0.72	+0.32
Low embeddedness	+0.18	+0.26	+0.30	+0.28
Brokers	+0.23	+0.38	+0.65	+0.30
Whole teams	+0.02	+0.21	+0.24	+0.07

$\Delta\beta_{50}$ is the increase over the no-intervention tipping point in the same structure (first row). Positions retain the practice across hires. 40 replicates per β value; some intervals are degenerate because the β grid is coarse near the tipping point (see Section 5.2). Results for the 10% budget are in Fig. 5.

Targeting mattered for the tipping point but hardly for ordinary turnover. At β = 3, below every tipping point, treating 10% of positions reduced the turnover rate by 4–8% and treating 20% by 10–15%, almost regardless of rule. The rules differed instead in how far they pushed the tipping point out (Table 3, Fig. 5). Treating hubs or brokers was the strongest rule in every structure, and its advantage was largest in the scale-free network, where treating hubs moved β_{50} by 0.72 against 0.27 for random targeting (a factor of 2.6); in the clustered WS network the factor was 2.2. Targeting the least embedded employees performed close to random targeting (differences of at most 0.06 in $\Delta\beta_{50}$ in all four structures; WS 0.18 against 0.12). Treating whole teams was the weakest rule in the clustered structures ($\Delta\beta_{50}$ = 0.02 in WS and 0.07 in SBM, against 0.12 and 0.23 for random targeting): concentrating protection inside a few cohesive groups leaves the remaining groups, in which contagion spreads fastest, unprotected.

The result depends on whether the practice stays with the position. When the boost was attached to the person and lost on quitting, a 20% budget in the WS network raised β_{50} by only 0.03 (random), 0.03 (hubs) and 0.03 (low embeddedness), compared with 0.12, 0.27 and 0.18 when it persisted across

hires, and the advantage of targeting hubs disappeared. Practices that are embedded in the structure of work are far more effective in this model than practices tied to the people who leave first.

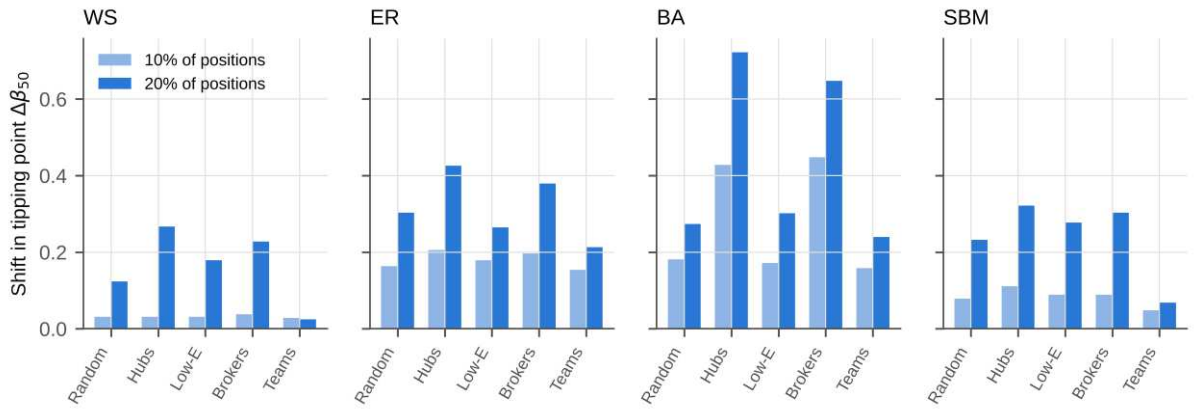


Fig. 5 Gain in tipping point $\Delta\beta_{50}$ by targeting rule, budget and network structure (light bars 10% of positions, dark bars 20%)

4.5 Global sensitivity

Contagion strength dominated the variance of both mean prevalence (first-order index 0.64, total-order 0.87) and turnover (first-order 0.52, total-order 0.76) when 11 parameters varied jointly over the ranges in Table 2 (9,984 runs; Fig. 6). The next most influential parameters by total-order index for mean prevalence were θ_0 (0.17), p_{shock} (0.15), μ_E (0.15), m_{shock} (0.13), λ (0.12). Mean degree (k) and embeddedness concentration (κ_E) contributed little (0.05 and 0.00). The first-order indices of mean prevalence sum to 0.77, so about 23% of its variance arises from interactions, as expected in a system with a tipping point, where the effect of one parameter depends on where the others place the system relative to the threshold. Indices for the overdispersion index are not interpretable: their total-order estimates exceed 1 with very wide intervals, because that statistic is heavy-tailed when runs collapse, and we do not report them as findings. With a base sample of 256 and three replicates per design point, small indices (below about 0.1) are indistinguishable from zero.

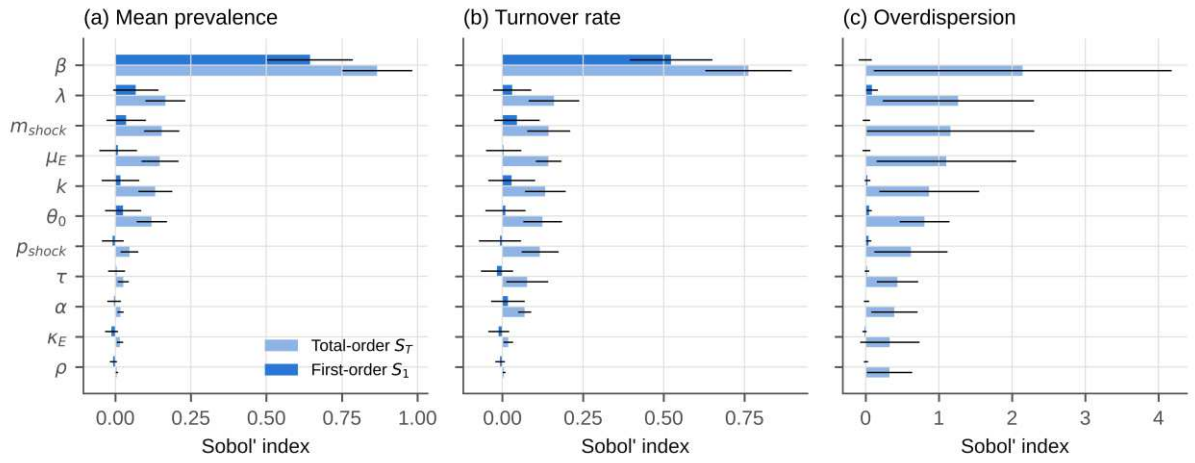


Fig. 6 Sobol' sensitivity indices (WS network; 11 parameters varied jointly). Dark bars are first-order and light bars total-order indices with 95% intervals. Panel (c), the overdispersion index, is unreliable (total-order indices above 1 with wide intervals) and is shown only for completeness

4.6 Robustness

Five modeling choices were varied (Fig. 7). (a) Replacing the share of signaling ties by a count scaled by mean degree lowered all tipping points (ER 4.24 to 3.72; WS 3.70 to 3.52; BA 4.24 to 3.75; SBM 3.75 to 3.31) because hubs then receive more total signal, but the clustered structures (SBM, WS) still tipped before the random graph, and BA was no longer distinguishable from ER; the clustering result is robust to this choice, whereas the claim that hubs do not matter is specific to the share specification. (b) Organization size from 250 to 2,000 positions changed the WS tipping point by only 0.04 (3.59 to 3.62). (c) A longer vacancy period lowered the tipping point, because vacancies signal for longer ($\tau = 2$: 3.75; $\tau = 4$: 3.70; $\tau = 8$: 3.25); the $\tau = 8$ estimate sits at the edge of the β grid and has no interval. (d) The system shows hysteresis: at $\beta = 3.5$ only 20% of runs starting from a healthy workforce collapsed, but 100% of runs starting with 30% of employees already searching did, whereas at $\beta = 4.0$ all runs collapsed from either start, so the tipping point is better understood as the edge of a bistable band than as a single threshold. (e) Defining collapse by a prevalence of 0.3 or 0.7 instead of 0.5 moved tipping points by at most 0.06. Several of these intervals are degenerate because the grid used in E5 is coarse; estimates should be read to the nearest 0.1 or so.

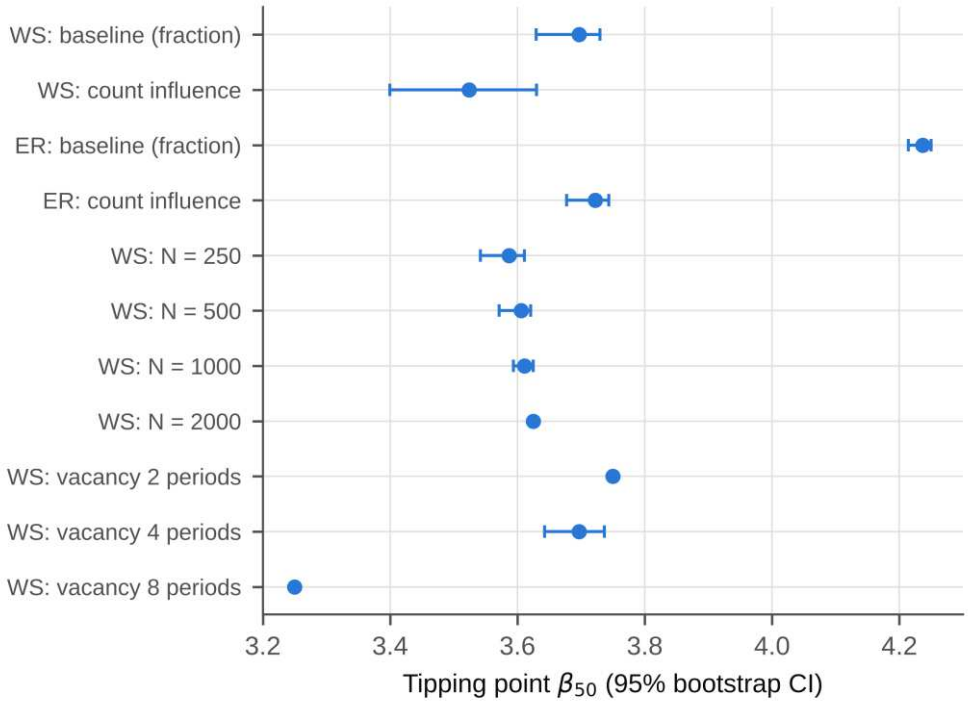


Fig. 7 Robustness of the tipping point β_{50} to influence specification, organization size and vacancy duration. Points without an interval have degenerate bootstrap distributions because the β grid is coarse near the estimate

5. Discussion

5.1 Findings and theoretical implications

Four findings follow from the experiments. First, turning a verified empirical mechanism, coworker contagion (Felps et al. 2009), into a dynamic system reveals behavior that cross-sectional estimates of an average coworker effect cannot show: modest amplification over a wide range of contagion strengths, then an abrupt transition to a self-sustaining regime, with a bistable band in which the outcome depends on the starting state. In threshold terms (Granovetter 1978), the same average sensitivity to peers can be benign or catastrophic depending on how close the system stands to the cascade condition, and the overdispersion of departures gave no warning before the transition, so it was not a usable early-warning indicator here, even though such indicators have been proposed for critical transitions in other complex systems (Scheffer et al. 2009). Second, structure matters through local overlap rather than through hubs: clustered structures, in which signals reach the same employees from several directions, tipped earlier, in line with the small-world argument that local clustering with short paths favors spread (Watts and Strogatz 1998), while degree heterogeneity alone did not lower the tipping point under proportional influence. Third, embeddedness operates as a system-level buffer in a way that individual-level studies of embeddedness (Mitchell et al. 2001; Lee et al. 2004) do not capture: the mean level sets the distance to the tipping point, and, at the same mean, a more unequal distribution of embeddedness protected the system slightly, because highly embedded employees act as firebreaks. Fourth, retention practices should be judged by where they sit in the

network and whether they stay with the position: spreading a limited budget across hubs and brokers, and building practices into the position or team structure rather than the individual, raised the tipping point most.

These results sharpen the literature on collective turnover. Meta-analytic work shows that collective turnover is associated with worse performance (Heavey et al. 2013) and that the relationship varies with context (Hancock et al. 2013); the present model adds a mechanism by which context could matter, namely the structure of ties and the distribution of embeddedness which determine whether a given average level of departures is stable or on the edge of a cascade. The same turnover rate can therefore carry different risk, consistent with the argument that turnover levels alone can mislead (Hausknecht and Holwerda 2013).

5.2 Limitations

The model is stylized and must be read as a theory-building instrument. (i) It is not calibrated to data: the contagion strength β has no empirical counterpart, and the tipping points reported are properties of the model, not of any real organization. (ii) Management does not respond: the collapse regime, in which about nine in ten positions are searching or vacant, is a runaway regime that real organizations would counter, and an endogenous managerial response could eliminate it or shift the tipping point. (iii) The network is static; employees do not form or drop ties, and departures do not rewire the network. (iv) Employees are identical apart from embeddedness; there is no heterogeneity in susceptibility, in the relationship of ties to tenure or in job type, and updating is synchronous. (v) Contagion is operationalized as a signal from searching and vacant positions; Felps et al. (2009) studied coworkers' search behaviors and embeddedness, and the present model does not separate the channels. (vi) Interventions act only on embeddedness and have no cost beyond a fixed budget share. (vii) Rewiring in the Watts–Strogatz sweep changes clustering and path length together. (viii) Statistical uncertainty is partly numerical. Bootstrap intervals for tipping points are narrow and conditional on the β grid, several are degenerate where the grid is coarse, the sensitivity analysis uses a modest sample, and the measurement window shows a small decaying transient (Section 4.1). Claims should be read at a resolution of about 0.1 in β. (ix) The simulation code was written with the assistance of a large language model, which is why it was checked against an independent implementation and a battery of tests, but verification cannot rule out every defect.

5.3 Implications for practice, and an empirical agenda

If the mechanism is realiztic, three tentative implications follow. Managers who track only the turnover rate may not see how close the organization is to a cascade; indicators of structure (clustered teams, concentrated embeddedness) and of the share of employees who are visibly searching are more informative. Retention effort is better spread across highly connected people and across teams than concentrated in a few teams, and structural practices outlast individual ones. Finally, building

embeddedness has a system-wide benefit in addition to its individual one. These are hypotheses for empirical work, not prescriptions.

The natural next step is empirical calibration. Longitudinal human resource records combined with a measured coworker network (for example from collaboration or reporting data) would allow β to be estimated from observed individual hazards in the manner of Felps et al. (2009), after which the simulated tipping points could be compared with the distance of real units from instability. Extending the model with endogenous management responses, dynamic ties, heterogeneous susceptibility and costed interventions would test whether the qualitative conclusions survive.

6. Conclusion

An agent-based model of turnover contagion on workplace networks yields a tipping point between modest amplification and self-sustaining turnover. Clustered networks tip earlier than random or scale-free ones, higher and more unequally distributed embeddedness delays the tipping point, and a fixed retention budget delays it most when it is spread over hubs and brokers and attached to positions rather than to individuals. All of this is shown under stylized assumptions and without empirical calibration; the contribution is a documented, verified model whose mechanisms can now be confronted with data.

Appendix A. ODD protocol

The model is described following the ODD protocol (Grimm et al. 2006; Grimm et al. 2010).

A.1 Purpose and patterns

Purpose: to examine how the structure of coworker ties, the level and heterogeneity of job embeddedness, and the targeting of retention practices shape the dynamics of voluntary turnover when departures are socially contagious. Patterns used to judge adequacy: a stable low turnover rate without contagion; turnover that rises with contagion strength; network independence at $\beta = 0$; and, at the individual level, search onset that falls with own embeddedness and rises with the share of signaling coworkers.

A.2 Entities, state variables and scales

Entities are positions (agents) and an undirected network of coworker ties. State variables per position: state $\in$ {employed, searching, vacant}; embeddedness $E_i \in [0,1]$; threshold θ_i; withdrawal pressure $P_i \geq 0$; vacancy timer; treated flag. Global variables are the parameters of Table 2. One time step is one abstract period; a run lasts 400 periods, the first 100 of which are discarded for outcome measurement (collapse is detected over all 400 periods). The default organization has 500 positions.

A.3 Process overview and scheduling

In each period, with all updates synchronous and based on the state at the start of the period: (1) pressure is updated, $P_i \leftarrow \rho P_i$ + shock (employed and searching positions only; vacant positions have zero pressure); (2) signals σ_j are read (searching or vacant) and the social input I_i computed; (3) effective pressure $Q_i = P_i + \beta I_i$ is compared with the threshold; (4) employed agents with $Q_i > \theta_i$ start searching; searching agents with $Q_i < h\theta_i$ revert to employment with probability δ, and the remaining searching agents quit with probability α; (5) quitting agents create vacancies (pressure reset, timer set to τ); (6) vacancies that were open at the start of the period count down and, when the timer is exhausted, are filled by a new employee with a fresh embeddedness draw and zero pressure; (7) outcomes are recorded.

A.4 Design concepts

Basic principles: threshold models of collective behavior (Granovetter 1978); shocks as triggers of withdrawal (Lee and Mitchell 1994); embeddedness as a cost of leaving (Mitchell et al. 2001); contagion of coworker search and quitting (Felps et al. 2009). *Emergence:* the tipping point, bistability and the turnover rate emerge from individual threshold crossings. *Adaptation and objectives:* none; agents follow rules. *Learning and prediction:* none. *Sensing:* agents observe the signaling state of their direct coworkers. *Interaction:* through the signal share along ties. *Stochasticity:* network generation, embeddedness draws, shocks, reversion and quit draws, and the selection of tie-breaks. *Collectives:* teams (stochastic block model) and Louvain communities (Blondel et al. 2008) are used for the team intervention only. *Observation:* turnover rate, onset rate, prevalence, peak prevalence, overdispersion, clustering, time to prevalence thresholds.

A.5 Initialization

A network of the specified type is generated; embeddedness is drawn independently from the Beta distribution; the treated positions are selected and their embeddedness is increased by ΔE (capped at 1); all agents start employed with zero pressure (the hysteresis check starts a share of agents in the searching state with pressure equal to their threshold).

A.6 Input data

None. The model uses no external data.

A.7 Submodels

Networks: ER with edge probability $k/(N - 1)$; WS ring lattice with k neighbors and rewiring probability p_{ws}; BA with $k/2$ edges per new node; SBM with 25 teams of 20 and 75% of expected ties inside teams. *Interventions:* random, highest degree, lowest E, highest sampled betweenness (100 sources) and whole Louvain communities, with ties broken at random; the boost applies to the position

(replacements are also boosted) or, in the robustness variant, only to the initial occupant. *Estimating the tipping point:* the share of runs with sustained collapse in each β cell is smoothed by isotonic regression and β_{50} is the interpolated crossing of 0.5; intervals are percentile bootstraps over runs within cells.

Declarations

Competing interests. The authors have no relevant financial or non-financial interests to disclose.

Funding. No funding was received to assist with the preparation of this manuscript.

Data availability. All results are generated by simulation; no empirical data were collected or used. The raw simulation outputs, random seeds and summary tables underlying every table and figure are archived in a public repository (Zenodo) under an open license. To preserve anonymity during double-anonymous review, the repository identifier is withheld here and the same files are supplied for review as Online Resource 1; the persistent identifier will be inserted in the accepted version.

Code availability. The model, experiment, analysis and test code (Python) and instructions to regenerate all results are archived in a public repository (Zenodo) under an open license. To preserve anonymity during double-anonymous review, the repository identifier is withheld here and the same files are supplied for review as Online Resource 1; the persistent identifier will be inserted in the accepted version.

Ethics approval and consent to participate. Not applicable. The study uses no human participants and no personal or organizational data.

Use of generative AI. Claude Sonnet 5.5 (Anthropic) was used to draft and debug the simulation code. The authors reviewed and edited the content as needed and take full responsibility for it.

References

Barabási A-L, Albert R (1999) Emergence of scaling in random networks. Science 286:509–512. https://doi.org/10.1126/science.286.5439.509

Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. J Stat Mech Theory Exp 2008:P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008

Davis JP, Eisenhardt KM, Bingham CB (2007) Developing theory through simulation methods. Acad Manag Rev 32:480–499

Erdős P, Rényi A (1959) On random graphs I. Publ Math Debrecen 6:290–297. https://doi.org/10.5486/PMD.1959.6.3-4.12

Feeley TH, Barnett GA (1997) Predicting employee turnover from communication networks. Hum Commun Res 23:370–387. https://doi.org/10.1111/j.1468-2958.1997.tb00401.x

Felps W, Mitchell TR, Hekman DR, Lee TW, Holtom BC, Harman WS (2009) Turnover contagion: How coworkers' job embeddedness and coworkers' job search behaviors influence quitting. Acad Manag J 52:545–561. https://doi.org/10.5465/AMJ.2009.41331075

Granovetter M (1978) Threshold models of collective behavior. Am J Sociol 83:1420–1443

Grimm V, Berger U, Bastiansen F et al (2006) A standard protocol for describing individual-based and agent-based models. Ecol Model 198:115–126. https://doi.org/10.1016/j.ecolmodel.2006.04.023

Grimm V, Berger U, DeAngelis DL, Polhill JG, Giske J, Railsback SF (2010) The ODD protocol: A review and first update. Ecol Model 221:2760–2768. https://doi.org/10.1016/j.ecolmodel.2010.08.019

Hagberg AA, Schult DA, Swart PJ (2008) Exploring network structure, dynamics, and function using NetworkX. In: Varoquaux G, Vaught T, Millman J (eds) Proceedings of the 7th Python in Science Conference (SciPy2008). Pasadena, CA, pp 11–15

Hancock JI, Allen DG, Bosco FA, McDaniel KR, Pierce CA (2013) Meta-analytic review of employee turnover as a predictor of firm performance. J Manag 39:573–603. https://doi.org/10.1177/0149206311424943

Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. Acad Manag Rev 32:1229–1245. https://doi.org/10.5465/amr.2007.26586485

Hausknecht JP, Holwerda JA (2013) When does employee turnover matter? Dynamic member configurations, productive capacity, and collective performance. Organ Sci 24:210–225. https://doi.org/10.1287/orsc.1110.0720

Hausknecht JP, Trevor CO (2011) Collective turnover at the group, unit, and organizational levels: Evidence, issues, and implications. J Manag 37:352–388

Heavey AL, Holwerda JA, Hausknecht JP (2013) Causes and consequences of collective turnover: A meta-analytic review. J Appl Psychol 98:412–453. https://doi.org/10.1037/a0032380

Herman J, Usher W (2017) SALib: An open-source Python library for sensitivity analysis. J Open Source Softw 2:97. https://doi.org/10.21105/joss.00097

Holland PW, Laskey KB, Leinhardt S (1983) Stochastic blockmodels: First steps. Soc Netw 5:109–137

Hom PW, Lee TW, Shaw JD, Hausknecht JP (2017) One hundred years of employee turnover theory and research. J Appl Psychol 102:530–545. https://doi.org/10.1037/apl0000103

Lee TW, Mitchell TR (1994) An alternative approach: The unfolding model of voluntary employee turnover. Acad Manag Rev 19:51–89. https://doi.org/10.5465/AMR.1994.9410122008

Lee TW, Mitchell TR, Sablynski CJ, Burton JP, Holtom BC (2004) The effects of job embeddedness on organizational citizenship, job performance, volitional absences, and voluntary turnover. Acad Manag J 47:711–722. https://doi.org/10.5465/20159613

Mitchell TR, Holtom BC, Lee TW, Sablynski CJ, Erez M (2001) Why people stay: Using job embeddedness to predict voluntary turnover. Acad Manag J 44:1102–1121. https://doi.org/10.5465/3069391

Rahmandad H, Sterman JD (2008) Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. Manag Sci 54:998–1014. https://doi.org/10.1287/mnsc.1070.0787

Rand W, Rust RT (2011) Agent-based modeling in marketing: Guidelines for rigor. Int J Res Mark 28:181–193. https://doi.org/10.1016/j.ijresmar.2011.04.002

Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M, Tarantola S (2010) Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. Comput Phys Commun 181:259–270. https://doi.org/10.1016/j.cpc.2009.09.018

Scheffer M, Bascompte J, Brock WA et al (2009) Early-warning signals for critical transitions. Nature 461:53–59. https://doi.org/10.1038/nature08227

Watts DJ (2002) A simple model of global cascades on random networks. Proc Natl Acad Sci USA 99:5766–5771. https://doi.org/10.1073/pnas.082090499

Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. Nature 393:440–442. https://doi.org/10.1038/30918

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.zip