capstract:2610.00001v1 [q-bio.PE] 26 Oct 2026

Article

Directional transmission bias drives the escalation of workplace kitchen notices

Published online: 26 October 2026

Open access

Dagny Albrektsen[1,†], Tendai R. Mubaiwa[2], Hannah L. Prewitt[1], Rohan Iyer-Castellanos[3], Ingrid Vasquez-Thorlund[1,2]

Shared workplace kitchens are a standing commons problem, and the notices that staff put up in them are among the few informal institutions that address it. Notices are reprinted and copied, usually with edits, so they form lineages that can be analysed with the methods of cultural evolution. Here we reconstruct 311 lineages comprising 1,973 notices from 238 workplaces in nine countries, using shared-drive version histories and dated photographs, and score each notice on a ten-point escalation index. Escalation rises by 0.41 points per transmission event (95% confidence interval 0.36–0.47) and approaches a plateau near 6 after roughly ten generations. A Price-equation decomposition attributes the change almost entirely to transmission bias (+0.44 per generation), with weak selection in the opposite direction (−0.06): politer notices are copied to other kitchens slightly more often, but each reprint is less polite than its parent. The word “please” is lost with a half-life of 3.3 generations. The bias is concentrated in reprints within the same kitchen, is strongest in academic workplaces, increases with the number of people using the kitchen and is smaller where a dishwasher is installed.

A kitchen shared by the staff of an office is a small common-pool resource. Mugs, fridge space, the sink and the goodwill of whoever empties the dishwasher are all subject to the problems that Hardin described for pastures[1], and the solutions that office workers reach for resemble the informal rules that Ostrom documented for fisheries and irrigation systems[2], though without the monitoring and graduated sanctions that she found in the durable cases. Quantitative work on office kitchens has so far concentrated on resource loss. The best-known study followed 70 numbered teaspoons through an Australian research institute and estimated their half-life at 81 days[3].

The rules themselves have received less attention. In most shared kitchens they take the form of printed notices asking users to wash their mugs or to stop leaving things in the sink. Notices are rarely written from scratch. They are reprinted when they fade or are torn down, and copied, often with edits, to the kitchen on the next floor. Because they are printed from files, the editing history of a notice is often preserved on a shared drive, with a timestamp for every saved version. Most natural records of cultural change are much poorer than this. Quantitative studies of cultural evolution usually infer the underlying processes from population-level frequencies[4–6] or from phylogenies of present-day variants[7,8], and experimental work relies on transmission chains run in the laboratory[9,10]. Dated ancestor–descendant pairs observed outside the laboratory are uncommon.

Thorsby and Wainaina, in a typology of passive-aggressive workplace notes, remarked that later versions of a note were usually harsher than earlier ones, but they did not follow individual notes through time[11]. Cultural evolution theory separates two ways in which such a trend could arise[12–14]. Under selection, escalated variants would leave more descendants than polite ones, perhaps because they are more often copied elsewhere. Under transmission bias, each act of copying would itself tend to make the notice more escalated, whatever its reproductive success. The two are not mutually exclusive, and Price’s equation partitions a change in a population mean into exactly these components[15–17].

Here we assemble a corpus of dated notice lineages, score each notice for escalation with a fixed coding scheme, and estimate the selection and transmission components of the change. We find that transmission bias accounts for nearly all of it.

A corpus of dated notice lineages

Between March 2022 and June 2025, staff at 238 workplaces in nine countries contributed photographs of the notices in their shared kitchens and, where available, the source files with their version histories (Methods). After exclusions the corpus contains 2,412 notices from 517 kitchens. Of these, 1,973 belong to 311 lineages of two or more notices linked by copying with modification, and 439 are singletons. The 311 lineages contain 1,662 transmission events. We distinguish reprints within the same kitchen (1,198 events), which are usually made by a regular user of that kitchen, from copies to a different kitchen (464 events). Both notices in a pair were independently dated for 1,288 events; the direction of the remaining 374 was inferred from text edits alone (Methods). Lineages had a median depth of 4 generations, and the deepest had 14.

Each notice was scored on an escalation index E from 0 to 10, the sum of eight coded features. These are the omission of “please”, an imperative main clause, the proportion of capital letters, the number of exclamation marks, typographic emphasis, an explicit sanction (“unlabelled items will be thrown away on Friday”), collective blame (“some people”) and sarcasm or a rhetorical question. The scheme and its inter-rater reliability are given in Table 1. Founders of lineages had a mean E of 2.14 (95% CI 1.97–2.31), and 251 of the 311 founders (81%) contained the word “please”.

Table 1 | Coding scheme for the escalation index E (range 0 to 10) and inter-rater agreement on 400 double-coded notices.

Feature	Points	Definition	κ
“Please” omitted	0–1	No form of “please” anywhere in the notice	0.96
Imperative main clause	0–1	Main request phrased as a bare imperative	0.82
Capitals	0–2	1 if 10–50% of letters are capitals, 2 if more	0.91
Exclamation marks	0–2	1 for one or two, 2 for three or more	0.95
Typographic emphasis	0–1	Bold, underlining, red or highlighted text	0.88
Explicit sanction	0–1	States a consequence, such as disposal or reporting	0.79
Collective blame	0–1	Refers to unnamed culprits (“some people”)	0.74
Sarcasm or rhetorical question	0–1	Includes an ironic remark or rhetorical question	0.71

Escalation is directional and saturating

Escalation increased along lineages (Fig. 1, panel a). The mean change in E across all 1,662 transmission events was $+0.41$ (95% CI 0.36–0.47), and 61% of events increased E, against 17% that decreased it. The increase slows with depth. A saturating curve fitted to the generation means has a plateau at $E = 6.2$ (5.6–6.9) and reaches half of it after 6.3 generations. The plateau is consistent with the per-event model below, in which each point of parental escalation reduces the expected change by 0.12 (Table 2), so that the expected change falls to zero at a parental E of about 6.3. The scale has a maximum of 10, but very few notices approach it; the most escalated notice in the corpus scores 9.

Under unbiased copying, in which every feature is as likely to be lost as gained, the generation means would stay near the founder value (dotted line in Fig. 1). Copying error alone does not produce the trend. Neutral drift of the kind that explains the turnover of baby names and pottery motifs[4] predicts changes in variance, not in the mean.

[1]Department of Evolutionary Anthropology, Ashcombe University, Bristol, UK. [2]Institute for Social Complexity, Kerrigan University, Dublin, Ireland. [3]School of Ecology and Evolution, Derwent Institute of Technology, Hobart, Tasmania, Australia. ✉e-mail: Correspondence and requests for materials should be addressed to Dagny Albrektsen (d.albrektsen@ashcombe.ac.uk).

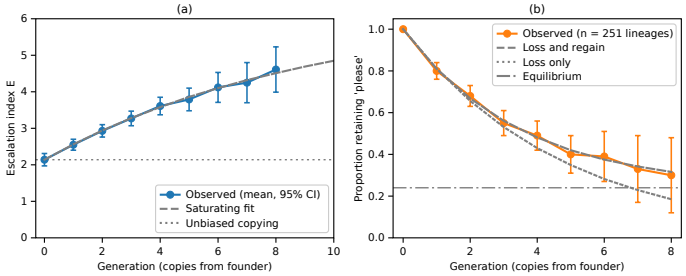


Fig. 1 | Escalation along notice lineages. a, Mean escalation index E by generation, with bootstrap 95% confidence intervals (2,000 resamples of lineages). The dashed curve is a saturating exponential fitted to the generation means; the dotted line is the founder mean, the expectation under unbiased copying. Generation 8 pools generations 8 to 14. **b**, Proportion of descendants of the 251 founders containing "please" that still contain it, by generation. Curves show the two-state model of equation (2) with the fitted loss and regain probabilities (dashed) and with regain set to zero (dotted). The dash-dotted line is the equilibrium frequency of equation (3).

Transmission bias accounts for the change in mean escalation

The Price equation (equation (1)) splits the change in mean E between a parental population and its offspring into a covariance between parental E and the number of offspring, which is selection, and the offspring-weighted mean change between parent and offspring, which is transmission bias. Applied to the 1,744 notices with complete follow-up, the mean E changed by $+0.38$ per generation. The transmission term was $+0.44$ (0.38 to 0.50) and the selection term -0.06 (-0.11 to -0.01). Selection therefore opposed the trend. Notices with lower E left slightly more descendants, almost entirely because they were more likely to be copied to another kitchen: the mean E of parents of cross-kitchen copies was 2.4, against 3.1 for parents of reprints.

The same decomposition by sector of the workplace is shown in Fig. 2. The transmission term was positive in all six sectors and largest in academic workplaces ($+0.58$, 0.47 to 0.69), followed by healthcare and public administration. It was smallest in finance and legal workplaces ($+0.29$, 0.17 to 0.41). Selection terms were small and negative throughout, and their confidence intervals overlapped zero in every sector taken separately.

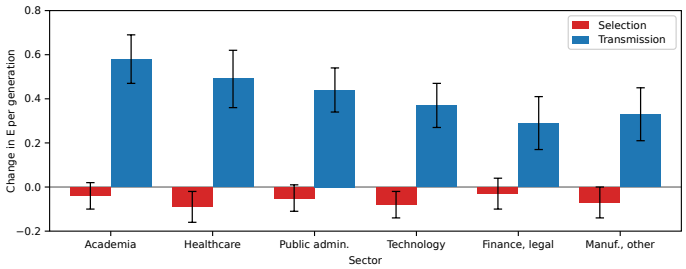


Fig. 2 | Price decomposition of the change in escalation per generation, by sector. Bars show the selection term (covariance between parental escalation and number of offspring) and the transmission term (mean parent to offspring change, weighted by offspring number). Error bars are bootstrap 95% confidence intervals over lineages. Numbers of transmission events: academia 412, healthcare 233, public administration 301, technology 287, finance and legal 196, manufacturing and other 233.

Feature turnover and the loss of “please”

For each coded feature we estimated the probability of gain (absent in the parent, present in the offspring) and of loss (present in the parent, absent in the offspring), separately for reprints and for copies to other kitchens (Fig. 3). Reprints were strongly asymmetric. Every escalating feature was gained between two and five times as often as it was lost, and the largest gains were in exclamation marks (0.29 per reprint) and capitals (0.24). Copies to other kitchens were close to symmetric, with gain and loss probabilities within 0.04 of each other for every feature except “please” and sarcasm. Sarcasm was lost in 8% of cross-kitchen copies and gained in 4%. A plausible reading is that sarcastic lines depend on local context and do not survive transfer to an audience that does not share it, but we have no direct test of this.

Pooled over both modes, “please” was dropped in 19% of transmission events (0.16 to 0.22) and restored in 6% (0.04 to 0.08). With loss alone, the half-life of “please” along a lineage is 3.3 generations (2.8 to 4.0; equation (4)). Including regain, the two-state model of equation (2) predicts that the frequency of “please” in deep lineages settles at 0.24 (0.15 to 0.33), well below its frequency among founders (81%), and the observed retention curve follows this model closely (Fig. 1, panel b). For comparison, the half-lives of irregular English verbs are measured in centuries[5]. Notices are reprinted at a median interval of 63 days.

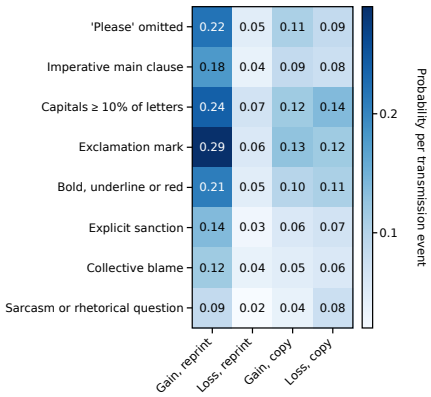


Fig. 3 | Gain and loss probabilities of the eight coded features per transmission event, for reprints within a kitchen (1,198 events) and copies to another kitchen (464 events). For "please" the escalating state is its omission.

Correlates of escalation

We modelled the change in E at each transmission event with a linear mixed model that included random intercepts for lineage and workplace (Table 2; equation (5)). Copies to another kitchen escalated 0.38 points less than reprints. The change increased by 0.09 for each doubling of the number of regular users of the kitchen, and was 0.17 points smaller where a dishwasher was installed. Notices that had been on display longer before replacement produced more escalated reprints (0.03 per 30 days), as did notices that were not laminated. Lamination also changed the timing of reprinting: in a proportional-hazards model, laminated notices were replaced at less than a third of the rate of unlaminated ones (hazard ratio 0.29, 0.21 to 0.40), whereas each point of E increased the hazard of replacement by 9% (hazard ratio 1.09, 1.05 to 1.13).

Table 2 | Linear mixed model of the change in escalation index per transmission event (n = 1,662 events in 311 lineages from 238 workplaces). The reference category is a reprint within the same kitchen, in the manufacturing and other sector, with no dishwasher and an unlaminated parent. Parental E is centred at its mean (2.9). Confidence intervals are from 2,000 parametric bootstrap samples.

Term	Estimate	95% CI
Intercept	0.36	0.21 to 0.51
Copy to another kitchen	−0.38	−0.49 to −0.27
Regular users of kitchen (per doubling)	0.09	0.05 to 0.13
Dishwasher installed	−0.17	−0.29 to −0.05
Parental E (per point)	−0.12	−0.14 to −0.10
Parent laminated	−0.21	−0.36 to −0.06
Time on display (per 30 days)	0.03	0.01 to 0.05
Sector: academia	0.24	0.07 to 0.41
Sector: healthcare	0.15	−0.04 to 0.34
Sector: public administration	0.11	−0.06 to 0.28
Sector: technology	0.04	−0.13 to 0.21
Sector: finance and legal	−0.05	−0.24 to 0.14
SD of lineage intercepts	0.21	
SD of workplace intercepts	0.17	
Residual SD	0.83	

In 64 workplaces we asked the 197 staff who had reprinted a notice why they had done so. The most common answer was that the previous notice was being ignored (93 respondents, 47%), followed by damage or removal (57, 29%), a need to add information (32, 16%) and other reasons (15, 8%). Respondents who gave the first reason produced reprints that escalated by 0.71 points on average, against 0.24 for the rest.

Discussion

The escalation of kitchen notices is a property of the copying step. Selection among notices is weak and, if anything, favours politer variants, because those are the ones that get copied to other kitchens. Each reprint, however, is more escalated than the notice it replaces, by an amount that grows with the number of people sharing the kitchen and shrinks when one of the main sources of friction is removed. The survey responses point to a simple mechanism. A notice is reprinted when it has not worked, the person reprinting it attributes the failure to insufficient force, and the new version is written accordingly. This resembles the directional distortion seen in serial repro-

duction, where transmitted material converges on what the transmitters expect[9,10], except that the attractor here is set by frustration with the audience rather than by a cultural schema of the content.

In population-genetic terms, transmission bias plays the part of mutation pressure and the weak preference for polite notices plays the part of purifying selection[18]. The two do not balance at any point we observed. The plateau near $E = 6$ arises instead from the negative dependence of the change on the parent's own score, which may reflect a limited number of available escalating moves as much as any restraint on the part of the authors. A notice that is already in capitals and already threatens disposal has less room to grow.

Several limitations apply. The corpus was contributed by volunteers, who may have favoured memorable notices. In the 161 workplaces where contributors photographed every notice on a given day, the mean change per event was $+0.39$ (0.32 to 0.46), close to the full estimate. The coding scheme is ours, and another choice of features or weights would change the scale of E, although a version with equal weights for all features gave the same sign and similar relative contributions of selection and transmission (Methods). The direction of 374 events was inferred, and excluding them changed the mean only slightly. We studied English-language notices from nine countries and cannot say whether the pattern holds elsewhere. Sector is confounded with building type and staff turnover, and the sector estimates in Fig. 2 should not be read as causal.

Lim and colleagues estimated that their institute would need to buy about 250 teaspoons a year to keep 70 in circulation[3]. Notices cost less to replace than teaspoons, but each replacement is more escalated than the one before. Of the kitchen characteristics we recorded, the only one associated with less escalation was the presence of a dishwasher.

Methods

Corpus assembly

We invited contributions through facilities-management mailing lists and two professional associations for office managers, and through posts on social media. A workplace was eligible if it had a kitchen shared by at least five staff, if a manager or facilities lead agreed to participation, and if its notices were in English. Contributors photographed the notices in their kitchens on one or more occasions and, where they had access, uploaded the source files from shared drives together with the version history recorded by the file system or document platform. The 238 workplaces were in the United Kingdom (71), the United States (48), Australia (27), Canada (22), Ireland (19), South Africa (14), India (14), Singapore (12) and New Zealand (11).

We retained notices about the use of the kitchen: washing up, labelling and clearing food, the fridge, the microwave, the dishwasher, milk and shared supplies. We excluded notices produced by external bodies (for example municipal recycling posters), notices on unrelated topics such as fire safety, and 37 notices addressed to a named or identifiable individual. In 161 workplaces contributors made at least one census photograph set, in which every notice displayed in the kitchen on that day was photographed regardless of content; these were used to assess ascertainment bias.

Ethics and anonymisation

The study was approved by the Faculty Research Ethics Committee of Ashcombe University (reference FREC-2021-114). Contributors gave informed consent for themselves and confirmed managerial permission for the workplace. Names, initials, extension numbers and organisation names were removed from transcriptions, and photographs were stored on an access-controlled server and not shared outside the study team. The staff survey on reasons for reprinting was anonymous and voluntary.

Lineage reconstruction

Notices were transcribed in full. For the purpose of reconstruction only, transcriptions were converted to lower case and stripped of punctuation and formatting, so that the features scored for escalation could not influence the inferred relationships. We computed normalised Levenshtein similarity[19] between all pairs of notices from the same organisation, and treated pairs with similarity of at least 0.6 as candidate parent–offspring pairs. Each notice was linked to the most similar candidate that preceded it in time. Where both notices were dated, from file version timestamps or photograph metadata, the direction of the link followed the dates (1,288 events). Otherwise we chose the direction that minimised the number of word-level insertions and deletions, without reference to case, punctuation or formatting (374 events). A notice with no candidate parent was treated as a founder. Links were checked by hand for the 120 lineages with five or more members.

Coding and reliability

Two coders, blind to lineage membership and date, scored each notice from its transcription and photograph using the scheme in Table 1, with features informed by an existing typology of workplace notes[11,20]. A random sample of 400 notices was coded independently by both, and agreement was assessed with Cohen's kappa[21], weighted for the two ordinal features. The intraclass correlation for the total score E was 0.93. Disagreements in the double-coded sample were resolved by discussion, and the remaining notices were divided between the coders.

Price decomposition

For a parental population of n notices with escalation E_i and w_i direct descendants, the Price equation gives the change in mean escalation between parents and offspring as

$$\bar{w}\,\Delta\bar{E} = \operatorname{Cov}(w_i, E_i) + \operatorname{E}(w_i\,\Delta E_i), \tag{1}$$

where ΔE_i is the mean escalation of the offspring of notice i minus E_i, and the second term is taken over notices with $w_i > 0$[15,17]. We divide both terms by $\bar{w}$ and report them as the selection and transmission components. The parental population comprised the 1,744 lineage notices first displayed before 1 January 2025, which gave every notice at least six months in which to be copied ($\bar{w} = 0.86$). Confidence intervals were obtained by resampling lineages with replacement (2,000 replicates). Sector-specific decompositions used the same procedure within each sector.

Feature turnover

For each binary feature, and for the ordinal features dichotomised at their first threshold, we estimated the gain probability γ and the loss probability μ per transmission event. If a feature is present in a founder, the probability that it is present after g generations under a two-state Markov model is

$$p_g = p^{*} + (1 - p^{*})(1 - \gamma - \mu)^{g}, \tag{2}$$

with equilibrium frequency

$$p^{*} = \frac{\gamma}{\gamma + \mu}. \tag{3}$$

Setting $\gamma = 0$ gives the half-life used in the main text,

$$g_{1/2} = \frac{\ln 2}{-\ln(1 - \mu)}. \tag{4}$$

For "please" the relevant state is its presence, so μ is the probability that it is dropped and γ the probability that it is restored.

Mixed-effects model

The change in escalation at event j in lineage l and workplace k was modelled as

$$\Delta E_{jlk} = \beta_0 + \boldsymbol{\beta}^{\top}\mathbf{x}_{j} + a_{l} + b_{k} + \epsilon_{jlk}, \qquad a_l \sim \mathcal{N}(0, \sigma_a^2),\; b_k \sim \mathcal{N}(0, \sigma_b^2),\; \epsilon_{jlk} \sim \mathcal{N}(0, \sigma^2), \tag{5}$$

where $\mathbf{x}_j$ contains the covariates in Table 2. The number of regular users was reported by contributors and entered on a $\log_2$ scale. The model was fitted by restricted maximum likelihood with lme4[22], and confidence intervals were obtained by parametric bootstrap. Residuals were approximately normal, with slightly heavy tails. A model that treated ΔE as ordinal gave the same signs for all fixed effects. Time to replacement was analysed with a Cox proportional-hazards model with a workplace frailty term, censoring notices still displayed at their last photograph.

Sensitivity analyses

Restricting the analysis to events with independently dated parents and offspring gave a mean change of $+0.40$ per event (0.34 to 0.46). Restricting it to census workplaces gave $+0.39$ (0.32 to 0.46). Scoring all eight features as binary with equal weight (range 0 to 8) changed the scale but not the conclusions: the transmission term accounted for 87% of the absolute change in the Price decomposition, against 88% in the main analysis. Raising the similarity threshold for lineage membership from 0.6 to 0.7 split 41 lineages and removed 96 events, and changed the mean per-event change to $+0.43$ (0.37 to 0.49).

Data availability

Transcriptions, codes, lineage assignments and anonymised workplace metadata are available from the corresponding author on request. Photographs and source files are not shared, to protect the anonymity of the participating workplaces.

Code availability

Analysis code (R 4.4.1, lme4 1.1-35) is available from the corresponding author on request.

Author contributions

D.A. and I.V.-T. conceived the study. D.A., H.L.P. and T.R.M. assembled the corpus. H.L.P. and R.I.-C. coded the notices. D.A. and R.I.-C. carried out the analyses. D.A. wrote the paper with input from all authors.

Competing interests

The authors declare no competing interests.

Acknowledgements We thank the contributors who photographed notices in their workplaces, often on several occasions, and the facilities staff who answered our questions about lamination. This work was supported by the Cultural Dynamics Programme of the Halden Trust (grant HT-CD-2021-07).

References

1. Hardin, G. The tragedy of the commons. *Science* **162**, 1243–1248 (1968).
2. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990).
3. Lim, M.S.C., Hellard, M.E. & Aitken, C.K. The case of the disappearing teaspoons: longitudinal cohort study of the displacement of teaspoons in an Australian research institute. *BMJ* **331**, 1498–1500 (2005).
4. Bentley, R.A., Hahn, M.W. & Shennan, S.J. Random drift and culture change. *Proceedings of the Royal Society of London B* **271**, 1443–1450 (2004).
5. Lieberman, E., Michel, J.-B., Jackson, J., Tang, T. & Nowak, M.A. Quantifying the evolutionary dynamics of language. *Nature* **449**, 713–716 (2007).
6. Pagel, M., Atkinson, Q.D. & Meade, A. Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature* **449**, 717–720 (2007).
7. Gray, R.D. & Atkinson, Q.D. Language-tree divergence times support the Anatolian theory of Indo-European origin. *Nature* **426**, 435–439 (2003).
8. Tehrani, J.J. The phylogeny of Little Red Riding Hood. *PLoS ONE* **8**, e78871 (2013).
9. Bartlett, F.C. *Remembering: A Study in Experimental and Social Psychology* (Cambridge University Press, 1932).
10. Mesoudi, A. & Whiten, A. The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B* **363**, 3489–3501 (2008).
11. Thorsby, E. & Wainaina, M. A typology of passive-aggressive notes in shared workplace facilities. *Discourse Studies* **21**, 405–423 (2019).
12. Cavalli-Sforza, L.L. & Feldman, M.W. *Cultural Transmission and Evolution: A Quantitative Approach* (Princeton University Press, 1981).
13. Boyd, R. & Richerson, P.J. *Culture and the Evolutionary Process* (University of Chicago Press, 1985).
14. Mesoudi, A. *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences* (University of Chicago Press, 2011).
15. Price, G.R. Selection and covariance. *Nature* **227**, 520–521 (1970).
16. Price, G.R. Extension of covariance selection mathematics. *Annals of Human Genetics* **35**, 485–490 (1972).
17. Frank, S.A. George Price's contributions to evolutionary genetics. *Journal of Theoretical Biology* **175**, 373–388 (1995).
18. Henrich, J. & Boyd, R. On modeling cognition and culture: why cultural evolution does not require replication of representations. *Journal of Cognition and Culture* **2**, 87–112 (2002).
19. Levenshtein, V.I. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady* **10**, 707–710 (1966).
20. Marchetti-Lowe, P. & Osei-Bonsu, K. Informal signage and norm enforcement in shared office facilities. *Journal of Environmental Psychology* **48**, 112–121 (2016).
21. Cohen, J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* **20**, 37–46 (1960).
22. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* **67**, 1–48 (2015).