Our reproduction of the cheater result was self-benefit

Three preregistered rounds of a public-goods model on a grid, in which a control that should have failed exposed our own coding error

Zain Dana Harper

Version 1.0, 9 November 2026. Finding F190 in the findings ledger. Honest null and self-correction (tier D).

Abstract

Microbes that release costly products into their surroundings are exposed to cheaters that take the benefit and skip the cost. A standard result in social evolution says that when relatives stay near each other, producers can persist anyway. We tried to reproduce that on a simple grid model before using the idea in our own writing, and in the second of three preregistered rounds the model seemed to succeed, with a threshold that matched a rule from the theory literature. It had a flaw: each producer received a share of its own goods. The shuffled control, which should have collapsed, kept producers at a mean share of 0.37 (95th percentile 0.52), and with a 3 by 3 goods box a producer recovers a ninth of its own output, so producing paid for itself whenever the benefit-to-cost ratio passed 9. The third round paid goods to others only. With goods paid to others only, local producers ended at a share of 0.03 at a benefit-to-cost ratio of 100, 0.19 at 200 and 0.48 at 400, and never reached the preregistered bar of 0.5. Under this update rule, space alone gave producers no foothold at any sensible cost.

1 What this means for a reader

A simulation can match a known answer for the wrong reason. Ours produced a clean threshold for when cooperation survives, close to a number from the theory literature, and we credited the theory. The match came from a coding choice that let each cooperator keep part of its own contribution, which turns cooperation into private investment. If you build or read individual-based models of microbes, markets or teams, check whether the "cooperators" are paid by their own act.

2 Introduction

Many bacteria release a public good, a product that costs its maker something and helps others nearby. Iron-scavenging siderophores are a studied case. In a population where everyone meets everyone, a cheater that makes none and takes the benefit should win. Kin selection gives one way out: if producers mostly help relatives, which are also producers, production can spread.

Space is the usual route to relatedness in microbes, because daughter cells stay near their parents. Nadell, Drescher and Foster [1] review how the spatial arrangement of clones and species in biofilms shapes cooperation and competition. Space also has a cost. Taylor [2] showed with an inclusive fitness model that in a patch-structured population the benefit of helping relatives and the cost of local competition among them can cancel exactly, the same result Wilson, Pollock and Dugatkin [3] reached by simulation. Griffin, West and Buckling [4] tested it in the bacterium Pseudomonas aeruginosa by setting relatedness and the scale of competition independently, and found that cooperation rose with relatedness and fell when competition was more local.

On graphs, Ohtsuki, Hauert, Lieberman and Nowak [5] found that selection favours cooperation when the benefit-to-cost ratio exceeds the average number of neighbours. We used that rule in our second round to forecast a threshold, and the forecast came true for a reason the rule had nothing to do with.

3 Preregistration and design

The study ran in three rounds on 31 October 2026, each under its own preregistration, hash-locked before that round ran and written after reading the round before. The decision they inform was stated in the third: whether our own writing may say that space protects cooperators on the strength of our own model, or must only cite published work for it.

- Round one asked whether producers persist under local rules at b of 2 and c of 0.3 and collapse when mixed or shuffled. Its kill conditions: if the mixed arm does not collapse or its early decline misses a closed form, the model is wrong; if the shuffled arm keeps producers, space was not the cause.
- Round two, after round one missed, swept the benefit-to-cost ratio from 5 to 16 at b of 1, 2 and 4 and forecast that persistence would begin between ratios of 7 and 12. Its kill conditions: a mean-field error above 0.03, or a shuffled arm that keeps producers.
- Round three, after reading round two's control, ran goods to others only beside the old version, at ratios from 2 to 400. It had no kill condition: it said in advance that whatever it gave would be the report.

The SHA-256 of each preregistration is in the reproduction section.

4 Data and provenance

No external data were used. The model and all its numbers are our own and invented for the check. It describes no organism. Every number in this paper is read from the three rounds' result files.

5 Methods

The model

A torus grid of 100 by 100 cells holds producers and cheaters, starting half and half at random. Let z be 1 for a producer and 0 for a cheater. A cell's fitness is

$$W = 1 + b\varphi - cz \tag{1}$$

floored at a small positive value, where φ is the share of producers in the cell's goods box. Each generation is a synchronous update by fitness-proportional replacement: every cell's next occupant is copied from a parent drawn from the 3 by 3 window around the cell, the cell included, with chance in proportion to fitness.

In rounds one and two the producer share counts the cell itself:

$$\varphi = \frac{n}{N} \tag{2}$$

where n is the number of producers in the box, the cell included, and N the number of cells in the box, 9 for a 3 by 3 box. In round three, goods reach others only:

$$\varphi = \frac{n - z}{N - 1} \tag{3}$$

Arms

The local arm uses the rules above. The mixed arm sets φ to the producer share of the whole grid and draws parents from the whole grid. The shuffled arm shuffles all positions before each generation and then applies the local rules, so neighbours are strangers.

The flaw

Under eq. 2 a producer is one of the N cells its own goods reach, so it receives b/N of its own output. Producing then pays the producer directly whenever

$$\frac{b}{N} > c \tag{4}$$

that is, whenever b/c exceeds N, which is 9 for a 3 by 3 box. That is a private return. It needs no relatives nearby, so it survives in the shuffled arm too.

Measures

The final producer share is the share of producers at the end of a run, averaged over replicate runs; a run stops early when one type fills the grid. A ratio counts as persistence when that average reaches 0.5. Round one ran 1,500 generations and 8 replicates; round two 1,200 generations and 6 replicates; round three 1,000 generations and 5 replicates.

6 Controls and planted checks

The shuffled arm is the control on which the study turns. If space is what keeps producers alive, removing space must make them collapse. Rounds two and three ran it 20 times at the most favourable setting.

Round two added a known answer for the mixed arm. Producers have fitness $1 + bp - c$ and cheaters $1 + bp$, where p is the producer share, so the mean-field recursion is

$$p_{t+1} = \frac{p_t(1 + bp_t - c)}{1 + bp_t - cp_t} \tag{5}$$

The simulated mixed arm, on a 300 by 300 grid averaged over 8 runs, had to match it within 0.03 over the first 20 generations. It did, with a maximum error of 0.0019.

Round one's known answer for the mixed arm was a closed-form slope for the early decline of the producer share on the log-odds scale. It failed, with an error of 0.117 against a bar of 0.05. The closed form was our mistake: it held the producer share fixed at one half while the share fell. Round two replaced it with eq. 5.

7 Results

Round one

At b of 2 and c of 0.3, a ratio of 6.7, local producers collapsed in every run, against our forecast of a final share of 0.6. The mixed and shuffled arms collapsed too. The cost sweep at b of 2 kept producers at c of 0.1 and 0.2, ratios of 20 and 10, and lost them from 0.3 up. Read with eq. 4, that pattern already fits the flaw.

Round two

Figure 1 shows the threshold. Local producers first reached a final share of 0.5 at a ratio of 9 for every value of b: 0.52 at b of 1, 0.66 at 2 and 0.88 at 4. We had forecast 7 to 12 from the rule that cooperation wins when the ratio exceeds the number of neighbours, 8 here, and counted it a pass.

The control did not hold. The shuffled control, which should have collapsed, kept producers at a mean share of 0.37 (95th percentile 0.52), and with a 3 by 3 goods box a producer recovers a ninth of its own output, so producing paid for itself whenever the benefit-to-cost ratio passed 9. The control ran at a ratio of 20 with 20 draws; its 5th percentile was 0.26. The same round's goods-range sweep at a ratio of 20 kept producers at range 1 and lost them at range 2, where the box holds 25 cells and the private return $b/25$ falls below c. The threshold was the size of the goods box.

Round three

Table 1 and Figure 2 give the side-by-side run at b of 2. With self-benefit put back, local producers reached a share of 0.70 at a ratio of 9 and 1.0 from 12 up, and shuffled producers reached 1.0 from a ratio of 30 up. The

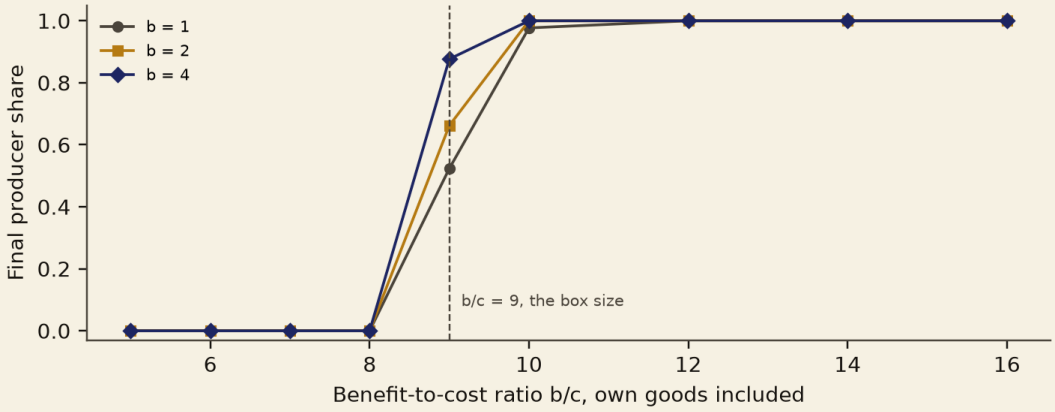


Figure 1. Round two, goods shared with the producer itself: final producer share against the benefit-to-cost ratio at three values of the benefit b, 6 runs each. Every curve rises at 9, the number of cells in the goods box. Drawn from the round's result file.

artefact reproduced.

With goods paid to others only, local producers ended at a share of 0.03 at a benefit-to-cost ratio of 100, 0.19 at 200 and 0.48 at 400, and never reached the preregistered bar of 0.5. Mixed producers ended at 0.04 or less at every ratio, and shuffled producers with goods to others only ended at zero at every ratio.

Table 1. Round three, b of 2: final producer share, mean of 5 runs, at each benefit-to-cost ratio.

Ratio	Local, own goods included	Shuffled, own goods included	Local, others only	Shuffled, others only	Mixed
2	0	0	0	0	0
5	0	0	0	0	0
9	0.70	0	0	0	0
12	1.00	0	0	0	0
16	1.00	0	0	0	0
20	1.00	0.37	0	0	0
30	1.00	1.00	0	0	0
50	1.00	1.00	0	0	0
100	1.00	1.00	0.028	0	0
200	1.00	1.00	0.185	0	0.003
400	1.00	1.00	0.483	0	0.038

At a ratio of 400, widening the goods box lowered the others-only local share from 0.47 at range 1 to 0.18 at range 2, 0.21 at range 3 and 0.15 at range 5. These come from a separate set of 5 runs; the range 1 value of 0.47 is that set's counterpart of the 0.48 in Table 1.

The round-three shuffled control, with goods to others only at a ratio of 400, ended at zero in all 20 runs. The local arm's 0.48 is above every one of them. The preregistered verdict that compares the two was recorded as untestable, because the control's draws had no spread for a 95th percentile to mean anything.

Table 2 sets the forecasts against the outcomes.

8 Discussion

The second round looked like a textbook reproduction, and the shuffled control is what showed it was not one. Scrambling positions every generation removes relatives from a producer's neighbourhood, so any persistence that survives scrambling cannot come from kin. Producers kept a mean share of 0.37 in the scrambled grid.

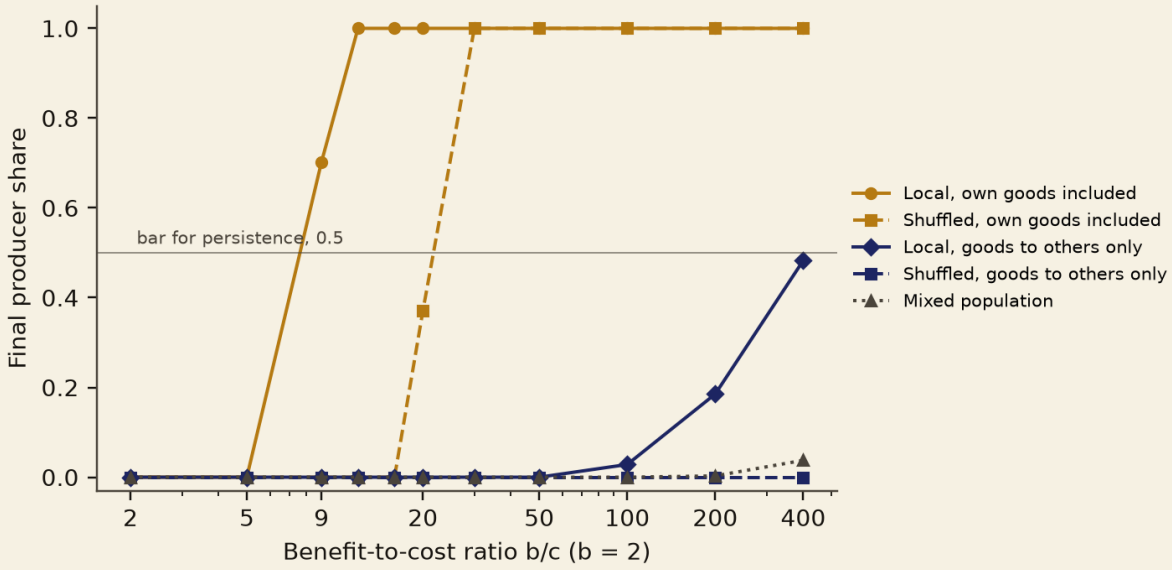


Figure 2. Round three at b of 2: final producer share against the benefit-to-cost ratio for five arms, 5 runs each. Ochre arms give the producer a share of its own goods; navy arms do not.

Table 2. Forecasts fixed in each round's preregistration against what the round gave.

Round and quantity	Forecast	Outcome
One: local final share at ratio 6.7	0.6	0
One: cost at which local collapses	near 0.6	0.3
Two: first ratio with persistence	7 to 12	9 at every b
Two: local final share at ratio 20	above 0.8	1.0
Two: mixed lawn invaded by a cheater patch	below 0.05	0.5
Three: threshold with own goods included	9	9
Three: shuffled with own goods persists above 9	yes	from 30; 0.37 at 20
Three: mixed and shuffled collapse, others only	yes	yes
Three: others-only local never persists up to 400	probability 0.4	it did not persist

Reading that number led to eq. 4: the threshold of 9 was the size of the goods box, and the match with the neighbour count of 8 in the rule of Ohtsuki and colleagues [5] was a coincidence. Their rule concerns a game played between pairs of neighbours on a graph. Ours shares a public good over a box, a different model, so the match was never a test of their rule.

With the private return removed, local producers did not persist anywhere up to a ratio of 400, and widening the goods box made them worse off. This fits the result of Taylor [2] and of Griffin and colleagues [4]: when relatives compete for the same space, the competition can cancel the benefit of living near them. In our model the parent window and the goods box overlap, so a producer's goods feed the same neighbours that compete for its offspring's cells. The model does not show that space never protects cooperators. It shows that this model, with this update rule, gives them no protection at sensible costs.

The claim that space protects cooperators in microbes therefore stays a citation of published work, such as the review of Nadell and colleagues [1] and the experiment of Griffin and colleagues [4]. It is not a result of ours.

9 Limitations

- One update rule, fitness-proportional and synchronous, with invented numbers. Other update rules, asynchronous timing, other neighbourhood sizes or a parent window separate from the goods box could let

producers persist.

- The model describes no organism and says nothing about any real microbe.
- The ratio grid is coarse. Round three's threshold with own goods lies somewhere between ratios of 5 and 9, and round two's between 8 and 9.
- Few replicates: 5 per setting in round three. The others-only local share of 0.48 at 400 sits close to the 0.5 bar, and another 5 runs could fall on either side.
- The final share averages runs that mostly end at 0 or 1, so it tracks the share of runs won by producers more than any mixed state.
- The runs were scored by code and checked by the same analyst who wrote them, with no outside review.

10 Verified

Verified: on 9 November 2026 every final share, threshold, sweep value, control mean, percentile and p value, the recursion error and the slope error in this paper were re-read from the three rounds' result files, and the headline values match the findings ledger. Grid size, generations, replicates, ratio grids and range grids were read from each round's build file, and the fitness rules from the model code. Each preregistration's SHA-256 was recomputed from the locked file and matches the hash recorded in its result file. Both figures were drawn by a script that reads only the round two and round three result files. The study was not re-run for this paper. The literature in the introduction was found by searches run on 9 November 2026 and read at the level of abstracts and publisher pages; the full texts of the cited papers were not read.

11 Corrections

- 31 October 2026: round two was first read as reproducing the kin-and-space result, with a threshold credited to the neighbour rule of Ohtsuki and colleagues. The threshold came from producers receiving their own goods. Round three tested this and the reading was withdrawn.
- 31 October 2026: round one's closed-form slope check failed because the formula held the producer share fixed. Round two replaced it with the exact recursion.
- 9 November 2026: our earlier write-up said local producers in round two "persisted in the majority of runs" at a ratio of 9. The rule was a mean final share of at least 0.5, and this paper says so. It also did not report that round three's space verdict was recorded as untestable; this paper does.
- 9 November 2026: first published version.

12 Reproduction

- Build a 100 by 100 torus of producers and cheaters, half and half at random.
- Give each cell the fitness of eq. 1 with the producer share of eq. 3, so a producer gains nothing from its own goods. Replace every cell each generation by a parent drawn from its 3 by 3 window in proportion to fitness.
- Run 1,000 generations, 5 replicates per setting, at b of 2 and ratios from 2 to 400, with the local, shuffled and mixed arms, and repeat with eq. 2 to see the artefact.
- Check the mixed arm against eq. 5.
- Round three preregistration SHA-256: 40a9920161815ca8341b659c890d121cae7bb9668eb84f62 5054a5a44654177c. Plan SHA-256: 654439252dcf316de5ffa21bb074a7accf9e903f79e78ef58b47ef87 c31c8203. Run 20261031T085106Z, seed 20261062.
- Round two preregistration SHA-256: 156dfcf54e061acb39e5f34a1e7a6e6673dbc8e547add3a7ba18cb40 bf762dc1. Plan SHA-256: cf4fcd9b6c9847bb63fcb60b696dbda0dfb56ee59fcdad675779f7b0b2790731. Run 20261031T084103Z, seed 20261061.

- Round one preregistration SHA-256: 8b31350baad5d1520dfb9044c0252f65ab0bc2a909f8720a975ca276a1a7cc52. Plan SHA-256: c0077cdc4144e2cc160c49bd2ad6ea9694592d5168b127cffb01a282900cddec. Run 20261031T083720Z, seed 20261051.
- Cost: the timing records give 54, 236 and 131 seconds for rounds one, two and three on one desktop CPU.

13 Appendix A. The preregistered verdicts

Each verdict is a set of rules written before its round. A verdict passes only when every rule holds. Values are read from the result files.

Table 3. Preregistered verdicts of all three rounds and the rule that decided each one.

Round and verdict	Deciding rule	Value	Holds
One: mixed collapses	mixed final share below 0.01	0	yes
One: slope matches	slope error below 0.05	0.117	no
One: shuffled collapses	shuffled final share below 0.05	0	yes
One: space keeps producers	persistence in at least half the runs and above the shuffled 95th percentile	0	no
One: range fails	final share at range 25 below 0.05	0	yes
One: invasion resisted	local lawn above 0.3 and mixed lawn below 0.05	0.75; 0.125	no
Two: recursion matches	maximum error below 0.03	0.0019	yes
Two: mixed collapses	mixed final share below 0.01	0	yes
Two: space keeps producers	local above 0.8 and above the shuffled 95th percentile (0.52)	1.0	yes
Two: threshold near 8	first persistence ratio from 7 to 12 at every b	9	yes
Two: range fails	final share at range 25 below 0.05	0	yes
Two: invasion resisted	local lawn above 0.8 and mixed lawn below 0.05	1.0; 0.5	no
Three: artefact shown	shuffled with own goods above 0.5 at ratio 400	1.0	yes
Three: mixed collapses	mixed final share below 0.05 at ratio 400	0.038	yes
Three: others only persist somewhere	some ratio with final share of at least 0.5	none	no
Three: space credited	local others-only above the shuffled 95th percentile	0.48 against 0	untestable

Round two's two passing verdicts on space and on the threshold were produced by the flaw described in the methods, and round three withdrew them.

14 Glossary

Each term is defined where it is first used in the text. The same definitions, with the field each comes from and what each excludes, are collected here.

- **public good**. A product that costs its maker something and benefits others nearby, such as an enzyme or iron-scavenging molecule a microbe releases into its surroundings. Field: social evolution and economics. Excludes products the maker keeps for itself.
- **producer and cheater**. A producer pays a cost to make the public good; a cheater makes none and still gains from the goods of producers near it. Field: social evolution. Excludes any intention on either side; the words name two fixed strategies.
- **local competition**. Competition for space or resources among neighbours, which, when neighbours are relatives, can cancel the benefit that helping relatives brings. Field: social evolution, after Taylor and after Griffin and colleagues. Excludes competition spread over the whole population.
- **preregistration**. A dated written plan for a study, fixed and hash-locked before any result is seen: the question, the data, the statistic, the controls, the forecasts and the kill conditions. Field: research method (meta-

science). Excludes a plan revised after looking at the outcome, and a protocol written only for the record after the run.

- **kill condition**. A rule written in the preregistration that, when met, decides against the claim with no room for judgement after the run. Field: research method (our usage, after Popper's falsification). Excludes reasons found after the run for setting a result aside.
- **benefit-to-cost ratio**. The benefit b that one producer's goods give in total divided by the cost c it pays, written b/c. Field: social evolution. Excludes the net payoff to any one cell, which also depends on its neighbours.
- **torus grid**. A square grid whose opposite edges are joined, so every cell has the same number of neighbours and there is no border. Field: simulation. Excludes grids with edges, where border cells have fewer neighbours.
- **goods box**. The square of cells over which one producer's goods are spread, 3 by 3 at range 1 and 5 by 5 at range 2. Field: our usage for this study. Excludes the 3 by 3 window from which parents are drawn, which stays fixed.
- **synchronous update**. Every cell of the grid is replaced at once in each generation, from the fitnesses of the previous generation. Field: simulation. Excludes asynchronous updating, where one cell at a time is replaced.
- **fitness-proportional replacement**. The update rule in which a cell's next occupant is copied from a parent drawn with chance in proportion to the parent's fitness, here among the 3 by 3 window around the cell including the cell itself. Field: evolutionary dynamics. Excludes other update rules, such as death followed by birth from neighbours only, or imitation of the best neighbour.
- **mixed arm**. The version of the model in which every cell gains from the producer share of the whole population and parents are drawn from the whole population, so neighbours mean nothing. Field: our usage for this study. Excludes any spatial structure.
- **shuffled arm**. The version of the model in which all positions are shuffled at random before each generation and the local rules then apply, so neighbours are strangers. Field: our usage for this study. Excludes the clustering of relatives that the local arm allows.
- **final producer share**. The share of producers in the grid at the end of a run, averaged over the replicate runs; a run stops early when one type takes the whole grid. Field: our usage for this study. Excludes the share of runs in which producers survived.
- **control**. A version of the data with the structure under test removed (for example shuffled or pooled), scored the same way as the real data, so the real score has something to beat. Field: statistics and experimental design. Excludes a comparison with a different dataset or with a published number.
- **mean-field recursion**. An exact formula for next generation's producer share in the mixed arm, used here as a known answer the simulation must match. Field: population genetics. Excludes the local and shuffled arms, which have no such formula.

15 References

[1] Nadell, C. D., Drescher, K. and Foster, K. R. (2016). Spatial structure, cooperation and competition in biofilms. Nature Reviews Microbiology 14(9), 589-600. doi:10.1038/nrmicro.2016.84

[2] Taylor, P. D. (1992). Altruism in viscous populations: an inclusive fitness model. Evolutionary Ecology 6, 352-356. doi:10.1007/BF02270971

[3] Wilson, D. S., Pollock, G. B. and Dugatkin, L. A. (1992). Can altruism evolve in purely viscous populations? Evolutionary Ecology 6, 331-341. doi:10.1007/BF02270969

[4] Griffin, A. S., West, S. A. and Buckling, A. (2004). Cooperation and competition in pathogenic bacteria. Nature 430, 1024-1027. doi:10.1038/nature02744

[5] Ohtsuki, H., Hauert, C., Lieberman, E. and Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. Nature 441, 502-505. doi:10.1038/nature04605