SCIENCE UTOPIA? CLOSED-LOOP LLM SIMULATION OF ACADEMIC RESEARCH ECOSYSTEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

Scientific progress emerges from a longitudinal ecosystem in which researchers, institutions, funding agencies, collaboration networks, and the scientific literature co-evolve. As AI becomes increasingly involved throughout the scientific research cycle, understanding these interconnected and evolving processes becomes increasingly important. We introduce **SUTO**, a persistent, closed-loop LLM-agent simulation framework for studying academic research ecosystems. SUTO models interconnected scientific processes such as research-direction choice, collaboration, submission, peer review, resubmission, citation, funding, and researcher attrition, while maintaining evolving states across simulated years. Its configurable institutional mechanisms and information channels provide a controlled testbed for matched counterfactual experiments and targeted interventions. Across 61 simulation worlds, SUTO simulates over 40,000 researchers from 8,000 institutions, producing ∼400,000 publication decisions and 1.2 million LLM-generated peer reviews. Using these longitudinal simulations, we find that rejection-driven resubmission substantially amplifies reviewer burden beyond population growth alone, cautious exploration balances citation impact with career success and long-term topic diversity, and resource inequality can emerge even without detectable cumulative advantage from narrowly winning early funding. Code is available at anonymous.4open.science/r/ScienceUtopia.

1 INTRODUCTION

Artificial intelligence (AI) is rapidly reshaping the scientific research landscape, with large language models (LLMs) increasingly participating in research ideation, experimentation, scientific writing, and peer review (Lu et al., 2024; Yamada et al., 2025; Jin et al., 2024b; Yu et al., 2025). As AI adoption extends across interconnected research activities (Li et al., 2026; Weng et al., 2025), its implications reach beyond individual discoveries to the organization and evolution of the scientific ecosystem. At the *micro level*, AI-mediated decisions can reshape how researchers choose research topics, form collaboration networks, pursue publication, and evaluate submissions, as well as how funding agencies allocate resources. At the *macro level*, these decisions collectively reshape the direction and organization of science: which ideas gain influence, which research directions receive sustained support, and who remains able to contribute. Crucially, these outcomes feed back into subsequent research, reshaping the knowledge, incentives, and opportunities that guide future decisions. Such feedback can amplify or counteract the effects of individual choices, so improvements in isolated research activities need not translate into a more productive, diverse, or sustainable scientific ecosystem. This transformation creates an urgent need for a *closed-loop, longitudinal perspective* to understand not only how AI changes *individual* research activities, but also how these changes propagate over time through an evolving scientific environment to produce *ecosystem-level* benefits and unintended consequences.

Challenges. Tracing these effects and causal pathways, however, is difficult in real academia, where controlled institutional interventions are often infeasible and counterfactual outcomes under alternative policies are unobservable. The researcher population, available resources, collaboration networks, and scientific record evolve with prior decisions and, in turn, shape subsequent research. Although LLM agents provide a natural substrate for studying increasingly AI-mediated scientific processes in controlled settings through simulation (Appendix G), faithfully capturing this co-evolution poses three challenges. First, *scientific fidelity* requires more than plausible agent responses: re-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

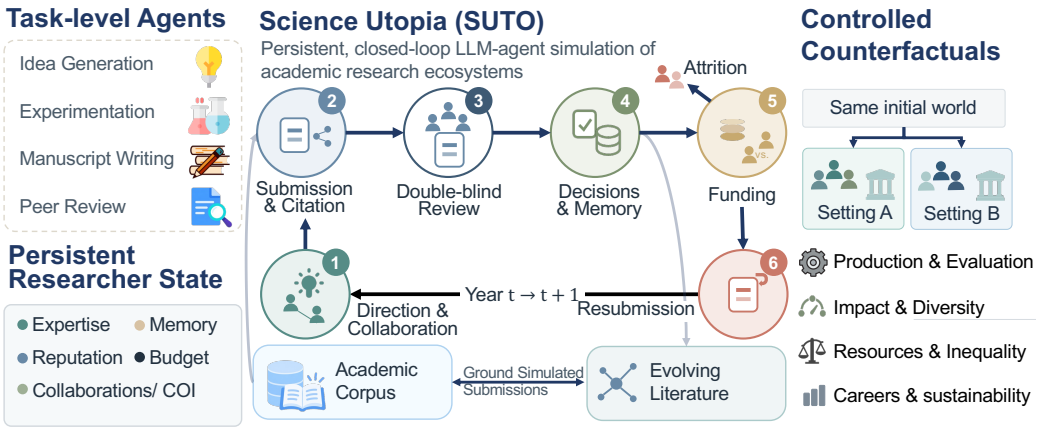


Figure 1: SUTO simulates each year as a six-phase cycle whose outcomes update persistent researcher states and inform subsequent decisions.

search choices and evaluations must be grounded in scientific content and reflect differences in researchers' expertise and strategies. Second, *closed-loop simulation* requires research outcomes to reshape the conditions for subsequent research. Lower acceptance rates may increase resubmissions, consuming authors' resources and adding review demand. Researchers who exhaust their budgets may leave, reducing both future submissions and the reviewer pool. Accepted papers expand the citable literature and update the publication records used in funding decisions. These changes must carry forward into the next year's decisions, so that both researcher behavior and the community's capacity to sustain research evolve within the simulation. Third, maintaining effective *experimental control* requires separating the mechanism being tested from the changes it produces, whereas ineffective design would suppress the feedback effects the simulation is intended to study.

This Work. We introduce SUTO (**S**cience **Uto**pia), a persistent, closed-loop LLM simulation framework for the longitudinal study of research ecosystems. SUTO integrates research-direction choice, collaboration, submission, peer review, resubmission, citation, funding, and researcher attrition through persistent researcher states and a shared scientific record. This structure supports matched counterfactual experiments that vary an institutional mechanism and follow its downstream effects under the same simulation rules. We investigate three questions. First, *how do publication and funding rules shape the long-term dynamics of research ecosystems?* Second, *how do individual research strategies shape career outcomes and the diversity of science?* Third, *under what conditions do cumulative advantage and bias emerge in funding and peer review?* Our analyses span 61 simulation worlds containing over 40,000 researcher agents across 8,000 institutions and approximately 390,000 simulated researcher-years. Overall, SUTO generates about 400,000 publication decisions and 1.2 million peer reviews. Our findings offer insights for future academic systems:

- **Repeated resubmission sustains a significantly greater manuscript backlog and substantially amplifies per-reviewer workload compared with researcher population growth.** Population growth alone expands the reviewer pool alongside submissions, leaving review demand near baseline (2.07 vs. 2.25 reviews per active reviewer). Resubmission alone raises this demand to 6.64. Stricter acceptance criteria compound this burden by shifting review work into future rounds, inducing 14% more review workload, even when the volume of new submissions changes little.
- **Publication growth can mask declining scientific participation.** Doubling the submission limit per project increases accepted output by 63–69% in university-only worlds, yet reduces long-term participation even when funding capacity scales with demand. Under fixed grant capacity, the share of initial researchers active at year ten falls from 39% to 24%.
- **Exploring novel but related topics balances citation impact with publication and funding success.** Researchers exploring new topics close to expertise (*cautious explorers*) achieve greater citation impact than those extending prior work (*exploiters*) and those moving into distant areas (*explorers*), with only a modest decline in paper acceptance and funding. When most researchers follow this strategy, long-term research activities remain broadly distributed across topics over time.

- **Narrow early funding wins show no significant career advantage over the subsequent three years.** However, resources become concentrated over time and nearly half of researchers exhaust their resources.
- **Revealing author information modestly raises review scores.** Revealing the first author's pseudonymous name and institution raises scores by 0.055–0.077 points relative to anonymous review on the 1–5 scale. Disclosing shared collaborators or other network connections has little further effect once author information is visible.

We hope this work will lay a foundation for understanding and reshaping research ecosystems and, more importantly, help shape the future of scientific research in the age of AI.

2 THE SUTO FRAMEWORK

As shown in Appendix Figure 13, SUTO simulates academic research as an LLM-powered, closed-loop cycle, with research, peer review, publication, and funding outcomes shaping subsequent decisions across simulated years.

2.1 RESEARCHERS, INSTITUTIONS, AND VENUES

Researchers. A researcher i maintains a state in year t, $s_i^t = \left(E_i, \rho_i^t, b_i^t, M_i^t, C_i^t\right)$, where E_i, ρ_i^t, and b_i^t are the researcher's expertise topics, reputation, and available budget, respectively. C_i^t is the set of conflicts of interest induced by affiliation and co-authorship. M_i^t is the agent's memory that records previous events, such as publications, peer-review experience, and funding outcomes. The memory connects past outcomes to subsequent decisions: prior reviews can guide revision and resubmission, while publication outcomes can influence whether a researcher continues or changes research direction. This feedback enables us to study how institutional rules shape research trajectories over time beyond their immediate effects on acceptance and funding. The budget b_i^t links research and funding outcomes to continued participation in the ecosystem. At the end of each year t, it is updated as $b_i^{t+1} = b_i^t + f_i^t - \sum_{a \in \mathcal{A}_i^t} \mathrm{cost}(a)$, where f_i^t is funding received during the year, $\mathcal{A}_i^t$ is the set of actions taken by researcher i, and $\mathrm{cost}(a)$ is the cost of action a. Researchers whose budget is exhausted become inactive and leave the active population. The budget can be interpreted as an abstract resource constraint, analogous to AI-usage tokens.

To examine how funding arrangements shape researchers' strategies and participation, we distinguish two institutional settings, motivated by differences between academic and industrial science. *University researchers* rely on external competitive grants, whereas *industry researchers* receive internal funding tied to accepted papers. Meanwhile, strategies between exploiting familiar topics and exploring new ones shape both individual career outcomes and the collective development of science (Foster et al., 2015). Prior work characterizes this trade-off along two dimensions: *exploration propensity*, how often researchers enter unfamiliar areas, and *exploration distance*, how far they depart from their prior work (Yang et al., 2023). To examine how institutional rules reward different *exploration strategies*, we consider three representative approaches: *explorers* favor distant moves into new areas; *cautious explorers* frequently explore nearby areas, balancing novelty with continuity in prior expertise; and *exploiters* preferentially continue along established research directions.

Funding Programs. SUTO includes two families of funding programs. Foundation-oriented programs prioritize theoretical and foundational work, while mission-oriented programs favor application-driven areas such as autonomous systems and security. This distinction is motivated by research-support and mission-focused programs such as NSF CAREER and DOE's Genesis Mission (National Science Foundation, n.d.; U.S. Department of Energy, n.d.; Braun, 1998; Mazzucato, 2018). Each program specifies eligible topics, evaluates submitted applications, and funds the top-ranked fraction of applications assigned to its panel.

Venues and scientific artifacts. Scientific artifacts are grounded in real literature so that ecosystem dynamics are not dominated by the factual reliability of fully generated papers. Each conference c covers a subset of research topics and has an acceptance rate α_c. A submission must be topically compatible with its venue. Given its selected research direction, an agent retrieves temporally valid papers from the corpus and selects top papers for submission to a compatible venue. This

retrieval-grounded design provides scientific material for review and citation while restricting retrieved artifacts by year. By default, each completed project allows at most one new submission, which may cite previously published work, including the researcher's own papers. Accepted papers and citations are tracked, making the evolving literature available for citation in subsequent years.

2.2 SIMULATION PIPELINE

Each simulated year follows the same execution order. Simulated year t is mapped to a real calendar year, so agents can only submit or cite papers that would have been available by that point in time. The phases below describe the default loop. Individual experiments may invoke additional controlled interventions without changing the pipeline. The motivation is in Appendix F.4.

Phase 1: Research Direction and Collaboration. Active agents whose previous projects have concluded (if any) select a new research direction $d \in \mathcal{D}$, where $\mathcal{D}$ is the fixed set of available research directions, based on their expertise, resources, and memory. A direction d lasts for T_d simulated years and costs κ budget units per active year, so its total cost is κT_d. This duration makes resource pressure visible: longer directions consume more budget before producing new outputs. Agents can choose to form co-authorship ties when collaboration is enabled. Agent i scores a candidate collaborator j using $\psi(i,j) = \sum_m w_m \phi_m(i,j)$, where each feature $\phi_m(i,j)$ captures one collaboration signal, such as topical similarity, recent-publication overlap, preferential attachment, or triadic closure, and w_m is its weight.

Phase 2: Paper Submission and Citation. Authors may submit various numbers of new papers in a year depending on their bandwidth. For each potential submission, SUTO retrieves candidate arXiv papers from SciEvo that match the agent's direction and are available in the current simulated year; the agent then chooses one candidate artifact and a compatible venue. Research costs and resubmission fees enter the budget update. After choosing the artifact, the agent cites prior papers from the accessible literature. Self-citation behavior is analyzed in Appendix D.2.

Phase 3: Peer Review. Each submission is assigned three reviewers R_p. Each reviewer j uses the title, abstract, and topics to provide a recommendation u_{jp} and written justification for paper p, following prior work (Jin et al., 2024b). Reviews are double-blind by default.

Phase 4: Paper Decisions and Memory Update. For a paper p reviewed by panel R_p, the decision score is the mean recommendation $\bar{u}_p = \frac{1}{|R_p|} \sum_{j \in R_p} u_{jp}$. Within each year, conference c ranks its submissions S_c by $\bar{u}_p$ and, for nonempty S_c, accepts the top $\max(1, \mathrm{round}(\alpha_c |S_c|))$ papers, where α_c is the acceptance rate. Authors retain experiences such as reviews and publication decisions in memory, which shape future behavior. For example, negative reviews may affect resubmission decisions and prompt harsher evaluations of others' work, while publication success may reinforce a research direction. Accepted papers enter the citable literature and can be cited by submissions in later years.

Phase 5: Funding Allocation and Attrition. University researchers submit short proposals grounded in their expertise and prior papers. Funding agents evaluate applications for program fit, research potential, and track record, considering both accepted and rejected papers.[1] Under rate-based funding, each panel ranks its n applications and funds the top $k = \max(1, \lfloor rn \rfloor)$ applicants, where r is the program's funding rate. To study how *funding conservatism* shapes downstream research trajectories, we introduce a controlled intervention that penalizes topical departure while leaving the underlying evaluation unchanged. We convert researcher i's within-panel rank into a normalized score $e_i \in [0,1]$, with higher scores indicating better ranks. We measure topical departure by $q_i \in [0,1]$, the percentile rank of the researcher's mean paper distance among the given year's applicants. We then adjust the funding score as $\tilde{e}_i = e_i - \lambda q_i$, where $\lambda \geq 0$ controls the penalty strength, and award funding to applicants with the highest adjusted scores. Using normalized ranks makes the penalty independent of the raw score scales. The reported experiments use $\lambda = 0$, preserving the original ranking.

Phase 6: Resubmission. From the second year onward, agents with previously rejected papers first decide whether to abandon, revise, or resubmit those papers. The decision is conditioned on

[1] We do not simulate the "proposal-funding" pipeline by assuming that the publications are positively associated with peer-review evaluations and funding outcomes of research proposals (Cabezas-Clavijo et al., 2013).

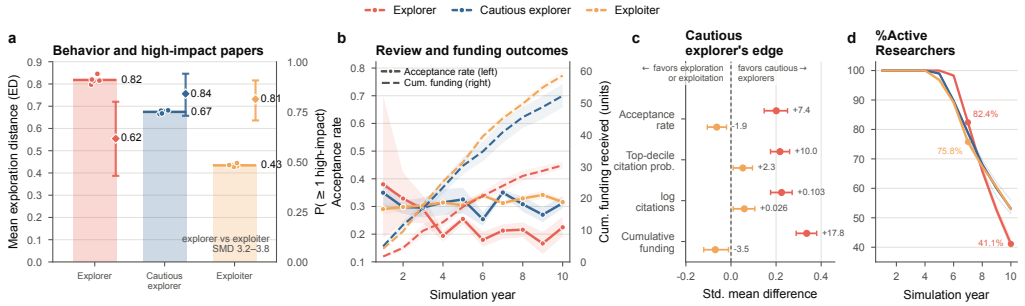


Figure 2: **Career outcomes across research strategies, including explorers, exploiters, and cautious explorers. (a)** Exploration distance and high-impact-paper probability across five calibration worlds. **(b)** Acceptance and cumulative funding in the same calibration worlds. **(c)** Cautious-explorer contrasts in three 5,000-researcher confirmatory worlds: standardized mean differences with 95% CIs. Here citation probability uses researcher-level citation totals. **(d)** Active researchers in the confirmatory worlds.

the received reviews, the paper's prior submission history, the agent's current resources, and the resubmission cost. Resubmitted papers re-enter the review process alongside new submissions.

3 Experiments and Findings

We organize our findings across three levels of the scientific ecosystem: researchers, examining how research strategies shape careers (§3.1); research outcomes, connecting topical departure to impact and disruption and examining how resubmission and peer review shape publication outcomes (§3.2); and institutional mechanisms, examining how funding capacity and allocation shape resource access and continued participation (§3.3). Finally, we examine the ecosystem-level dynamics emerging over time (§3.4). Experimental details are in Appendix F.

3.1 Human Factors: Exploration Strategies and Careers

We track *explorers*, *exploiters*, and *cautious explorers* (§2.1) over 10 years while holding initial funding and reputation constant. We use two complementary distances to quantify researchers' movement through the knowledge space (Huang et al., 2022; Yang et al., 2023; Zhang et al., 2024). For paper p initially submitted by researcher i in year t, we compute its *paper distance* from researcher i's recent work as $\nu^t_{p,i} = 1 - \cos(\mathbf{p}, \mathbf{a}^t_i) \in [0, 2]$, where $\mathbf{p}$ is the paper embedding and $\mathbf{a}^t_i$ is the mean embedding of the author's work over the preceding four years. A researcher's *exploration distance (ED)* is the mean annual cosine distance from each paper to their most recent paper. Thus, paper distance measures departure from recent expertise, whereas ED measures successive movement through knowledge space. The resulting trajectories match the intended strategies: explorers have the largest ED, exploiters the smallest, and cautious explorers lie between them (Figure 2a).

Cautious exploration balances novelty with career reward. Distinct exploration strategies lead to different career outcomes. Explorers have the lowest acceptance rate (29.8%), compared with 37.2% for cautious explorers and 39.2% for exploiters (Figure 2b), and are also less likely to produce a high-impact paper (0.62 versus 0.84 and 0.81 in Figure 2a). Cautious explorers outperform explorers in acceptance, citation outcomes, and cumulative funding (Figure 2c), while relative to exploiters they trade modestly lower acceptance and funding for stronger citation impact. Thus, exploration is most effective when it extends existing expertise without losing institutional support.

The costs of distant exploration accumulate over time. Explorers' early advantage in remaining active reverses as publication and funding outcomes accumulate. By year 10, only 41.1% of explorers remain active, compared with ∼53% of cautious explorers and exploiters. This delayed reversal suggests that early persistence can mask accumulating disadvantages in publication and funding.

Cautious exploration sustains long-run scientific diversity. To test whether individual-level advantages scale to the ecosystem, we construct matched counterfactual worlds that differ only in researcher composition—explorer-heavy, exploiter-heavy, or cautious-heavy—while holding institutions, venues, and all other simulation settings fixed. We measure ecosystem diversity using nor-

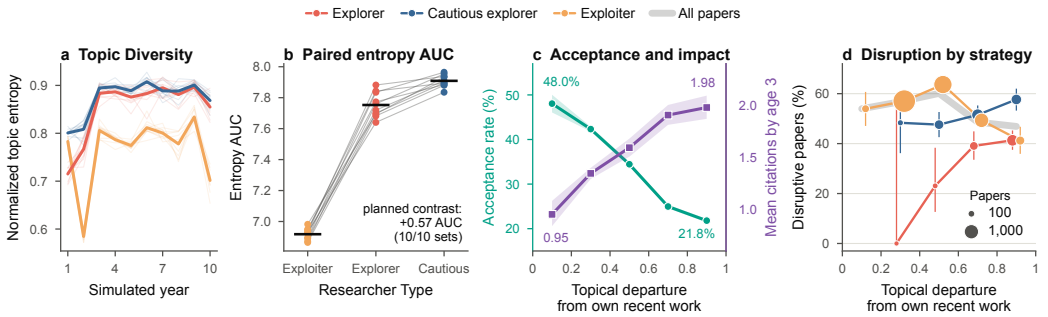


Figure 3: **(a)** Normalized topic entropy over ten simulated years. **(b)** Ten-year entropy AUC (Eq. 1) across matched counterfactual worlds; lines connect worlds within each matched set. **(c)** Acceptance rate and mean citations by age three across topical-departure bins. **(d)** Share of disruptive papers (CD > 0 at age three) by strategy and topical departure; marker area denotes the number of papers and gray pools all strategies. (a, b) compare ten matched sets of explorer-heavy, exploiter-heavy, and cautious-heavy worlds. (c, d) pool papers from three large-scale ten-year simulation worlds.

malized topic entropy and summarize over the full simulation by the area under the entropy curve:

$$H_t = -\frac{\sum_{d\in\mathcal{D}} p_{t,d}\log p_{t,d}}{\log|\mathcal{D}|} \in [0,1], \qquad \mathrm{AUC}_H = \sum_{t=1}^{T-1}\frac{H_t + H_{t+1}}{2} \in [0,\, T-1], \tag{1}$$

Here, $p_{t,d}$ is the fraction of all direction choices made in year t that select d, with $\sum_{d\in\mathcal{D}} p_{t,d} = 1$. Higher H_t indicates a more even distribution of research exploration across topics. With one-year spacing between observations, AUC_H summarizes diversity over years 1 through $T = 10$. Varying the proportions of the three strategies substantially changes the resulting research landscape (Figure 3). Cautious-heavy populations maintain 0.11 higher normalized topic entropy than exploiter-heavy populations and end with 8.5 more active research directions across all ten matched counterfactual sets (Figure 3). They also outperform explorer-heavy populations in ten-year AUC_H by 0.156 on average, despite explorers making larger individual departures from prior expertise. Explorer-heavy worlds initially spread research directions more broadly, but manifest greater long-run attrition and lower diversity. Policies that encourage exploration should therefore also help exploratory researchers remain active. Otherwise, research may gradually concentrate on topics with safer career prospects, leaving fewer people to pursue risky but potentially valuable directions.

3.2 Research Outcomes: Impact, Resubmission, and Peer Review

The researcher-level results raise a natural question: is exploratory work less influential, or simply harder to publish? We first relate topical departure to acceptance, impact, and disruption, then examine how resubmission and reviewer knowledge shape publication outcomes and review workload.

Greater topical departure faces a stronger publication gate but higher impact once accepted. Papers that depart further from an author's recent research are substantially harder to publish (Figure 3c). From the lowest to the highest departure bin, acceptance falls from 48.0% to 21.8%, while mean review scores decrease from 2.45 to 2.28. Among accepted papers, however, downstream impact moves in the opposite direction. Mean citations by age three increase from 0.95 to 1.98, and the 90th percentile increases from 3 to 5. These results reveal a selection–impact trade-off: work farther from an author's recent trajectory is less likely to pass the publication gate, but the subset that is accepted receives greater subsequent impact.

Topical departure does not imply disruptiveness. Do papers that move farther from their authors' prior trajectories also redirect subsequent research away from existing research? Prior work therefore distinguishes between *disruptive* contributions, which make prior approaches less central to subsequent work, and *consolidating* contributions, which reinforce and extend existing streams of knowledge (Park et al., 2023; Funk & Owen-Smith, 2017). We quantify this distinction using the CD index (Funk & Owen-Smith, 2017). For a focal paper p, let n_{focal} denote later papers that cite p but not its references, n_{both} papers that cite both p and its references, and n_{prior} papers that cite p's references but not p: $\mathrm{CD}_p = \frac{n_{\mathrm{focal}} - n_{\mathrm{both}}}{n_{\mathrm{focal}} + n_{\mathrm{both}} + n_{\mathrm{prior}}}$. Positive values indicate *disruption*, where subsequent work increasingly builds on the focal paper while bypassing its predecessors. Negative values indicate *consolidation*, where the focal paper and its predecessors form coherent themes and

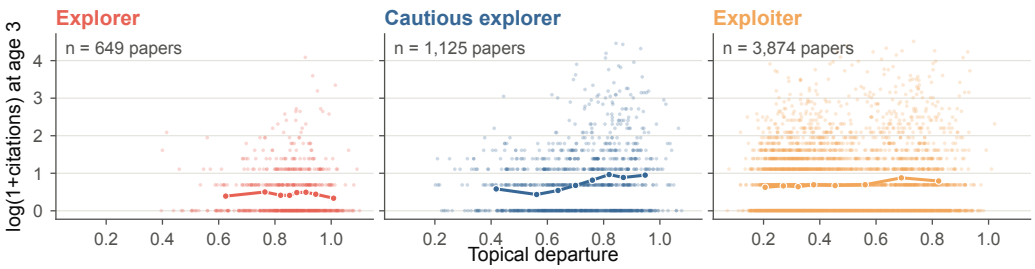


Figure 4: **Topical departure is associated with the clearest citation gains for cautious explorers.** Each point represents an accepted paper, with topical departure from the author's recent work on the x-axis and citations within three years on the y-axis ($\log(1 + \text{citations})$). Lines show mean citations within departure octiles. Citation gains rise most consistently with departure for cautious explorers, while the relationship is weak for explorers and exploiters.

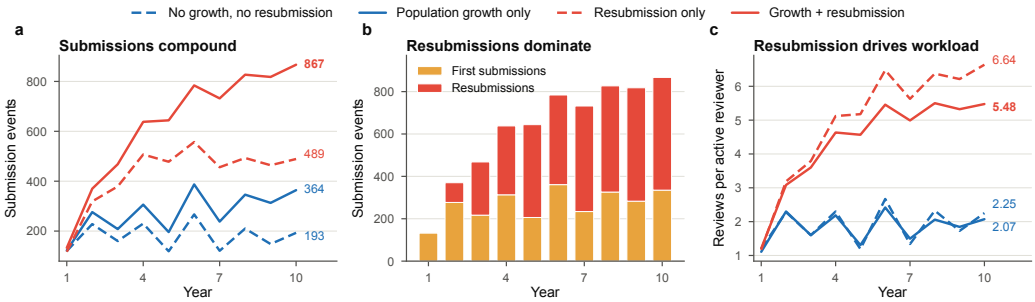


Figure 5: **Population growth and resubmission compound submission volume but have different effects on reviewer workload.** We cross researcher population growth and resubmission over ten simulated years. Blue / red denotes disabling / enabling resubmission, and dashed / solid lines denote without / with population growth, respectively. **(a)** Annual submission volume. The joint condition produces more submissions than the sum of the increases in either mechanism alone. **(b)** Submission composition in the joint condition. Resubmissions account for 61% of year-10 submissions. **(c)** Per-reviewer workload. Population growth alone leaves per-reviewer workload close to baseline, whereas conditions with resubmission show substantially higher workload.

continue to be referenced jointly. Across strategies, greater topical departure does not correspond to a higher probability of disruption: the share of papers with positive CD falls from 54.0% in the lowest-departure bin to 46.9% in the highest (Figure 3d). However, this relationship differs by research strategy. From intermediate to high departure, this share falls from 63.6% to 41.2% for exploiters but rises from 47.6% to 57.7% for cautious explorers. The latter increase appears in all three simulated worlds. Explorers have the lowest share overall (39.8%), never exceeding 41.3% in any departure bin. Changing one's own research direction is therefore not equivalent to changing the direction of the field. Exploration grounded in prior expertise is more likely to produce departures that subsequent research builds on.

Resubmission pressure intensifies under greater selectivity. We compare matched ten-year simulations that allow rejected papers to be resubmitted: one uses a fixed 30% acceptance rate, and the other follows ACL main-conference acceptance rates from 2016 (28%) to 2025 (20.3%) (Xu, 2026). The two regimes produce similar numbers of first submissions over the decade (1,706 vs. 1,741), but diverge substantially in total submission growth: 282% under fixed acceptance versus 358% under declining acceptance. Rejected papers also accumulate more venue attempts under the declining regime (3.93 vs. 3.78 on average), while mean annual reviews among reviewers assigned at least one paper rise from 1.7 to 5.5, compared with 1.8 to 4.6 under fixed acceptance (Figure 8). Among papers accepted after resubmission, prior rejections reach ∼3 by the final year in both regimes. This comparison shows how greater selectivity can increase review workload without a corresponding increase in new manuscripts: rejected papers shift review activity into future rounds. Further results on repeated submission attempts are provided in Appendix B.3 (Figure 8).

Revealing author names and institutions slightly increases LLM review scores. We evaluate each reviewer–paper pair under three conditions: 1) double-blind review, 2) disclosure of the first author's pseudonymous name and institution, and 3) disclosure of (2) plus a description of

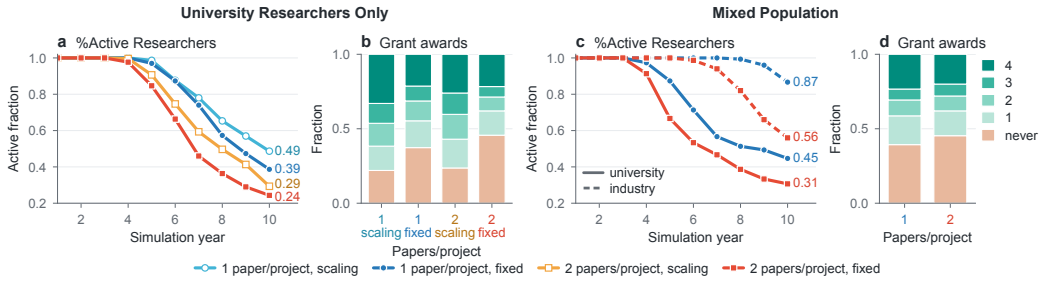


Figure 6: We vary submission volume per completed project. **(a,b)** University-researcher-only simulations compare scaling versus fixed grant capacity. **(c,d)** Mixed-population simulations (university and industry researchers) use fixed grant capacity and a fixed industry compensation pool. **(a,c)** Fraction of the initial researcher population remaining active over years. In (c), solid and dashed lines denote university and industry researchers. **(b,d)** Distribution of cumulative grant awards among the initial university researchers. "4" denotes four or more awards and "never" denotes no awards.

the reviewer–author relationship in the collaboration network. Author-information disclosure raises scores by 0.055 points on the 1–5 scale. A separate 120-paper follow-up, with this contrast prespecified as its primary comparison, finds a 0.077-point increase (95% CI $[0.026, 0.128]$). Additional studies are in Appendix B.5.

3.3 Institutional Mechanisms: Funding Capacity and Allocation

We examine how institutional resources support scientific production and continued participation. First, we test whether funding capacity keeps pace with higher research output. Then, we ask whether narrowly securing an early grant creates a lasting funding advantage.

Higher research output expands publication volume faster than ecosystem capacity. AI-assisted research may enable researchers to produce more papers from each completed project, but the resources needed to sustain this additional activity may not expand accordingly. Does increasing individual research output necessarily benefit the scientific ecosystem? To separate the effect of higher output from differences in research funding modes, we first study university-researcher-only worlds by increasing the submission limit from one to two papers per completed project, under grant capacity that either scales with demand or remains fixed. Doubling submission intensity increases accepted papers by 69% under scaling capacity and 63% under fixed capacity, but sharply reduces long-run participation. By year ten, the fraction of initial researchers still active falls from 49% to 29% under scaling capacity and from 39% to 24% under fixed capacity (Figure 6a). These results illustrate a potential risk of increasing research output without adapting the mechanisms that sustain researcher participation. Even when grant capacity scales with demand, higher submission intensity produces more accepted papers but leaves fewer researchers active, with little improvement in citation coverage or topic diversity.

Fixed grant capacity primarily reduces access to funding, with smaller changes in inequality among recipients (Figure 6b). Compared with capacity that scales with demand, the share of researchers who never receive an award rises from 22% to 37% with one submission per project and from 24% to 46% with two submissions. The Gini coefficient among funded researchers increases only modestly, from 0.36 to 0.39 and from 0.40 to 0.43, respectively. Resource scarcity thus manifests mainly in more researchers receiving no funding, rather than greater concentration among recipients. The same pattern appears in the mixed university–industry population (Figure 6c,d). Doubling submission intensity increases accepted papers by 77%, but reduces year-ten participation from 45% to 31% among university researchers and from 87% to 56% among industry researchers. Larger initial resource endowments delay industry attrition but do not prevent it under a fixed compensation pool. Among university researchers, the share never funded rises from 39% to 45%, while inequality among recipients changes modestly. Citation coverage decreases from 60% to 57%, and topic entropy changes little (0.69 to 0.71). Overall, publication growth and modest changes in inequality among funding recipients can conceal a broader contraction in participation and funding access.

3.4 Emergent Ecosystem-Level Dynamics

The study in Appendix C.2 finds no detectable subsequent-funding advantage. This raises a broader question: *can persistent inequality nevertheless emerge endogenously from the coupled dynamics of scientific production, citation, funding, and researcher attrition?* To examine this question, we run three large confirmatory worlds, each with $1{,}000$ institutions and $5{,}000$ researchers, for a full simulated decade without experimental intervention. We track the evolution of resource inequality, topic diversity, active population size, and peer-review dynamics (Appendix D.1; Figure 11).

Substantial stratification emerges over time. Funding-resource inequality increases steadily over time: the Gini coefficient rises from near parity at initialization (≈ 0.04) to ≈ 0.36 by year ten (Figure 11a). Citation inequality remains consistently high throughout the simulation (≈ 0.56 to ≈ 0.54). Importantly, this concentration of resources does not coincide with a collapse in research diversity. New research-direction choices remain broadly distributed across topics, with topic entropy remaining high at $\sim$0.81 to $\sim$0.85. By the end of the simulation, all 53 available research directions attract active researchers (Figure 11b). This reveals a decoupling between resource inequality and intellectual diversity: resources become increasingly concentrated even as the research landscape remains broadly diversified. At the same time, inequality is accompanied by substantial researcher attrition. The active population declines from $5{,}000$ researchers to approximately $2{,}530$ by year ten as resource-depleted researchers exit the system (Figure 11c). Peer-review dynamics remain comparatively stable. Across 455K individual reviews, the mean score remains between 2.29 and 2.40, with 62–72% of scores falling between 2.0 and 2.5. The standard deviation across individual reviews increases slightly from 0.50 to 0.55, while mean disagreement among reviewers of the same paper increases from 0.29 to 0.35 (Figure 11d).

4 Related Work

Large language model agents and agent-based models of science. Recent systems use LLM agents for idea generation (Lu et al., 2024; Tang et al., 2025), experimentation (Yamada et al., 2025), manuscript preparation (Weng et al., 2025), and peer review (Jin et al., 2024b; Yang et al., 2026). In parallel, a longer tradition of agent-based models studies how individual incentives and institutional rules generate collective outcomes (Šešelja, 2022), including publication and resubmission (Kovanis et al., 2016), resource-constrained publishing and reviewing (Bianchi & Squazzoni, 2022), long-term effects of publication incentives (Smaldino & McElreath, 2016), funding and cumulative advantage (Merton, 1968; Bol et al., 2018), and risks from adversarial participants (Jiang et al., 2026). SUTO couples research choice, collaboration, publication, peer review, resubmission, citation, funding, resources, and career survival within a persistent, literature-grounded LLM ecosystem, enabling controlled studies of how local mechanisms propagate into longer-run outcomes.

Large language model-based social simulation. LLM-based simulations study complex systems that are difficult to manipulate directly (Anthis et al., 2025). These provide controlled tests of how differences emerge and accumulate over time. This is particularly important as LLM generations increasingly influence and mutually reinforce human writing and communication (Kobak et al., 2025; Yakura et al., 2024; Zhu & Cong, 2024). Closer to our setting, recent work moves toward research-community simulation: AgentRXiv (Schmidgall & Moor, 2025) lets agents build on prior work across repeated interactions, while ResearchTown (Yu et al., 2025) models researchers and papers as an agent–data graph for reading, writing, and reviewing.

5 Conclusion

We presented SUTO, a holistic, LLM-based framework for simulating academic research as an evolving ecosystem. By integrating research-direction choice, collaboration, submission, peer review, resubmission, citation, funding, memory, and attrition into a single literature-grounded loop, SUTO helps policy-makers understand how local decisions accumulate into long-term scientific dynamics. It provides a testbed for asking not only how scientific agents behave, but how the rules of academia shape what research survives, spreads, and ultimately succeeds.

AI USE STATEMENT

We used generative AI tools to assist with language refinement, organization, translation, limited feedback on research methodology and experiments, limited code assistance, and the derivation of auxiliary metadata from existing data, such as keyword tags for papers. Separately, Qwen3-32B serves as the LLM backbone of the agents in our simulated scientific ecosystem. This constitutes part of the experimental methodology rather than research or writing assistance. The core conceptual framework, hypotheses, experimental design, methodological decisions, and interpretation of results were developed and finalized by the authors. All AI-assisted manuscript content and code were reviewed and verified by the authors, and AI-assisted implementations were tested for correctness by at least two authors. The authors take full responsibility for the final content, claims, and artifacts presented in this work.

ETHICS STATEMENT

We adhere to the ICLR Code of Ethics. This study does not involve human subjects, private user data, or interventions in real-world scientific decision making. Our experiments use the publicly available SciEvo dataset, released under the Apache 2.0 License. The source papers and bibliographic metadata are drawn from real publications. The additional trajectories, reviews, and institutional outcomes are generated within the simulation through LLM-agent behavior and explicit computational rules; they should not be interpreted as factual assessments of the original authors or institutions. All experimental LLM inference is performed using locally hosted models.

Our SUTO framework studies scientific processes including research-direction choice, collaboration, peer review, funding, citation, and career dynamics in controlled simulated environments. As with any LLM-based study, the resulting outcomes depend on the underlying data, agent behavior, and specified institutional mechanisms. We therefore focus on controlled comparisons across matched simulation conditions and on the mechanisms that emerge under these settings, rather than treating absolute simulated quantities as direct estimates of real-world outcomes. SUTO is intended to complement empirical studies by enabling controlled investigation of mechanisms that are difficult to isolate in observational data. Its findings can inform reasoning about scientific institutions and policy, while real-world decisions should additionally incorporate empirical evidence and domain expertise. Model and corpus biases may affect the simulated outcomes, which should not be used alone to evaluate individuals or determine real-world funding or hiring decisions.

REPRODUCIBILITY STATEMENT

Section 2 describes the simulation framework, Appendices A–D describe the experiments and analyses, and Appendix F describes the data and model setup. An anonymous code release is currently available at anonymous.4open.science/r/ScienceUtopia. Upon acceptance, we will release the code, processed data, experiment configurations, and instructions required to reproduce the main results of this work. The repository will include scripts for constructing and running the simulated worlds, reproducing the experimental conditions, and generating the reported analyses. Experimental settings and simulation parameters are specified through configuration files, allowing each reported simulation condition to be reconstructed and rerun under the same setup.

REFERENCES

Jacy Reese Anthis, Ryan Liu, Sean M Richardson, Austin C Kozlowski, Bernard Koch, Erik Brynjolfsson, James Evans, and Michael S Bernstein. Position: Llm social simulations are a promising research method. In *ICML*, 2025.

Federico Bianchi and Flaminio Squazzoni. Can transparency undermine peer review? a simulation model of scientist behavior under open peer review. *Science and Public Policy*, 49(5):791–800, 2022.

Thijs Bol, Mathijs De Vaan, and Arnout Van De Rijt. The matthew effect in science funding. *Proceedings of the National Academy of Sciences*, 115(19):4887–4890, 2018.

Dietmar Braun. The role of funding agencies in the cognitive development of science. *Research policy*, 27(8):807–821, 1998.

Alvaro Cabezas-Clavijo, Nicolás Robinson-García, Manuel Escabias, and Evaristo Jiménez-Contreras. Reviewers' ratings and bibliometric indicators: Hand in hand when assessing over research proposals? *PloS one*, 8(6):e68258, 2013.

Jacob G Foster, Andrey Rzhetsky, and James A Evans. Tradition and innovation in scientists' research strategies. *American sociological review*, 80(5):875–908, 2015.

Russell J Funk and Jason Owen-Smith. A dynamic network measure of technological change. *Management science*, 63(3):791–817, 2017.

Shengzhi Huang, Wei Lu, Yi Bu, and Yong Huang. Revisiting the exploration-exploitation behavior of scholars' research topic selection: Evidence from a large-scale bibliographic database. *Information Processing & Management*, 59(6):103110, 2022.

Fengqing Jiang, Yichen Feng, Yuetai Li, Luyao Niu, Basel Alomair, and Radha Poovendran. Badscientist: Can a research agent write convincing but unsound papers that fool llm reviewers? In *ACL*, pp. 24712–24727, 2026.

Yiqiao Jin, Yijia Xiao, Yiyang Wang, and Jindong Wang. Scievo: A 2 million, 30-year cross-disciplinary dataset for temporal scientometric analysis. *arXiv:2410.09510*, 2024a.

Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, Yijia Xiao, and Jindong Wang. Agentreview: Exploring peer review dynamics with llm agents. In *EMNLP*, pp. 1208–1226, 2024b.

Dmitry Kobak, Rita González-Márquez, Emőke-Ágnes Horvát, and Jan Lause. Delving into llm-assisted writing in biomedical publications through excess vocabulary. *Science Advances*, 11(27): eadt3813, 2025.

Michail Kovanis, Raphaël Porcher, Philippe Ravaud, and Ludovic Trinquart. Complex systems approach to scientific publication and peer-review system: development of an agent-based model calibrated with empirical journal data. *Scientometrics*, 106(2):695–715, 2016.

Fenghai Li, Zihan Tang, Haofei Yu, Yining Zhao, and Jiaxuan You. Can large language models forecast what researchers study next? *arXiv:2609.00747*, 2026.

Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv:2408.06292*, 2024.

Mariana Mazzucato. Mission-oriented innovation policies: challenges and opportunities. *Industrial and corporate change*, 27(5):803–815, 2018.

Robert K Merton. The matthew effect in science: The reward and communication systems of science are considered. *Science*, 159(3810):56–63, 1968.

National Science Foundation. Faculty Early Career Development Program (CAREER), n.d. URL `https://www.nsf.gov/funding/opportunities/career-faculty-early-career-development-program`. Accessed: September 20, 2026.

Michael Park, Erin Leahey, and Russell J Funk. Papers and patents are becoming less disruptive over time. *Nature*, 613(7942):138–144, 2023.

Samuel Schmidgall and Michael Moor. Agentrxiv: Towards collaborative autonomous research. *arXiv:2503.18102*, 2025.

Dunja Šešelja. Agent-based models of scientific interaction. *Philosophy Compass*, 17(7):e12855, 2022.

Paul E Smaldino and Richard McElreath. The natural selection of bad science. *Royal Society open science*, 3(9):160384, 2016.

Jiabin Tang, Lianghao Xia, Zhonghang Li, and Chao Huang. Ai-researcher: Autonomous scientific innovation. *NeurIPS*, 38:9481–9520, 2025.

U.S. Department of Energy. The Genesis Mission, n.d. URL https://www.energy.gov/undersecretaryforscience/genesis-mission/genesis-mission. Accessed: September 20, 2026.

Yang Wang, Benjamin F Jones, and Dashun Wang. Early-career setback and future career impact. *Nature Communications*, 10(1):4331, 2019.

Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo Zhang, Jindong Wang, Yue Zhang, and Linyi Yang. Cycleresearcher: Improving automated research via automated review. In *ICLR*, 2025.

Xovee Xu. Computer science conference acceptance rates. https://csconfstats.xoveexu.com/conferences/, 2026. Accessed: 2026-03-28.

Hiromu Yakura, Ezequiel Lopez-Lopez, Levin Brinkmann, Ignacio Serna, Prateek Gupta, Ivan Soraperra, and Iyad Rahwan. Empirical evidence of large language model's influence on human spoken communication. *arXiv:2409.01754*, 2024.

Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *arXiv:2504.08066*, 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv:2505.09388*, 2025.

Jing Yang, Qiyao Wei, and Jiaxin Pei. Paper copilot: Tracking the evolution of peer review in ai conferences. In *ICLR*, 2026.

Xingsheng Yang, Zhaoru Ke, Qing Ke, Haipeng Zhang, and Fengnan Gao. Cautious explorers generate more future academic impact. *arXiv preprint arXiv:2306.16643*, 2023.

Haofei Yu, Zhaochen Hong, Zirui Cheng, Kunlun Zhu, Keyang Xuan, Jinwei Yao, Tao Feng, and Jiaxuan You. Researchtown: Simulator of human research community. In *ICML*, 2025.

Lin Zhang, Fan Qi, Gunnar Sivertsen, Liming Liang, and David Campbell. Gender differences in the patterns and consequences of changing research directions in scientific careers. *Quantitative Science Studies*, 5(4):882–905, 2024.

Wu Zhu and Lin William Cong. Divergent llm adoption and heterogeneous convergence paths in research writing. *Cornell SC Johnson College of Business Research Paper Forthcoming*, 2024.

A Human Factors: Exploration Strategies and Careers

A.1 Cautious exploration balances novelty with career reward

Reported worlds. The individual-strategy study uses five calibration worlds and three confirmatory worlds with 1,000 institutions over ten years. Each institution has five researchers: one explorer, three exploiters, and one cautious explorer, initialized with equal funding and reputation. Figure 2a,b uses the calibration worlds, and c,d uses the confirmatory worlds.

Metrics. For research strategy, we measure exploration distance (ED), direction-switching, acceptance rate, funding, and citation outcomes. Figure 2a reports the probability of producing at least one paper in its cohort's top citation decile. Panel c reports whether a researcher's aggregated citation total reaches the within-world 90th percentile (ties included).

Strategy manipulation check. Figure 2a shows the calibration-world manipulation check (mean ED 0.82, 0.67, and 0.43 for explorers, cautious explorers, and exploiters). The confirmatory worlds show the same ordering: explorers move farthest from their preceding papers (mean ED 0.83), cautious explorers occupy the middle (0.69), and exploiters stay closest to their preceding papers (0.49).

A.2 Cautious exploration sustains long-run scientific diversity

Reported worlds. The study that varies the proportions of the three strategies uses ten matched triples of 100-institution worlds with explorer-heavy $(3/1/1)$, exploiter-heavy $(1/4/0)$, or cautious-heavy $(1/1/3)$ compositions.

Metrics. For worlds that differ in the proportions of the three strategies, we measure the area under the ten-year normalized topic-entropy curve (defined in Eq. 1), the number of occupied research directions, and total submissions.

B Research Outcomes: Impact, Resubmission, and Peer Review

B.1 Disruption Depends on Uptake

Figure 7 supports the disruption analysis in §3.2. All three citation counts in CD_p use the same observation window, from the focal paper's publication year through age three. We report CD only for papers with references and at least three papers in the denominator.

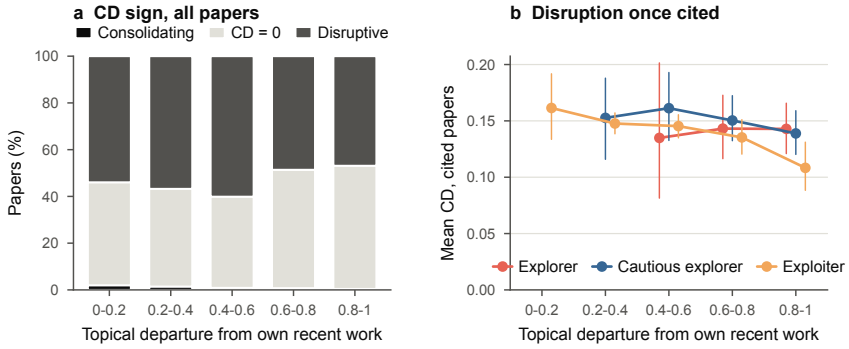


Figure 7: **Disruption tracks whether a paper is cited at all. (a)** CD of the 8,507 papers at age three, pooled over three confirmatory worlds. Almost all papers with CD $= 0$ (98.5%) have not yet been cited, and consolidating papers (CD < 0) never exceed 2.0% of a departure bin. **(b)** Mean CD among cited papers with bootstrap 95% confidence intervals. Once cited, papers from all three strategies are similarly disruptive.

B.2 Submission growth and reviewer workload under population change and resubmission

Rapid growth in conference submissions has increased pressure on peer-review systems, especially in AI and machine learning. Two mechanisms are commonly implicated: growth in the researcher population and the resubmission of rejected manuscripts (Yang et al., 2026). Although both can increase submission volume, they imply different dynamics: population growth introduces more new manuscripts and more researchers, whereas resubmission repeatedly returns existing manuscripts to review. We therefore ask whether review pressure arises primarily from growth in new submissions or from repeated resubmission.

We run a 2×2 counterfactual experiment crossing researcher population growth with resubmission, while holding the initial population, venues, literature corpus, and acceptance rate fixed. Under population growth, institutions add new researchers each year equal to 10% of the initial population. Under resubmission, rejected papers may return to review at a fixed resource cost (Appendix F). Both mechanisms increase submission volume, but their combination produces more submissions than the sum of their separate effects (Figure 5a). By year 10, compared with 193 in the baseline, annual submissions reach 364 under population growth alone and 489 under resubmission alone. When both mechanisms are enabled, submissions rise further to 867. Resubmissions account for 61% of year-10 submissions in this joint condition (Figure 5b).

Total submission volume, however, does not directly determine reviewer workload. Despite substantially more submissions, population growth alone leaves reviews per active reviewer close to

baseline (2.07 vs. 2.25 at year 10), as the larger researcher population also distributes the review workload across more reviewers (Figure 5c). With resubmission, workload instead rises to 6.64 reviews per active reviewer without population growth and 5.48 when both mechanisms are enabled. These ratios divide all annual reviews by the year-end active reviewer pool. Thus, similar increases in submission volume can have very different reviewer-level consequences depending on whether they come from new researchers or manuscripts repeatedly returning after rejection.

To quantify this persistence, we define the resubmission multiplier M_t as the ratio of total submissions to first submissions in year t. By year 10, M_t reaches 2.78 under resubmission alone and 2.59 when combined with population growth (Figure 8a), i.e., each first submission is accompanied by roughly 1.6–1.8 additional resubmissions at the system level. Among papers rejected at least once, the average number of venue attempts reaches 3.84 under resubmission alone and 3.58 in the joint condition, while 84.2% and 81.8% of rejected papers, respectively, are resubmitted rather than abandoned. Under the modeled resubmission cost, these repeated attempts are also accompanied by greater researcher attrition: without population growth, 221 of 302 researchers remain active at year 10 under resubmission versus 257 without it; with population growth, 475 versus 527 of the final 600 researchers remain active.

Reported worlds. The resubmission study uses four matched 10-year worlds crossing researcher population growth and rejection recycling, plus a separate pair comparing fixed and declining acceptance rates. Under population growth, institutions alternately add junior researchers, producing annual entry equal to 10% of the initial population. Entrants receive standard initial resources and expertise sampled from the founding distribution.

Metrics. For resubmission, we measure submission volume, submission composition (first submissions vs. resubmissions), the resubmission multiplier, cascade depth, time to publication, reviewer workload (reviews per researcher active at year end), venue redistribution, and review-score differences between first submissions and resubmissions. Each submission receives three reviews. We track the frequency and persistence of resubmission, their interaction with population growth, and the changing active reviewer population.

B.3 Additional Resubmission-Cascade Results

Figure 8 supports the resubmission analysis in §B.2.

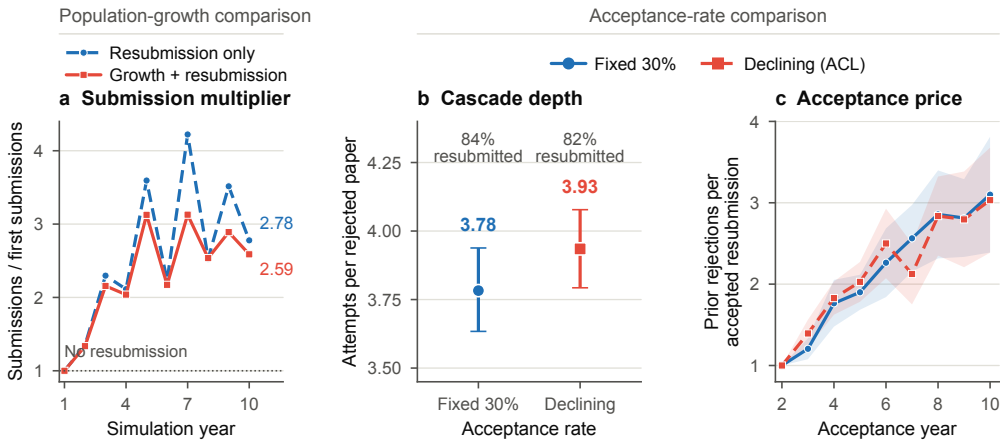


Figure 8: **Recycling multiplies submissions, and resubmission cascades deepen under stricter acceptance.** **(a)** Annual multiplier M_t: all submission events divided by first submissions. Year-ten values are 2.78 without population growth and 2.59 with growth; the dotted line marks $M_t = 1$ without resubmission. **(b, c)** A separate matched pair of ten-year worlds compares fixed 30% acceptance with declining ACL-calibrated rates. **(b)** Mean venue attempts per paper rejected at least once, before acceptance or abandonment. Percentages indicate the fraction resubmitted; bars show paper-level bootstrap 95% confidence intervals. **(c)** Mean prior rejections among resubmitted papers accepted in each year; shading shows the mean ± 1.96 standard errors. Each condition has one world; intervals do not estimate variability across worlds.

B.4 No-influx resubmission comparison

This separate comparison uses one complete ten-year world per condition with the same 1,202 founders: five researchers at each of 240 universities and two industry researchers. Condition A disables resubmission and C enables it. Neither admits new researchers. Both use Qwen3-32B, three reviews per submission, and panelized sequential funding selection from the remaining application IDs. Panels contain at most 25 applications, with acceptance rates of 0.23/0.10 for NSF/DARPA and quota $\max(1, \lfloor n_{\text{panel}} r \rfloor)$.

We measure review demand relative to the active research community as R_t/N_t, where R_t is the total number of reviews in year t and N_t is the number of researchers active at year end. We also normalize review counts by the initial population of 1,202 researchers to distinguish changes in review volume from the effects of attrition. The annual resubmission multiplier M_t is the ratio of total submissions to first submissions within the same year. Figure 9 traces these measures and submission composition over ten years.

In year ten, C records 3,660 individual reviews versus A's 1,608. Division by the year-end active reviewer pool gives 5.666 versus 1.926, while division by the full 1,202-person pool gives 3.045 versus 1.338. C has fewer first submissions (432 versus 536), plus 788 repeat submissions. C has higher review volume and both normalized measures in years 2–10, but fewer reviews in year one (1,287 versus 1,320), when neither world has any resubmission. Shared founders do not imply identical realized pre-resubmission trajectories. C's review volume and active-pool normalization peak in year six (5,904 and 6.031). Neither grows monotonically.

Over ten years, C has 6,040 first submissions versus A's 6,566, plus 7,912 repeats, yielding 13,952 versus 6,566 submission events and 41,856 versus 19,698 reviews. Each event receives three reviews, so these volumes are mechanically linked, not independent evidence of greater research production or quality. The smaller year-ten active pool in C (646 versus 835) amplifies the primary ratio, while the fixed-pool difference remains positive. These counts do not establish that reviewing caused inactivity. This pair supports a descriptive review-demand contrast, not an influx interaction, an isolated funding mechanism, or independent world-level uncertainty estimates.



Figure 9: **Annual submission and review demand in the A/C comparison. (a)** Paired annual bars (A left, C right) separate first submissions (solid) from repeats (hatched). Bar heights count all submission events. **(b)** Year-end reviewer pool N_t, counting researchers satisfying both `is_active` and `can_review`. **(c)** All individual reviews R_t divided by the active pool (solid lines) or the full 1,202-person pool (dashed lines), on a shared axis. The numerator includes reviews by researchers inactive at year end, so active-pool ratios are not mean assigned workload among survivors. **(d)** Annual total/first-submission multiplier M_t. Each submission event receives three reviews, so review demand is mechanically linked to the submission totals in (a). Points are saved annual observations from one ten-year world per condition (no influx). Lines connect observations without smoothing. No across-world uncertainty interval is available from this pair.

B.5 Details of the Author-Information Experiment

Reported worlds. The network-proximity study generates valid collaboration networks and replays matched peer-review decisions under controlled information conditions.

Review conditions. We separate the effect of revealing author information from the additional effect of explicitly disclosing the reviewer–author relationship in the collaboration network. Each reviewer evaluates the same paper under three conditions:

1. **Condition 1: Blind.** No author name, institution, or network relationship is shown.
2. **Condition 2: Author information.** The first author's name, represented by a stable pseudonymous identifier, and institution are shown, but the network relationship is not disclosed.
3. **Condition 3: Author information plus relationship.** The same author information is shown alongside a neutral statement of the actual network relationship, such as sharing an intermediary collaborator or belonging to disconnected components.

Across conditions, we hold fixed the paper's title, abstract, and topic labels, the reviewer's expertise, the scoring instructions, and the 1–5 scale.

Papers and reviewer matching. For each paper, we select three reviewers with similar expertise but different relationships to the first author: a distance-2 reviewer shares a collaborator with the author (distance 2); a distance-3 reviewer is connected through two intermediary researchers; and a stranger is connected to the author through a shortest-path with distance ≥ 4 or in a disconnected component. We exclude direct collaborators, coauthors, researchers affiliated with any author's

institution, and researchers with recorded conflicts of interest. Within each triplet, paper–reviewer cosine similarities differ by at most 0.10, and each reviewer is assigned at most ten papers.

Collaboration network. Network distances are computed using only collaborations formed before each paper's review year. The source simulations allow both new and repeated partnerships: in each agent-year, collaborator search targets new partners with probability 0.70 and existing partners otherwise. New-partner scores combine topic similarity, recent-publication overlap, log degree, and shared collaborators, with respective weights of 0.40, 0.30, 0.20, and 0.10. The resulting network supplies the relationships used for reviewer matching and disclosure.

Score comparisons. For network-proximity bias, the primary outcome is the paper-level matched score shift caused by disclosing the reviewer's true network relationship to the author. We compute two within-reviewer, within-paper differences. The *author-information effect* is the score under author-information disclosure (Condition 2) minus the score under blind review (Condition 1). The *additional relationship effect* is the score under author-information-plus-relationship disclosure (Condition 3) minus the score under author-information disclosure alone (Condition 2). For each effect, we first average the three reviewers' score changes within each paper, then average across papers. We report changes on the 1–5 scale, with 95% confidence intervals from $10{,}000$ bootstrap resamples of whole papers.

We also subtract the stranger reviewer's relationship-induced score change from that of the reviewer who shares a collaborator with the author. This tests whether disclosing a closer relationship produces a larger score change.

Results and interpretation. Revealing author information alone increased the mean score by 0.055 points (95% CI $[0.014, 0.097]$). Once author information was visible, adding a description of the relationship had little average effect, with a mean score change of 0.000 points (95% CI $[-0.038, 0.037]$). Reviewers who shared a collaborator with the author did not show a detectably larger response to relationship disclosure than stranger reviewers.

C Institutional Mechanisms: Funding Capacity and Allocation

C.1 Higher research output expands publication volume faster than ecosystem capacity

Reported worlds. The output-scaling study uses four university-only worlds with 300 researchers each, crossing one or two submissions per completed project with scaling or fixed grant capacity. Two mixed-population worlds, each with 150 university and 150 industry researchers, compare the submission limits under fixed capacity. All six worlds run for ten years. Scaling capacity applies each funding program's configured award rate to its application volume; fixed capacity reserves annual award slots equal to 20% of the initial university population, giving 60 slots in university-only worlds and 30 in mixed worlds. The mixed worlds also hold the industry compensation pool fixed.

Output intervention. The intervention varies the submission limit per completed project while retaining the same cost settings, project-duration rules, and manuscript-selection procedure. Table 2 reports the resource parameters.

Metrics. For output scaling, active fractions use the initial population as the denominator. Grant coverage and award-count inequality include all initial university researchers, including those inactive by year ten; inequality among recipients conditions on receiving at least one award. Three-year citation coverage is the share of accepted papers receiving at least one citation from a paper with a disjoint author set within three years, restricted to papers accepted by year seven. Topic entropy is rarefied to a common accepted-paper sample size within each population comparison. Citation opportunities also depend on submission volume and the size of the accepted-paper pool, so these citation measures are descriptive rather than direct estimates of scientific quality.

C.2 Narrow funding wins show no detectable increase in later funding

Prior work compares applicants near grant cutoffs to study how early funding success or failure shapes later outcomes (Bol et al., 2018; Wang et al., 2019). We use SUTO to trace these effects

over three years following researchers' first application year. We call applicants ranked just above a panel's funding cutoff *narrow winners* and those just below it *narrow losers*. Their adjacent panel ranks provide a local comparison of subsequent funding, publications, citations, and research choices.

Narrow winners show no detectable increase in subsequent funding (Figure 10a): the estimated effect is -1.50 resource units (95% CI $[-14.5, +11.5]$).

Publication output, citations, future award probability, and exploration likewise show no detectable differences (Figure 10b). All near-cutoff applicants reapply, and none exits the simulation during follow-up. The comparison follows applicants who remain in funding competition throughout the three-year window.

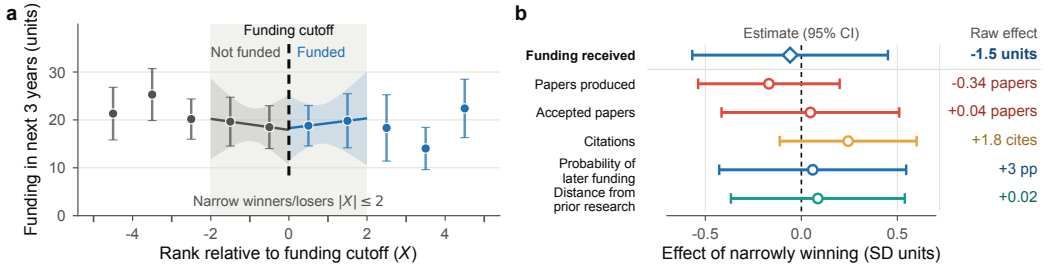


Figure 10: **Narrow funding wins show no detectable advantage over the next three years. (a)** Applicants just above the funding cutoff (*narrow winners*) receive similar subsequent funding to those just below it (*narrow losers*), with no clear jump at the cutoff. Points show mean funding at each rank. Lines are descriptive local fits, and shading marks the near-cutoff comparison window ($N = 444$ applications). **(b)** Estimated effects on later funding, publications, citations, award probability, and exploration all have confidence intervals spanning zero. Colors group outcome families, and the diamond marks the primary funding outcome. Effects are scaled by the narrow-loser standard deviation, with raw-unit estimates at right (pp: percentage points). Error bars and bands show 95% confidence intervals.

Reported worlds. The funding-cutoff study uses three 500-agent, six-year worlds and records within-panel ranks for regression-discontinuity analysis.

Estimation and diagnostics. The running variable is $x = k - \text{rank} + 0.5$, where k is the panel's award count, so winners have $x > 0$. We fit a local-linear model with separate slopes on either side, triangular weights, and competition-panel fixed effects. The primary bandwidth $h = 2$ is selected by leave-one-competition-out cross-validation on pre-application publication count. The sample contains 444 applications from 265 applicants in 183 competitions. Confidence intervals use two-way cluster-robust standard errors by applicant and competition, with finite-sample corrections. Bandwidth checks at $h \in \{2, 3, 4, 6, 8\}$ yield funding estimates from -1.50 to $+0.47$, all with intervals spanning zero. Balance checks detect no cutoff differences in prior paper or acceptance counts, concurrent applications, or strategy indicators.

Metrics. For funding RD, the primary outcome is cumulative subsequent funding over the next three simulated years; secondary outcomes include publication, acceptance, citation, reapplication, attrition, and exploration behavior.

D Emergent Ecosystem-Level Dynamics

D.1 Resource Inequality and Researcher Attrition

Figure 11 supports the ecosystem-level analysis in §3.4.

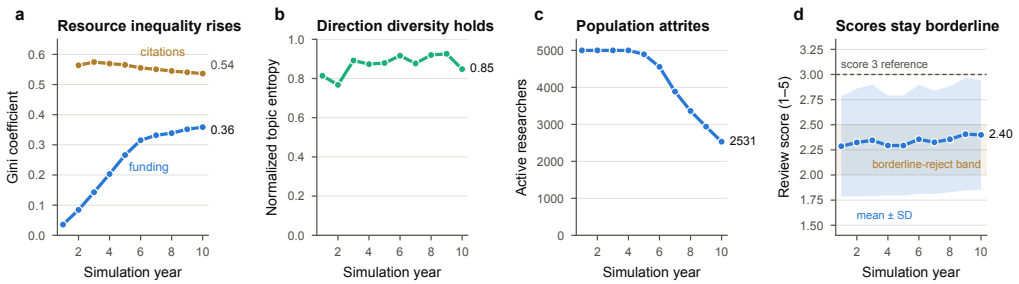


Figure 11: **Resource inequality compounds while topic diversity holds.** Pooled over three confirmatory worlds (1,000 institutions, 5,000 researchers, ten years); lines show means and bands span the minimum–maximum across worlds. **(a)** Funding-resource Gini (rising) with the citation Gini overlaid. **(b)** Normalized topic entropy, where 1 denotes research directions being occupied evenly. **(c)** Number of still-active researchers. **(d)** Mean individual review score with a ± 1 SD envelope. Scores remain between 2.29 and 2.40 on the 1–5 scale (median 2.0 every year), while the envelope widens slightly. The dashed line at 3 is a score reference. Acceptance follows the ranking-and-quota rule in §2.

D.2 Self-citation does not explain citation inequality

Because agents may cite their own prior work, we test whether self-citation contributes materially to the citation inequality observed above. Across the three large confirmatory worlds spanning 1,000 institutions, 5,000 researchers, and ten years, only 308 of 75,661 citation edges (0.41%) are self-citations (Figure 12). Meanwhile, self-citation differs by research strategy. Exploiters self-cite in 0.57% of outgoing citation edges, compared with 0.13% for cautious explorers and around 0.0% among explorers (0 of 10,861 edges). This ordering holds in all three worlds. The pattern is consistent with topical continuity: exploiters remain close to their prior work, whereas explorers move farther away and therefore have fewer relevant prior papers of their own to cite. Despite these differences, self-citation is too rare to explain the ecosystem's citation concentration, indicating that citation inequality in SUTO arises from broader citation-allocation dynamics.

Citation Inequality Persists After Removing Self-Citations

Figure 12 supports the self-citation analysis in §3.4.

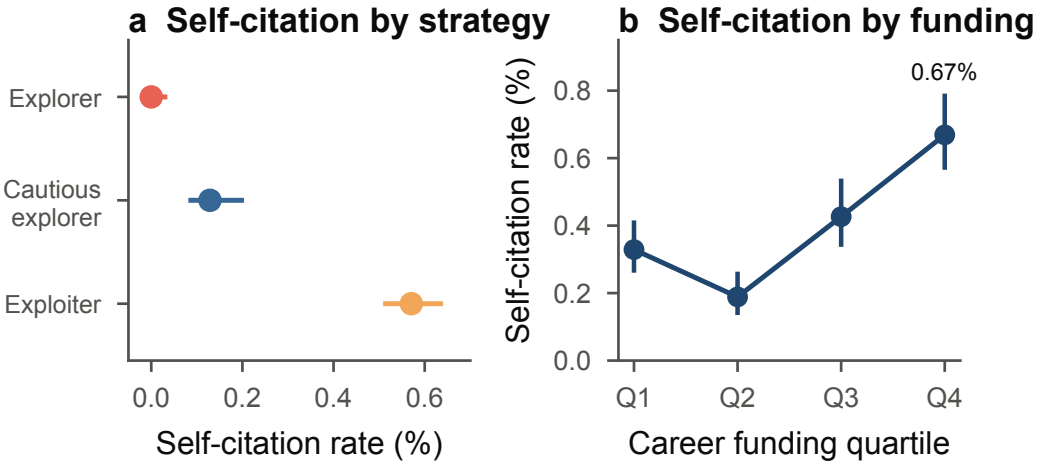


Figure 12: **Self-citation across strategies.** Across three confirmatory worlds (75,661 citation edges), exploiters self-cite more frequently than their peers. Removing all self-citation edges changes the citation Gini by at most 0.0005 in any world.

E NOTATION

Table 1 lists the symbols used in the paper, grouped by where they are introduced.

Table 1: Notation used in the paper.

Symbol	Meaning
Researchers and budgets (§2)	
i, j	Researcher indices (j also denotes a candidate collaborator or a reviewer)
t	Simulated year
$s_i^t = (E_i, \rho_i^t, b_i^t, M_i^t, C_i^t)$	State of researcher i in year t: expertise topics, reputation, available budget, memory, and conflict-of-interest set
f_i^t	Funding received by researcher i in year t
$\mathcal{A}_i^t$, $\mathrm{cost}(a)$	Actions taken by researcher i in year t and the cost of action a
Venues, directions, and collaboration (§2)	
c, α_c, S_c	A conference, its acceptance rate, and its submissions in a year
$d \in \mathcal{D}$	A research direction from the fixed set $\mathcal{D}$ ($\lvert\mathcal{D}\rvert = 53$)
T_d, κ	Duration of direction d in years and its budget cost per active year
$\psi(i,j)$, $\phi_m(i,j)$, w_m	Collaboration score of the pair (i,j), its m-th feature, and the feature's weight
p, R_p	A paper and its reviewer panel
u_{jp}, $\bar{u}_p$	Recommendation of reviewer j on paper p and the panel mean used for the decision
n, n_{panel}, r, k	Panel application count (n or n_{panel}), program funding rate, and award count under rate-based funding, $k = \max(1, \lfloor rn \rfloor)$ for nonempty panels
e_i, $\tilde{e}_i$, λ	Researcher i's normalized panel-rank score $e_i \in [0,1]$, adjusted score $\tilde{e}_i = e_i - \lambda q_i$, and penalty strength $\lambda \geq 0$
q_i	Percentile of researcher i's mean paper distance among that year's applicants with scored papers, using all recorded papers. Applicants without scored papers receive $q_i = 0$
Evaluation (§3)	
$\mathbf{p}$, $\mathbf{a}_i^t$	Paper embedding and researcher i's mean paper embedding from years $t-4$ through $t-1$, with initial expertise as the fallback
$\nu_{p,i}^t$	Paper distance $\nu_{p,i}^t = 1 - \cos(\mathbf{p}, \mathbf{a}_i^t) \in [0,2]$ (§3.1)
ED	Mean annual distance to the author's most recent paper from an earlier year, excluding years without a comparison (§3.1)
CD_p; n_{focal}, n_{both}, n_{prior}	CD index and its three citation counts for focal paper p, observed at age three: focal-only, focal-and-reference, and reference-only citing papers (§3.2)
H_t, $p_{t,d}$	Normalized topic entropy in year t and the share of that year's direction choices selecting d, with $0 \log 0 = 0$ (Eq. 1)
AUC_H, T	Trapezoidal area under H_t from year 1 to T, bounded by 0 and $T-1$ (Eq. 1)
M_t	Resubmission multiplier: total submissions divided by first submissions in year t (§B.2)
R_t, N_t	All individual review events in year t and the year-end active reviewer-capable population, used in R_t/N_t (Appendix B.4)
N	Number of observations in an analysis sample

F EXPERIMENTAL DETAILS

F.1 DATA

To ground the simulation in realistic scientific literature, we use SciEvo (Jin et al., 2024a) as our literature corpus. SciEvo is a longitudinal scientometric dataset containing 2.1 million arXiv papers published from 1991 to 2025, spanning eight broad subject groups and 156 categories. It provides paper titles, abstracts, publication timestamps, subject categories, and LLM-extracted keywords, together with citation relations, author information, and publication venues retrieved from Semantic

Scholar. For our experiments, we select computer science papers within our 2016–2025 simulation window and organize them by publication year.

F.2 Model and literature grounding

All simulations use Qwen3-32B (Yang et al., 2025) with thinking mode. Our experiments used NVIDIA A100 GPUs with 80 GB of memory. Submissions are based on retrieval grounded in SciEvo. Given an agent's selected research direction and submission cycle, the agent retrieves contemporaneous arXiv papers and selects a topically compatible paper and venue. Simulated years correspond to 2016–2025, with year-based filtering of the retrieved corpus determining which papers are available for submission and citation. Unless otherwise specified, conferences use a fixed 30% acceptance rate to ensure comparable selectivity across counterfactual worlds. The resubmission-cascade study (§B.2) uses this regime, while Appendix B.3 additionally evaluates a schedule based on observed yearly acceptance rates.

All families use the same set of research directions $\mathcal{D}$.

F.3 Implementation settings

Table 2 records the resource settings used in the reported studies. The capacity treatments, population compositions, and numbers of worlds are specified with each experiment.

Table 2: Resource and funding settings. Budgets and costs are in abstract resource units.

Setting	Value and scope
Uniform initial budget	100 per researcher in the strategy, composition, funding-cutoff, and university-only output studies
Default initial endowments	University: 80 or 160. Industry: 160 or 320, depending on institution tier
Research cost	10 per active project-year
Project duration	1–3 years, determined by the selected direction
Resubmission fee	5 per attempt
Grant amount and rate	20 per award. Foundation-oriented programs fund 23% and mission-oriented programs 10% of panel applicants under rate-based allocation, subject to the quota rule in §2
Funding-conservatism penalty	Effective $\lambda = 0$ in the reported experiments

Retrieval and semantic distance. Retrieval uses normalized `thenlper/gte-small` embeddings of abstracts and queries to rank eligible candidates by cosine similarity. A submission query combines the agent's research intention and direction keywords. Retrieval returns up to 1,000 unsubmitted candidates from the current calendar year, and the selection prompt displays the top-k with their titles, abstracts, and topics. Citation selection retrieves previously accepted papers in the same world. Exploration and paper-distance metrics use the same embeddings.

Decision inputs and review rubric. Research prompts condition on expertise, strategy, resources, and available directions. Paper-selection prompts include the research focus, candidate artifacts, venue topics, and recent paper history. Review prompts include the paper's title, abstract, and topic labels, reviewer expertise, and memory. The rubric requests a 1–5 score and a one-sentence justification (decimals allowed). It also asks reviewers to place more than 60% of papers at 2.0–2.5 or below. Actual acceptance follows within-venue ranking and quotas. Funding evaluation uses applicants' proposals, expertise, paper records, and program priorities.

Decoding and memory. Qwen3-32B runs in thinking mode with temperature 0.7 and structured JSON responses for decision prompts. The default completion budget is 2,048 tokens, with an 8,192-token request for batched funding rankings. Memory is an append-only record of dated events, feedback summaries, and decision rationales, saved with researcher state. When memory is included in a prompt, entries are concatenated in order.

F.4 Motivation for closed-loop simulation

The loop is closed by carrying updated researcher states and scientific records into subsequent decisions. Rejected manuscripts and their reviews remain available for later resubmission decisions, while accepted papers enter the citable literature and the authors' track records used in funding evaluation. Funding and action costs update researchers' budgets, determining who can remain active and hence contribute submissions and reviews in later years. An intervention can therefore change both the outcomes of the current round and the population, resources, and literature on which later rounds operate.

LLMs provide context-dependent choices and evaluations within this loop, including research-direction and paper selection, peer review, and funding assessment. Explicit institutional rules govern acceptance quotas, resource accounting, and attrition. This separation allows institutional mechanisms to be varied while retaining the same agent decision framework; the resulting counterfactual trajectories describe their downstream consequences under the specified simulation assumptions.

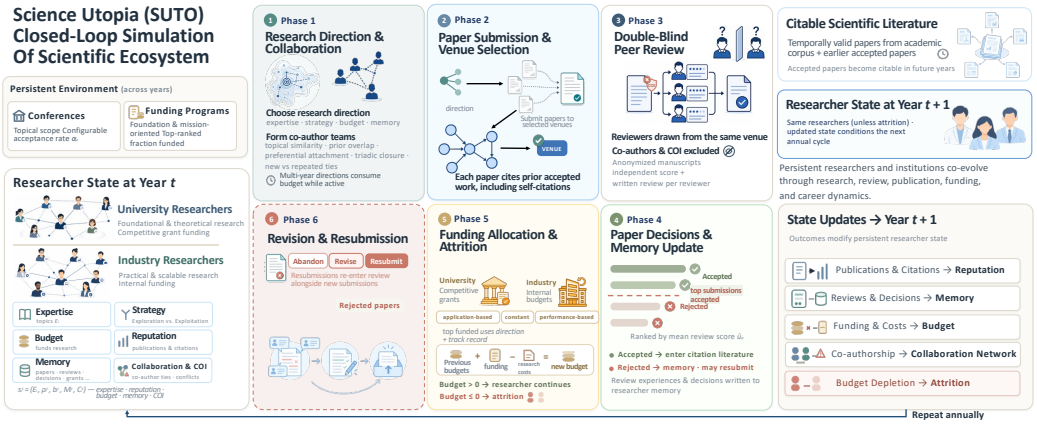


Figure 13: **Overview of SUTO's LLM-powered research ecosystem.** SUTO models scientific research as a closed-loop ecosystem where research-direction choice, collaboration, submission, peer review, resubmission, citation, and funding are connected through annual feedback loops, with outcomes updating persistent researcher states and shaping subsequent decisions.

F.5 Simulation corpus and released dataset

Across the 61 reported worlds, SUTO simulates over 40,000 researcher agents affiliated with 8,000 institutions over ∼390,000 researcher-years. These runs produce ∼177,000 submitted manuscripts, ∼400,000 accept or reject decisions, and ∼1.2M individual peer reviews with written justifications. The simulations also record ∼595,000 researcher memory entries, including ∼194,000 research-direction decisions with written rationales. We plan to release this corpus, covering per-year snapshots of researcher states, submissions, reviews, decisions, citations, memory records, and funding outcomes, as a resource for research on peer review, research evaluation, and the science of science.

G Broader Impact

LLMs are increasingly involved throughout scientific workflows, including writing, reviewing, and research decision-making (Kobak et al., 2025). Their impact on science, however, may not be captured by evaluating these tasks in isolation. Repeated interactions among researchers, institutions, funding mechanisms, and publication systems can create feedback loops whose long-term consequences only emerge at the ecosystem level.

Our framework provides a controlled testbed for studying these dynamics. By simulating persistent researchers and institutions under explicit rules, it enables counterfactual experiments on mechanisms that are difficult to isolate in real scientific systems, revealing how seemingly local changes—to review criteria, funding incentives, research strategies, or AI-mediated behavior—can accumulate into system-level changes in workload, resource allocation, scientific impact, and diversity. More

broadly, we view such ecosystem simulation as a complementary tool for anticipating how increasingly AI-mediated science may evolve and for studying the institutional mechanisms that shape its long-term trajectory.