The Paper and the Room: Artificial Intelligence and Inequality at Academic Conferences

Ying Hoi Kwan
University of Nottingham

July 6, 2026

Abstract

Will artificial intelligence democratise academic careers? The answer, we argue, turns on a feature of conferences that is easily missed: a conference is at once a place to have a paper improved and a place to stand in the room with the senior scholars who shape careers. Artificial intelligence transforms the paper—and, because it substitutes for the mentoring the disadvantaged lack, helps them most—but it cannot manufacture the room. The two are complements, and by the logic of unbalanced growth the effect inverts the democratising hope: cheapening the paper raises rather than lowers the room's relative worth, so the returns migrate to the one good artificial intelligence cannot supply—physical proximity to the people who cite, referee, and hire. We study this shift in the economics of science: a decision-theoretic model of how researchers choose which conferences to target, carried to a rational-expectations equilibrium, with the empirically relevant case located by eliciting adoption behaviour from large language models. Four results follow. The distributional effect of artificial intelligence is a *knife-edge*, set by the correlation between capability and prestige: mildly equalising when access is broad, sharply widening when access tracks standing. The access the technology appears to open is an *enclosure*, competed away as adoption spreads, so that non-adopters are crowded out of the dense rooms though their work is no worse. The obvious remedy—competing more strategically for the room—*backfires*, widening the gap it was meant to close. Most striking, artificial intelligence can leave measured inequality unchanged while rebuilding its *composition*: at the calibrated operating point the level of inequality is indistinguishable from the world before the technology, yet the advantage of position gives way not to talent but to luck. Judged by an inequality

index, artificial intelligence appears to change nothing; judged by what decides who succeeds, it changes almost everything.

1 Introduction

A recurring hope about artificial intelligence is that it will democratise expert work. If a large language model can draft, polish, and critique a research paper at almost no cost, the advantages of a well-resourced environment—a tireless and wide-ranging supervisor, a deep bench of colleagues, the accumulated craft of an elite department—might be placed within everyone's reach, and the stratification that characterises academic science might begin to dissolve. The hope has an empirical footing: in settings from professional writing to customer support, artificial-intelligence assistance raises the output of less-expert workers far more than that of experts (Noy and Zhang, 2023; Brynjolfsson et al., 2025). If the technology compresses skill gaps at the level of the task, the thought runs, perhaps it will compress them at the level of the career.

This paper argues that the hope is half right, and that the missing half is decisive. The argument turns on what a research conference—the institution through which early-career scholars enter a field—actually provides. A conference does two things at once. It is a place to have a paper read and improved, and it is a place to stand in the same room as the senior scholars who shape careers: those who cite, co-author, referee, and hire. Artificial intelligence transforms the first of these and leaves the second untouched. Improving the paper is a task that software now performs cheaply and, because it substitutes for the mentoring that the disadvantaged lack, performs best for those who began with least. Securing the attention of a busy senior scholar, in the one session a researcher can physically attend, is not a task but a scarce and congestible good, rationed by physical proximity and by the prestige-linked capacity to turn a brief encounter into a citation or a job. The paper has become cheap to produce; being heard has not. It is this asymmetry—cheap to submit, costly to be heard—that organises everything that follows.

This is, at its core, a question in the economics of artificial intelligence, and the logic that answers it is an old one. Digital technologies advance by relentlessly cheapening a particular set of costs (Goldfarb and Tucker, 2019); what governs their consequences is rarely the cheapened input itself but its scarce *complement*. When one input to a joint product turns abundant, the return to the whole comes to rest on whatever input remains scarce (Kremer, 1993)—cheapen the writing of a paper, and the binding constraint, with the reward that attaches to it, shifts to the attention the paper must still win. This has a precise and underused analogue in Baumol (1967). The paper side of a conference is a *progressive*

sector, its productivity transformed by artificial intelligence; the room is a *stagnant* one, a senior scholar's hours no more plentiful than before. Baumol's lesson is that when the progressive sector races ahead, the *relative* value of the stagnant one rises rather than falls: artificial intelligence does not devalue the room but makes it the most valuable thing a conference has left. We pursue this logic within the economics of science (Dasgupta and David, 1994), asking what becomes of the reputation-based reward system on which science runs once the means of producing the paper are nearly free.

That the equilibrium is already shifting is visible in the institutions themselves. Documenting internal editorial data, Gartenberg et al. (2026) report that submissions to a leading management journal rose by roughly two-fifths after large language models became publicly available, even as average writing quality fell—a drift toward "more rather than better" work. When the cost of preparing a submission collapses, more submissions follow; but the senior attention that a submission ultimately competes for does not multiply. The bottleneck simply moves downstream, from producing the paper to being heard above the rising noise.

The asymmetry sharpens the question we ask. Rather than asking whether artificial intelligence will level the academic field, we ask: given that it levels the paper but not the room, does artificial intelligence narrow or widen the inequality of academic careers—and, if it widens that inequality, can the researchers it disadvantages compensate by competing more strategically for the scarce room? And beneath it lies a subtler question that a summary measure of inequality cannot answer: even where the level of inequality is unchanged, has artificial intelligence altered what that inequality is made of—the parts played by talent, by standing, and by luck in deciding who gets ahead? The question matters beyond science. The structure we study—a task whose cost artificial intelligence collapses, sitting beside a complementary form of access that remains physically scarce—recurs wherever advancement depends jointly on producing work and on proximity to the people who allocate rewards. But it is especially stark, and especially tractable, in science, where the conference is a designed institution that a field can in principle choose to organise differently.

We answer with a model rather than a dataset, for a reason the empirical record cannot yet settle: the question concerns an equilibrium that artificial intelligence is only beginning to produce. We build a decision-theoretic model of how a single doctoral researcher chooses which conferences to target, and then carry a population of such researchers to a rational-expectations equilibrium in which the selectivity each anticipates coincides with the selectivity the market produces. The individual problem is a well-posed expected-utility choice: a researcher with an imperfectly known paper selects a portfolio of submissions under correlated acceptance risk, anticipating which offers to accept and which to decline. Capability with artificial intelligence is modelled as a weakest-link technology—the least of access, skill,

and institutional permission—so that subsidising a tool a department still forbids changes nothing (Kremer, 1993; Bresnahan et al., 2002). The room is modelled as a congestible commons: the chance of meeting a relevant senior scholar rises with the density of such scholars in the session one attends and falls with the number of parallel sessions, and the contact a venue can supply is capped, so that crowding erodes it (Falkinger, 2008; Gartenberg et al., 2026). Prestige enters not as a separate force but through the very channels artificial intelligence might displace—sharper internal feedback and a higher rate of converting contact into outcomes—reflecting the cumulative advantage long documented in science (Merton, 1968; DiPrete and Eirich, 2006). We solve the model as an agent-based simulation, in the tradition that lets a striking aggregate pattern emerge from many interacting agents who begin, in the dimension that matters, from equal footing (Pluchino et al., 2018).

Four results follow, and none is assumed. The first is a *knife-edge*. The distributional effect of artificial intelligence is not a fixed quantity; it depends on a single feature of the world—how capability is distributed relative to prestige. When access to capable tools is independent of institutional standing, artificial intelligence is mildly equalising, because its paper-side levelling helps the disadvantaged most. When capability instead tracks prestige—when elite institutions also command the better access, skill, and permission—the two advantages compound, and the gap between elite and non-elite careers more than doubles. Strikingly, this widening occurs even as average capability in the population falls, because entrenchment concentrates the technology in a privileged minority: what drives inequality is not how much artificial intelligence is used but who uses it.

The second result is an *enclosure*. As adoption spreads, the researchers who do not adopt are progressively crowded out of the dense rooms—their access to elite contact collapses and the acceptance bar at the selective venues climbs—even though the quality of their papers is unchanged. Their work is no worse; they are simply shut out. The access that artificial intelligence appears to open is competed away as more researchers use it to chase the same scarce seats.

The third result is a *backfire*, and it concerns the obvious remedy. The natural advice to a disadvantaged researcher is to be strategic—to target the small, dense workshops where the room is richest rather than the large generalist conference. We find that offering this option widens the very gap it was meant to close. Because strategic targeting sends everyone toward the same scarce rooms, and because the artificial-intelligence-equipped clear the resulting higher bar and capture the rationed contact, the adopter–non-adopter gap is larger when researchers compete strategically than when they do not, at every level of adoption. The room cannot expand to absorb the extra competition, so competing harder for it rewards those already ahead.

The fourth result is a *recomposition*, and it is the most striking because it hides from the statistic most often used to detect it. When we calibrate the model to where the world appears to sit—eliciting, from large language models standing in for a population of researchers, how readily scholars would use artificial intelligence against an institutional ban or a regional restriction—the level of inequality barely moves: at that operating point the dispersion of career outcomes is statistically indistinguishable from the world before artificial intelligence. Yet beneath that unchanged number the composition of success has been rebuilt. The share of outcomes governed by *position*—the prestige-linked advantage of sharper feedback and easier conversion of contact—falls by roughly half, while the share governed by *luck*, the chance of who is in the room, rises to dominate; the share owed to *talent* stays small throughout. Artificial intelligence holds the level of inequality fixed while migrating its source from a hierarchy one can name to a lottery no one controls. Measured by an inequality index alone it would seem to have changed nothing; measured by what decides who succeeds, it changes almost everything.

The paper makes four contributions. To the economics of science in the tradition of Dasgupta and David (1994), it carries the reputation-based reward system into the era of cheap production, supplying a mechanism—the complementarity between the levelable paper and the congestible room—that reconciles two facts otherwise in tension: that artificial intelligence demonstrably compresses skill gaps on tasks (Noy and Zhang, 2023; Brynjolfsson et al., 2025), yet need not compress, and may widen, inequality in careers. To the study of artificial intelligence and inequality, it shows that the technology's distributional sign is a property not of the technology but of its allocation, reframing a yes-or-no debate as a conditional one and identifying the doubly disadvantaged—non-elite and low in capability—as the population of concern. To the design of scientific institutions, it issues a counterintuitive warning: the intuitive responses to artificial-intelligence-driven crowding, whether encouraging strategic venue choice or trusting that cheaper participation will broaden access, can entrench the hierarchy they were meant to loosen, because the binding constraint is the finite room, not the now-abundant paper. To the measurement of artificial intelligence's distributional effects, it adds a caution with reach beyond this application: because the technology can leave the level of inequality untouched while transforming its composition, a summary index of dispersion can register no change precisely where what determines success has been most rewritten. Methodologically, the paper joins decision-theoretic discipline at the level of the individual to an agent-based equilibrium at the level of the population, so that its distributional claims rest on the choices of agents who each solve a well-posed problem rather than on reduced-form behavioural rules.

These results are the output of a calibrated model, not an empirical estimate, and we

treat them as such. The contribution is a set of mechanisms and the testable predictions they generate—chiefly, that the distributional effect of artificial intelligence on academic careers should vary with the correlation between tool access and institutional prestige, and that the selectivity of small workshops should rise with adoption while the measured quality of those it excludes does not; and that, as adoption deepens, institutional pedigree should lose power to predict early-career success relative to less systematic variation—the signature of a hierarchy giving way to a lottery. We return to these predictions, and to what would falsify them, in the conclusion.

The remainder of the paper proceeds as follows. Section 2 situates the argument in the economics of digitisation and of science, in work on artificial intelligence and inequality, and on the role of conferences and physical proximity in scientific careers. Section 3 develops the model and states the results as propositions. Section 4 solves the model as a rational-expectations equilibrium and reports the four central findings together with their robustness. Section 5 draws out the implications for doctoral training, conference design, and the governance of artificial-intelligence access, and concludes.

2 Related Literature

This paper is a contribution to the economics of science (Dasgupta and David, 1994) in the age of artificial intelligence, and it approaches that subject through two lenses. The first is the economics of digitisation, which asks how a technology that drives a particular set of costs toward zero reshapes the activity built upon them (Goldfarb and Tucker, 2019). The second is the logic of complementarity—the proposition, static in Kremer (1993) and dynamic in Baumol (1967), that when a technology makes one input abundant the return migrates to the scarce input it cannot supply. Together they locate our question precisely: what becomes of a reputation-based reward system when the means of producing the paper, but not the means of being heard, fall to near zero.

This paper draws on, and connects, three bodies of work: the economics of academic dissemination and the submission decision; the causal value of conferences and temporary colocation in scientific careers; and the emerging evidence on how generative artificial intelligence reshapes knowledge work and the inequalities that attend it. The motivation is a development that all three now confront. Documenting internal editorial data, Gartenberg et al. (2026) report that submissions to a leading management journal rose by roughly two-fifths following the public release of large language models, even as average writing quality fell and a growing share of both manuscripts and referee reports bore the marks of machine generation; they read this as a drift toward "more rather than better" research. We take

that finding as a point of departure. The congestion observed at the journal is, we argue, the aggregate footprint of individual researchers re-optimising their dissemination strategy once the cost of preparing a submission collapses. Our object is that individual decision, applied to the choice of conferences. What the surge cannot expand is the scarce resource on the other side of the ledger: generative models multiply the supply of manuscripts at near-zero cost but cannot multiply the human attention required to evaluate it, a constraint that binds the reviewing commons and, equally, the finite attention of senior scholars in a conference room (Falkinger, 2008). The asymmetry between a cheaply expandable supply of papers and a fixed supply of attention is the thread that runs through what follows.

The economics of dissemination and the submission decision. Where to send one's work has long been modelled as sequential search under uncertainty. Oster (1980) characterises the optimal order in which to submit a manuscript to a ranked set of outlets—the seminal statement of submission as an ordered, one-at-a-time gamble. Ellison (2002a) documents that the process governing these gambles has slowed markedly over decades, as escalating revision demands lengthen the path from submission to print, and his companion theory of evolving quality and refereeing standards (Ellison, 2002b) rationalises that slowdown as a shift in the standards outlets impose. Two features of this tradition are central to what follows. First, it studies *journal* submission, where the terminal payoff is publication; we study *conference* selection, where the payoff is a bundle of feedback and professional contact rather than a line on a publication list. Second, and more consequentially, the sequential logic depends on a large per-submission friction: because adapting and submitting a paper is costly, the researcher targets one venue, waits, and only then moves on. Once that friction falls toward zero—as it does when a manuscript can be re-tailored to a venue's framing and format at negligible cost—the ordered gamble gives way to a simultaneous portfolio problem, in which the researcher submits broadly at once and selects ex post among the acceptances. That transition is the analytical heart of our model.

Artificial intelligence as a search and adaptation technology. A recent literature in the economics of innovation models artificial intelligence as a technology that reorganises search. Agrawal et al. (2024) represent AI-assisted discovery as *prioritised search* over a vast combinatorial space, in which a predictive model reorders where researchers look and so compresses the expected effort to a result; relatedly, Agrawal et al. (2023) consider how progress toward highly capable AI may alter the rate and direction of the knowledge it helps produce. This work locates AI's effect at the *discovery* stage—which hypotheses to pursue. Our contribution carries the same logic to the *dissemination* stage. The mechanism that

makes AI powerful in our setting is not hypothesis generation but the near-elimination of adaptation cost: when tailoring a paper to a venue is cheap, the researcher's search over venues widens, just as predictive search widens the researcher's reach over hypotheses. The two are complementary descriptions of a single shift in which AI lowers the cost of exploring a large space of options.

Conferences, colocation, and the value of the room. The value of a conference is not exhausted by the dissemination of the manuscript; a substantial body of credible empirical work shows that physical presence carries causal returns of its own. Exploiting the cancellation of a major political-science meeting by a hurricane as a natural experiment, Lopez de Leon and McQuillin (2020) find that presenting a paper raises its subsequent citations, isolating an effect of presentation distinct from selection. The mechanism is face-to-face contact and the ties it forms. Studying attendees across a series of Gordon Research Conferences, Chai and Freeman (2019) show that temporary colocation causally raises new collaboration, with junior researchers among the principal beneficiaries. Catalini (2018) and Boudreau et al. (2017) supply the microgeographic logic behind these returns: colocation and reduced search costs sharply raise the probability that two researchers begin to work together, with the largest effects where finding one another was otherwise costly. Crucially, the returns are conditional on the intensity of participation rather than mere attendance—Baruffaldi and Poege (2025) show that firms learn from scientific conferences only when they engage closely, through the personal connections of their researchers. Two implications matter for our model. The returns to the room are real and causal, and they are not reducible to access to information—precisely the thing artificial intelligence supplies cheaply. Physical proximity to the right people is therefore the component of a conference that AI cannot substitute for, and it is the component our model places at the centre of the attendance decision.

Generative AI, productivity, and the leveling of skill. On the paper side, the evidence on generative AI and knowledge work is more settled, and points consistently in one direction. In a pre-registered experiment with several hundred professionals, Noy and Zhang (2023) find that access to a large language model raised output quality and cut completion time while compressing the distribution of performance: the largest gains accrued to initially weaker writers. Brynjolfsson et al. (2025), studying the staggered roll-out of an AI assistant to several thousand customer-support agents, document the same pattern in the field—sizeable average productivity gains concentrated among less experienced and lower-skilled workers, with little effect on the most able. This is the empirical basis for treating AI

feedback as a partial *substitute* for scarce human expertise: where a researcher's own support is weak or narrow, AI raises the quality of the work most; where it is already strong and broad, AI adds little. We import this leveling pattern into the feedback channel of the model, but cautiously: the evidence concerns professionals and service workers rather than scientists or doctoral students, so we treat scientist-specific leveling as a modelled hypothesis rather than an established fact. The same literature also disciplines how we model the capability to use AI at all. A long tradition on technology and skill shows that the returns to a new tool are realised only in combination with complementary inputs—skill and organisational practice as much as the tool itself (Bresnahan et al., 2002). This complementarity, in which the scarcest input governs the value of the rest, motivates our representation of AI capability as a weakest-link technology, and connects to the classic statement that output in a chain of complementary tasks is capped by its weakest link (Kremer, 1993). Where Kremer (1993) fixes the static logic—value set by the scarcest complement—Baumol (1967) supplies its dynamic counterpart: when a progressive sector (here, the production of the paper) is transformed by a technology while a stagnant one (the room) is not, the *relative* value of the stagnant sector rises. On this reading artificial intelligence does not diminish the room's place in a career; by cheapening all that surrounds it, it magnifies it.

Cumulative advantage and inequality in science. Whether a productivity shock compresses or widens inequality depends on whom it reaches. Merton (1968) named the tendency of recognition to flow disproportionately to the already-eminent the Matthew effect, and DiPrete and Eirich (2006) formalise cumulative advantage as a general mechanism by which small initial differences widen over a career. Dasgupta and David (1994) give that reward system its economic foundation—science as a reputation economy in which recognition is the coin—and the conference room is one of the places that coin is earned, so a technology reshaping the room reshapes the distribution of scientific reward. A complementary mechanism works through scale rather than time: Rosen (1981) shows that where a technology lets the ablest reach a large audience at low marginal cost, small differences in talent become large differences in reward—a superstar logic the cheap, broadcastable paper might be expected to amplify, yet one that the rivalrous room, which cannot be served at scale, structurally resists. These accounts frame the distributional question our model makes precise. If artificial intelligence is cheap and broadly available—uncorrelated with institutional standing—it acts like a subsidy to the weakly resourced and compresses the gap. If instead capability tracks prestige, because elite institutions command better access, training, or permission to use frontier tools, the same technology amplifies existing hierarchies. The novelty is not to assert one outcome but to locate the knife-edge: the sign of AI's distributional

effect is governed by the correlation between capability and prestige, and the worst-placed researcher is the one who is both outside the elite and unable to deploy AI. They also frame a question subtler than the sign of the gap—what the gap is made of. We will show that artificial intelligence *erodes* the paper-side cumulative advantage Merton (1968) named, the compounding feedback edge of the well-placed, only for the space it vacates to be refilled, in the room, by luck rather than by merit.

Decision-theoretic and methodological foundations. The model rests on expected-utility foundations (von Neumann and Morgenstern, 1944; Savage, 1954), with beliefs about venue quality treated as objects updated by Bayesian learning across repeated years. Two further strands shape specific components. The role of costly, observable adaptation in the acceptance process is in the spirit of signalling (Spence, 1973): an acceptance conveys information about quality, and AI, by raising fit, both makes that signal cheaper to obtain and blunts its informativeness. The repeated choice among venues of uncertain value is an explore–exploit problem of the kind formalised by Gittins (1979), and the shift it induces—from concentrating on a familiar niche to spreading across less-explored subfields—maps onto the exploration–exploitation tension in organisational learning (March, 1991). For the quantitative analysis that accompanies the theory we adopt an agent-based approach in the tradition of Pluchino et al. (2018), whose model of "talent versus luck" shows how stochastic opportunity, not ability alone, shapes who succeeds; that lens lets us express the value of strategy and of AI in interpretable units of luck. To locate the model's one free distributional parameter in the world, we draw on the emerging practice of *silicon sampling*, in which language models prompted with personas reproduce population-level response distributions (Argyle et al., 2023; Bail, 2024); we use it only to calibrate, and treat its outputs with the caution its critics urge, since such samples can compress the true variance of human response (Bisbee et al., 2024).

Contribution. Against this literature the paper makes four contributions. First, it supplies a micro-foundation for the dissemination congestion that Gartenberg et al. (2026) and others observe: the surge in submissions is the rational consequence of many researchers treating venues as cheap, near-costless options once AI collapses the cost of adaptation. Second, it separates two value channels of a conference that prior work has treated as a single bundle—a *paper side* (feedback, fit, and the breadth of audiences a paper can reach), which artificial intelligence cheaply augments, and a *room side* (physical proximity to senior scholars), which it cannot—and derives the strategy this separation implies: submit widely, attend few, and choose small, low-concurrency venues spread across diverse subfields. Third,

it characterises when AI compresses and when it widens academic inequality, identifying the correlation between AI capability and institutional prestige as the decisive parameter; and it shows this distributional effect to be doubly fragile—the access the technology appears to open is self-limiting, competed away as adoption spreads so that non-adopters are crowded out of the dense rooms even though their papers are no worse, while the obvious remedy of competing more strategically for those rooms widens the gap rather than closing it. Fourth, it shows that this distributional effect can be invisible to the measure most often used to detect it: at the operating point our calibration locates, the level of inequality is unchanged from the pre-artificial-intelligence world even as its *composition* is transformed, the prestige-linked advantage of position giving way not to talent but to luck. Measuring inequality by its level alone can therefore miss what artificial intelligence does to it—an observation with methodological force beyond the present application. Although we develop the argument for the concrete and economically consequential case of a doctoral researcher choosing conferences, the structure is more general: it is a model of how a costless adaptation technology reshapes search over a set of opportunities when those opportunities deliver both a scalable, information-like return and a rivalrous, access-like one. The same tension arises wherever dissemination, networking, and technology adoption meet—among entrepreneurs choosing which competitions and demo days to enter, or organisations deciding where to place scarce researchers—settings to which the framework naturally extends.

3 Theory

This section builds a model of how a single doctoral researcher decides which conferences to target, and then asks what happens when an entire cohort of researchers decides in the same way. One asymmetry organises everything that follows. A conference does two things for a researcher. It is a place to get a paper read and improved—call this the *paper* side—and it is a place to stand in the same room as the senior scholars who shape careers—the *room* side. Artificial intelligence can do a great deal for the paper side: it drafts, polishes, and critiques cheaply, and, as we will show, it does most for those who begin with the least. It can do nothing for the room side. Being physically present where a limited number of senior people are, in the one session a researcher can actually attend, is a scarce good that no software manufactures. The model makes this asymmetry precise, and the paper's three central results—about *who* artificial intelligence helps, about *how its benefit erodes* as everyone adopts it, and about *why the obvious remedy backfires*—all follow from it.

A note for the reader. Every symbol is explained in words where it first appears, and all are collected for reference in Table 1. The argument can be followed without manipulating

any algebra: the equations simply record, precisely, what the surrounding words say.

3.1 Setup

Consider a doctoral researcher, labelled i, with a single paper to disseminate through conferences. Five features describe the researcher. The first is the paper's *true quality*, written θ_i (theta); a paper counts for nothing unless it clears a minimum standard $\bar{\theta}$ (theta-bar). The second is the quality of feedback the researcher already receives at home—in practice, how good and how broad the supervisor is—written τ_i^{int} ("tau-internal"); a higher value means sharper, more useful feedback. The third is the researcher's *effective AI capability*, written α_i (alpha), a number between 0 (artificial intelligence is of no use to them) and 1 (it is fully effective); it is defined in (1) below. The fourth is how well the researcher converts a meeting with a senior scholar into something that matters—a citation, a co-authorship, a job lead—written ω_i (omega), again between 0 and 1; this is social capital, the ability to turn contact into outcome. The fifth is how much the researcher values learning relative to career advancement, written η (eta).

Institutional prestige enters through these features rather than as a separate force: a researcher at an elite institution tends to have a stronger, broader supervisor (high τ_i^{int}) and to convert contacts more effectively (high ω_i). That such head-starts tend to compound over a career is the cumulative-advantage logic of Merton (1968) and DiPrete and Eirich (2006). Throughout, the researcher weighs uncertain options by their expected payoff, the standard criterion of von Neumann and Morgenstern (1944) and Savage (1954).

Conferences are labelled c, each belonging to a subfield j. Every conference is summarised by features that can be read off the previous year's programme: the number of attendees N_c; the fraction δ_c (delta) of those attendees who are *relevant* senior scholars in the researcher's own subfield—we will call such scholars "elites" for brevity; the number of sessions running at the same time m_c; the quality of the discussant format d_c, between 0 and 1; the value to a curriculum vitae of presenting there, V_c; how well the paper naturally fits the venue, f_c^0; and the score a paper must reach to be accepted, T_c.

3.2 AI capability as a weakest-link technology

Whether artificial intelligence actually helps a researcher depends on three things at once: *access* to frontier tools a_i (shaped by geography, sanctions, and price), *skill* in using them s_i, and institutional *permission* to use them p_i, each between 0 and 1. The key point is that these are not substitutes but complements: a failure in any one disables the other two. The best tools and the sharpest prompting skill are worth nothing where a department forbids

their use; permission and skill are worth nothing without access. The capability that matters here is not the power to run a quick edit—anyone can do that by ignoring a ban from a home computer—but the power to use the technology deeply and safely on the unpublished, often proprietary material from which a paper's genuine contribution is built; it is that deeper use, not casual polishing, that access and permission gate. Capability is therefore set by whichever input is scarcest—a chain is only as strong as its weakest link, the pattern known in agriculture as Liebig's law of the minimum:

$$\alpha_i = \min\{\, a_i,\ s_i,\ p_i \,\}. \tag{1}$$

The lesson for policy is immediate. Capability rises only when the *binding* input rises, and the binding input differs across groups: subsidising tools that an institution still forbids changes nothing.[1] That the value of a tool is governed by its scarcest complement has a long pedigree, from the weakest-link production of Kremer (1993) to the evidence that information technology raises productivity only alongside complementary skill and reorganisation (Bresnahan et al., 2002). Of the three inputs, *skill* is the researcher's own—the talent dimension of capability—whereas *access* and *permission* are conferred by where the researcher sits. The distinction lies dormant here but becomes decisive in the final result, when we ask what, once the tools are everywhere, still decides who succeeds.

3.3 The room: elite contact and concurrency

The distinctive value of showing up in person is the chance that a relevant senior scholar actually hears the work—a channel whose career returns, working through face-to-face contact and the collaborations it seeds, are well documented (Lopez de Leon and McQuillin, 2020; Chai and Freeman, 2019; Catalini, 2018; Boudreau et al., 2017). How likely is such an encounter? The expected number of relevant elites at the venue is $\delta_c N_c$ (the relevant fraction times the headcount), but they are spread across m_c simultaneous sessions, so the expected number in the one session the researcher attends is

$$\lambda_c = \frac{\delta_c N_c}{m_c}. \tag{2}$$

[1] The strict minimum is the Leontief limit of a constant-elasticity-of-substitution aggregator. Writing $\alpha_i = \left(w_a a_i^{\rho} + w_s s_i^{\rho} + w_p p_i^{\rho}\right)^{1/\rho}$ with non-negative weights summing to one, the elasticity of substitution among inputs is $1/(1-\rho)$, and $\alpha_i \to \min\{a_i, s_i, p_i\}$ as $\rho \to -\infty$. All qualitative results below survive any sufficiently low elasticity; we adopt the limiting case because the three inputs are close to non-substitutable—prompt skill cannot offset an outright ban—and because it gives the sharpest statement of the binding-constraint logic.

Treating who turns up in that session as a matter of chance (a Poisson event), the probability that *at least one* relevant elite is in the room is

$$p_c^E = 1 - e^{-\lambda_c}. \tag{3}$$

Here is the crucial comparison, and it is worth making concrete. A large generalist conference has a high headcount N_c but a low relevant fraction δ_c and many parallel tracks m_c; a small specialised workshop has a modest headcount, a high relevant fraction, and few tracks. Take two illustrative venues. A workshop with $N_c = 60$ attendees, a tenth of them relevant ($\delta_c = 0.1$), in a single track ($m_c = 1$), delivers $\lambda_c = 0.1 \times 60/1 = 6$ expected relevant elites in the room. A flagship with $N_c = 600$, a fiftieth relevant ($\delta_c = 0.02$), across $m_c = 24$ parallel sessions, delivers $\lambda_c = 0.02 \times 600/24 = 0.5$. The flagship contains twelve relevant elites against the workshop's six in absolute terms—yet the researcher is far more likely to share a room with one at the workshop. The room is dense where the field is narrow and the doors are few. (Ordinary, non-elite contacts are useful too, but they are available at every venue in equal measure; being the same everywhere, they do not affect the choice, and we set them aside.)

3.4 Feedback: the discussant and the supervisor

Feedback sharpens what the researcher knows about the paper and improves it. It arrives from three sources, and we measure each by its *precision*—loosely, how much it improves the work. Home feedback has precision τ_i^{int}. Artificial-intelligence feedback has precision $\tau_i^{\text{AI}} = \alpha_i \, \bar{\tau}^{\text{AI}}$: it is the most a capable assistant could add, $\bar{\tau}^{\text{AI}}$, scaled down by how effectively the researcher can actually use it, α_i. Conference feedback depends on format: a venue with a serious discussant who reads the paper delivers $\tau_c^{\text{ext}} = d_c \, \bar{\tau}^{\text{ext}}$, whereas a present-and-leave slot ($d_c \approx 0$) delivers almost nothing. After attending a set of conferences D, the three sources add up to a total precision

$$\tau_i^{\text{tot}}(D) = \tau_i^{\text{int}} + \tau_i^{\text{AI}} + \sum_{c \in D} \tau_c^{\text{ext}}. \tag{4}$$

One asymmetry between artificial intelligence and the supervisor matters for the whole argument. Supervisor feedback is deep but narrow—excellent within the supervisor's specialty, thin outside it—while artificial-intelligence feedback is broad and even. For a student with a world-class, wide-ranging advisor (typically the well-resourced student), artificial intelligence mostly *adds to* feedback that is already good. For a student whose advisor is junior or narrow, it can *exceed and replace* the supervisor, and supply breadth the advisor

never could. This is precisely the pattern found in field and laboratory studies of artificial-intelligence assistance, which raise the output of less-expert workers most while adding little for the most expert (Noy and Zhang, 2023; Brynjolfsson et al., 2025).

3.5 The payoff from an attended set

The value of attending a set of conferences D has two parts. The first is *knowledge*, which rises with total feedback but with diminishing returns—each extra unit of feedback helps a little less than the last, the familiar concavity of learning:

$$\text{Knowledge}_i(D) = \eta\, g\big(\tau_i^{\text{tot}}(D)\big), \tag{5}$$

where g is an increasing, concave function. The second part is *career value*, which combines professional contact with the line on the curriculum vitae. Here a point about breadth is essential. Senior scholars in *different* subfields open *independent* doors, so contacts in different subfields do not duplicate one another—their value adds up across subfields—but contacts piled into a single subfield quickly become redundant. How many subfields a researcher can credibly work across is limited, however, both by their own reach and by the market's tolerance for breadth, which is the long-standing tension between exploiting one niche and exploring several (March, 1991). Writing Λ_{ij} for the expected elite contacts a researcher gathers in subfield j (the sum of λ_c over the attended venues in that subfield), $J(\alpha_i)$ for the number of subfields the researcher can reach (which artificial intelligence widens), and $\bar{J}$ for the number the field rewards before it begins to read breadth as a lack of focus, the *effective breadth* is $\tilde{J}_i = \min\{J(\alpha_i), \bar{J}\}$ and career value is

$$\text{Career}_i(D) = \omega_i R^E \sum_{j=1}^{\tilde{J}_i} \Lambda_{ij} + \sum_{c \in D} V_c, \tag{6}$$

where R^E is the gross value of one elite contact. Career value is thus the contacts that count—converted at the researcher's own rate ω_i and capped at the rewarded breadth $\tilde{J}_i$—plus the curriculum-vitae value of the venues attended. The total payoff puts the two parts together, switched off entirely if the paper fails to clear the quality bar:

$$U_i(D) = q_i \left[\eta\, g\big(\tau_i^{\text{tot}}(D)\big) + \omega_i R^E \sum_{j=1}^{\tilde{J}_i} \Lambda_{ij} + \sum_{c \in D} V_c \right], \qquad q_i = \mathbf{1}[\theta_i \geq \bar{\theta}]. \tag{7}$$

The term q_i is simply a switch: it equals 1 when the paper clears the bar ($\theta_i \geq \bar{\theta}$) and 0 when it does not.

3.6 Submission, acceptance, and the cost collapse

To attend, a researcher must first be accepted, and to submit they must first adapt the paper to the venue, at a cost

$$\kappa_i = (1 - \alpha_i)\,\kappa_0, \tag{8}$$

where κ_0 is the cost before any artificial-intelligence help. As capability α_i rises toward 1 this cost falls toward zero, and the set of venues worth targeting, $\mathcal{C}_i(\alpha_i)$, widens accordingly: when submitting is nearly free, it is worth trying for many more venues. Capability also improves how well the paper fits a venue, $f_{ic} = f_c^0 + \alpha_i \bar{\phi}$, which raises the chance of acceptance,

$$p_{ic}^{\text{acc}} = \Phi\big(\theta_i + f_{ic} - T_c\big), \tag{9}$$

where Φ is the standard normal curve familiar from statistics, here turning a paper's margin over the bar (quality plus fit, less the threshold) into a probability between 0 and 1.

Two consequences of acceptance deserve note. First, a researcher does not perfectly know their own quality: they hold a belief $\theta_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$—a best guess μ_i surrounded by uncertainty σ_i^2, with the uncertainty larger for non-elite students, who get noisier signals about how good their work is. Because every venue's acceptance is driven by the same underlying quality θ_i, acceptances are *correlated*: a strong paper tends to be accepted everywhere and a weak one nowhere, so the pattern of acceptances is itself a signal of quality in the sense of Spence (1973). Second, because artificial-intelligence fit lowers the effective bar $T_c - \alpha_i \bar{\phi}$, an acceptance tells one less about quality than it used to: artificial intelligence makes the quality test cheap to run but blunts what each result reveals.

3.7 The decision problem

The researcher now chooses which venues to submit to. Formally they pick a submission set S from the affordable venues $\mathcal{C}_i(\alpha_i)$; acceptances A then arrive, each venue in S succeeding with its own probability p_{ic}^{acc}; and, limited by time and money to attending at most k_i conferences, the researcher goes to the most valuable affordable subset of the offers, declining the rest. Putting this together, the researcher solves

$$\max_{S \subseteq \mathcal{C}_i(\alpha_i)} \; \mathbb{E}_{A|S}\left[\max_{\substack{D \subseteq A \\ |D| \leq k_i}} U_i(D) \right] - \kappa_i\,|S|. \tag{10}$$

In words: choose the submission set that maximises the expected value of the best affordable bundle of acceptances, minus the cost of all the submissions made. The inner maximisation is the freedom to turn offers down (the withdraw option); the outer expectation averages over which submissions happen to land. Artificial-intelligence capability α_i pushes on this problem in five separate places at once: it lowers the submission cost κ_i, raises feedback precision τ_i^{AI}, improves fit f_{ic}, widens the affordable set $\mathcal{C}_i$, and widens breadth $J(\alpha_i)$.

3.8 Learning across years

A researcher does not know a venue's relevant density δ_c for certain; they estimate it from last year's programme, holding a belief $\delta_c \sim \mathrm{Beta}(a_c, b_c)$ with average $a_c/(a_c + b_c)$, whose confidence $a_c + b_c$ is greater for more capable researchers (artificial intelligence makes a programme cheap to scrape and read). Having attended in year t and seen e_{ct} relevant elites and n_{ct} others, they revise the belief to $\mathrm{Beta}(a_c + e_{ct},\, b_c + n_{ct})$ and solve (10) again the next year. The choice is thus a repeated balance between exploring unfamiliar venues and exploiting known ones, of the kind formalised by Gittins (1979); more capable researchers start with sharper beliefs and need to explore less.

3.9 From the individual to the population

Equation (10) is the problem of a single researcher who takes the conference world as fixed: each venue's bar T_c, its capacity N_c, and its supply of elite attention look like constants from one person's vantage point. The paper's distributional claims, however, concern what happens when an *entire population* behaves this way at once—and then two of those "constants" are no longer fixed. Two aggregate adjustments therefore need stating before the results.

The bar rises with adoption. Suppose organisers fill a fixed capacity N_c from a pool of P_c applicants. The bar that clears the market, T_c^*, is the one at which exactly enough applicants get in:

$$\mathbb{E}_i\big[\Phi(\theta_i + f_c^0 + \alpha_i\bar{\phi} - T_c^*)\big] = \frac{N_c}{P_c}, \tag{11}$$

the average taken over everyone who applies. The left-hand side—the average chance of acceptance—rises with every applicant's fit, and fit rises with capability; so when average capability $\bar{\alpha}$ rises, the only way to keep the same fraction N_c/P_c getting in is for the bar T_c^* to climb. The fit advantage that helps one researcher in isolation is therefore partly competed away once everyone has it: when seats are scarce, the bar simply rises to meet the crowd, and the true gain in the chance of acceptance is smaller than (9) on its own would suggest.

The room has a ceiling. The room is finite in a second sense too. A senior scholar's attention is limited, so a venue can convert proximity into genuine contact only up to a point. This same scarcity is what places the room beyond artificial intelligence's reach: the binding constraint is senior scholars' attention, and no assistant, however capable, adds a senior scholar to the field or an hour to the day. Tools that promise to broker introductions only redirect the same fixed attention—they lengthen the queue rather than enlarge the supply. The room is therefore beyond artificial intelligence not as a modelling convenience but because the good it rations—being heard by someone who matters—is not one that software can produce. This is the local form of an older dynamic (Baumol, 1967): as artificial intelligence transforms the progressive, paper-producing side of a conference while leaving the stagnant room untouched, the room does not lose importance but gains it—the scarce complement whose relative value rises precisely because all that surrounds it has grown cheap.

Assumption 1 (Elite-contact capacity)**.** Each venue c supplies at most $\bar{\Lambda}_c$ units of effective elite contact per period. When a mass M_c of researchers competes for contact at c, each receives expected contact $\min\{\lambda_c,\ \bar{\Lambda}_c/M_c\}$.

Without this ceiling, contact at a venue would be unrationed and the room would never fill up. Assumption 1 adds the congestion that, as we will see, makes the access advantage of small venues *self-limiting*: the more researchers crowd in, the less attention each can win. Both adjustments—the rising bar and the contact ceiling—are easy to state but cannot be solved by hand once capability α_i and quality θ_i differ from one researcher to the next. Their quantitative behaviour is exactly what the simulation of Section 4 measures.

3.10 Results

The results come in two groups. The first four propositions establish how the pieces of the model behave—they are the *mechanism*. The last four are the paper's *central results*: they describe what the mechanism implies for inequality, for access, and for the composition of success itself, and they are the claims a reader should carry away. Crucially, the central results turn out to rest on the single structural fact that the room is finite, not on the particular numbers used to illustrate the model—a robustness we establish in Section 4.

The mechanism. The first proposition is the structural advantage of the small venue from which much of the rest follows.

Proposition 1 (Concurrency ranking)**.** *The probability of elite contact* $p_c^E = 1 - e^{-\delta_c N_c/m_c}$ *falls with concurrency* m_c *and rises with relevant density* δ_c*. A small, specialised, single-track*

venue can therefore beat a large, prestigious, multi-track venue in expected elite contact, even though the latter holds more elites in absolute terms.

(Proof in Appendix A.)

Proposition 2 (AI repurposes the conference)**.** *As capability α_i rises, artificial-intelligence feedback τ_i^{AI} rises, and because knowledge is concave the marginal value of one more piece of conference feedback falls. A venue's feedback then justifies the trip only where its discussant precision τ_c^{ext} still exceeds what artificial intelligence already supplies. Weak-discussant venues lose their feedback rationale, and the reason to attend in person shifts toward networking and toward the few venues with strong discussants.*

(Proof in Appendix A.)

Proposition 3 (Bounded feedback democratisation)**.** *Hold capability α fixed and equal across groups. Because knowledge is concave and elite students start with sharper home feedback ($\tau_E^{\mathrm{int}} \gg \tau_N^{\mathrm{int}}$), the same artificial-intelligence feedback adds more knowledge for non-elite researchers. Artificial intelligence substitutes for, and can exceed, weak and narrow supervision, while it merely tops up supervision that is already strong and broad.*

(Proof in Appendix A.)

Proposition 4 (From depth to breadth)**.** *As capability α_i rises, effective breadth $\tilde{J}_i = \min\{J(\alpha_i), \bar{J}\}$ rises until the field's tolerance $\bar{J}$ binds. The best networking strategy therefore shifts from concentrating on one subfield, where extra contacts are redundant, to spreading across several, where each contact is an independent door—as far as the field's reward for focus allows.*

(Proof in Appendix A.)

The central results. The remaining four propositions are the paper's contribution. The first says the distributional effect of artificial intelligence is not a fixed quantity but a knife-edge that turns on a single feature of the world.

Proposition 5 (Prestige and access: the knife-edge)**.** *Advantage is the* interaction *of prestige—which sets home feedback τ_i^{int} and conversion ω_i—with artificial-intelligence capability α_i, which sets cost, feedback, fit, and breadth. The effect of artificial intelligence on the gap between elite and non-elite researchers therefore depends on how capability is distributed. Where capability is unrelated to prestige—cheap and widely available—artificial intelligence narrows the gap, through the levelling of Proposition 3. Where capability instead tracks prestige—elite*

institutions commanding better access, skill, or permission—the two advantages compound and artificial intelligence widens the gap. The doubly disadvantaged researcher is the one who is both non-elite and low in capability.

(Proof in Appendix A.)

The second central result shows that even this access is not durable: as artificial intelligence spreads, the very crowding it produces erodes the advantage of the dense rooms.

Proposition 6 (Self-limiting access)**.** *Under Assumption 1, expected elite contact at a venue is* $\min\{\lambda_c,\ \bar{\Lambda}_c/M_c\}$*, which never rises with the competing mass* M_c *and strictly falls once* $M_c > \bar{\Lambda}_c/\lambda_c$*. Because rising adoption draws capable researchers toward the low-concurrency venues (Propositions 1 and 4), the mass* M_c *at those venues grows; their access advantage therefore survives while adoption is partial but erodes as adoption becomes widespread.*

(Proof in Appendix A.)

The third central result concerns the obvious piece of advice one might give a disadvantaged researcher—"be strategic about which venues you target"—and shows that, once the room is congestible, this advice backfires.

Proposition 7 (Strategic targeting amplifies the gap)**.** *Take two researchers alike in every respect except capability, the more capable having higher fit* f_{ic} *and lower cost* κ_i*. Let each choose between submitting only to a generalist venue and targeting the dense, low-concurrency workshops of Proposition 1. Because the workshop bar* T_c^* *rises with the mass of high-fit applicants who target it* (11)*, and because the room there is capacity-constrained (Assumption 1), the more capable researcher both clears the bar more often and captures more of the rationed contact. The gain from targeting workshops—the difference between targeting them and staying with the generalist venue—is therefore increasing in capability. Offering researchers the option to target venues strategically thus* widens *the career gap between high- and low-capability researchers rather than closing it: the strategic move sends everyone toward the same scarce seats, which the more capable win.*

(Proof in Appendix A.)

The final central result looks beneath the gap to the *composition* of success—to what, once the dust has settled, decides which researchers rise. It is the model's closest contact with the tradition of Pluchino et al. (2018), in which chance, not talent, dominates who ends up ahead.

Proposition 8 (The lottery of the room)**.** *Sort the determinants of the payoff* U_i *into three kinds.* Talent *is what the researcher brings—paper quality* θ_i *and using-skill* s_i*.* Position

is what prestige confers, through sharper home feedback τ_i^{int} *and a higher conversion rate* ω_i*.* Luck *is the chance realisation of elite contact in the room, the Poisson event of* (3)*. Artificial intelligence operates only through the paper: by* $\tau_i^{\text{AI}} = \alpha_i \bar{\tau}^{\text{AI}}$ *it supplies feedback to all and, by the concavity of g with Proposition 3, does most for those whose home feedback was weakest—compressing the advantage that standing used to buy on the paper side. It does not erode talent in the same way; by widening fit and breadth it lets a good paper travel a little further. On the room—on luck—it has no purchase (Assumption 1). Holding fixed how capability is distributed across the hierarchy, a rise in average capability* $\bar{\alpha}$ *therefore lowers the share of* $\text{Var}(U_i)$ *owed to* position*, leaves the already-small share owed to* talent *roughly unchanged, and raises the share owed to* luck*. The more capable artificial intelligence becomes, the less a researcher's success turns on standing—or, as ever, on talent—and the more it turns on who happens to be in the room.*

(Proof in Appendix A.)

Two cautions guard against misreading. First, this concerns the *composition* of success, not its level: Proposition 5 speaks to the average gap between elite and non-elite researchers, which can widen even as the present result holds, because a group may keep a higher mean outcome while the *dispersion* of individual outcomes comes to be ruled by chance. The two are complementary faces of one mechanism—artificial intelligence can enlarge the elite's average lead while making any single career more of a lottery—and Section 4 reports both, the mean gap and the variance shares, side by side. Second, the result follows from the model's central premise rather than slipping past it: precisely because the room is the one good artificial intelligence cannot manufacture, levelling the paper must let the residual weight of success fall on the room. The implication is not fatalism but a redirection—what decides whether chance or hierarchy prevails is access to the room itself, not the power of the tools, a theme we take up in Section 5.

Taken together, the central results say something specific. Artificial intelligence levels the paper but not the room; whether it narrows or widens academic inequality turns on whether capability is shared out with or against the existing hierarchy (Proposition 5); the access it appears to open is competed away as adoption spreads (Proposition 6); the natural remedy of strategic targeting makes matters worse (Proposition 7); and, beneath the gap itself, levelling the paper lets the composition of success settle on the room, so that as the tools grow more capable it is chance, more than standing or talent, that decides who rises (Proposition 8). The rising bar of (11) and the contact ceiling of Assumption 1, together with the population conditions in these propositions, cannot be solved in closed form once capability and quality vary across researchers. Their quantitative behaviour—and, importantly, how little of it depends on the calibration—is the subject of Section 4, which solves the model as an agent-

based rational-expectations equilibrium in the tradition of Pluchino et al. (2018).

Table 1: Notation used in the model (a reading key).

Symbol	Meaning	Symbol	Meaning
Indices and sets		***Conference (cont.)***	
i	Researcher	T_c	Acceptance bar
c	Conference (venue)	T_c^*	Market-clearing bar
j	Subfield	λ_c	Expected elites in session ($\delta_c N_c/m_c$)
$\mathcal{C}_i(\alpha_i)$	Set of venues worth targeting	$\bar{\Lambda}_c$	Venue's ceiling on elite contact
S	Submission set	M_c	Researchers competing for contact
A	Offers received	P_c	Applicants to conference c
D	Conferences attended		
D_j	Attended venues in subfield j	***Feedback, contact, and payoff***	
		τ_i^{AI}	AI feedback precision ($\alpha_i \bar{\tau}^{\text{AI}}$)
Researcher		τ_c^{ext}	Discussant feedback precision ($d_c \bar{\tau}^{\text{ext}}$)
θ_i	True paper quality	$\tau_i^{\text{tot}}(D)$	Total feedback precision
$\bar{\theta}$	Minimum quality to count	$g(\cdot)$	Knowledge function
q_i	Quality switch (1 if $\theta_i \geq \bar{\theta}$, else 0)	p_c^E	Prob. of ≥ 1 elite contact ($1 - e^{-\lambda_c}$)
μ_i	Guess of own quality	Λ_{ij}	Expected elite contacts in subfield j
σ_i^2	Uncertainty of quality guess	R^E	Value of one elite contact
τ_i^{int}	Precision of internal feedback	$J(\alpha_i)$	Reachable breadth
ω_i	Contact-to-career outcome rate	$\bar{J}$	Breadth the field rewards
η	Weight on learning vs. career	$\tilde{J}_i$	Effective breadth
k_i	Max conference attendance limit	$U_i(D)$	Payoff from attending set D
AI capability		***Submission, cost, and learning***	
α_i	Effective AI capability: $\min(a_i, s_i, p_i)$	κ_0	Pre-AI submission cost
a_i	Tool access	κ_i	Cost with AI ($(1-\alpha_i)\kappa_0$)
s_i	Using skill	$\bar{\phi}$	Max fit improvement via AI
p_i	Institutional permission	$\bar{\tau}^{\text{AI}}$	Max AI feedback addition
$\bar{\alpha}$	Population-average capability	$\bar{\tau}^{\text{ext}}$	Max discussant feedback addition
		p_{ic}^{acc}	Chance of acceptance
Conference		Φ	Standard normal curve
N_c	Attendees (admission capacity)	a_c	Density-belief prior (prior elites)
δ_c	Fraction of senior scholars	b_c	Density-belief prior (prior others)
m_c	Concurrent sessions running	e_{ct}	Relevant elites seen at c in year t
d_c	Quality of discussant format	n_{ct}	Others seen at c in year t
V_c	CV value of presenting		
f_c^0	Natural fit of paper to venue		
f_{ic}	Fit with AI ($f_c^0 + \alpha_i \bar{\phi}$)		

4 Quantifying the Equilibrium: An Agent-Based Analysis

4.1 Why an agent-based analysis is required

The propositions of Section 3 are stated as comparative statics, and each can be signed analytically for a single researcher facing fixed beliefs. The substantive questions of the paper, however, are distributional and equilibrium questions: how the dispersion of career outcomes across thousands of heterogeneous researchers responds to a change in the joint distribution of artificial-intelligence capability and institutional prestige (Proposition 5), and whether the crowding that artificial intelligence induces is a transient error or a stable resting point of the system (Proposition 6). Neither question admits a closed-form answer once capability α_i and latent quality θ_i vary across the population, because the individual decision is a discrete portfolio choice under correlated acceptance risk, and the market-clearing thresholds in equation (11) are themselves endogenous to the choices that researchers make.

We therefore solve the model numerically. The simulation is not an illustration layered on top of the theory; it *is* the theory, evaluated at a population of agents whose decisions and beliefs are required to be mutually consistent. Agent-based methods of this kind have an established role in the economics of science and in management research, where the object of interest is the equilibrium of many interacting, boundedly informed actors rather than the behaviour of a representative agent (Pluchino et al., 2018). Our contribution within that tradition is to insist on decision-theoretic discipline at the level of the individual: each researcher solves a well-posed expected-utility problem, and the population is carried to a rational-expectations fixed point. The full source, including the calibration and every experiment reported here, is provided in the replication files `conf_sim_v3.py` (the simulation and every experiment reported here), together with `llm_adoption_stage1.py` and `llm_stage2_calibration.py` (the language-model calibration of Section 4.8).

4.2 The computational environment

The environment instantiates the primitives of Section 3 directly. A population of N researchers is partitioned into a prestigious minority (a share of one in five) and a non-prestigious majority. Each researcher i is endowed with a latent paper quality θ_i, drawn from a common distribution, and observes only a noisy private signal μ_i of that quality; the signal is sharper for prestigious researchers, reflecting the better mentoring and internal feedback that a strong research group provides. Two further advantages are tied to prestige and are independent of artificial intelligence: the precision of internal supervision

τ_i^{int} and the rate ω_i at which a contact with a senior scholar converts into a career outcome. These channels are the model's representation of the "old" sources of advantage that artificial intelligence may or may not displace.

A researcher's artificial-intelligence capability $\alpha_i \in [0,1]$ is constructed as the weakest of three components—access, skill, and permission—so that capability is governed by its binding constraint rather than by an average, as in equation (1). The experiments vary a single parameter, the correlation between capability and prestige, between two poles: a *democratised* world, in which capability is distributed independently of institutional standing, and an *entrenched* world, in which capability is high for prestigious researchers and low for the rest. This correlation is the central object of the analysis.

The venue universe consists of one large generalist conference and a set of specialised workshops, arranged across subfields. The asymmetry between them is the heart of Proposition 1 and deserves emphasis. The expected number of useful encounters with relevant senior scholars in the single session a researcher can attend is

$$\lambda_c \;=\; \frac{\delta_c \, N_c}{m_c}, \tag{12}$$

the density of relevant senior scholars times venue size, divided by the number of parallel sessions. The generalist venue is penalised on both counts: it runs many concurrent sessions, and—because senior scholars are a *limited pool*—a larger venue dilutes the density of those relevant to any one researcher. We model this dilution explicitly, setting the relevant-scholar density at a venue to

$$\delta_c \;=\; \bar{\delta}\left(\frac{N_{\text{small}}}{N_c}\right)^{\gamma}, \tag{13}$$

where $\bar{\delta}$ is the density at a small workshop, N_{small} a reference workshop size, and $\gamma \in [0,1]$ a dilution elasticity ($\gamma = 0$ leaves density constant; $\gamma = 1$ holds the absolute number of relevant scholars fixed, so that doubling venue size halves density). At our calibration the generalist venue contains more senior scholars in absolute terms than any workshop, yet its density of *relevant* scholars is less than half a workshop's, and its per-session encounter rate is roughly one-fifth as large. The big conference is where one acquires a line on the curriculum vitae, not where one meets the right people.

4.3 The submission decision and its justification

Each researcher chooses a portfolio of submissions S to maximise the expected value of the conferences ultimately attended, net of submission costs, exactly as in the individual

problem (10):

$$\max_{S} \; \mathbb{E}_{A|S}\left[\max_{D\subseteq A,\,|D|\le k} U_i(D)\right] \; - \; \kappa_i\,|S|.$$

This is a discrete optimisation under uncertainty with two features that any credible solution must respect. First, acceptances are *correlated*: because a researcher's realised review score at every venue is driven by the same latent quality θ_i, a strong researcher tends to be accepted everywhere and a weak one nowhere. Second, the researcher does not know θ_i; the decision must be taken under the posterior implied by the private signal μ_i.

We address both features without resorting to a crude heuristic. The researcher's posterior over θ_i is the conjugate Gaussian update of the prior by the signal, and the value of any candidate submission is integrated over that posterior on a Gauss–Hermite grid. Integrating over θ_i is what reinstates the correlation in the decision: conditional on a value of θ_i, acceptances at distinct venues are independent (only idiosyncratic review noise remains), so the conditional problem is tractable, and averaging over the posterior restores the common-quality dependence. The same integration handles the quality gate $\mathbf{1}[\theta_i \ge \bar{\theta}]$ exactly, since the gate, the acceptance probability, and the value of attendance are all functions of the one quantity θ_i and are evaluated jointly at each grid point.

The portfolio itself is selected by greedy forward inclusion: venues are considered in descending order of value, and a venue is added whenever its marginal contribution to expected attended value exceeds the submission cost κ_i. This is not an expedient. The objective is *submodular* in S—the cap of k attendable venues and the concavity of the knowledge function together generate diminishing returns to additional submissions—while the cost is modular, so greedy selection is the standard, provably near-optimal procedure for this class of problem. Processing venues in value order, and tracking the distribution of filled attendance slots with a small dynamic program, makes the marginal "option value" of each venue—its probability of being accepted *and* of displacing a weaker offer into the attended set—exact given the posterior. The single inclusive venue is therefore selected on its merits, as insurance for researchers whose posterior places weight on rejection at selective venues, and is never imposed; a researcher for whom no venue clears its cost rationally submits nothing.

For transparency we state what is exact and what is approximate. The realised market is exact: correlated acceptances, capacity-based clearing, the congestion of the room, the breadth cap, the quality gate, and submission costs all operate without approximation. The decision is exact in its treatment of acceptance correlation and the quality gate; it is approximate in two conservative respects, namely that the career value of a venue is treated additively at the decision stage (the breadth cap of equation (6) binds only in the realised payoff) and that the portfolio is chosen greedily rather than by exhaustive search. Section 4.9

reports that the conclusions are insensitive to these choices.

4.4 Equilibrium concept and convergence

The acceptance thresholds at which venues clear are endogenous. When artificial intelligence raises the apparent fit of submissions, scores inflate, more researchers apply, and the bar required for admission rises—which in turn alters who finds it worthwhile to apply. A one-shot simulation in which researchers hold fixed beliefs about selectivity would conflate the genuine, cost-driven response to cheaper submission with a mechanical artefact of stale expectations. We therefore solve for a *rational-expectations equilibrium*, in which the selectivity researchers anticipate coincides with the selectivity the market produces.

Equilibrium is computed by iterating beliefs to a fixed point. Researchers enter a period with a believed clearing threshold for each venue, make their submission and attendance decisions, and the market clears to produce a realised threshold; beliefs are then updated part of the way toward the realised value, and the period repeats. Figure 1 shows that this process converges quickly and stably from a naive prior: the believed thresholds settle within roughly seven iterations and fluctuate within a few hundredths of a threshold unit thereafter. At the fixed point the selective workshops clear at a score of about nine and a half and the generalist venue at about seven and a half, on a quality scale whose population mean is five—a quantitative measure of how far artificial-intelligence-driven competition has raised the bar, and a first sign that the workshops, not the generalist venue, are where the binding scarcity lies. Because the comparative statics that follow are computed at this fixed point, the crowding they describe is a property of equilibrium, not a transient.

4.5 Calibration

The values in Table 2 are chosen in the spirit of the talent-versus-luck literature (Pluchino et al., 2018), whose force comes not from finely tuned numbers but from keeping the substantive parameters few and letting the striking outcome *emerge*. We follow three disciplines. First, parameters that merely fix the *scale* of an arbitrary quantity—the unit in which quality or career value is measured—are *normalisations*: what matters is the ratios among payoff components and the signs of the asymmetries, not these levels—just as the equal starting capital in the talent-versus-luck model is a unit choice, not an assumption. Second, parameters that could tilt the model toward one channel are set to *neutral* values—equal, where two forces are symmetric—so that no result is smuggled in by assumption. Third, the *only* parameters carrying the paper's substantive content are the prestige head-starts: the advantages in supervision, in converting contacts, and in self-knowledge that elite researchers hold

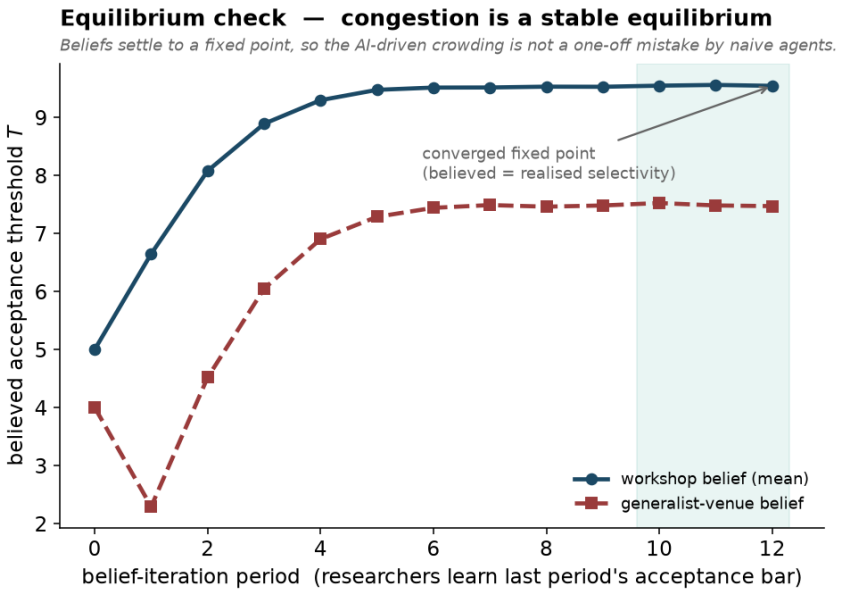


Figure 1: Convergence of believed acceptance thresholds to a fixed point. Researchers replay the year and update beliefs toward realised selectivity; thresholds settle within about seven iterations, establishing that the crowding induced by artificial intelligence is a stable rational-expectations equilibrium rather than a one-off mistake by naive agents.

independently of artificial intelligence. For these we justify the *sign and rough magnitude*, not a precise figure, and Section 4.9 shows the central *widening* result depends only on their sign—that elite exceeds non-elite—not on their size.

One feature mirrors the talent-versus-luck design directly and is worth stating plainly. In the democratised world (the left-hand end of the experiment), artificial-intelligence capability is drawn from a *common* distribution: every researcher, elite or not, has the same expected capability. Inequality is therefore not written into the capability distribution; whatever inequality the model produces there must *emerge* from the scarce room and the prestige head-start—exactly as the inequality in the talent-versus-luck model emerges from luck acting on equal starting points rather than from unequal talent.

4.6 Results

The results fall into two groups, as in the theory. The first four—the structural ranking, the levelling on the paper side, and the repurposing and breadth shifts—establish that the model's *mechanism* behaves as Section 3 claims. Three of the paper's *central results* follow in this subsection—the distributional knife-edge, the enclosure of the room, and the backfire of strategy; the fourth, the lottery of the room, is taken up in Section 4.7. All standard errors

Table 2: Calibration, grouped by the *role* each parameter plays. Most values are normalisations or neutral baselines; only the prestige head-start (group C) carries substantive content, and the central widening result depends on its sign, not its magnitude (Section 4.9).

Parameter	Value	Role and rationale
(A) Units and scale — fix the units of quality and value; only ratios and signs matter, not these levels		
Quality θ (mean, s.d.)	5, 1.5	Quality has no natural unit; the mean is a free location, the spread is read relative to the floor below.
Contact value R^E	0.5	Fixes the unit of career value (the counterpart of the equal starting capital in Pluchino et al., 2018).
CV value V_c	0.4	A small fixed benefit per venue; sets a minor unit.
Review noise	1.0	The neutral reference level for luck in acceptance.
(B) Neutral baselines — set so that no channel is favoured by assumption		
Quality floor $\bar{\theta}$	2.0	Two s.d. below the mean, so the gate screens only clearly poor work: a floor, not the contest.
Submission cost κ_0	0.10	Small, encoding the premise that submitting is now cheap; too small to drive choices on its own.
AI fit $\bar{\phi}$ / AI feedback $\bar{\tau}^{\text{AI}}$	2.0 / 2.0	Set *equal*, so the model favours neither AI-as-polish nor AI-as-reviewer.
Discussant feedback $\bar{\tau}^{\text{ext}}$	1.5	Just below $\bar{\tau}^{\text{AI}}$, so the repurposing result (Prop. 2) *emerges* from concavity rather than being imposed.
Knowledge weight η	2.0	Makes learning and career payoffs comparable in size; neither dominates by fiat.
Rewarded breadth $\bar{J}$ / budget k	3 / 3	Small integers from real limits: a field rewards a few areas, a student attends a few venues a year.
(C) Prestige head-start — the only substantive asymmetry; signs matter, sizes do not		
Prestigious share	0.20	One in five: a top-tier minority, roughly the concentration of candidates from leading departments.
Supervision τ^{int}	$0.5 \to 2.5$ $(5\times)$	A star advisor and a deep group are a real, AI-independent edge. The multiple is illustrative; results survive its variation.
Conversion ω	$0.30 \to 0.70$ $(\sim 2\times)$	Elite students are better coached and connected, so the same meeting pays off more.
Self-knowledge (signal noise)	$1.2 \to 0.5$ $(\sim 2\times$ sharper)	Better mentoring yields more accurate self-assessment; this is why some low-confidence non-elites rationally opt out.
(D) AI capability — the experimental variable, egalitarian at the democratised end		
corr(capability, prestige)	$0 \to 1$	*The* experiment. At 0, capability is drawn from a common distribution (equal for all in expectation), so any widening must emerge.
Component mean, s.d.	0.55, 0.18	A moderate bell curve of ability to use AI, in the spirit of normally distributed talent in Pluchino et al. (2018).
Entrenched endpoints	0.15 / 0.90	The far end of the experiment: non-elite low, elite high.
(E) Venue technology — the paper/room asymmetry		
Workshop density $\bar{\delta}$	0.12	Fraction of a focused room who are relevant experts.
Size dilution γ	0.35	Experts are a limited pool, so larger venues are less dense (0, none; 1, fixed headcount).

are computed across the one hundred replications and are small; we report them where they bear on a marginal claim.

The structural ranking (Proposition 1). The encounter rate of equation (12) is an order of magnitude higher at a specialised workshop than at the generalist venue—roughly 7.2 against 1.3 relevant encounters per attended session at the calibrated values. This holds even though the generalist venue contains more senior scholars in absolute terms: its size and its many parallel sessions, compounded by the density dilution of equation (13), make any one of them unlikely to be the right one in the room a researcher actually occupies. The advantage of the focused workshop for accumulating elite contact is thus a structural feature of the venue technology, prior to any behaviour.

Levelling on the paper side (Proposition 3). When capability is held equal across the two groups, the knowledge that artificial intelligence adds is larger for non-prestigious researchers than for prestigious ones—a gain of about 0.84 against 0.40, a ratio of roughly two to one. The mechanism is the concavity of the knowledge function combined with the fact that prestigious researchers already enjoy strong internal feedback: a tireless artificial reviewer is worth more to someone who lacks one. On the paper side of the conference, then, artificial intelligence is a levelling force. The decisive question is whether that levelling survives once the room is taken into account.

Repurposing and breadth (Propositions 2 and 4). Figure 2 traces two consequences of rising capability. As capability increases across its range, the marginal value of a human discussant falls steadily—by more than half at the calibrated values, from about 1.0 to 0.4—because artificial feedback substitutes for the feedback a discussant would have provided. Over the same range the effective breadth of a researcher's portfolio rises, from a single subfield to roughly two and a half, approaching the rewarded maximum of three. The two movements are the quantitative content of the claim that artificial intelligence repurposes the conference: its value as a source of feedback erodes while its value as a site of networking and breadth-building persists. Importantly, this shift emerges from the concavity of the technology; it is not built into the relative weights, a point we verify in Section 4.9.

The central results. The three results that follow, together with the recomposition of Section 4.7, are the paper's contribution. The first establishes that the distributional sign of artificial intelligence is not fixed but turns on a single feature of the world.

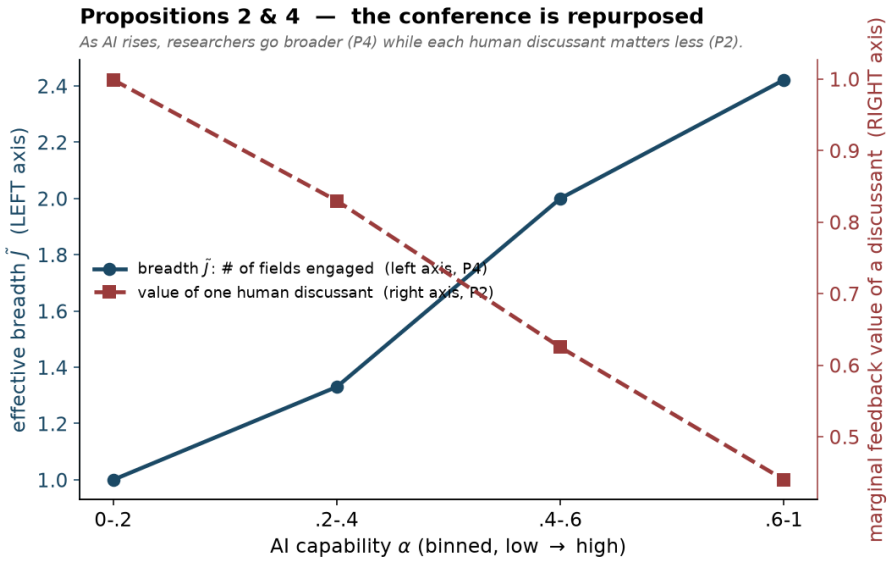


Figure 2: The conference is repurposed. As artificial-intelligence capability rises (horizontal axis), effective breadth $\tilde{J}$ increases (left axis, Proposition 4) while the marginal feedback value of a human discussant falls (right axis, Proposition 2).

The distributional knife-edge (Proposition 5). Figure 3 reports it: the dispersion of utility, measured by its Gini coefficient, plotted against the correlation between artificial-intelligence capability and prestige, with the no-artificial-intelligence world as a reference. The distributional effect of artificial intelligence is not a fixed quantity. When capability is distributed independently of standing, inequality is, if anything, slightly below its pre-artificial-intelligence level (0.245 against a baseline of 0.248): the paper-side levelling of Proposition 3 marginally compresses outcomes. As capability begins to track prestige, however, inequality rises steadily above the baseline, reaching 0.315 when capability is fully entrenched, and the elite–non-elite gap *more than doubles*, from 1.40 to 3.30 (standard errors below 0.006). The same researchers who command the dense rooms also wield the more effective tools, and the two advantages compound. One feature of the figure sharpens the interpretation: average capability actually *falls* as the correlation rises, from 0.40 to 0.19, because entrenchment concentrates capability in the prestigious minority and denies it to the majority. The widening is therefore not driven by more artificial intelligence in aggregate but by *who holds it*—the distribution, not the level. The honest reading is asymmetric. The widening when capability tracks prestige is large and, as the next subsection shows, robust; the compression when capability is democratised is genuine but thin, and the curve crosses its baseline very early. Artificial intelligence does not reliably level the academic field: at best, under full democratisation of capability, it is mildly equalising; the moment access

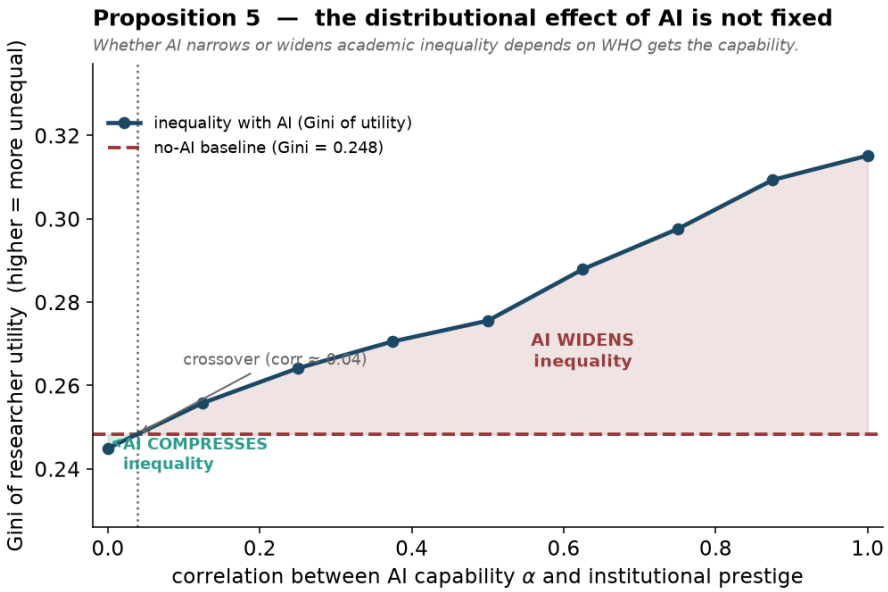


Figure 3: The distributional effect of artificial intelligence depends on who holds the capability. Inequality (Gini of utility) is plotted against the correlation between capability and prestige; the dashed line is the pre-artificial-intelligence baseline. Compression under democratised capability is thin; widening under entrenched capability is large.

to capability mirrors the existing hierarchy, it widens the very inequality it was hoped to dissolve.

The enclosure of the room (Proposition 6). Figure 4 shows what happens to those left behind as artificial-intelligence adoption spreads. As the adopting share rises from a small minority toward near-universality, the acceptance bar at the selective workshops climbs from about 8.8 to 9.5—the equilibrium effect of equation (11)—and the elite contact accruing to non-adopters falls steadily, from roughly 0.63 to 0.17 encounters (replication standard errors below 0.01), while their workshop attendance rate collapses from about nine per cent to two. The decisive observation is what does *not* change: the non-adopters' paper quality holds steady at the population mean of five throughout. Their papers are no worse than before; they are simply, and progressively, shut out of the room. This is enclosure rather than competition on the merits: as more researchers deploy artificial intelligence to crowd the dense venues, the congestion bound of Assumption 1 binds ever more tightly, and the access that artificial intelligence appears to offer is competed away for those who lag. Note that adopters' contact also falls in absolute terms as the room fills, but far less, so the gap between the two groups *widens* as adoption rises.

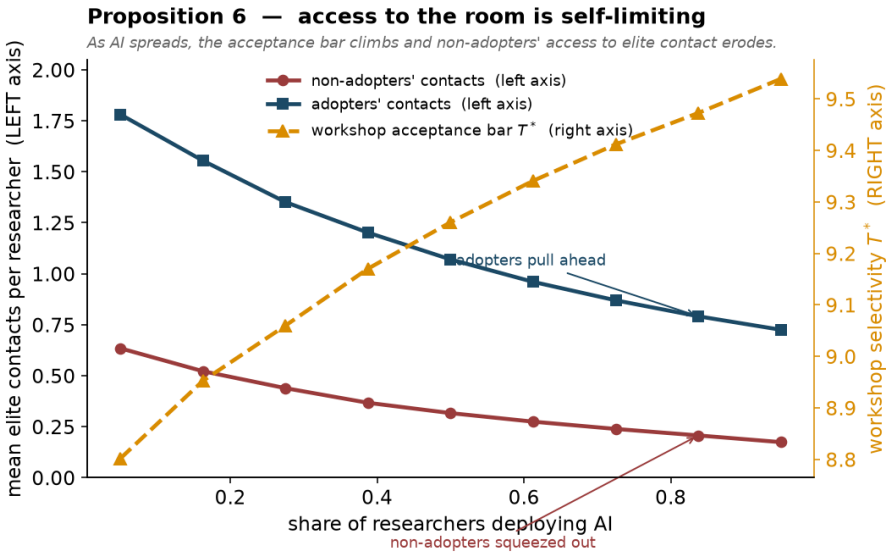


Figure 4: The enclosure of the room. As the share deploying artificial intelligence rises (horizontal axis), the workshop acceptance bar T^* climbs (right axis) and non-adopters' elite contact erodes (left axis), even though their paper quality is unchanged; adopters pull ahead.

The backfire of strategy (Proposition 7). The natural advice to a disadvantaged researcher is to be strategic—to aim for the dense, low-concurrency workshops of Proposition 1 rather than the generalist venue. Figure 5 asks whether that advice helps. It compares, on the same population at each adoption level, the adopter–non-adopter utility gap in a world where everyone competes strategically for the workshops against the gap in a world where everyone plays it safe and submits only to the generalist venue. The strategic gap is larger at *every* level of adoption: the option to target widens the gap by between roughly 0.09 and 0.26 of a utility unit, a difference that is always positive and several times the replication standard error (below 0.013), and is largest when adoption is partial, narrowing only as it approaches universality. The mechanism is the rising bar of equation (11) interacting with the contact ceiling of Assumption 1: strategic targeting routes everyone toward the same scarce rooms, where the higher fit and lower cost of the artificial-intelligence-equipped let them clear the inflated bar and capture the rationed contact. The obvious remedy thus backfires. Telling researchers to compete harder for the room does not close the gap that artificial intelligence opens; it widens it, because the room cannot expand to absorb the extra competition.

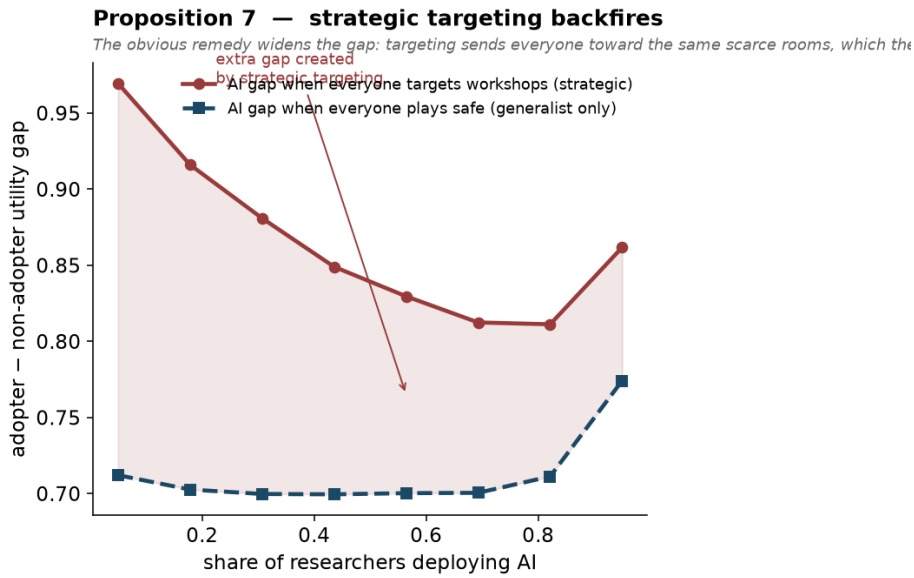


Figure 5: Strategic targeting backfires. At every level of artificial-intelligence adoption, the adopter–non-adopter utility gap is larger when everyone competes strategically for the dense workshops (solid) than when everyone submits only to the generalist venue (dashed); the shaded region is the extra gap created by strategy.

4.7 The lottery of the room

The three results so far concern the *gap* between groups. The model permits a sharper question: as artificial intelligence grows more capable, what *kind* of advantage decides who succeeds? Proposition 8 predicts a recomposition—that levelling the paper must let the residual weight of success fall on the room—and the simulation lets us measure it. We decompose the variance of career outcomes across researchers into the three sources of Proposition 8: *talent* (a researcher's own paper quality and using-skill), *position* (the prestige-linked advantages of internal feedback and of converting contact into outcome), and *luck* (the chance realisation of elite contact in the room). The decomposition uses the Lindeman–Merenda–Gold averaging of each source's incremental contribution over all orderings, which apportions explained variance fairly when the sources are correlated; together they account for between 0.84 and 0.88 of the variance in utility across the capability range, the remainder being curvature in the payoff.

We trace the shares as capability strengthens from absent to full, with capability distributed independently of prestige—the democratised pole, so that any shift is driven by capability itself and not by entrenchment. The pattern is unambiguous and, against the hope that artificial intelligence rewards merit, sobering. The share of outcome variance owed to *position* collapses, from 0.47 in the pre-artificial-intelligence world to 0.27 at full

capability: the levelling of Proposition 3 made quantitative, as artificial intelligence erodes the paper-side advantage that standing used to buy. But the freed weight does not pass to talent. Talent's share stays small throughout, between roughly 0.07 and 0.12; a good paper matters a little more when artificial intelligence lets it travel further, but it is never what mainly decides outcomes. The weight passes, instead, to *luck*: the share governed by the chance of who is in the room rises from 0.47 to 0.61, becoming the single largest determinant of who gets ahead. As artificial intelligence grows more capable, success turns less on standing, no more on talent, and ever more on fortune.

This result must be read alongside, not against, the knife-edge of Figure 3; the two speak to different objects. The knife-edge and the elite–non-elite gap concern the *average* difference between groups, which artificial intelligence widens when capability tracks prestige. The decomposition concerns the *composition* of the variance across all researchers. A group can retain a higher mean outcome while the spread of individual outcomes comes to be ruled by chance; the two are complementary faces of one mechanism. Artificial intelligence can enlarge the elite's average lead (Figure 3) and, at the same time, make any single researcher's fate more of a lottery (Figure 6). Both follow from the one fact that the room does not scale.

The individual-level signature is visible directly. Figure 6 plots realised success against talent for every researcher, at low and at high capability, with luck shown by colour. At both levels the most successful researchers are not the most talented but the luckiest: the correlation between success and luck (0.68 at low capability, 0.76 at high) far exceeds that between success and talent (0.34 and 0.42), and the gap between the two *grows* with capability. Among the top tenth of performers at high capability, the median researcher sits at the 78th percentile of luck but only the 73rd of talent. This is the pattern Pluchino et al. (2018) made famous in a generic model of success—that fortune, not merit, sorts the winners—and the present model shows that artificial intelligence does not dissolve it but, by neutralising the one channel effort could improve, deepens it.

Stacking the deck for talent. One might object that talent's small share is an artefact of the modest dispersion we assign to it. It is not. Tripling the standard deviation of paper quality—a ninefold increase in its variance—while re-centring the quality floor so that the fraction of papers it screens is held fixed, raises talent's share of outcome variance at full capability only from 0.12 to 0.24, even as luck's share remains above one-half (falling only from 0.61 to 0.54); across the sweep talent overtakes luck at no point. Even with the deck stacked heavily in talent's favour, it is the scarce room, not the paper, that continues to sort who succeeds.

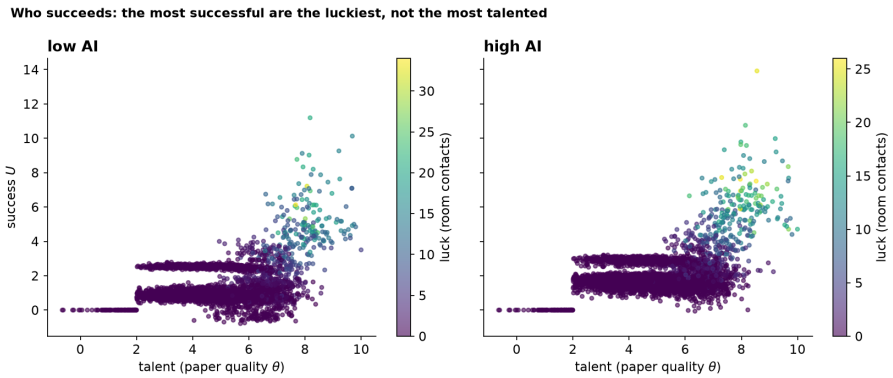


Figure 6: Success tracks luck, not talent. Realised career outcome (vertical axis) against talent (horizontal), at low (left) and high (right) artificial-intelligence capability, coloured by luck. The most successful researchers are the luckiest rather than the most talented, and the dominance of luck over talent strengthens as capability rises.

4.8 Locating the operating point: a calibration from large language models

The knife-edge of Figure 3 makes the distributional effect of artificial intelligence depend on one quantity—the correlation ρ between capability and prestige—without saying where the actual world sits along it. That correlation cannot be read from a published table, because it is generated by the choices of tens of thousands of individual researchers: whether a student behind an institutional prohibition uses artificial intelligence regardless, whether a researcher facing restricted regional access finds a way through. To locate the operating point we calibrate ρ by *silicon sampling*—eliciting those choices from large language models prompted to stand in for a heterogeneous population of researchers, a technique shown to reproduce population-level response distributions across the social sciences (Argyle et al., 2023; Bail, 2024), though with the documented caution that such samples can understate true variance (Bisbee et al., 2024). Eliciting this quantity by surveying junior researchers directly is, in any case, barely feasible—asking scholars to report how readily they would breach an institutional prohibition invites acute social-desirability bias and asks the vulnerable to confess to potential misconduct—so silicon sampling is less a shortcut than a way to recover a baseline willingness no honest survey easily could. We use it only to place an operating point on a curve the theory has already derived; the propositions and the knife-edge stand whether or not one accepts the calibration. We treat it, accordingly, as an *illustrative parameterisation*—a way to mark a plausible point on a derived curve—rather than the estimation of a structural behavioural parameter, which an elicited willingness, sensitive to the very access regime the analysis varies, could not in any case provide.

The procedure has two stages, both supplied as replication files. In the first, three open-weight models from different families each respond, in the guise of thirty researcher personas spanning fields, attitudes to rules, and competitive pressure, to how likely they would be to use artificial intelligence in the face of an institutional ban and in the face of restricted access. The elicited willingness is high but, decisively, roughly *flat* across prestige: researchers at strong and weak institutions report a similar readiness to route around a constraint. In the second stage that willingness is combined with the grounded structure of who actually *faces* each constraint—field-specific permission, bans on technical work being far rarer than on prose, and the provision of sanctioned tools by resourced institutions—to construct each researcher's capability as the weakest of access, skill, and permission. Because willingness is prestige-flat, the correlation that emerges is *structural*: it arises from who is granted sanctioned access, not from who is willing to seek it. The language models therefore carry a deliberately light load: they supply a single qualitative fact—that willingness does not climb with prestige—while the magnitude of ρ^* is set by the field- and institution-level access structure, which is disciplined by external evidence rather than by the models. Were the elicited willingness wrong in level but right in being flat, ρ^* would be unchanged.

The calibration places the operating point at $\rho^* \approx 0.09$ (elite capability 0.51 against non-elite 0.46), just past the knife-edge's crossing. Two findings follow, and the second matters more. First, at this point the dispersion of outcomes is statistically *indistinguishable* from the pre-artificial-intelligence world: a Gini of 0.248 against a baseline of 0.248, a difference far inside the replication standard error of 0.002 (Figure 7). Measured by inequality alone, one would conclude artificial intelligence had changed nothing. Second, and against that conclusion, the *composition* of success at the operating point has been transformed: the variance decomposition there returns shares of 0.15 to talent, 0.24 to position, and 0.61 to luck—position roughly halved from its pre-artificial-intelligence weight of 0.47, luck risen to dominate. The headline statistic is unmoved while the engine beneath it has been rebuilt. This is the paper's sharpest methodological claim: to gauge the distributional effect of artificial intelligence by an inequality index is to risk missing it entirely, because artificial intelligence can hold the level of inequality fixed while migrating its source from a hierarchy one can name to a lottery no one controls.

That the operating point sits at the crossing is a contingent fact about how access is currently distributed, not a law, and it is fragile in a definite direction. Figure 8 varies the single structural quantity the calibration rests on—the gap in sanctioned access between resourced and under-resourced institutions—and traces the consequence. When access is equal, artificial intelligence compresses inequality (a Gini of 0.236, below baseline); as the access gap widens, ρ^* rises from near zero toward 0.11 and the compression erodes, the

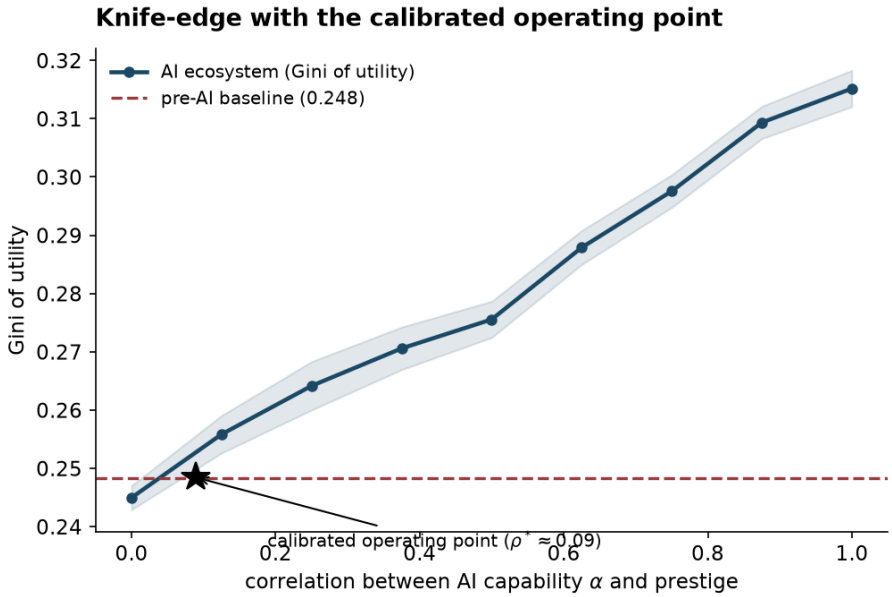


Figure 7: The calibrated operating point. The knife-edge of Figure 3, with the empirically calibrated point ($\rho^* \approx 0.09$) marked: inequality there is indistinguishable from the pre-artificial-intelligence baseline (dashed), even though the composition of success (Section 4.7) has by then shifted decisively toward luck.

Gini climbing back to and through its baseline. The neutrality observed today is therefore conditional: it holds while sanctioned access is only mildly unequal, and tips toward widening if the divergence in institutional provision grows.

Finally, the enclosure and backfire results survive the move from the stylised capability distribution to the calibrated one, confirming they are not artefacts of how capability was modelled. Under calibrated capability the elite contact accruing to non-adopters still erodes as adoption spreads—from about 0.59 to 0.09 encounters as the adopting share rises from a minority to near-universality—while the workshop bar climbs; and strategic targeting still backfires, the adopter–non-adopter gap remaining wider in the strategic world than the safe one at every adoption level, by a margin largest (0.46 of a utility unit) when adoption is partial. The paper's central results thus hold under a capability distribution disciplined by elicited behaviour, not only under the analytical sweep.

4.9 Robustness and validation

Sensitivity of the distributional result. Because Proposition 5 is the paper's headline, we re-estimate it under a battery of alternative calibrations; Table 3 reports, for each, the dispersion of utility under democratised and entrenched capability against that calibration's

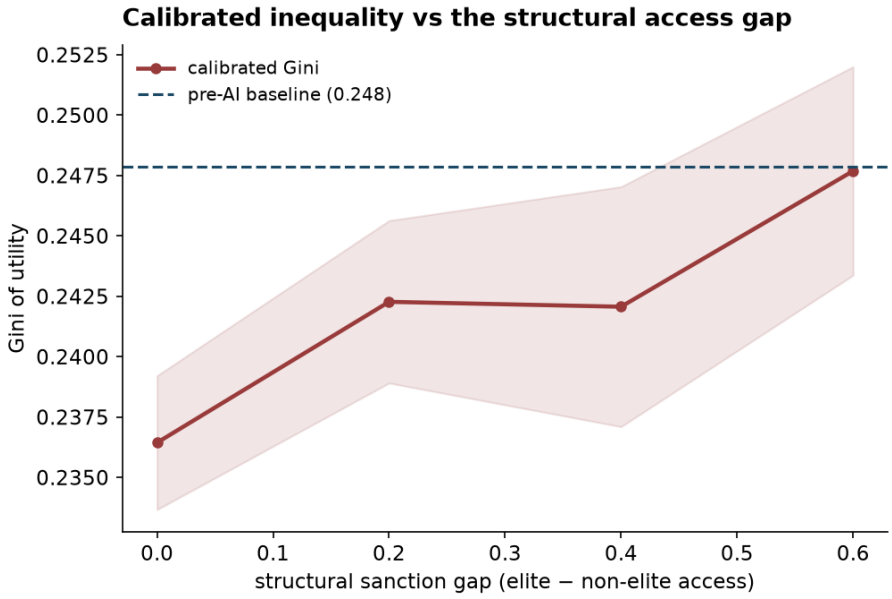


Figure 8: Inequality at the operating point is contingent on the structural access gap. As the gap in sanctioned artificial-intelligence access between resourced and under-resourced institutions widens (horizontal axis), the calibrated correlation ρ^* rises and the dispersion of outcomes climbs from below the pre-artificial-intelligence baseline (dashed) toward and beyond it.

own baseline. Two patterns emerge. The *widening* of inequality under entrenched capability is universal: it appears in every calibration we examined, including ones that weaken the knowledge channel, strengthen the room, raise submission costs, and reallocate generalist contacts across subfields. The *compression* under democratised capability is conditional: it holds at the baseline and when the knowledge channel is at least as important as the career channel, but it disappears when the knowledge weight is halved or submission is made costly; a further check not reported in Table 3—enriching the venue universe with additional subfields, which adds room capacity—removes the compression as well. This is not a weakness of the result but its precise content: artificial intelligence compresses inequality only to the extent that its paper-side levelling outweighs the prestige-linked career channel, and widens it otherwise. The qualitative knife-edge of Figure 3 is robust; the location of the crossing is calibration-dependent, and under several reasonable calibrations there is no compression at all.

Behavioural validation. That the decision rule is well-posed is confirmed by the behaviour it produces in equilibrium, none of which is imposed. The distribution of submissions is bimodal. A majority of researchers—about two-thirds—rationally submit nothing

Table 3: Robustness of the distributional result. Each row reports the Gini of utility under democratised ($\rho = 0$) and entrenched ($\rho = 1$) capability against that calibration's own no-artificial-intelligence baseline. Each row is re-estimated on its own set of replications, so the baseline row differs from the production values underlying Figure 3 by at most 0.003, inside the replication standard errors reported in the text. Widening is universal; compression is conditional.

Calibration	Baseline	$\rho = 0$	$\rho = 1$	Compresses?
Baseline	0.245	0.242	0.318	yes
Knowledge weaker ($\eta = 1$)	0.193	0.230	0.267	no
Knowledge stronger ($\eta = 4$)	0.308	0.219	0.369	yes
Room stronger ($R^E = 1$)	0.275	0.270	0.348	yes
Costlier submission ($\kappa_0 = 0.3$)	0.175	0.178	0.260	no
Generalist contacts dispersed	0.244	0.240	0.316	yes

at all: those whose posterior places them near the quality floor, for whom no venue clears its cost once artificial-intelligence-inflated competition has driven the equilibrium bar above the quality of a marginal applicant. Among the roughly one-third who do submit, portfolios are broad—about five venues—and nearly all, some 99%, include the generalist venue as a safety; averaged over the whole population, submissions run to 1.7 per researcher. This abstention is sharply sorted by confidence, climbing from about 7% in the most confident quartile of researchers to near-universal in the least, so that the least confident withdraw almost entirely while the confident continue to enter. The model thus reproduces, without being instructed to, the empirically familiar pattern in which inflated competition prices the least confident out of the market. The level of abstention is sensitive to the capacity of the venue system and should be read as qualitative.

Methodological checks. Finally, the conclusions do not depend on the approximations disclosed in Section 4.3. Reallocating the contacts made at the generalist venue across subfields, rather than attributing them to the researcher's home subfield, leaves every result unchanged (Table 3, final row). The repurposing result of Proposition 2 is a property of the concavity of the knowledge function and is unaffected by the relative weight on the career channel, as the comparison across knowledge weights in Table 3 confirms. And the equilibrium of Section 4.4 is reached from a naive prior, so the crowding it describes cannot be an artefact of optimistic initial beliefs.

4.10 Synthesis

Taken together, the experiments support a single, two-sided conclusion. On the paper side of the conference, artificial intelligence is a levelling technology: it substitutes for the feedback that privileged researchers already possess, and it does so most for those who have least. On the room side, it is not. Physical proximity to senior scholars remains scarce and congestible, the focused workshop remains structurally superior to the generalist venue for acquiring it, and access to it is governed by the prestige-linked capacity to convert contact into outcomes. This is complementarity at work: by making the paper abundant, artificial intelligence shifts the return to the one input it cannot supply, so that the scarce complement—the room—rather than the cheapened task comes to govern who succeeds. Whether artificial intelligence narrows or widens academic inequality therefore turns on a single empirical quantity—the degree to which capability is distributed with, or against, the existing hierarchy. Four findings make this concrete. The distributional knife-edge: the widening that follows when capability mirrors prestige is robust and large, while the compression that democratised capability might deliver is thin and conditional. The enclosure of the room: rising adoption squeezes non-adopters out of the dense venues even though their papers are no worse. And the backfire of strategy: the obvious remedy—competing harder for the room—widens the gap rather than closing it, because the room cannot expand to absorb the competition. And the lottery of the room: as artificial intelligence grows more capable it erodes the advantage of position but hands the freed weight to luck rather than to talent, so that at the operating point our calibration locates—where the measured level of inequality is itself unchanged—the composition of success has nonetheless shifted decisively from standing toward chance. The room, in short, does not scale; and because it does not, artificial intelligence redraws not only how wide academic inequality is, but what it is made of.

5 Discussion and Conclusion

The paper began from an asymmetry. A research conference is at once a place to improve a paper and a place to be heard by the people who shape careers, and artificial intelligence acts on these two functions in opposite ways: it makes the paper cheap to produce and improve, levelling that dimension in favour of those who start with least, while leaving wholly untouched the scarce, congestible access to senior attention on which durable advantage rests. From this single asymmetry four results followed. The distributional effect of artificial intelligence is a knife-edge, equalising when capability is shared independently of prestige and sharply widening when capability tracks it. The access that artificial intelligence seems to

open is an enclosure, competed away as adoption spreads, so that non-adopters are crowded out of the dense rooms though their papers are no worse. The obvious remedy backfires, because telling researchers to compete more strategically for the room sends everyone toward the same scarce seats, which the artificial-intelligence-equipped win. And beneath the level of inequality, the composition of advantage recomposes: as capability grows, the prestige-linked advantage of position gives way not to talent but to luck, until success turns on who happens to be in the room. We close by drawing out what these findings imply for the people who must act on them, what the analysis cannot settle, and how it might be tested.

What academic inequality is made of. The fourth result is the one we think most consequential, because it reframes the first three. When the model is calibrated to elicited adoption behaviour and read at that operating point, the dispersion of outcomes barely moves: measured inequality is statistically indistinguishable from the world before artificial intelligence, and an inequality index alone would report that nothing had happened. Yet beneath that unchanged number the *composition* of success has been rebuilt. The share of outcome variance owed to *position*—the prestige-linked advantage of sharper feedback and easier conversion of contact—falls by roughly half, the share owed to *luck*, the chance of who is in the room, rises to dominate, and the share owed to *talent* stays small throughout. Artificial intelligence holds the level of inequality fixed while migrating its source from a hierarchy one can name to a lottery no one controls. The methodological lesson is immediate and, we think, general beyond this setting: to judge the distributional effect of artificial intelligence by a summary measure of dispersion is to risk missing the effect entirely, because the technology can leave inequality's level untouched while changing what it is made of. The substantive lesson cuts across two traditions at once. It sharpens the claim of Pluchino et al. (2018) that fortune, not talent, sorts the successful: artificial intelligence does not dissolve the role of luck but enlarges it, precisely by neutralising the one channel—the paper—that effort and ability could improve. And it qualifies the cumulative-advantage tradition of Merton (1968): on the paper side, cheap capability *erodes* the Matthew effect, stripping away the compounding feedback advantage that standing once conferred—but what fills the vacated space is not merit. It is chance, mediated by a room that no technology has yet learned to enlarge.

Doctoral training and the advice we give. The first implication is cautionary, and it is addressed to advisors and to the students they counsel. The instinct to tell a disadvantaged student that artificial intelligence will level the field is correct on the paper side and incomplete in a way that matters: a polished manuscript clears the first hurdle but

does nothing for the second, and our results show the second is where the hierarchy is decided. Worse, the most natural strategic advice—aim for the selective workshops where the right people gather—is collectively self-defeating, because when everyone follows it the bar at those venues rises and the advantage flows to the already-advantaged. What helps the disadvantaged is not better paper-side tooling, which they will share with everyone, but structured access to the room and a higher rate of converting the contact they do make: deliberate introductions, sponsored attendance, formal mentoring that does not depend on winning a congested competition. The model locates the durable advantage in the conversion of contact into outcomes and in proximity itself; interventions that leave those untouched will not move careers, however much they improve papers.

The design of conferences. The second implication is constructive, and it is addressed to the scholarly societies and programme committees that design venues. If the room is a congestible commons, it can be designed—expanded, protected, and rationed more equitably than informal proximity allows. Our structural result that small, single-track, topically dense workshops dominate large multi-track conferences for accumulating elite contact (Lopez de Leon and McQuillin, 2020; Chai and Freeman, 2019; Boudreau et al., 2017; Catalini, 2018) suggests that a field which cares about access should invest in the former rather than concentrate prestige in a handful of mega-conferences whose parallel sessions dilute every encounter. Because artificial intelligence repurposes the conference—eroding its value as a source of feedback while preserving its value as a site of contact—conferences can lean deliberately into the function that remains scarce, designing for structured interaction (assigned discussants, mentoring sessions, scheduled access to senior scholars) rather than for presentation alone. And because the acceptance bar will rise as cheaper submission inflates competition (Gartenberg et al., 2026), fields that wish to preserve early-career access may need to protect it explicitly, through dedicated junior tracks or reserved capacity, rather than trusting an open contest whose threshold artificial intelligence is driving upward.

The governance of artificial-intelligence access. The third implication concerns institutional policy toward the technology itself, and it follows from two features of the model. Because capability is a weakest-link of access, skill, and permission (Kremer, 1993; Bresnahan et al., 2002), a policy that subsidises tools while leaving a departmental prohibition or a skills deficit in place changes nothing; the binding constraint differs across institutions, and only relieving the binding one raises capability. More fundamentally, our knife-edge result identifies the decisive quantity as the *correlation* between capability and prestige. Whether artificial intelligence equalises or entrenches is therefore not fixed by the technology but

chosen, collectively, through the pattern of adoption: if elite institutions adopt faster, train better, and permit more freely while others lag, capability comes to mirror prestige and inequality widens; if access, training, and permission are extended first to the less-resourced, capability is decoupled from standing and the equalising side of the knife-edge prevails. Both tendencies are visible in practice. Frontier models are cheap and broadly accessible, which pushes toward the democratised pole; yet better-resourced institutions can also purchase the secure, privacy-compliant enterprise licences that grant high access and permission, train their members to exploit them, and face fewer of the blanket prohibitions that data-protection or research-integrity concerns prompt at less-resourced institutions—each raising a different one of the three binding inputs precisely where prestige is already high. The decisive difference is therefore not price—a personal subscription costs little—but compliance: data-protection and research-integrity rules bar feeding unpublished or proprietary material into a consumer tool, so the cheap model can lawfully be used only for surface tasks, while the deep, data-intensive use that elevates a paper's substance requires exactly the sanctioned, privacy-compliant access the well-resourced enjoy. Which tendency dominates is the empirical question our knife-edge identifies as decisive. The policy goal that follows is specific and unusual—not to maximise adoption, but to keep capability independent of prior advantage.

Limitations. These conclusions are the output of a calibrated model, not an empirical estimate, and several of its abstractions bound their reach. We treat artificial intelligence as improving the paper and its fit to a venue but never as substituting for the room; if artificial-intelligence-mediated or virtual proximity were to become a genuine substitute for face-to-face contact, the central asymmetry would weaken—and with it the recomposition of success toward luck, which in that sense is a consequence of the premise that the room lies beyond the technology's reach as much as a finding about the world. The degree to which that premise holds is an empirical question we have assumed rather than answered, and the natural test is to let artificial intelligence partially substitute for the room and trace how far the dominance of luck survives. The room itself is modelled as a single congestible commons, whereas real scholarly networks are repeated, multiplex, and partly online. The decision is a rich single-period portfolio choice with beliefs that settle across iterations, but it is not a full model of a career: we do not endogenise the quality of the paper, the entry and pricing decisions of conference organisers, or the long-run feedback from placement back to prestige. Prestige enters through two channels—internal feedback and the conversion of contact—that stand in for a richer reality. In particular, prestige shapes how contact is *converted* but not, in our model, how it is *made*: elite attention is sought through congestion and realised by chance, whereas senior scholars who used artificial intelligence to triage

whom they engage would make contact itself prestige-assortative, giving the room a second positional channel on the demand side. That the room remains serendipitous is consistent with the colocation evidence we rely on (Lopez de Leon and McQuillin, 2020; Chai and Freeman, 2019)—conferences matter precisely because they create contact the networked default does not—but the strength of any such filtering is unmeasured. Were it strong, the recomposition of Section 4.7 would carry position rather than luck; what would not change is that talent's share stays small, so artificial intelligence would still fail to make success meritocratic, altering only which non-merit advantage prevails. The composition of the room is in this sense itself a knife-edge—an open room enthrones luck, a filtered room entrenches position, and neither rewards talent—and tracing it as artificial intelligence reaches the demand side is a natural companion to the substitution test already noted. A second demand-side force points the same way: by making polish cheap, artificial intelligence degrades polish as a signal of quality, so the scarce attention in the room comes to be rationed ever more tightly by the deep substance the technology cannot fake—substance that, as argued above, is itself gated by sanctioned access. Far from easing the room, artificial intelligence may then make it more ruthless, and ruthless in a way that again rewards position over talent. The operating point, finally, is located not from field data but by eliciting adoption choices from large language models, a calibration that inherits the limits of that method—few models, assumed frequencies for the traits that gate access, and a stated likelihood read as a probability—and that may understate the true spread of behaviour (Bisbee et al., 2024); we therefore treat it as locating a plausible point on a theory-derived curve, not as measuring where the world sits. None of these simplifications is innocent, but each is conservative with respect to our central claim: they concern the magnitude and the precise location of the knife-edge, not the existence of the asymmetry that produces it, and the robustness analysis of Section 4.9 shows the widening result to be insensitive to the calibration even where the compression result is not.

Predictions and future work. Because the analysis is mechanistic, it yields predictions that data can confront, and stating them is the natural next step. The knife-edge implies that the distributional effect of artificial intelligence on early-career outcomes should vary systematically with the correlation, across fields and institutions, between access to capable tools and institutional prestige—equalising where the two are decoupled and entrenching where they coincide. The enclosure implies that the selectivity of small workshops should rise with adoption while the measured quality of the work they exclude does not, a pattern testable with submission and acceptance records combined with independent quality measures. The recomposition implies a third, subtler prediction: as adoption deepens, the power

of institutional pedigree to predict early-career success should wane relative to unsystematic variation—the statistical signature of a hierarchy giving way to a lottery. The growing body of evidence on artificial intelligence in scientific production (Agrawal et al., 2024; Baruffaldi and Poege, 2025) and the editorial data now accumulating at journals (Gartenberg et al., 2026) make such tests feasible in a way they were not even recently. Beyond testing, the framework extends naturally to hybrid and virtual venues as partial substitutes for the room, to multi-period career dynamics, and to the endogenous supply of conferences—each a setting in which the same question recurs: what happens to a scarce, access-like good when the complementary, information-like good beside it suddenly becomes free.

Conclusion. Artificial intelligence has not abolished scarcity in academic life; it has relocated it. When the paper was the binding constraint, levelling the paper would have levelled the field. But the paper is no longer the binding constraint—the room is, and the room does not scale. And because it does not, levelling the paper does not put merit in the room's place; it puts chance there: the advantage of pedigree, dissolved on the page, reappears as the luck of the encounter. The promise that cheaper production will democratise science is, on the evidence of our model, conditional at best and self-undermining at worst: it holds only where access to the technology is kept independent of prior advantage, and it is reversed by the very strategies that individuals will rationally adopt. A field that wishes artificial intelligence to widen access rather than narrow it cannot simply trust the technology to do so. It must build access to the room—deliberately, and against the grain of what unaided competition will produce.

There is a deeper way to see what artificial intelligence has done, and it inverts the hope with which we began. In the manner Baumol (1967) taught us to expect of a progressive sector racing ahead of a stagnant one, cheapening the paper has not diminished the room but raised its relative worth—making the one good a scholarly career cannot manufacture the most valuable thing it has left. Democratising the paper, in other words, is precisely what makes the room decisive. The frontier of academic inequality has moved accordingly: it no longer runs between those who can produce good work and those who cannot, for the machine has narrowed that gap, but between those who can enter the room and those who cannot. Whether science grows more open or more closed in the age of artificial intelligence will be settled there, and not on the page.

A Proofs

This appendix proves the propositions of Section 3. Propositions 1–4 and Proposition 6 are established in full. For Propositions 5 and 7 we prove the stated qualitative claim in full and name explicitly the single quantity each leaves to calibration—the location of the knife-edge, and the magnitude of the general-equilibrium amplification, respectively—which Section 4 supplies. Proposition 8 is directional by construction: we give its mechanism analytically and leave the variance-share magnitudes, including the robustness of talent's small share to a large increase in talent dispersion, to the simulation. Two elementary inequalities do most of the work.

Two difference inequalities

Throughout, g is increasing and concave, and we write $\Sigma_i(\alpha) = \sum_{j=1}^{\tilde{J}_i(\alpha)} \Lambda_{ij}$ for a researcher's counted elite contact, where $\tilde{J}_i(\alpha) = \min\{J(\alpha_i), \bar{J}\}$ and $J(\cdot)$ is nondecreasing.

(A) Room complementarity. Because each $\Lambda_{ij} \geq 0$ and $\tilde{J}_i(\alpha)$ is nondecreasing in α_i, the sum Σ_i is nondecreasing in α_i. Hence for any $\alpha_H \geq \alpha_L$ and $\omega_H \geq \omega_L$,

$$\big(\omega_H - \omega_L\big)\, R^E \left(\Sigma_i(\alpha_H) - \Sigma_i(\alpha_L)\right) \;\geq\; 0, \tag{14}$$

so the career payoff $\omega_i R^E \Sigma_i(\alpha_i)$ has *increasing differences* in (α_i, ω_i): capability is worth more, through the room, to a researcher who converts contact better.

(B) Paper substitution. Since g is concave, for every $t \geq 0$ the increment $x \mapsto g(x + t) - g(x)$ is nonincreasing in x. Applying this with $t = \alpha_i \bar{\tau}^{\mathrm{AI}}$ to the knowledge payoff $\eta\, g\big(\tau_i^{\mathrm{int}} + \alpha_i \bar{\tau}^{\mathrm{AI}} + c\big)$ (with $c \geq 0$ the attended discussant precision), the artificial-intelligence knowledge gain

$$G(\tau_i^{\mathrm{int}}) = \eta\Big[g\big(\tau_i^{\mathrm{int}} + \alpha_i \bar{\tau}^{\mathrm{AI}} + c\big) - g\big(\tau_i^{\mathrm{int}} + c\big)\Big] \tag{15}$$

is nonincreasing in τ_i^{int}: knowledge has *decreasing differences* in $(\alpha_i, \tau_i^{\mathrm{int}})$. Capability is worth less, through the paper, to a researcher whose home feedback is already sharp.

The mechanism (Propositions 1–4)

Proof of Proposition 1 (Concurrency ranking). The rate $\lambda_c = \delta_c N_c / m_c$ is strictly decreasing in m_c and strictly increasing in δ_c. As $e^{-\lambda}$ is strictly decreasing, $p_c^E = 1 - e^{-\lambda_c}$ is strictly increasing in λ_c; composing, p_c^E is strictly decreasing in m_c and strictly increasing in δ_c. For the ranking, the workshop $(\delta, N, m) = (0.1, 60, 1)$ has $\lambda = 6$ while the flagship $(0.02, 600, 24)$ has $\lambda = 0.5$; since $6 > 0.5$ and p^E increases in λ, the small single-track venue delivers the

higher contact probability, even though the flagship holds $\delta N = 12 > 6$ relevant elites in absolute terms. □

Proof of Proposition 2 (AI repurposes the conference). Let venue c add discussant precision $\tau_c^{\text{ext}} = d_c \bar{\tau}^{\text{ext}} \geq 0$ to a feedback stock $x_i = \tau_i^{\text{int}} + \alpha_i \bar{\tau}^{\text{AI}} + (\text{other attended})$; its marginal knowledge value is $\Delta_c = \eta\big[g(x_i + \tau_c^{\text{ext}}) - g(x_i)\big]$. Since x_i increases with α_i and, by concavity, $x \mapsto g(x + \tau_c^{\text{ext}}) - g(x)$ is nonincreasing, Δ_c is nonincreasing in α_i: the feedback value of every venue falls as capability rises. Moreover $g(x + \tau_c^{\text{ext}}) - g(x) \to 0$ as $\tau_c^{\text{ext}} \to 0$, so the decline bites first and hardest at weak-discussant venues (small d_c), whose feedback rationale vanishes; the trip is justified on feedback grounds only where τ_c^{ext} is large enough that Δ_c still exceeds the opportunity cost of attending, and elsewhere the reason to attend shifts to the career channel. □

Proof of Proposition 3 (Bounded feedback democratisation). Immediate from (15): with capability held fixed and equal across groups, the artificial-intelligence knowledge gain $G(\tau^{\text{int}})$ is nonincreasing in τ^{int}. Because elite students begin with sharper home feedback, $\tau_N^{\text{int}} < \tau_E^{\text{int}}$ gives $G(\tau_N^{\text{int}}) \geq G(\tau_E^{\text{int}})$, with strict inequality wherever g is strictly concave on the relevant range. The same feedback therefore adds (weakly) more knowledge for the non-elite researcher, and can exceed weak, narrow supervision while merely topping up supervision that is already strong. □

Proof of Proposition 4 (From depth to breadth). $\tilde{J}_i(\alpha) = \min\{J(\alpha_i), \bar{J}\}$ is the minimum of a nondecreasing function and a constant, hence nondecreasing in α_i: strictly increasing on $\{\alpha_i : J(\alpha_i) < \bar{J}\}$ and equal to $\bar{J}$ once $J(\alpha_i) \geq \bar{J}$. Because contacts in distinct subfields enter (6) additively while within-subfield contacts are redundant, career value is maximised by placing attendance in the $\tilde{J}_i$ highest-Λ_{ij} subfields; as $\tilde{J}_i$ rises with α_i this set widens, until the field's tolerance $\bar{J}$ caps it. □

The central results (Propositions 5–8)

Proof of Proposition 5 (Prestige and access: the knife-edge). Let ρ index the correlation between capability α and prestige, holding the marginal distributions of α and of prestige fixed; raising ρ is then a mean-preserving reallocation of capability toward the prestigious. Prestige raises both conversion ω and home feedback τ^{int}. Consider shifting a unit of capability from a non-elite researcher to an elite one. Its effect on the elite–non-elite mean gap $G(\rho) = \mathbb{E}[U \mid E] - \mathbb{E}[U \mid N]$ is the difference in the marginal value of capability between the two researchers. Capability enters U through a common part (fit f_{ic}, cost κ_i, curriculum-vitae value), which cancels in the difference, plus the room term and the paper term. By (14) the room term contributes a nonnegative amount to this difference (the elite convert

at higher ω); by (15) the paper term contributes a nonpositive amount (the elite start with higher τ^{int}). Hence

$$\frac{\partial G}{\partial \rho} = \underbrace{(\text{room complementarity})}_{\geq 0} + \underbrace{(\text{paper substitution})}_{\leq 0}, \tag{16}$$

and the distributional sign is the balance of the two. Where capability is independent of prestige ($\rho = 0$) no correlation engages the room complementarity, and the paper-side levelling of Proposition 3 dominates, so artificial intelligence narrows the gap; where capability tracks prestige (ρ large) the room complementarity compounds the standing advantage and artificial intelligence widens it. The crossover—the knife-edge—is the calibration-dependent value of ρ at which the two terms in (16) cancel; it cannot be signed from the primitives, and Section 4 locates it. □

Proof of Proposition 6 (Self-limiting access). By Assumption 1 each competitor at venue c receives $m(M_c) = \min\{\lambda_c, \bar{\Lambda}_c/M_c\}$. As a function of the competing mass M_c, the term $\bar{\Lambda}_c/M_c$ is strictly decreasing, so $m(\cdot)$ is nonincreasing everywhere and strictly decreasing once $\bar{\Lambda}_c/M_c < \lambda_c$, i.e. for $M_c > \bar{\Lambda}_c/\lambda_c$. By Propositions 1 and 4, rising capability draws researchers toward the low-concurrency (high-λ_c) venues, so the mass M_c at those venues increases with adoption. Their per-capita contact therefore holds while adoption is light ($M_c \leq \bar{\Lambda}_c/\lambda_c$) and strictly erodes once adoption pushes M_c past $\bar{\Lambda}_c/\lambda_c$: the access advantage of the dense venue is self-limiting. □

Proof of Proposition 7 (Strategic targeting amplifies the gap). Take two researchers identical except in capability, $\alpha_H > \alpha_L$ (hence higher fit and lower cost), with common conversion ω, choosing between a dense workshop W and the generalist G, with $\lambda_W > \lambda_G$. Holding bars fixed, the expected room payoff from venue c is $\omega R^E \lambda_c \, p_c^{\text{acc}}(\alpha)$, where $p_c^{\text{acc}}(\alpha) = \Phi\big(z_c(\alpha)\big)$ and $z_c(\alpha) = \theta + f_c^0 + \alpha\bar{\phi} - T_c$. Its capability gradient is $\omega R^E \lambda_c \, \bar{\phi}\, \phi\big(z_c(\alpha)\big)$, with $\phi = \Phi'$. The gain from targeting the workshop rather than the generalist, $\Gamma(\alpha)$, is therefore increasing in α whenever the workshop's gradient exceeds the generalist's, that is whenever

$$\lambda_W \, \phi\big(z_W(\alpha)\big) \;>\; \lambda_G \, \phi\big(z_G(\alpha)\big) \qquad \Longleftrightarrow \qquad \frac{\lambda_W}{\lambda_G} \;>\; \frac{\phi\big(z_G(\alpha)\big)}{\phi\big(z_W(\alpha)\big)}. \tag{$\star$}$$

Condition ($\star$) says the workshop's contact-density advantage outweighs any adverse curvature from its higher bar, and its domain can be located exactly. Writing $\Delta \equiv z_G(\alpha) - z_W(\alpha)$, a constant independent of α that is positive whenever the workshop's bar is the higher (at the calibrated venues $\Delta = T_W - T_G \approx 2$, baseline fit being common), condition ($\star$) is equivalent

to

$$z_W(\alpha) \;>\; -\frac{\ln(\lambda_W/\lambda_G)}{\Delta} \;-\; \frac{\Delta}{2},$$

which at the calibrated venues ($\lambda_W/\lambda_G \approx 5.5$) is $z_W(\alpha) > -1.85$. The condition therefore holds precisely for researchers whose quality-plus-fit lies within roughly two units of review noise of the workshop bar—the population for whom targeting the workshop is a live option—and fails far below the bar, where the targeting option is in any case worth little. Under ($\star$) the more capable researcher gains strictly more from targeting, so offering the strategic option widens the capability gap among researchers within reach of the workshop. In general equilibrium two forces sharpen this: the workshop bar T_W^* rises with the mass of high-fit applicants who target it (11), and its contact is rationed toward those who clear it (Assumption 1). The partial-equilibrium amplification under ($\star$) is thus a lower bound on the full effect, whose magnitude Section 4 quantifies. □

Proposition 8 (The lottery of the room)—mechanism. Sort the determinants of U_i into *talent* (θ_i, using-skill s_i), *position* (τ_i^{int}, ω_i), and *luck* (the Poisson realisation of elite contact, (3)). Artificial intelligence acts only through capability α_i, and hence only on the paper. Through $\tau_i^{\text{AI}} = \alpha_i \bar{\tau}^{\text{AI}}$ it raises feedback and, by (15) and Proposition 3, adds most where home feedback is weakest, compressing the paper-side dispersion that position generates. It touches talent only through the bounded fit and breadth channels ($f_{ic} = f_c^0 + \alpha_i \bar{\phi}$ and $\tilde{J}_i$), which let a good paper travel marginally further but do not scale with θ_i. On the room it has no purchase (Assumption 1). Consequently, holding fixed how capability is distributed across the hierarchy, a rise in average capability $\bar{\alpha}$ reduces the part of $\text{Var}(U_i)$ contributed by position and, since it cannot act on the room, raises the share contributed by luck; the part contributed by talent is small to begin with, because talent enters U_i only through the gated, bounded paper channel. The direction of the shift—from position toward luck—is thus analytic; its magnitude, and in particular the finding that talent's share stays small even when talent's dispersion is inflated several-fold, is established numerically in Section 4.

References

Ajay Agrawal, John McHale, and Alexander Oettl. Superhuman science: How artificial intelligence may impact innovation. *Journal of Evolutionary Economics*, 33(5):1473–1517, 2023. doi: 10.1007/s00191-023-00845-3.

Ajay Agrawal, John McHale, and Alexander Oettl. Artificial intelligence and scientific discovery: A model of prioritized search. *Research Policy*, 53(5):104989, 2024. doi: 10.1016/j.respol.2024.104989.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351, 2023.

Christopher A. Bail. Can generative AI improve social science? *Proceedings of the National Academy of Sciences*, 121(21):e2314021121, 2024. doi: 10.1073/pnas.2314021121.

Stefano Baruffaldi and Felix Poege. Like stars: How firms learn at scientific conferences. *Management Science*, 71(3):2056–2078, 2025. doi: 10.1287/mnsc.2022.01373.

William J. Baumol. Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review*, 57(3):415–426, 1967.

James Bisbee, Joshua D. Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M. Larson. Synthetic replacements for human survey data? The perils of large language models. *Political Analysis*, 32(4):401–416, 2024. doi: 10.1017/pan.2024.5.

Kevin J. Boudreau, Tom Brady, Ina Ganguli, Patrick Gaule, Eva Guinan, Anthony Hollenberg, and Karim R. Lakhani. A field experiment on search costs and the formation of scientific collaborations. *Review of Economics and Statistics*, 99(4):565–576, 2017.

Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt. Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1):339–376, 2002. doi: 10.1162/003355302753399526.

Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942, 2025. doi: 10.1093/qje/qjae044.

Christian Catalini. Microgeography and the direction of inventive activity. *Management Science*, 64(9):4348–4364, 2018. doi: 10.1287/mnsc.2017.2798.

Sen Chai and Richard B. Freeman. Temporary colocation and collaborative discovery: Who confers at conferences. *Strategic Management Journal*, 40(13):2138–2164, 2019. doi: 10.1002/smj.3062.

Partha Dasgupta and Paul A. David. Toward a new economics of science. *Research Policy*, 23(5):487–521, 1994. doi: 10.1016/0048-7333(94)01002-1.

Thomas A. DiPrete and Gregory M. Eirich. Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology*, 32:271–297, 2006. doi: 10.1146/annurev.soc.32.061604.123127.

Glenn Ellison. The slowdown of the economics publishing process. *Journal of Political Economy*, 110(5):947–993, 2002a. doi: 10.1086/341868.

Glenn Ellison. Evolving standards for academic publishing: A q-r theory. *Journal of Political Economy*, 110(5):994–1034, 2002b.

Josef Falkinger. Limited attention as a scarce resource in information-rich economies. *The Economic Journal*, 118(532):1596–1620, 2008. doi: 10.1111/j.1468-0297.2008.02182.x.

Claudine Gartenberg, Sharique Hasan, Alex Murray, and Lamar Pierce. More versus better: Artificial intelligence, incentives, and the emerging crisis in peer review. *Organization Science*, 37(3):795–812, 2026. doi: 10.1287/orsc.2026.ed.v37.n3.

John C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B*, 41(2):148–164, 1979. doi: 10.1111/j.2517-6161.1979.tb01068.x.

Avi Goldfarb and Catherine Tucker. Digital economics. *Journal of Economic Literature*, 57 (1):3–43, 2019. doi: 10.1257/jel.20171452.

Michael Kremer. The O-Ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575, 1993. doi: 10.2307/2118400.

Fernanda L. Lopez de Leon and Ben McQuillin. The role of conferences on the pathway to academic impact: Evidence from a natural experiment. *Journal of Human Resources*, 55 (1):164–193, 2020. doi: 10.3368/jhr.55.1.1116-8387R.

James G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87, 1991. doi: 10.1287/orsc.2.1.71.

Robert K. Merton. The Matthew effect in science. *Science*, 159(3810):56–63, 1968. doi: 10.1126/science.159.3810.56.

Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023. doi: 10.1126/science.adh2586.

Sharon Oster. The optimal order for submitting manuscripts. *American Economic Review*, 70(3):444–448, 1980.

Alessandro Pluchino, Alessio Emanuele Biondo, and Andrea Rapisarda. Talent versus luck: The role of randomness in success and failure. *Advances in Complex Systems*, 21(3–4): 1850014, 2018. doi: 10.1142/S0219525918500145.

Sherwin Rosen. The economics of superstars. *American Economic Review*, 71(5):845–858, 1981.

Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973. doi: 10.2307/1882010.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.