Growing Downward: Persistent Inequality in Publishing in Economics*

Ricardo Dahis†

June 8, 2026

Abstract

This paper documents new facts about concentration in publishing in economics. First, the profession grows *downward*. The number of economists grew almost sixfold since 1990, but new entrants publish in lower-tier journals while incumbents hold the top. Second, there is high and persistent concentration at the top. Along with the downward growth, the top-1% authors accounted for 38.4% of top-5 publication credit in 1990 and for 78.3% in 2025. Third, the persistence is widespread within cohorts, within subfields, and within gender. Fourth, new journals only slightly dilute concentration. Fifth, elite authors diversify on topics faster than the rest of the profession. We interpret the findings with a screening model of attention under information overload. The evidence is consistent with the model: as the field grows, citations concentrate on established work and the conditional citation premium of top-author papers narrows.

JEL Codes: A11, A14, I23, J24, O33

Keywords: economics of science, publishing, concentration, entry, coauthorship, scientometrics

---

*I thank Simon Angus, Christian Düben, Maxime Gravoueille, Rafe Meager and participants in seminars at Monash University for suggestions. All errors are my own.

†Monash University, Department of Economics. Email: ricardo.dahis@monash.edu.

1 Introduction

It is widely observed that publishing in economics is highly concentrated. While in principle concentration per se does not imply inefficiency in the production of knowledge,[1] the field as a whole shows concerning signs. For instance, output and citations concentrate in a small set of researchers and departments (Conley and Önder, 2014; Petersen and Penner, 2014). Authorship clusters in a few wealthy-country institutions and postal codes (Aigner, Greenspon, and Rodrik, 2025). Editorial boards interlock and overrepresent a handful of elite addresses (Bantman-Masum et al., 2024; Angus et al., 2021). Economics is more hierarchical and more inward-citing than its social-science neighbors (Fourcade, Ollion, and Algan, 2015). And these advantages compound: a publication in the five leading journals, held by few, carries outsized weight in tenure and promotion (Hamermesh, 2018; Heckman and Moktan, 2020). However, despite the accumulated evidence, it is less clear how concentration is evolving over time, what forces drive the trends, and what we can learn from these patterns.

This paper closes these gaps by documenting recent trends in publishing in economics. We trace how concentration in economics publishing evolved over thirty-five years and introduce attention under information overload as an interpretation. We source publication records for 596 economics journals over 1990–2025 to build a researcher- and institution-level panel. Our headline measure of concentration is the top-share, computed as the share of publication credit held by the top 1% of economists ranked within each year over the full active population (Piketty and Saez, 2003).[2] We document five sets of key facts.

First, economics grew *downward*: the field expands by adding capacity and entrants at the bottom of the journal hierarchy while the top is held by incumbents. The number of economists

---

[1] For instance, knowledge is close to a public good: one definitive paper can settle a question, and the field gains little from many researchers reproducing the same result (Romer, 1990). Scientific production rewards complementarities and is convex at the top, so small differences in talent can efficiently generate large differences in output (Rosen, 1981; Jones, 2009). A hierarchy can function as a tournament that sorts effort and screens quality (Lazear and Rosen, 1981).

[2] We measure inequality in *predistribution*, the distribution of publication credit before any corrective policy, in the sense of Bozio et al. (2024). Unlike income or wealth, publication credit and citations cannot be redistributed after the fact, so the predistribution is the entire object of interest.

publishing at least one article rose from 10,751 in 1990 to 60,521 in 2025, a 5.6-fold increase, and the institutions hosting them grew 5.2-fold, while the top-5 journals expanded capacity only 1.3-fold. Recent entrants concentrate at the bottom of the hierarchy: an over-representation index, entrants' share of a tier's credit divided by their share of all credit, fell in the Top-5 from 0.84 in 1995 to 0.56 in 2025, so entrants now carry roughly half their proportional weight at the top while remaining near parity at the bottom. New entrants overwhelmingly join the bottom: 95.9% of 2025 entrants published outside the upper tiers, up from 82.9% in 1990.

Second, accordingly, the top of economics has concentrated sharply. The top-1% share of publication credit in the five leading journals rose from 38.4% in 1990 to 78.3% in 2025, on a trailing five-year stock with fractional author credit. In other words, by 2025 one economist in a hundred holds nearly four-fifths of the credit in the Top-5. While the mean number of authors on a top-5 publication rose from 1.57 to 3.08 in the same period, the rise in concentration is not explained by team-size growth. Similar patterns broadly hold within a larger set of journals including second-tier general-interest and top-field journals.

Despite some movement along the distribution, concentration is persistent both at the author- and more so at the institution-level. Relative to the distribution in 2000, incumbent top authors gain about 4 percentage points of credit share by 2020. Top institutions lose about 15 percentage points of credit share to the middle-half of the distribution of institutions. Nevertheless, concentration is persistent. The correlation between an author's rank today and their rank ten years later is stable at about 0.3-0.4 between 1995 and 2015. The analogous correlation at the institution level is stable at about 0.7. A similar pattern holds when decomposing the distribution into quintiles: a researcher in the top incumbent quintile stays there with probability 0.48, and an institution stays with probability 0.84.

The evidence suggests that institutions are aggregators of productive individuals rather than having a direct effect on output. We run a two-way fixed-effects decomposition of output and find that author effects account for 38% of the variance while institution effects account for only 1%, with no assortative matching after correcting for limited-mobility bias. Moreover, we find that elite authors increasingly coauthor with other elite authors *across* institutions.

The remaining three facts concern persistence and the field's horizontal change. Third, the concentration is persistent within entry cohorts, within subfields, and within gender: the same names hold the top across each cut. Fourth, new general-interest journals dilute the top-1% share only slightly, absorbing entrants at the bottom rather than redistributing credit at the top. Fifth, elite authors broaden their topics faster than the rest of the profession, so the top diversifies on subject matter even as it acquires a higher share of total credit in the top journals. To the extent that topic diversity correlates with research quality, the evidence suggests that average quality at the top is increasing.

We interpret these facts with a model of attention under information overload. As the active population of economists grows, the attention available to screen papers does not. Each editor, referee, and reader now faces more candidates with the same fixed cognitive budget, so the read of any one paper becomes coarser. A Bayesian evaluator facing a noisier signal leans less on the paper and more on what is already known about its author; because reputation favors incumbents, screening tilts toward them as load rises, even with no change in preferences or the distribution of talent. The model nests the standard explanations such as real quality, proximity, and cumulative advantage, and isolates overload as a new force that scales with the number of submissions. It predicts that a fixed credential gap converts into rising concentration as the field grows.

The model yields six testable predictions that we take to the data. As the field grows, references concentrate on a few established names: the share of citations going to the top-1% of cited authors roughly doubles, from 5.8% to 11.3%, and citations tilt toward older work. Conditional on the same journal and year, the citation premium of top-author papers narrows from about 28% in the 1990s to 7% in the 2010s, consistent with the marginal top paper being accepted on reputation rather than measured quality. Where relevant, we note that the evidence is also consistent with other mechanisms, such as cumulative advantage, selection, and changing production technology.

Our contribution is threefold. First, we build the first researcher-level panel of economics publishing concentration spanning 1990–2025 on the full 596-journal universe, where prior work measured a cross-section. Second, we decompose the 38.4%-to-78.3% rise and show it is growing downward, driven by entry into the bottom rather than a narrowing of the top, with the within-incumbent series rising only to about 0.45. Third, we relate the patterns to the overload model, which yields six predictions, which the data largely support. The model amplifies a fixed credential gap as the field grows: rising load widens the acceptance gap that established authors hold over entrants, which concentrates the top, while the average quality of published work moves ambiguously.

This paper relates to the literature on the economics of science. One strand documents the structure of top-journal publishing: rising selectivity, length, and team size (Card and DellaVigna, 2013), the career weight of the Top-5 (Hamermesh, 2018; Heckman and Moktan, 2020), and the cross-sectional concentration of output across researchers and departments (Conley and Önder, 2014; Petersen and Penner, 2014). A second strand measures inequality trends in science at large: rising citation and impact concentration across disciplines (Nielsen and Andersen, 2021; Sinatra et al., 2023) and the cumulative-advantage processes that produce it (Merton, 1968; Azoulay, Stuart, and Wang, 2014; Bol, Vaan, and Rijt, 2018). A third documents the rise of teams and its consequences for credit (Wuchty, Jones, and Uzzi, 2007; Jones, Wuchty, and Uzzi, 2008). We add to all three a long-run, researcher-level time series for economics and a mechanism that scales with entry.

The closest paper to ours is Conley and Önder (2014), who measure the cross-sectional concentration of economics output across departments and individuals. We add the time series, 1990–2025, and apply a top-share method on a researcher-level panel rather than a single cross-section. We differ from Aigner, Greenspon, and Rodrik (2025), who study citation inequality across countries, by using a within-field denominator and assigning credit at the level of the individual researcher. And where Jones, Wuchty, and Uzzi (2008) document rising cross-institutional stratification, we find that institutional concentration falls as author concentration rises, so the two margins move in opposite directions. The broader tension throughout is that concentration at the top can rise even as the field opens at the bottom: economics admits many more people, yet credit at the top concentrates, and the force we identify is the scarcity of attention and space in outlets.

The remainder of the paper proceeds as follows. Section 2 describes the data and the concentration measure. Section 3.1 establishes how the field grew over time. Section 3.2 documents that concentration is persistent. Section 3.4 estimates a two-way variance decomposition and finds that institutions explain only a tiny fraction of it. Section 3.5 traces out the team collaboration and elite assortativity channels. Section 3.3 shows that persistent concentration happens within subfields, within gender, and within cohorts. Section 3.6 documents the role of new journals. Section 3.7 shows how the field diversified across topics. Sections 4 and 5 develop and test the model of attention under information overload. Section 6 concludes. The Appendix contains additional results and robustness checks.

2 Data and Measurement

Our main data source is OpenAlex, which provides comprehensive coverage of economics journals and their articles, authors, and institutions from 1990 to 2025. We define the journal universe using the 2025 ABDC Journal Quality List for fields of research 3801, 3802, and 3803, which encompass economics. This approach avoids the need for backfilling year-by-year membership decisions over the thirty-five-year period. We categorize these journals into four tiers: the *Top-5* (AER, AER:Insights, Econometrica, JPE, QJE, REStud), *second-tier general-interest* journals, *top-field* journals leading specific subfields, and *all others*. We report the full list of journals with their entry year and paper counts in Table A.1. The first three tiers together comprise 29 journals, which we refer to as the *upper tiers*.

We construct an annual panel of every economist with at least one publication in the universe, linking each article to its authors and to their institutional affiliations. We restrict observations to journal articles, dropping editorials, comments, and reprints.

Our main measure of concentration is the top share (Piketty and Saez, 2003). Let R_t denote the reference population, all economists with at least one publication in the universe in year t. Each economist i earns publication credit y_{it} in year t. The share held by the top $p\%$ is the credit accruing to economists ranked in the top $p\%$ by y_{it}, divided by total credit:

$$S_p^t = \frac{\sum_{i \in \text{top-}p\%} y_{it}}{\sum_{i \in R_t} y_{it}}, \tag{1}$$

where the ranking is computed within year t. The denominator covers the full active population. The statistic therefore registers both how much the top earns and how many enter at the bottom. This is the field-structure object we want: it asks what fraction of credit the leading researchers hold relative to everyone active, so a rise can come either from the top earning more or from the field growing beneath it, and the decomposition in Section 3.1 separates the two.

To reduce single-year sampling noise our preferred measure is the top share over a trailing five-year stock. At year t the reference population R_t is the set of economists active in the window $[t-4, t]$, credit accumulates over the five years,

$$\tilde{y}_{it} = \sum_{s=t-4}^{t} y_{is}, \qquad \bar{S}_p^t = \frac{\sum_{i \in \text{top-}p\%} \tilde{y}_{it}}{\sum_{i \in R_t} \tilde{y}_{it}}, \tag{2}$$

and the top $p\%$ is ranked on the windowed stock $\tilde{y}_{it}$. We report $\bar{S}_1^t$, the trailing-five-year top-1% share, as the headline; Figure H.2 shows the trend is present at every window length.

Credit can be assigned in two ways. Under fractional counting an article with k authors gives each author $1/k$ of a publication; under whole counting each author receives one. Our preferred outcome is the fractional counting, as it tracks the rising prevalence of coauthorship without double-counting. We report results with whole counting in the Appendix. We compute standard errors by bootstrap.

The dataset has some limitations. The main concern is author disambiguation, since OpenAlex identifiers can fragment a single economist across diacritic and name-order variants. We stress-test the identifiers under worst-case consolidation. The headline top-1% share of 78.3% falls to a bounded range under the most aggressive merging, and we report the full procedure in Appendix C. Institutional affiliation is missing on roughly 13–26% of fractional credit, with worse coverage in the early years. Region coverage is incomplete, reaching about 84% of credit, which we address in Appendix D.[3]

Table 1 reports the scope of the panel at five-year marks. The active population rose from 10,751 economists in 1990 to 60,521 in 2025, and active institutions from 1,911 to 9,958, while annual articles grew only from 9,630 to 29,019. Coauthorship rose in parallel: the mean number of authors per Top-5 article climbed from 1.57 to 3.08. The Top-5 share of all articles fell from 3.4% to 1.5% as the all-others tier expanded from 86.0% to 92.2%.

3 Main Results

3.1 A Field Growing Downward

The field as a whole grew far faster than its top tier of journals. Figure 1 shows the two series side by side. The active population rose 5.6-fold between 1990 and 2025 and the count of active institutions rose 5.2-fold, while article capacity in the set of leading journals grew only 1.3-fold. Capacity widens as one moves down the hierarchy: second-tier general-interest journals expanded 1.6-fold and top-field journals 1.9-fold, still far short of the population growth.

Another way to document this pattern is by plotting entry across different tiers of journals. Measuring this is not as simple as decomposing each tier's output by entry cohort, because the entering population grows so fast that recent cohorts mechanically dominate every tier in proportion to their size, masking where entrants actually land. To avoid this issue we build an entrant *over-representation index*, defined as the share of a tier's output produced by economists within five years of their first publication, divided by those entrants' share of all output that year. A value of one means entrants publish in a tier in proportion to their overall footprint; a value below one means they are under-represented. We plot the results in Figure 2. In 2025 the index forms a monotone ladder — 0.56 for Top-5, 0.64 for the second tier, 0.76 for top-field journals, and 1.03 for all other

---

[3] Coverage rises with credit, so the unassigned skew toward the bottom and the top-1% statistics are well covered. Among authors with positive Top-5 credit, the top 1% have an assigned gender for 91.2% of cases, a country for 100%, and an institution for 99.1%, against 81.5%, 94.9%, and 95.0% for the rest.

journals. Entrants are at parity at the bottom and carry roughly half their expected weight at the top. The Top-5 index has fallen from 0.84 in 1995 to 0.56 in 2025, while the all-other index has held at about one. In other words, the top has become relatively less accessible to entrants over time as the field has grown.

Relatedly, the downward growth happened through intensifying competition for existing journals, not with new outlets. Table 2 and Figure 3 describe the entry margin. Total annual entry rose from 10,751 in 1990 (mechanically equal to the active population that year, as the panel is left-censored at 1990) to 28,864 in 2025, but entry into the upper tiers fell from 1,839 to 1,194 while entry into all-others journals rose from 8,912 to 27,670. The all-others share of entry climbed from 82.9% to 95.9%. That is, the field did not absorb its newcomers by opening enough new journals: the number of journals rose modestly from 238 to 465, while the average number of active authors per journal grew from 45 to 130. We separately study the geography of entry in Appendix D.

Figure 4 separates the actual top-1% share from a comparison that holds the incumbent set fixed. For researchers, the actual share rises from 38.4% to 78.3%, while the incumbent-only series rises only mildly, to about 45%. For institutions, the actual share rises from 46% to 80% while the incumbent-only series rises to about 55%. The gap between the two series reflects the contribution of entry to the denominator.

Within the incumbent core, the distribution of output de-polarizes, but authors and institutions move in different directions. Figure 5 arranges incumbents by percentiles of performance in 2000 and tracks where credit moves over time. For researchers, the apparent across-the-board decline in the top share is a composition effect: the top gains about two percentage points twenty years later. For institutions, the pattern is genuine spreading-down: the top bin loses roughly fifteen percentage points to the upper-middle even with the population held fixed.

3.2 Persistence

Not only is concentration high in economics, it is also persistent. Institutional positions are far stickier than individual ones. Figure 6 plots the rank–rank persistence of authors and institutions across the field. The Spearman rank–rank correlation for authors is essentially flat, 0.42 in the early window and 0.44 in the late one. For institutions it is far higher and rising, 0.70 to 0.73.

The top quintile rarely cedes its place. Figure 7 reports quintile transitions from 2015 to 2025 within the incumbent set, with the population held fixed. A researcher in the top incumbent quintile stays there with probability 0.48, and an institution stays with probability 0.84. Movement is local: the off-diagonal mass concentrates in the neighboring quintile, and the bottom-left and top-right corners are nearly empty, so units seldom cross the distribution within a decade. Cutting quintiles within incumbents nets out entry, so the stickiness is not manufactured by entrants flooding the lower ranks.

3.3 Decompositions

We decompose the rise in concentration across three horizontal margins: subfields, cohorts, and gender. We investigate whether the rise in concentration is a story of redistribution across these margins, or whether it is a story of rising concentration within them.

First, we find that the trend is within subfields. A Theil decomposition of publication-credit inequality on a 2000 base puts the within-subfield component at 95.7% of the total in 2025, as the total Theil rose from 0.108 to 0.187, with the between-subfield share near 3 to 4% throughout (Appendix B).

Second, the trend is within cohorts, which is the entry-time face of growing downward rather than the established names tightening their grip. A symmetric Das Gupta decomposition splits the 35.7-point rise in the top-1% Top-5 share into a within-cohort term of $+58.2$ points and a between-cohort term of -22.6 points (Appendix B).[4] The within term measures concentration building inside fixed five-year entry cohorts; the negative between term means the arrival of large young cohorts, on its own, dilutes the top. Each cohort enters internally top-heavy, a thin elite above a long tail, and the field stacks one bottom-heavy cohort on the next. Restricting the decomposition to authors

---

[4] The headline series rises 39.9 points (38.4 to 78.3). This decomposition re-ranks the top 1% within the population assignable to a five-year entry cohort, which excludes left-censored pre-1990 authors of unknown entry year and shifts the endpoints to a 35.7-point rise. The within- and between-cohort shares are unaffected by this level difference.

active by 1999 confirms the reading: the rise falls from 35.7 to 5.9 points and the incumbent top share reaches only 44.3%, the ceiling near 45% that the incumbent-only series hits in Figure 4. What rise survives among incumbents is itself within-cohort, $+24.1$ against -18.2 points, so the within-cohort pattern holds whether or not the cohort set is fixed; it is the succession of new cohorts, not the old guard, that supplies the magnitude. Recent cohorts also reach the top faster, with median years-to-top falling from five or six for cohorts entering through 2000 to two for the 2016–2020 cohort. The top nonetheless stays concentrated as it renews: of authors first active in 1991–1999, 62.7% of the surviving 2025 incumbent top-1% come from that cohort, with five-year persistence of 8.6% (Table 3).

Finally, the trend is also within gender, even as female representation in the profession roughly doubled. The female share of credit rose from 9.2% in 1990 to 21.6% in 2025 among authors with assigned gender, yet of the 25.5-point rise in the top-1% share over the gendered population, $+26.0$ percentage points are within-gender and only -0.6 between-gender. The changing gender mix does not contribute to the general trend (Appendix B).[5]

3.4 Individuals Produce, Institutions Aggregate

We study to what extent research output is explained by researcher- or institution-level characteristics. We separate the two with a two-way fixed-effects decomposition in the tradition of Abowd, Kramarz, and Margolis (1999). We model log quality-weighted annual output as the sum of an author effect, an institution effect, and a year effect, estimated on the 458,551 author-year observations of the largest mover-connected set (112,412 authors, 8,079 institutions; Appendix F). The author

---

[5] The headline series rises 39.9 points (38.4 to 78.3); restricting to authors with an assigned gender and re-ranking the top 1% within them gives the 25.5-point rise decomposed here, with the within- and between-gender shares unaffected by the level difference. Gender is missing for 34.0% of authors. Our algorithm assigns gender based on first name and country of affiliation, so it is more likely to assign gender with high confidence to Western names.

effect θ_i measures portable individual productivity;[6] the institution effect ψ_j measures the institution's output premium, net of individuals' performance. Separating the two requires researchers who move, of whom there are 35,268. Because raw variance components are biased when many institutions are linked by few movers, we report the leave-out correction of Kline, Saggio, and Sølvsten (2020) alongside the plug-in estimates.

We find that output is overwhelmingly a property of the researcher. Table F.2 reports the decomposition. After the leave-out correction, the author component accounts for 38.2% of the variance in log output, the institution component for 1.3%, and the sorting covariance for -0.3%. The author effect is roughly thirty times the institution effect. The plug-in estimates overstate the institution share at 3.7% and the sorting term at -4.1%. The bias correction collapses both, as expected given the relatively low number of moves per institution on average (Andrews et al., 2008). The corrected correlation between author and institution effects is -0.04, statistically and economically indistinguishable from zero.

We read these results as follows. A top institution leads because it holds productive economists, not because it makes the economists it holds more productive. This reconciles the two facts of Figure 6 and Figure 5: the institutional ranking is sticky because the underlying talent is sticky and recurrently re-hired by the same names, yet institutional concentration of output can still disperse as portable talent spreads across a growing set of hosts. The pattern stands in contrast to the canonical labor-market application, where firm effects explain on the order of a fifth of wage variance and sorting is positive (Card, Heining, and Kline, 2013). In economics publishing the institution explains only a small share of the variance.

---

[6] We test reputation portability (Definition 3) directly. If reputation is owned by the researcher, a mover's Top-5 output should be flat across an institutional move and should not depend on the destination's prestige. Across 7,582 clean single movers, Top-5 fractional credit is flat in event time around the move, with only a small transient bump in the move year that decays within three years, and the path does not rise with destination prestige: authors who move up and down the institutional ranking trace the same profile (Figure H.5). A caveat to this analysis is that we observe realized output, not acceptance, since submissions are unobserved; and the decomposition speaks to output flow conditional on having having papers accepted.

3.5 Team Production and Elite Assortativity

Part of the explanation for the persistent concentration documented in previous Sections is that top authors are also changing their paper production technology. In this Section we document whether persistent concentration operates through team production channels. We start by tracing the mean number of authors per paper, output per author, and coauthor set size by tier. We find that elite authors do not run larger teams or publish more per year, but they do draw on a larger rotating set of coauthors. Figure 8 reports three series by tier. Team size rose across the board, from 1.59 to 2.73 authors per paper field-wide and from 1.68 to 2.61 for the top-1%, so the top writes on teams no larger than the field (panel A). Output per author fell at the top: the top-1% published 2.49 economics papers per year in 1990 and 1.49 in 2025, after a peak of 2.83 in 2006, against a field mean near 1.3 throughout (panel B). What rose is the coauthor set. The mean number of distinct coauthors over a rolling five-year window climbed from 1.49 to 6.27 for the top-1% and from 1.66 to 5.37 field-wide, with the top tiers running above the rest of the field (panel C). The team-size growth does not drive the concentration: a whole-count decomposition that holds the team-size distribution at its 1990 level (Appendix E) attributes the 30.7-point whole-count rise almost entirely to own-concentration, with the team-size matching terms staying within ± 0.1 percentage point of zero.

Elite authors increasingly partner with other elites, and they do so across institutions. Figure 9 plots the share of Top-5 papers that are elite and cross-institutional, with at least two top-1% authors at different institutions. That share rose from 1.8% in 1990 to 38.6% in 2025. Raw elite homophily rose too: the share of a top-1% author's Top-5 coauthors who are themselves top-1% went from 0.289 to 0.694. Much of this raw rise is mechanical, since a growing elite pool makes elite-elite pairing more likely by chance. A degree-preserving rewiring counterfactual that randomizes who coauthors with whom while holding each author's paper count fixed reproduces about 71.6% of the rise (Figure H.4); the residual 28.4% is excess sorting of the top toward the top, beyond what paper counts alone imply. Elite collaboration intensifies across universities rather than within them.

This reconciles rising author concentration with dispersing institutional concentration, and it is the cross-institutional elite pairing the model predicts in Extension 2, read against the excess sorting rather than the raw series.

3.6 New Journals

We investigate the role of new journals in the trends in concentration of academic credit. Several top journals opened mid-period, and if they pulled credit away from the established elite, the measured rise could overstate how much the top tightened. We test this by adding each newly launched general-interest or top outlet to the universe in event time and tracking the top-1% share around the launch. The data are the 1990–2025 author-by-paper-by-tier panel that underlies the rest of the paper, across four launch cohorts: the *Journal of the European Economic Association* in 2003, the four *American Economic Journals* in 2009, *AER: Insights* in 2019, and the two *Journal of Political Economy* field titles in 2023. Because OpenAlex cover-dates some titles before their official start, we keep only rows from the official launch year onward.

The specification holds the top-1% set fixed and recomputes its credit share under two scenarios. For each cohort and event time $\tau \in [-5, 5]$ we compute the top-1%'s share of credit in the incumbent Top-5 alone, Y_{T5}, in the incumbent Top-5 plus the new outlet, Y_{T5+}, and in the new outlet by itself, Y_{NJ}. The top-1% set is the canonical panel-wide rank, so mechanically membership does not respond to the launch. Before the launch Y_{T5+} equals Y_{T5} by construction. After it, the gap $Y_{T5+} - Y_{T5}$ measures whether folding in the new journal changes the trend in concentration.

Five years out the gap is negative, and its size tracks who publishes in the outlet (Figure 11). Including *AER: Insights* moves the top-1% share by -1.0 percentage points, the *JEEA* by -6.2, and the 2009 *American Economic Journals* by -13.3. The top-1% capture 56% of *AER: Insights* credit, so adding it barely moves the share: the elite reproduce in the new elite outlet. They capture only 7% of *American Economic Journal* credit, so adding those titles dilutes the measure: the broader outlets absorb entrants at the bottom. In the model this is $\eta_J \approx 0$ (Proposition 5): new journals add capacity at the bottom without relieving the load that concentrates the top.

3.7 Topics

Persistent concentration of credit is in principle orthogonal to the precise content of papers along the distribution. The top could be concentrating on the same topics as the rest of the field, or it could be concentrating on a narrower set of topics. We investigate which is the case.

We employ a BERTopic model to measure topic breadth on the abstracts of Top-5 papers. The model sorts them into 195 topics, and we drop the model's residual no-topic category. For each year we form the distribution of papers across topics and summarize its breadth with a Herfindahl index and a normalized Shannon entropy, computed for the field as a whole and separately for papers with at least one top-1% author and for the rest. We also measure how far the two groups' topic mixes differ with a Jensen–Shannon divergence.

We find that the field widened its coverage of topics, and that elite researchers are a driving force behind the trend (Figure 10). The pooled topic Herfindahl fell from 0.0151 in 1990 to 0.0114 in 2025. The decline is the top's doing: the top-1% Herfindahl fell 44%, from 0.0224 to 0.0125, while the rest of the field concentrated, from 0.0179 to 0.0228. In normalized entropy, the gap between the top-1% and the rest flipped from -0.042 in 1990 to $+0.145$ in 2025, so the top began slightly less diverse than the rest and ended notably more diverse. The two groups also study increasingly different things, with the Jensen–Shannon divergence between their topic distributions rising from 0.31 to 0.36. Concentration of credit thus coincides with a broadening, not a narrowing, of the top's subject matter. Elite researchers hold a large share of the credit while ranging over more topics.

4 A Model of Information Overload

Persistent inequality already has many explanations. The literature offers at least five families of forces, all predicting that the top stays on top. First, talent and resources: elite departments hire and fund the strongest, through assortative matching (Becker, 1973), compounding funding (Bol, Vaan, and Rijt, 2018), tournament incentives (Lazear and Rosen, 1981; Heckman and Moktan, 2020), and a rising knowledge burden (Jones, 2009). Second, proximity: incumbents sit closer to editors and seminars (Colussi, 2018; Brogaard, Engelberg, and Parsons, 2014). Third, cumulative advantage: success breeds success, the Matthew effect (Merton, 1968; Azoulay, Stuart, and Wang, 2014), with rational herding on famous types as one micro-foundation (Banerjee and Fudenberg, 2004). Fourth, networks and superstars: the elite coauthor with the elite (Ductor et al., 2014; Jones, Wuchty, and Uzzi, 2008), and small talent gaps scale into large output gaps (Rosen, 1981; Frank and Cook, 1995). Fifth, gatekeeping: incumbents raise barriers to entry (Azoulay, Fons-Rosen, and Graff Zivin, 2019). Almost all operate through real quality or accumulated advantage, and none requires the number of submissions to change.

This paper introduces a distinct force that does scale with the number of submissions, which we call *information overload*. As the active population grew fivefold, the attention available to screen it did not. Each editor, referee, and reader now faces far more candidates with the same fixed cognitive budget, so the read of any one paper becomes coarser. Information overload is this fall in attention per candidate as the field grows. A noisier read of quality leads a Bayesian evaluator to lean less on the paper and more on what is already known about its author. Because reputation favors incumbents, the screen tilts toward them as load rises, with no change in tastes, talent, or merit.

The model nests these known channels and isolates overload as a distinct margin. Real quality maps to the talent distribution F_θ; proximity and cumulative advantage map to the reputation prior $\mu(r)$; the five families above all enter through these objects. Overload enters only through the signal noise $\sigma_\varepsilon^2(L_t)$, which rises with the population. The known channels raise concentration by changing fundamentals or priors. Overload raises it through inference alone, holding talent, proximity, and reputation fixed, and converts a fixed credential gap into growing concentration as N_t rises. If reputations were equal across groups the load would only add noise, not concentrate it, a restriction we take to the data. This account rationalizes the descriptive facts of Section 3 as one equilibrium response to a single shock: the top share rises as the screen weights incumbents more, institutional concentration falls while author concentration rises because reputation is carried by the person not the address, and topics broaden because reputation clears the screen. We show that one friction generates six testable predictions, taken to the data in Section 5.

4.1 Setup

Time is indexed by t. A mass $N_t > 0$ of researchers is active in period t; this is the same active population whose top share Equation (1) measures. Each researcher i has talent $\theta_i \sim F_\theta$, fixed across t, and writes $\bar{n}$ papers of true quality $q_p = \theta_i + \nu_p$ with ν_p idiosyncratic. The field therefore produces $P_t \equiv \bar{n} N_t$ papers in period t, the publication flow that gatekeepers screen and citers draw on; it grows with the active population N_t. The talent distribution F_θ is fixed *in shape*; the geographic composition of who is drawn may shift, as the entry facts of Section 3.1 and Appendix D show, without changing that shape. Three types of agent act on these papers: authors, who choose effort and coauthors; gatekeepers, the editors and referees who screen submissions; and citers, who choose what to read and cite.

Reputation. Let $r_{i,t}$ denote the reputation of researcher i at t, measured by cumulative publication credit. Reputation is a sufficient statistic for the public prior on talent: $\mathbb{E}[\theta_i \mid r_{i,t}] = \mu(r_{i,t})$ with μ strictly increasing, and the prior precision at the average reputation is $\tau_q > 0$. We write r for reputation throughout, to keep it distinct from the reference population R_t of Section 2 and from the top-share statistic S^t_α.

The signal technology. Each submitted paper p generates a public signal observed by gatekeepers,

$$s_p = q_p + \varepsilon_p, \qquad \varepsilon_p \sim \mathcal{N}\big(0,\, \sigma^2_\varepsilon(L_t)\big), \tag{3}$$

whose noise rises with attention load. Load is candidates per unit of gatekeeper attention, $L_t \equiv N_t/A$, where $A > 0$ is the attention budget per gatekeeper, fixed across t. We adopt the one-parameter overload law

$$\sigma^2_\varepsilon(L_t) = \sigma^2_0\, L_t, \qquad \tau_\varepsilon(L_t) \equiv \frac{1}{\sigma^2_\varepsilon(L_t)} = \frac{1}{\sigma^2_0\, L_t}, \tag{4}$$

the reduced form of rational inattention: a gatekeeper spreads a fixed budget A across N_t candidates, so per-candidate signal precision falls in N_t (Sims, 2003; Mackowiak and Wiederholt, 2009; Mackowiak and Wiederholt, 2015; Caplin and Dean, 2015). We assume a linear form but that does not affect the sign of the comparative statics results.

Gatekeeping. A gatekeeper observing $(s_p, r_{i,t})$ forms the posterior mean

$$\mathbb{E}[q_p \mid s_p, r_{i,t}] = w_t\, s_p + (1 - w_t)\, \mu(r_{i,t}), \qquad w_t \equiv \frac{\tau_\varepsilon(L_t)}{\tau_q + \tau_\varepsilon(L_t)}, \tag{5}$$

and accepts the paper if and only if this posterior clears a threshold $\bar{q}_t$, where $\bar{q}_t$ adjusts to fill the fixed capacity κ of the tier. The weight w_t is the reliance on the signal; $1 - w_t$ is the reliance on reputation. The acceptance rule is equivalent to a reputation-conditional signal cutoff $s_p \geq c_t(r) \equiv [\bar{q}_t - (1 - w_t)\mu(r)]/w_t$, so a more reputable author clears the screen at a lower signal.

4.2 Definitions

Definition 1 (Bayesian signal weight)**.** *The equilibrium weight a gatekeeper places on the signal relative to reputation is*

$$w_t \;\equiv\; \frac{\tau_\varepsilon(L_t)}{\tau_q + \tau_\varepsilon(L_t)} \;\in\; (0, 1).$$

$w_t = 1$ *is full trust in the signal;* $w_t = 0$ *is decision on reputation alone.*

Definition 2 (Top-α share)**.** *For a population fractile* $\alpha \in (0, 1)$*, the top-*α *share* S_α^t *is the share of top-tier publication credit accruing to the top* α *fraction of active researchers ranked by credit, defined exactly as in Equation (1). The headline object is the top-1% share* S_1^t*.*

Definition 3 (Reputation portability)**.** *Reputation* $r_{i,t}$ *is* individual-portable *if it is owned by the researcher and accumulates with* i *across institutional moves. Equivalently, the conditional distribution of talent given reputation,* $F_{\theta|r}$*, does not depend on the host institution of* i*.*

4.3 Assumptions

Assumption 1 (Gaussian conjugacy)**.** $q_p \mid r \sim \mathcal{N}(\mu(r), \tau_q^{-1})$ *and* $\varepsilon_p \sim \mathcal{N}(0, \sigma_\varepsilon^2(L_t))$*, independent.*

Assumption 2 (Attention rationing)**.** *Per-candidate signal precision satisfies Equation (4) with* $\sigma_0 > 0$*: gatekeeper attention is rival across candidates and fixed in aggregate at* A*.*

Assumption 3 (Reputation informativeness)**.** μ *is strictly increasing in* r *and* $\mathrm{cov}(\theta, r) > 0$.

Assumption 4 (Capacity rationing)**.** *Each tier has fixed capacity (*κ *at the Top-5). The threshold* $\bar{q}_t$ *clears the slot constraint.*

Assumption 5 (Stable talent shape)**.** F_θ *is fixed across* t *in shape, though the composition of who is drawn may change.*

Assumption 6 (Log-concavity and MLRP)**.** F_θ *and the induced reputation distribution* F_r *are log-concave. We further impose, as a high-level condition, that the admitted-author map shifts in the monotone-likelihood-ratio order as* w_t *falls.*

Assumptions 1 and 2 deliver Lemma 1. Assumptions 1 to 6 deliver Propositions 2, 4 and 5 and Extension 1 and Extension 2. Continuity of μ and finite moments of F_θ are imposed throughout.

4.4 Results

The mechanism rests on one elementary comparative static, which we state as a lemma rather than feature, because it follows directly from differentiating Definition 1. We state each result with its conditions; all proofs are collected in Appendix G.

Lemma 1 (Overload lowers the signal weight)**.** *Under Assumptions 1 and 2,*

$$\frac{\partial w_t}{\partial N_t} < 0, \qquad \frac{\partial w_t}{\partial A} > 0.$$

As the field grows, the screen trusts the signal less and reputation more. The substantive content is what this does to the acceptance bar, to concentration, sorting, and citing.

Proposition 1 (Overload widens reputation reliance and the acceptance gap)**.** *For any two reputation levels $r_H > r_L$, the gap in acceptance cutoffs is*

$$\Delta c_t(r_H, r_L) \;\equiv\; c_t(r_L) - c_t(r_H) \;=\; \frac{1-w_t}{w_t}\left[\mu(r_H) - \mu(r_L)\right] \;>\; 0.$$

Under Assumptions 1 to 3, the reputation reliance $1 - w_t$ and the cutoff gap both rise with the population, $\partial(1-w_t)/\partial N_t > 0$ and $\partial \Delta c_t/\partial N_t > 0$, and the acceptance-probability gap $a_t(r_H) - a_t(r_L)$ widens with them.

A more reputable author always clears the screen at a lower signal, and overload widens that advantage. As load rises the editor leans on reputation, so an entering author with a thinner record must clear a higher signal bar than an established author for the same posterior. A higher bar means a lower acceptance rate, $\pi_2 < \pi_1$, so entrants are admitted less often than incumbents and the wedge grows with the field. This is the screen-level reason entrants almost never reach the top in Section 3.2. This is the amplifier: the credential gap $\mu(r_H) - \mu(r_L)$ is held fixed, yet the acceptance gap it produces widens purely because w_t falls. If reputations were equal, $\mu(r_H) = \mu(r_L)$, the gap would stay at zero at every load, the testable restriction that overload concentrates only through a pre-existing credential gap. Proposition 2 is the distributional image of this widening gap.

Proposition 2 (Concentration rises with overload)**.** *Under Assumptions 1 to 6, the top-α share is strictly decreasing in the signal weight,*

$$\frac{d\,S^t_\alpha}{d\,w_t} \;<\; 0, \qquad \textit{equivalently} \qquad \frac{d\,S^t_\alpha}{d\,N_t} \;>\; 0 \textit{ at fixed } A \textit{ and } F_\theta.$$

Conditioning on publication, the same fall in w_t *lowers* the true quality of the marginal admitted top-1% paper, because the screen now admits it on reputation rather than signal. The conditional citation residual of admitted top-1% papers therefore narrows. This is the gate margin, and it is what prediction P4 identifies. It moves opposite to the citer margin of Extension 1, a point we return to below.

The average quality of all published work, by contrast, moves ambiguously.

Proposition 3 (Average admitted quality moves ambiguously)**.** *Let* $\bar{Q}_t \equiv \mathbb{E}[q_p \mid \text{accepted at } t]$ *be the average true quality of admitted papers, and* $\pi_t \equiv \kappa/P_t$ *the acceptance rate at fixed capacity* κ*. Under Assumptions 1 to 5, the response of* $\bar{Q}_t$ *to the population is the sum of two terms of opposite sign,*

$$\frac{d\bar{Q}_t}{dN_t} = \underbrace{\frac{\partial \bar{Q}_t}{\partial w_t}\frac{\partial w_t}{\partial N_t}}_{\text{reliance penalty} < 0} + \underbrace{\frac{\partial \bar{Q}_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial N_t}}_{\text{truncation premium} > 0},$$

so the model does not sign it.

Two forces pull average published quality in opposite directions. A lower w_t admits more papers on reputation, including reputable papers with weak signals, which lowers the average realized quality of what is published, the reliance penalty. Against it, fixed capacity against a growing population lowers the acceptance rate π_t and raises the bar a paper must clear, which lifts the average, the truncation premium. Because the net is unsigned, average published quality is not a clean test of overload. This is why Proposition 2 and prediction P4 read the screen through the *conditional* residual of the admitted top-1%, holding the journal-by-year cell fixed, rather than through the quality of all output.

Proposition 4 (Institutional concentration falls under portable reputation)**.** *Suppose reputation is individual-portable (Definition 3) and inter-institutional moves are costless. Under Assumptions 1 to 5, institutional concentration of top-tier credit is non-increasing in the reputation weight* $1 - w_t$*.*

This is the divergence of Figure 5: authors concentrate while their institutions disperse. Place matters less precisely because the person carries the reputation that clears the screen. The portability premise of Definition 3 is not assumed here but measured: the two-way fixed-effects decomposition of Section 3.4 finds institution effects explaining 1% of output variance against 38% for author effects, with a near-zero author-institution correlation. Reputation and productivity travel with the researcher, which is exactly the condition under which Proposition 4 predicts institutional dispersion.

Proposition 5 (Opening journals: a signed condition)**.** *Let the extended-tier journal count* J_t *rise at fixed* A *and* κ*, and define the load elasticity of journals* $\eta_J \equiv \partial \log L_t / \partial \log J_t$*. Under Assumptions 1 to 6,*

$$\operatorname{sign}\left(\frac{\partial S^t_\alpha}{\partial J_t}\right) = \operatorname{sign}(\eta_J).$$

If $\eta_J > 0$*, new journals raise per-gatekeeper load and worsen concentration; if* $\eta_J < 0$*, they screen pre-publication and alleviate it.*

The sign of η_J is in principle testable. Descriptively, the launch evidence of Figure 11 is consistent with η_J near zero: new general-interest outlets barely move the top-1% share, absorbing entrants at the bottom rather than redistributing load at the top. With only four launch cohorts we read this as descriptive rather than a formal test of the sign.

Extension 1 (Defensive citing is Bayes-optimal at the citer margin)**.** *Suppose a citer with attention budget* A_c *selects a reading list of size* K *from a candidate pool of mass* P_t *to inform a quality judgment, and cites at most* $M \leq K$ *of them, the number of references the paper has room for. Under Assumptions 1 to 3, the Bayes-optimal list over-samples the high-reputation tail: the optimal inclusion weight on candidate* k *is proportional to the prior precision* $\rho(r_k)$ *that the candidate's reputation injects, which is increasing in* r_k*, where* $\rho(\bar{r}) = \tau_q$ *at the average reputation. The cited set is then the* M *candidates of highest* $\rho(r_k)$*, so when the space limit binds against a growing pool, the share of citations going to the high-reputation tail rises.*

Neither limit grows with the literature: as the pool P_t expands, a fixed reading budget and a fixed paper's length both force the citer to select, and reputation is the margin of selection.

Extension 2 (Elite cross-institutional pairing)**.** *Suppose each author may choose one coauthor to maximize expected acceptance, and the team's screen signal enters through the maximum reputation on the paper. Under Assumptions 1 to 5 and Definition 3, the best reply is to pair with the highest-reputation available author, independent of institution; as* w_t *falls, the return to that maximum rises, so top-1%×top-1% pairing intensifies and is increasingly cross-institutional.*

This is the cross-institutional elite pairing of Figure 9: a falling w_t raises the return to the most reputable name on the paper.

4.5 Counterfactuals

The model signs four interventions. First, opening more journals helps only if it lowers per-gatekeeper load, that is, if $\eta_J < 0$ in Proposition 5; the launch evidence of Figure 11 suggests η_J near zero, so new outlets neither worsen nor alleviate concentration at the top. Second, expanding capacity κ at fixed J_t lowers the threshold $\bar{q}_t$ and mechanically reduces the top share. Third, an attention-technology shock that raises A, such as machine-assisted screening, raises w_t and shrinks both the top share and defensive citing. Fourth, relaxing either citer constraint weakens defensive citing at the citer margin: a reading-budget subsidy that raises attention A_c, or longer reference lists that raise the number of slots K, each lets citers reach past the most established names.

4.6 Testable predictions

The model yields six predictions, which we take to the data in Section 5. They are (P1–P2) defensive citing, the rise in unconditional reference concentration on the famous-old and recent-star at the citer margin; (P3) editor gatekeeping that scales with load; (P4) a narrowing conditional quality residual of the admitted top-1$\%$ at the gate margin; and (P5–P6) cohort dynamics: a static Matthew clock and decaying returns to early Top-5 placement. The six follow from the single shock of a fivefold rise in N_t against fixed space and attention.

5 Evidence for the Model

We take the model's predictions to the data. We test each in turn, mirroring the gate-versus-citer split.

P1–P2: Defensive citing at the citer margin

The model predicts that references concentrate on reputation when load rises. We test this on two cuts. Figure 12 plots the concentration of cited authors in Top-5 reference lists (panel a) and the elite-citation share by vintage of the cited work (panel b). On the first cut, the share of reference-citations going to the top 1% of cited authors rose from 5.8% in 1991 to 11.3% in 2025, and the Gini of the cited-author frequency distribution rose from 0.166 to 0.482. Reference lists became markedly more concentrated on a few names. On the second cut, the elite-citation share tilted toward old work: citations to old (11+ year) work rose from 1.2% in 2002 to 25.2% in 2025, while citations to recent (0–3 year) work fell from 17.4% in 1992 to 8.7% in 2025. Both cuts move as the model predicts.

P3: The editor channel

The model assigns part of the persistent concentration to the gate, where editors screen submissions. Editorial power is itself geographically concentrated: editorial teams at leading journals are less geographically diverse and more North American than the authors they publish (Angus et al., 2021). We test whether editor appointments concentrate publishing toward an editor's home journal and home institution.

We assemble the editor roster from four sources: the replication data of Brogaard, Engelberg, and Parsons (2014), which cover thirty journals through 2011; editor lists curated by the journals themselves, including the AEA portfolio, the *Econometrica* editor and associate-editor histories, and the *Review of Economic Studies* managing-editor history; an open editorial-board collection (GOELD); and archived editorial-board pages retrieved from the Internet Archive's Wayback Machine to fill the remaining gaps. This yields a full dataset of 430 editor-turnover events. We resolve each editor-in-chief to an OpenAlex author identifier by name matching, and read the home institution in the year before appointment from the author-by-institution panel. After resolving editors and imposing data requirements to construct the panel our sample reduces to 104 editor-turnover events. This estimation sample is selected rather than a random draw: relative to the 326 events that drop out, the 104 retained are overwhelmingly Top-5 (90% versus 21%), all have a recorded home institution (100% versus 13%, the binding requirement), and skew toward later appointments (Table H.1). We therefore read the editor estimates as evidence for institution-identified editors at the leading journals, not for the universe of editorial turnovers. We estimate dynamic effects of editor tenure using the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), with not-yet-treated editors as the control group. Figure 13 reports the author-level event study (panel a) and the per-period editor home-bias estimates against the load the editor faced (panel b).

We find that the author-level effect is positive and small. Echoing Brogaard, Engelberg, and Parsons, 2014, the overall effect of an editor's appointment on a matched author's publishing at the editor's home journal is 0.044 (standard error 0.005) for publication counts, 0.022 (0.003) for citation credit, and 0.045 (0.004) for any publication; on the Top-5 subsample the three effects are 0.052, 0.027, and 0.054. The by-period pattern scales up with load. As the active-population-to-Top-5 ratio falls from 22.4 in the earliest window to 11.3 in the latest, the home-institution effect moves from -0.013 toward $+0.004$. While the effects are small, they are not zero, and they scale with load as the model predicts.

P4: The conditional quality residual at the gate margin

The model predicts that the conditional quality residual narrows at the gate, because screening at journals admits the marginal top paper on reputation rather than measured quality. Figure 14 plots the conditional top-1% citation premium by decade, estimated with journal-by-year fixed effects so that the comparison is within outlet and year. The premium fell from $+0.283$ (standard error 0.059) in the 1990s to $+0.071$ (0.101) in the 2010s. Conditional on publication in the same journal and year, papers with a top-1% author once earned a 28% citation premium; by the 2010s that edge had narrowed to about 7%. Binning by decade understates the precision of the trend. Pooling all years and interacting the top-1% indicator with a linear time trend, the premium falls by 0.008 log points per year (95% confidence interval $[-0.015, -0.0002]$, $p = 0.04$, across 26 journal clusters), so the narrowing is distinguishable from zero in the pooled specification even though the single-decade estimates are imprecise. Restricting to the Top-5 alone leaves only five journal clusters and the trend is no longer distinguishable from zero; the narrowing is identified from the full cross-tier sample, not the Top-5 in isolation. The narrowing is the gate-margin prediction, and it moves opposite to the rising unconditional concentration of P1–P2, exactly as the two-margin structure of Section 4 predicts.

P5–P6: Cohort dynamics

The framework expects two things of how growth moves researchers through the field. First, the strength of the link from an early hit to later top status should not rise across cohorts: the path that turns early success into standing is the same one, only more people travel it. Second, an early Top-5 placement should keep sorting authors toward the top even as its citation payoff fades, because attention thins as the field grows. Both follow from the same shock, a roughly fivefold rise in N_t against fixed attention. We take them in turn.

The link from an early hit to top status held steady. We assign each author to a cohort by year of first publication, define an early signal as at least one Top-5 paper at career age 0–2, and ask whether that signal predicts top-1% status at career age 10. A per-cohort logit traces the coefficient on the signal across cohorts entering between 1990 and 2015, the latest for which age ten is observed by 2025. Figure 15 plots it. The coefficient is flat, ranging from 3.1 to 3.8 with overlapping standard errors and no trend. The path from early hit to a place at the top did not get stronger; what changed is the number of researchers walking it. Under one reading of cumulative advantage the link might instead tighten as the field grew; the flat profile is at least consistent with a fixed path that more researchers travel.

The payoff to an early Top-5 placement splits across two margins. Using the same cohorts, we widen the early-signal window to career age 0–5 and follow two outcomes through career age 10: cumulative citation credit, how much an early hit pays in attention, and the log-odds of reaching the top 1%, whether it still sorts an author upward. Figure 16 plots both by cohort. The citation payoff fell from 0.852 for the 1990–95 cohort to 0.603 for the 2011–15 cohort, a one-time drop after the first cohort and flat thereafter. The sorting payoff did not decay; it was flat to rising, from 3.87 to 4.05. An early Top-5 paper still lifts an author toward the top, even as the citations that paper draws erode. Placement keeps sorting; its citation reward does not hold.

Discussion

While the model's predictions are mostly supported by the data, other explanations other than overload also fit some of the patterns. For instance, a rising top share while the field grew downward could reflect cumulative advantage or a widening spread of talent. Concentrating references on established work could equally reflect genuine quality sorting or citation habits; the overload reading adds that the shift is toward reputation conditional on content and grows with load.

The sharpest test of overload against these alternatives would exploit variation in load that is independent of time and of accumulated reputation, ideally an attention shock that moves submissions per gatekeeper without moving the talent pool. The present data do not contain such variation: load rises monotonically with the field's growth, so its effect cannot be cleanly separated from cumulative advantage in the time series alone. We flag the discriminating experiment, such as machine-assisted refereeing, a desk-reject rule, or a referee-pool expansion, as the task left for future work (Section 6).

The narrowing quality edge is the test that most directly separates the accounts: cumulative advantage predicts a stable or widening edge for top-author papers within a journal and year, while the screen predicts it shrinks as the marginal top paper is admitted on reputation. The pooled trend supports the screen, though we read it as suggestive rather than decisive, since the edge is precisely estimated only when all tiers contribute identifying variation, and the top-1% indicator and the citation outcome are constructed from related quantities.

6 Conclusion

This paper shows that the field of economics grew downward. The top-1% credit share rose from 38.4% in 1990 to 78.3% in 2025, but the rise is mostly entry at the bottom of the hierarchy and reallocation among incumbents, not a tightening of access to the top. New entrants fill the journals below the leading tiers, and the share held by incumbents alone climbs only to about 45%. Institutions spread out while authors concentrate, which is partly explained by author movement across institutions; author identity explains 38% of the variance in output and the institution explains 1%. Elite assortativity grows substantially as the field expands. These patterns fit a model of attention under information overload, where screening leans harder on reputation as the number of submissions climbs.

The paper contributes a researcher-level panel of economics-publishing concentration over 1990–2025 on the full 596-journal universe, a decomposition showing the rise is growth at the bottom, and a screening model that organizes the facts into testable predictions. It leaves five questions open. First, beyond the descriptive evidence presented here, one could test the mechanism causally: a shock that moves attention per submission without moving the talent pool, such as machine-assisted refereeing, a desk-reject rule, or an expansion of the referee pool, would identify the channel the model turns on. Second, it is theoretically unclear in what direction the rise of artificial intelligence will take concentration in publishing. On the one hand they may ease the attention bottleneck, lowering the noise per submission and loosening the top. However, it will also lower the cost of producing research, writing and submitting, which will enlarge the pool of researchers. The net effect on concentration is an open empirical question. Third, we measure economics in isolation and do not replicate the series across fields; whether economics concentrates faster than science at large, which documents rising impact inequality across dozens of disciplines (Nielsen and Andersen, 2021; Sinatra et al., 2023), is left open. The panel and the screen carry over to other fields with comparable tier structures and to other rewards reputation may capture, including citations, editorships, grants, and prizes. Fourth, the question of whether the research itself improved or worsened as concentration persisted is open. Topic breadth at the top rose, but breadth is neither good nor bad on its own, and citation counts track attention more than quality. The science-of-science literature offers measures closer to substance, among them how far a paper departs from the work it builds on, its disruptiveness (Wu, Wang, and Evans, 2019; Park, Leahey, and Funk, 2023); tracking such measures by tier over 1990–2025 would show whether a more concentrated top produces more or less original work. Finally, future research could measure whether this reputation-weighted screen misallocates attention towards elite authors' papers and away from other good work.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). “High Wage Workers and High Wage Firms”. In: *Econometrica* 67.2, pp. 251–333.

Aigner, E., J. Greenspon, and D. Rodrik (2025). “The Global Distribution of Authorship in Economics Journals”. In: *World Development* 189, p. 106876. DOI: 10.1016/j.worlddev.2025.106876.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008). “High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?” In: *Journal of the Royal Statistical Society: Series A* 171.3, pp. 673–697.

Angus, S. D., K. Atalay, J. Newton, and D. Ubilava (2021). “Geographic Diversity in Economic Publishing”. In: *Journal of Economic Behavior & Organization* 190, pp. 255–262. DOI: 10.1016/j.jebo.2021.08.005.

Azoulay, P., C. Fons-Rosen, and J. S. Graff Zivin (2019). “Does Science Advance One Funeral at a Time?” In: *American Economic Review* 109.8, pp. 2889–2920. DOI: 10.1257/aer.20161574.

Azoulay, P., T. Stuart, and Y. Wang (2014). “Matthew: Effect or Fable?” In: *Management Science* 60.1, pp. 92–109. DOI: 10.1287/mnsc.2013.1755.

Banerjee, A. and D. Fudenberg (2004). “Word-of-Mouth Learning”. In: *Games and Economic Behavior* 46.1, pp. 1–22. DOI: 10.1016/S0899-8256(03)00048-4.

Bantman-Masum, E. et al. (2024). “Who Are the Gatekeepers of Economics? Geographic Diversity, Gender Composition, and Interlocking Editorship of Journal Boards”. In: *Review of Political Economy* 37.2. DOI: 10.1080/09538259.2024.2303654.

Becker, G. S. (1973). “A Theory of Marriage: Part I”. In: *Journal of Political Economy* 81.4, pp. 813–846.

Bol, T., M. de Vaan, and A. van de Rijt (2018). “The Matthew Effect in Science Funding”. In: *Proceedings of the National Academy of Sciences* 115.19, pp. 4887–4890. DOI: 10.1073/pnas.1719557115.

Bozio, A., B. Garbinti, J. Goupille-Lebret, M. Guillot, and T. Piketty (2024). "Predistribution versus Redistribution: Evidence from France and the United States". In: *American Economic Journal: Applied Economics* 16.2, pp. 31–65. DOI: `10.1257/app.20220023`.

Brogaard, J., J. Engelberg, and C. A. Parsons (2014). "Networks and productivity: Causal evidence from editor rotations". In: *Journal of Financial Economics* 111.1, pp. 251–270.

Callaway, B. and P. H. C. Sant'Anna (2021). "Difference-in-Differences with Multiple Time Periods". In: *Journal of Econometrics* 225.2, pp. 200–230.

Caplin, A. and M. Dean (2015). "Revealed Preference, Rational Inattention, and Costly Information Acquisition". In: *American Economic Review* 105.7, pp. 2183–2203. DOI: `10.1257/aer.20140117`.

Card, D. and S. DellaVigna (2013). "Nine Facts about Top Journals in Economics". In: *Journal of Economic Literature* 51.1, pp. 144–161. DOI: `10.1257/jel.51.1.144`.

Card, D., J. Heining, and P. Kline (2013). "Workplace Heterogeneity and the Rise of West German Wage Inequality". In: *The Quarterly Journal of Economics* 128.3, pp. 967–1015.

Colussi, T. (2018). "Social Ties in Academia: A Friend Is a Treasure". In: *Review of Economics and Statistics* 100.1, pp. 45–50.

Conley, J. P. and A. S. Önder (2014). "The Research Productivity of New PhDs in Economics: The Surprisingly High Non-Success of the Successful". In: *Journal of Economic Perspectives* 28.3, pp. 205–216. DOI: `10.1257/jep.28.3.205`.

Ductor, L., M. Fafchamps, S. Goyal, and M. J. van der Leij (2014). "Social Networks and Research Output". In: *Review of Economics and Statistics* 96.5, pp. 936–948. DOI: `10.1162/REST_a_00430`.

Fourcade, M., E. Ollion, and Y. Algan (2015). "The Superiority of Economists". In: *Journal of Economic Perspectives* 29.1, pp. 89–114. DOI: `10.1257/jep.29.1.89`.

Frank, R. H. and P. J. Cook (1995). *The Winner-Take-All Society*. New York: Free Press.

Hamermesh, D. S. (2018). "Citations in Economics: Measurement, Uses, and Impacts". In: *Journal of Economic Literature* 56.1, pp. 115–156. DOI: `10.1257/jel.20161326`.

Heckman, J. J. and S. Moktan (2020). “Publishing and Promotion in Economics: The Tyranny of the Top Five”. In: *Journal of Economic Literature* 58.2, pp. 419–470. DOI: 10.1257/jel.20191574.

Jones, B. F. (2009). “The Burden of Knowledge and the “Death of the Renaissance Man”: Is Innovation Getting Harder?” In: *Review of Economic Studies* 76.1, pp. 283–317. DOI: 10.1111/j.1467-937X.2008.00531.x.

Jones, B. F., S. Wuchty, and B. Uzzi (2008). “Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science”. In: *Science* 322.5905, pp. 1259–1262. DOI: 10.1126/science.1158357.

Kline, P., R. Saggio, and M. Sølvsten (2020). “Leave-Out Estimation of Variance Components”. In: *Econometrica* 88.5, pp. 1859–1898.

Lazear, E. P. and S. Rosen (1981). “Rank-Order Tournaments as Optimum Labor Contracts”. In: *Journal of Political Economy* 89.5, pp. 841–864. DOI: 10.1086/261010.

Mackowiak, B. and M. Wiederholt (2009). “Optimal Sticky Prices under Rational Inattention”. In: *American Economic Review* 99.3, pp. 769–803. DOI: 10.1257/aer.99.3.769.

Mackowiak, B. and M. Wiederholt (2015). “Business Cycle Dynamics under Rational Inattention”. In: *Review of Economic Studies* 82.4, pp. 1502–1532. DOI: 10.1093/restud/rdv027.

Merton, R. K. (1968). “The Matthew Effect in Science”. In: *Science* 159.3810, pp. 56–63. DOI: 10.1126/science.159.3810.56.

Nielsen, M. W. and J. P. Andersen (2021). “Global Citation Inequality Is on the Rise”. In: *Proceedings of the National Academy of Sciences* 118.7, e2012208118. DOI: 10.1073/pnas.2012208118.

Park, M., E. Leahey, and R. J. Funk (2023). “Papers and Patents Are Becoming Less Disruptive over Time”. In: *Nature* 613.7942, pp. 138–144. DOI: 10.1038/s41586-022-05543-x.

Petersen, A. M. and O. Penner (2014). “Inequality and Cumulative Advantage in Science Careers: A Case Study of High-Impact Journals”. In: *EPJ Data Science* 3, p. 24. DOI: 10.1140/epjds/s13688-014-0024-y.

Piketty, T. and E. Saez (2003). “Income Inequality in the United States, 1913–1998”. In: *Quarterly Journal of Economics* 118.1, pp. 1–41. DOI: 10.1162/00335530360535135.

Romer, P. M. (1990). “Endogenous Technological Change”. In: *Journal of Political Economy* 98.5, S71–S102. ISSN: 0022-3808. DOI: 10.1086/261725.

Rosen, S. (1981). “The Economics of Superstars”. In: *American Economic Review* 71.5, pp. 845–858.

Sims, C. A. (2003). “Implications of Rational Inattention”. In: *Journal of Monetary Economics* 50.3, pp. 665–690. DOI: 10.1016/S0304-3932(03)00029-1.

Sinatra, R., P. Deville, M. Szell, D. Wang, and A.-L. Barabási (2023). “Ranking Mobility and Impact Inequality in Early Academic Careers”. In: *Proceedings of the National Academy of Sciences* 120.36, e2305196120. DOI: 10.1073/pnas.2305196120.

Wu, L., D. Wang, and J. A. Evans (2019). “Large Teams Develop and Small Teams Disrupt Science and Technology”. In: *Nature* 566, pp. 378–382. DOI: 10.1038/s41586-019-0941-9.

Wuchty, S., B. F. Jones, and B. Uzzi (2007). “The Increasing Dominance of Teams in Production of Knowledge”. In: *Science* 316.5827, pp. 1036–1039. DOI: 10.1126/science.1136099.

Tables

Table 1: Scope of the panel by tier and year.

	1990	2000	2010	2020	2025
Counts					
Active economists	10,751	14,643	27,378	47,325	60,521
Active institutions	1,911	2,670	5,277	7,801	9,958
Articles	9,630	13,015	19,643	25,391	29,019
Coauthorship					
Mean authors per article	1.59	1.65	1.94	2.50	2.73
Mean authors per Top-5 article	1.57	1.73	2.17	2.91	3.08
Share of articles by tier (%)					
Top-5	3.4	2.9	2.0	1.7	1.5
Second-tier general-interest	3.4	1.1	2.5	1.8	1.8
Top-field	7.2	4.9	4.3	4.6	4.5
All others	86.0	91.1	91.2	91.9	92.2

Notes: The table reports the active population, institution count, and article flow of the journal universe at five-year marks, the mean number of authors per article overall and within the Top-5, and the share of articles in each tier. An economist is active in a year with at least one publication in the universe. Read the count rows against the article row to see the population outgrowing capacity, and the tier shares to see output shifting toward the bottom of the hierarchy. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Table 2: Anatomy of entry by margin, selected years.

Year	Entrants (total)	Upper-tier entry	All-others entry	All-others %
1990	10,751	1,839	8,912	82.9
2000	6,043	314	5,729	94.8
2010	11,039	477	10,562	95.7
2025	28,864	1,194	27,670	95.9

Notes: The table reports the annual flow of economists publishing for the first time in the universe, split by whether the entry journal is in the upper tiers (Top-5, second-tier general-interest, or top-field) or in the all-others tier. "All-others %" is the share of entrants joining outside the upper tiers. Read down the final column to see entry shifting toward the bottom of the hierarchy. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Table 3: Cohort composition and persistence of the incumbent top-1%.

Cohort (first econ pub)	Authors in 2025	Share of top-1% (%)	Persistence 2020–2025 (%)
Pre-1991	824	37.3	7.4
1991–1999	2,700	62.7	8.6

Notes: The table reports incumbents (first active by 1999) split into the pre-1991 and 1991–1999 entry cohorts, giving for each the number of authors active in 2025, their share of the incumbent top-1%, and persistence in the top-1% from 2020 to 2025. Read the share column to see recent incumbent cohorts dominating the top, and the persistence column to see year-to-year turnover within it. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figures

Figure 1: Growth of the field relative to its top tier, 1990–2025.

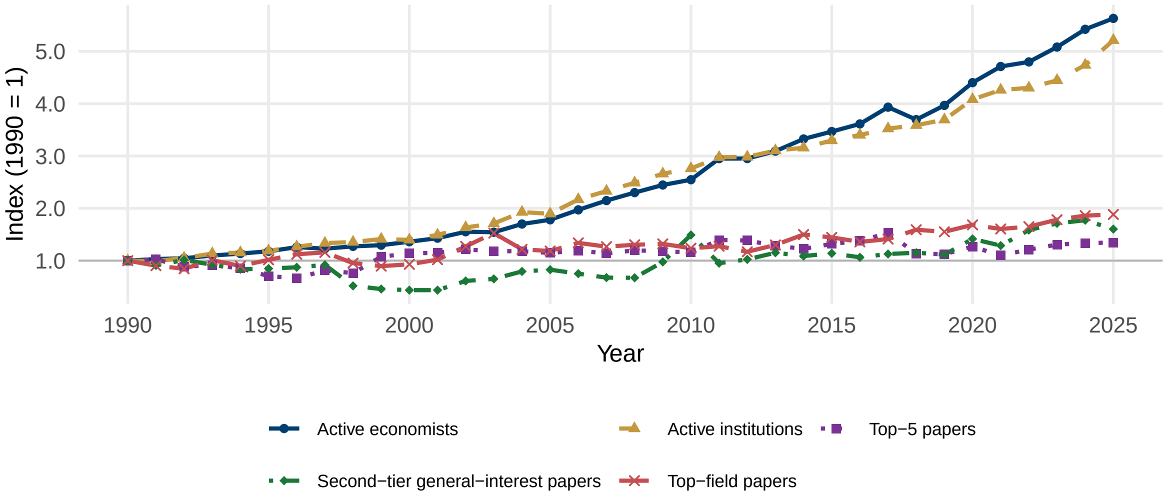

Notes: The figure plots the number of active economists and active institutions against the article capacity of the leading journal tiers, each indexed to its 1990 level. An economist is active in year t with at least one publication in the universe; capacity is the count of articles published in a tier. Read the gap between the population and capacity lines as the widening competition for top-tier space. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 2: Entrant over-representation across journal tiers, 1995–2025.

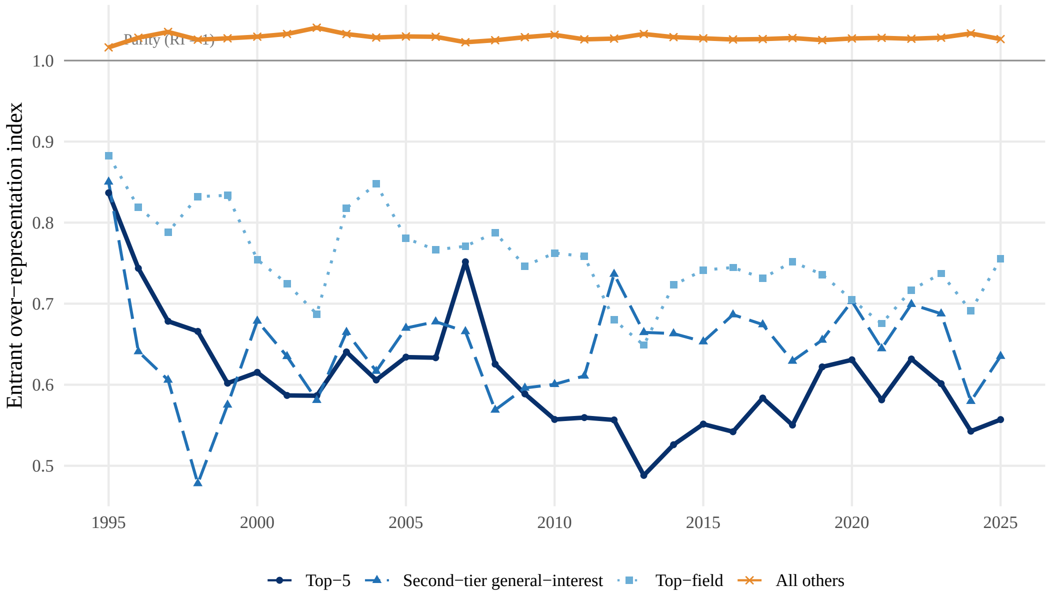

Notes: The figure plots an entrant over-representation index by journal tier. An economist is an entrant in year t if t minus the year of first economics-journal publication is below five. The index is the entrant share of a tier's publication credit (fractional) divided by the entrant share of all publication credit that year, so a value of one means entrants publish in the tier in proportion to their overall footprint and a value below one means they are under-represented. Dividing by the overall entrant share nets out the growth of the entering population, which would otherwise inflate every tier in step. The series begins in 1995 because entry is left-censored at 1990. Read the gap between the Top-5 line and the all-other line as the degree to which entry concentrates at the bottom of the hierarchy. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 3: Where entrants publish and the density of journals, 1990–2025.

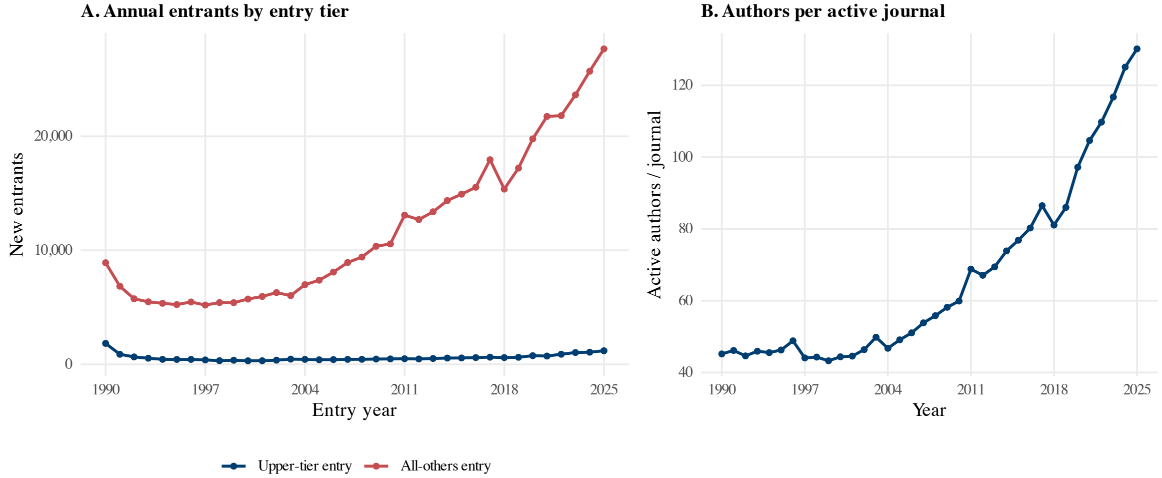

Notes: The figure traces the entry flow by tier alongside the average number of active authors per journal. Read the rising authors-per-journal line as intensification of existing outlets rather than expansion through new ones. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 4: Top-1% publication share, actual versus incumbent-only, for researchers and institutions, 1990–2025.

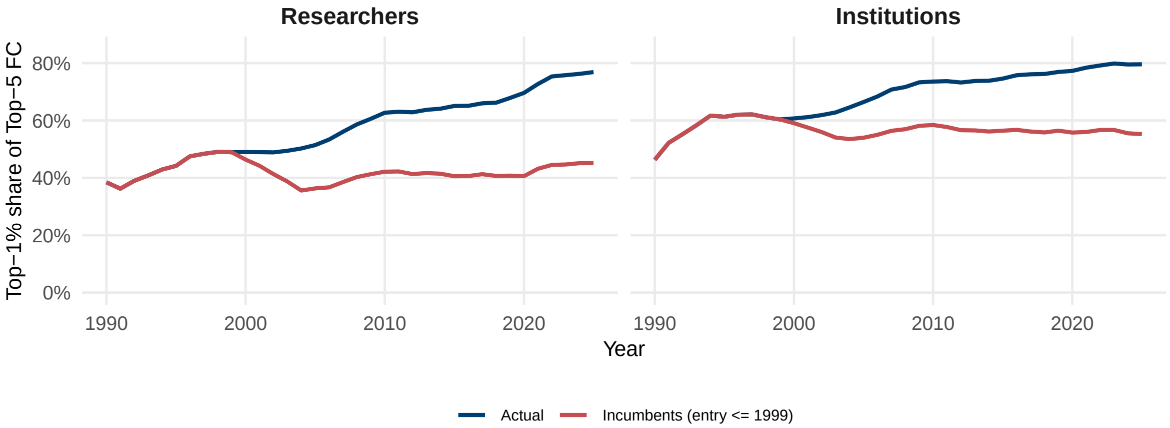

Notes: Each panel plots the actual top-1% fractional-count share against a series that holds the incumbent set fixed at its pre-2000 membership. Read the gap between the two lines as the share of the rise attributable to entry into the bottom of the hierarchy rather than to changes among incumbents. Institution-based series should be read as slopes given affiliation coverage; see Section 2. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 5: Change in top-share concentration among incumbents, by seniority, for researchers and institutions.

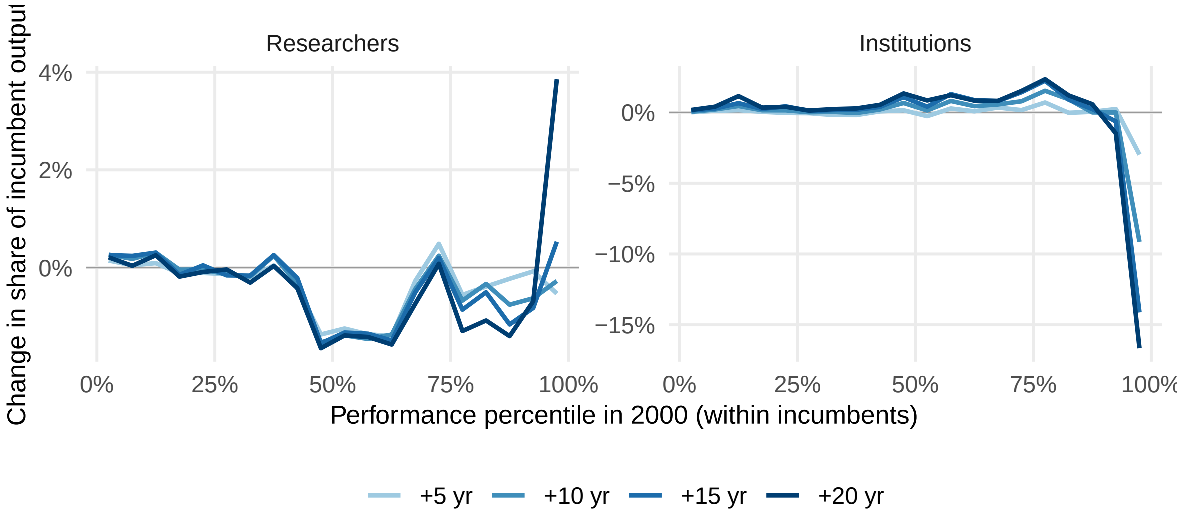

Notes: The figure plots the change in the top share across the incumbent distribution, conditioning on years of seniority so that the population is held fixed. Read the researcher panel as showing the top holding or gaining once composition is netted out, and the institution panel as showing the top bin ceding share to the upper-middle. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 6: Rank–rank persistence of authors and institutions, early versus late window.

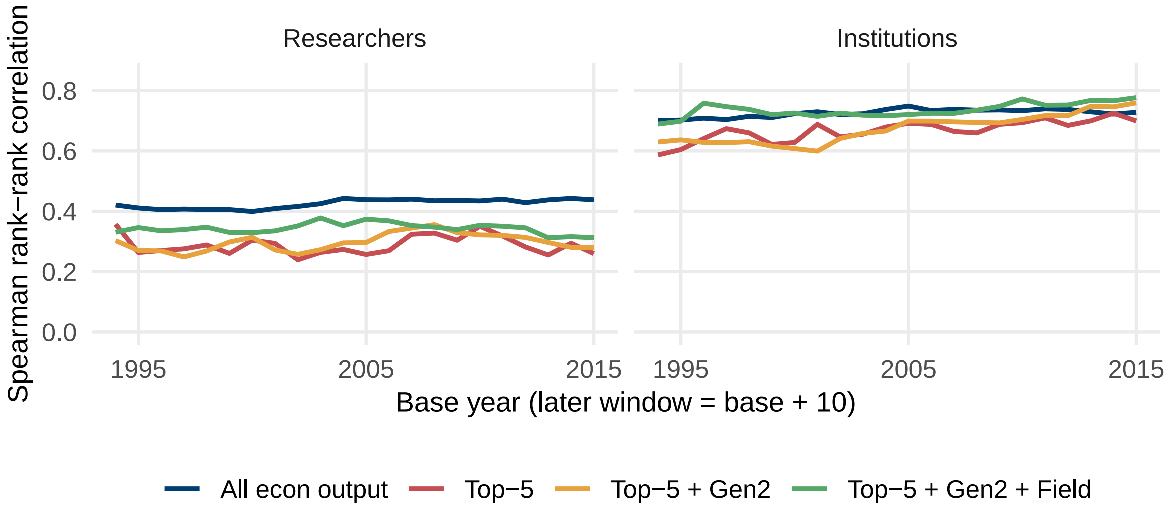

Notes: The figure plots late-window rank against early-window rank in publication credit, separately for authors and institutions, with the Spearman correlation reported in each panel. Persistence is measured field-wide over the full active population, not restricted to incumbents. Read the steeper institution slope as greater stickiness of place than of person. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 7: Quintile persistence among incumbents, 2015 to 2025.

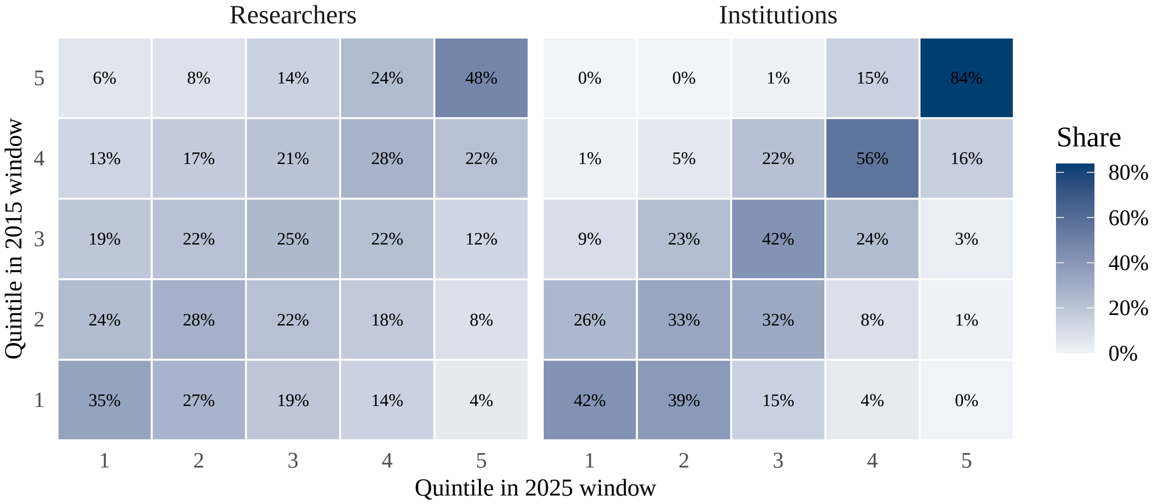

Notes: The figure plots the quintile transition matrix from 2015 to 2025 for incumbents (first active by 1999), with quintiles cut within the incumbent set so that the population is held fixed. Each cell is the share of units in a 2015 origin quintile (vertical axis) that occupy each 2025 destination quintile (horizontal axis), shaded by intensity; rows sum to 100%. The left panel covers researchers, the right institutions, with units ranked on five-year rolling fractional-count output. Read the top-right cell of each panel for the top-quintile stayer rate. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 8: Team size, productivity, and the coauthor pool by tier, 1990–2025.

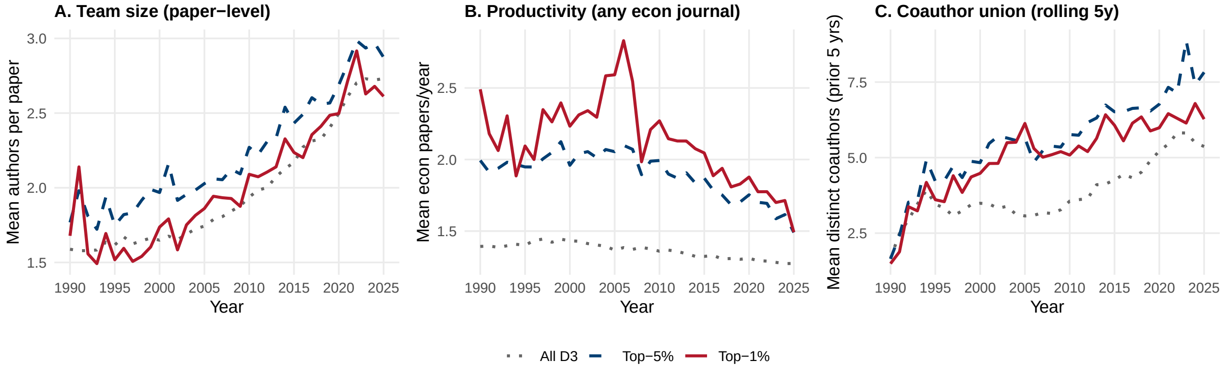

Notes: The figure reports three series for the field (all), the top-5%, and the top-1%. Panel A is the mean number of authors per paper; panel B is mean economics papers per author per year; panel C is the mean number of distinct coauthors over a rolling five-year window. Tiers are defined on five-year cumulative fractional-count output. Read panels A and B as the top not writing larger teams or publishing more per year, and panel C as the top drawing on a far larger rotating coauthor pool. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 9: Elite, cross-institutional coauthorship at the Top-5, 1990–2025.

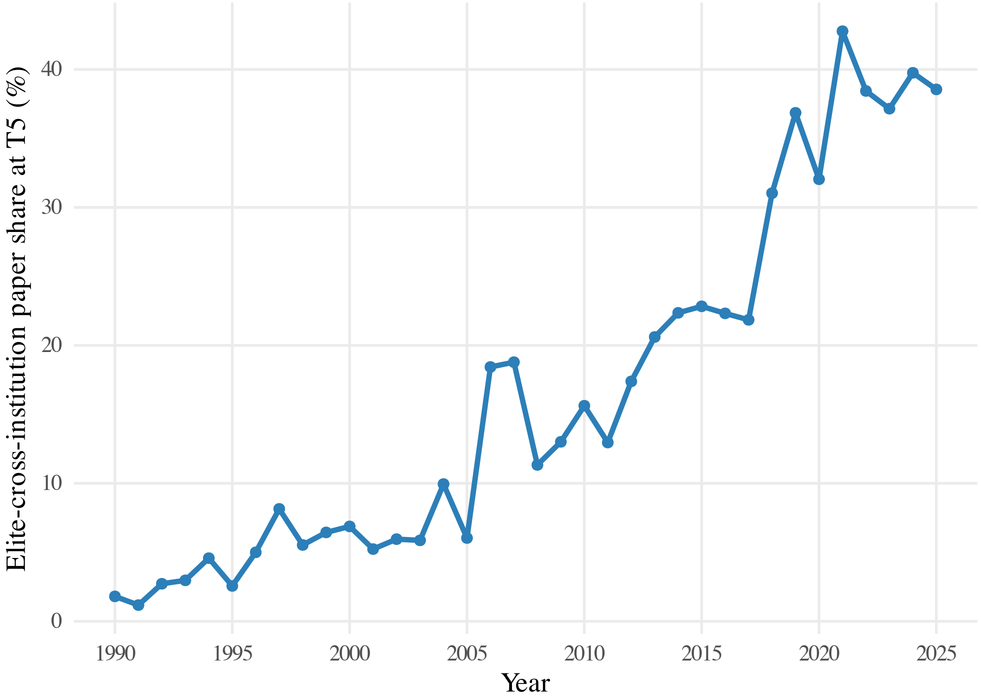

Notes: The figure plots the share of Top-5 papers that are elite and cross-institutional, defined as having at least two top-1% authors affiliated with different institutions. Elite status is assigned on five-year cumulative fractional-count output among authors active at the Top-5 in the year. Read the rising share as elite coauthorship intensifying across universities rather than within them. Institutional affiliation is recorded for more than 80% of credit from 1999 onward. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 10: Topic diversity of the field and of the top-1%, 1990–2025.

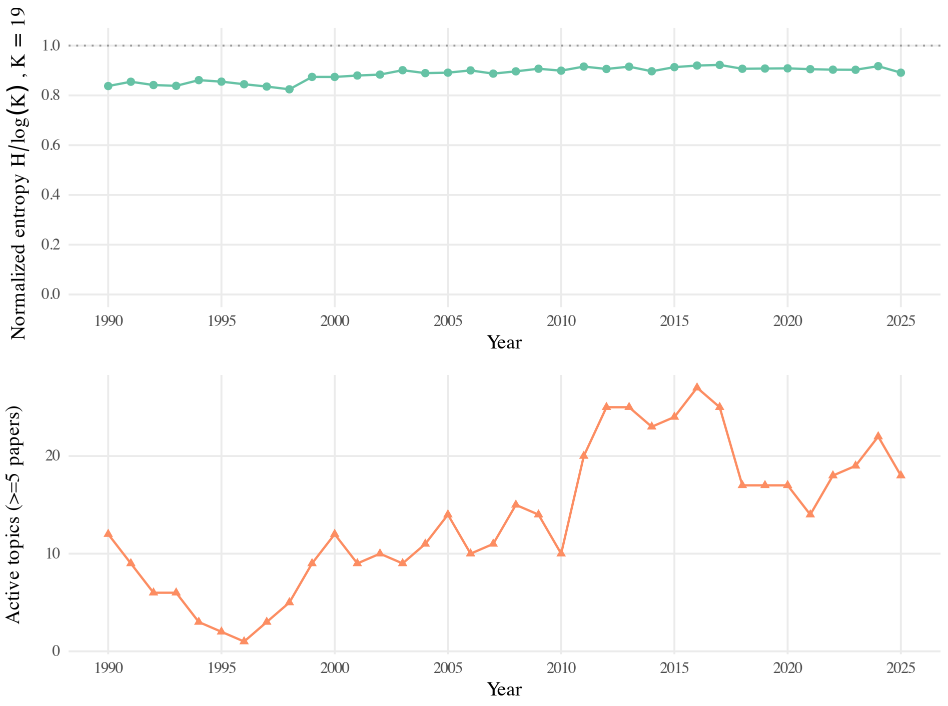

(a) Topic diversity, pooled

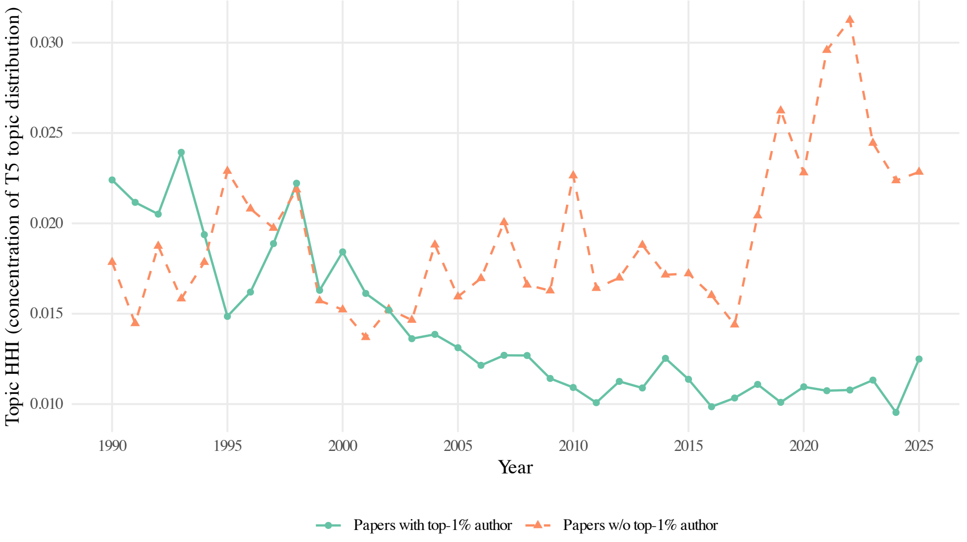

(b) Topic concentration, top-1% versus rest

Notes: Panel (a) plots the pooled topic Herfindahl and normalized entropy for the whole field. Panel (b) plots the topic Herfindahl separately for the top-1% and for the rest of the field. Read the falling top-1% Herfindahl against the rising Herfindahl of the rest as the top diversifying its topics faster than the field. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 11: Elite share around the launch of new general-interest journals.

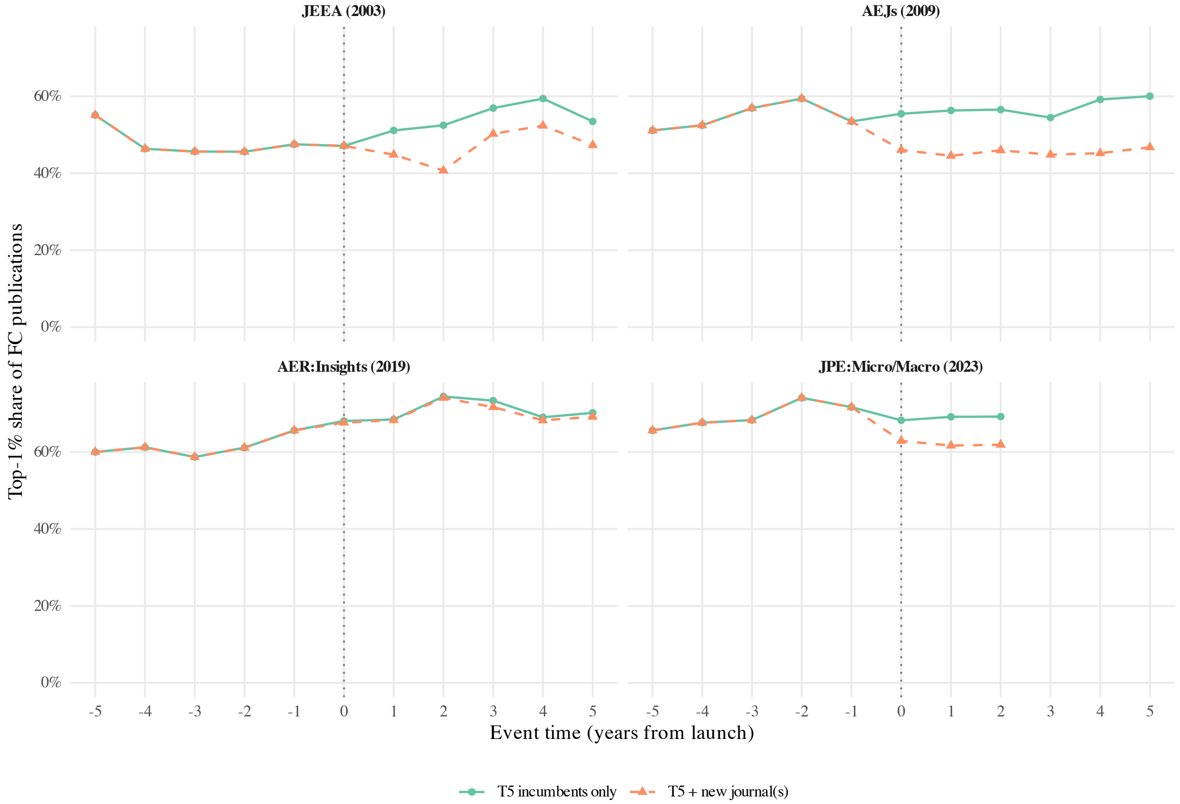

Notes: The figure plots the top-1% share of publication credit in event time around each journal launch, computed over the incumbent Top-5 alone and over the Top-5 plus the newly launched outlet, with the launch year normalized to zero. Read the small gap between the two paths as the limited dilution of the elite share by new outlets. The series are descriptive, without inference. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 12: Reference concentration and the tilt toward old work in Top-5 reference lists.

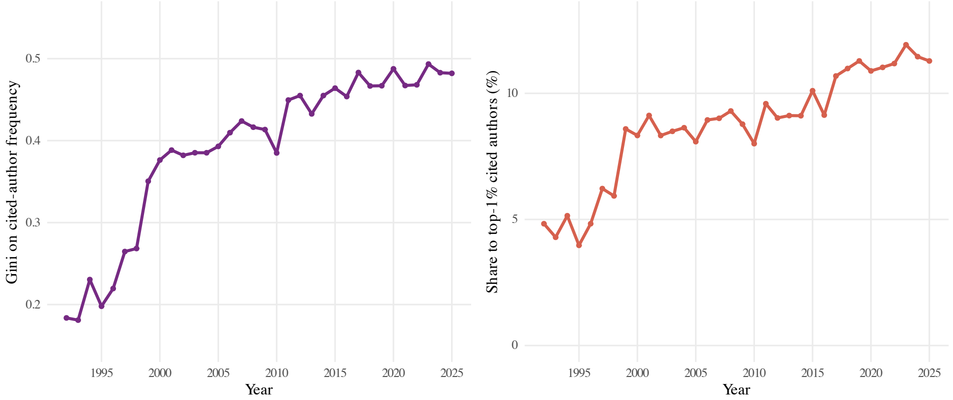

(a) Concentration of cited authors

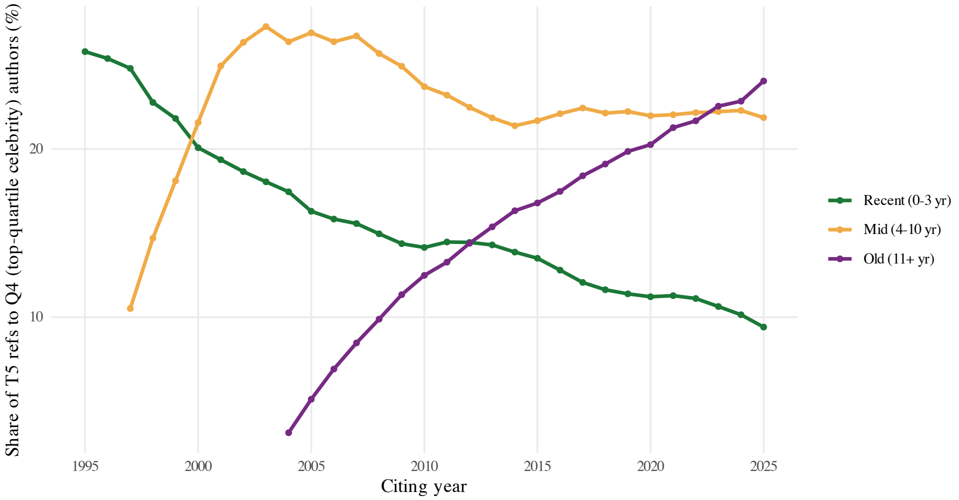

(b) Elite-citation share by vintage

Notes: Panel (a) plots the share of all reference-citations in Top-5 papers going to the top 1% most-cited authors, and the Gini of the cited-author frequency distribution, by citing year. Panel (b) plots the share of elite (top celebrity quartile) citations going to cited work of each vintage, recent (0–3 years old), mid (4–10 years), and old (11+ years), relative to the citing year. Read the rising lines in panel (a) as references concentrating on a few names, and the crossing lines in panel (b) as the elite tilting away from recent work and toward old work. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 13: The editor channel: author-level home-bias and its relation to load.

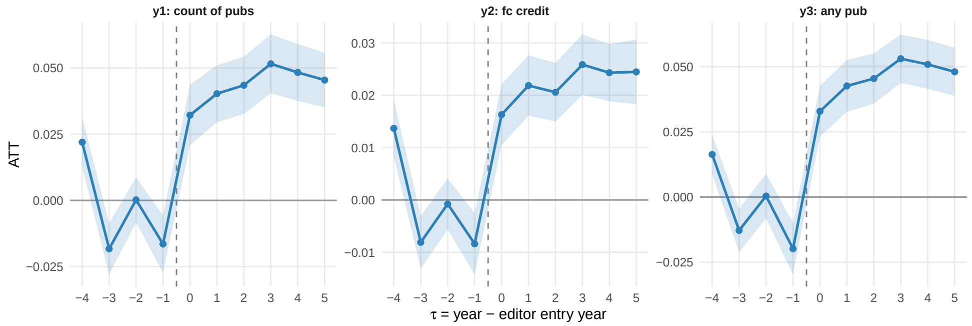

(a) Author-level event study

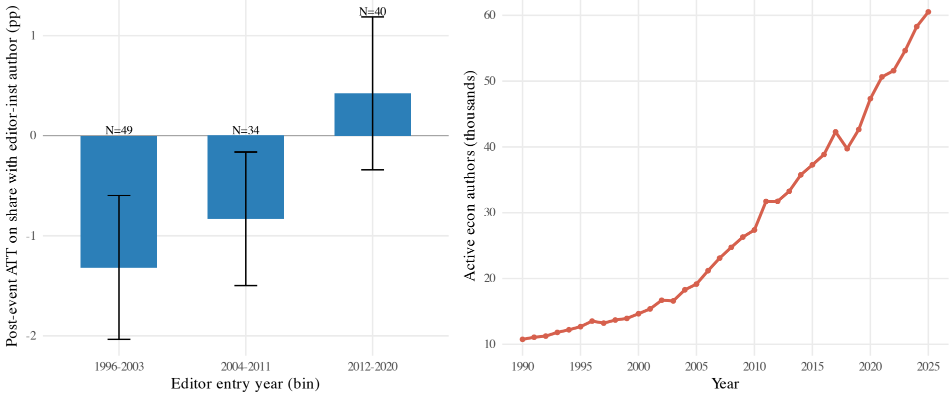

(b) Home-bias by period against load

Notes: Panel (a) plots the dynamic effect of an editor's appointment on a matched author's publishing at the editor's home journal, by event time, estimated by the staggered difference-in-differences estimator of Callaway and Sant'Anna with not-yet-treated editors as controls. Panel (b) plots the per-period editor home-institution effect against the active-population-to-Top-5 load ratio the editor faced. Read panel (a) as a small positive effect with clean pre-trends, and panel (b) as a home-bias effect that rises with load. Standard errors are clustered at the editor level. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 14: Conditional top-1% citation premium by decade.

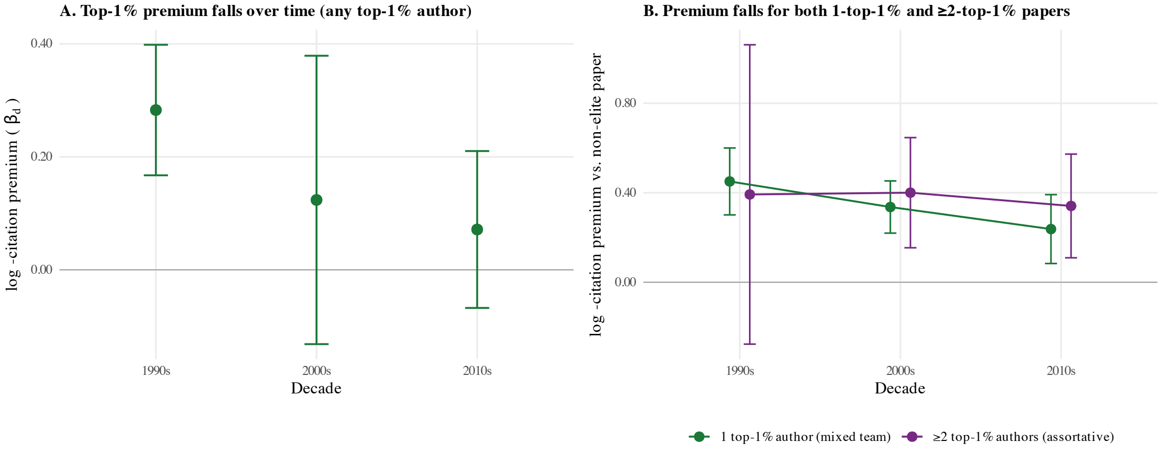

Notes: The figure plots the conditional citation premium of papers with a top-1% author, by decade of publication, from a regression of log ten-year citations on a top-1% indicator interacted with decade, with journal-by-year fixed effects, on papers published 1990–2015. Standard errors are clustered by journal. Read the falling premium as the within-journal advantage of the top narrowing over time. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 15: The link from an early Top-5 paper to top-1% status, by cohort.

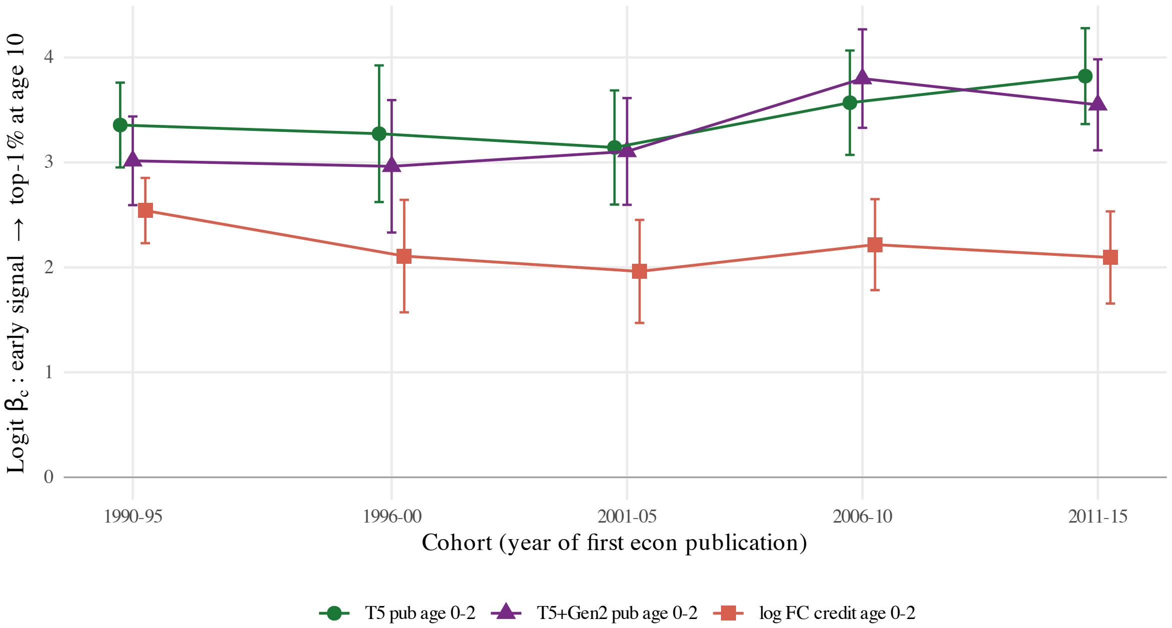

Notes: The figure plots the per-birth-cohort logit coefficient on an early-career signal, at least one Top-5 publication at career age 0–2, predicting top-1% status at career age 10, for cohorts 1990–2015. Read the flat profile as the strength of this link holding constant across cohorts, even as the field grows. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure 16: Returns to an early Top-5 placement.

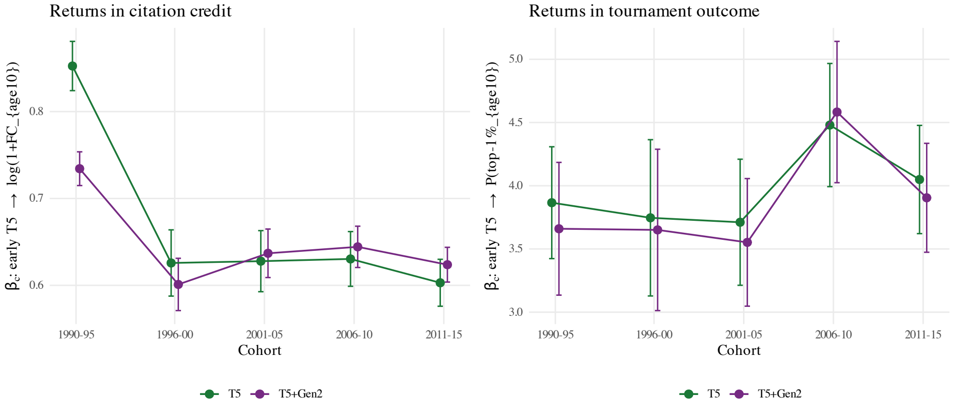
Returns in citation credit

Returns in tournament outcome

Notes: The figure plots the per-cohort return to an early Top-5 placement, at least one Top-5 publication at career age 0–5: the OLS coefficient on log field-citations at age 10 (the citation-magnitude return) and the logit coefficient on top-1% status at age 10 (the rank return). The magnitude return decays once while the rank return holds. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Appendix A The Upper-Tier Journals

Table A.1 lists the 29 upper-tier journals: the Top-5, the second-tier general-interest outlets, and the top-field journals. For each it reports the tier, the first year the journal appears in the panel, the paper count over 1990–2025, and the average papers per year in the first and last five-year windows. Nineteen of the 29 are indexed continuously from 1990; the other ten enter mid-panel, the American Economic Journals between 2006 and 2009, JEEA in 2001 (OpenAlex cover-dates it two years before the 2003 official launch used in Section 3.6), Theoretical Economics in 2007, AER: Insights in 2019, the two JPE field journals in 2022, and JAERE in 2014.

Table A.1: The 29 upper-tier journals: tier, panel entry, and paper counts.

Journal	Tier	1st yr	N	'90-94	'21-25	Journal	Tier	1st yr	N	'90-94	'21-25
AER	Top-5	1990	5,929	115	104	JEconometrics	Top-field	1990	4,409	81	156
ECMA	Top-5	1990	2,074	51	69	JET	Top-field	1990	3,699	74	117
JPE	Top-5	1990	1,979	62	79	JPubE	Top-field	1990	3,645	64	147
REStud	Top-5	1990	1,845	34	91	JDevE	Top-field	1990	3,225	70	154
QJE	Top-5	1990	1,546	49	45	JIE	Top-field	1990	2,624	54	100
AER:I	Top-5	2019	211	0	31	JME	Top-field	1990	2,573	48	84
EJ	2nd-tier	1990	4,520	211	106	JHE	Top-field	1990	2,281	22	83
REStat	2nd-tier	1990	3,010	98	120	JEEM	Top-field	1990	2,202	44	108
JEEA	2nd-tier	2001	1,375	0	60	ET	Top-field	1990	2,037	70	40
AEJ:Policy	2nd-tier	2009	754	0	58	JUE	Top-field	1990	1,762	47	52
AEJ:Micro	2nd-tier	2007	709	0	56	RandJE	Top-field	1990	1,279	38	30
AEJ:Applied	2nd-tier	2006	700	0	48	JoLE	Top-field	1990	1,279	30	48
AEJ:Macro	2nd-tier	2007	582	0	43	TE	Top-field	2007	551	0	43
JPE:Micro	2nd-tier	2022	110	0	22	JAERE	Top-field	2014	462	0	46
JPE:Macro	2nd-tier	2022	56	0	11						

Notes: The table lists the 29 upper-tier journals in two panels. "Tier" is Top-5, 2nd-tier (second-tier general-interest), or Top-field. "1st yr" is the first year the journal appears in the panel; "N" is the count of distinct papers over 1990–2025; "'90-94" and "'21-25" are average papers per year in those five-year windows. Journals are sorted within tier by paper count. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Appendix B Decompositions: Cohort, Gender, and Subfield

This appendix collects the decompositions summarized in Section 3.3. Each shows that the rise in the top-1% share is a within-group phenomenon: it occurs inside entry cohorts (Figures B.1, B.2a and B.2b), within gender (Figure B.3), and within subfields (Figure B.4), rather than through shifts in the cohort, gender, or subfield mix.

Figure B.1: Within- versus between-cohort decomposition of the top-1% share, 1990–2025.

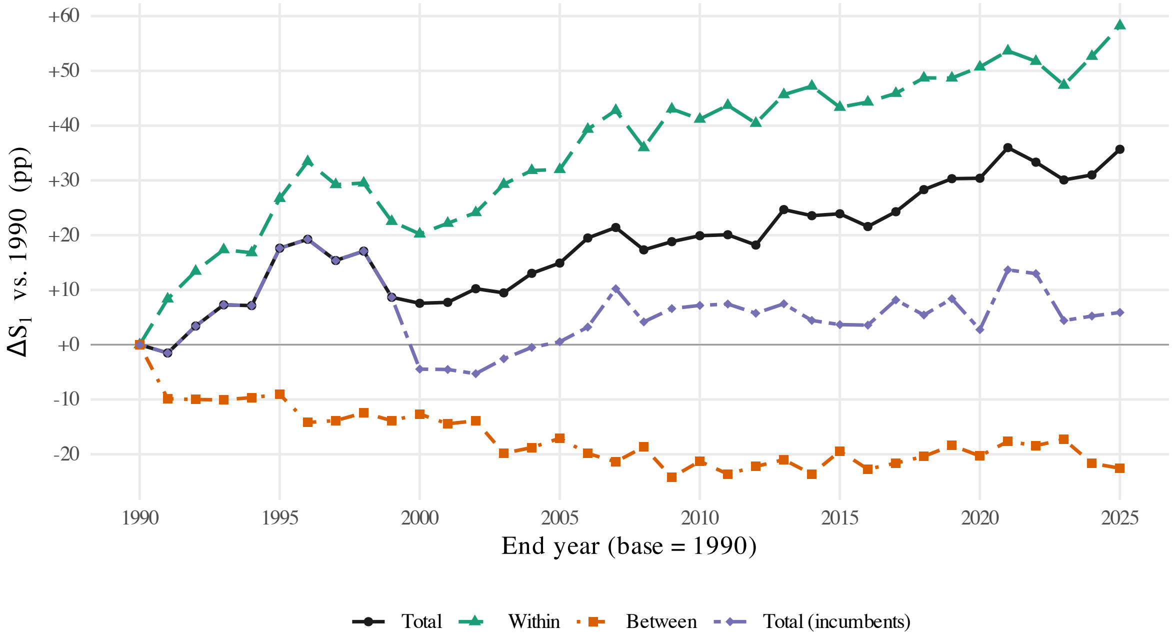

Notes: The figure decomposes the change in the top-1% fractional-count share into a within-cohort component (concentration inside fixed five-year entry cohorts) and a between-cohort component (shifts in cohort size). Read the within line rising and the between line falling as concentration occurring inside cohorts while entry of large young cohorts dilutes the aggregate. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure B.2: Cohort composition and time-to-top of the top-1%.

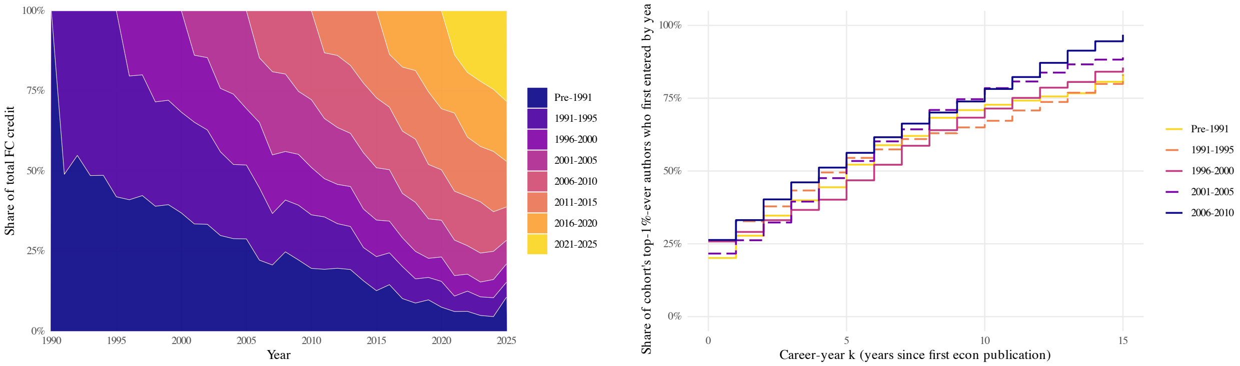

(a) Cohort composition of Top-5 output

(b) Time from first publication to top-1%

Notes: Panel (a) gives the share of annual Top-5 output produced by each five-year entry cohort. Panel (b) gives the median number of years from an author's first publication in the universe to the first year in the top-1%, by cohort; the 2021–2025 value is right-censored. Read the falling time-to-top as recent cohorts reaching the top more quickly. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure B.3: Female representation and the gender decomposition of the top-1% share.

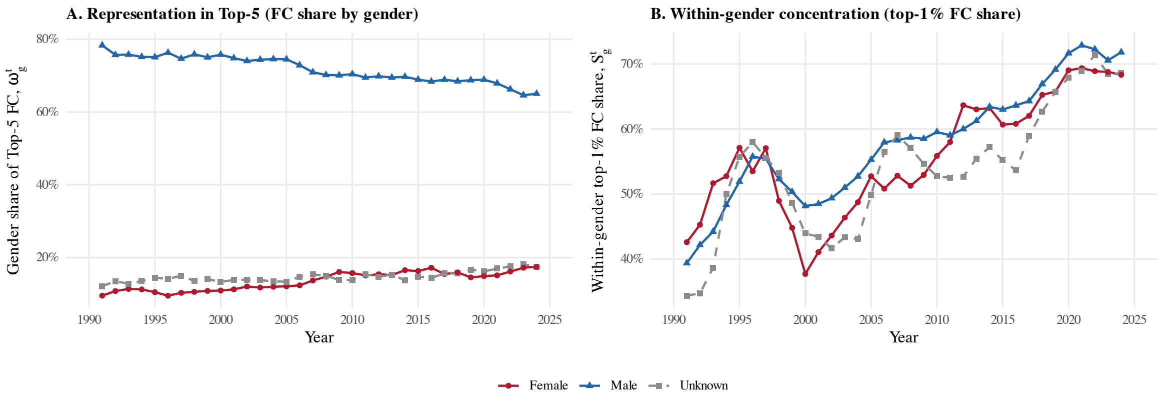

(a) Female share of credit

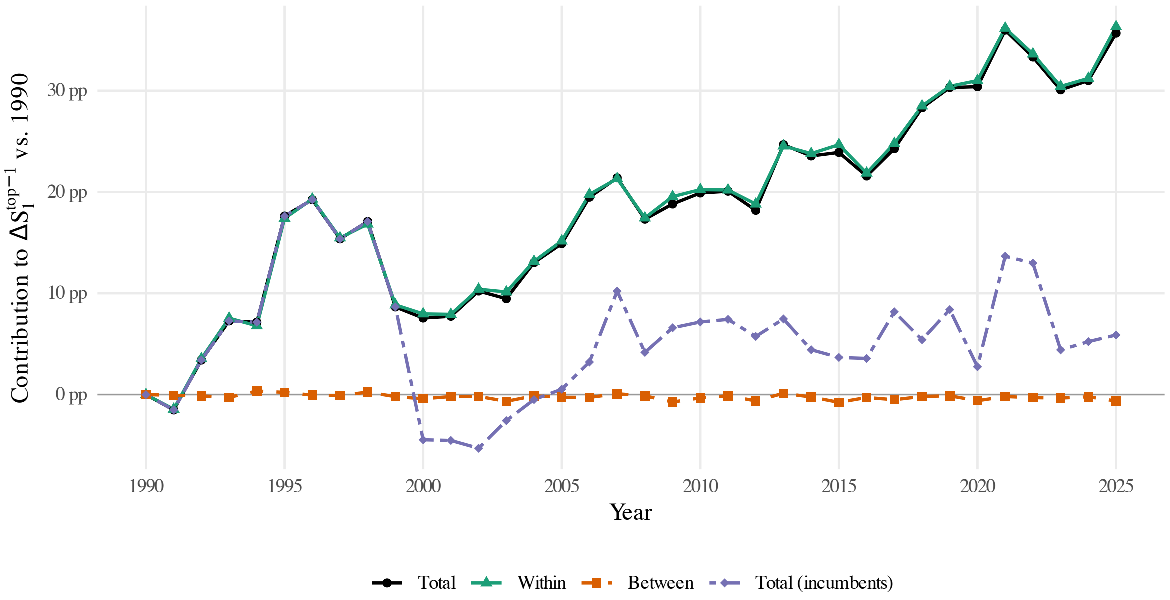

(b) Within- versus between-gender decomposition

Notes: Panel (a) plots the female share of publication credit among authors with assigned gender. Panel (b) decomposes the rise in the top-1% share into within-gender concentration and between-gender composition. Read the near-zero between-gender term as the rise occurring within gender rather than through a change in the gender mix. Gender is unassigned for 34.0% of authors. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure B.4: Within- versus between-subfield decomposition of inequality in publication credit, 2000–2025.

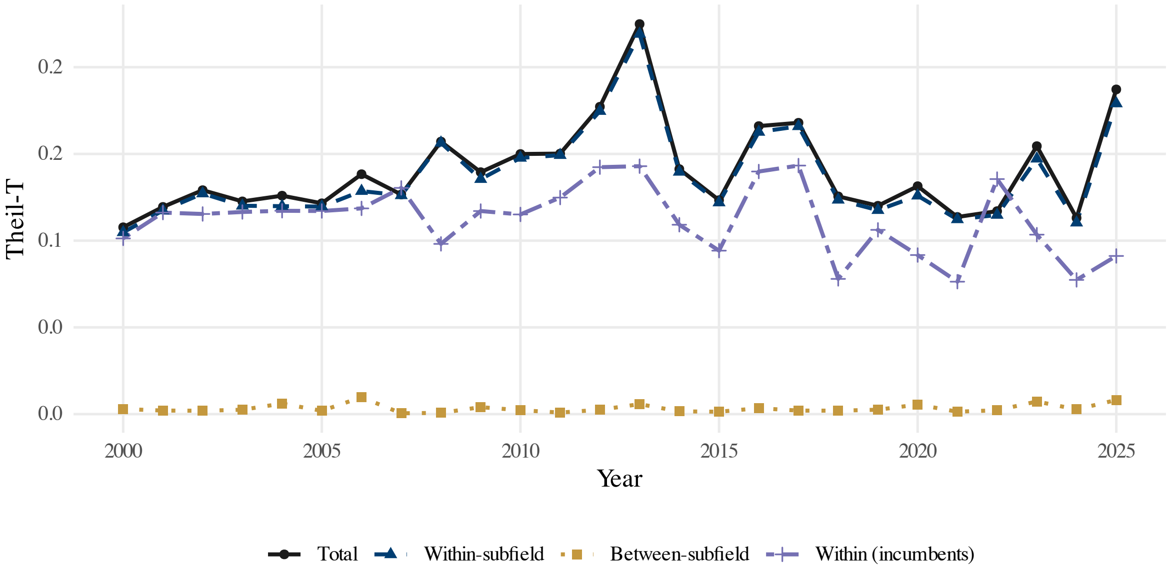

Notes: The figure decomposes the total Theil index of publication-credit inequality into within-subfield and between-subfield components, indexed to a 2000 base. Read the within-subfield share staying above 95% as concentration rising inside subfields rather than across them. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Appendix C Author Disambiguation: Validation and Robustness

Author disambiguation is the central measurement risk. The top-share statistic counts credit per distinct author, so a single economist split across name variants would understate the top and a false merge of two economists would overstate it. OpenAlex assigns each authorship to a disambiguated author identifier; we take that identifier as the unit and stress-test it from two directions. We first validate the identifiers against external signals on a stratified sample. We then re-estimate the headline top-1% share under aggressive de-merging.

The identifiers carry the signals one would expect of a working disambiguation. Table C.1 reports a stratified sample of 439 OpenAlex author identifiers. Half (50.8%) are linked to an ORCID record, and the source declares no false merge: 0.0% carry a `merged_into_id` flag, so OpenAlex has not itself marked a single sampled identifier as a duplicate to be folded into another. Every sampled author carries at least one alternative name spelling and 74.5% carry at least three, the alias multiplicity that diacritic and name-order variation produce. The reverse risk, a single person fragmented across identifiers, appears in 4.9% of cases: that share of ORCID identities maps to two or more OpenAlex author identifiers, a split signal. We have not benchmarked the sample against RePEc or EconLit, so that row is left blank rather than reported as a coverage rate.

Table C.1: Validation of OpenAlex author identifiers on a stratified sample.

Statistic	Value
Sample size (stratified)	439
ORCID-linked	50.8%
False-merge declared (`merged_into_id`)	0.0%
≥ 1 alternative name	100.0%
≥ 3 alternative names	74.5%
ORCID $\rightarrow \geq 2$ OpenAlex IDs (split)	4.9%
Benchmark coverage (RePEc/EconLit)	*not loaded*

Notes: The table reports validation statistics for a stratified sample of 439 OpenAlex author identifiers drawn from the universe. “ORCID-linked” is the share matched to an ORCID record; “False-merge declared” is the share carrying an OpenAlex `merged_into_id` flag, a self-reported duplicate; the alternative-name rows give the share with at least one and at least three recorded name spellings; “ORCID $\rightarrow \geq 2$ OpenAlex IDs” is the share of ORCID identities mapping to two or more identifiers, a fragmentation signal. The benchmark-coverage row is left blank: the sample is not benchmarked against RePEc or EconLit. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

The headline rise survives aggressive de-merging. Table C.2 and Figure C.1 re-estimate the top-1% share under two stress scenarios. The baseline series rises from 38.4% in 1990 to 78.3% in 2025; these are the same numbers the body reports as the fractional shares 0.384 and 0.783, stated here as percentages. Dropping names flagged as high merge-risk lowers the 2025 level to 73.1%, with the 1990 level essentially unchanged at 38.2%. Forcing the worst-case consolidation of identifiers that share a name leaves the 2025 level at 76.0% from a 1990 level of 38.5%. The 2025 top-1% share is therefore bounded in $[73.1, 78.3]$ percent across these scenarios, and the rise of roughly forty percentage points holds in every one. The concentration we document is not an artifact of fragmented identifiers.

Table C.2: Top-1% publication share under disambiguation stress scenarios.

Scenario	Top-1% share 1990	Top-1% share 2025
Baseline (published)	38.4	78.3
Drop high merge-risk names	38.2	73.1
Worst-case ID consolidation	38.5	76.0

Notes: The table reports the top-1% share of publication credit in 1990 and 2025 under the baseline disambiguation and two de-merging stress scenarios. “Drop high merge-risk names” removes identifiers whose name spellings are most prone to collision; “Worst-case ID consolidation” forces same-name identifiers to merge. Shares are in percent; the baseline values equal the body’s fractional shares 0.384 and 0.783. Read the 2025 column as a bound of $[73.1, 78.3]$ percent on the top-1% level. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure C.1: Top-1% publication share, 1990–2025, under disambiguation stress scenarios.

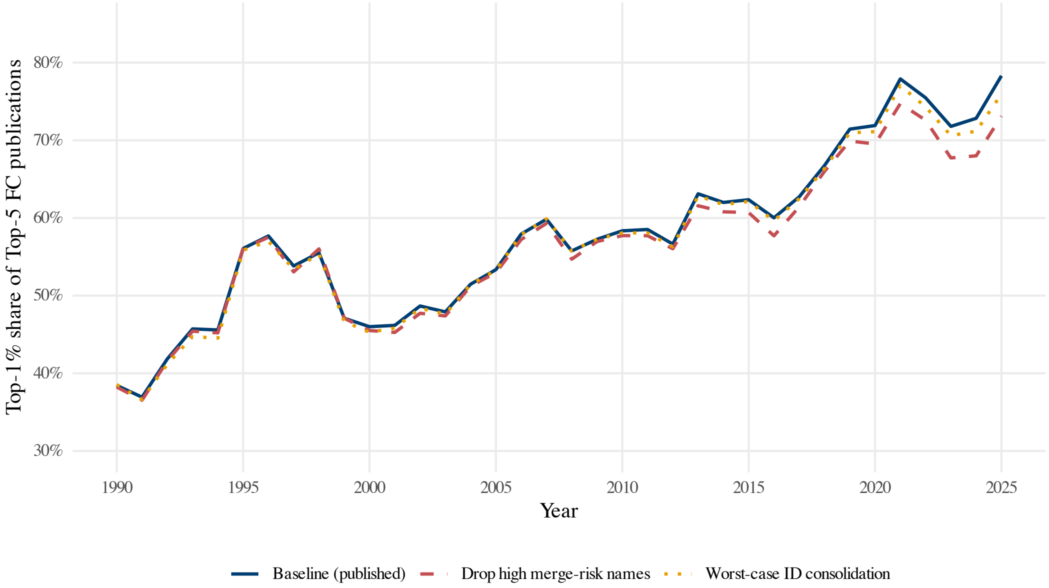

Notes: The figure plots the top-1% share of publication credit over 1990–2025 under the baseline disambiguation and the two de-merging scenarios of Table C.2. Read the three series staying close together as the rise being robust to aggressive de-merging. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

The 2016 back-catalog correction

OpenAlex ingested the full back-catalogs of several Top-5 journals through CrossRef and JSTOR deposits and stamped many of the resulting records with `publication_year = 2016` and a January-1 date, often with no DOI or an obsolete DOI format. Left uncorrected, this produces one anomalous journal-year cell: the *American Economic Review* shows 1,000 Top-5 papers in 2016 against a median of 174.5 across its other years, a 5.7-fold spike. A scan of every journal-year cell over 1990–2025 flags this single cell and no other, so the spike is specific to the back-catalog dump rather than a general feature of the data. We drop the affected records with a fixed rule applied before every aggregation: a Top-5 record dated January 1 with no DOI or an obsolete volume-issue-page DOI. The rule removes 978 work identifiers in all (807 from the *AER*, of which 747 fall in 2016, the remainder pre-DOI January-1 records scattered across *Econometrica*, the *Review of Economic Studies*, the *QJE*, and the *JPE*), which is 6.9% of all distinct Top-5 papers and 0.17% of all works in the universe. After the correction no journal-year cell is anomalous. The correction matters only locally: Figure C.2 overlays the headline top-1% series with and without the rule, and the two coincide everywhere except at 2016. The uncorrected 2016 share is in fact *lower* than the corrected one (0.41 against 0.61), because the spurious records carry many distinct historical authors who inflate the active-population denominator and dilute measured concentration; removing them returns 2016 to the trend set by its neighboring years.

Figure C.2: Top-1% publication share, 1990–2025, with and without the 2016 back-catalog correction.

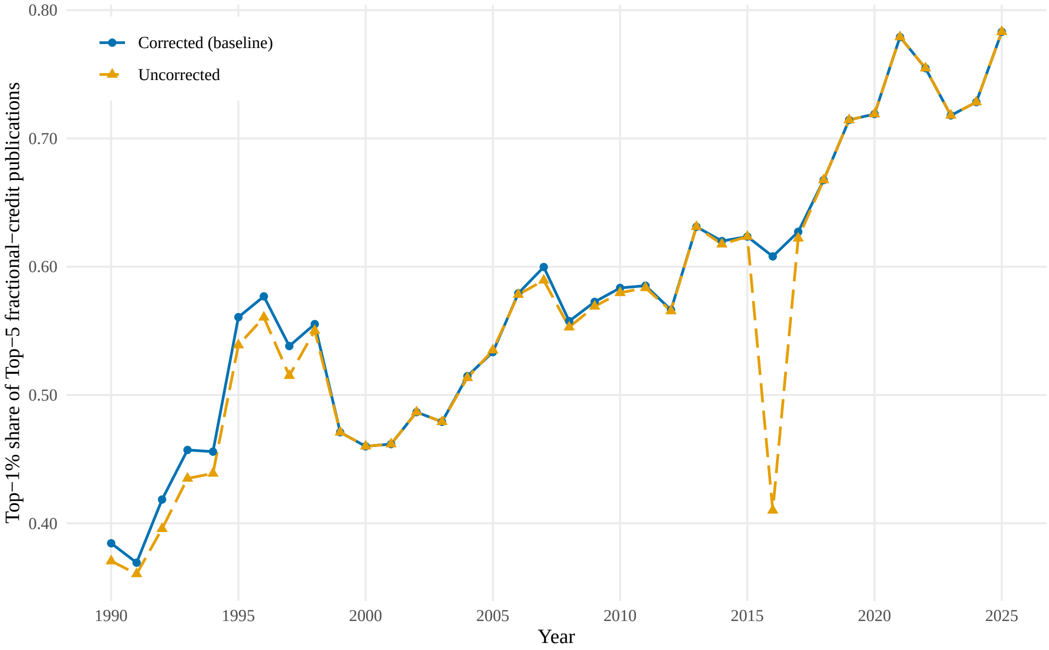

Notes: The figure plots the headline top-1% share of Top-5 publication credit over 1990–2025 computed two ways: with the back-catalog correction applied (as in the body) and without it. The series coincide except at 2016, where the uncorrected series dips because spurious back-catalog records inflate the active-population denominator. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Appendix D The Geography of Growing Downward

The growing downward has a geography. Entry into the field internationalized while access to its top tier did not. This appendix gives the researcher-level counterpart to the country-level account of Aigner, Greenspon, and Rodrik (2025): where Aigner, Greenspon, and Rodrik measure shares across countries, we track individual economists from first publication to top-tier conversion. Region is recoverable for about 84% of credit, so we read the geographic series as slopes rather than exact levels.

Entry internationalized sharply. The non-US share of entrants rose from 47% in 1990 to 87% in 2025. The newcomers filling the bottom of the hierarchy, documented in Section 3.1, are increasingly drawn from outside the United States. The widening of the talent draw that the model of Section 4 assumes thus has a concrete geographic reading: the field grew by admitting economists from a less national pool.

Conversion to the top did not follow entry. Table D.1 and Figure D.1 compare US and non-US entrants by whether they ever reach the Top-5. Across all entry cohorts, 7.73% of US entrants ever publish in the Top-5 against 1.48% of non-US entrants, a ratio above five to one. The gap narrows but persists when the bar is the broader upper tiers rather than the Top-5 alone: 19.1% of US entrants versus 7.1% of non-US entrants ever reach it. The gap also holds within each entry decade. Among 1990s entrants, US Top-5-ever conversion is 12.1% against 5.2% for non-US; among 2000s entrants, 9.4% against 2.6%; among 2010s entrants, 6.2% against 1.1%. Conversion falls for both groups in later decades, the left-censoring of recent cohorts, but the US advantage is present throughout.

Table D.1: Top-5-ever conversion of US versus non-US entrants, by entry decade.

Entry decade	US T5-ever %	Non-US T5-ever %
1990s	12.1	5.2
2000s	9.4	2.6
2010s	6.2	1.1

Notes: The table reports the share of entrants in each entry decade who ever publish in the Top-5, separately for US and non-US economists, where entrants are classified by the region of affiliation at first publication. Later decades show lower conversion for both groups because recent cohorts have had less time to reach the top. Region is recoverable for about 84% of credit. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure D.1: Entry and top-tier conversion of US versus non-US economists, 1990–2025.

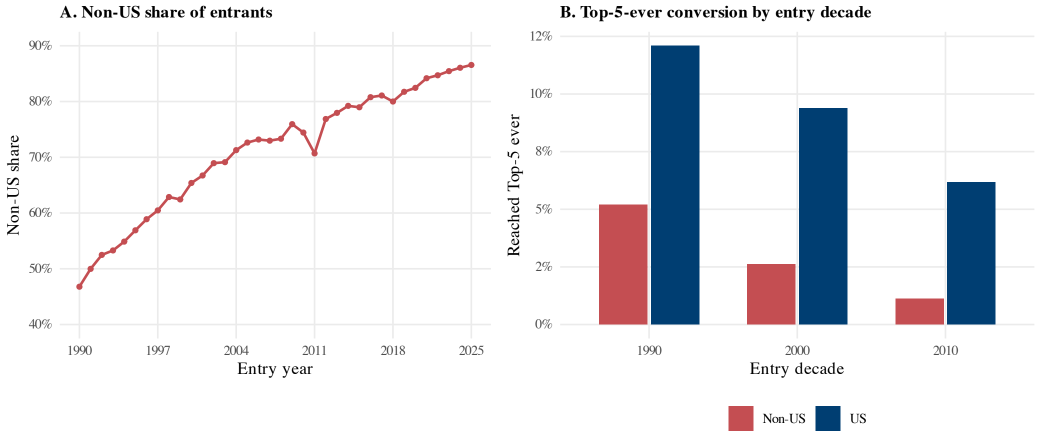

Notes: The figure traces the non-US share of entrants alongside the Top-5-ever conversion rates of US and non-US entrants. Read the rising non-US entry share against the persistent US conversion advantage as internationalization of entry without internationalization of the top. Region is recoverable for about 84% of credit. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

The composition of the periphery, the economists who never reach the upper tiers, flips over the period. Within this group the US share falls from 52% in 1990 to 13% in 2025, while the Asia and Middle East share rises to 40%. The bottom of the hierarchy, where the field's growth concentrated, became predominantly non-US. The growing downward is in this sense an internationalization of the field's base alongside a stable, US-tilted top.

Appendix E Kitagawa–Oaxaca–Blinder Decomposition

The rise in concentration could in principle be team-size arithmetic. Whole-count credit assigns a full unit to each author on a paper, so a shift toward larger teams could raise the top share mechanically, without any change in who writes the papers. This appendix isolates that channel with a Kitagawa–Oaxaca–Blinder decomposition. We split the change in the whole-count top-p% share into an own-concentration term and two matching terms, then show the matching terms are negligible.

Setup. Index papers by j with team size n_j. Let $\tau_t^i = \mathbf{1}[i \text{ is a top-}p\% \text{ author by whole-count ranking in year } t]$, and let $\bar{\tau}_j^t = \frac{1}{n_j}\sum_{i\in j}\tau_t^i$ be the share of top-p% authors on paper j. The whole-count top-p% share is the share of authorships captured by top-p% authors,

$$S_p^{WC,t} = \frac{\sum_{j\in\mathcal{J}_t} n_j\,\bar{\tau}_j^t}{\sum_{j\in\mathcal{J}_t} n_j} \in [0,1], \tag{6}$$

where $\mathcal{J}_t$ is the set of papers in year t. We hold the membership pattern $\bar{\tau}_j^t$ fixed and vary the team-size machinery through two counterfactuals. The B1 counterfactual replaces each n_j with an independent draw $\tilde{n}_j \sim F_{n,1990}$ from the 1990 team-size distribution, giving $S_p^{WC,t,B1}$. The B2 counterfactual additionally rescales the team sizes of top-author papers by ρ_0/ρ_t, where $\rho_t = \mathbb{E}[n_j \mid \bar{\tau}_j^t > 0]/\mathbb{E}[n_j]$ records whether top-p% authors sit on larger ($\rho_t > 1$) or smaller ($\rho_t < 1$) teams than average, giving $S_p^{WC,t,B2}$. The three terms are

$$\underbrace{A = S_p^{WC,t,B2} - S_p^{WC,1990}}_{\text{own concentration}}, \qquad \underbrace{B_1 = S_p^{WC,t} - S_p^{WC,t,B1}}_{\text{random matching}}, \qquad \underbrace{B_2 = S_p^{WC,t,B1} - S_p^{WC,t,B2}}_{\text{assortativity}}, \quad (7)$$

which sum to the observed change $\Delta S = S_p^{WC,t} - S_p^{WC,1990}$. Term A is the change in concentration that would obtain if team sizes had stayed at their 1990 distribution. Term B_1 is the contribution of random matching through team-size composition; term B_2 is the contribution of the top-author propensity ρ_t, the assortativity of authors to team sizes. Term A nets out team-size arithmetic but not author entry, so it remains an upper bound on within-core concentration.

The rise is own concentration. Table E.1 reports the decomposition of the 2025 top-1% share against the 1990 base under four specifications. The observed change is $\Delta S = +30.74$ percentage points. In the canonical specification the own-concentration term is $A = +30.77$, with matching terms $B_1 = -0.26$ and $B_2 = +0.23$ summing to $+30.74$. The matching terms stay within ± 0.1 percentage points of zero in every specification: $B_1 + B_2$ equals -0.03 under the canonical decomposition, -0.03 under the conditional-B_1 stratification (which redistributes the split as $B_1 = -1.91$, $B_2 = +1.88$), -0.03 under the symmetric path-averaged decomposition, and -0.06 under the inverse-propensity-weighted B_2. The own-concentration term is between $+30.77$ and $+30.80$ throughout. The rise in the top share is not team-size arithmetic; it is a change in who writes the papers.

Table E.1: Kitagawa–Oaxaca–Blinder decomposition of the change in the top-1% share, 1990–2025.

Specification	Term A	B1	B2	Sum
Canonical (B1 random + B2 rho)	+30.77	-0.26	+0.23	+30.74
Conditional B1 (stratified by $\bar{\tau}$)	+30.77	-1.91	+1.88	+30.74
Symmetric / path-averaged	+30.77	+0.02	-0.05	+30.74
IPW B2 (propensity reweight)	+30.80	-0.26	+0.20	+30.74
Observed ΔS		+30.74 pp		

Notes: The table decomposes the change in the whole-count top-1% share between 1990 and 2025 into the own-concentration term A and the matching terms B_1 (random matching through team-size composition) and B_2 (assortativity through the top-author propensity ρ_t), per Equation (7). Rows vary the construction of the counterfactuals; columns report the three terms and their sum, all in percentage points. The observed change is $\Delta S = +30.74$. Read the matching terms staying within ± 0.1 of zero as the rise being own concentration, not team-size arithmetic. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Base-year robustness. The own-concentration term does not depend on the choice of base year. Table E.2 recomputes term A for the top-1% share against four base years. The term is $+31.08$ against a 1990 base, $+32.41$ against 1992, $+23.69$ against 1995, and $+29.33$ against 2000. For the top-0.5% share the term ranges from $+11.0$ to $+14.7$ across the same base years. The top-5% column is not identified: the top-5% set of the active population exceeds the entire Top-5 publisher pool in every year, so the whole-count top-5% share saturates and term A is degenerate. We report its value as $+0.00$ in the table only to mark the degeneracy, not as an estimate. Read down the top-1% column to see the own-concentration term holding near thirty percentage points regardless of where the base is set.

Table E.2: Own-concentration term A by base year and top-share threshold.

Base year	Top-0.5%	Top-1%	Top-5%
1990	+12.65	+31.08	+0.00
1992	+14.67	+32.41	+0.00
1995	+11.03	+23.69	+0.00
2000	+11.11	+29.33	+0.00

Notes: The table reports the own-concentration term A of Equation (7), in percentage points, for the top-0.5%, top-1%, and top-5% whole-count shares against four base years. The top-5% column is not identified because the top-5% set of the active population exceeds the Top-5 publisher pool, so the share saturates; the reported $+0.00$ marks this degeneracy rather than an estimate. Read the top-1% column as the own-concentration term being robust to the base-year choice. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Appendix F Two-Way Fixed Effects: Author and Institution Components

This appendix details the decomposition of Section 3.4, separating portable individual productivity from the output premium of an institution. The exercise treats annual publication output as a worker-firm matched panel in the tradition of Abowd, Kramarz, and Margolis (1999), with economists as workers and institutions as firms.

Specification. For economist i affiliated with institution $J(i,t)$ in year t, we model log quality-weighted output

$$y_{it} = \theta_i + \psi_{J(i,t)} + \delta_t + \varepsilon_{it}, \tag{8}$$

where θ_i is the author effect, ψ_j the institution effect, and δ_t a year effect. The outcome is $y_{it} = \log(1+\sum_p c_{ip}\, \omega_{\text{tier}(p)})$, summing over papers p published by i in year t, where c_{ip} is fractional author credit and ω is a tier weight ($\omega = 10$ for Top-5, 5 for second-tier general-interest, 3 for top-field, 1 for all others). The outcome is a flow: a paper contributes to its publication year alone, with weight fixed by its journal's tier, so the measure carries no citation dynamics. Affiliation is the author's modal institution, carried forward between publication years and lagged two years to align the address with the period when the work was produced.

Identification. The author and institution effects are separately identified, up to one normalization, only within a set of institutions connected by economists who move between them (Abowd, Kramarz, and Margolis, 1999). We estimate on the largest connected set: 458,551 author-year observations, 112,412 authors, 8,079 institutions, and 35,268 movers, with a mean of 10.6 movers per institution. Table F.1 reports the panel scope. Identification rests on the assumption that mobility is uncorrelated with the transitory component ε_{it}, that is, that economists do not move in response to idiosyncratic output shocks at the destination. We assess this assumption below.

Table F.1: Scope of the two-way fixed-effects panel.

	Value
Author-year observations	458,551
Authors	112,412
with ≥ 2 observations	71,481
Institutions (largest connected set)	8,079
Movers (≥ 2 institutions)	35,268
Mean movers per institution	10.6
Median years observed per author	2

Notes: The table reports the scope of the panel used to estimate Equation (8), restricted to the largest set of institutions connected by movers. A mover is an economist observed at two or more institutions. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Variance decomposition and the leave-out correction. We decompose $\mathrm{Var}(y_{it})$ into the author component $\mathrm{Var}(\theta)$, the institution component $\mathrm{Var}(\psi)$, the sorting term $2\,\mathrm{Cov}(\theta, \psi)$, and a residual. Plug-in estimates of these components are biased when institutions are linked by few movers: estimation error in $\hat{\theta}$ and $\hat{\psi}$ inflates the variance terms and pushes the covariance negative, the limited-mobility bias of Andrews et al. (2008). We correct it with the leave-out estimator of Kline, Saggio, and Sølvsten (2020), which removes the bias without assuming homoskedasticity, using a Johnson–Lindenstrauss random projection with 300 draws to compute the leverage terms at this scale. We validate the projection against an exact computation on a 25-institution subsample: the correlation between the approximate and exact statistical leverages is 0.997.

Table F.2 reports both estimates. The leave-out correction roughly halves the author share, from 46.9% to 38.2%, and collapses the institution share from 3.7% to 1.3% and the sorting term from

-4.1% to -0.3%. The implied correlation between author and institution effects falls from -0.16 to -0.04. The collapse of the institution and sorting terms under the correction is the diagnostic signature of limited-mobility bias: a genuine negative sorting would survive it.

Table F.2: Variance decomposition of log quality-weighted output.

	Plug-in	KSS
Share of variance in log quality-weighted output		
Author component Var(θ)	0.469	0.382
Institution component Var(ψ)	0.037	0.013
Sorting 2 Cov(θ, ψ)	-0.041	-0.006
Correlation corr(θ, ψ)	-0.158	-0.040

Notes: The table decomposes the variance of log quality-weighted annual output from Equation (8) into the author component Var(θ), the institution component Var(ψ), and the sorting term 2 Cov(θ, ψ), each as a share of the total. "Plug-in" is the raw decomposition; "KSS" is the leave-out correction of Kline, Saggio, and Sølvsten (2020), which removes limited-mobility bias. The bottom row reports the implied correlation between author and institution effects. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Institutional dominance is composition. Decomposing the between-institution variance in mean output into a part from the average author effect of an institution's members and a part from the institution effect ψ, the member-composition part accounts for essentially all of the between-institution variance and ψ for almost none, with a small negative covariance between them. A leading institution leads because it hosts high-θ economists, not because the address raises their output. This is the formal content of the claim that places aggregate while people produce.

Match effects and the exogenous-mobility assumption. The decomposition assumes no idiosyncratic author-institution complementarity correlated with moves. We assess this with the event-study and residual-cell diagnostics of Card, Heining, and Kline (2013). Figure F.1 plots output around an economist's first move, grouped by the institution-effect quartile of origin and destination. Trajectories are flat before the move; a transitory dip in the move year appears in all directions alike, including lateral moves, so it reflects relocation disruption rather than a directional complementarity. Movements up and down the institution-effect ladder produce small and roughly symmetric output changes, consistent with the small ψ variance. Figure F.2 plots the mean regression residual by origin and destination quartile, split into pre- and post-move periods; all cells lie within ± 0.015, under 3.5% of the residual standard deviation, with no systematic pattern. Adding a full set of match (author-by-institution) effects raises the R^2 by only 6.6 percentage points, part of which is mechanical overfitting of short spells. The additive model is an adequate description.

Figure F.1: Output around an economist's first institution move, by origin and destination institution-effect quartile.

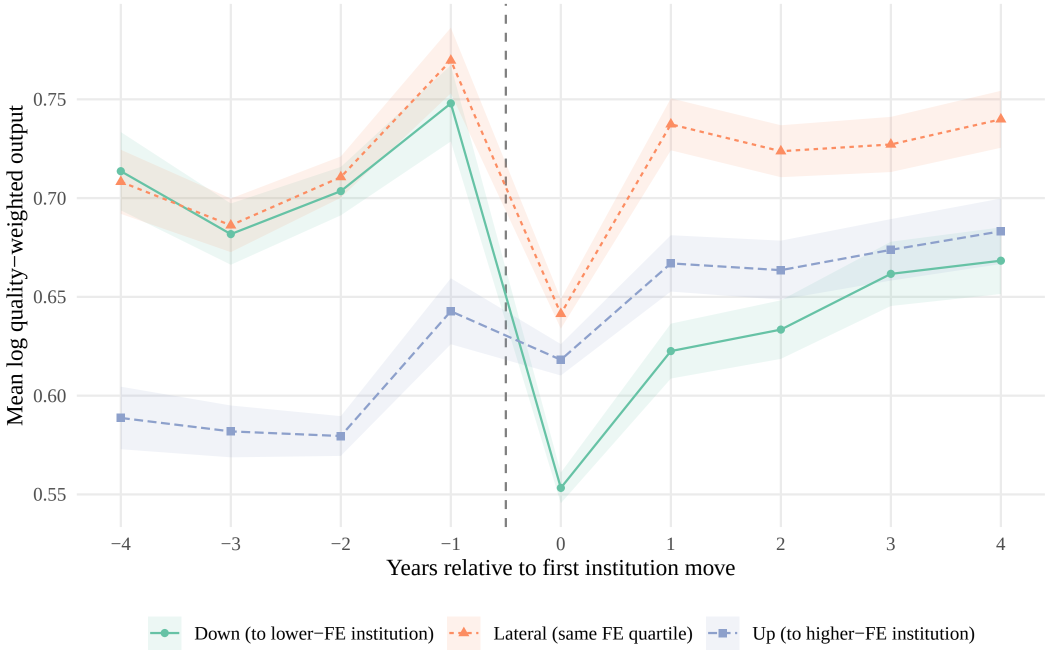

Notes: The figure plots mean log quality-weighted output in event time around an economist's first observed institution move, grouping movers by whether the destination institution-effect quartile is above (up), below (down), or equal to (lateral) the origin quartile. The vertical dashed line marks the move. Bands are 95% confidence intervals. Read the flat pre-move paths and the common move-year dip across all directions as consistent with exogenous mobility rather than match-specific complementarities. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure F.2: Mean regression residual by origin and destination institution-effect quartile.

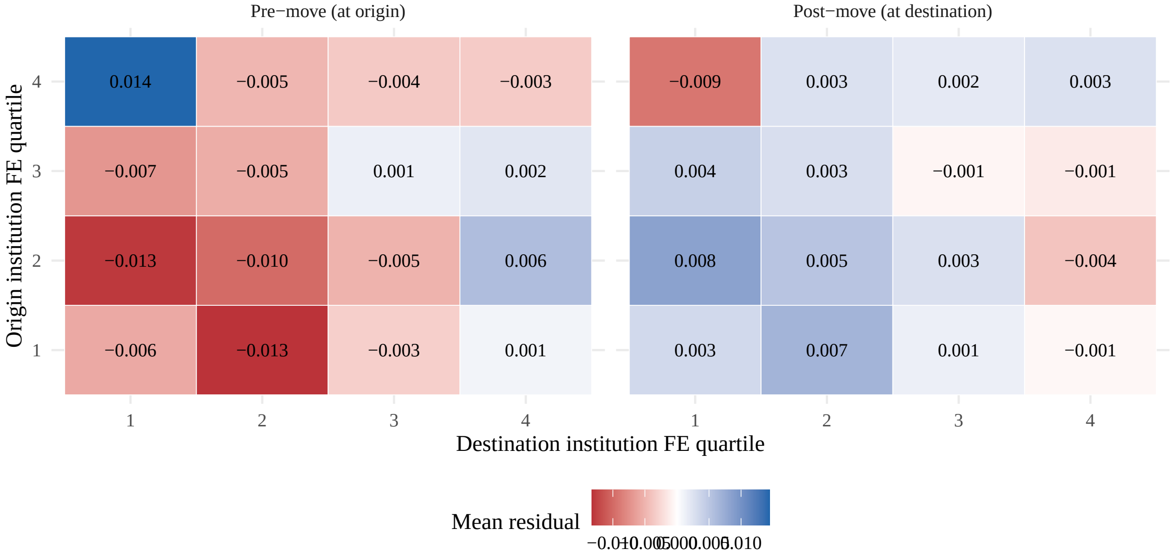

Notes: The figure plots the mean residual from Equation (8) for first-move economists, by the institution-effect quartile of their origin and destination, split into pre-move (at origin) and post-move (at destination) periods. Splitting by period avoids the mechanical within-author zero. Read all cells lying within ± 0.015, under 3.5% of the residual standard deviation, as the absence of a systematic match-effect pattern. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Who moves, and when. Mobility is structured by the career clock but symmetric in direction, which is why it carries no net sorting. Figure F.3 plots experience at move, by move ordinal and direction. First moves cluster at about three years of experience, consistent with the doctoral-to-first-job transition given that first publication typically occurs during the doctorate; second moves cluster at six to nine years, the tenure horizon; later moves spread further out. Within every ordinal, moves up and down the institution-effect ladder are close to equal in number and identically timed. A concern that the institution effect is driven by economists demoted after an unsuccessful tenure case would predict a downward skew at the tenure-clock horizon; the data show no such asymmetry. Direction-symmetric mobility is the micro-foundation of the near-zero sorting term.

Figure F.3: Experience at move, by move ordinal and direction.

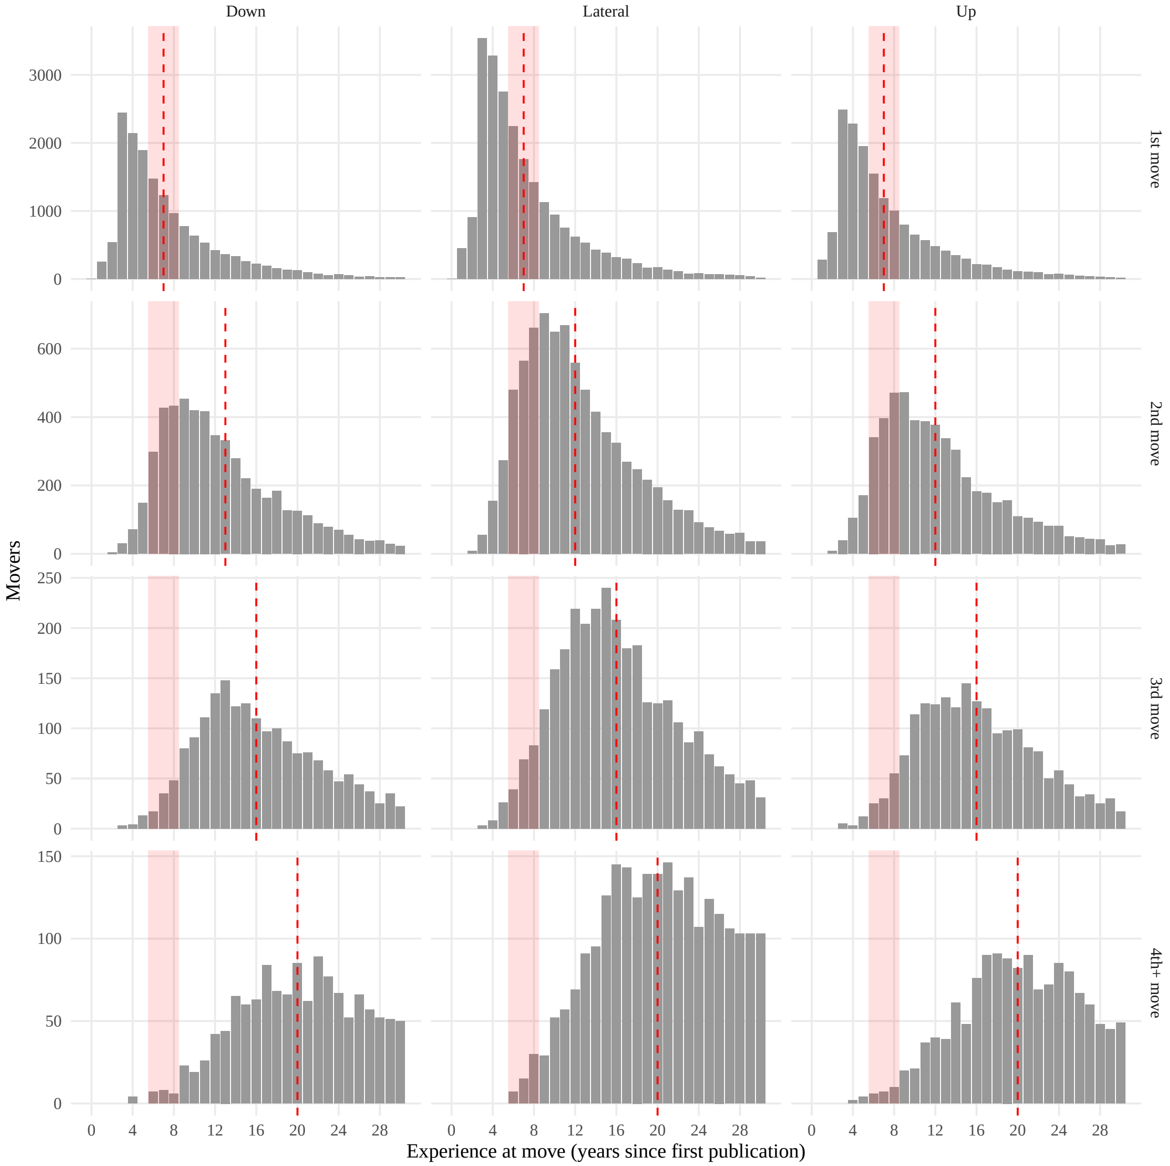

Notes: The figure plots the distribution of experience (years since first publication) at the time of a move, faceted by move ordinal (first through fourth-or-later) and by direction in the institution-effect ranking (down, lateral, up). Dashed lines mark medians; the shaded band marks the six-to-eight-year tenure horizon. Read the near-identical down and up columns within each ordinal as direction-symmetric mobility. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Limitations. Three bound the reading. The institution effect is a host effect, not a training effect: it conditions on the author's accumulated productivity and so does not capture whether doctoral institutions build the talent that later moves. The outcome is an output flow weighted by journal tier, so the decomposition speaks to productivity, not to access or to citation impact. And affiliation coverage is incomplete, as Section 2 notes, so the institution effect is estimated on the subset of author-years with a recoverable address and should be read as a slope.

Appendix G Proofs for Section 4

This appendix gives the proofs of the results in Section 4. Lemma 1 and propositions 1, 3 and 5 are complete; Propositions 2 and 4 rely on the high-level conditions stated in their assumptions (Assumption 6 and symmetric host sizes), flagged at the point of use; the citer argument of Extension 1 is a sketch, as noted. Throughout, $\tau_\varepsilon(L_t) = (\sigma_0^2 L_t)^{-1}$ with $L_t = N_t/A$, and $w_t = \tau_\varepsilon/(\tau_q + \tau_\varepsilon)$ as in Definition 1.

Proof of Lemma 1. Treat w_t as a function of the signal precision $\tau_\varepsilon > 0$. By the quotient rule,

$$\frac{\partial w_t}{\partial \tau_\varepsilon} = \frac{(\tau_q + \tau_\varepsilon) - \tau_\varepsilon}{(\tau_q + \tau_\varepsilon)^2} = \frac{\tau_q}{(\tau_q + \tau_\varepsilon)^2} > 0,$$

since $\tau_q > 0$. Under Assumptions 1 and 2, $\tau_\varepsilon(L_t) = (\sigma_0^2 N_t/A)^{-1}$, so $\partial\tau_\varepsilon/\partial N_t = -A/(\sigma_0^2 N_t^2) < 0$ and $\partial\tau_\varepsilon/\partial A = 1/(\sigma_0^2 N_t) > 0$. The chain rule gives $\partial w_t/\partial N_t = (\partial w_t/\partial\tau_\varepsilon)(\partial\tau_\varepsilon/\partial N_t) < 0$ and $\partial w_t/\partial A = (\partial w_t/\partial\tau_\varepsilon)(\partial\tau_\varepsilon/\partial A) > 0$. □

Proof of Proposition 1. From the reputation-conditional cutoff $c_t(r) = [\bar{q}_t - (1 - w_t)\mu(r)]/w_t$, the $\bar{q}_t$ terms cancel in the difference, so for $r_H > r_L$,

$$\Delta c_t = c_t(r_L) - c_t(r_H) = \frac{(1 - w_t)\big[\mu(r_H) - \mu(r_L)\big]}{w_t} = \frac{1 - w_t}{w_t}\big[\mu(r_H) - \mu(r_L)\big],$$

positive because μ is increasing (Assumption 3) and $w_t \in (0, 1)$. Write the reliance ratio in precision terms: since $w_t = \tau_\varepsilon/(\tau_q + \tau_\varepsilon)$, we have $1 - w_t = \tau_q/(\tau_q + \tau_\varepsilon)$ and $(1 - w_t)/w_t = \tau_q/\tau_\varepsilon = \tau_q\sigma_0^2 N_t/A$ under Assumption 2. This is increasing in N_t, so $\partial\Delta c_t/\partial N_t > 0$; and $1 - w_t$ is increasing in N_t because τ_ε falls in N_t (Lemma 1). For the acceptance probabilities, $a_t(r) = \Pr(s \geq c_t(r) \mid r)$

is decreasing in the cutoff $c_t(r)$; with a Gaussian, hence log-concave, signal density the map from cutoff to acceptance probability is monotone, so a wider cutoff gap Δc_t produces a wider acceptance-probability gap $a_t(r_H) - a_t(r_L)$. $\square$

Proof of Proposition 2. Fix A and F_θ (Assumption 5). A paper with signal s from an author with reputation r is accepted iff $w_t s + (1 - w_t)\mu(r) \geq \bar{q}_t$, equivalently iff $s \geq c_t(r) \equiv [\bar{q}_t - (1 - w_t)\mu(r)]/w_t$. The per-author acceptance probability is $a_t(r) = \Pr(s \geq c_t(r) \mid r)$, decreasing in $c_t(r)$ and hence increasing in $\mu(r)$, which is increasing in r (Assumption 3). Differentiating $c_t(r)$ in w_t,

$$\frac{\partial c_t(r)}{\partial w_t} = \frac{-\,\mu(r)\,w_t - [\bar{q}_t - (1 - w_t)\mu(r)]}{w_t^2} \cdot (-1) = \frac{\mu(r) - \bar{q}_t}{w_t^2},$$

so $\partial c_t / \partial w_t$ is increasing in $\mu(r)$: at fixed $\bar{q}_t$, a fall in w_t lowers the cutoff more (or raises it less) for high-reputation authors, pivoting the iso-acceptance frontier in (s, r)-space toward the reputation axis. This pivot holds $\bar{q}_t$ fixed. Capacity (Assumption 4) then forces $\bar{q}_t$ to re-clear $\int a_t(r) f_r(r)\, dr = \kappa$, so the fixed-$\bar{q}_t$ pivot is not by itself the general-equilibrium shift. Assumption 6 closes the step: it imposes that the net admitted-author density $g_t(r) \propto a_t(r)\, f_r(r)$ shifts up in the monotone-likelihood-ratio order as w_t falls, redistributing the fixed mass of κ toward high r. An MLRP-up shift of g_t raises every upper quantile of accepted reputation (this is where log-concavity of F_r is used: it guarantees the shift raises the top fractile rather than fattening an interior tail). Because credit is increasing in $\mu(r)$ (Assumption 3), the credit-weighted top-α share S_α^t rises. Thus $dS_\alpha^t / dw_t < 0$. Composing with Lemma 1 ($dw_t / dN_t < 0$) gives $dS_\alpha^t / dN_t > 0$. $\square$

Proof of Proposition 3. Average admitted quality is $\bar{Q}_t = \mathbb{E}[q_p \mid w_t s_p + (1 - w_t)\mu(r) \geq \bar{q}_t]$, a function of N_t through two channels. First, the weight: $\partial \bar{Q}_t / \partial w_t > 0$. Raising w_t tilts the screen toward the signal $s_p = q_p + \varepsilon_p$, which is increasing in true quality q_p, so the admitted set selects more on q_p and less on reputation; equivalently a lower w_t admits more high-$\mu(r)$ papers with low realized q_p, lowering the mean. With $\partial w_t / \partial N_t < 0$ (Lemma 1), this channel contributes a negative term, the reliance penalty. Second, capacity: by Assumption 4 the acceptance rate is $\pi_t = \kappa / P_t$, and $\partial \pi_t / \partial N_t < 0$. For a fixed admitted distribution, lowering the acceptance rate is an upward truncation, which raises the conditional mean, so $\partial \bar{Q}_t / \partial \pi_t < 0$ and this channel contributes a positive term, the truncation premium. The total derivative is the sum of the two,

$$\frac{d\bar{Q}_t}{dN_t} = \frac{\partial \bar{Q}_t}{\partial w_t}\frac{\partial w_t}{\partial N_t} + \frac{\partial \bar{Q}_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial N_t},$$

with the first term negative and the second positive; their relative magnitudes depend on σ_0, τ_q, and the shape of F_θ, so the sign of $d\bar{Q}_t / dN_t$ is not determined. □

Proof of Proposition 4. Let reputation be individual-portable (Definition 3) and moves costless. The acceptance probability $a_t(r)$ derived above depends on the author's reputation $r_{i,t}$ alone, not on the host institution. Two authors with equal r at institutions $j \neq j'$ therefore have equal expected accepted credit. Write accepted credit as $y_{i,t} = \beta_t\, h(r_{i,t}) + u_{i,t}$, where h is increasing and $u_{i,t}$ collects institution-specific amenities and wage effects on placement that are mean-independent of r. The institutional Herfindahl is $\mathrm{HHI}_t = \sum_j (Y_{j,t}/Y_t)^2$ with $Y_{j,t}$ the credit summed over j's members. As $1 - w_t$ rises, the elasticity of $y_{i,t}$ to individual $r_{i,t}$ rises and the residual institution component shrinks in relative variance; since portable reputation is dispersed across hosts, the share of total credit variance attributable to between-institution differences falls. Under symmetric host sizes this lowers $\sum_j (Y_{j,t}/Y_t)^2$, so HHI_t is non-increasing in $1 - w_t$. □

Proof of Proposition 5. The journal count J_t enters the gatekeeper problem only through the load $L_t = L_t(N_t, J_t)$, with $\eta_J = \partial \log L_t / \partial \log J_t$. By the chain rule,

$$\frac{\partial w_t}{\partial J_t} = \underbrace{\frac{\partial w_t}{\partial \tau_\varepsilon}}_{>0} \cdot \underbrace{\frac{\partial \tau_\varepsilon}{\partial L}}_{<0} \cdot \underbrace{\frac{\partial L_t}{\partial J_t}}_{\mathrm{sign}(\eta_J)},$$

so $\mathrm{sign}(\partial w_t / \partial J_t) = -\,\mathrm{sign}(\eta_J)$. By Proposition 2, $dS^t_\alpha / dw_t < 0$, so $\mathrm{sign}(\partial S^t_\alpha / \partial J_t) = -\,\mathrm{sign}(\partial w_t / \partial J_t) = \mathrm{sign}(\eta_J)$. Thus $\eta_J > 0$ worsens concentration and $\eta_J < 0$ alleviates it. □

Proof of Extension 1. The citer chooses a list S of size K from the pool to minimize the posterior variance of a quality judgment, subject to the attention constraint $\sum_{k \in S} A_c/K \leq A_c$. With Gaussian signals, posterior precision is additive: including candidate k adds precision $\rho(r_k)$, the prior precision its reputation injects, with ρ increasing in r (Assumption 3) and $\rho(\bar{r}) = \tau_q$ at the average reputation. We assume reading a high-reputation work informs the focal judgment through this additive-precision channel; we do not derive the channel from a primitive common-factor structure, which is why this argument is a sketch. The marginal value of including k, per unit of attention A_c/K, is proportional to $\rho(r_k)$. The constrained optimum therefore admits candidates in decreasing order of $\rho(r_k)$, so the optimal inclusion weight is proportional to $\rho(r_k)$ and increasing in r_k. The space cap retains the M candidates of highest $\rho(r_k)$, so the cited set is the top-M reputation tail of the pool. High-reputation authors crowd in and mid-reputation authors thin out; as the pool P_t grows against a fixed M, the cited share captured by the high-reputation tail rises. This is the unconditional citer margin, distinct from the conditional gate margin of Proposition 2. $\square$

Proof of Extension 2. Expected acceptance of a team paper is increasing in the team posterior mean $w_t s + (1 - w_t)\mu(r^{\text{max}})$, where $r^{\text{max}} = \max_k r_k$ over team members under the assumed maximum-reputation team rule. Holding the author's own reputation fixed, this is increasing in the partner's reputation, so the best reply is $\arg\max_j r_j$ over the available pool. Under portable reputation (Definition 3), the partner's institution does not enter the objective, so the choice is independent of institution. For an elite author the maximizer is another elite, who by the dispersion of Proposition 4 is with high probability at a different host. The return to r^{max} is $1 - w_t$, increasing as w_t falls (Lemma 1), so the incentive to pair with the top sharpens with overload, and elite pairing intensifies and becomes more cross-institutional. The maximum-reputation team rule is assumed, not derived; its microfoundation is left open. $\square$

Table G.1: Mapping model objects to data moments.

Model object	Data moment	Margin	Tested in
w_t (Definition 1)	Conditional top-1% citation residual, journal×year FE (narrows)	Gate	P4
S^t_α (Proposition 2)	Top-1% credit share, $38.4\% \rightarrow 78.3\%$	Gate	§3.1
Institutional Herfindahl (Proposition 4)	Institution top share spreads down; ψ explains 1% of output	Gate	§3.2, F
Cross-inst. elite pairing (Extension 2)	Top-1%×top-1% pair share rising	Gate	§3.5
η_J (Proposition 5)	Top-1% share around journal launches	Gate	§3.6
Reading list (Extension 1)	Reference concentration, unconditional	Citer	P1–P2
Opportunity set (Assumption 5)	Non-US entrant share, $47\% \rightarrow 87\%$	—	Appendix D

Notes: The table maps each object of the information-overload model to the empirical moment it rationalizes and to the testable prediction of Section 5. "Margin" records whether the object operates at the gate (screening at submission) or at the citer (reading and citing); the two margins move the conditional and unconditional concentration measures in opposite directions. The calibration matches the rise in the top-1% share, the rise in cross-institutional elite pairing, and the fall in institutional concentration jointly, using the observed active population and the single noise-load elasticity σ_0. The cross-institutional pairing row assumes the maximum-reputation team rule of Extension 2. Source: model of Section 4; moments from Section 3.

Appendix H Supplementary Figures

This appendix collects robustness displays for the headline concentration trend. Each varies one measurement choice and leaves the rise in the top share unchanged. Figure H.1 varies the credit rule, Figure H.2 varies the smoothing window, and Figure H.3 varies the journal universe.

Figure H.1: Top-share trend under fractional versus whole counting of credit.

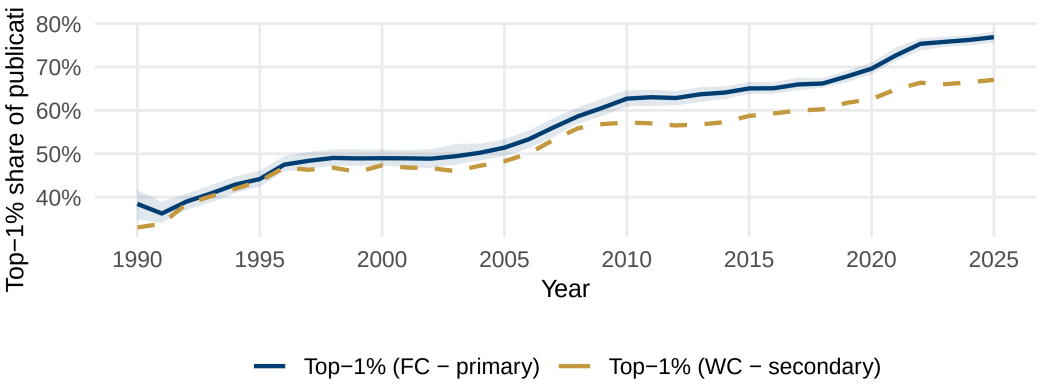

Notes: The figure plots the top-share trend under fractional counting, which splits a paper's credit across its authors, against whole counting, which assigns a full unit to each author. The level differs across rules but the upward trend is common to both. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure H.2: Top-share trend under alternative rolling-window lengths.

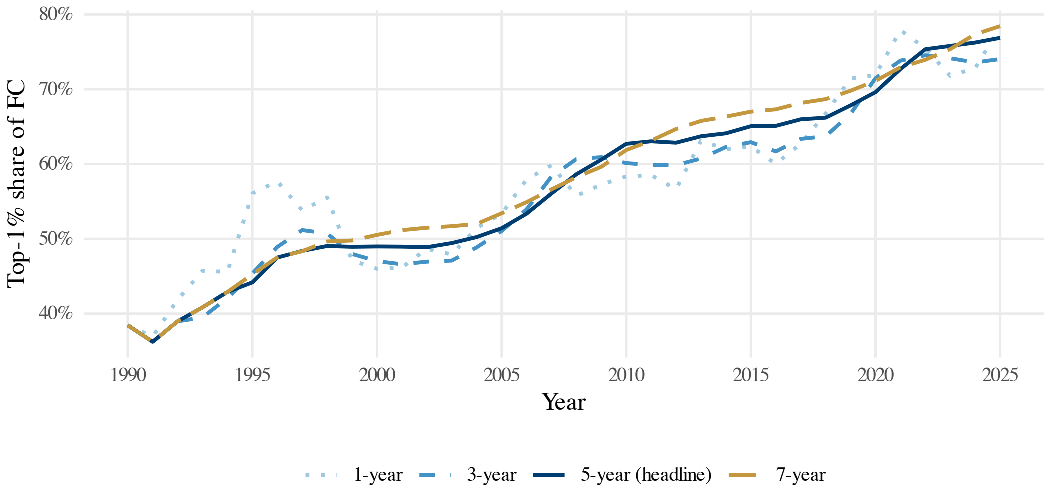

Notes: The figure plots the top-share trend computed over alternative rolling-window lengths used to define the active population. The rise is present at every window length, so it is not an artifact of the smoothing choice. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure H.3: Top-share trend under extended journal universes.

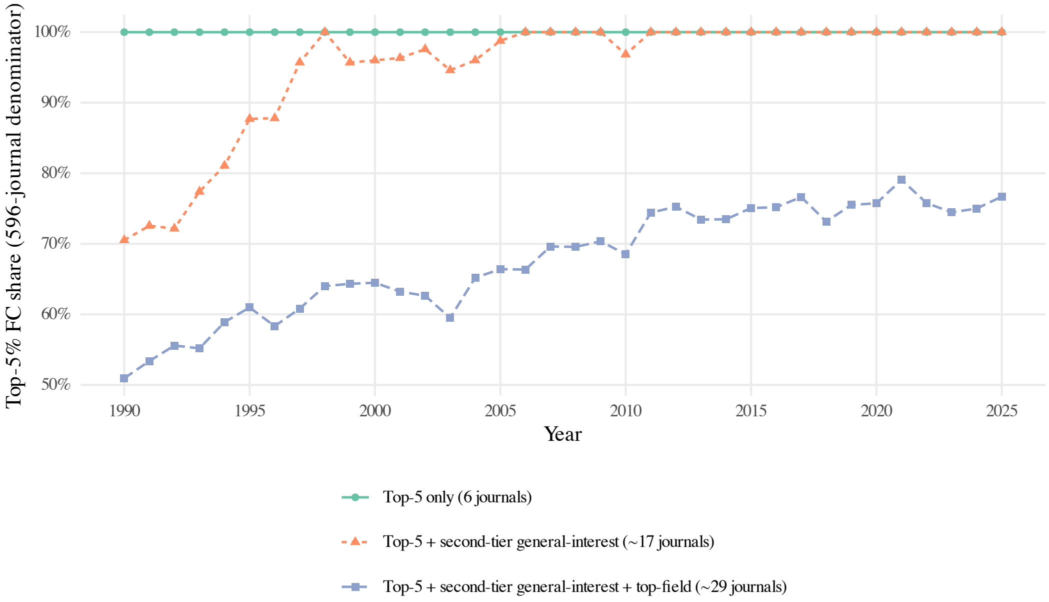

Notes: The figure plots the top-share trend under journal universes broader than the 596-journal baseline. The trend tracks the baseline across the extended universes, so it does not depend on the journal set. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803) and extended universes.

Figure H.4 reports the degree-preserving rewiring counterfactual for elite matching discussed in Section 3.5.

Figure H.4: Observed elite-matching share versus a degree-preserving rewiring counterfactual, 1990–2025.

Notes: The figure plots the observed share of coauthorship links that join two top authors against a counterfactual that rewires coauthorship at random while holding each author's paper count fixed, with a 95% confidence band on the counterfactual. Read the gap between the two lines as the part of the rise in elite pairing that goes beyond what paper counts alone would produce. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Figure H.5 reports the direct test of reputation portability discussed in Section 3.4: movers' Top-5 output around an institutional move, by the prestige direction of the move.

Figure H.5: Top-5 output around an institutional move, by destination prestige.

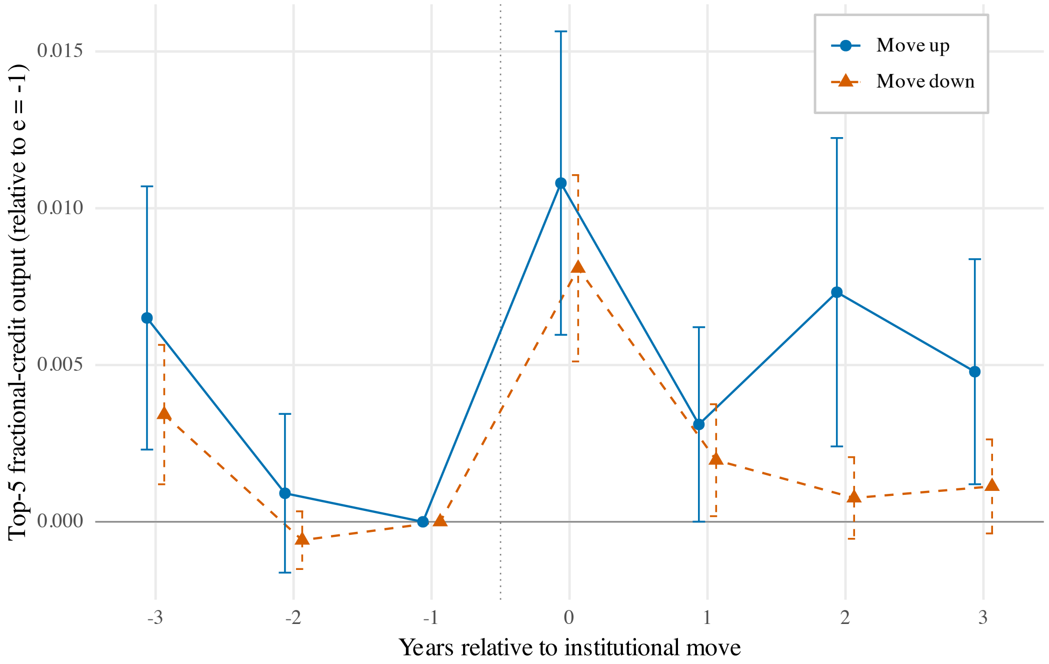

Notes: The figure plots a mover's Top-5 fractional-credit output in event time around a single institutional move, relative to the year before the move, separately for moves up and down the institutional prestige ranking (institutions ranked by total Top-5 credit). The sample is 7,582 clean single movers with at least three active years on each side. Bands are 95% confidence intervals. Read the flat profiles, common across move directions, as Top-5 output traveling with the researcher rather than with the destination, the condition reputation portability requires. The object is realized output, not acceptance, since submissions are unobserved. Source: OpenAlex; 596-journal 2025 ABDC universe (fields of research 3801–3803).

Table H.1 compares the editor-turnover events that enter the difference-in-differences estimation sample with those that drop out, discussed in Section 5.

Table H.1: Editor-turnover events: resolved (estimation sample) versus unresolved.

	Resolved	Unresolved	Difference
Appointment decade (%)			
1990s	8.7	30.1	-21.4
2000s	48.1	57.7	-9.6
2010s	39.4	11.0	28.4
2020s	3.8	1.2	2.6
Journal tier (%)			
Top-5	90.4	21.2	69.2
Other-top	9.6	78.8	-69.2
Editor-name match quality (%)			
Exact	28.8	20.6	8.3
Fuzzy	0.0	1.8	-1.8
Ambiguous (top FC)	71.2	61.7	9.5
Unmatched	0.0	16.0	-16.0
Institution recorded (%)	100.0	13.2	86.8
Events (N)	104	326	

Notes: The table compares the 104 editor-turnover events that enter the staggered difference-in-differences estimation sample with the 326 that drop out, by appointment decade, journal tier (Top-5 versus other upper-tier), name-match quality, and whether a home institution is recorded. A recorded home institution is the binding panel requirement. Region is omitted: no institution-level country lookup is available. Read the large tier and institution-recorded gaps as the estimation sample being a selected, Top-5, institution-identified subset rather than a random draw. Source: editor roster from Brogaard, Engelberg, and Parsons (2014), journal records, GOELD, and the Internet Archive; OpenAlex.

Further displays are available on request and in the replication package: the coauthor ladder of conversion, hazard, and step-down rates, and the homophily and cross-institutional pairing series underlying Section 3.5.