FEDERAL RESERVE BANK of ST. LOUIS

ECONOMIC RESEARCH

FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

Specialization vs. Diversification: Research Strategy and Performance in Economics

Authors	Harris Dellas; Manu García; Carlos Garriga; Christian Zimmermann
Working Paper Number	2026-025A
Release Date	November 2026
Citable Link	https://doi.org/10.20955/wp.2026.025

Federal Reserve Bank of St. Louis, Research Division, P. O. Box 442, St. Louis, MO, 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

Specialization vs. Diversification: Research Strategy and Performance in Economics*

Harris Dellas,[†c,b,a] Manu García,[a] Carlos Garriga,[a]
and Christian Zimmermann[a]

[a] Federal Reserve Bank of St. Louis

[c] University of Surrey

[b] Karl Brunner Institute

November 4, 2026

Abstract

The tension between diversification and specialization is a central theme in economics, shaping discourse in fields as disparate as international trade, portfolio theory, and corporate scope. This paper examines this trade-off within the realm of research strategy. Using the RePEc panel database of approximately 60,000 economists, we demonstrate that the pursuit of a diversified research portfolio yields superior performance as measured by citations and publications as well as lower research attrition rates. This advantage applies to both senior and, in contrast to conventional wisdom, early-career researchers. This strategy is significantly facilitated by joint work with new co-authors.

JEL Classifications: D83, D85, J24, O31.
Keywords: Research strategy, diversification, specialization, publications, citations, co-authorship, research attrition

*We acknowledge the useful comments of Marios Angeletos, David Argente, Pedro Gomis-Porqueras, John Leahy, Francesco Lippi, Franck Portier, José-Victor Ríos-Rull, Robert Shimer, Harald Uhlig, Ludo Vischers. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

[†]Corresponding Author: Harris Dellas; E-mail: harris.dellas@gmail.com

1 Introduction

Research performance in economics, typically measured by citations, publications, and rank, is widely understood to be primarily a function of relevant innate ability, acquired skill, effort, and luck. However, it plausibly also depends on research strategy: the choice of co-authors, topic selection (variety vs niche), professional mobility, networking, and so on. Despite its potential importance, the question of whether and how research strategy matters for research performance remains unexplored.

This paper investigates a so far understudied strategic element, namely, researchers' choice of the degree of *output* diversification/specialization.[1] We study the effects of this choice on research performance utilizing the Research Papers in Economics (RePEc) database. Covering approximately 60,000 economists, RePEc provides data on authorship, publication year and venue, field classification of the paper's research topic, citations (quality-adjusted and unadjusted), author rankings, and various author characteristics (academic age, graduating university, affiliation, etc.). The field classification is done on the basis of codes from the *Journal of Economic Literature* (JEL) classification.[2] Each research item carries one or more JEL codes. A code contains a letter (indicating the primary research field) and may also contain one or two digits which represent the secondary and tertiary subdivisions of the primary field.

The fact that a researcher's body of work at any point in time is a collection (a "portfolio") of JEL classification codes can be exploited to construct a specialization-diversification index (SDI) of research output. Our index follows the Herfindahl-Hirschman Index (HHI), a common measure of market concentration. The SDI is essentially the HHI with the researcher's portfolio replacing the market and the individual JEL codes replacing the firms. So instead of the distribution of the market share of firms, we have the distribution of the portfolio's JEL codes. Moreover, the multi-layered structure of JEL codes (letter, 1-digit, and 2-digit subdivisions) makes it possible to consider alternative diversification "geographies" that differ in terms of the amount of local variation they include (such as variation in the JEL digits, in particular, the second one).

The proposed diversification index is novel and more granular than existing ones, which all rely on the concept of a researcher's "primary" field. Diversification in these alternative measures corresponds to the number or share of works outside one's "primary" field. In contrast, ours takes into account the full distribution of fields and sub-fields within a researcher's portfolio and moreover it does not require the identification of one's primary field.

[1] Note that there exists a different concept of specialization which concerns the input, that is, the "tools" or methods used, rather than the output of a researcher's production function; we are not aware of any data that could make that alternative concept operational.

[2] See `https://www.aeaweb.org/econlit/jelCodes.php`.

We find that research performance is positively and significantly related to the degree of diversification. This finding holds across cohorts, gender, and institutional "pedigree." Additional patterns obtained include: *Junior Scholars*: Contrary to the "specialize early for tenure" mantra, conditional on portfolio size, junior researchers with more diversified portfolios exhibit superior performance at the end of the "tenure window" (about five years from the time of the first publication) as well as 10 and 15 years from that date; *Research Survival*: The pursuit of higher diversification reduces the rate of research attrition; *Collaborative Engine*: Joint work contributes to performance not only by increasing the labor input and creating synergies (complementarities) among co-authors but also by facilitating diversification, of which it is a primary engine.

A challenge that permeates all work on the determinants of performance is the possible dependence of research strategy choices on other, unobserved determinants of performance, such as ability. In the context of this paper, if high-ability individuals favored diversification, say, because of preferences or a higher expected rate of return, not accounting for such dependence would produce an exaggerated estimate of the causal impact of diversification. We address this issue using a battery of empirical strategies. Namely, we include as controls in the regressions a proxy for ability; we consider groups that are likely to exhibit little ability heterogeneity, for instance, the set of individuals who have comparable numbers of top-five publications; we run panel regressions with fixed effects that account for constant unobserved individual ability; and also perform instrumental variable estimation. The estimated impact of diversification diminishes when controlling for ability but it always remains substantial and highly significant. Note that if diversification provides an essential channel for high ability researchers to fulfill their potential, successfully controlling for ability may lead to an understatement of its contribution. In any case, the empirical evidence seems to strongly suggest that diversification is not merely a byproduct of talent, but a potent strategy for enhancing research performance.

The success of the described empirical strategy becomes apparent when comparing the panel estimates with and without fixed effects in various groups in the population of researchers. For instance, in the full sample, the introduction of individual fixed effects reduces the estimated elasticity of performance with regard to diversification by 40%, to a value of -0.98. But in the panels of top 5 authors, not only fixed effects make little if any difference -especially for more narrow subsets- but also the estimated coefficients are almost identical to those –fixed effect corrected– in the full sample. This constitutes prima facie evident that the cross-sectional estimates in some ability corrected samples are free from omitted variable bias.

We complement the empirical analysis with a simple model of research strategy that is

estimated on the RePEc data using the Simulated Method of Moments (SMM). We use the estimated model to assess important aspects of the diversification-collaboration-performance nexus that cannot be tested empirically due to data limitations. While the model assumes that the pursuit of diversification enhances human capital, it also assumes that it varies with innate ability and luck, hence it cannot be used to compute its causal effect. But it can be productively used to shed light on mechanisms. For instance, the model implies that the impact of diversification on output does not arise, as one may have speculated, from a "learning-by-diversifying" channel (mastering new methods and tools, learning from co-authors, etc.). While the finding of "doing better by diversifying without getting better" could be due to model mis-specification, it is also consistent with the work of Simonton (1997) on the implications of output variety for success in the realm of creative arts. According to Simonton, a diversified portfolio of works contains more "lottery tickets" than a non-diversified one, which increases the probability of sustained success (producing multiple hits).

The model can help decompose the total effect of collaboration on research output. Collaboration matters through various channels: by facilitating diversification, by increasing labor input, by exploiting complementarities among co-authors, etc.[3] According to the model, shutting down collaboration's diversification channel reduces median lifetime research output in both the population of researchers and the top 5 group (by 16% and 18% respectively) ; shutting down the non-diversification channels has a modest impact in the population (a reduction of about 6%) but proves quite consequential for top 5 authors (a 42% reduction). According to the model, collaboration is more beneficial to high ability individuals but mostly for reason not related to the pursuit of output variety.

The remainder of this paper is organized as follows. Section 2 provides a summary of the existing literature in this area and clarifies its connection to our work. Section 3 describes the RePEc database utilized in our analysis and the SDI metric we use. The empirical analysis is presented in Section 4, followed by the model in Section 5. Finally, Section 6 offers concluding remarks. The appendix collects additional empirical results and the details of the model's calibration and estimation.

2 Related Literature

On the determinants of performance

Little research has been done on the determinants of the research performance (as mea-

[3]An important historical example is the Manhattan Project. Robert Oppenheimer instituted weekly colloquia that brought scientists from disparate branches together. This fostered an environment where, as Hans Bethe observed, a problem in one area would often intrigue a specialist from another, leading to "unexpected solutions."

sured by citations and publications) of economists. Laband and Piette (1994) and Brogaard et al. (2014) examine network proximity to editors, finding that it influences publications, and attribute this to information channels rather than favoritism. Conley and Önder (2014) investigate the effect of pedigree (the rank of the graduating department) on publication records. They report that graduating from a top-tier department is neither necessary nor sufficient for success; in fact, most graduates from even the highest-ranked departments produce little published research. In contrast, there exists a large literature that relates research output to career success (e.g., placement, tenure, and promotion) as a function of researcher characteristics. For instance, Sarsons et al. (2021) find that, for comparable portfolios, joint work carries a smaller positive impact on the probability of tenure for female authors of joint papers relative to males. Similarly, Stansbury and Rodriguez (2025) demonstrate that socioeconomic background significantly affects institutional placement and tenure achievement.

On research diversification

We are not aware of any research in economics specifically addressing research output diversification. Outside of economics, with the exception of Berg (2022) and Dubois et al. (2014), the existing literature focuses almost exclusively on *interdisciplinary work in the sciences*. We briefly review representative studies from this literature.

Abramo et al. (2017) document stylized facts about the publication patterns of science faculty in Italy. Using data on Italian science professors' publications from 2004 to 2008, they examine field-specific propensities to contribute across different disciplines. They find that chemists are the most likely to publish outside their primary field, while mathematicians are the least likely. Their measures of diversification include the number of publications outside one's main field (defined as the field containing the largest share of an author's publications) and the share of total publications accounted for by such cross-field work. In follow-up studies, Abramo et al. (2019) assess whether publishing outside one's main field is "citation-profitable" by comparing the impact of primary-field and non-primary-field publications. Finally, Abramo et al., 2019b, utilize a 2012–2016 dataset to demonstrate that diversification is primarily achieved through collaborative work, highlighting joint authorship as a key driver of output variety.

Kelchtermans et al. (2022) examine the nature of exploration beyond a researcher's primary field and the characteristics of those who attempt it. They find that entry into new fields tends to be "local," as long jumps are rare, and it is most prevalent among highly productive, senior researchers with expansive co-author networks. Dubois et al. (2014) provide a comprehensive analysis of mathematicians, emphasizing the nexus between mobility, co-authorship, and productivity. They find that research diversification (measured by the count of distinct fields in a portfolio) is a critical, positive determinant of long-term performance.

There is also a significant body of work in psychology that explores how diversified portfolios drive success in creative industries. Berg (2022) finds that creators who achieve initial "hits" while having a varied portfolio are more likely to generate subsequent successes. This aligns with the "Equal-Odds Rule" of Simonton (1997), which posits that creative quality is a stochastic function of quantity, so a more diversified portfolio contains more "lottery tickets".

On corporate scope

Diversification within corporations has been the subject of extensive theoretical work, especially in the context of strategic management. Production functions (technology), transaction costs, organizational frictions and market structure are the main determinants studied.

The empirical patterns of specialization-diversification in the corporate world are quite mixed. The tech world seems to be dominated by ever expanding, highly *diversified* conglomerates. McAfee and Brynjolfsson (2017) attribute this to the fact that once a firm builds a successful platform, the marginal cost of adding a new side or product to that platform is nearly zero (platform pivot). The research analog to this is a researcher mastering a particular tool (econometric, computational, ...) and applying it repeatedly to diverse topics, usually in joint work with researchers more familiar with those topics. But outside the digital world, Ekerdt and Wu (2024) find that in the industrial economy, single-industry (specialized) firms have become larger and more dominant during the last few decades.

Dominance create a presumption that the chosen diversification/specialization strategy is effective. But it does not by itself directly inform on the diversification/performance nexus. On this front, we are aware of only one piece of work, namely, Rumelt (1974), that does so by providing evidence that "related diversification" (diversifying into similar fields) outperforms both "single business" and "unrelated conglomerates."

3 Data and Measurement

We scrutinize the production of research in economics by utilizing RePEc (Research Papers in Economics), the primary bibliographic databank of the economics profession. While RePEc has served for decades as the global repository for impact metrics and in particular rankings, its potential to reveal the underlying mechanics of how economic research is actually produced has remained largely untapped. By integrating publication histories, field classifications, and co-authorship records with citation networks for nearly 60,000 economists over half a century, we assemble a uniquely comprehensive dataset that tracks researchers throughout their careers and allows us to study the "economics of doing economics."

3.1 The RePEc Database

RePEc aggregates metadata contributed by over 2,000 publishers. The database indexes 4.9 million research works from 4,100 journals and 5,500 working paper series, along with 64,000 books and 230,000 book chapters. Coverage is comprehensive: any journal or working paper series of consequence in economics is represented. Each indexed item is linked to authors, citations, publication dates, and subject classifications.

The RePEc Author Service provides the foundation for individual-level analysis. Approximately 68,000 researchers have registered, receiving unique identifiers that link to their complete publication histories. The system aggregates across versions of the same work -—a journal article that was preceded by one or more working papers appears as a single item with its full history-— avoiding the double-counting that plagues bibliometric studies relying on raw publication records. We compute academic age from the year of the appearance of an author's first publication, whatever its form. Graduation year, available through the RePEc Genealogy project for a subset of authors, is used only to construct pedigree, not to date career length.

For each research item, RePEc records the year of publication and its JEL (Journal of Economic Literature) classification codes. The JEL system provides a standardized, hierarchical taxonomy of research fields in economics. There are 20 primary categories (each identified by a single letter), 141 secondary categories (letter plus one digit), and 867 tertiary categories (letter plus two digits). Over one million items carry JEL codes, averaging three codes per item, which is roughly a third of all indexed works: a researcher's portfolio is therefore measured on the coded subset of her output. Because authors assign these codes themselves, they reflect their judgment about content rather than algorithmic inference or editorial categorization. This self-classification is essential: it captures what researchers believe their work is about, measured consistently across the discipline and over time.

The CitEc project within RePEc extracts references from indexed works and matches them to other items in the database, generating citation links for approximately two million research items. CitEc computes recursive impact factors for journals and serials using a methodology similar to Google's PageRank algorithm: citations are weighted by the quality of citing sources, creating a measure of influence that accounts for both quantity and provenance. Self-citations are excluded. These impact factors allow for the construction of quality-adjusted output measures that distinguish a highly cited paper in a leading journal from one with similar raw citation counts in an obscure outlet.

We restrict our sample to economists with at least one publication and valid JEL code information, and to those whose index lies strictly between zero and one. The second restriction excludes 1,559 researchers whose entire recorded output carries a single JEL code, for whom

the index equals one and its logarithm is undefined; they are the most specialized researchers in the database. Their exclusion means the elasticity we estimate is identified off researchers with at least two codes, so it should not be extrapolated to complete specialization; it does not affect the comparisons across levels of diversification within the retained sample, which is where all of our estimates lie. This yields 58,479 researchers whose first publications span 1951 to 2025, with 99% falling after 1970, effectively the population of publishing economists over the past five decades.

The JEL classification system provides a common language for measuring research portfolios. Unlike text-based topic models, which require arbitrary choices about the number of topics and may conflate methodology with subject matter, JEL codes reflect authors' own characterization of what their work contributes. Unlike citation-based field definitions, which shift as literatures evolve and merge, the JEL taxonomy offers a stable benchmark for tracking portfolio composition over decades, though not all codes were present from the start of the period under study. And unlike ad hoc categorizations based on publication venue, JEL codes capture research content independently of where it appears. The hierarchical structure from broad primary fields down to specific three-digit topics allows measurement at multiple levels of aggregation, from diversification across the major areas of economics to fine-grained variation in research focus within subfields.

The longitudinal dimension is quite long. We observe researchers from first publication through the end of their careers, with a median span of 16 years and many extending well past 30. This contrasts with the cross-sectional snapshots or short panels typical of bibliometric research, which cannot distinguish life-cycle dynamics from cohort effects. The career-length records permit within-researcher analysis: we can ask not only whether diversified researchers are more productive but also whether the timing of diversification over the career matters for its rate of return. The panel structure also allows for the identification of persistence in portfolio composition which may be related to the existence of portfolio adjustment costs.

The co-authorship data is comprehensive. RePEc records all authors on each work and links them to their own registered profiles when available. We observe not just whether researchers collaborate, but with whom, how frequently, and how these patterns evolve alongside changes in research focus. This network structure covering hundreds of thousands of collaborative relationships enables the study of the complementarity between diversification and co-authorship that our model emphasizes.

Finally, the RePEc Genealogy project provides institutional linkages for a substantial subset of authors, recording doctoral institutions and graduation years. It is a wiki, so records are contributed and amended by RePEc account holders rather than centrally verified. As RePEc provides a ranking of universities, this can help construct a measure of the researcher's

academic pedigree.

3.2 Measuring Research Diversification

Previous work on research portfolios has relied primarily on simple counts of fields (see, for instance, Dubois et al., 2012), treating a researcher with publications in ten fields as more diversified than one with publications in five. This approach has an obvious limitation: it ignores the distribution of effort across fields. An economist with 95 percent of output in macroeconomics and token contributions to nine other areas is not meaningfully diversified, yet a simple count would suggest otherwise.

We measure portfolio concentration using the full distribution of publications across fields. The Specialization-Diversification Index (SDI) adapts the Herfindahl-Hirschman Index to research portfolios:

$$\text{SDI}_i = \sum_{j=1}^{J} s_{ij}^2 = \sum_{j=1}^{J} \left(\frac{n_{ij}}{N_i} \right)^2 \tag{1}$$

where N_i is the sum of all JEL codes appearing in researcher i's portfolio of works, counting repetitions, and n_{ij} is the number of appearances of JEL code j. The SDI ranges from values approaching zero, when output is spread uniformly across many fields, to one, when all output concentrates in a single field. A researcher with equal output in k fields has SDI $= 1/k$; one whose work concentrates in a single category despite scattered appearances elsewhere will have a large value of SDI (low diversification).

We use the inverse of the SDI to measure the effective number of fields. The model in Section 5 works directly with this inverse, $n = 1/\text{SDI}$, so a higher n means a broader portfolio while a higher SDI means a more concentrated one. An SDI of 0.10 implies an effective portfolio breadth of 10 fields; an SDI of 0.25 implies 4 fields. This transformation makes magnitudes easier to interpret: when we find that a researcher's SDI falls from 0.20 to 0.10 over a career, we can claim that the effective breadth has doubled from 5 to 10 fields.

The hierarchical structure of the JEL system allows for the construction of diversification measures at different levels of aggregation. Our baseline SDI uses the complete three-character codes (letter plus two digits), capturing fine distinctions in research focus, for example the difference between E31 (Price Level; Inflation; Deflation) and E32 (Business Fluctuations; Cycles). We also construct broader measures using only the primary letter (20 categories) or the letter plus first digit (141 categories). The former captures diversification across the major areas of economics, microeconomics versus macroeconomics, agricultural economics versus economic history. For robustness, we additionally employ an SDI based on the NEP (New Economic Papers) classification, which uses volunteer editors to assign working papers to

approximately 100 fields, some introduced later than others and overlapping by construction, as well as the simple count of distinct JEL codes appearing in a researcher's portfolio.

3.3 Measuring Research Output

RePEc combines 32 criteria to produce composite scores and rankings for each registered author. The criteria span citations, publication outlets, centrality in coauthorship networks, traffic to abstract pages, downloads, page counts, breadth of citations across fields, and the publication success of doctoral students. Authors are ranked according to each criterion, and a harmonic mean of ranks determines overall standing, with the two best and two worst criteria dropped to reduce sensitivity to outliers. The rankings are internally consistent as the same source is used for both the journal impact factors and the author rankings themselves.

For our analysis, we focus mostly on measures of publications and citations. We use quality-adjusted pages which weigh journal article pages by the recursive impact factor of the publishing journal, capturing both the quantity of output and its market value (the selectivity of the outlet). And quality-adjusted citations which are constructed analogously but weight each citation by the recursive impact factor of the *citing* publication rather than of the publishing journal. In addition we consider the number of NEP fields from which an author's work has received citations, a measure of research reach across fields. And we also use the number of publications in the top five journals[4] (*American Economic Review*, *Econometrica*, *Journal of Political Economy*, *Quarterly Journal of Economics*, and *Review of Economic Studies*). The AER publications do not include the *AER Papers and Proceedings*.

3.4 Variable Construction and Data Assembly

To characterize the production function of economic research, we supplement the core bibliographic records with several researcher attributes. We define *academic age* as the years that have elapsed since an author's first recorded publication. We measure *pedigree*, a proxy for ability, as the RePEc ranking of one's doctoral program (this information exists for roughly 25% of researchers through the RePEc Genealogy project). *Gender* is imputed using the NamSor algorithm, which leverages name-origin patterns to provide high-confidence classification for international scholar populations. We also compute the variable HHI^{co_auth} which measures the diversification of one's set of co-authors and use it as an instrument for own diversification in the panel regressions.

[4]Our main use of top five publications, though, is not on the output but on the input side of the production function, namely, as a proxy for unobserved ability.

Our primary measure of collaboration is the total count of *distinct* RePEc-registered coauthors up to any given date, which RePEc reports including the author herself, so a researcher who has never collaborated takes the value one. We also consider the average number of authors per paper. New co-authors may facilitate the pursuit of diversified research while a larger number of co-authors per paper may not necessarily matter for diversification if it involves the same set of co-authors.

3.5 Summary Statistics: The Landscape of Economic Research

Table 1 characterizes the research trajectories of nearly 60,000 economists. The average career in our sample spans 18 years and yields approximately 26 total publications (median 13), though the distribution is heavily right-skewed; the top decile of researchers has produced more than 60 items, reflecting the high-productivity "superstar" phenomenon common in knowledge-intensive industries (Azoulay et al., 2010).

Our primary object of interest, the Specialization-Diversification Index (SDI), reveals substantial heterogeneity. Using the full three-character code, the mean SDI is 0.15 (median 0.10), suggesting that the typical economist maintains an effective portfolio breadth of approximately 7 to 10 distinct research fields. However, the dispersion is large: researchers in the 10th percentile (SDI = 0.04) spread effort across 25 effective fields, whereas those in the 90th percentile (SDI = 0.33) concentrate their output in roughly 3 fields.

Two features of the table bear on the estimation samples that follow. About a tenth of researchers record zero quality-adjusted pages or citations, so the tenth percentile of both measures is effectively zero; these observations cannot enter specifications in logs, which is the main reason the regression samples below are smaller than 58,479. And the number of distinct JEL codes averages 20 while the index implies an effective breadth of under seven fields. That gap is precisely the point made above: a simple count treats a token contribution as a full field, whereas the index weights each by the effort devoted to it.

Table 1: Summary Statistics

Variable	Mean	SD	P10	Median	P90	N
Panel A: Research Output						
Total Publications	25.80	38.78	2	13	62	58,479
Journal Articles	15.11	24.46	1	7	37	58,479
Quality-Adjusted Pages	113.46	341.39	0.00	10.62	282.10	58,479
Quality-Adjusted Citations	158.08	758.40	0.01	8.92	264.16	58,479
Panel B: Portfolio Concentration						
SDI (letter + 2 digits)	0.15	0.12	0.04	0.10	0.33	58,479
SDI_L1 (letter + 1 digit)	0.23	0.17	0.08	0.18	0.50	58,479
SDI_L (letter only)	0.39	0.21	0.18	0.33	0.64	58,479
Distinct JEL Codes	20.00	21.18	3	13	45	58,479
Panel C: Career and Collaboration						
Career Length (years)	17.63	10.98	6	16	33	58,478
Registered co-authors (incl. self)	9.28	12.90	1	5	21	58,479
Authors per Paper	3.29	1.72	1.56	3.00	5.29	58,479

Notes: Sample includes all RePEc-registered economists with at least one publication and valid JEL codes, 1970–2023. Quality-adjusted measures weigh by recursive journal impact factors. SDI is the Herfindahl-Hirschman Index computed over the distribution of publications across JEL codes.

As field definitions coarsen from the full code to the letter alone, the mean SDI increases from 0.15 to 0.39. This suggests that economists do not often venture far away from their favorite fields. Finally, the collaboration data shows a discipline characterized by teamwork: the average researcher works with eight distinct collaborators over their career, with a typical paper involving 3.3 authors. This collaborative intensity is a key margin we explore as a complement to research diversification.

4 Empirical Analysis

4.1 An Overview

We first run cross-sectional OLS regressions of academic performance on diversification:

$$y_i = \alpha + \beta \cdot \log(\text{SDI}_i) + \mathbf{X}_i'\gamma + \epsilon_i \tag{2}$$

where y_i is a measure of performance (log quality adjusted citations/pages or number of NEP fields citing the author), SDI_i is research output specialization-diversification, and $\mathbf{X}_i$ is a vector that contains two sets of variables: exogenous *controls* such as career age and cohort fixed effects; and *proxies* for unobserved variables that matter for both the diversification choice and performance (such as the number of publications and the impact factor of the first published article).

The coefficient of interest, β, captures the association between specialization and performance. A negative β indicates that more diversified researchers (lower SDI) perform better.

Identifying the Causal Effect: Challenges The identification of the causal effect of diversification on research performance faces two primary sources of bias. a) Omitted (ability) variable bias: If higher ability researchers, who naturally perform better, also favor diversification due to personal preference or higher returns,[5] and this dependence is not adequately controlled for, the OLS estimate of β will overstate the causal impact. b) Reverse causality: If the pursuit of diversification actually enhances ability (e.g., through learning), then inclusion of ability proxies in the regression will lead to an OLS estimate of β that understates the true causal effect. We mostly focus on the former (type I error).

As ability is unobserved, we employ three empirical strategies to minimize the first bias. The first includes *proxies* for an individual's initial and current ability level as independent variables: the impact factor of the journal of a researcher's first publication is an example of the former; academic age as well as the total number of one' publications (research items) at a point in time is an example of the latter. Similarly, the use of individual *fixed effects* in panel regressions helps eliminate the bias that arises from time invariant omitted research characteristics (but it is not effective when the omitted variables vary stochastically over time).

The second involves a reduction in sample heterogeneity. We focus on data subsets characterized by allegedly less heterogeneity in ability levels. For instance, we group researchers by academic "pedigree" (e.g., the rank of the doctoral institution) or by performance thresholds, such as a specific number of "top five" publications or placement within a certain percentile of the performance distribution, etc.

The third approach deploys an instrument for the researcher's diversification index. We have experimented, mostly in the context of our panel estimation, with various candidate instruments: Lagged values of the individual's own diversification; the combined diversification

[5]Note that at low numbers of publications, adding another publication is likely to mechanically increase the SDI if the choice of the JEL codes represents a random draw from a uniform distribution. This effect dissipates, though, quickly with the increase in the number of items. And there is a presumption that the SDI may actually decrease with additional publications if the researcher favors doing research in her "primary" area of expertise.

indexes of one's *co-authors*; the diversification index of the researcher's *doctoral department* (calculated from the JEL codes of all faculty members); and finally, an individual-specific index derived from the diversification profile of the journals where the researcher has published. Each journal's diversification is computed from the JEL codes distribution of all its published papers, then weighted by the researcher's publication share. We discuss the underlying motivation as well as the respective pros and cons of these alternative instruments in Section 4.2.2.

4.2 Empirical Results

Table 2 presents the OLS estimates of Equation (2), with age, age squared, the total number of items (works) and, in column (4), the RePEc recursive impact factor of the venue of one's first publication included as controls. Items serves as a crude proxy of *lifetime* ability *and* effort; and the first-publication impact factor proxies *initial* ability. Adding a cohort effect makes no difference. Columns (1)–(4) report the effect of diversification on quality-adjusted pages; columns (5)–(6) on quality-adjusted citations; and column (7) on the number of NEP fields citing the author.

Greater diversification is associated with better performance. The results are statistically and economically significant: a 10% increase in diversification (a 10% reduction in SDI), is associated with about a 10% increase in quality-adjusted pages or citations.

To assess robustness to omitted variable bias, we report sensitivity bounds for columns (2) and (3). For column (2), the Oster (2019) bounds under proportional selection ($\delta = 1$) yield an identified set of $[-1.11, -0.97]$. A value of 7.80 for δ^* indicates that unobservable confounders would need to be nearly 8 times as important as the observable controls (age) to drive the coefficient to zero, a high bar, which suggests that the reported relationship is robust. For columns (3) and (4), where the control variables *items* and the first-publication impact factor are not exogenous, we report the Diegert et al. (2022) bounds that relax the exogenous controls assumption. The results again indicate a robust causal relationship from diversification to performance.

Table 2: Diversification and Performance

	(1)	(2)	(3)	(4)	(5)	(6)
	log(Quality-Adjusted Pages)			log(Quality-Adjusted Citations)		
log(SDI)	-1.308^{***}	-1.013^{***}	-0.786^{***}	-1.329^{***}	-0.974^{***}	-0.727^{***}
	(0.014)	(0.011)	(0.014)	(0.014)	(0.011)	(0.014)
Age		0.1017^{***}	0.0920^{***}		0.1610^{***}	0.1505^{***}
		(0.0023)	(0.0024)		(0.0025)	(0.0026)
Age^2		-0.000821^{***}	-0.001004^{***}		-0.001632^{***}	-0.001831^{***}
		(0.000043)	(0.000045)		(0.000049)	(0.000049)
log(firstpub IF)		0.6043^{***}	0.6151^{***}		0.5322^{***}	0.5440^{***}
		(0.0040)	(0.0039)		(0.0040)	(0.0039)
Items			0.0113^{***}			0.0123^{***}
			(0.0005)			(0.0006)
R^2	0.163	0.548	0.573	0.164	0.542	0.571
N	46,156	46,156	46,156	46,156	46,156	46,156

Notes: Robust standard errors in parentheses. firstpub IF is the RePEc recursive impact factor of the venue of the researcher's first publication. All columns use the same sample. Oster (2019) bounds for the age-controlled specification, assuming proportional selection ($\delta = 1$) and $R_{max} = 1.3\tilde{R}$: identified set $[-1.108, -0.966]$, $\delta^* = 7.80$. Diegert, Masten and Poirier (2022) bounds: [pending]. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

How are the results affected when local variation (variation in the 2nd or 3rd element of the JEL code) is ignored? Table 3 reports the results from a regression of performance on SDI_L (letter level variation) and SDI_L1 (letter plus one digit) in addition to the baseline SDI measure. For completeness, we use the RePEc rank of the author as the performance measure, of quality adjusted publications or citations produces identical results. It can be seen that the impact of diversification on performance is decreasing in its "ambition": the pursuit of research in unrelated or distant fields (represented by different JEL letters) impacts positively on performance but the estimated elasticity indicates a smaller effect than that of more narrow diversification. That is, for the average researcher, "visiting" nearby fields has a higher payoff than venturing across the spectrum of fields. This result seems to be in line with a finding in industrial organization by Rumelt (1974), who reports that "related diversification" (diversifying into similar fields) in business activity outperforms both "single business" and "unrelated conglomerates."

Table 3: Non-local Diversification and Performance

	(1) SDI	(2) SDI_L1	(3) SDI_L	(4) NEP SDI
log(SDI)	-0.3426^{***}	-0.2535^{***}	-0.1482^{***}	-0.4785^{***}
	(0.0066)	(0.0065)	(0.0070)	(0.0077)
Age	0.0247^{***}	0.0315^{***}	0.0362^{***}	0.0335^{***}
	(0.0020)	(0.0021)	(0.0022)	(0.0025)
Age^2	-0.0001	-0.0002^{***}	-0.0003^{***}	-0.0002^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Items	0.0121^{***}	0.0137^{***}	0.0147^{***}	0.0114^{***}
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.646	0.615	0.596	0.642
N	58,478	58,478	58,478	45,050

Notes: The dependent variable is the log of the RePEc composite author score. SDI_L, letter only; SDI_L1, letter and one digit; SDI, letter and two digit; NEP SDI, New Economic Papers based diversification index. NEP sample is smaller due to incomplete coverage. All specifications include cohort fixed effects. Robust standard errors in parentheses. $^{***}p < 0.01$.

We have also examined the relation between diversification and performance in subgroups corresponding to gender and "pedigree" (the rank of the graduating department). Conditional on academic age, the number of items, the size of the co-author network and cohort effects, the estimated elasticity is somewhat larger for females than for males, -0.54 against -0.42. Differences across pedigree are smaller and not monotone in rank: -0.39 for graduates of top ten programs, -0.41 for ranks eleven to fifty, and -0.32 for those outside the top fifty. Table 19 in the appendix reports these regressions.

4.2.1 Addressing Bias

"Neutralizing" ability We saw in Table 2 that inclusion of proxies for ability, such as academic age, items and firstpub makes diversification less consequential (from -1.31 to -0.79) while also substantially increasing the explanatory power of the regression. Heterogeneity in ability and thus the size of potential bias can be further reduced by working with subjects that are likely to be comparable in terms of ability. We speculate that this may be the case for people who have a certain publication record in top five journals, e.g., 3-5, 5-10, ≥ 10 and so on. Table 4 reports the unconditional distribution of researchers with top five publications.

Note that about one in ten in the population of researchers has a top 5 and only about 3% have more than 3 top 5s!

Table 4: Full Distribution of Researchers with Top-5 Publications

# Top 5	Count	Cum. %	# Top 5	Count	Cum. %
1	2,924	44.16	21	16	98.29
2	1,202	62.31	22	13	98.49
3	642	72.00	23	9	98.63
4	363	77.48	24	11	98.79
5	299	82.00	25	11	98.96
6	241	85.64	26	4	99.02
7	171	88.22	27	6	99.11
8	111	89.90	28	2	99.14
9	102	91.44	29	3	99.18
10	80	92.65	30	8	99.31
11	58	93.52	31	5	99.38
12	66	94.52	32	2	99.41
13	43	95.17	33	4	99.47
14	39	95.76	34	5	99.55
15	35	96.29	35	3	99.59
16	29	96.72	36	1	99.61
17	30	97.18	37	2	99.64
18	27	97.58	38	4	99.70
19	18	97.86	39	3	99.74
20	13	98.05	≥ 40	17	100.00

Table 5 reports results for groups of authors who have at least one top five publication (column 1), two to five (column 2), five to ten (column 3) and more than ten (column 4) publications. We consider ranges rather than a specific number to reduce the effect of a lucky draw and to preserve sample size. We have also included academic age and items (within each of the binned columns the correlation between age and the number of top 5s is small, between 0.09 and 0.21, and between 0.03 and 0.13 net of items; see Table 13 in the appendix). Diversification matters. Moreover, the estimated effect is comparable across the subgroups of top 5s in the Table, and much smaller than that in the general population of authors. The similarity suggests that the lower values of the coefficient are not due to diversification becoming less consequential as the level of ability increases[6] but rather to the reduction of omitted ability bias that comes from the suppression of ability heterogeneity. We also think that while the estimation of the effect in the 2-5 group may still contain some ability bias

[6]Note that such an interpretation would also be contradicted by Figure A.2 in the appendix, which suggests that if anything performance is a convex function of diversification.

as luck maybe a factor in producing a couple of top 5s, the estimated coefficient in the 10+ group likely *understates* the true effect because of reverse causality: it seems plausible that undertaking research in diverse areas is *necessary* for producing a large number of top 5 publications. The possibility of publishing in leading journals numerous papers that all fall in a narrow area must be limited.

Table 5: Diversification and Performance: Researchers with Top-5 Publications

	(1) ≥ 1 Top 5	(2) 2–5 Top 5	(3) 5–10 Top 5	(4) > 10 Top 5
log(SDI)	-0.427^{***}	-0.231^{***}	-0.132^{***}	-0.205^{***}
	(0.018)	(0.022)	(0.022)	(0.036)
Age	0.0336^{***}	0.0374^{***}	0.0338^{***}	0.0248
	(0.0078)	(0.0098)	(0.0122)	(0.0191)
Age^2	-0.00047^{***}	-0.00068^{***}	-0.00059^{***}	-0.00055^{**}
	(0.00013)	(0.00016)	(0.00018)	(0.00024)
Items	0.0056^{***}	0.0040^{***}	0.0034^{***}	0.0036^{***}
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.323	0.262	0.411	0.512
N	6,647	2,511	1,006	487

Notes: Dependent variable is log(Quality-Adjusted Pages). Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

We have also used other partitions of the data to reduce ability heterogeneity. For instance, we have considered individuals who belong to particular percentiles of the performance distribution (say, top 10% for citations) and have a certain pedigree (say, graduates of a top 10 program). The results are very similar to those reported in Table 5, both in terms of the estimated impact of diversification and its explanatory power. Table 14 in the appendix reports these results.

Panels In panels, the use of individual fixed effects provides an effective way of eliminating bias arising from the possible dependence of the diversification choice on time invariant, unobserved researcher characteristics that also matter for performance (such as exogenous, inherent ability). We estimate the following equations

$$\log(y_{it}) = \beta \cdot \log(\text{SDI}_{it}) + \alpha_i + \delta_t + \epsilon_{it} \tag{3}$$

where y_{it} is researcher i's performance (quality-adjusted citations or pages) in year t, SDI_{it} is her diversification SDI index in t α_i is an individual fixed effect, and δ_t is a year fixed effect.

Table 6 reports results for citations; Table 15 in the appendix reports the same specifications for pages. Three groups are considered: the full set of RePEc authors; those with at least one top five publication; and those with 3-7 top 5 (the results are similar for other partitions of the authors with top 5 publications).

The table reveals important information about potential ability omission bias. In particular, notice that while FE estimation makes a big difference -relative to OLS- in the full sample, an indication of omitted variable bias in that sample, it makes very little difference in the full *top 5* sample; and no difference whatever in the top 5 subsample of 3 to 7 publications. It seems that a strategy of focusing on samples that suppress ability heterogeneity can be as effective in containing omitted ability bias as using FE (perhaps at the expense of neglecting the influence of diversification on human capital accummulation). This finding strongly suggests that the results in the cross section regressions with subsets of top 5 publishing authors (Table 5) or other more homogeneous partitions of the data (Table 14 in the appendix) most likely capture the direct, causal impact of diversification on performance.

Another important feature of Table 6 is that the estimated coefficient of SDI in the regressions that include FE is the same in all three samples, in spite of the fact that ability differs greatly across them. This indicates that the return to the pursuit of diversification is comparable across different ability levels, that is, a diversification strategy is equally beneficial to low and high ability researchers.

Table 6: Panel Estimation: Diversification and Citations

	Full Sample		Top 5 Sample		3–7 Top 5	
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	FE	OLS	FE	OLS	FE
log(SDI)	-1.606^{***}	-0.982^{***}	-1.114^{***}	-0.958^{***}	-0.949^{***}	-0.954^{***}
	(0.004)	(0.009)	(0.005)	(0.019)	(0.008)	(0.035)
Individual FE	No	Yes	No	Yes	No	Yes
Year FE	No	Yes	No	Yes	No	Yes
R^2	0.150	0.950	0.222	0.865	0.237	0.816
R^2 Within		0.157		0.188		0.203
N	863,018	863,018	151,944	151,944	42,368	42,368

Notes: The dependent variable is log(Quality-Adjusted Citations). Standard errors clustered at individual level. $^{***}p < 0.01$. FE columns include individual and year fixed effects.

4.2.2 Addressing Bias: IV Estimation

Panel, fixed effects estimation helps eliminate the bias arising from the existence of time-invariant unobserved heterogeneity affecting both performance and diversification. But it may be less effective in the presence of unobserved heterogeneity in ability that is not driven by the diversification choice and varies over time in a way that is not accounted for by the time effects variable.

Identifying an instrument for diversification that does not depend on a researcher's ability represents a major challenge. We have considered a number of candidates, some of which are only applicable in the cross section regressions and some that can also be used in the panels. The results from these regressions are reported in Appendix A.6.

Use of lagged values, say in $t-k$, of the dependent variable is a commonly use instrument. The problem is that it may eliminate the effect of shocks to ability between t and k but is less effective when shocks to ability have persistent effects.

A more promising alternative is to use a weighed average of the SDI of *one's co-authors*. This can also be used in the panel regressions but may not be exogenous enough to the unobserved drivers of one's performance. The problem arises from the possibility that a positive shock to one's ability may lead him to work with more able and thus more diversified people. So an increase in the SDI of the co-authors is related to one's increased ability.

The results from using these two instruments are reported in Table 18 in the appendix.

Another possible instrument is the diversification index of the department that granted the doctoral degree. We construct it by pooling all the publications of the extant faculty to create a portfolio of department-wide JEL codes.[7] Nevertheless, this instrument may fail the exogeneity requirement. If the higher ranked departments have more able faculty and students, and if ability drives diversification, then replacing own SDI with that of the graduating school does not suffice to eliminate the endogeneity problem.

One final candidate is an individual-specific variable that relies on the diversification indexes of the journals where the researcher has published. The diversification index of journal x can be constructed from the distribution of JEL codes it contains. And these indexes can be used in one's portfolio of works, with weights corresponding to the portfolio shares of works in the particular journals. We have constructed a time series of this variable for each individual from 2000 to the present on the basis of all extant journals that report JEL codes to RePEc.This variable predicts own diversification strongly but does not satisfy the exclusion

[7]This is available for use only in the cross section regressions. The reason that this is that RePEc only contains the most recent list of faculty rather than that its time evolution, including the one at the time of one's doctoral studies. An additional problem is that this measure of diversification is common to all graduates of that particular program. As a few universities dominate production of PhD degrees, this limits variation. Note that this variable could be part of a graduating university FE in the panel regressions.

restriction. Because our performance measures weight output by the recursive impact factor of the publishing journal, the instrument and the outcome are built from the same set of journals, so publishing in a broad, high-impact outlet raises both directly, without passing through own diversification. The journal mix is in any case a consequence of the diversification choice rather than a shifter of it, and individual fixed effects do not repair either problem. We therefore do not use this variable as an instrument.

In view of the preceding discussion of the issues plaguing the instruments available, and in light of the findings reported in Table 6, we believe that the simple strategy of "neutralization" of ability through the judicious partitioning of the population of researchers represents the most compelling method and delivers the most reliable results regarding the causal effects of diversification on performance.

4.3 Early diversification

4.3.1 Early career diversification and performance

For young researchers, research strategy may play an important role in surviving the tenure process. The advice they often get from senior colleagues is to avoid spreading their interests. The conventional wisdom seems to be that specialization (establishing a "niche") early on helps gain recognition and supports stronger tenure letters.[8] While plausible, this thinking ignores the potential benefits that the pursuit of research diversification variety may bring.

In order to address this issue we compute a researcher's SDI at the end of the five year period following the year of the appearance of one's first work (using 6 makes no difference to the results). We then relate this SDI to cumulative research performance at the end of that five year period (short-run relationship) and at the end of 10 (medium-term) and 15 years (long-term). We also include the cumulative number of items (papers produced) during these five years as a proxy for effort and ability during this time interval, and in columns (4)–(6) the impact factor of the venue of the first publication as a proxy for initial ability. Table 7 shows that for a given number of research items, a more diverse portfolio at the end of this early window is associated with better performance in both the short and longer term, with the footprint being larger in the medium-longer term. Adding the initial-ability proxy raises the R^2 substantially but leaves the estimated effect of early diversification close to unchanged, by 13% at the five-year horizon and by 5% or less at fifteen years.

[8] "Have a theme: You should think of your research as writing a book, not articles. The idea is that your first five years' research should involve one or two unifying themes. The reason for sticking to a theme is that people evaluating you need to be able to say "This young scholar is one of the world's experts at ___." You get to fill the blank." David Levine's Cheap Advice for New Faculty, `https://faculty.haas.berkeley.edu/levine/cheap_advice.html#dissertation`.

Table 7: Early-Career Diversification and Short and Long Run Performance: 2000–2005 Cohort

	(1) Pages at 5 yrs	(2) Pages at 15 yrs	(3) Cites at 15 yrs	(4) Pages at 5 yrs	(5) Pages at 15 yrs	(6) Cites at 15 yrs
log(SDI 5yr)	-0.2096^{***}	-0.4947^{***}	-0.6808^{***}	-0.1814^{***}	-0.4711^{***}	-0.6554^{***}
	(0.0547)	(0.0512)	(0.0608)	(0.0404)	(0.0405)	(0.0501)
Items 5yr	0.0524^{***}	0.0455^{***}	0.0616^{***}	0.0486^{***}	0.0423^{***}	0.0582^{***}
	(0.0051)	(0.0047)	(0.0059)	(0.0041)	(0.0039)	(0.0050)
log(firstpub IF)				0.7931^{***}	0.6636^{***}	0.7143^{***}
				(0.0098)	(0.0100)	(0.0123)
R^2	0.095	0.113	0.147	0.552	0.463	0.432
N	8,082	8,082	8,082	8,082	8,082	8,082

Notes: Sample restricted to the 2000–2005 cohort. Dependent variables are log quality-adjusted pages or citations cumulated to the stated horizon. log(SDI 5yr) is the log diversification index at the end of the fifth year and Items 5yr the number of publications by that point. Columns (4)–(6) add the recursive impact factor of the venue of the first publication as a proxy for initial ability. All six columns use the same sample, which is ten researchers smaller than in earlier drafts because that variable is missing for them. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.3.2 Early Diversification and Attrition

The cessation of research activity (attrition) is common across all academic ages. It occurs for a number of reasons, including the revelation of low ability and changes in preferences or opportunities. The question we address here is whether lack of variety can trigger attrition too. This could arise either because diversification enhances ability and performance or because the pursuit of output variety brings satisfaction and keeps the research flame alive.

Figure A.3 in the appendix shows the patterns of research attrition for cohorts of males and females who first published in 1995 and later. For each cohort (according to the year of the appearance of the first paper), it depicts the number of those who produced no further papers after five, ten and 15 years.[9]

We run probit regressions to compute the probability of attrition. An individual is assigned the value of 1 if attrition (no further publications) has occurred and zero otherwise.[10] The RHS variables are the logarithm of the SDI variable at the 5 years mark, SDI_5; and the number of items after 5 years. The general specification is:

[9]Dubois et al., 2013, also report substantial attrition for academic mathematicians at all ages.

[10]Using a weaker criterion, such as fewer than two or fewer than three further publications, makes no difference to the properties of the results.

$$P(attr_t = 1) = F(\beta_0 + \beta_1 \text{SDI_5} + \beta_2 items_5 + \varepsilon_i),$$

where $F(\cdot)$ is the standard normal cumulative distribution function and $t = 5, 10, 15$.

Table 8 reveals that early diversification is associated with a lower probability of attrition. To get a sense of the quantitative effect consider two individuals, both with five publications at the end of the five year window, one with six distinct JEL codes and another with eight. For the 2000–2005 cohort, the latter's probability of remaining research active is 1.2 percentage points higher after five years and 2.8 points higher after fifteen.

Table 8: Specialization and Attrition (Probit)

	Full Sample			2000–2005 Cohort		
	(1) Exit 5yr	(2) Exit 10yr	(3) Exit 15yr	(4) Exit 5yr	(5) Exit 10yr	(6) Exit 15yr
log(SDI 5yr)	0.157***	0.118***	0.018	0.207***	0.213***	0.241***
	(0.025)	(0.018)	(0.015)	(0.047)	(0.032)	(0.029)
Items (5yr)	−0.0454***	−0.0262***	−0.0198***	−0.0558***	−0.0340***	−0.0177***
	(0.0034)	(0.0019)	(0.0012)	(0.0064)	(0.0031)	(0.0026)
Pseudo R^2	0.101	0.061	0.046	0.097	0.063	0.037
N	24,225	24,225	24,225	7,700	7,700	7,700

Notes: Dependent variable (cessation of research) is unity if the individual has produced no further papers 5, 10, or 15 years since the first publication; and zero otherwise. log(SDI 5yr) is the log diversification index at the of five years from the first publication. Items (5yr) is the total number of publications by the end of the fifth year. Columns (1)–(3) concern the full sample while columns (4)–(6) concern the 2000–2005 cohorts. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.4 Joint work, diversification and performance

Joint work plays an increasing role in research production and nowadays represents dominant practice in economics. For instance, Card and DellaVigna (2013) find that the number of authors per paper in economics increased from 1.3 in 1970 to 2.3 in 2012. More importantly from the point of view of the theme of this paper, joint work is closely connected to both diversification and performance, with all three being strongly, positively correlated. As discussed in the literature review, Abramo et al. (2019) find that diversification in the sciences is almost exclusively pursued through joint work with researchers from other fields. Kelchtermans et al. (2022) report that besides the positive link between joint work and diversification, the better published researchers have a greater presence in other fields, that is, there is also a positive relationship between diversification and performance. On the latter association, Card and DellaVigna (2013) find that citation counts are significantly higher for papers written by

more co-authors. Dubois et al. (2013) report that while the total output of mathematicians who have more collaborating authors tends to be *lower* than the output of those who work separately (who, unlike in economics, are the majority in mathematics), collaborating with authors of a different specialty has positive effects. In their words, interdisciplinary work (within mathematics) spurs productivity.

Our analysis of this nexus makes use of two measures of joint work: The lifetime stock of one's *distinct* RePEc-registered co-authors, *Coauthors*, which includes also the author;and the average number of authors per paper, whether repeated or not. We lean more on the former as it seems conceptually more relevant for the diversification channel.

The first column in Table **??** establishes a link between diversification and co-authorship: Working with new co-authors increases the variety of research output. Columns (2)-(3) reveal that joint work is an important determinant of performance even when holding diversification constant. The comparison of column (3) to its counterpart in Table 2 (column 3) shows an increase in R^2 from 0.3 to 0.38 and a diminution of the estimated direct effect of diversification on performance. The latter suggests that joint work represents a key channel through which output variety impacts on research performance.

Table 9: Diversification, Co-authorship, and Performance

	(1)	(2)	(3)
	log(SDI)	log(Quality-Adjusted Pages)	
log(SDI)		-0.171^{***}	-0.178^{***}
		(0.011)	(0.011)
log(Coauthors)	-0.201^{***}	0.998^{***}	0.984^{***}
	(0.002)	(0.005)	(0.005)
log(Age)	0.064^{***}		
	(0.006)		
Age			0.0045^{***}
			(0.0007)
Items			-0.0001^{***}
			(0.0001)
R^2	0.216	0.612	0.612
N	45,904	45,904	45,904

Notes: Coauthors is the count of distinct registered co-authors including the author, so the regressor is the log of one plus the number of collaborators. Robust standard errors in parentheses.$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The strong interconnection between diversification and joint work makes it difficult to separately identify the causal impact of these two variables: joint work is not exclusively driven by the desire to pursue output variety, and variety can also be obtained without any joint

work. What we learn from Table 9 is that the combination of co-authorship and diversification accounts for a large fraction of performance variation across researchers: those with more co-authors and greater output variety enjoy substantially better performance.

In general, joint work impacts on research performance through three channels. First, it increases raw labor input with more people working on the same paper. Second, it increases the effectiveness of the raw labor input through synergies. The improvement in effectiveness comes from either complementarities in skills and expertise between co-authors that allow for greater input specialization within the team on the basis of comparative advantage (one person doing the data part, another the theory part, another the computation part, and so on); or, greater knowledge transfer and other spillovers within the team that increase one's human capital. And third, by facilitating a more diverse research output, and through that, more/better output.

We do not have data that could help us identify the individual channels. What we can do, though, is to compute the effect of joint work on raw-quality unadjusted and quality adjusted output. The second and third channels presumably impact on the latter while the first one primarily on the former. Table 10 reports the results of the relevant regressions, for pages in columns (1) and (2) and for citations in columns (3) and (4). Co-authorship has a much bigger effect on quality adjusted than on "raw" page counts, 1.03 against 0.26, a finding consistent with the notion that joint work represents something much more consequential than a simple increase in the quantity of the labor input in the research production function. The corresponding gap for citations is far smaller, 1.26 against 0.96, which indicates that the raw citation count embodies quality whereas the raw page count does not.

Table 10: Diversification, Co-authorship and Performance: Quality-Adjusted vs. Unadjusted Output

	(1) log(WScPages)	(2) log(Pages)	(3) log(WScCites)	(4) log(Cites)
log(SDI)	-0.4474^{***}	-0.3383^{***}	-0.2925^{***}	-0.2871^{***}
	(0.0155)	(0.0075)	(0.0150)	(0.0106)
Age	0.0529^{***}	0.0419^{***}	0.0809^{***}	0.0623^{***}
	(0.0032)	(0.0014)	(0.0031)	(0.0022)
log(1 + Coauthors)	1.0257^{***}	0.2614^{***}	1.2582^{***}	0.9560^{***}
	(0.0131)	(0.0085)	(0.0129)	(0.0102)
Items	0.0002	0.0096^{***}	-0.0006^{**}	0.0061^{***}
	(0.0003)	(0.0004)	(0.0003)	(0.0004)
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.398	0.583	0.484	0.590
N	47,950	47,950	47,950	47,950

Notes: All specifications include cohort fixed effects. WScPages and WScCites are weighted by the recursive journal impact factor; Pages and Cites are unweighted. Coauthors is the count of distinct registered co-authors including the author, so the regressor is the log of one plus the number of collaborators. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5 The Model

We have presented evidence in the previous section that the pursuit of research output diversification improves performance (better publications-citations); represents a profitable strategy at all ability levels and all stages of one's career; and it is greatly facilitated by joint work.

Due to data limitations, in particular unobserved human capital/ability, the empirical section is silent on *how* diversification and collaboration strategies operate. Do they work through human capital accumulation? Is joint work's main contribution the facilitation of the pursuit of output variety? And so on. We attempt to gain insights on these issues through the lenses of a simple, structural model of research strategy.

The model contains the following features: the research output depends on human capital, research variety and the number of distinct co-authors. Joint work matters for research output both directly and indirectly through its facilitation of diversification. While productivity enhancing, both co-authorship and diversification are subject to costs that limit their levels. Researchers differ in their initial level of human capital which evolves as a function of the diversification choice and researcher specific productivity shocks.

The model's parameters are obtained by combining estimation (simulated methods of moments) and calibration. It fits the data well and almost all its key parameters are estimated precisely. A discussion of the main results follows the presentation of the model.

5.1 Model Details

Production: Researcher's i output in period t, Y_{it} is

$$Y_{it} = h_{it} \cdot n_{it}^{\psi} \cdot C_{it}^{\nu} \tag{4}$$

where h_{it} is human capital, n_{it} is the degree of diversification (breadth), and C_{it} is the stock of distinct co-authors. The parameter $\nu \geq 0$ captures channels through which joint work raises output that are independent of breadth (such as labor input and skill complementarity).

Human capital accumulation:

$$h_{i,t+1} = z_{it} \cdot h_{it} + \delta_0 \cdot n_{it}^{\omega} \tag{5}$$

where z_{it} is a productivity shock. The second term captures the "learning-by-diversifying effect" discussed previously. This effect is subject to decreasing returns if $\omega < 1$. Note that the distinction between ψ and ω is temporal: ψ governs how breadth raises output *in the present*; ω governs how it raises output *in the future* through skill accumulation.

A researcher's initial human capital, h_{i0}, is drawn from

$$h_{i0} \sim \text{Fréchet}(\alpha_h, s_h) \tag{6}$$

where h_{i0} contains all possible determinants of initial human capital, such as innate ability, training, size of network, etc. We adopt the Fréchet distribution because success in academia is typically judged in terms of relative rather than absolute performance and moreover success breeds success (a right fat tail); Lotka (1926) represents a famous early application.

The idiosyncratic productivity shock evolves according to:

$$\log z_{it} = \rho_z \log z_{i,t-1} + (1 - \rho_z)\mu_z + \varepsilon_{it}, \quad \varepsilon_{it} \sim \text{Gumbel}(0, \beta_{\text{ev}}) - \beta_{\text{ev}}\gamma_E \tag{7}$$

where $0 < \rho_z < 1$, γ_E is the Euler–Mascheroni constant (ensuring that the expected value $\mathbb{E}[\varepsilon] = 0$) and $\mu_z = \log(1 - \delta)$ where δ is the deterministic rate of depreciation of human capital. Again we opt for a right fat tail distribution, the Gumbel this time, to capture the asymmetry of research productivity: "negative" shocks (administrative work, failed projects, child care) are common, while large increases in ability are more rare. Note that the Fréchet

is "fatter" than the Gumbel. This reflects our view that a researcher has a higher probability to be born with extraordinary ability than to acquire it later in life: superstars seem to be born not made. Equations (5) and (7) enable the model to generate a coefficient of variation in the cross sectional research output that matches well that in the data.

Co-authorship: The stock of co-authors is

$$C_{i,t+1} = C_{it} + c_{it} \tag{8}$$

where c_{it} is the new set of distinct co-authors.

The degree of diversification: In the empirical analysis, a researcher's degree of diversification, SDI, in period t is the HHI of the JEL codes in his portfolio of works. In the model, we work directly with the inverse of that index, which we construct in a simpler but equivalent way. Let researcher i have a time invariant, unit endowment of "effort." ℓ, identical for all i, that can be allocated across J topics (JELs). Each topic $j \in J$ draws an amount of effort equal to $\ell_j \in [0, 1]$ with $\sum_{j \in J} \ell_j = \ell = 1$. We define a researcher's degree of diversification as $n = 1/\sum_{j \in J} \ell_j^2$ (the inverse of the HHI index). This measure equals the count of distinct topics when effort is spread equally across n fields, is unity if effort is concentrated in a single field (perfect specialization) and achieves its maximum when $n = J$ and effort is uniformly spread. Having the researcher directly choose the level of n rather than its change (as we do for co-authors) greatly simplifies the analysis and comes with no loss of generality.

Cost of collaboration and diversification: We assume that the management of one's network of co-authors is subject to convex costs.

$$\mathcal{C}_c(c) = \frac{\kappa}{2} c^2 \tag{9}$$

The cost of diversification (the choice of n) is assumed to be individual specific:

$$\mathcal{C}_n(n, C; \gamma_i) = \frac{\gamma_i}{\xi} \cdot \frac{(n-1)^\xi}{(1 + \eta\, C)^\lambda} \tag{10}$$

The cost is assumed to be convex in breadth ($\xi > 1$). This could be the case if people first venture in neighboring fields, close to one's main expertise, before pursuing more distant ones and thus more costly to master. It could also be that working with a large set of fields is challenging. The cost is also decreasing in the stock of distinct co-authors ($\eta > 0$), that is, new co-authors help increase research variety. The assumption that researchers with a larger

network of co-authors can afford to be more interdisciplinary is in line with the evidence presented by Kelchtermans et al. (2022).

The researcher-specific cost γ_i is drawn from a log-normal distribution:

$$\log \gamma_i \sim \mathcal{N}(\log \bar{\gamma},\ \sigma_\gamma^2) \tag{11}$$

This captures variation in field-entry barriers across researchers as a function of researcher characteristics (such as one's primary field) that are not included in the model. Cost heterogeneity, σ_γ, generates SDI dispersion that is orthogonal to that generated by ability heterogeneity (h) and can be calibrated directly from the data once ability has been accounted for.

Optimal research strategy In each period t, researcher i maximizes the expected present value of output net of costs by selecting portfolio diversification $n_{it} \geq 1$ (the effective number of research fields) and collaboration intensity $c_{it} \geq 0$ (the number of new coauthors):

$$V_t(h, C, z) = \max_{n \geq 1,\, c \geq 0} \left\{ Y - \mathcal{C}_n - \mathcal{C}_c \ +\ \rho\, E_{z'|z}[V_{t+1}(h', C', z')] \right\} \tag{12}$$

The choices are subject to equations (5)–(8) and the terminal condition $V_{T+1} \equiv 0$. ρ is the discount factor.

Proposition 1 (Optimal diversification and collaboration). At an interior optimum, the first-order conditions (FOC) are:

$$n: \quad \underbrace{\psi \cdot \frac{Y}{n}}_{\text{direct return}} + \underbrace{\rho\, \mathbb{E}\left[\frac{\partial V_{t+1}}{\partial h'}\right] \cdot \delta_0 \omega\, n^{\omega-1}}_{\text{learning return}} = \underbrace{\gamma_i \frac{(n-1)^{\xi-1}}{(1+\eta C)^{\lambda}}}_{\text{marginal cost}} \tag{13}$$

$$c: \quad \rho\, \mathbb{E}\left[\frac{\partial V_{t+1}}{\partial C'}\right] = \kappa\, c \tag{14}$$

FOC (13) equates two benefits to the marginal diversification cost: a static production gain and a dynamic learning gain. FOC (14) equates the shadow value of a marginal coauthor (which operates through a reduction in future diversification costs) to the marginal collaboration cost. The two choices interact: a larger network of co-authors, C, supports greater diversification (higher n); anticipation of this raises the shadow value of co-authorship, $\partial V/\partial C'$, raising optimal c.

Solution. The model is solved by backward induction in a discretized state space with $T = 35$ periods and $J = 15$ fields, on a grid of 15 human capital points, 12 co-author points, 9

productivity points and 35 diversification choices, with 9 γ-types. It is then simulated forward with $N = 10{,}000$ researchers drawing h_0 from the Fréchet, z from the Gumbel and γ_i from the log-normal distributions. Estimation uses this grid; the reported moments are computed on a finer grid of 25 human capital points, 18 co-author points, 9 productivity points, 50 diversification choices and 11 γ-types, simulated with $N = 100{,}000$ researchers.

5.2 Empirical Implementation

We divide the model's 17 parameters into three groups. The five parameters in the first group are set exogenously: the discount factor, $\rho = 0.95$, the persistence of the productivity shock, $\rho_z = 0.85$, the deterministic rate of human capital depreciation, $\delta = 0.05$, the curvature of the diversification cost, $\xi = 2$, and the elasticity of the diversification cost with respect to the co-author network, $\lambda = 1$. The six parameters in the second group, mostly about heterogeneity, are calibrated: the Gumbel scale, β_{ev}, the Fréchet shape, α_h, the standard deviation of the cost of diversification, σ_γ, the Fréchet scale, s_h, the parameter δ_0 in the human capital accumulation equation, and the initial stock of co-authors, C_0. The remaining six parameters, $\{\psi, \nu, \omega, \bar{\gamma}, \eta, \kappa\}$, which constitute the main determinants of the choice and the impact of research strategy (diversification and co-authorship) are estimated using the simulated method of moments, SMM.

Appendix A.10 provides a detailed discussion of the calibration/estimation procedure. Here, we focus on the economics of the main results. The parameter values are shown in Table 11. They enable the model to fit the data well. In particular, the model matches 11 of the 13 targeted moments within 20%, and generates realistic diversification profiles, collaboration trajectories, and impact of diversification on research output. The two misses are the correlation between early diversification and late-career output, and the within-career slope of diversification.

The model reproduces the targeted cross-sectional OLS coefficient on diversification closely, 0.96 against 1.02 in the data. The production elasticity ψ is smaller than that coefficient because the regression does not hold human capital and the co-author stock fixed, whereas ψ is the partial elasticity of output with respect to breadth alone. It shows that the direct effect of diversification on performance is substantial, while its indirect effect through the process of human capital accumulation (Equation 5) seems quite small as ω is close to zero. The reverse is true for co-authorship. Its direct effect on output, ν, is negligible while its indirect effect through the facilitation of diversified work, η, (Equation 10) is substantial. In other words, a more diversified research portfolio is supported by a network of distinct co-authors and leads to superior research outcomes.

A noteworthy feature is that diversification impacts output without directly affecting the

Table 11: Structural Parameters

Parameter	Estimate	SE (bootstrap)	Description
ψ	0.270	0.001	Diversification elasticity
$\bar{\gamma}$	0.071	0.007	Diversification cost
ω	0.001	0.000	Learning by diversifying
ν	0.017	0.000	Coauthor–output elasticity
η	0.427	0.027	Coauthor–diversification cost
κ	2.573	0.074	Collaboration cost
α_h	1.80		Fréchet shape
s_h	0.30		Fréchet scale
σ_γ	0.80		SD diversification cost
β_{ev}	0.125		Gumbel scale
δ_0	0.15		Human capital–diversification
C_0	1.5		Initial co-author stock
ξ	2		Diversification cost
ρ	0.95		Discount factor
ρ_z	0.85		Productivity persistence
δ	0.05		Human capital depreciation
λ	1		Co-author elasticity of cost
Moments fit ($< 20\%$)	11/13		

human capital of the researcher (the second term in Equation 5). To the extent that the coefficient ω is correctly identified in the SMM, that is, the model is not severely misspecified, such a pattern of "performing better by diversifying but without becoming better" seems puzzling as one may have reasonably thought that the mere act of diversification helps people improve their research skills.

The "Equal-Odds Rule" of Simonton (1997) resolves the puzzle. Recall that his mechanism has precisely the property of doing better without getting better by producing a large portfolio of works. The items of the portfolio serve as "lottery tickets" for achieving breakthroughs.

All in all, the model corroborates the regression results that a diversified body of work pays off; and that a network of co-authors is important for achieving greater diversification. The main discrepancy with the empirical analysis concerns other, non-diversification related, benefits of joint work. In the full sample they are negligible in the model ($\nu = 0.017$) but substantial in the regressions (see Table 9); as we show below, the gap closes in the top 5 sub-sample. Such a discrepancy might arise when ability is more strongly correlated with the formation of co-authorships than with the propensity to diversify. If this were the case, the OLS regression of research output on SDI and co-authorship would lead to a greater bias in the estimate of the latter and the model would provide the more accurate estimate.

But the discrepancy might alternatively arise if the model estimation procedure targeted the wrong, biased, value of a crucial moment. We consider this possibility below.

Top 5 sub-sample: The estimate of SDI in the *full sample OLS regressions* is likely upward biased due to omitted ability, even when controls such as age and number of items are included. Consequently, targeting this value when estimating the model is bound to lead to an exaggerated impact of diversification at the expense of the other production inputs, such as the stock of co-authors. In order to circumvent this problem, we have repeated the estimation/calibration exercise in a sub-sample of the top 5 authors, namely, 972 researchers with 3-5 top 5 journal publications. We believe that doing so goes a long way in controlling for the effects of unobserved ability variation, and better suits the estimation of the model which controls for the interaction of ability and the diversification choice.[11] Note that the values of the moments targeted in the SMM are now those pertaining to the top 5 sub-sample.

In order to facilitate comparisons, Table 12 reports the values of the model parameters obtained in both the full and the top 5 samples. ψ is now smaller as the value of the targeted, corresponding OLS estimate of SDI is smaller, falling from 1.02 to 0.71 (see Appendix A.10). The most noteworthy findings are about the effects of co-authorship. First, the network of

[11] In the model higher ability individuals naturally select a more diversified portfolio as human capital and SDI are complements in the production function. But unlike in the data, in the model this influence can be accounted for, rendering the estimate of SDI, ψ, free of ability bias.

co-authors now has a powerful *direct* effect on production outcomes (a much larger ν), that is, it acts as a productivity booster in a researcher's production function. This suggests the existence of complementarities in the co-authors skill sets.[12] Second, the network of collaborators is now even more consequential in two senses: it has a larger effect on output variety (a higher η) and it makes collaboration more efficient (a lower κ). More collaboration makes also diversification more consequential for this set of researchers in spite of the lower value of ψ through three channels: the complementarity between SDI and the size of the collaboration network in the production function; the greater ease of doing collaborative work; and the greater effect that collaboration has on output variety. The two calibrations also differ in the ability distribution. The Fréchet shape rises from 1.80 to 3.00 and the scale from 0.30 to 1.00, so the top-5 sample is calibrated to a markedly thinner ability tail, consistent with a coefficient of variation of output of 0.63 against 2.76 in the full sample. This is by construction rather than a finding: the sub-sample was chosen precisely to compress ability heterogeneity.

We can use the model to compute the total impact of collaboration on research output through the direct channel, ν, and the indirect channel that operates through the cost of diversification, η. Shutting down the indirect channel by setting $\eta = 0$ lowers median lifetime output by 16.1% in the full sample and by 18.4% in the top-5 sample. Shutting down the direct channel by setting $\nu = 0$ lowers it by 5.7% and 42.4% respectively. In the full sample, then, collaboration raises performance mainly by making diversification cheaper, whereas among top-5 researchers the direct productivity effect dominates. The two channels interact, so the individual losses do not sum to the loss from removing both.

The fact that the model-based results on the effects of diversification are now consistent with those in the corresponding regressions (Table 16 in the appendix) is reassuring because it suggests that neither unobserved ability nor a causal effect of co-authorship on diversification may have tainted the estimates in the top 5 sub-sample OLS regressions.

Before concluding this section let us clarify one important difference between the model and the empirical analysis. A main objective of the empirical analysis has been the identification of the causal contribution to performance of exogenous changes in the degree of diversification. In the model, ability (human capital), diversification and co-authorship are all endogenous and interdependent. So the model cannot be used to answer the main question posed in the empirical section. What the model can be used for is to complement the empirical analysis by shedding light on mechanisms: whether diversification-collaboration's effect on performance involves the accumulation of human capital; whether joint work's main

[12] While new co-authors seem more relevant for greater diversification (Equation 10), the average per paper number of co-authors seems a more relevant input in the production function, Equation 4. Because this adds another choice variable and an additional trade-off, we have opted to not make this distinction.

contribution to performance is through its facilitation of the pursuit of output variety; and so on. In particular, because n and C are deterministic functions of the state and of the cost type γ_i, the model cannot be used to decompose the cross-sectional variance of output into a contribution of research strategy separate from that of the fundamentals. The single source of strategy variation orthogonal to ability and luck is the idiosyncratic diversification cost γ_i, whose dispersion is disciplined by the residual SDI dispersion that we target, and it accounts for less than one percent of the variance of log lifetime output.

Table 12: Estimated parameters

Parameter	Full sample	Top-5 sample
ψ	0.270 (0.001)	0.162 (0.001)
$\bar{\gamma}$	0.071 (0.007)	0.064 (0.001)
ω	0.001 (0.000)	0.001 (0.000)
ν	0.017 (0.000)	0.170 (0.003)
η	0.427 (0.027)	0.518 (0.008)
κ	2.573 (0.074)	1.393 (0.020)
α_h	1.80	3.0
s_h	0.30	1.0
σ_γ	0.80	0.50
β_{ev}	0.125	0.025
δ_0	0.15	0.15
C_0	1.5	2.5
ξ	2	2
ρ	0.95	0.95
ρ_z	0.85	0.85
δ	0.05	0.05
λ	1	1
Objective	0.418	0.250
Moments fit ($< 20\%$)	11/13	12/13

Notes: Bootstrap standard errors (500 replications) in parentheses.

6 Conclusions

Specialization is ubiquitous and consequential in nature and human activity. In economics, it is considered to be a major determinant of performance, and its properties have been studied extensively in areas as varied as international trade, portfolio theory, and corporate scope. In this paper, we explore its role in the domain of research strategy, utilizing the RePEc database

of about 60,000 economists. We propose a novel metric of specialization that relies on the Journal of Economic Literature classification system. We address three core questions: Does the pursuit of research output variety improve performance? Does the answer to this question depend on the researcher's level of ability and career stage? And what role does collaborative work play in this process?

Conventional wisdom claims that diversification of research output is detrimental to early-career researchers. We find no support for this claim. In contrast, it seems to enhance performance (citations/ publications) and also to reduce research attrition among junior scholars. We also establish that joint work is a significant determinant of both the degree of diversification and research performance.

Our model reveals some interesting properties of the diversification-collaboration-performance nexus that are not empirically detectable due to the lack of suitable data. Two properties stand out. First, collaborative work's most significant contribution to research output comes from its facilitation of the pursuit of diversification. And second, the effect of diversification on research output does not owe to human capital gains arising from learning-by-diversifying. Instead, the estimated model supports Simonton's (1997) "Equal-Odds Rule" which argues that a diversified research portfolio's main benefit comes from the fact that it contains more "lottery tickets" for success.

It would be of interest to examine whether our main findings apply also to other scientific fields, such as medicine.

References

Abramo, Giovanni, Ciriaco Andrea D'Angelo, and Flavia Di Costa, "Diversification versus specialization in scientific research: Which strategy pays off?," *Technovation*, 2019, *82–83*, 51–57.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang, "Superstar Extinction," *Quarterly Journal of Economics*, 2010, *125* (2), 549–589.

Berg, Justin M., "One-Hit Wonders versus Hit Makers: Sustaining Success in Creative Industries," *Administrative Science Quarterly*, 2022, *67* (3). Published online March 24, 2022.

Brogaard, Jonathan, Joseph Engelberg, and Christopher A. Parsons, "Networks and Productivity: Causal Evidence from Editor Rotations," *Journal of Financial Economics*, 2014, *111* (1), 251–270.

Card, David and Stefano DellaVigna, "Nine Facts about Top Journals in Economics," *Journal of Economic Literature*, 2013, *51* (1), 144–161.

Conley, John P. and Ali Sina Önder, "The Research Productivity of New PhDs in Economics: The Surprisingly High Non-success of the Successful," *Journal of Economic Perspectives*, 2014, *28* (3), 205–216.

Diegert, Paul, Matthew A. Masten, and Alexandre Poirier, "Assessing Omitted Variable Bias when the Controls are Endogenous," Technical Report 2022. Working paper.

Dubois, Pierre, Jean-Charles Rochet, and Jean-Marc Schlenker, "Productivity and mobility in academic research: evidence from mathematicians," *Scientometrics*, 2014, *98* (3), 1669–1701.

Ekerdt, Lorenz K. F. and Kai-Jie Wu, "The Rise of Specialized Firms," Working Paper 24-06, Center for Economic Studies, U.S. Census Bureau 2024.

Kelchtermans, Stijn, Daniel Neicu, and Reinhilde Veugelers, "Off the beaten path: what drives scientists' entry into new fields?," *Industrial and Corporate Change*, 2022, *31* (3), 654–680.

Laband, David N. and Michael J. Piette, "Favoritism versus Search for Good Papers: Empirical Evidence Regarding the Behavior of Journal Editors," *Journal of Political Economy*, 1994, *102* (1), 194–203.

Lotka, Alfred J., "The frequency distribution of scientific productivity," *Journal of the Washington Academy of Sciences*, 1926, *16* (12), 317–323.

McAfee, Andrew and Erik Brynjolfsson, *Machine, Platform, Crowd: Harnessing Our Digital Future*, New York: W. W. Norton, 2017.

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.

Rumelt, Richard P., *Strategy, Structure, and Economic Performance*, Boston: Division of Research, Harvard Business School, 1974.

Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben, and Arthur Schram, "Gender Differences in Recognition for Group Work," *Journal of Political Economy*, 2021, *129* (1).

Simonton, Dean Keith, "Creative Productivity: A Predictive and Explanatory Model of Career Trajectories and Landmarks," *Psychological Review*, 1997, *104* (1), 66–89.

A Empirical Appendix

A.1 Career Age and the Top-Five Partition

A natural objection to the partitions of Table 5 is that they sort researchers by career age rather than by ability. Within each binned column the number of top-five publications is only weakly related to academic age, and more weakly still net of the number of items, which is the relevant conditional because both enter the regressions.

Table 13: Career Age and Top-Five Count within the Estimation Samples

Sample	N	Mean age	ρ(age, top 5)	ρ net of items
≥ 1 Top 5	6,647	28.6	+0.334	+0.153
2–5 Top 5	2,511	29.3	+0.162	+0.125
5–10 Top 5	1,006	34.7	+0.088	+0.061
> 10 Top 5	487	42.0	+0.210	+0.031

Notes: Academic age is years since the first publication. The final column reports the correlation between the residuals of age and of the top-five count, each projected on the number of items. The ≥ 1 row pools the entire range of top-five counts, which is why its correlation is the largest of the four.

A.2 Other Homogeneous Partitions

Table 14 partitions on the performance distribution and on pedigree rather than on top-five counts. The estimated elasticities, between -0.32 and -0.41, are close to those obtained from the top-five partitions of Table 5.

Table 14: Diversification and Performance: Percentile and Pedigree Partitions

	(1) Top 10% cites	(2) Top 25% cites	(3) Top 10% cites & top-10 PhD	(4) Top 10% cites & top-50 PhD
log(SDI)	-0.3169^{***}	-0.4077^{***}	-0.3239^{***}	-0.3177^{***}
	(0.0221)	(0.0183)	(0.0346)	(0.0274)
Age, Age^2, Items	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.152	0.158	0.280	0.248
N	4,622	11,555	1,442	2,486

Notes: Dependent variable is log(Quality-Adjusted Pages). Percentiles are of quality-adjusted citations within the estimation sample of Table 2. Pedigree is the RePEc rank of the PhD-granting department. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.3 Panel Estimation: Pages

Table 15: Panel Estimation: Diversification and Pages

	Full Sample		Top 5 Sample		3–7 Top 5	
	(1) OLS	(2) FE	(3) OLS	(4) FE	(5) OLS	(6) FE
log(SDI)	-1.367^{***}	-0.879^{***}	-0.686^{***}	-0.799^{***}	-0.535^{***}	-0.796^{***}
	(0.004)	(0.009)	(0.005)	(0.021)	(0.007)	(0.038)
Individual FE	No	Yes	No	Yes	No	Yes
Year FE	No	Yes	No	Yes	No	Yes
R^2 Within		0.111		0.109		0.119
N	805,059	805,059	145,970	145,970	40,950	40,950

Notes: Dependent variable is log(Quality-Adjusted Pages). The estimation sample restricts on pages rather than on citations, which is why it is smaller than that of Table 6. Standard errors clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.4 Top-Five Partitions with Co-authorship

Table 16 repeats the top-five partitions of Table 5 with the co-author stock included, which is the empirical counterpart to the direct co-author channel ν in the model of Section 5.

Table 16: Diversification and Co-authorship: Researchers with Top-5 Publications

	(1) ≥ 1 Top 5	(2) 2–5 Top 5	(3) 5–10 Top 5	(4) >10 Top 5
log(SDI)	-0.246^{***}	-0.128^{***}	-0.082^{***}	-0.158^{***}
	(0.019)	(0.022)	(0.023)	(0.036)
log(1 + Coauthors)	0.408^{***}	0.245^{***}	0.188^{***}	0.170^{***}
	(0.020)	(0.024)	(0.036)	(0.037)
Age, Age^2, Items	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.375	0.299	0.439	0.532
N	6,647	2,511	1,006	487

Notes: Dependent variable is log(Quality-Adjusted Pages). Samples are those of Table 5. Robust standard errors in parentheses. $^{***}p < 0.01$.

A.5 Panel Estimation with Lagged Diversification

Table 17: Panel Fixed Effects with Lagged Diversification

	(1) Pooled OLS	(2) FE	(3) FE, Lag 1	(4) FE, Lag 2	(5) FE, Lag 3
Panel A: Dependent variable is log(*Quality-Adjusted Citations*)					
log(SDI)	-1.606^{***}	-0.982^{***}	-0.756^{***}	-0.577^{***}	-0.439^{***}
Panel B: Dependent variable is log(*Quality-Adjusted Pages*)					
log(SDI)	-1.367^{***}	-0.879^{***}	-0.745^{***}	-0.618^{***}	-0.500^{***}
Individual FE	No	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes
Observations (Panel A)	863,018	863,018	811,216	759,494	707,927
Researchers	51,802	51,802	51,802	51,802	51,802

Notes: Lags of the diversification index are constructed within the estimation sample. Standard errors clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.6 Instrumental Variables

Table 18 instruments a researcher's own diversification with the diversification of her co-authors. The first stage is strong. The instrumented estimate is larger in magnitude than the corresponding fixed-effects estimate.

Table 18: Co-author Diversification as an Instrument

	(1) First stage log(SDI)	(2) FE log(Cites)	(3) FE + IV log(Cites)	(4) FE + IV log(Cites)
log(SDI of co-authors)	0.2172***			
	(0.0046)			
log(SDI)		−0.9885***	−1.2915***	−1.1990***
		(0.0090)	(0.0472)	(0.0559)
Instrument	—	—	Current	Lagged one year
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
First-stage F	2,184			
N	798,174	798,174	798,174	750,826

Notes: The dependent variable in columns (2)–(4) is log(Quality-Adjusted Citations). The instrument is the weighted average diversification index of a researcher's co-authors. The sample is that of Table 6 restricted to author-years for which the co-author index is defined. Standard errors clustered at the individual level. ***$p < 0.01$. The journal-based instrument discussed in Section 4.2.2 is not reported because it does not satisfy the exclusion restriction: our performance measures weight output by the recursive impact factor of the publishing journal, so the instrument and the outcome are constructed from the same set of journals.

A.7 Heterogeneity by Gender and Pedigree

Table 19: Diversification and Performance by Gender and Pedigree

	(1) Male	(2) Female	(3) Top 10 PhD	(4) Rank 11–50	(5) Below 50
log(SDI)	−0.4161***	−0.5439***	−0.3857***	−0.4059***	−0.3240***
	(0.0181)	(0.0342)	(0.0512)	(0.0516)	(0.0337)
log(1 + Coauthors)	1.0744***	1.0527***	0.7507***	0.6617***	0.9556***
	(0.0152)	(0.0297)	(0.0489)	(0.0472)	(0.0270)
Age	0.0530***	0.0653***	0.0295***	0.0452***	0.0469***
	(0.0037)	(0.0071)	(0.0087)	(0.0086)	(0.0067)
Items	−0.0002	−0.0012	0.0009	0.0019***	−0.0007**
	(0.0003)	(0.0009)	(0.0006)	(0.0007)	(0.0003)
Cohort FE	Yes	Yes	Yes	Yes	Yes
R^2	0.422	0.334	0.428	0.387	0.387
N	34,217	11,110	2,618	3,679	8,404

Notes: Dependent variable is log(Quality-Adjusted Pages). Gender is assigned only where the name-based classifier score exceeds 0.5 in absolute value, so researchers with ambiguous scores are excluded rather than allocated, which is why columns (1) and (2) do not sum to the estimation sample. Pedigree is the RePEc rank of the PhD-granting department. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.8 Quality-Adjusted versus Raw Output

Table 20: Diversification Effect: Quality-Adjusted versus Raw Outcomes

	(1) log(WScPages)	(2) log(Pages)	(3) log(WScCites)	(4) log(Cites)
log(SDI)	-0.1995^{***}	-0.3326^{***}	0.0911^{***}	-0.0520^{***}
	(0.0112)	(0.0080)	(0.0066)	(0.0054)
Age	-0.0108^{***}	0.0313^{***}	-0.0050^{***}	0.0025^{**}
	(0.0025)	(0.0014)	(0.0014)	(0.0011)
log(Coauthors)	1.0015^{***}	0.1808^{***}	1.3314^{***}	0.9387^{***}
	(0.0051)	(0.0035)	(0.0036)	(0.0030)
Items	-0.0007^{***}	0.0101^{***}	-0.0027^{***}	0.0052^{***}
	(0.0001)	(0.0004)	(0.0001)	(0.0002)
Cohort FE	Yes	Yes	Yes	Yes
R^2	0.640	0.601	0.894	0.908
N	47,950	47,950	47,950	47,950

Notes: WScPages and WScCites are weighted by the recursive journal impact factor; Pages and Cites are unweighted. The raw counts are each researcher's maximum over her panel observations. The co-author measure is strongly correlated with quality-weighted citations, which attenuates the diversification coefficient in column (3) and reverses its sign. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.9 Figures

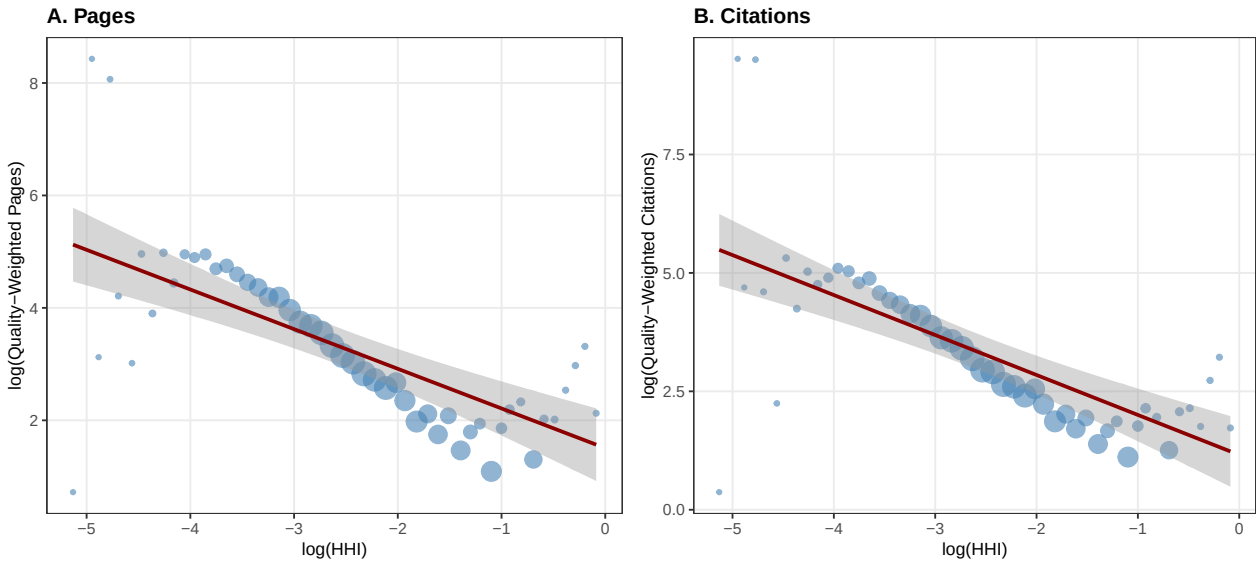
Figure A.1 plots binned means of quality-adjusted output against the diversification index in logs, and Figure A.2 plots the same relationship in levels. Each point is one of fifty equal-width bins of the index, sized by the number of researchers it contains, and the line is a linear fit.

A.10 Calibration and Estimation of the Model

This appendix describes how the parameters in Tables 11 and 12 are obtained and reports the fit of the targeted moments.

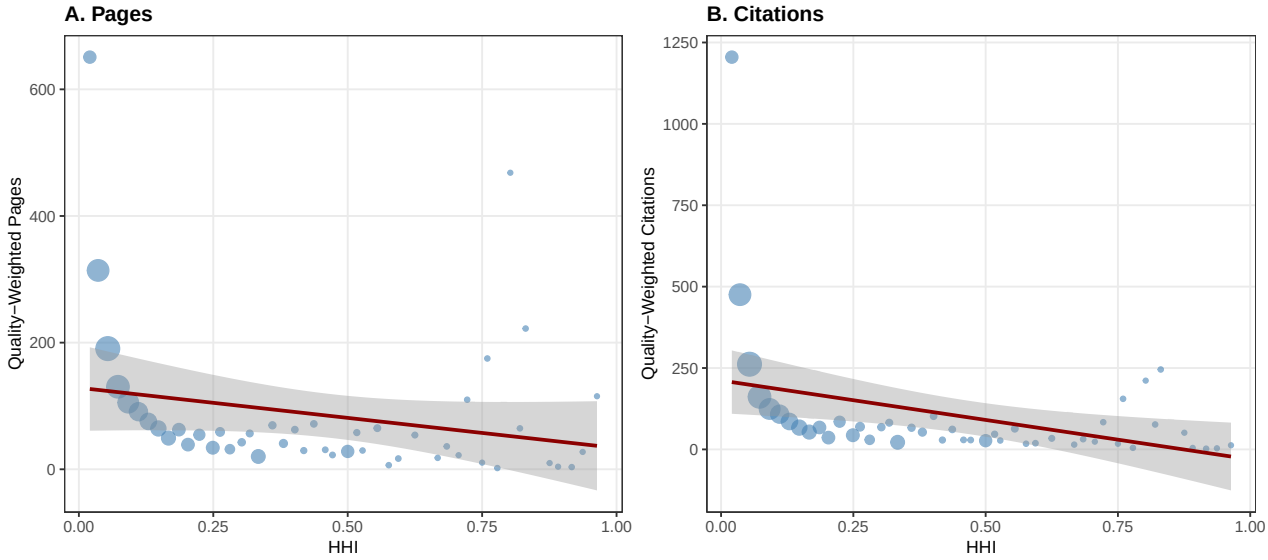
The diversification variable used in estimation. The moments are constructed from $1 - \text{SDI}_{it}$, where SDI_{it} is the Herfindahl index of Section 3.2, so that the variable increases in breadth. In the model the counterpart is $1 - 1/n_{it}$, since $n = 1/\text{SDI}$. This is why the targeted cross-sectional coefficient is positive, while the elasticities with respect to log(SDI)

Figure A.1: Quality-adjusted output and diversification, in logs



Source: RePEc. Panel A plots log(Quality-Adjusted Pages) and Panel B log(Quality-Adjusted Citations) against log(SDI). Binned means over fifty bins of the index; point size is proportional to the number of researchers in the bin.

Figure A.2: Quality-adjusted output and diversification, in levels

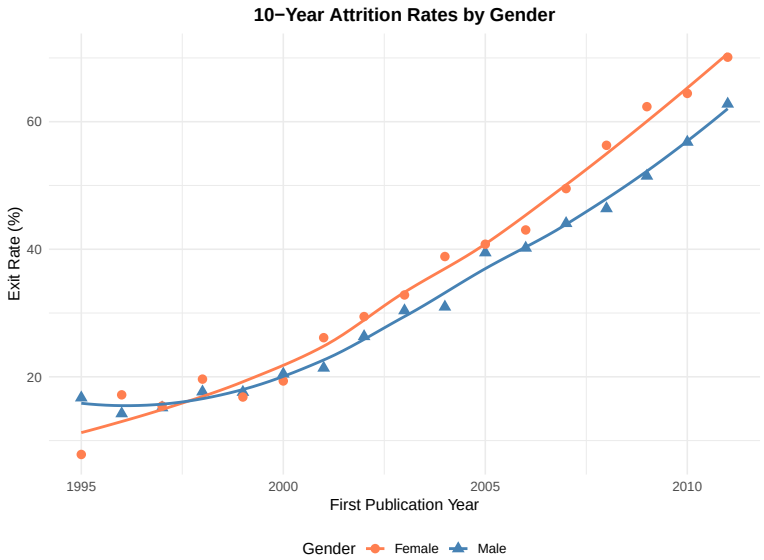


Source: RePEc. As Figure A.1, with the index and the outcomes in levels rather than in logs.

reported in Section 4 are negative. Age effects are removed from the variable on both sides, by a quadratic in academic age in the data and by period demeaning in the simulation.

Samples. The full-sample moments are computed on the panel of researchers with positive quality-adjusted output and academic age between one and thirty-five. The top-5 moments are computed on the researchers whose lifetime count of top-five publications lies between

Figure A.3: Research attrition



Source: RePEc. Attrition after five, ten and fifteen years from the first publication, by cohort and gender. For the male cohort of 2004 (1,603 individuals), 136 (8.5%) dropped out after five years, 199 (13.5% of survivors) after ten years and 291 (23% of survivors) after fifteen.

three and five.

Estimated and calibrated parameters. Six parameters, $\{\psi, \bar{\gamma}, \omega, \nu, \eta, \kappa\}$, are estimated by the simulated method of moments, minimising the sum of squared proportional deviations between the thirteen targeted moments and their model counterparts. Six further parameters are calibrated: α_h and s_h govern the ability distribution, σ_γ the dispersion of the diversification cost, δ_0 the scale of human capital accumulation, C_0 the initial co-author stock, and β_{ev} the Gumbel scale. The last is set by a grid search so that the coefficient of variation of lifetime output matches its value in the data, 2.76 in the full sample and 0.63 in the top-5 sample. The remaining five parameters are fixed a priori.

Numerical implementation. The model is estimated on a grid of 15 human capital points, 12 co-author points, 9 productivity points, 35 diversification choices and nine γ-types, simulated with 10,000 researchers under common random numbers. The optimiser is Nelder–Mead with thirty random starting values, 1,000 iterations per start and a 2,000-iteration refinement from the best start. The parameters so obtained are then evaluated on a finer grid of 25 human capital points, 18 co-author points, 9 productivity points, 50 diversification choices and eleven γ-types, simulated with 100,000 researchers; all reported moments and decompositions use this finer grid. Standard errors are computed by bootstrap with 500 replications.

Fit. Table 21 reports the thirteen targeted moments. Eleven of the thirteen fall within 20%

of their data counterparts in the full sample and twelve of thirteen in the top-5 sample. The full-sample misses are the correlation between early diversification and late-career output, which the model overstates, and the life-cycle slope of diversification, which it overstates by 23%. The single top-5 miss is the dispersion of residual diversification, which the model understates by 43%: with ability heterogeneity compressed by construction in that subsample, the model generates less strategy variation orthogonal to ability than the data display.

Table 21: Targeted Moments and Model Fit

	Full sample			Top-5 sample		
Moment	Data	Model	Err. (%)	Data	Model	Err. (%)
Cross-sectional OLS coefficient on diversification	1.016	0.964	−5.1	0.707	0.682	−3.6
SD of residual diversification	0.108	0.100	−7.2	0.078	0.045	−42.7
Corr. (log co-author stock, diversification)	0.351	0.351	0.0	0.357	0.364	1.8
Corr. (early diversification, late output)	0.097	0.148	53.4	0.177	0.205	15.9
Mean diversification, age 5	0.812	0.698	−14.0	0.836	0.783	−6.3
Mean diversification, age 10	0.854	0.749	−12.2	0.886	0.823	−7.1
Mean diversification, age 15	0.877	0.779	−11.2	0.909	0.847	−6.9
Mean diversification, age 20	0.893	0.799	−10.6	0.917	0.862	−6.1
Diversification slope, age 20 less age 5	0.081	0.100	23.4	0.081	0.079	−3.3
Mean co-author stock, age 5	3.902	3.799	−2.6	6.175	6.589	6.7
Mean co-author stock, age 10	6.323	6.824	7.9	9.910	10.913	10.1
Mean co-author stock, age 20	12.405	12.988	4.7	17.298	17.803	2.9
Mean annual change in co-author stock	0.463	0.508	9.6	0.705	0.595	−15.7
Moments within 20%		11 of 13			12 of 13	

Notes: Diversification is 1 − SDI, so higher values indicate broader portfolios. Residual diversification is net of a quadratic in academic age. Model values are computed on the fine grid with 100,000 simulated researchers. Err. (%) is the proportional deviation of the model moment from its data counterpart.