

Testing for path dependence in economic history publications*

Johan Fourie[†]

Abstract

This paper asks whether the concept of path dependence applies to economic history’s own publication record. Using annual citation histories for 1,262 articles published between 2012 and 2021 in four economic history journals, I document strong persistence: conditional on journal and publication year, the elasticity of long-run citations with respect to age-two citations is about 0.78. The relationship holds when the outcome is restricted to citations accumulated strictly after the early window, and it is stable across journals, topics and publication cohorts. Observable author and article characteristics explain about a fifth of the variation in early citations, yet the unexplained component predicts long-run citations almost as strongly as the explained component. Articles placed earlier in a journal issue receive more early and more long-run citations, which is consistent with a visibility channel. Together these patterns are consistent with cumulative advantage and difficult to attribute to quality differences alone, although unobserved quality cannot be ruled out. Half of all citations come from economics, and two fifths arrive through formats other than journal articles. Conference presentation shows no reliable citation premium. The findings are relevant to the use of early citation counts in research evaluation.

Keywords: path dependence; cumulative advantage; citations; bibliometrics

JEL codes: A14; N01; O33

*The citation data were collected by hand over more than a decade by successive cohorts of research assistants at LEAP (the Laboratory for the Economics of Africa’s Past), Stellenbosch University. I thank two anonymous referees for comments that substantially improved the paper. The usual disclaimer applies. I used a generative artificial-intelligence assistant (a large language model) in the course of this work, and beyond routine copy-editing: the assistant helped to write and debug the analysis code, execute and tabulate the empirical results, draft and revise sections of the text, and produce the figures. I specified the research questions and design, made all substantive analytical and interpretive decisions, independently verified the results and the underlying code, and take full responsibility for the content of the paper. The AI tool is not listed as an author, as it cannot be accountable for the work. Cite this paper as: Fourie, Johan. 2026. “Testing for path dependence in economic history publications.” Working Paper, Department of Economics, Stellenbosch University.

[†]Department of Economics, Stellenbosch University. Email: johanf@sun.ac.za.

1 Introduction

Path dependence is one of economic history’s most widely adopted concepts. In the canonical formulation, early events that contain little information about underlying quality can be amplified by positive feedback, producing persistent and sometimes permanent differences in outcomes (David, 1985; Arthur, 1989; David, 2007; Puffert, 2002, 2024). The concept has influenced empirical work on spatial development (Bleakley and Lin, 2012), institutional change (Mahoney, 2006) and the broader persistence-studies agenda in economics (Cioni, Federico and Vasta, 2022). This paper asks whether the same concept applies to the discipline’s own publication record: whether differences in the early reception of economic history articles are amplified into persistent differences in their long-run influence.

The object of study is broader than citation counts. Scholarly knowledge is produced (an article is published), disseminated (it circulates as a working paper, is presented at conferences, is reprinted or discussed in books), recognised (it is cited) and eventually either absorbed into the field’s canon or forgotten. Citations are the measurable record of this process. If the allocation of scholarly attention is self-reinforcing – if being read early makes a paper more likely to be read later, independently of its content – then the canon of economic history partly records early reception rather than only the merits of contributions. This possibility matters in practice because citation counts are increasingly used in hiring, promotion and research evaluation.

The paper makes four contributions to the cliometric study of knowledge production in economic history. First, it documents strong persistence in citation trajectories within the four core economic history journals: conditional on journal and publication year, the elasticity of long-run (age five to eight) citations with respect to age-two citations is approximately 0.78, and the relationship holds when the outcome is restricted to citations accumulated strictly after age two, which removes the mechanical overlap between a cumulative outcome and a cumulative regressor. Second, it tests whether this persistence can be reduced to observable fundamentals, the main empirical alternative to cumulative advantage. It cannot: author and article characteristics, including author seniority, h-index and prize records, explain about a fifth of the variation in early citations, yet the four fifths they leave unexplained predict long-run citations almost as strongly as the part they explain. Third, the paper organises the possible reinforcement channels – search visibility, dissemination formats, conferences, citation networks and teaching – into a single framework, states the testable implication of each and reports the corresponding evidence. Visibility receives direct reduced-form support: articles placed earlier in a journal issue, a placement only weakly related to content, receive more early and more long-run citations. Conferences receive none: matched to hand-verified programmes of the two annual association meetings, presented papers show no reliable citation premium, in contrast to earlier findings by Di Vaio, Waldenström and Weisdorf (2012). Fourth, the paper documents where the field’s citations come from: half from economics, three in ten from other social sciences and two fifths through formats other than journal articles, mainly books, chapters and working papers.

Throughout, I distinguish between what the data establish and what they suggest. The data

establish persistence. Whether that persistence constitutes path dependence – the amplification of differences in early reception that are unrelated to quality – cannot be established conclusively with observational citation data, because the unexplained component of early reception contains unobserved quality as well as contingency. My claim is therefore deliberately bounded: the evidence is consistent with cumulative advantage and difficult to reconcile with quality sorting alone, and the title’s ‘testing’ is meant in that spirit.

The paper proceeds as follows. Section 2 defines persistence and path dependence, sets out the reinforcement model and presents the channel framework. Section 3 describes the data, defines all variables and accounts for the sample. Section 4 presents the persistence results and their robustness. Section 5 contains the central test: the decomposition of early citations into an explained and an unexplained component. Section 6 presents the channel-level evidence. Section 7 interprets the findings and concludes.

2 Persistence, cumulative advantage and channels

2.1 Two concepts

It is useful to distinguish two concepts at the outset. *Persistence* is a statistical property: a positive elasticity of long-run citations with respect to early citations. *Path dependence* is a claim about mechanism: differences in early reception that are not attributable to a paper’s fundamentals are amplified into long-run differences. Persistence is consistent with path dependence, but it is equally consistent with quality sorting, in which better papers attract citations both early and late and early citations merely reveal quality sooner. Liebowitz and Margolis (1995) made the corresponding point about technology adoption: observed persistence does not by itself establish lock-in. Later work distinguishes lock-in from weaker forms of path dependence, in which early events influence but do not determine the outcome (Brenner and zu Jeddelloh, 2024; Roos, 2024). Section 4 establishes persistence; Section 5 tests how much of it survives conditioning on fundamentals. I claim path dependence only in the bounded sense defined there.

2.2 A model of reinforcement

Suppose the expected flow of new citations for paper i at age t satisfies

$$\mathbb{E}[\Delta C_{it} \mid C_{i,t-1}, X_i] = \lambda_t \cdot (C_{i,t-1} + 1)^\theta \cdot \exp(X_i' \gamma), \quad (1)$$

where C_{it} is cumulative citations, $\Delta C_{it} = C_{it} - C_{i,t-1}$, λ_t captures the citation life cycle, and X_i are paper characteristics. When $\theta > 0$ the process is self-reinforcing: a higher existing stock raises expected new citations, so early shocks persist and the log-log elasticity between citations at two ages is positive.¹ Skewed citation distributions and persistent inequality of attention are standard

¹David (1985) formalised QWERTY lock-in as a generalised Pólya urn, in which each draw of a colour raises the probability of drawing it again. Equation (1) is the citation analogue, preferential attachment in scholarly attention

findings in bibliometrics, in economics as elsewhere (Hamermesh, 2018; Bornmann and Wohlrabe, 2019; Aistleitner, Kapeller and Steinerberger, 2019).

Equation (1) also states the identification problem. A positive elasticity can reflect $\theta > 0$ (reinforcement) or persistent variation in $\exp(X_i'\gamma)$ that the econometrician does not observe (quality). My strategy is to estimate the elasticity conditioning on the fundamentals that can be observed (Section 4) and to test whether the component of early citations orthogonal to those fundamentals still predicts long-run outcomes (Section 5); I do not claim a single causal parameter. Publication and career incentives in economics are concentrated on a small set of journals (Heckman and Moktan, 2020), and economic history has documented patterns of stratification and co-authorship (Seltzer and Hamermesh, 2018; Cioni, Federico and Vasta, 2020); both encourage coordination on work that is already visible.

2.3 Channels

If attention is self-reinforcing, it operates through specific channels, each acting on what is already visible. The first is search and visibility. Readers find papers through search engines, reference lists and tables of contents, which rank or expose papers partly by existing citations and placement, so early attention raises the probability of later discovery. The concept's own history illustrates the point: David (1985) appeared in the *American Economic Review* Papers and Proceedings, a short and widely read format, and remains the canonical citation for path dependence, cited several times as often as its best-known critique (Liebowitz and Margolis, 1990). The implication tested here is that a within-issue placement advantage, which is only weakly related to content, predicts both early and long-run citations (Section 6.1).

A second channel is the variety of dissemination formats. Books, book chapters, dissertations and working papers make results available to readers who do not follow the journals, and their authors cite what is already visible. Fogel and Engerman (1974) reached its audience as a book and was reviewed in the general press, retaining its canonical position although several of its findings were contested. The implication is a large share of citations from formats other than journal articles (Section 6.2).

A third channel is conference presentation, which puts a paper before an audience before or around publication. Conrad and Meyer (1958), conventionally credited with founding cliometrics, was presented at the 1957 joint meeting of the Economic History Association and the NBER Conference on Research in Income and Wealth before its journal publication. The implication is that presented papers show a citation premium (Section 6.3).

A fourth channel is the citation network itself. A citation is made by a paper that is itself cited, so attention propagates through the network of citing works, and reference lists are partly copied from other reference lists. North and Weingast (1989) is the most-cited article in the *Journal of Economic History* on the publisher's own count, while sustained empirical challenges, for example Sussman and Yafeh (2006), are cited orders of magnitude less often. The implication is that early

(Merton, 1968; Price, 1976); the analogy is heuristic, and θ is not estimated structurally below.

citations predict subsequent citation flows, not only cumulative totals (Sections 4 and 6.4).

The fifth channel is teaching and canonisation. Syllabi, textbooks and survey articles consolidate attention on a small set of references. These data contain no measure of teaching use, so this channel is not tested.

3 Data and measurement

3.1 Corpus and journal selection

The dataset contains 3,250 research articles published between 1997 and 2025 in the four generalist economic history journals: the *Journal of Economic History* (JEH), the *Economic History Review* (EHR), *Explorations in Economic History* (EEH) and the *European Review of Economic History* (EREH). The citation records were hand-coded at LEAP (Stellenbosch University) since 2013. These are the four established general-interest journals of the field and the standard corpus in quantitative studies of economic history publishing (Di Vaio and Weisdorf, 2010; Di Vaio, Waldenström and Weisdorf, 2012; Cioni, Federico and Vasta, 2020). Younger or more specialised outlets (including *Cliometrica*, founded in 2007, and the regional economic history journals) are not included, because the design requires long, complete citation histories collected on a consistent protocol from before 2014, and the data collection began before several of these outlets had accumulated the required history. The conclusions apply to the field’s four core journals.

3.2 Citation measurement

Citation counts come from Google Scholar. The author and several research assistants recorded the cumulative citation count of every article in the corpus once a year, in February–March, from 2014 through 2026, giving thirteen annual snapshots per article. The citation count of article i at age k is defined as its count in the snapshot of calendar year (publication year + k). Because snapshots are taken early in the year, age k corresponds to roughly k years of exposure for an article published early in its publication year and somewhat less for one published late in the year; journal and year fixed effects absorb systematic differences in this timing across journals and cohorts, and the within-issue analysis of Section 6.1 compares articles in the same issue, which share a publication date.

Early citations are the age-two count. *Long-run citations* are the age-eight count where the snapshots permit it (publication years 2012–2018) and the age-five count otherwise (2019–2021). Regressions use $\log(1 + \text{count})$. Google Scholar is broader but noisier than curated databases; it counts citations from working papers, books and non-indexed outlets, which is an advantage for a field whose citations often arrive in those forms (Section 6.2). Counts occasionally decline between snapshots; where a growth measure is used, negative growth is set to zero before taking logs. As an external check, the corpus was matched to OpenAlex: 3,241 of 3,250 articles (99.7 percent) match, and among matched articles the OpenAlex citation count agrees closely with the Google Scholar count (correlation of log counts 0.71; rank correlation 0.78; Appendix A).

3.3 Variable definitions

The following variables are used throughout; Appendix A gives full construction detail.

Fast starter: an indicator equal to one if the article’s early citations are above the 75th percentile of its journal-year cohort. The indicator is a discretisation of the continuous early-citation variable and is used descriptively and in the placebo test; Section 4.4 shows it has no content beyond the continuous variable.

Top institution: an indicator equal to one if any of the first five authors is affiliated with one of seventeen named research universities in the United States and the United Kingdom (the list is in Appendix A). The list is an ad hoc proxy for institutional visibility in the two countries that dominate the corpus, not a ranking; Section 4.3 shows the main estimate is unchanged when the control is dropped.

Topic: one of sixteen categories assigned by keyword matching in article titles, with a residual ‘other’ category for the 19 percent of estimation-sample titles matching no keyword. The full keyword dictionary is in Appendix A. Results are reported both excluding and retaining the ‘other’ category.

Article controls: number of authors, log article length in pages, title length in characters, position of the article within its issue (by first page), and issue number.

Author quality (available for the 58 percent of the estimation sample matched to RePEc): the team’s maximum author seniority (years since first recorded publication) and maximum h-index; indicators for an article-level prize (Cole, Ashton or Figuerola) and for an author having won a dissertation prize (Gerschenkron or Nevins) before publication.

Conference presentation: an indicator equal to one if the article was matched to a presentation at the annual meetings of the Economic History Association (EHA) or the Economic History Society (EHS). The EHA programmes (1,006 papers, 2006–2025, with session times) were transcribed by hand because the original multi-column PDFs cannot be parsed reliably; EHS programmes (2003–2024) were recovered from structured HTML. Articles are linked to presentations by a fuzzy title match validated by author surnames (Appendix E). The variable measures presentation at the field’s two annual association meetings, not conference exposure in general: the biennial European Historical Economics Society (EHES) meeting and smaller workshops are not covered, because their programmes are not archived in a consistently recoverable form. This scope restriction is a limitation, noted again in Section 6.3.

3.4 Sample construction

The construction of the estimation sample is documented in full in Appendix Table A1. Of the 3,250 articles in the four core journals, 1,672 have an observable age-two count (publication year 2012 or later, given the first snapshot in 2014), and 1,262 also have an observable age-five or age-eight count (publication year 2021 or earlier, given the last snapshot in 2026). No observations are lost to missing author counts or negative citation counts. The estimation sample is therefore the

1,262 core-journal articles published 2012–2021; the binding constraint is the snapshot window, not data quality.

3.5 Summary statistics

Table 1 reports summary statistics for the estimation sample. The mean article has 7.1 citations at age two and 34.4 at age five to eight; both distributions are strongly right-skewed, which motivates the log specification. About 26 percent of articles are fast starters by construction (the within-cohort threshold), and 23 percent have a top-institution author.

Table 1: Summary statistics (estimation sample)

Variable	N	Mean	SD
Early citations (age 2)	1,262	7.11	10.90
Long-run citations (age 5–8)	1,262	34.43	44.38
Citation growth after age 2	1,262	27.31	37.64
Number of authors	1,262	1.79	0.85
Top institution	1,262	0.23	0.42
Fast starter	1,262	0.26	0.44

Notes: The estimation sample contains the 1,262 articles published 2012–2021 in the four core journals with complete citation trajectories (Section 3.4). All variables are defined in Section 3.3.

3.6 Empirical strategy

The core regression is

$$\log(1 + C_i^{LR}) = \alpha + \beta \log(1 + C_i^E) + X_i' \gamma + \mu_j + \tau_y + \varepsilon_i, \quad (2)$$

where C_i^{LR} is long-run citations, C_i^E is age-two citations, X_i are the article controls, and μ_j and τ_y are journal and publication-year fixed effects. The coefficient β is an elasticity: a one percent higher early-citation stock is associated with β percent higher long-run citations. Cross-sectional regressions report heteroskedasticity-robust standard errors; panel regressions cluster by article.

4 Persistence in citation trajectories

4.1 Graphical evidence

Figure 1 presents the raw pattern. Panel (a) plots long-run against early citations on log scales; the relationship is close to log-linear across the full range. Panel (b) plots mean citation trajectories for fast starters and other articles: fast starters are ahead at age two by construction, and the gap widens, rather than closes, between ages two and eight.




Figure 1: Early and long-run citations. Panel (a): age-two against age five-to-eight citations, log scales. Panel (b): mean cumulative citations by age and fast-starter status.

4.2 Regression evidence

Table 2 reports estimates of equation (2). The bivariate elasticity is 0.758. Adding article controls gives 0.717; adding journal and year fixed effects gives the preferred estimate, $\hat{\beta} = 0.784$ (robust SE 0.020); adding topic fixed effects gives 0.757. The estimate therefore lies between 0.717 and 0.784 in every specification.²

The magnitude is large. Doubling the early-citation stock is associated with $2^{0.784} - 1 \approx 72$ percent more long-run citations. In levels: two articles in the same journal and year, one with 10 and one with 20 early citations, are predicted to differ by $0.784 \times [\log 21 - \log 11] \approx 0.51$ log points, or about 66 percent, in long-run citations.³

4.3 Robustness

Six checks support the estimate; details and figures are in Appendix B.

First, the outcome contains the regressor by construction, since long-run citations are cumulative. Re-estimating equation (2) with citations accumulated *strictly after* age two as the outcome (a flow, not a stock; only 1 percent of articles record no such growth) gives an elasticity of 0.671 (robust SE 0.023). Persistence is a property of subsequent citation flows, not an accounting identity. The within-author panel of annual new citations (Appendix D) gives the same answer: the

²The sample falls to 1,025 in columns (4) and (5) because articles in the residual ‘other’ topic category are excluded when topic fixed effects are added. Retaining ‘other’ as its own fixed-effect category keeps all 1,262 articles and gives $\hat{\beta} = 0.770$.

³Percentage changes reported here and below from log-point differences refer to 1+ citations, the quantity actually logged; they approximate percentage changes in raw citations.

Table 2: Early citations and long-run citations

	<i>Dependent variable:</i>				
	Log(long-run citations)				
	(1)	(2)	(3)	(4)	(5)
Log(early citations)	0.758*** (0.021)	0.717*** (0.023)	0.784*** (0.020)	0.757*** (0.021)	
Fast starter					1.018*** (0.046)
N authors		0.012 (0.022)	0.012 (0.018)	0.022 (0.019)	0.082*** (0.025)
Top institution		0.105** (0.041)	0.064* (0.034)	0.087** (0.039)	0.208*** (0.050)
Log(article length)		0.163** (0.065)	0.226** (0.091)	0.087 (0.099)	0.306*** (0.115)
Title length (characters)		−0.0003 (0.001)	0.0001 (0.001)	−0.001* (0.001)	−0.001 (0.001)
Article position in issue		−0.026*** (0.006)	−0.023*** (0.006)	−0.019*** (0.006)	−0.036*** (0.008)
Issue number		−0.011 (0.015)	0.004 (0.014)	0.001 (0.015)	−0.045** (0.019)
Constant	1.887*** (0.044)	1.562*** (0.216)			
Journal FE	No	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Topic FE	No	No	No	Yes	Yes
Observations	1,262	1,262	1,262	1,025	1,025
R ²	0.558	0.572	0.692	0.707	0.476
Adjusted R ²	0.558	0.569	0.687	0.697	0.458

Note:

*p<0.1; **p<0.05; ***p<0.01

Heteroskedasticity-robust standard errors in parentheses.

coefficient of new citations on the lagged citation stock, with author, calendar-year and article-age fixed effects, is 0.55.

Second, the estimate is stable across subsamples: leave-one-journal-out estimates range from 0.766 to 0.792 (Appendix Figure B1), and leave-one-year-out estimates from 0.767 to 0.802. A paper-level bootstrap reproduces the analytical standard error (0.018 against 0.018).

Third, a permutation test reshuffles the fast-starter indicator within journal-year cohorts 1,000 times. The true coefficient (1.084) exceeds every permuted coefficient (empirical $p < 0.001$; Appendix Figure B2), so the persistence result is not an artefact of the within-cohort ranking.

Fourth, the log specification is not driving the result: a Poisson pseudo-maximum-likelihood regression of the citation level on log early citations, with the same controls and fixed effects, gives an elasticity of 0.845 (Appendix B).

Fifth, the elasticity is stable across publication cohorts. Because the long-run outcome mixes age-eight (2012–2018 cohorts) and age-five (2019–2021 cohorts) counts, the cohort comparison holds the outcome age fixed at five years, which the snapshots permit for every cohort. Cohort-specific elasticities range from 0.65 to 0.91 with no cohort statistically distinguishable from the pooled estimate; a linear cohort trend in the elasticity is small and marginally positive (interaction 0.011 per year, robust SE 0.006; Appendix B, Figure B4). If anything, persistence is slightly stronger for more recent cohorts, a period in which online search and repositories expanded; the trend is too imprecise to support a stronger statement.

Sixth, the estimate does not depend on individual controls: dropping the top-institution indicator leaves it at 0.787, and an extended sample that accepts age-three counts as the long-run outcome for 2022–2023 cohorts ($N = 1,518$) gives 0.815.

The elasticity is also similar across article types: it exceeds 0.65 in every title-based topic with at least thirty sample articles (Appendix Figure B3), is 0.75 for solo-authored and 0.80 for co-authored articles (difference not significant), and its interaction with the top-institution indicator is essentially zero. Persistence is a general property of citation dynamics in these journals, not a feature of particular subfields or author groups.

4.4 The fast-starter threshold has no independent content

A binary fast-starter indicator produces a large coefficient on its own: 1.084 with journal and year fixed effects, equivalent to roughly 196 percent more long-run citations. Table 3 shows that this threshold premium is a relabelling of the continuous variable. Adding $\log(1 + C^E)$ to the regression reduces the fast-starter coefficient to 0.084; adding the author-quality controls reduces it to 0.067, statistically indistinguishable from zero. Crossing the cohort threshold has no predictive content beyond the continuous early-citation level itself. The elasticity, not the threshold, is the finding; a double-machine-learning variant of this exercise was explored and set aside because the author-quality controls are missing for 42 percent of the sample and the required imputation destabilises it (Appendix B).

Table 3: Attenuation of the fast-starter threshold

Specification	Coefficient	Robust SE	N
Fast starter, no early-citation control	1.084	0.041	1,262
+ continuous $\log(1 + \text{early citations})$	0.084	0.043	1,262
+ author quality (seniority, h-index, prizes)	0.067	0.056	724

Notes: Dependent variable is log long-run citations; the reported coefficient is on the binary fast-starter indicator. All rows include journal and year fixed effects and article controls. The sample falls to 724 in the last row because the h-index is available for 724 of the 732 RePEc-matched articles.

5 Amplification versus quality sorting

By the definition of Section 2.1, the persistence just documented indicates path dependence only if early reception that is unrelated to fundamentals is amplified. This section reports the paper’s central test in two steps: I first ask how well fundamentals predict early citations, and then whether the unpredictable part propagates.

5.1 Early citations are largely unpredictable from fundamentals

Panel A of Table 4 regresses log early citations on progressively richer sets of predetermined characteristics. Journal, year and topic fixed effects explain 11 percent of the variation. Adding the article controls (team size, top institution, article and title length, within-issue position, issue number) raises the R^2 to 0.222. Adding author seniority, h-index and prize records raises it to 0.243 on the RePEc-matched subsample, where the article controls alone explain 0.214, so the author-quality measures add little. Observable fundamentals – who wrote the article, where they work, what it is about, where and how it was published – account for about one fifth of the variation in early reception. This is not a weakness of the citation data: the same observables explain long-run citations much better (the R^2 in Table 2 exceeds 0.69 once early citations are included). Early reception is simply not well predicted by anything recorded about the article at publication.

5.2 The unexplained component propagates

The second step splits log early citations into the component fitted by the article-level fundamentals of Panel A (journal, year and topic fixed effects and the article controls, estimated on the full sample) and the residual, and regresses long-run citations on both (Panel B of Table 4). Both components predict long-run citations strongly: the explained component with a coefficient of 1.068 and the unexplained component with a coefficient of 0.770. Because the residual is a generated regressor, Panel B also reports bootstrap standard errors that re-estimate both steps on each draw; they are essentially identical to the analytic ones. The same decomposition with post-age-two citation growth as the outcome, which removes any mechanical overlap, gives 1.014 and 0.653. A one-standard-deviation difference in the unexplained component is associated with 85 percent more long-run citations. Figure 2 shows both components against long-run citations in decile means; the

two slopes are of the same order.

A concrete pair from the corpus illustrates the pattern. Two articles appeared in the *Economic History Review* in 2018, both on finance-related topics by the title classification, both single-authored, neither with a top-institution affiliation and of similar length. One had 4 citations at age two and 6 at age eight; the other had 23 and 72. The example establishes nothing about the two articles themselves (their content differs, and readers may judge one more important), but it shows the scale on which early differences among observably similar articles translate into long-run differences. The decomposition quantifies this pattern for the whole sample.

This interpretation has limits. The residual is early reception *not explained by these article-level fundamentals*; it is not measured chance. It contains unobserved quality, novelty, pre-publication circulation, author networks and measurement error, the RePEc author-quality measures that the full-sample first stage does not include (which add little to the predictability of early citations, Section 5.1), and any genuinely contingent component of early attention. The decomposition therefore establishes two facts: most of early reception is unpredictable from what is recorded about an article, and the unpredictable part accounts for most of the long-run elasticity. These facts are what a cumulative-advantage process produces. A pure quality-sorting account requires the unexplained component to be, in effect, a better measure of quality than the observable characteristics that predict early citations – possible, but a demanding requirement, and one that is inconsistent with the visibility evidence of the next section. The data cannot separate unmeasured merit from contingency conclusively, and I do not claim that they can.

Table 4: Predictability and propagation of early citations

<i>Panel A. Predicting log early citations: variance explained</i>			
Journal, year and topic fixed effects		$R^2 = 0.110$	
+ article controls		$R^2 = 0.222$	
+ author quality (seniority, h-index, prizes; $N = 724$)		$R^2 = 0.243$	
<i>Panel B. Long-run citations on the two components of early citations</i>			
	Coefficient	Robust SE	Bootstrap SE
Explained component (fundamentals)	1.068	0.054	0.060
Unexplained component (residual)	0.770	0.020	0.020
<i>outcome: citation growth after age two</i>			
Explained component	1.014	0.063	0.071
Unexplained component	0.653	0.024	0.024

Notes: $N = 1,262$. Panel A reports $1 - \text{SSR}/\text{SST}$ from fixed-effects regressions of log early citations; the last row uses the RePEc-matched subsample. Panel B regresses log long-run citations (upper rows) and log post-age-two citation growth (lower rows) on the fitted and residual components from the Panel A article-controls specification, with journal and year fixed effects. Bootstrap standard errors re-estimate both steps on 999 article-level resamples.




Figure 2: Explained and unexplained components of early citations against long-run citations (centred decile means; both axes in logs, so the plotted slopes approximate the elasticities of Table 4, Panel B; the figure plots decile means and does not partial out the journal and year fixed effects).

6 Channel-level evidence

This section reports the evidence for each testable channel set out in Section 2.3.

6.1 Visibility: within-issue position

Within a journal issue, articles that appear earlier are more visible in tables of contents and browsing, and lead articles are sometimes flagged by the journal itself. Position within the issue is therefore a visibility difference that is at most weakly related to content. Two reduced-form facts follow the visibility channel’s prediction. Conditional on journal and year, each one-position move towards the back of an issue is associated with 5.4 percent fewer early citations and 6.5 percent fewer long-run citations (Table C1). Both gradients are essentially unchanged when the conditioning is tightened to journal-by-year and then to within-issue fixed effects, and the long-run gradient is negative in all four journals (Appendix C).

Position is not, however, randomly assigned, and I do not use it as an instrument. Editors set running order, and a falsification battery in the spirit of Danieli et al. (2026) and Chyn, Frandsen and Leslie (2025) shows that predetermined author-quality proxies are jointly related to position ($F = 6.3$, $p = 0.012$): articles by higher h-index and prize-winning authors are placed slightly earlier. In addition, position retains a small but statistically significant association with long-run citations conditional on early citations (Appendix C), so even under conditional random assignment an exclusion restriction would fail. For these two reasons no instrumental-variables estimate is reported; Appendix C documents the tests. Two features nevertheless make the reduced form informative about visibility: conditioning on the observable quality proxies leaves the gradients unchanged, and other candidate visibility measures with no editorial discretion (the size of the issue, the alphabetical rank of the first author’s surname) have no explanatory power for early citations, which is consistent with the position gradient reflecting placement rather than author characteristics.

The other candidate measures also reinforce Section 5.1: early reception is largely unpredictable from observables. Taken together, the position gradients provide reduced-form support for the visibility channel, but they do not identify a causal elasticity from visibility to long-run citations.

6.2 Formats and audiences

Classifying the 37,853 works that cite a corpus article by their OpenAlex discipline (72,563 of 72,695 citation links could be classified; Appendix A) shows where the field’s citations originate (Figure 3, panel a). Economics supplies 50 percent of citation links; other social sciences (sociology, political science, demography, geography) supply 30 percent; business and management 5 percent. Counting any citing work that has economics or history among its root disciplinary tags, 78 percent of citations involve the field or its immediate neighbours, and about one citation in nine comes from a distant domain (health and life sciences 6 percent, other sciences and engineering 5 percent). Economic history is read mainly by economists and neighbouring social scientists.

Panel (b) classifies the same links by document type: 59 percent of citations come from journal articles, 28 percent from books and book chapters, 8 percent from preprints and working papers (SSRN, NBER and RePEc are among the most frequent citing venues), and the remainder from dissertations and reports. Two fifths of the field’s citation impact therefore arrives through formats outside the journal system, which supports the dissemination-formats channel and explains why the broad coverage of Google Scholar is appropriate for this corpus.

Early success does not change the composition of a paper’s audience. On the subsample of 597 estimation-sample articles with classified citers, the cross-field share of citations is statistically unrelated to fast-starter status, and early citations predict within-field citation counts somewhat more strongly than cross-field counts (elasticities 0.69 and 0.50). Whatever reinforcement operates here operates within economic history and its neighbours; it does not export early successes to distant fields.




Figure 3: Sources of citations to economic history articles. Panel (a): share of citation links by the citing work’s discipline. Panel (b): share by document type. Based on 72,563 classified citation links from 37,853 citing works.

6.3 Conferences

Di Vaio, Waldenström and Weisdorf (2012) report that conference and workshop presentation predicts higher citations for economic history articles. In these data, using hand-verified programmes and the author-validated match of Section 3.3, 85 estimation-sample articles are linked to an EHA or EHS presentation, and the association is small and unreliable. The cross-sectional premium is

0.07 log points (robust SE 0.056; permutation $p = 0.16$). The within-author panel gives the same answer more precisely: comparing articles by the same author at the same article age, presented articles receive no more annual citations than the author’s other articles (coefficient 0.017; Table D1). A stricter specification with author-by-publication-year fixed effects produces a larger coefficient, but it is identified from only 13 author-years in which the same author has both presented and non-presented articles, and Appendix D explains why it is not relied upon. Presenters are positively selected on predetermined characteristics (they are more often at top institutions), which is the most plausible interpretation of the small positive cross-sectional coefficient (Appendix E).

Three qualifications bound this null. The sample of 85 presented articles can rule out large premia but not detect small ones; the estimates are consistent with a premium of a few percent. The variable covers the two annual association meetings only, so the null does not speak to workshops, seminar series or the EHES biennial meeting. And a sharper test based on session timing within the EHA programme is infeasible on power grounds: too few matched articles fall in the relevant time slots (Appendix E). Subject to those limits, the data provide no support for the conference channel.

6.4 Networks, flows and self-citation

The reinforcement model of Section 2.2 operates through flows: the post-age-two growth elasticity of 0.671 (Section 4.3) and the within-author panel coefficient of 0.55 on the lagged citation stock (Appendix D) both show that a larger existing stock predicts larger subsequent inflows, conditional on age and author. Network-position measures add little beyond this. Point estimates go in the direction cumulative advantage predicts – fast starters have deeper citation cascades (coefficient 0.19, SE 0.14) and more indirect citers (0.15, SE 0.09), conditional on early citations – but neither is statistically significant, and the temporal concentration of citations (a Herfindahl index across citing years) shows no difference at all (0.011, SE 0.013). The flow evidence supports reinforcement while the network-position estimates are too imprecise to add independent support.

Self-citation does not drive the results. Among within-corpus citation links, 16 percent are self-citations by the matching used here; removing them changes the early-citation coefficient by less than half a percent.

7 Interpretation and conclusion

The evidence establishes three findings. First, early citations strongly predict long-run citations in economic history’s four core journals: an elasticity of about 0.78 conditional on journal and publication year, present in citation flows after the early window, stable across journals, cohorts, topics and estimation methods. Second, early citations are largely unpredictable: observable author and article fundamentals explain about one fifth of their variation, and the unexplained four fifths predict long-run outcomes almost as strongly as the explained fifth. Third, articles with a within-issue placement advantage (a visibility difference at most weakly related to content) receive more

citations both early and in the long run, and the association survives conditioning on observable author quality.

What these findings support, rather than establish, is path dependence in the sense of Section 2.1. The unexplained component of early reception contains unobserved quality as well as contingency, and no test in this paper separates the two conclusively. The interpretation I find most consistent with the evidence is that the citation process in economic history contains a substantial cumulative-advantage component: a quality-only account would have to attribute nearly the full long-run elasticity to quality variation invisible in seniority, h-index, prizes, affiliation, topic, length and placement, and would also need to explain why placement itself, which does not change content, predicts long-run citations. The findings do not imply that the field’s canon is wrong or that inferior work is locked in; they show that early reception, much of it unexplained by fundamentals, durably predicts which contributions accumulate influence.

Two implications follow for research evaluation. First, early citation counts conflate reception with fundamentals: because unexplained differences in early reception persist and widen, a count observed a few years after publication partly measures how early attention arrived rather than assessed quality, and normalised or distribution-aware measures (Hamermesh, 2018; Bornmann and Wohlrabe, 2019; D’Ippoliti, 2021) are preferable to raw early counts. Second, evaluation windows matter: assessments made at short horizons embed early-visibility differences that the subsequent citation process does not correct.

The low predictability of early reception also has implications for authors and editors, within the same limits. For authors, measured credentials – seniority, h-index, prizes, affiliation – neither guarantee nor preclude wide reception. Whether the unexplained component reflects judgement about which questions matter, timing, exposition or chance is what these data cannot distinguish; what they do show is that it is not reducible to status. For editors and referees, the corresponding fact is that which articles will be widely cited is largely not predictable from what is observable at publication. One reading of this is welcome: the allocation of attention in these journals is not controlled by the measured status of authors, so influential work can, and does, come from outside the field’s established centres. The other reading repeats the paper’s caution: if early reception cannot be predicted from fundamentals, it should not be presumed to measure quality, and whatever advantage arises early persists.

The same reasoning is relevant to the formation of the field’s canon. Canonisation itself is not observed in these data: the teaching channel of Section 2.3 is stated but not tested, and the outcome window ends at age eight. What the results indicate is a route through which the canon can form: if early reception compounds, then the set of contributions the field treats as central partly records how attention was allocated in their first years. For editors, referees and hiring committees, the corresponding reading of an early citation count is that it records reception as well as fundamentals, which is a further reason to prefer the normalised measures discussed above.

The results raise more questions than they settle, and several are tractable. The central one is what generates the unexplained component of early reception. Cleaner exogenous variation in early

visibility, such as editorial placement rules or changes in search and recommendation systems, would permit causal estimates of the amplification these data suggest; parallel corpora for neighbouring fields would show whether the dynamics are specific to economic history.

The study is related to the concept of path dependence itself in three ways. First, it provides evidence consistent with cumulative advantage in the production of knowledge, a setting different from the technology-adoption cases in which the concept was formulated. Second, it indicates which form of the concept is useful in this setting: not lock-in, but the weaker form in which early events influence without determining the outcome (Brenner and zu Jeddelloh, 2024; Roos, 2024). The analysis shows how that weaker form can be tested against mere persistence, by conditioning on observable fundamentals and following the unexplained component, while the separation of contingency from unobserved quality remains open. Third, scholarly communication is an informative setting in which to study self-reinforcing processes: the analyst observes dense annual outcome trajectories for a defined corpus, together with candidate fundamentals, at a completeness that is rarely available for technology adoption or institutional change.

Using the complete citation histories of the field’s four core journals, this paper establishes strong persistence, shows that observable fundamentals account for about a fifth of early reception, and shows that the unexplained remainder durably predicts which contributions accumulate influence. The findings are consistent with early reception contributing to later influence, rather than serving only as an early signal of it. For a field that studies persistence and contingency, its own publication record is evidence worth analysing.

A Data construction

A.1 Sample attrition

Table A1 documents the construction of the estimation sample from the corpus.

Table A1: Sample construction

Step	Papers	Dropped	Publication years
Four core journals (JEH, EHR, EEH, EREH), 1997-2025	3,250	–	1997-2025
Age-2 citations observable (published 2012 or later)	1,672	1,578	2012-2024
Age-5 or age-8 citations observable (published 2021 or earlier)	1,262	410	2012-2021
Author count non-missing	1,262	0	2012-2021
Citation counts non-negative	1,262	0	2012-2021

Notes: Age-2 availability requires publication in 2012 or later (first snapshot 2014); age-5/8 availability requires publication in 2021 or earlier (last snapshot 2026).

A.2 Citation snapshots

The author and several research assistants recorded each article’s cumulative Google Scholar citation count once a year in February–March, 2014 through 2026. The count at age k is the count in the snapshot of calendar year (publication year + k). Counts occasionally decline between snapshots because Google Scholar re-indexes sources; negative implied growth is set to zero where a growth measure is used. The 2026 snapshot was collected in February–March 2026.

A.3 OpenAlex linkage

Articles were matched to OpenAlex works by title search with year verification: 3,241 of 3,250 corpus articles (99.7 percent) matched, including 99.4 percent of the estimation sample. Among matched articles, the correlation between the OpenAlex citation count (2026 download) and the Google Scholar count (2025 snapshot) is 0.71 in logs and 0.78 in ranks ($N = 3,088$ articles with both counts). All citing works of matched articles were retrieved (72,695 citation links), of which 72,563 could be classified by discipline using the citing work’s OpenAlex primary field and root-level concept tags; the remaining 132 links lack the required metadata. RePEc author data match 732 of the 1,262 estimation-sample articles (58 percent); the h-index component is available for 724 of the 732, which is why some author-quality specifications report $N = 724$. The main elasticity on the matched subsample is 0.800 (robust SE 0.027), against 0.784 in the full sample, so RePEc coverage does not select on the relationship of interest.

A.4 Top institutions

The top-institution indicator equals one if any of the first five authors’ affiliation strings contains one of: Harvard, MIT, Stanford, Berkeley, Yale, Princeton, Chicago, Northwestern, Columbia, Penn, UCLA, Michigan, NYU, Oxford, Cambridge, the London School of Economics (matched

by both the full name and the abbreviation ‘LSE’) and Warwick. This is a fixed, ad hoc set of seventeen large US and UK research universities chosen for visibility in the corpus (US and UK institutions dominate authorship in these journals); it is a control variable, not a ranking exercise. The main estimate is unchanged when the control is dropped (Section 4.3).

A.5 Topic dictionary

Table A2 reproduces the keyword dictionary. A title is assigned to the topic whose keywords it matches most often; ties are resolved by dictionary order; titles matching no keyword are ‘other’.

Table A2: Topic classification dictionary

Topic	Keywords (matched anywhere in the title)	Papers
agriculture	agricultur, farm, crop, land, rural, grain	136
labor	labor, labour, wage, worker, employment, education, skill	117
finance	bank, credit, debt, finance, money, monetary, currency	114
industry	industr, manufactur, factory, technology, innovation	92
institutions	institution, state, government, politic, law, democracy	92
trade	trade, export, import, tariff, globalization, commerce	82
demography	population, demographic, fertility, migration, family, gender, women, woman, marriage, marital	61
war	war, warfare, military, conflict, revolution	60
growth	growth, gdp, development, productivity, divergence, convergence, national income, national product, per capita	54
inequality	inequality, living standard, welfare, poverty, income, wealth	51
financial markets	stock, equity, bond, securit, share, investor, insurance, financial market, capital market, exchange	41
colonial	colonial, colony, empire, imperial	39
crises	depression, crisis, crises, recession, boom, recovery, business cycle, panic, bubble, default, deflation, inflation	31
health	health, mortality, disease, epidemic, pandemic, nutrition, height, stature, anthropometr, numeracy, heaping, biological, morbidity, medical, hospital, sanitat	29
slavery	slave, slavery, enslaved, abolition	13
urban	urban, city, cities, agglomeration	13
other	no keyword matched	237

Notes: Keyword stems are matched anywhere in the lower-cased title (so ‘industr’ matches ‘industrial’ and ‘industrialization’). Paper counts refer to the estimation sample.

B Additional robustness

Cohort stability. Figure B4 plots the cohort-specific estimates underlying Section 4.3: with the outcome fixed at age-five citations, the elasticity ranges from 0.65 (2015 cohort) to 0.91 (2016 cohort) with overlapping confidence intervals. The pooled linear trend in the elasticity is 0.011 per cohort year (robust SE 0.006); splitting the period in two gives 0.79 (2012–2016) against 0.86 (2017–2021), a difference of 0.067 (robust SE 0.034). The evidence indicates stability, with weak signs of a mild strengthening in recent cohorts.




Figure B1: Leave-one-journal-out estimates of the early-citation elasticity.




Figure B2: Permutation test: fast-starter status reshuffled within journal-year cohorts 1,000 times. The vertical line marks the true estimate.




Figure B3: Early-citation elasticity by topic (dashed line: pooled estimate).




Figure B4: Early-citation elasticity by publication cohort, outcome fixed at age-five citations for all cohorts (95 percent confidence intervals; dashed line: sample-size-weighted mean).

Functional form. A Poisson pseudo-maximum-likelihood regression of the level of long-run citations on log early citations with the full controls and fixed effects gives 0.845 (robust SE 0.035), close to the log-log estimate.

Topic fixed effects with the full sample. Retaining ‘other’ as a topic category keeps $N = 1,262$ and gives 0.770 (0.020), against 0.757 (0.021) when ‘other’ articles are dropped.

Extended sample. Accepting age-three counts as the long-run outcome for 2022–2023 cohorts adds 256 articles ($N = 1,518$; 879 articles contribute age-eight, 383 age-five and 256 age-three outcomes) and gives 0.815 (0.017).

Double machine learning. A partially linear DML estimator (random-forest nuisance functions, five-fold cross-fitting) for the fast-starter effect was explored and set aside: the author-quality controls are missing for 42 percent of the estimation sample and the OpenAlex topic side-data do not cover the 2012–2021 cohorts, so the required median imputation destabilises the nuisance estimates and produces a coefficient well outside the range of every OLS specification. The transparent attenuation sequence of Table 3 is reported instead.

C Within-issue position: design checks

Section 6.1 reports the position gradients in reduced form only. This appendix documents why no instrumental-variables estimate is constructed from them, following the negative-control approach of Danieli et al. (2026) and the examiner-design checklist of Chyn, Frandsen and Leslie (2025).

Panel A shows both gradients are stable as the conditioning tightens from journal and year to journal-by-year to within-issue fixed effects, so they are not artefacts of composition across issues. Panel B shows why position is nevertheless not usable as an instrument: predetermined quality proxies are jointly related to position ($F = 6.3$, $p = 0.012$), with higher h-index and prize-winning authors placed slightly earlier; conditional random assignment is rejected. In addition, position retains an association with long-run citations conditional on early citations (-0.023 , robust SE 0.006), consistent with lasting visibility of prominent slots, so the exclusion restriction required by an instrument would fail even under random assignment. Conditioning the gradients on the observable quality proxies leaves them essentially unchanged, and Panel C shows the long-run gradient is negative in all four journals. The gradients are therefore informative reduced-form evidence of a visibility channel, and nothing more is claimed of them. Among alternative visibility measures, issue size and the first author’s alphabetical rank have no power for early citations ($F = 0.03$ and 0.2); a lead-article indicator behaves like position itself.

D Within-author panel

The annual snapshots permit a panel of new citations per article per year; after removing article-years with negative recorded inflows and articles without usable author names, the estimation sample contains 26,999 article-year observations. Table D1 regresses log new citations on the log

Table C1: Within-issue position: gradients and design checks

<i>Panel A. Position gradients under tighter conditioning (N = 1,262)</i>		
	Position → log early	Position → log long-run
Journal and year fixed effects	−0.056 (0.010)	−0.067 (0.010)
Journal × year fixed effects	−0.056 (0.010)	−0.066 (0.010)
Within-issue fixed effects	−0.060 (0.010)	−0.072 (0.010)
<i>Panel B. Negative-control tests (quality proxies on position, journal × year FE)</i>		
	Coefficient	Robust SE
Team maximum h-index	−0.347	0.113
Team maximum seniority	0.134	0.907
Top institution	−0.008	0.005
Article prize	−0.002	0.001
Dissertation prize	−0.002	0.002
Joint test	$F = 6.3, p = 0.012$ ($N = 724$)	
<i>Panel C. Long-run position gradient by journal</i>		
	Coefficient	Robust SE
<i>Journal of Economic History</i>	−0.054	0.021
<i>Economic History Review</i>	−0.067	0.012
<i>European Review of Economic History</i>	−0.040	0.034
<i>Explorations in Economic History</i>	−0.078	0.032
Exclusion check: position → long-run early	−0.023	0.006

Notes: Heteroskedasticity-robust standard errors in parentheses. Panel B regresses each predetermined quality proxy on position within journal $\times$ year cells; the joint test regresses position on all proxies and tests their joint significance. The exclusion check conditions the long-run gradient on log early citations (journal $\times$ year fixed effects).

lagged citation stock and the conference indicator, with author, calendar-year and article-age fixed effects (column 1) and author-by-publication-year fixed effects (column 2), clustering by article. The article-age fixed effects absorb a common age profile in citation flows (the empirical counterpart of λ_t in equation (1)); without them, differences in article age would enter the estimates as if they were treatment effects. The lagged-stock coefficient of 0.55 (cluster SE 0.012) shows strong persistence in citation flows within an author's own portfolio; because it conditions on the author rather than the article, it does not control for article-specific quality, and is a reduced-form measure of persistence rather than an estimate of θ . The conference coefficient is 0.017 (0.034): within author and age, presented articles do not accumulate more citations than the author's other articles. The stricter column-2 specification gives a conference coefficient of 0.545 (cluster SE 0.171), but it is identified from only 13 author-years in which the same author has both presented and non-presented articles; with so few effective comparisons the estimate is reported for completeness and no conclusion depends on it. Column 2's sample is larger by one observation because it omits the author-count control, which is missing for one article-year.

Table D1: Within-author comparisons

	<i>Dependent variable:</i>	
	Log(new citations)	
	(1)	(2)
Presented at conference	0.017 (0.034)	0.545*** (0.171)
Log(lagged citation stock)	0.550*** (0.012)	0.099*** (0.037)
Number of authors	0.027* (0.014)	
Author FE	Yes	No
Author-by-year FE	No	Yes
Calendar-year FE	Yes	Yes
Article-age FE	Yes	Yes
Observations	26,999	27,000
R ²	0.583	0.642
Adjusted R ²	0.552	0.598
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

Notes: Standard errors clustered by article. Column (2) omits the author-count control, which is missing for one article-year, and is identified from the 13 author-years with variation in presentation status.

E Conference measurement, selection and timing

Matching. Conference entries are linked to corpus articles by Jaro–Winkler title distance in three tiers: distance below 0.10 accepted outright; distance below 0.25 accepted when a presenter surname matches an article author; distance below 0.15 accepted for entries without usable author information. This recovers retitled papers (spelling variants, subtitle changes) without admitting topical false positives; 146 conference entries match, giving 141 distinct articles (43 via EHA and 102 via EHS, four of them presented at both meetings), of which 85 are in the estimation sample.

Selection. Presenters differ from non-presenters in the expected direction: more early citations (9.3 against 7.0, $p = 0.05$), more top-institution affiliation (0.32 against 0.22, $p = 0.07$), more fast starters (0.36 against 0.25, $p = 0.03$), and fewer co-authors (1.5 against 1.8, $p = 0.003$). Only the institutional difference is unambiguously predetermined, since a presentation near publication can itself raise age-two citations. A probit of presentation on citations gives the same picture: fast-starter status predicts presenting (coefficient 0.338, SE 0.122), and long-run citations add little conditional on early citations (0.199, SE 0.115, $p = 0.08$). The small positive cross-sectional premium of Section 6.3 is therefore best read as selection. The permutation distribution of the conference coefficient is in Figure E1.

Session timing. The hand-transcribed EHA programmes include session start times, which would in principle support a comparison of adjacent time slots. The design is infeasible here: of the EHA-matched estimation-sample articles, 11 fall in the last pre-lunch slot and 3 in the first post-lunch slot. The constraint is the rate at which EHA papers subsequently appear in these four journals, not the programme data.

Data availability

The hand-coded journal and conference data, all derived data files needed to reproduce every result offline, and the full analysis code are available in a public replication package at <https://github.com/johanfourieza/research/tree/main/2026/path>. The package includes a README with reproduction instructions and a variable-level codebook.

References

- Aistleitner, Matthias, Jakob Kapeller, and Stefan Steinerberger. 2019. “Citation Patterns in Economics and Beyond.” *Science in Context*, 32(4): 361–380.
- Arthur, W. Brian. 1989. “Competing Technologies, Increasing Returns, and Lock-In by Historical Events.” *Economic Journal*, 99(394): 116–131.
- Bleakley, Hoyt, and Jeffrey Lin. 2012. “Portage and Path Dependence.” *Quarterly Journal of Economics*, 127(2): 587–644.




Figure E1: Permutation test: conference presentation reshuffled within journal-year cohorts 1,000 times. The vertical line marks the true estimate.

Bornmann, Lutz, and Klaus Wohlrabe. 2019. “Normalisation of Citation Impact in Economics.” *Scientometrics*, 120(2): 841–884.

Brenner, Thomas, and Sonja zu Jeddeloh. 2024. “Path Dependence in an Evolving System: A Modeling Perspective.” *Cliometrica*, 18(1): 1–36.

Chyn, Eric, Brigham Frandsen, and Emily Leslie. 2025. “Examiner and Judge Designs in Economics: A Practitioner’s Guide.” *Journal of Economic Literature*, 63(2): 401–439.

Cioni, Martina, Giovanni Federico, and Michelangelo Vasta. 2020. “The Long-Term Evolution of Economic History: Evidence from the Top Five Field Journals (1927–2017).” *Cliometrica*, 14(1): 1–39.

Cioni, Martina, Giovanni Federico, and Michelangelo Vasta. 2022. “Persistence Studies: A New Kind of Economic History?” *Review of Regional Research*, 42(3): 227–248.

Conrad, Alfred H., and John R. Meyer. 1958. “The Economics of Slavery in the Ante Bellum South.” *Journal of Political Economy*, 66(2): 95–130.

Danieli, Oren, Daniel Nevo, Itai Walk, Bar Weinstein, and Dan Zeltzer. 2026. “Negative Control Falsification Tests for Instrumental Variable Designs.” *American Economic Review*, 116(4): 1380–1414.

- David, Paul A.** 1985. “Clio and the Economics of QWERTY.” *American Economic Review*, 75(2): 332–337.
- David, Paul A.** 2007. “Path Dependence: A Foundational Concept for Historical Social Science.” *Cliometrica*, 1(2): 91–114.
- D’Ippoliti, Carlo.** 2021. ““Many-Citedness”: Citations Measure More Than Just Scientific Quality.” *Journal of Economic Surveys*, 35(5): 1271–1301.
- Di Vaio, Gianfranco, and Jacob Louis Weisdorf.** 2010. “Ranking Economic History Journals: A Citation-Based Impact-Adjusted Analysis.” *Cliometrica*, 4(1): 1–17.
- Di Vaio, Gianfranco, Daniel Waldenström, and Jacob Weisdorf.** 2012. “Citation Success: Evidence from Economic History Journal Publications.” *Explorations in Economic History*, 49(1): 92–104.
- Fogel, Robert W., and Stanley L. Engerman.** 1974. *Time on the Cross: The Economics of American Negro Slavery*. Boston: Little, Brown.
- Hamermesh, Daniel S.** 2018. “Citations in Economics: Measurement, Uses, and Impacts.” *Journal of Economic Literature*, 56(1): 115–156.
- Heckman, James J., and Sidharth Moktan.** 2020. “Publishing and Promotion in Economics: The Tyranny of the Top Five.” *Journal of Economic Literature*, 58(2): 419–470.
- Liebowitz, Stan J., and Stephen E. Margolis.** 1990. “The Fable of the Keys.” *Journal of Law and Economics*, 33(1): 1–25.
- Liebowitz, Stan J., and Stephen E. Margolis.** 1995. “Path Dependence, Lock-In, and History.” *Journal of Law, Economics, and Organization*, 11(1): 205–226.
- Mahoney, James.** 2006. “Analyzing Path Dependence: Lessons from the Social Sciences.” In *Understanding Change: Models, Methodologies, and Metaphors.*, ed. Andreas Wimmer and Reinhart Kössler, 129–139. London: Palgrave Macmillan.
- Merton, Robert K.** 1968. “The Matthew Effect in Science.” *Science*, 159(3810): 56–63.
- North, Douglass C., and Barry R. Weingast.** 1989. “Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England.” *Journal of Economic History*, 49(4): 803–832.
- Price, Derek J. de Solla.** 1976. “A General Theory of Bibliometric and Other Cumulative Advantage Processes.” *Journal of the American Society for Information Science*, 27(5): 292–306.
- Puffert, Douglas J.** 2002. “Path Dependence in Spatial Networks: The Standardization of Railway Track Gauge.” *Explorations in Economic History*, 39(3): 282–314.

- Puffert, Douglas J.** 2024. “Path Dependence.” In *Handbook of Cliometrics.* , ed. Claude Diebolt and Michael Hauptert. Cham:Springer.
- Roos, Michael.** 2024. “History and Path Dependence.” In *Principles of Complexity Economics: Concepts, Methods and Applications.* 581–645. Cham:Springer Nature.
- Seltzer, Andrew J., and Daniel S. Hamermesh.** 2018. “Co-Authorship in Economic History and Economics: Are We Any Different?” *Explorations in Economic History*, 69: 102–109.
- Sussman, Nathan, and Yishay Yafeh.** 2006. “Institutional Reforms, Financial Development and Sovereign Debt: Britain 1690–1790.” *Journal of Economic History*, 66(4): 906–935.