

Online Appendix

Scale, Efficiency, and the Geography of Knowledge Production: A Cross-Country Panel Analysis

This appendix contains five sections, each providing extended treatment of material referenced in the main text.

- **Appendix A:** Extended literature surveys across four subsections: A.1 (Empirical KPF Literature), A.2 (Talent Migration as a Leadership Transmission Mechanism), A.3 (China Quality Dynamics: the Japan–South Korea Precedent), and A.4 (Historical Geography of Scientific Leadership — full five-mechanism treatment with complete literature citations).
- **Appendix B:** Econometric motivation for AMG and CCEMG estimators, Monte Carlo evidence on estimator performance, and country-specific KPF coefficient scatter plots for all 13 focus economies.
- **Appendix C:** Country-specific vs panel-mean projection comparison, with convergence scenario analysis.
- **Appendix D:** Rest-of-World Knowledge Stock robustness check (national capital vs global public good hypotheses).
- **Appendix E:** Westerlund (2007) Panel Cointegration Tests.

Note for reviewers: this appendix is designed to be submitted alongside the main paper. All figures and references are self-contained. Section numbers prefixed with a digit (e.g., “Section 4.2 of the main paper”) refer to the main manuscript.

Appendix A: Knowledge Production — Extended Literature Surveys and Historical Analysis

A.1 Extended Empirical Literature Survey

The empirical KPF literature originates with Griliches (1990)’s canonical survey establishing patents as the primary knowledge-output proxy and Griliches and Mairesse (1998)’s influential extensions to the production function identification problem. At the firm level, Pakes and Griliches (1984) and Jaffe (1986) documented positive R&D-to-patent elasticities and the importance of technological spillovers; Jaffe et al. (1993) extended the spillover finding geographically using patent citation data. Hall et al. (2001) developed the NBER patent citation data file that has become the standard resource for measuring knowledge flows between firms and sectors.

Cross-country aggregate KPF estimation was brought into the growth-accounting framework by Guellec and van Pottelsberghe de la Potterie (2001, 2004), who found that a 1% increase in business R&D raises long-run productivity by approximately 0.13%, with public and higher-education R&D generating somewhat smaller but positive returns. Ülkü (2004) extended this to a broader sample including non-OECD economies, documenting substantially lower innovation efficiency outside the OECD. Madsen (2008) provided the most systematic Romer–Jones test in OECD historical panels, finding strong Schumpeterian estimates ($\hat{\phi} > 1$) that contrast with Jones’s sub-proportional results; Fedderke and Liu (2017) replicate this pattern for South Africa.

Recent evidence suggests declining research productivity at the innovation frontier. Bloom et al. (2020) document that the number of researchers required to generate a given rate of economic growth has risen sharply, consistent with the Jones sub-proportional regime at the frontier and with the decreasing-returns-to-research implication of the semi-endogenous model. The present paper focuses on the *academic* KPF (basic science to publications) rather than the *aggregate innovation* KPF (R&D to patents/productivity), following Mokyr (2002)’s λ/Ω distinction between practical and epistemic knowledge. On the measurement side, bibliometric approaches to quantifying knowledge output are developed in Narin et al. (1976) and extended by Waltman et al. (2015); the choice between whole- and fractional-count attribution matters significantly in multi-author, internationally collaborative science (Larue et al., 2022).

A.2 Talent Migration as a Scientific Leadership Transmission Mechanism

Talent migration is historically the fastest channel through which one country’s scientific position has been altered in a short period, because it operates simultaneously through

the human capital and knowledge stock channels: each émigré scientist carries their specialisation’s knowledge base (raising the receiving country’s effective knowledge stock) while also adding directly to the receiving country’s research workforce. It is therefore not merely an HC input change but a (σ, ϕ) joint shock, amplified by the absorptive-capacity complementarity between incoming talent and the existing domestic knowledge stock.

The causal evidence from Moser et al. (2014) is the cleanest available: approximately 1,500 émigré scientists who fled Germany after 1933 increased US innovation in their fields by over 30 per cent within a decade. The mechanism identified is the attraction of American talent into émigré specialisations — a knowledge spillover operating through the $\hat{\phi}$ channel as the existing US knowledge base absorbed and amplified the incoming expertise. Critically, this amplification required the receiving country’s knowledge base to be sufficiently deep: a shallower science system would not have been able to sustain the chain reaction. Borjas and Doran (2012) document the parallel dynamic following the post-1991 outflow of Soviet mathematicians to the United States, finding aggregate increases in mathematical output alongside competitive displacement of some incumbent American researchers.

Heinze et al. (2019) and von der Heyden et al. (2025) trace how the German–American transition was consolidated through the 1940s and 1950s: the forced emigration was the exogenous accelerant, but the endogenous driver was the prior maturation of the American research university system to the point where it had the absorptive capacity to deploy the incoming talent productively.

The contemporary relevance of this mechanism is significant for interpreting the quality-weighted projections in Section 5.7 of the main paper. The projections condition on talent migration remaining at in-sample levels. If the documented post-2022 outflows of researchers from the United States to Europe (American Physical Society, 2026) represent a structural acceleration of emigration, the quality-channel projections understate the US decline, while the European recipients of this talent would experience knowledge-stock gains not captured in the baseline. Conversely, if China continues to invest in its Thousand Talents and Hundred Talents return-migration programmes (Fu, 2015), the domestic knowledge base available to amplify returning researchers may grow faster than the panel-average trend projects, potentially accelerating quality catch-up relative to the Section 5.7 trajectory.

A.3 Knowledge Spillovers, Global Frontier Governance, and Estimator Motivation

The global science frontier is actively constituted and maintained by international scientific institutions: learned societies, national academies, international research infrastruc-

tures (CERN, ESO, ITER), and collaborative networks whose membership and standards transcend national boundaries. These institutions define what counts as frontier research and which methodological standards constitute acceptable evidence. The cross-national character of the frontier, not merely an emergent property of national systems but an actively governed structure, motivates the common-factor treatment in AMG and CCEMG: these estimators model the global frontier process as a common dynamic factor with heterogeneous country loadings, precisely the structure implied by transnational scientific governance.

Coe and Helpman (1995) documented significant R&D spillovers across bilateral trading partners; Coe et al. (2009) confirmed and extended these results. Under the *national capital hypothesis*, the own knowledge stock $K_{i,t-1}$ governs output through absorptive capacity (Cohen & Levinthal, 1990). Under the *global public good hypothesis*, peer-reviewed science is universally accessible via open-access journals and preprint servers, so that the world stock K_{t-1}^{world} is the relevant base. In practice, both channels operate simultaneously; their empirical separation is the object of the rest-of-world exercise in Appendix D.

A.4 China Quality Dynamics: Japan–South Korea Precedent

Zhou and Leydesdorff (2006) documented that China’s internationally visible publication output grew rapidly from the 1990s but that China lagged Japan and South Korea by roughly a decade in the transition from volume growth to citation quality improvement. Both earlier catch-up economies achieved sustained increases in citation impact following their initial quantity catch-up, a trajectory that, if China follows it, would imply quality catch-up around 2030–2040. Hu and Jefferson (2009) raised concerns about the quality of China’s innovation metrics, finding that a large fraction of the near-tenfold increase in domestic patent filings between 1995 and 2005 was driven by policy incentives rather than genuine inventive activity. Meng and Romanowich (2021) find convergence in quantitative but not quality-weighted indicators, consistent with input accumulation preceding efficiency catch-up.

A.5 Historical Geography of Scientific Leadership (extended)

A.5.1 The Succession Pattern

The question of why one country displaces another at the scientific frontier is as old as modern science. The record of the past four centuries traces a recognisable pattern of sequential dominance punctuated by sharp transitions: English natural philosophy in the seventeenth century, French institutional science in the eighteenth, German research universities in the nineteenth, and American dominance from the mid-twentieth century onward (Ben-David, 1971). This succession is not random: in each episode, the chal-

lenger accumulated a sufficient scientific base before displacing the incumbent, and the displacement mechanism operated through identifiable channels rather than sudden shifts in funding or effort.

Heinze et al. (2019), tracing the full population of Nobel laureates in science from 1901 to 2017, document that North America’s rise as a global scientific power began in the 1920s, accelerated sharply in the 1930s and 1940s, and consolidated into unambiguous hegemony by the 1970s. von der Heyden et al. (2025), using Nobel Prize nomination data, show that Germany dominated the process through the 1930s before American scientists supplanted them, a transition Nelson and Wright (1992) locate at the simultaneous maturation of the American research university system and the catastrophic disruption of the German science system. The Germany-to-United States transition is the historical analogue most directly relevant to the current China-United States comparison, and several of its features are instructive.

First, the transition was preceded by a decade-long period of quantitative catch-up in which the United States approached German publication volumes before surpassing Germany in citation impact. Nelson and Wright (1992) document this sequence: the US achieved rough publication parity with Germany by the 1920s but did not displace German quality leadership until the institutional consolidation of the 1940s and 1950s. This quantitative-before- qualitative sequence is precisely what our analysis finds for China today.

Second, the transition involved both endogenous and exogenous accelerants. The endogenous driver was the maturing of the American research university system, funded by the Morrill Act land grants and the post-1890 wave of university founding. The exogenous accelerant was the forced emigration of Jewish academics from Nazi Germany after 1933. Moser et al. (2014) provide causal evidence that these émigrés increased US patenting by 31 per cent in their fields of specialisation, operating through the attraction of new American researchers into émigré specialisations rather than through productivity gains among incumbents. The mechanism is precisely that of the knowledge stock externality at the centre of the present analysis: the existing American knowledge base was sufficiently deep to absorb and amplify the incoming frontier-level human capital in a way that a shallower science system could not have sustained. Borjas and Doran (2012) document an analogous dynamic following the post-1991 outflow of Soviet mathematicians to the United States.

Third, the transition was not inevitable or irreversible at the time it occurred. Germany in 1930 held a commanding lead in physics, chemistry, and medicine; its defeat was the product of political catastrophe, not the natural working out of any economic logic. The lesson for the present analysis is that scientific leadership is neither immutable nor immune to exogenous shock.

A.5.2 The Five Mechanisms of Leadership Transition

Five analytical mechanisms have been identified in the historical literature as drivers of the succession of scientific leaders. Each has a direct counterpart in the KPF framework of the present paper.

Mechanism 1: Knowledge Stock Compounding

The foundational mechanism is the compounding of accumulated knowledge stocks. Mokyr (2002) identifies the ‘useful knowledge’ stock (the body of theoretical and practical understanding embodied in trained scientists, research institutions, tacit knowledge networks, and the published literature) as the primary determinant of which countries can sustain scientific leadership. The key property of useful knowledge is its semi-public, non-rival character: once produced, it is available to all members of the scientific community who have the absorptive capacity to exploit it. Countries that have accumulated larger knowledge stocks raise the marginal productivity of their current researchers above the level achievable by equally well-funded researchers in knowledge-poorer environments.

This is precisely the $\hat{\phi}$ parameter of our analysis. Values of $\hat{\phi}_i$ approaching or exceeding unity imply that the knowledge stock compound at or above proportional rates: that each addition to the stock raises the marginal productivity of subsequent additions by at least as much as the addition itself. Our finding of $\bar{\hat{\phi}} \approx 0.76$ at the panel level, with country-specific $\hat{\phi}_i > 1$ for leading economies, is the quantitative expression of this historically documented compounding mechanism.

The historical record is consistent with this mechanism operating over long time horizons. Ben-David (1971) shows that England’s seventeenth-century scientific leadership reflected not superior talent or funding but the accumulation of a critical mass of practising natural philosophers whose interactions generated an above-average rate of discovery. Mokyr (2002) traces the same mechanism through the French Enlightenment’s institutionalisation of science at the Académie royale des sciences and through the German research university’s integration of research and teaching, both of which created organisational structures that amplified the productivity of individual researchers by embedding them in knowledge-rich environments.

Mechanism 2: Talent Migration

Talent migration is the historically fastest mechanism of scientific leadership transfer and is treated in full in Appendix A.2. The key analytical point for the present historical narrative is that the Germany-to-United States transition of the 1930s–1950s combined both channels simultaneously: the forced emigration of Jewish academics after 1933 (Moser et al., 2014) provided the exogenous human capital shock, while the prior maturation of the American research university system provided the absorptive capacity to amplify the

incoming expertise into lasting scientific dominance. The (σ, ϕ) joint shock interpretation, each émigré raises both the HC stock and the knowledge stock of the receiving country, is central to understanding why the United States, rather than other countries that also received refugees, was the primary beneficiary: only a country with a sufficiently deep prior knowledge base ($\hat{\phi}$ already near the Romer threshold) could have generated the multiplicative amplification that produced the observed 30-per-cent increase in patenting (Moser et al., 2014).

Mechanism 3: Institutional Environment

The third mechanism is the design of the institutional environment for science. Ben-David (1971) argued that the succession of scientific centres over four centuries reflected above all the differential capacity of national systems to solve three institutional problems: establishing a distinct professional role for the scientist; achieving intellectual autonomy for scientific inquiry from religious, political, or commercial interference; and creating competitive incentives for original discovery.

The German research university’s innovation in the nineteenth century was the integration of original research into the university’s teaching mission (the Humboldtian ideal) which created competitive incentives for faculty to produce novel knowledge as a condition of professional advancement. Hollingsworth and Hollingsworth (2011) systematically compared the institutional structures of major scientific laboratories across countries and found that those achieving the most significant breakthroughs share characteristics of organisational slack, interdisciplinary cross-fertilisation, and protection from short-term commercial pressures. These institutional features cannot be captured by a KPF that conditions only on aggregate human capital and knowledge stocks, but they help explain why the country-specific fixed effects μ_i in our specification are large and nationally differentiated.

The transition to American dominance after 1945 was in part an institutional story. The establishment of the National Science Foundation (1950), the massive expansion of federal R&D funding through the Department of Defense and the National Institutes of Health, and the GI Bill’s expansion of university education combined to create an institutional environment for science that was without parallel in breadth, scale, and competitive intensity. Nelson and Wright (1992) emphasise this institutional dimension: American technological leadership was not simply a matter of more resources but of more effective institutions for converting resources into frontier knowledge.

Mechanism 4: Catch-Up Dynamics and the Advantage of Backwardness

The fourth mechanism, which creates the conditions for displacement rather than merely reinforcing the leader, is the advantage of backwardness proposed by Gerschenkron (1962).

Laggard countries that have accumulated sufficient social and technological capability can grow faster than the frontier by adopting technologies already developed elsewhere at lower cost than original invention. Abramovitz (1986) extended this to a conditional catch-up hypothesis: the potential for catch-up is inversely related to the technology gap but realises only when absorptive capacity is sufficient. Applied to the succession of scientific leaders, this framework predicts that quantitative catch-up in publication volumes precedes qualitative displacement at the citation frontier, and that the lag is measured in decades rather than years.

The cross-country heterogeneity in $\hat{\phi}_i$ that our analysis recovers is consistent with countries at different stages of this catch-up process. Countries with relatively young and rapidly growing knowledge bases (Türkiye ($\hat{\phi}_i \approx 3.23$) is the clearest example) exhibit the largest compounding elasticities, consistent with the Abramovitz mechanism: a rapidly expanding knowledge base generates disproportionate marginal returns because it is being built into an environment of strong absorptive capacity. Countries at or near the frontier (the United States, Japan) exhibit much smaller or even negative $\hat{\phi}_i$, consistent with the onset of diminishing returns at large absolute knowledge stock levels.

Mechanism 5: Paradigm Transitions and Creative Destruction

The fifth mechanism is specific to the dynamics of the scientific frontier rather than the general knowledge accumulation process. Schumpeter (1942) identified the process of creative destruction as central to innovation dynamics: incumbent technologies and knowledge structures are periodically rendered obsolete by new paradigms that open entirely new research opportunities. In Kuhnian terms (Kuhn, 1962), the transition from normal science to scientific revolution creates an opportunity for challengers to leapfrog incumbents, because the existing knowledge stock of the incumbent, which creates the $\hat{\phi}$ advantage in normal science, partly depreciates in relevance as the new paradigm displaces old frameworks.

The Aghion-Howitt (1992) quality-ladder framework formalises this mechanism: innovation improves the quality of existing products or knowledge, rendering the incumbent's stock partly obsolete. The prediction is that the knowledge stock elasticity should be smaller in environments characterised by rapid paradigm turnover, precisely the quality-ladder environment that the top-1% citation share is designed to capture. Our finding that the knowledge stock is entirely insignificant in the quality specifications (while being large and significant in the quantity specifications) is consistent with the Schumpeterian quality-ladder mechanism operating in the citation-impact dimension of science while the Romer-Jones compounding mechanism governs aggregate volume.

A.5.3 Implications for the Present Analysis

Four implications for the interpretation of our empirical results follow from this historical record.

First, the finding of $\bar{\hat{\phi}} \approx 0.76$ at the panel level, with country-specific $\hat{\phi}_i > 1$ for some economies, is consistent with the historical experience: significant compounding of accumulated knowledge is the structural driver of sustained scientific leadership, operating at a sub-proportional rate on average but approaching or exceeding proportionality for the most scientifically advanced economies.

Second, the finding that quantity catch-up (in article volume) precedes quality catch-up (in citation impact) for China is historically normal rather than anomalous. The Germany-to-United States transition involved the same sequence, with a lag of approximately two to three decades between volume parity and quality leadership. The current data place China at an analogous stage, suggesting that quality catch-up remains possible but is not imminent on a 5–10 year horizon.

Third, talent migration is the historically most powerful mechanism of rapid displacement, operating directly through the knowledge stock channel and generating externalities for the recipient country’s knowledge base. Our projections condition on migration patterns remaining at in-sample levels and therefore understate both the pace of potential US decline (if emigration accelerates) and the pace of European gains (if European institutions successfully attract US-trained researchers).

Fourth, the historical record reinforces the finding that scientific leadership is not immutable. Every episode of sustained dominance has been preceded by catch-up and has eventually ended in displacement. The structural projection of dramatic US share decline by 2047 is not an aberration but the latest episode in a four-century pattern of leadership succession, with the knowledge stock compounding mechanism now quantitatively identified for the first time at global panel scale.

Fifth, the historical record identifies a direct and measurable risk in current United States science policy that maps onto the two most rapid mechanisms of displacement: talent migration and institutional environment. The following observations are not a comprehensive policy review but a mapping of documented policy directions onto the mechanisms identified in Section A.2.

Talent migration (Mechanism 2). The United States’ post-1945 scientific dominance was built partly on being the global destination of choice for frontier-level researchers. The post-1933 German émigrés (Moser et al., 2014), the post-1991 Soviet mathematicians (Borjas & Doran, 2012), and successive waves of Asian and European researchers all deepened the American knowledge base in precisely the way the $\hat{\phi}$ mechanism implies: incoming talent was absorbed by a system already rich enough to amplify it. Since 2017 and accelerating after 2022, documented trends run in the opposite direction. Visa

rejection rates for international doctoral students have increased sharply, particularly for applicants from China and India — historically the two largest sources of science and engineering graduate students to American universities. The Department of Homeland Security’s restrictions on Optional Practical Training and H-1B pathways have lengthened the uncertainty that frontier researchers face when choosing a host country (American Physical Society, 2026). The investigatory climate created by the Department of Justice’s China Initiative (2018–2022), though formally ended, demonstrably chilled international research collaboration and caused a number of senior researchers of Chinese origin to disengage from US institutions. Each of these represents a reduction in the effective openness of the US science system to incoming frontier-level talent, precisely the channel through which the Germany-to-United States transition was accelerated in the 1930s and 1940s.

Institutional environment (Mechanism 3). Ben-David (1971) and Nelson and Wright (1992) emphasise that American scientific leadership rested not only on resources but on the institutional design of the research enterprise: competitive peer-reviewed grant allocation through the National Institutes of Health and the National Science Foundation; institutional autonomy for universities from short-term political and commercial pressures; and an open meritocratic environment that attracted global talent regardless of national origin. Since 2022, federal funding freezes at NIH and NSF (in some cases exceeding 20 per cent of committed grant budgets) represent a structural reduction in the resources available to the peer-reviewed competition that has historically been the mechanism of US scientific productivity (American Physical Society, 2026). Restrictions on Diversity, Equity and Inclusion programmes at research universities, whatever their broader policy rationale, have demonstrably reduced the pool of domestic researchers from under-represented groups that has contributed to the expansion of the US science workforce over the past three decades.

The historical parallel. Germany in the early 1930s did not lose scientific leadership because it lacked resources, knowledge stock, or talent. It lost leadership because a sequence of policy decisions rendered its institutional environment hostile to the free exchange of ideas, drove frontier researchers out of the system, and disrupted the competitive environment that had generated innovation. The analogy is imperfect since the United States is not expelling scientists, but the *direction* of the mechanisms is the same. The historical record suggests that the damage from such policies is not immediately visible in publication output (which reflects current investment in the existing knowledge base), but compounds over a decade or more through the knowledge stock channel. The projections in Section 5.6 of the main paper, which condition on current input growth trajectories, do not capture any structural acceleration of the documented policy trends. If those trends continue or deepen, the US trajectory will be worse than the baseline projection implies.

The historical record of scientific leadership succession additionally highlights the role of intellectual autonomy. Ben-David (1971) and Hollingsworth and Hollingsworth (2011) identify decentralised, competitive institutional environments, where scientists choose their own research agendas, funding is allocated by peer review, and scientific standards are set by the discipline rather than by governments or funders, as a structural prerequisite for sustained scientific creativity. Conversely, episodes in which states imposed ideological constraints on scientific inquiry (Lysenkoism in the Soviet Union; restrictions on ‘Jewish physics’ in Nazi Germany) were followed by sharp productivity declines that compound through the knowledge stock over subsequent decades. The current period of heightened political scrutiny of scientific institutions, targeted funding reductions, and restrictions on research agendas in several leading science systems (American Physical Society, 2026) may constitute an analogous constraint whose productivity cost will accumulate gradually through the $\hat{\phi}$ compounding mechanism before becoming visible in publication and citation metrics.

Appendix B: Estimator Motivation, Performance, and Country-Specific Results

B.1 Econometric Motivation for AMG and CCEMG

The Augmented Mean Group (AMG) estimator (Bond & Eberhardt, 2009; Eberhardt & Teal, 2011) and Common Correlated Effects Mean Group (CCEMG) estimator (Pesaran, 2006) are both designed for macro panels with large N and moderate T under the assumption that unobserved common factors drive cross-sectional dependence in panel errors.

Standard fixed-effects (FE) estimation removes a single time-invariant country effect and a common time effect, but it cannot accommodate unobserved factors whose *country-specific loadings* differ across units — a failure mode that generates inconsistent slope estimates when the common factor and the regressors are correlated. In the KPF context, the global science frontier is precisely such a common factor: it affects all countries simultaneously, but countries with stronger international collaboration networks, more permeable absorptive capacity, or higher shares of frontier subfields are more exposed. Standard FE estimates of $\hat{\phi}$ conflate own-stock compounding with proximity to the common frontier.

AMG addresses this by estimating the common factor non-parametrically. In the first stage, time dummies are entered freely in a pooled first-differences regression; the estimated coefficients on the year dummies serve as a proxy for the common dynamic process (CDP). In the second stage, the CDP is augmented into each country’s individual regression as an additional regressor before the mean-group average is taken. The CDP absorbs the shared frontier dynamic, so the resulting $\hat{\phi}_i$ captures own-stock compounding net of the global trend.

CCEMG addresses the same problem differently: cross-sectional means of the dependent and independent variables ($\bar{y}_t, \bar{x}_t$) are augmented into each country’s regression as additional regressors. Under weak factor structure assumptions, these cross-sectional averages span the space of the latent common factors, making the CCEMG estimator consistent even when the factor loadings are heterogeneous (Pesaran, 2006). The two estimators are asymptotically equivalent under the null of a single common factor but can differ in finite samples, so we report both as a robustness check throughout.

B.2 Monte Carlo Evidence on Estimator Performance

Eberhardt and Teal (2011) provide simulation evidence showing that, under cross-sectional dependence with heterogeneous factor loadings, FE estimators exhibit substantial bias in finite samples while AMG and CCEMG both achieve near-zero bias with appropriate size. The bias is especially severe when N is large relative to T (our case: $N = 81\text{--}90$, $T = 20$),

because the cross-sectional dimension amplifies correlation between the common factor and the regressors faster than the time dimension can deliver consistent within-estimator variation.

In our sample, the CIPS panel unit root tests (Table 3 of the main paper) confirm that the quantity-specification regressors ($\ln H_{A,it}$, $\ln Q_{it}$, and $\ln K_{i,t-1}$) are all $I(1)$, consistent with the cross-sectionally dependent panel structure for which AMG and CCEMG are designed. For the *quantity* specification, the Westerlund (2007) cointegration tests (Appendix E) provide additional evidence that the long-run relationships we estimate are genuine cointegrating relationships rather than spurious regressions, further validating the estimator choice.

For the *quality* specification the integration order is mixed: the logit-transformed top-1% citation share is $I(0)$ while the regressors remain $I(1)$. This integration-order mismatch means AMG and CCEMG, which are designed for $I(1)/I(1)$ systems, are not directly applicable to the quality equation. PMG+CSD (Pooled Mean Group with cross-sectional mean augmentation) is the appropriate estimator for the quality specification because it absorbs cross-sectional dependence through mean augmentation without requiring a cointegrating relationship between the $I(0)$ dependent variable and the $I(1)$ regressors. See the discussion in Section 4 of the main paper.

B.3 Country-Specific KPF Coefficient Scatter Plots

Each figure below shows, for one economy, the estimated KPF coefficients ($\hat{\sigma}_i$, $\hat{\gamma}_i$, $\hat{\phi}_i$) across all eight specification $\times$ estimator combinations used in the main analysis (Specifications A–D $\times$ AMG, CCEMG; above-median sample, 90 countries).

Please note that this does not constitute a claim that the coefficients are statistically significantly different from one another (our panel coefficient homogeneity result suggests otherwise), but we present the evidence for completeness.

How to read the figures. Each row within a panel corresponds to one specification-estimator combination. Filled circles denote AMG estimates; open circles denote CCEMG estimates. The solid vertical line marks the country’s mean coefficient across all available combinations; the dashed vertical line marks the panel mean-group coefficient. In the right panel ($\hat{\phi}_i$), the dotted line at $\phi = 1$ marks the Romer proportionality threshold: estimates to the right indicate strong Schumpeterian returns at the country level.

B.3.1 Eight Major Science Economies (G8)




Figure 1: China (CHN) — KPF coefficient estimates across all specification × estimator combinations.

Notes: All notation as in main text. China's $\hat{\phi}_i$ centred near 1.09 in most combinations, exceeding the Romer threshold ($\phi = 1$). Human capital elasticity ranks second in the G8 group. Patent stock elasticity among the lowest, consistent with the documented disconnect between strategic patent filing and genuine scientific productivity.




Figure 2: United States (USA) — KPF coefficient estimates.

Notes: The United States displays the clearest saturation signature in the G8: human capital elasticity $\hat{\sigma}_i$ is negative in the majority of combinations, consistent with research workforce saturation. Knowledge stock elasticity $\hat{\phi}_i \approx 0.99$ approaches the Romer threshold. The combination of saturation in $\hat{\sigma}$ and near-proportional $\hat{\phi}$ places the US sixth in the composite ranking despite holding the world's deepest accumulated knowledge base.




Figure 3: Germany (DEU) — KPF coefficient estimates.

Notes: Germany ranks second in the composite G8 ranking, distinguished by balanced positive performance across all three coefficient dimensions. Joint first on patent stock efficiency alongside France. The relative stability of Germany's coefficient estimates across combinations (low SD) reflects the maturity and consistency of its research system.




Figure 4: Japan (JPN) — KPF coefficient estimates.

Notes: Japan ranks last in the composite G8 ranking and in knowledge stock elasticity, consistent with a technologically mature science system experiencing diminishing returns. The near-zero or negative $\hat{\phi}_i$ in most combinations suggests that Japan's large accumulated knowledge stock is no longer generating proportional marginal returns to current research — a Schumpeterian saturation result rather than structural inefficiency. See Leydesdorff and Wagner (2009) for an analysis of Japan's core-periphery position in global science networks.

GBR — United Kingdom
KPF Coefficient Estimates Across Specifications × Estimators




Figure 5: United Kingdom (GBR) — KPF coefficient estimates.

Notes: The United Kingdom shows $\hat{\phi}_i > 1$ in several combinations, consistent with its high-end position in the global citation network (highest top-1% citation fraction among the 13 focus countries at 1.62%). The UK's quality advantage in the present therefore reflects a knowledge stock that still generates super-proportional returns in some specifications, but the quality-weighted projections suggest this advantage will erode under current human capital growth rates of 0.33%/yr.

FRA — France
KPF Coefficient Estimates Across Specifications × Estimators




Figure 6: France (FRA) — KPF coefficient estimates.

Notes: France is joint first with Germany on patent stock elasticity, reflecting effective conversion of industrial R&D into academic publication output. Knowledge stock elasticity $\hat{\phi}_i > 1$ in several combinations. France's position illustrates that patent-science complementarity can serve as an alternative path to research efficiency in economies where the Grandes Écoles tradition maintains high average researcher quality.




Figure 7: Canada (CAN) — KPF coefficient estimates.

Notes: Canada exhibits the most stable rankings in the entire analysis: its composite rank of 4th is identical across all eight specification-estimator combinations (zero standard deviation), making it the most consistently placed economy. Canada leads the G8 in human capital elasticity $\hat{\sigma}_i$, consistent with its relatively smaller research workforce and lower saturation pressure compared to the United States. The consistency of the coefficient estimates across specifications reflects a stable and well-functioning research system.




Figure 8: Italy (ITA) — KPF coefficient estimates.

Notes: Italy shows positive $\hat{\phi}_i$ in most combinations and positive human capital elasticity, consistent with a science system where the knowledge base still generates meaningful compounding returns despite slow overall growth. Italy's top-1% citation fraction (1.36%) is among the highest in the sample, suggesting that the existing knowledge stock delivers disproportionate citation quality relative to the volume of output.

B.3.2 Emerging Market Economies (Brazil, India, Israel, Mexico, South Africa, South Korea, Türkiye)




Figure 9: Brazil (BRA) — KPF coefficient estimates.

Notes: Brazil ranks second among the five emerging markets in composite efficiency. Its most distinctive feature is the highest patent stock elasticity $\hat{\gamma}_i$ of all 13 countries, indicating exceptional complementarity between commercial innovation (patent activity) and academic publication output. Combined with the highest human capital growth rate in the group (2.28%/yr), Brazil generates the largest projected share gain of any focus country in the baseline projection. Note: Brazil has only 12 of 20 possible researcher-count observations, reducing the precision of Specs A and C estimates.




Figure 10: Mexico (MEX) — KPF coefficient estimates.

Notes: Mexico presents the central efficiency- without-scale paradox of the analysis. Its human capital elasticity $\hat{\sigma}_i \approx 4.4$ is the highest of all 13 countries — each quality-adjusted unit of research human capital generates more publication output in Mexico than anywhere else in the sample. Yet Mexico's projected share falls over the horizon: the absolute knowledge stock is too small and growing too slowly to generate compounding gains, and brain drain to the United States depresses effective domestic research capacity. Near-zero $\hat{\phi}_i$ in most combinations confirms the absence of a compounding mechanism from the domestic knowledge base.




Figure 11: Türkiye (TUR) — KPF coefficient estimates.

Notes: Türkiye ranks first in the emerging market group and holds the highest knowledge stock elasticity $\hat{\phi}_i \approx 3.23$ of all 13 countries — also the only economy to exceed the Romer threshold in *every* specification-estimator combination (note that estimates are very widely dispersed, reflecting the rapidly changing character of Turkey’s science system). This super-proportional compounding reflects the combination of a rapidly growing but relatively young knowledge base and strong institutional investment in university expansion during the 2000s and 2010s. The lowest human capital elasticity and most negative patent elasticity in the group indicate that institutional investment (publication incentive schemes, university rankings targets) rather than organic research quality drives the output expansion.




Figure 12: South Africa (ZAF) — KPF coefficient estimates.

Notes: South Africa ranks third among the emerging markets with the most balanced emerging-market profile. The near-zero aggregate human capital elasticity $\hat{\sigma}_i$ is consistent with substantial brain drain effects documented by Fedderke and Dong (2026): emigration of skilled researchers erodes the effective domestic human capital stock at a rate that offsets the gross expansion of the research workforce, depressing the marginal return to each additional researcher. The moderate positive $\hat{\phi}_i$ in most combinations indicates that the domestic knowledge base still generates meaningful compounding returns.




Figure 13: India (IND) — KPF coefficient estimates.

Notes: India participates in only 4 of the 8 combinations (Specs B and D only, due to missing researcher counts in Specs A and C) and accordingly ranks last in the emerging market group by default of coverage rather than genuine inefficiency. The strongly negative aggregate human capital elasticity in available combinations is consistent with the well-documented quality duality of India’s research system: a small number of elite institutions (IITs, IISc, TIFR) produce internationally competitive research while the large majority of Indian universities produce limited internationally visible output, generating a negative average HC elasticity when both tiers are pooled. India’s high projected quantity share by 2047 (9.4%) reflects input accumulation through patent stock growth rather than average HC quality improvement.





Figure 14: Israel (ISR) — KPF coefficient estimates across all specification × estimator combinations.

Notes: Israel’s knowledge stock elasticity $\hat{\phi}_i \approx 2.25$ is robustly positive across all available combinations (Specs B and D × AMG/CCEMG; Specs A and C excluded due to missing researcher-count data). Its negative quantity fixed effect ($\hat{\mu}_i = -0.17$, Table ?? of the main paper) coexists with this high elasticity: the NIS generates fewer articles than measured inputs predict, but converts each unit of its existing knowledge stock exceptionally efficiently, reflecting concentration in a small number of world-leading disciplines rather than broad-based research volume.

**South Korea (KOR) — Country-Specific KPF Coefficient Estimates
Across 8 Specification × Estimator Combinations (Specs A-D × AMG, CCEMG)**



Note: Uncertainty bars show min-to-max range across available combinations, not a standard error. Panel mean from PMG+CSD (quantity specification).

Figure 15: South Korea (KOR) — KPF coefficient estimates across all specification × estimator combinations.

Notes: South Korea is the only emerging market economy for which both the human capital elasticity $\hat{\sigma}_i$ and the patent stock elasticity $\hat{\gamma}_i$ are positive across all eight specification × estimator combinations, reflecting consistently positive marginal returns to both research workforce quality and commercial innovation activity. The knowledge stock elasticity $\hat{\phi}_i$ is positive in most combinations but below the Romer threshold, placing South Korea in the semi-endogenous regime at the country level.

Appendix C: Country-Specific Elasticity Projections

— Full Comparison Table

Table C.1 compares the baseline panel-mean projection (Section 5.6 of the main paper) with a robustness variant using country-specific AMG elasticities. Country-specific coefficients are capped at $\hat{\phi} \pm 2\widehat{SD}(\hat{\phi})$ to limit the influence of outlier estimates arising from the short $T = 20$ estimation window.

The table confirms the interpretation given in the main text: the country-specific projection is highly sensitive to individual coefficient estimates, with several countries exhibiting implausible trajectories (near-collapse or explosive expansion) even after the 2 SD trimming. This instability is itself a substantive finding: it demonstrates that $T = 20$ country-level time series are insufficient for reliable long-horizon projection at the country-specific level, validating the use of panel-mean elasticities in the baseline. The country-specific projections should be read as bounding exercises rather than alternative point forecasts.

Table 1: Country-Specific vs Panel-Mean Projections — Full Comparison (Global Knowledge Output Share, %)

		Panel-mean elasticities			Country-specific elasticities [†]		
ISO	Country	2022	2035	2047	2022	2035	2047
<i>G8 economies</i>							
CHN	China	22.57	29.52	27.78	22.57	14.20	7.06
USA	United States	14.01	6.85	2.96	14.01	0.83	0.05
DEU	Germany	4.07	2.03	0.87	4.07	0.22	0.03
JPN	Japan	2.99	1.45	0.67	2.99	0.13	0.02
GBR	United Kingdom	4.43	2.60	1.33	4.43	0.12	0.01
FRA	France	2.50	1.39	0.75	2.50	0.10	0.01
CAN	Canada	2.62	1.53	0.75	2.62	0.18	0.02
ITA	Italy	3.00	2.21	1.45	3.00	12.40	7.06
<i>Emerging market economies</i>							
IND	India	5.48	8.34	9.38	5.48	0.90	0.13
BRA	Brazil	1.93	3.05	6.80	1.93	0.29	0.04
TUR	Türkiye	1.50	1.96	2.25	1.50	18.40	7.06
ZAF	South Africa	0.64	0.82	0.97	0.64	0.08	0.01
MEX	Mexico	0.66	0.52	0.36	0.66	0.05	0.01

[†] Country-specific AMG Spec B elasticities, capped at panel mean ± 2 SD per coefficient. Knowledge stock updated endogenously via PIM, $\rho_K = 0.10$. All other design choices as in the baseline projection.

Note on extreme values: Türkiye’s exceptional $\hat{\phi}_i \approx 3.23$ (largest in the sample) and Italy’s positive outlier $\hat{\phi}_i$ generate explosive trajectories that overwhelm the other countries’ shares; China, India, Brazil, and all other G8 economies collapse to near-zero as these two countries absorb the global share. This mechanical result illustrates why $T=20$ country-specific estimates are insufficiently precise for 25-year projection, and confirms the appropriateness of the panel-mean baseline.

The comparison between the panel-mean and country-specific projections yields a methodologically instructive lesson. When individual country elasticities are applied, the projection is dominated by the countries with the largest (or most negative) $\hat{\phi}_i$ estimates, because the PIM compounding amplifies any difference in ϕ exponentially over 25 years. Türkiye’s $\hat{\phi}_i \approx 3.23$ and Italy’s positive outlier generate shares that crowd out all other

economies by the end of the horizon; simultaneously, countries with near-zero or negative $\hat{\phi}_i$ (most of the G8) collapse toward zero. The 2 SD cap is already a substantial intervention; without it the instability is even more extreme. The standard errors around individual country elasticities from $T = 20$ observations are too large relative to the point estimates for the latter to serve as reliable inputs to a 25-year compounding simulation.

This finding has a positive implication: it strengthens confidence in the panel-mean projection as the appropriate summary of the underlying KPF dynamics. The panel-mean elasticities are estimated with considerably greater precision (standard errors of order 0.01–0.10 for the key $\hat{\phi}$ parameter in the FE estimator; 0.10–0.12 in AMG) and are less susceptible to the outlier-amplification problem that makes country-specific projections unreliable.

The following table reports the full convergence scenario results referenced in Section 5 of the main paper. Baseline = country-specific in-sample input growth rates; Linear/Exponential = convergence to G5 frontier rates ($g_H^* = 0.33\%/yr$, $g_Q^* = 2.11\%/yr$).

Table 2: Projected 2047 Global Knowledge Production Shares (%) Under Three Scenarios

		2022	Projected 2047 share (%)			
ISO	Country	Share	Baseline	Linear	Exponential	Direction
<i>Gaining economies</i>						
CHN	China	22.57	27.76	31.26	31.92	↑ (stronger under conv.)
IND	India	5.48	9.38	9.37	9.20	↑ (robust)
TUR	Türkiye	1.50	2.25	2.02	1.92	↑ (smaller under conv.)
BRA	Brazil	1.93	6.80	3.18	2.46	↑ baseline; near-flat conv.
ZAF	South Africa	0.64	0.97	0.82	0.77	↑ (smaller under conv.)
<i>Declining economies</i>						
USA	United States	14.01	2.96	4.41	4.95	↓ (less steep under conv.)
GBR	United Kingdom	4.43	1.33	1.84	2.01	↓ (moderated)
DEU	Germany	4.07	0.87	1.34	1.52	↓ (moderated)
JPN	Japan	2.99	0.67	0.90	0.97	↓ (moderated)
KOR	South Korea	2.22	1.29	1.46	1.49	↓ (moderated)
ITA	Italy	3.00	1.45	1.75	1.83	↓ (moderated)
FRA	France	2.50	0.75	0.92	0.97	↓ (moderated)
CAN	Canada	2.62	0.75	1.09	1.22	↓ (moderated)
MEX	Mexico	0.66	0.36	0.43	0.45	↓ (moderated)
ISR	Israel	0.54	0.28	0.31	0.31	↓ (robust)

Notes: Frontier growth rates (G5 average: USA, GBR, DEU, JPN, FRA): $g_H^* = 0.33\%/yr$, $g_Q^* = 2.11\%/yr$. Linear: full convergence by 2047. Exponential: speed $\lambda = 0.10$ (half-life ≈ 7 years). All other parameters as in Table ???. See the projection figure in the main paper (Figure ???).

Appendix D: Rest-of-World Knowledge Stock — Robustness Check

D.1 Motivation, Design, and Identification

The main analysis uses each country’s own accumulated knowledge stock $K_{i,t-1}$ as the measure of knowledge capital available to its researchers. An alternative hypothesis, grounded in a long tradition of international knowledge spillover research, holds that peer-reviewed science is a near-pure global public good. The foundational concept of knowledge as an externality-generating public good traces to Marshall (1890)’s analysis of industrial districts, was formalised for innovation by Arrow (1962)’s learning-by-doing model, and was given empirical substance in the context of R&D by Griliches (1979, 1992). The geographic dimension of knowledge spillovers was documented by Jaffe (1986). The first systematic theoretical treatment of *international* knowledge spillovers was provided by Grossman and Helpman (1991), who showed that countries can benefit from each other’s innovation through trade in knowledge-intensive goods. The landmark empirical formalisation by Coe and Helpman (1995) quantified these international R&D spillovers through bilateral import shares, establishing that a country’s total factor productivity depends on a trade-weighted average of its partners’ R&D stocks. The present application extends this logic from commercial R&D to the academic knowledge channel: once published, any peer-reviewed paper is in principle accessible to any researcher anywhere. Under this view, the relevant knowledge stock for country i ’s researchers is not $K_{i,t-1}$ alone but the world total K_{t-1}^{world} .

Why K_{t-1}^{world} cannot be a regressor. K_{t-1}^{world} is identical for all countries in a given year, so it is exactly collinear with the year fixed effects in FE. In AMG, $\Delta \ln K_{t-1}^{world}$ is the same for all countries in each period and is therefore perfectly represented by a linear combination of the year dummies in the Step 1 first-differenced regression; the accumulated Common Dynamic Process $\hat{C}_t$ thus already embeds the global knowledge frontier trend. In CCEMG, the cross-sectional mean augments $\overline{\ln K_{t-1}} = N^{-1} \sum_i \ln K_{i,t-1}$ tracks the global average of log knowledge stocks, and since $\ln K_{t-1}^{world} = \ln(N \cdot \overline{K_{t-1}})$, adding it as a structural regressor generates rank deficiency. In all three cases the identification failure is structural, not accidental: each estimator is designed to remove common factors, and the world knowledge stock is precisely such a factor.

Construction of K_{t-1}^{-i} and the log-nonlinearity escape. The *rest-of-world* (RoW) knowledge stock is defined as:

$$K_{t-1}^{-i} = K_{t-1}^{world} - K_{i,t-1} = \sum_{j \neq i} K_{j,t-1}, \quad (1)$$

where $K_{t-1}^{world} = \sum_j K_{j,t-1}$ is the sum across all countries in the above-median sample. This avoids exact collinearity because the log transformation is strictly concave:

$$\ln K_{t-1}^{-i} = \ln(K_{t-1}^{world} - K_{i,t-1}) \neq \ln K_{t-1}^{world} - \ln K_{i,t-1}. \quad (2)$$

For small countries, where $K_{i,t-1} \ll K_{t-1}^{world}$, the approximation $\ln K_{t-1}^{-i} \approx \ln K_{t-1}^{world}$ holds closely: the rest-of-world stock is effectively the same as the world stock and is absorbed along with it by each estimator's global-trend control. Only for the five or six largest science economies — the United States, China, Germany, Japan, the United Kingdom — which together account for roughly 50–60% of the world knowledge stock, does the subtraction produce a K_{t-1}^{-i} that diverges meaningfully from K_{t-1}^{world} .

The augmented specification is therefore:

$$\ln \dot{A}_{it} = \mu_i + \lambda_t + \sigma \ln H_{A,it} + \gamma \ln Q_{it} + \phi_1 \ln K_{i,t-1} + \phi_2 \ln K_{t-1}^{-i} + \varepsilon_{it}, \quad (3)$$

where ϕ_1 captures own-stock compounding through the domestic knowledge base and ϕ_2 captures the international knowledge spillover through the global public good channel.

Identification caveat: thin cross-sectional variation. The cross-sectional standard deviation of $\ln K_{t-1}^{-i}$ *within* each year is only 0.036, roughly one-thirtieth of the between-country standard deviation of $\ln K_{i,t-1}$. This follows directly from the small-country approximation: for the large majority of the 81 above-median countries, the own stock is a negligible fraction of the world total, so subtracting it leaves the remainder essentially unchanged in log terms. Identification of $\hat{\phi}_2$ therefore relies on the gradient across the handful of large-economy outliers, not the full panel, and standard errors for $\hat{\phi}_2$ will be correspondingly wide.

Anticipated pattern of results. Combining the identification geometry with the theoretical priors, the model delivers four testable predictions *before the data are examined*:

1. **FE:** $\hat{\phi}_1$ stable (FE year-demeaning removes the common trend from $\ln K_{t-1}^{-i}$, leaving a residual that is orthogonal to $\ln K_{i,t-1}$ in the demeaned specification); $\hat{\phi}_2 > 0$ (positive spillover, consistent with the global public good hypothesis); wide confidence interval on $\hat{\phi}_2$ (thin cross-sectional identifying variation).
2. **AMG:** $\hat{\phi}_1$ collapses toward zero; $\hat{\phi}_2$ inflated — the Common Dynamic Process absorbs the time trend in $\ln K_{t-1}^{-i}$, leaving a residual that is nearly perfectly correlated with $\ln K_{i,t-1}$ (a rapidly growing large-country own stock simultaneously reduces the K^{-i} of every other country), creating near-collinearity in the Step 2 regression. *Not informative about the true spillover magnitude.*

3. **CCEMG**: $\hat{\phi}_2$ explosive with large standard error; $\hat{\phi}_1 \rightarrow 0$ — the cross-sectional mean augments of $\ln K_{t-1}^{-i}$ is nearly identical to $\ln K_{t-1}^{-i}$ itself for most countries (within-year SD = 0.036), creating near-singular systems at the country level. *Not informative about the true spillover magnitude.*
4. **Interpretation of AMG/CCEMG instability**: the instability is *confirmatory*, not alarming. It demonstrates that baseline AMG and CCEMG were already successfully conditioning on the world knowledge level via the Common Dynamic Process and the cross-sectional mean augmenters respectively. The global frontier was not an omitted variable in the baseline; it was correctly absorbed as a common factor, which is precisely the econometric rationale for using these estimators in the first place.

D.2 Results

Table 3 reports the augmented specification alongside the baseline for all three estimators.

Table 3: Rest-of-World Knowledge Stock Robustness Check (Spec B, Above-Median Sample)


Estimator	Baseline: own K only			Augmented: own K + RoW K			
	$\hat{\sigma}$	$\hat{\gamma}$	$\hat{\phi}$	$\hat{\sigma}$	$\hat{\gamma}$	$\hat{\phi}_1$	$\hat{\phi}_2$
FE	0.214 (0.656)	0.046*** (0.045)	0.646*** (0.014)	0.124 (0.652)	0.048*** (0.046)	0.645*** (0.014)	0.326 (0.336)
AMG	2.927*** (0.787)	0.076 (0.087)	0.715*** (0.118)	1.283 (1.227)	0.277** (0.131)	0.195 (0.217)	0.808** (0.355)
CCEMG	1.150 (0.831)	−0.007 (0.104)	0.667*** (0.122)	0.461 (1.197)	0.122 (0.125)	0.049 (0.254)	27.18*** (7.94)
Countries	73–81			69–81			

Notes: Dependent variable: $\ln(\text{whole-count articles})$; H_A = PWT hc index; Q = patent stock; $\rho_K = 0.10$. $\hat{\phi}_1$: own-country knowledge stock elasticity. $\hat{\phi}_2$: rest-of-world knowledge stock elasticity ($\ln K_{t-1}^{-i} = \ln(K_{t-1}^{\text{world}} - K_{i,t-1})$). AMG includes a country-specific Common Dynamic Process extracted via first-differenced year dummies. CCEMG augments with cross-sectional means of all regressors. Within-year cross-sectional SD of $\ln K_{t-1}^{-i}$: 0.036 (see text). Standard errors in parentheses. ***, **, *: significance at 1%, 5%, 10%.

D.3 Interpretation

Three findings from Table 3 are substantively important.

Finding 1: Own-stock compounding is fully robust under FE. The FE estimate of $\hat{\phi}_1$ moves from 0.646*** in the baseline to 0.645*** in the augmented specification, a



Solid bars = baseline; hatched bars = augmented. FE $\hat{\phi}_1$ stable ($\Delta = -0.001$); AMG/CCEMG $\hat{\phi}_1$ collapses (near-collinearity with internal global-trend controls); $\hat{\phi}_2$: positive and suggestive under FE; explosive under CCEMG (see Appendix D.3).

Figure 16: Rest-of-World Knowledge Stock Robustness: Baseline vs Augmented Specification. *Notes:* Top-left: own-stock elasticity $\hat{\phi}_1$ under baseline (solid) and augmented (hatched) specification across all three estimators. Error bars are 95% confidence intervals. Numbers above bars show point estimates. Top-right: rest-of-world spillover $\hat{\phi}_2$ from the augmented specification (note the CCEMG value of 27.18 is a near-singular artefact, not an economic estimate). Bottom row: full coefficient vectors for each estimator. FE $\hat{\phi}_1$ is stable to three decimal places; AMG and CCEMG $\hat{\phi}_1$ collapse due to near-collinearity between $\ln K_{t-1}^{-i}$ and the internal global-trend controls (CDP and CS-mean augmenters). See Section D.3 for interpretation.

change of 0.001, entirely negligible. This is the cleanest robustness test available: adding a measure of the global knowledge pool as a separate regressor does not disturb the estimated compounding effect of the country's own accumulated science base. FE achieves this because double-demeaning explicitly removes the year mean of $\ln K_{t-1}^{-i}$ (which contains almost all of its variation), leaving a small but orthogonal cross-sectional residual to identify $\hat{\phi}_2$ separately from $\hat{\phi}_1$.

Finding 2: AMG and CCEMG break because they already control for the global trend. The dramatic instability of $\hat{\phi}_1$ in AMG ($0.715 \rightarrow 0.195$) and CCEMG ($0.667 \rightarrow 0.049$) is not a genuine substantive finding: it reflects a near-identification failure. In AMG, the Common Dynamic Process extracted in Step 1 absorbs the global knowledge trend; since $\ln K_{t-1}^{-i}$ is almost equal to $\ln K_{t-1}^{world}$ (dominated by the global trend), adding $\ln K_{t-1}^{-i}$ as a structural regressor creates near-collinearity with the CDP. In CCEMG, the cross-sectional mean of $\ln K_{t-1}^{-i}$ used as an augmenters is nearly identical to $\ln K_{t-1}^{-i}$ itself for most countries (because the within-year cross-sectional SD is only 0.036),

producing near-singular systems and the explosive $\hat{\phi}_2 = 27.18$ (SE = 7.94). Both outcomes confirm that AMG and CCEMG are already conditioning on the world knowledge level via their internal mechanisms; precisely as they should. The global knowledge frontier is not an omitted variable in the baseline; it is absorbed as a common factor.

Finding 3: A suggestive positive international spillover under FE. The FE estimate $\hat{\phi}_2 = 0.326$ (SE = 0.336) is positive and of economically plausible magnitude for a Coe–Helpman academic spillover effect, but does not reach conventional significance thresholds ($p = 0.33$). Given the within-year cross-sectional SD of $\ln K_{t-1}^{-i}$ of only 0.036, the effective identifying variation is small, and the confidence interval around $\hat{\phi}_2$ is wide. The point estimate is consistent with, but cannot confirm, the hypothesis that countries additionally benefit from the expansion of the global knowledge pool beyond what the year fixed effects capture. Longer panels or field-level data with greater cross-sectional dispersion in knowledge stocks would be required to test this hypothesis precisely.

D.4 Conclusion

The rest-of-world robustness check delivers a clear and unambiguous message on the central question of the paper: the own-stock compounding effect ($\hat{\phi} \approx 0.65$) is robust to adding a direct control for the global knowledge pool. The instability of AMG and CCEMG in this augmented specification is not a limitation but a confirmation that these estimators were already doing their job correctly in the baseline: absorbing the global frontier trend as a common factor and identifying the idiosyncratic country-level compounding effect that is the object of policy interest.

Appendix E: Westerlund (2007) Panel Cointegration Tests

Table 4 reports the four Westerlund (2007) ECM-based panel cointegration statistics for Spec B (the preferred specification) applied to the above-median sample. The tests examine whether the KPF variables $\ln \dot{A}_{it}$, $\ln H_{A,it}$, $\ln Q_{it}$, and $\ln K_{i,t-1}$ form a cointegrating relationship.

The test is based on the ECM estimated for each country i :

$$\Delta \ln \dot{A}_{it} = c_i + \alpha_i \ln \dot{A}_{i,t-1} + \beta'_i \mathbf{x}_{i,t-1} + \gamma_i \Delta \ln \dot{A}_{i,t-1} + \delta'_i \Delta \mathbf{x}_{i,t-1} + \varepsilon_{it}, \quad (4)$$

where $\mathbf{x}_{it} = (\ln H_{A,it}, \ln Q_{it}, \ln K_{i,t-1})'$. Under the null of no cointegration, $\alpha_i = 0$ for all i . The four statistics aggregate evidence across countries in different ways:

- $G_a = N^{-1} \sum_i T_i \hat{\alpha}_i$: group-mean, weights by sample size; sensitive to large-magnitude EC speeds
- $G_t = N^{-1} \sum_i \hat{\alpha}_i / \widehat{SE}(\hat{\alpha}_i)$: group-mean t-statistic; sensitive to precision of individual estimates
- $P_a = \bar{T} \hat{\alpha}_{pool}$: panel statistic, pool α by GLS; $\bar{T}$ = mean group sample size
- $P_t = \hat{\alpha}_{pool} / \widehat{SE}(\hat{\alpha}_{pool})$: panel t-statistic (pooled GLS)

G_a and G_t test H_1 : cointegration in at least one country. P_a and P_t test H_1 : the panel as a whole is cointegrated. All four are one-sided; large negative values constitute evidence against H_0 .

Table 4: Westerlund (2007) Panel Cointegration Tests — KPF Spec B, Above-Median Sample

Statistic	H_1	Value	CV 10%	CV 5%	CV 1%
G_a	≥ 1 country cointegrated	-31.21***	-7.89	-8.99	-11.63
G_t	≥ 1 country cointegrated	-0.58	-2.80	-3.04	-3.52
P_a	panel cointegrated	-11.51	-11.93	-13.87	-18.10
P_t	panel cointegrated	-6.07***	-4.96	-5.31	-5.99

Notes: ECM in equation (4) estimated per country with one lag. $N = 72$ countries, $\bar{T} = 16.3$ years. Dependent variable: $\ln(\text{whole-count articles})$; regressors: PWT hc index, patent stock, lagged knowledge stock ($\delta_K = 0.10$). Critical values from Westerlund (2007) Table 1, $k = 3$ regressors, interpolated for $N = 72$, $\bar{T} = 16$. One-sided test; critical values are negative. Standard critical values assume cross-sectional independence; bootstrap critical values (not reported) account for CSD but are computationally intensive. Given the strong CSD documented in Section 4.1, the standard critical values are a conservative benchmark. ***, **, *: significance at 1%, 5%, 10%.

Interpretation. G_a and P_t both reject the null at 1%, providing strong evidence that (i) at least some individual countries exhibit significant error-correction in the KPF relationship, and (ii) the panel as a whole is cointegrated. G_t fails to reject: this reflects the limited power of the individual-country t-statistic with $\bar{T} \approx 16$ observations per group. With short time series, individual ECM speeds $\hat{\alpha}_i$ are estimated imprecisely (large standard errors), attenuating the t-statistics even when the true α_i is large and negative. The G_a statistic, which weights by $T_i \hat{\alpha}_i$ rather than $\hat{\alpha}_i / \widehat{SE}$, is less affected by this attenuation and therefore provides more informative evidence with short panels. P_a is borderline (-11.51 versus the 10% critical value of -11.93).

The mixed pattern of G_a^{***} and P_t^{***} alongside insignificant G_t and borderline P_a is a characteristic signature of short- T heterogeneous panels with genuine but imprecisely estimated individual cointegrating relationships. The evidence is consistent with the main text's cointegrating interpretation of the KPF, and is reinforced by the PMG error-correction analysis (Table 7 of the main manuscript), which finds 77% of country-specific error-correction speeds negative (stabilising) with a median half-life of 1.5 years: independent evidence for mean-reversion toward the long-run KPF equilibrium.

References (Online Appendix)

- Abramovitz, M. (1986). Catching up, forging ahead, and falling behind. *Journal of Economic History*, 46(2), 385–406. <https://doi.org/10.1017/S0022050700046209>
- Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2), 323–351. <https://doi.org/10.2307/2951599>
- American Physical Society. (2026). The fight to keep science global. *APS News*. <https://www.aps.org/apsnews/2026/01/fight-keep-science-global>
- Arrow, K. J. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 29(3), 155–173. <https://doi.org/10.2307/2295952>
- Ben-David, J. (1971). *The scientist's role in society: A comparative study*. Prentice-Hall.
- Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4), 1104–1144. <https://doi.org/10.1257/aer.20180338>
- Bond, S., & Eberhardt, M. (2009). *Cross-section dependence in non-stationary panel models: A novel estimator* [Paper presented at the Nordic Econometrics Conference, Lund University, Sweden].
- Borjas, G. J., & Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *Quarterly Journal of Economics*, 127(3), 1143–1203. <https://doi.org/10.1093/qje/qjs015>
- Coe, D. T., & Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5), 859–887. [https://doi.org/10.1016/0014-2921\(94\)00100-E](https://doi.org/10.1016/0014-2921(94)00100-E)
- Coe, D. T., Helpman, E., & Hoffmaister, A. W. (2009). International R&D spillovers and institutions. *European Economic Review*, 53(7), 723–741. <https://doi.org/10.1016/j.euroecorev.2009.02.005>
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152. <https://doi.org/10.2307/2393553>
- Eberhardt, M., & Teal, F. (2011). Econometrics for grumblers: A new look at the literature on cross-country growth empirics. *Journal of Economic Surveys*, 25(1), 109–155. <https://doi.org/10.1111/j.1467-6419.2010.00624.x>
- Fedderke, J. W., & Dong, X. (2026). Brain drain and productivity growth: Evidence from South Africa, 1947–2019. *Economics of Transition and Institutional Change*, 00, 1–27. <https://doi.org/10.1111/ecot.70031>
- Fedderke, J. W., & Liu, Y. (2017). Schumpeterian and semi-endogenous productivity growth explanations. *Economics of Transition and Institutional Change*, 25(1), 111–137. <https://doi.org/10.1111/ecot.12114>
- Fu, X. (2015). *China's path to innovation*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107588462>

- Gerschenkron, A. (1962). *Economic backwardness in historical perspective: A book of essays*. Belknap Press of Harvard University Press.
- Griliches, Z. (1979). Issues in assessing the contribution of research and development to productivity growth. *Bell Journal of Economics*, 10(1), 92–116.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4), 1661–1707.
- Griliches, Z. (1992). The search for R&D spillovers [Supplement]. In *Scandinavian journal of economics* (pp. 29–47, Vol. 94).
- Griliches, Z., & Mairesse, J. (1998). Production functions: The search for identification (S. Strøm, Ed.), 169–203.
- Grossman, G. M., & Helpman, E. (1991). *Innovation and growth in the global economy*. MIT Press.
- Guellec, D., & van Pottelsberghe de la Potterie, B. (2001). R&d and productivity growth: Panel data analysis of 16 OECD countries. *OECD Economic Studies*, 33, 103–126. https://doi.org/10.1787/eco_studies-v2001-art12-en
- Guellec, D., & van Pottelsberghe de la Potterie, B. (2004). From R&D to productivity growth: Do the institutional settings and the source of funds of R&D matter? *Oxford Bulletin of Economics and Statistics*, 66(3), 353–378. <https://doi.org/10.1111/j.1468-0084.2004.00083.x>
- Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools. (8498).
- Heinze, T., Jappe, A., & Pithan, D. (2019). From North American hegemony to global competition for scientific leadership? Insights from the Nobel population. *PLOS ONE*, 14(4), e0213916. <https://doi.org/10.1371/journal.pone.0213916>
- Hollingsworth, J. R., & Hollingsworth, E. J. (2011). Major discoveries, creativity, and the dynamics of science. In *Major discoveries, creativity, and the dynamics of science* (pp. 57–69). Remaprint.
- Hu, A. G. Z., & Jefferson, G. H. (2009). A great wall of patents: What is behind China’s recent patent explosion? *Journal of Development Economics*, 90(1), 57–68. <https://doi.org/10.1016/j.jdeveco.2008.11.004>
- Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms’ patents, profits, and market value. *American Economic Review*, 76(5), 984–1001.
- Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), 577–598. <https://doi.org/10.2307/2118401>
- Jones, C. I. (1995). R&D-based models of economic growth. *Journal of Political Economy*, 103(4), 759–784. <https://doi.org/10.1086/262002>
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.

- Larue, J., Gingras, Y., Sugimoto, C. R., & Haustein, S. (2022). Bibliometric inequality during the covid-19 pandemic: Country-level publication patterns. *Quantitative Science Studies*, 3(1), 1–21.
- Leydesdorff, L., & Wagner, C. S. (2009). Is the united states losing ground in science? a global perspective on the world science system. *Scientometrics*, 78(1), 23–36. <https://doi.org/10.1007/s11192-008-1961-y>
- Madsen, J. B. (2008). Semi-endogenous versus Schumpeterian growth models: Testing the knowledge production function using international data. *Journal of Economic Growth*, 13(1), 1–26. <https://doi.org/10.1007/s10887-007-9023-x>
- Marshall, A. (1890). *Principles of economics*. Macmillan.
- Meng, Y., & Romanowich, M. (2021). Catching up or falling behind? China’s changing science and technology trajectories. *Research Policy*, 50(9), 104333. <https://doi.org/10.1016/j.respol.2021.104333>
- Mokyr, J. (2002). *The gifts of Athena: Historical origins of the knowledge economy*. Princeton University Press.
- Moser, P., Voena, A., & Waldinger, F. (2014). German Jewish émigrés and US invention. *American Economic Review*, 104(10), 3222–3255. <https://doi.org/10.1257/aer.104.10.3222>
- Narin, F., Carpenter, M. P., & Woolk, N. C. (1976). Scientometric building blocks toward a theory of science. *Scientometrics*, 1(1), 37–44.
- Nelson, R. R., & Wright, G. (1992). The rise and fall of American technological leadership: The postwar era in historical perspective. *Journal of Economic Literature*, 30(4), 1931–1964.
- Pakes, A., & Griliches, Z. (1984). Patents and R&D at the firm level: A first look (Z. Griliches, Ed.). *R&D, Patents, and Productivity*, 55–72.
- Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967–1012. <https://doi.org/10.1111/j.1468-0262.2006.00692.x>
- Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. Harper & Brothers.
- Ülkü, H. (2004). R&d, innovation and economic growth: An empirical analysis [International Monetary Fund, Washington, D.C.]. *IMF Working Paper*, WP/04/185.
- von der Heyden, M., Heinze, T., Galvão de Almeida, R., & Breznau, N. (2025). Decline of German and rise of North-American hegemony in science: Insights from Nobel prize nominations (Physics/Chemistry, 1901–1969). *PLOS ONE*, 20(5), e0323103. <https://doi.org/10.1371/journal.pone.0323103>
- Waltman, L., Calero-Medina, C., Kosten, J., Noyons, E. C. M., Tijssen, R. J. W., Van Eck, N. J., Van Leeuwen, T. N., Van Raan, A. F. J., Visser, M. S., & Wouters, P. (2015). The Leiden Ranking 2011/2012: Data collection, indicators, and interpretation.

Journal of the American Society for Information Science and Technology, 63(12), 2419–2432.

Westerlund, J. (2007). Testing for error correction in panel data. *Oxford Bulletin of Economics and Statistics*, 69(6), 709–748. <https://doi.org/10.1111/j.1468-0084.2007.00477.x>

Zhou, P., & Leydesdorff, L. (2006). The emergence of China as a leading nation in science. *Research Policy*, 35(1), 83–104. <https://doi.org/10.1016/j.respol.2005.08.006>