Mapping the Market for Ideas in Europe, 1450–1650: A Title-Embeddings Approach*

Noel Johnson† Alexander Taylor‡

August 26, 2026

Abstract

Between 1450 and 1650 European print grew from a Mainz craft producing about 3,000 Latin editions per decade into a market of nearly 800,000 editions across seven major languages and 1,300 cities. Religious titles were the largest single category of that output, 215,000 editions and about 29 percent of classified titles, spanning Lutheran polemic from Wittenberg, Catholic devotional manuals from Lyon, and Calvinist catechisms from Geneva. Over the two centuries the religious share of European output roughly halved, peaking at half of all editions in the Luther decade of the 1520s. Did the expansion of the market for books and ideas reorganize this landscape, and did the reorganization matter for growth? Using title-level embeddings of the Universal Short Title Catalogue and a gravity-style city-period market access panel for six languages (Latin, German, Italian, French, Dutch, Spanish), we find that within-city increases in market access made a city's print output more topically diverse, less distinct from the European centroid, and more even, a pattern consistent with monopolistic competition under spatial product differentiation rather than with Smithian specialization. Decomposing the response by subject, rising market access is associated with a shift in city output away from religious and occasional print and toward jurisprudence, medicine, classical scholarship, and periodicals. In Latin, the cross-border language of European scholarship, the useful-knowledge component of diversity predicts subsequent 50-year city population growth while the ceremonial and devotional component does not, consistent with Mokyr's Industrial Enlightenment operating through trans-European scholarly print.

JEL codes: N13, N73, F12, R12, O33, Z12.

Keywords: market access; monopolistic competition; spatial product differentiation; useful knowledge; Industrial Enlightenment; word embeddings; USTC; early modern Europe; print; knowledge diffusion; religion; Reformation.

*The authors wish to acknowledge helpful comments from participants in the Markets and Society 2025 conference, the Mercatus AI Working Group Seminar, Mark Koyama, Desiree Desierto, Thomas Stratmann, Garett Jones, and John V.C. Nye. All errors are our own.

†George Mason University. Contact author. Email: njohnsoL@gmu.edu.

‡University of Evansville. Email: at154@evansville.edu.

1 Introduction

Between 1450 and 1650 European print transformed from a Mainz craft producing about 3,000 Latin editions per decade into a market of nearly 800,000 editions across seven major languages and 1,300 cities. Across that same window the geography of European urban populations changed fundamentally. Antwerp, Amsterdam, Paris, and London grew into commercial hubs while older centers stagnated. Did the expansion of the market for books and ideas change the way knowledge itself was organized, and if so, did that re-organization in turn shape which cities grew? Two questions, and the empirical literatures bearing on them have built up in parallel without quite meeting. This paper documents a joint chain running from market access, through the composition of printed material, to subsequent city growth. We show that the relationship between print content and city growth is driven by titles that reflect Mokyr-style useful-knowledge that is prescriptive, instrumental, or scholarly.

Religion sits at the center of this transformation. Religious titles are the largest single category of European print in our data, 215,000 editions and about 29 percent of classified output, spanning Lutheran polemic from Wittenberg, Catholic devotional manuals from Lyon, and Calvinist catechisms from Geneva. That share was not constant. It peaked at roughly half of all editions in the 1520s, the decade in which cheap print carried the Luther affair across the Empire (Rubin, 2014), and then fell by roughly half, from 47 percent of European output in the 1480s to 21 percent by the 1640s. Our results speak directly to this secularization of the printed book market. The subject-level decomposition of our main findings (Appendix D) shows that rising market access is associated with precisely this compositional shift within cities, away from religious and occasional print and toward jurisprudence, medicine, classical scholarship, and periodicals, and Section 9 shows that it is the useful-knowledge component of a city's print mix, not the ceremonial and devotional component, that predicts subsequent city growth. Market integration, on this evidence, is associated with the reallocation of Europe's printed output away from religion, which places the mechanisms we study inside the economics of religion as much as inside the economics of knowledge.

Three theoretical traditions offer different answers as to the relationship between market expansion for print and how print content should evolve. Smith (1776), formalized by Stigler (1951),

predicts that "the division of labor is limited by the extent of the market". Larger markets enable finer specialization, and one expects high-market-access cities to develop distinctive niches in their print composition. Mokyr (2005, 2002, 2009, 2017) centers on a different mechanism. Variety and integration in the European knowledge stock, the "Republic of Letters", drive growth via knowledge spillovers, so cities whose print mix is more diverse or more transnational ought to grow faster. Direct empirical tests of the Mokyr useful-knowledge channel include de la Croix et al. (2018) (apprenticeship institutions and growth in preindustrial Europe), Kelly et al. (2023) (mechanical human capital and the British Industrial Revolution), and de la Croix et al. (2025), who reconstruct the scholar-affiliation network of ∼200 European universities and 150 academies (1084–1793) and show that simulated network exposure predicts city-level adoption of Botanical Realism (botanic gardens), mathematical astronomy (observatories), and Protestantism. New trade theory predicts the opposite of Smith on print composition, both the monopolistic-competition tradition of Dixit and Stiglitz (1977), Krugman (1979, 1980), and Helpman and Krugman (1985), and the spatial product-differentiation literature of Hotelling (1929), Lancaster (1979), and Anderson et al. (1992). High-market-access cities support more firms producing more varieties, but those varieties are positioned in product space toward the broad European mainstream of demand, not into distinctive local niches. The same framework predicts that, in equilibrium, market access is the sufficient statistic that jointly determines both a city's production composition (the variety, position, and evenness of its print output) and its growth dynamics. Composition and growth are co-determined by MA rather than causally chained through composition into growth.

We construct a city-level market access (MA) measure following Donaldson and Hornbeck (2016). A gravity-style, population-weighted index of a city's exposure to other European cities, with bilateral transport costs computed via least-cost paths over Roman roads, medieval routes, rivers, and seas. MA is conceptually access to trade outside the city. Own-city size enters our regressions separately through within-city demeaning or as an explicit log(pop) control. Identification exploits within-city variation in MA across four 50-year anchor years, in a two-way fixed-effects framework that absorbs unobserved time-invariant city characteristics and common period shocks. MA is itself the gravity-style sufficient statistic for a city's gains from trade in the modern quantitative-trade literature (Eaton and Kortum, 2002; Arkolakis et al., 2012; Allen and Arkolakis, 2014; Donaldson and Hornbeck, 2016; Donaldson, 2018), so its within-city variation supplies a

theoretically grounded treatment for print composition rather than an arbitrary regressor. We do not claim a water-tight causal identification, but the design approaches a reduced-form structural reading of the trade-and-production literature rather than a purely descriptive cross-section.

A paired example illustrates the patterns we document. Frankfurt am Main, 1580–1620, sat near the top of the European MA distribution. A free imperial city at the confluence of major Rhine and Main trade routes and home of the continent's leading book fair. Its print output in these decades was a consolidated humanist-scientific Latin academic mix (*Florilegium Magnum seu Polyanthea*, *Rerum Persicarum Historia*, *Imperatorum Romanorum Numismatum Series*, Paracelsus's *Operum Medico-Chimicorum*), content that could appear, with the printer's imprint changed, in Antwerp, Leiden, Basel, or Venice.[1] Frankfurt's print mix had homogenized toward a pan-European mainstream and the city's population grew through the period. Tübingen, 1550–1599, by contrast, sat in the lower tier of MA, a modest Swabian university town in the interior southwest, away from major commercial corridors. It printed roughly 800 German titles in this window plus substantial additional Latin output from its theological and medical faculties: Lutheran scholastic disputations from the Stift, Hebrew and biblical philology, university medical dissertations, and astronomical-mathematical works (Michael Mästlin, who taught the young Kepler, was on the faculty). The print mix was recognizably local, concentrated on a distinctive academic-confessional niche one does not find replicated in the Frankfurt, Antwerp, or Venice catalogs. Tübingen's population remained modest throughout. Two cities, different MA positions, two different print-composition signatures, and city growth that tracks size and MA rather than the distinctiveness of print composition.

The paper's empirical core makes this pattern systematic. We find three robust within-city print-composition findings (Section 8). As a city's MA rises within its own time series its print output becomes (1) more *diverse* in within-city topical content, (2) *less distinct* from the contemporaneous European centroid in title-embedding space, and (3) more *even* in its distribution across topical clusters. All three signs are consistent with new trade theory under spatial product differentiation and inconsistent with the Smithian extent-of-the-market prior. Section 8's pattern, convergence

[1] *Florilegium Magnum seu Polyanthea* is a humanist anthology of classical and patristic citations organized by topic, a genre common to early-modern European scholarship. *Rerum Persicarum Historia* is a history of Persia. The *Imperatorum Romanorum Numismatum Series* is a numismatic catalog of Roman emperors from Julius Caesar to Rudolph II (whose reign ended 1612). The *Operum Medico-Chimicorum* is an authoritative collection of Paracelsus's medical-alchemical writings. Each is representative of a genre that circulated across the major European print centers at the turn of the seventeenth century rather than being distinctively Frankfurt-specific.

of print composition toward a European mainstream as MA rises, is also qualitatively the kind of European-integration dynamic Mokyr's account emphasizes.

We also test whether print composition predicts subsequent 50-year city population growth (Section 9) conditional on market access and initial size. Two compositional signals come through Mokyr-positive. First, the share of titles in useful-knowledge classifications (medical texts, jurisprudence, philosophy and morality, academic dissertations, periodicals, educational books, history and chronicles, political tracts, classical authors) is a robust positive predictor of subsequent growth. Second, within-city diversity in the Latin subset, the language of trans-European scholarly print, predicts growth strongly and positively. Both findings are consistent with the Mokyr "Industrial Enlightenment" hypothesis that the composition of useful versus ceremonial content, rather than total print volume, carries information about subsequent growth. The magnitudes are modest (roughly twelve log-points of additional growth over the 50-year window when the useful share moves from zero to one), and the channel may operate primarily through the transnational Latin scholarly medium rather than through vernacular print, but the sign and the robustness pattern point in the same direction.

The first wave of printing-and-society research, exemplified by Elizabeth Eisenstein's account of print as an agent of intellectual change, argued that the press first broadened European print culture through the sheer volume and variety of titles it made available, then deepened the same culture as scholars and printers consolidated knowledge into canonical reference works during the late sixteenth and seventeenth centuries (Eisenstein, 1980). Our findings present a different empirical pattern. As a city's market access rose, its print output became broader (more diverse) and more mainstream (less distinctive), not narrower and more specialized. The "deepening" phase Eisenstein describes does not appear as a within-city tightening of topical concentration in our panel, but rather the opposite.

A more recent literature establishes the printing press as a driver of city-level economic outcomes. Dittmar (2011) shows that early adoption of the press predicted city-population growth between 1500 and 1600. Rubin (2014) shows that cities with a press by 1500 were substantially more likely to adopt the Reformation by 1600, the canonical demonstration that the print market shaped Europe's religious landscape. Dittmar and Seabold (2019) extend the analysis to printer competition and the diffusion of Reformation ideas. Boerner et al. (2021) link printing-press diffusion

to medium-run European economic growth. Taylor and Hall (2026) use a geographic regression-discontinuity around the French border to show that the 1539 Villers-Cotterêts ordinance accelerated the shift of French print from Latin to the vernacular and fostered French national identity.

A parallel "language as data" literature uses textual analysis of historical corpora to test cultural and intellectual change directly. Grajzl and Murrell (2024) apply structural topic modeling to 57,863 English print volumes from 1530 to 1700 to study the coevolution of ideas on religion, science, and institutions in the run-up to industrialization. Almelhem et al. (2026) apply latent Dirichlet allocation to 264,443 volumes published in England between 1500 and 1900 from the HathiTrust Digital Library and document that volumes using language at the nexus of science and political economy became progress-oriented during the Enlightenment, with industrial-flavoured works at that same nexus carrying the strongest Mokyr signal.

Our contribution extends both literatures along three margins. From the extensive margin (whether a city had a press) to the intensive / compositional margin (what each city printed). From single-language, national-level, case studies to a cross-language comparison of six European print markets. And from the eighteenth- and nineteenth-century post-Enlightenment data on which the textual literature concentrates to the 1450–1650 print-revolution period in which the Mokyr channel's trans-European Latin scholarly medium first took shape.

2 Why title embeddings?

In order to measure the content of printed material coming off the press, we construct title embeddings. Embeddings represent terms as vectors in a high-dimensional Euclidean space such that semantically related terms appear near one another. The basic idea, popularized by Mikolov et al. (2013) and Pennington et al. (2014) and reviewed in Levy and Goldberg (2014), is that the distribution of words across documents carries semantic information. Words that appear in similar contexts (e.g. *medicus* and *chirurgia*; *rex* and *imperator*) tend to share meaning. By representing each word as a vector whose dimensions encode this distributional information, embeddings turn a discrete vocabulary into a continuous geometric object that admits arithmetic and distance computations.

Word embeddings are now standard in computational linguistics and across the social sciences. They have been used in economics for firm and patent similarity (e.g., Hoberg and Phillips, 2010;

Kelly et al., 2021), in political science to track ideological positioning of speeches and platforms, and in history to trace conceptual change in large text corpora over time (e.g., Hamilton et al., 2016). They belong to the broader "text as data" methodological tradition reviewed in Gentzkow et al. (2019) which also includes topic-modeling methods (latent Dirichlet allocation, structural topic models) of the kind Grajzl and Murrell (2024) and Almelhem et al. (2026) apply to early modern English print. Our application takes the embeddings tool to an under-used historical setting, the 800,000 short titles of the USTC.

The pipeline we use is the standard one. Per language, we construct a co-occurrence matrix from the USTC title corpus using a 5-word skip-gram window, transform it to positive pointwise mutual information (PPMI) (Church and Hanks, 1990; Levy and Goldberg, 2014), and reduce it to 200 dimensions via truncated singular-value decomposition. We then aggregate word vectors to title vectors using the smooth-inverse-frequency (SIF) procedure of Arora et al. (2017), a weighted average of the title's constituent word vectors (down-weighting common words via a smoothing constant $a = 10^{-3}$) followed by removal of the first principal component (which captures shared syntactic structure rather than semantic content). The result is a 200-dimensional vector per title, comparable across titles within the same language. The full set of pipeline parameters (window size, vocabulary thresholds, SIF constant, k-means K) is collected in Appendix Table 8.

We measure similarity between titles in the embedding space using cosine similarity which is the cosine of the angle between two vectors: $\cos(u, v) = (u \cdot v)/(\|u\| \, \|v\|)$, with $\cos = 1$ for identical direction and $\cos = -1$ for opposite. We use cosine rather than Euclidean distance for two reasons. First, cosine similarity is invariant to vector magnitude. Two titles with the same topical mix but different total word counts (and hence different vector norms after SIF weighting) have the same cosine similarity, whereas their Euclidean distance differs mechanically with the norm. Second, cosine is the standard measure in the NLP literature that produced these embedding methods, so our metrics are directly comparable to the existing tooling. The cosine distance $d_{\cos}(u, v) = 1 - \cos(u, v)$ inherits these properties and ranges from 0 (identical) to 2 (opposite).

Why do we use embeddings instead of more traditional subject classifications to measure content? The USTC supplies a `classification1` field with approximately twenty top-level categories (Religious, Academic Dissertations, Ordinances and Edicts, Jurisprudence, Periodicals, News Books, Political Tracts, Poetry, Literature, Music, Drama, Calendars and Prognostications,

Funeral orations, Wedding pamphlets, History and Chronicles, Medical Texts, Classical Authors, Philosophy and Morality, Educational Books, and a small number of others). The largest, Religious, contains 215,000 titles (29.4% of those with non-missing classification) and includes Lutheran polemic from Wittenberg, Catholic devotional manuals from Lyon, Calvinist catechisms from Geneva, anti-Trinitarian treatises from Kraków, and Greek-Hebrew biblical commentaries from Basel, content that occupies radically different positions in any coherent measure of intellectual content. Treating these as a homogeneous category would erase precisely the kind of variation we are trying to measure.

Embeddings give traction on three issues that classification metadata cannot:

1. *Within-classification heterogeneity.* Consider two Latin USTC "Religious" titles from the famous Luther–Erasmus debate of the mid-1520s, Erasmus's *De libero arbitrio diatribe sive collatio* (Antwerp, 1524) and Luther's reply *De servo arbitrio* (Wittenberg, 1525). Both are tagged Religious. Both address theological disputes about the will. Yet they sit far apart in the Latin embedding space because the Lutheran-polemic vocabulary differs systematically from the Erasmian-humanist vocabulary. Within-classification cosine distances between such title pairs routinely exceed 0.5, an order of magnitude larger than distances between thematically tight pairs (e.g. two Wittenberg anti-papal pamphlets of the same year, $d_{\cos} \approx 0.05$). The discrete `classification1` field flattens this variation.

2. *Cross-language alignment.* Each language's embedding space is trained separately on that language's co-occurrence statistics, so the Latin space and the Italian space are not the same 200-dimensional vector space, and a Latin title's vector is not directly comparable to an Italian title's vector. What is comparable across languages is the set of city-level summary metrics defined below. Each metric is a distance computed within the relevant language's embedding space, and the resulting scalar is mathematically commensurable across languages once language fixed effects absorb cross-language scale differences. A city that printed in multiple languages (e.g. Paris, which printed in Latin, French, and small amounts of Italian and German) is treated as one observation per (city, language, period) cell. The regression's language fixed effect ensures that each language's metric distribution contributes on its own scale.

3. *Continuous-space metrics.* The three city-level summary metrics we construct (diversity, distinctiveness, and specialization-entropy, formally defined in Section 2 below) are well-defined, real-valued quantities at the city-decade level. Discrete classifications support only simple counts and proportions. The continuous metrics let us track how a city's overall topical mix evolves, not just whether one or another category was present.

Figure 1 makes (1) and (3) concrete for the Latin corpus. For three seed words (*anatomia*, *geometria*, and *ius*), the figure shows the ten nearest words in the 200-dimensional embedding space, ranked by cosine similarity to the seed. *Anatomia* lies closest to *corporis*, *humani*, *anatomicae*, *fabrica*, and *foetus*. *Geometria* lies closest to *arithmetica*, *euclidis*, *mathematica*, *astronomia*, and *optica*. *Ius* lies closest to *civile*, *feudale*, *iustinianeum*, *pontificium*, and *municipale*. The embedding has, from co-occurrence statistics alone, recovered the substantive structure of three intellectual domains: anatomical vocabulary, the mathematical sciences, and a typology of law.

None of these clusters is recoverable from the discrete USTC classification. Anatomical vocabulary is subsumed within `Medical Texts`, which mixes anatomical treatises with surgical manuals, pharmaceutical handbooks, and Hippocratic commentaries. The mathematical sciences are scattered across `Academic Dissertations`, `Educational Books`, and `Classical Authors` (Euclid lives in the last category alongside Cicero). The law typology folds canon law, civil procedure, and customary-law compilations together under a single `Jurisprudence` tag. A complementary appendix figure (Figure 17) projects a stratified sample of 5,000 Latin titles into two-dimensional UMAP space, colored by USTC `classification1`, and shows that classifications occupy partly-overlapping regions of the continuous space. Academic Dissertations form a distinctive cluster, but Religious titles scatter across most of the manifold rather than forming a tight region. The same nearest-neighbor exercise of Figure 1 carried out for the five vernacular languages (German, Italian, French, Dutch, Spanish) reproduces the Latin pattern in each (Figures 12, 13, 14, 15, and 16 in Appendix A.13). The embedding pipeline is not Latin-specific.

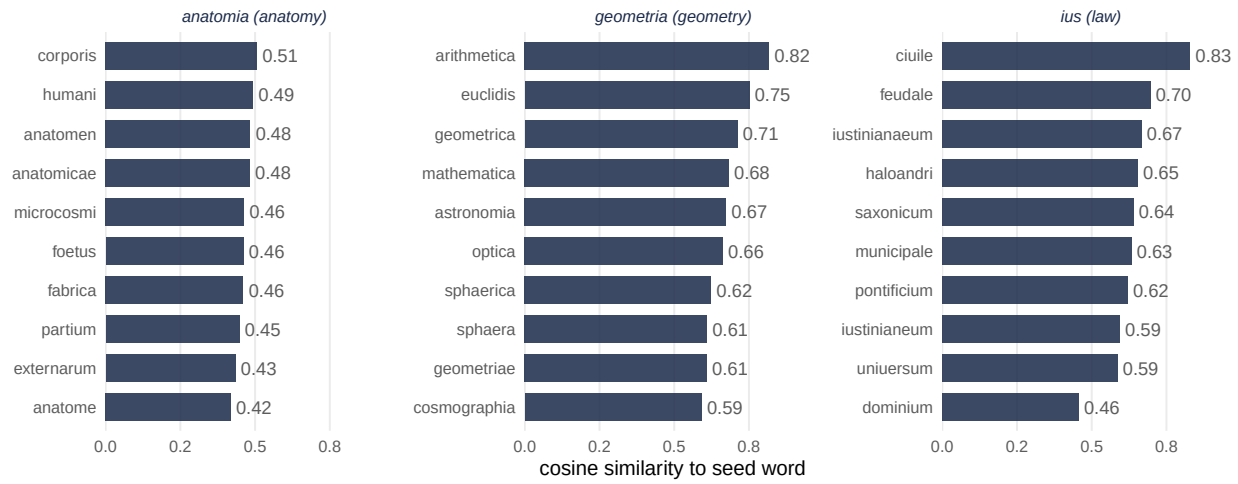


Figure 1: Ten nearest words to each of three Latin seed terms in the 200-dimensional embedding space, ranked by cosine similarity. The seeds are *anatomia* (anatomy), *geometria* (geometry), and *ius* (law). Each seed's neighborhood reproduces a recognizable intellectual domain—anatomical vocabulary, the mathematical sciences, and a typology of law—from co-occurrence statistics alone.

The same nearest-neighbor exercise carried out at the title level instead of the word level is also informative and makes explicit how the continuous metric complements the discrete USTC subject-classification field. Figure 2 shows the eight nearest titles in the Latin title-embedding space for two iconic seed works, with each neighbor shaded by its USTC classification. The top-panel seed is William Harvey's *Exercitatio anatomica de motu cordis* (Leiden, 1639), the founding text of the modern theory of blood circulation. The bottom-panel seed is Tycho Brahe's *Astronomiae Instauratae Progymnasmata* (Frankfurt, 1610), the posthumous synthesis of his observational program.

The Harvey neighborhood includes near-edition reprints of Harvey himself (Amsterdam 1645 and Leiden 1639) at cosine similarity ≈ 1, plus a Rotterdam 1648 edition at 0.96. These three USTC records are split across `classification1` categories. Two are tagged *Astrology and Cosmography* (a USTC misclassification, plausibly driven by "anatomica de motu" being keyed off "motu") and only the Rotterdam edition is correctly tagged *Medical Texts*. The remaining four neighbors are independent works by other authors on the same topic (Groningen and Leiden academic disputations on *motus cordis* and *circulatio sanguinis*, a 1647 synthesis volume *Recentiorum disceptationes de motu cordis*, and an explicit Anti-Harveian rebuttal) at similarity 0.78–0.83. These neighbors are tagged *Medical Texts* or *Academic Dissertations*. So the eight nearest titles to Harvey span three

different USTC classifications but are, substantively, all about the circulation of blood.

The Brahe neighborhood is similar. The first four neighbors are reprints of Brahe's *Progymnasmata* (Frankfurt 1619, Frankfurt 1610, Prague 1603, Uraniborg 1603) and the 1648 *Opera Omnia* that incorporates it, at similarity 0.84–0.99. These are split between *Astrology and Cosmography* and *Academic Dissertations*. The remaining three are independent astronomical works (a Leiden 1631 *Uranometriae libri tres*, a Leipzig 1499 *Tabulae solis et lunae coniunctionum*, and a Strasbourg 1582 *Ephemeris caelestium motuum*), classified as *Astrology*, *Science*, and *Calendars and Prognostications* respectively. So the eight nearest titles to Brahe span four different USTC classifications but are, again, all observational astronomy.

This exercise is a sanity check rather than evidence for any of the substantive claims that follow, but it makes the Section 2 argument concrete. Title-level cosine distance picks up both edition-equivalent reprints of the same work and topically related works by other authors, including across USTC classification boundaries that fragment what the embedding treats as a single intellectual neighborhood. Working from `classification1` alone the eight titles in either neighborhood would be three or four unrelated subject-buckets. Working from cosine similarity, they are visibly one substantive cluster.

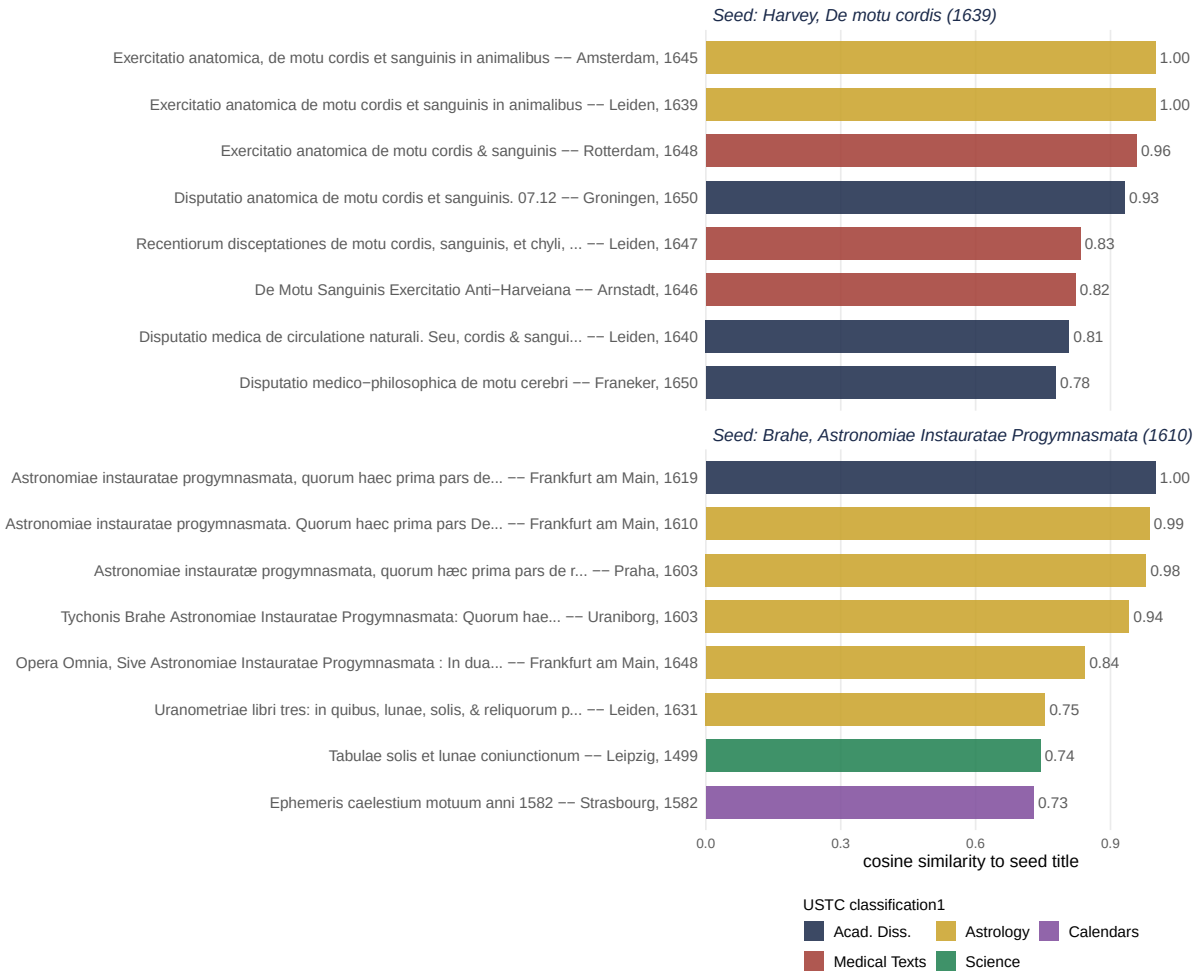


Figure 2: Eight nearest titles to two iconic Latin seed works in the 200-dimensional title-embedding space, ranked by cosine similarity to the seed. Bars are colored by USTC `classification1`. Top panel: William Harvey's *Exercitatio anatomica de motu cordis et sanguinis in animalibus* (Leiden, 1639). The eight neighbors span three different USTC classifications (Astrology, Medical Texts, Academic Dissertations) but are substantively all about the circulation of blood. Bottom panel: Tycho Brahe's *Astronomiae Instauratae Progymnasmata* (Frankfurt, 1610). The eight neighbors span four different USTC classifications (Astrology, Academic Dissertations, Science, Calendars and Prognostications) but are substantively all observational astronomy. The continuous metric finds the cluster the discrete classification splits.

Three composition metrics from title embeddings

Building on the title embeddings introduced above, the empirical analysis in the rest of the paper rests on three city-level summary metrics. They are not standard tools in the economic-history reader's toolkit, so we define each formally here, give an intuition for what it captures, and provide a concrete example from the data. Appendix Table 10 reports per-language and pooled descriptive

statistics for the six-language base panel. The same three metrics underlie every regression in Sections 8 and 9.

Throughout, let $v_i \in \mathbb{R}^{200}$ denote the SIF title embedding for title i in the embedding space of its language ℓ. For a (city, language, 50-year period) cell let $\mathcal{T}_{c,\ell,p}$ be the set of titles published in city c, language ℓ, and period p, and let

$$\bar{v}_{c,\ell,p} \;=\; \frac{1}{|\mathcal{T}_{c,\ell,p}|} \sum_{i \in \mathcal{T}_{c,\ell,p}} v_i$$

be the city-period centroid. Let $\bar{v}_{\text{Eur},\ell,p}$ be the analogous centroid computed over all titles in language ℓ printed anywhere in Europe in period p.

Diversity. The diversity of a city-period cell is the average cosine distance between its constituent titles and its own centroid:

$$D_{c,\ell,p} \;=\; \frac{1}{|\mathcal{T}_{c,\ell,p}|} \sum_{i \in \mathcal{T}_{c,\ell,p}} d_{\cos}(v_i,\ \bar{v}_{c,\ell,p}).$$

Substantively, D measures *within-city topical spread*. A city printing across many subject areas of ℓ will have titles scattered widely around its centroid and therefore high D, while a city printing a tight thematic niche will have its titles clustered near the centroid and low D. Across the six-language panel D has mean about 0.59–0.63 and standard deviation 0.11–0.15. For a concrete example, Paris in the 1600 decade printed 1,066 Latin titles spanning humanist literature, medical and natural-philosophical texts, jurisprudence, devotional and Counter-Reformation theology, and academic disputations, and posts $D_{\text{Paris, lat, 1600}} = 0.82$ (top of the per-decade distribution). Coimbra in the same decade printed 74 Latin titles dominated by Jesuit theology and Iberian scholasticism, with $D_{\text{Coimbra, lat, 1600}} = 0.45$ (near the bottom). The metric reflects what the casual reader would also see by skimming the catalogs.

Distinctiveness. The distinctiveness of a city-period cell is the cosine distance between its centroid and the contemporaneous European centroid for the same language:

$$X_{c,\ell,p} \;=\; d_{\cos}\big(\bar{v}_{c,\ell,p},\ \bar{v}_{\text{Eur},\ell,p}\big).$$

X measures *how far the city's typical title sits from what the European literate public in that language was reading on average* in that period. A city whose print mix matches the cross-European mean will have small X. A city whose print mix is positioned in a niche of its language's product space, far from where the demand-weighted mass of the continent lies, will have large X. Across the panel X has mean about 0.59–0.74 and standard deviation around 0.20, with observed values roughly in $[0, 1.4]$ (the theoretical range is $[0, 2]$). For a concrete example, take Latin in the 1600 decade. Wittenberg, the largest Latin-language printing center that decade (2,133 titles, heavily weighted toward Lutheran scholastic disputations and academic Latin), defines the bulk of the European Latin mass for that period and posts $X_{\text{Wittenberg, lat, 1600}} = 0.45$ (low). Venice, with 1,414 Italian-Catholic and humanist Latin titles in the same decade, sits in a quite different region of the Latin product space and posts $X_{\text{Venice, lat, 1600}} = 1.14$ (high). The metric is silent on which center is "representative" in any absolute sense. It asks only how far a given city sits from the demand-weighted European average.

Specialization-entropy. The third metric is the Shannon entropy of the city-period titles' assignments to $K = 10$ topical clusters obtained from k-means run on the full title-embedding corpus of language ℓ:

$$H_{c,\ell,p} \;=\; -\sum_{k=1}^{K} s_{c,\ell,p,k} \log s_{c,\ell,p,k},$$

where $s_{c,\ell,p,k}$ is the share of titles in $\mathcal{T}_{c,\ell,p}$ assigned to cluster k. H takes values on $[0, \log K] \approx [0, 2.30]$ and measures *how evenly the city's titles are distributed across topical clusters* (high H) versus how concentrated they are in one or two clusters (low H). Across the panel H has mean about 0.59–0.92 and standard deviation around 0.35. For a concrete example, take Latin in the 1600 decade. Coburg, a small Saxon imperial-city press, printed 197 Latin titles distributed evenly across nine of the ten clusters and posts $H_{\text{Coburg, lat, 1600}} = 1.61$ (near the maximum). Louvain, by contrast, printed 83 Latin titles overwhelmingly concentrated in one cluster (Counter-Reformation Catholic theology) and posts $H_{\text{Louvain, lat, 1600}} \approx 0$. The metric is independent of within-cluster heterogeneity, which D captures, and of position in the product space, which X captures. A city can have high H (broad even spread across clusters) while still having low D (tight thematic concentration within each cluster), or high X (far from the European mean) while having low H (one dominant cluster). The three metrics are conceptually complementary.

Whether these metrics capture intellectually meaningful structure beyond the formal definitions is itself a testable question, which we revisit in Section 6 (cross-language amplitudes, time evolution, and within-vs-between variance decomposition).

3 The economic geography of European print

The USTC corpus covers 788,299 titles across 1,339 European cities between 1450 and 1650, in seven major languages: Latin (314K), German (140K), French (105K), Italian (67K), English (52K), Dutch (45K), and Spanish (42K). See Figure 3 and Table 1.

The geography of print is highly heterogeneous across languages, and this matters for our research design (Table 1).

Table 1: Print regimes by language: title and city counts, top-city share, implied regime classification.

Language	N titles	N cities	Top-city share	Implied regime
Latin	301,216	873	Paris 9.0%	polycentric
German	117,223	459	Leipzig 10.4%	polycentric
Italian	65,124	259	Venice 34.7%	semi-polycentric
French	94,447	366	Paris 53.6%	semi-polycentric (Paris-dominant)
Dutch	39,931	167	Amsterdam 29.1%	semi-polycentric
Spanish	32,752	234	Madrid 24.0%	semi-polycentric
English	51,502	109	London 85.1%	strongly monocentric

Note: N titles counts editions with an identified place of printing, the relevant base for a city-level analysis; the full per-language totals including editions with no recorded place are larger (Latin 314,172, German 140,074, French 104,729, Italian 66,655, English 52,266, Dutch 44,783, Spanish 42,448). Latin, German, Italian, French, Dutch, and Spanish are the six panel-credible languages used in the regression analysis; English is excluded as monocentric (see text below the table).

Our econometric design identifies the within-city slope of each composition metric on $\log(\text{MA})$ from within-city variation across four 50-year periods. For a language with a single dominant producing city, two practical issues arise. First, the within-city slope is identified essentially from a single city's trajectory, which is informative about that city but not generalizable to the language's

panel. Second, the distinctiveness metric—cosine distance from the city centroid to the contemporaneous European centroid for the same language—loses much of its content when one city accounts for the bulk of titles, because the city centroid and the European centroid mechanically converge. London accounts for 85% of all English titles between 1450 and 1650, so the English centroid effectively is London's centroid, and the distinctiveness metric carries little independent variation. Paris's 54% share of French titles is high but well short of the English regime. French has 366 distinct printing cities, including Lyon (the second-largest French printing center after Paris), Geneva, Rouen, and a substantial Catholic-Reformation network of provincial presses.

The six polycentric and semi-polycentric languages—Latin, German, Italian, French, Dutch, and Spanish—give us enough cross-city variation to identify panel coefficients on log(MA) and enough within-language heterogeneity for the distinctiveness metric to carry meaningful content. We restrict the headline analysis to these six. English is excluded from the regression panel.[2]

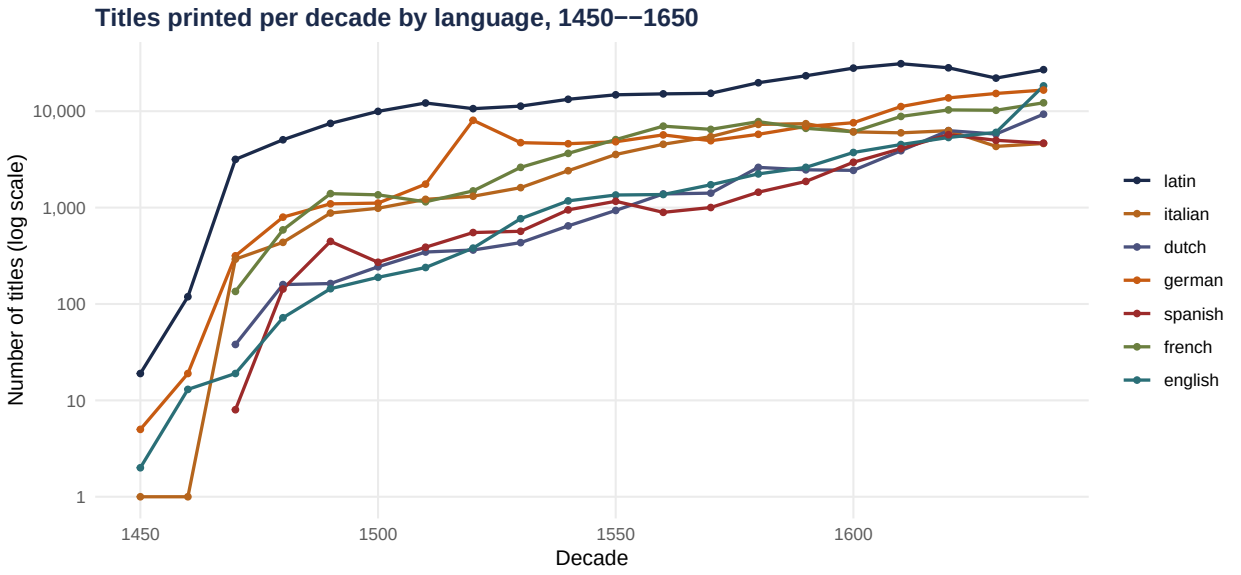


Figure 3: Titles printed per decade by language, 1450–1650. Latin and German rise steadily through the sixteenth century; Italian and French begin earlier and grow more modestly; English remains small until 1620 when it accelerates.

[2]London's 85% share is high enough that the English panel reduces, in effect, to a single-city case study with the distinctiveness metric mechanically near zero. English appears in cross-language descriptive panels where appropriate.

4 Why market access? Theory and literature

The treatment variable we focus on in this paper is market access. The market access framework, formalized in Donaldson and Hornbeck (2016) and Donaldson (2018), represents a city's economic centrality as the population-weighted sum of its accessibility to other places:

$$\text{MA}_i = \sum_{j \neq i} \text{pop}_j \, \tau_{ij}^{-\theta},$$

where τ_{ij} is the bilateral cost of transport between cities i and j and θ is a trade elasticity.[3] Recent work in economic history has used market access to explain U.S. railroad expansion (Donaldson and Hornbeck, 2016), Indian railroads (Donaldson, 2018), and early modern European trade (Pascali, 2017).

For the present paper, market access is an attractive treatment variable for three reasons.

It captures economic exposure to the rest of the European urban system. A city with high MA is connected to many populous places at low transport cost. It is "central" in the early modern urban network in the precise sense that gravity-trade theory makes operational. The book trade depended directly on this connectivity. Books were exchanged at fairs, transported by river and road, and sold by merchants who themselves traveled. Print output reflects the information environment a city inhabits and that environment is a function of market access.

It varies non-trivially over time as populations shift. Between 1450 and 1650, Paris grew from ∼225K to ∼415K, London from ∼50K to ∼400K, Amsterdam from ∼3K to ∼175K, while Bruges declined from ∼125K to ∼35K. These shifts feed mechanically into all other cities' MA, generating within-city variation that is plausibly less endogenous to a city's own knowledge composition than the city's own population is.

[3] The trade elasticity θ governs how responsive bilateral trade is to bilateral transport cost. A high θ means trade falls off sharply with cost (cost-sensitive goods, substantial substitution between trade partners), and a low θ means trade is relatively insensitive to cost (high-value, low-bulk goods or those with few substitutes). Printed books are at the low-θ end of this spectrum. They were high-value relative to bulk (a printed octavo was light and durable enough to ship by river barge or pack-mule with little cost penalty), and few close substitutes existed for a given title, so a small shift in transport costs did not redirect trade away from established publishing centers. We discuss the choice of θ and report robustness to alternative values in Appendix A. The MA framework is calibrated in the modern literature on commodities and labor rather than on high-value, low-bulk goods like printed books. We therefore interpret log(MA) here as a reduced-form proximity-to-trade measure rather than as a structural trade-cost shifter, and read the within-city slopes as associations rather than as parameters of a calibrated model.

It connects directly to the literature on print and growth. Dittmar (2011) showed that early adoption of the printing press predicted city growth between 1500 and 1600. Subsequent work (Dittmar and Seabold, 2019; Boerner et al., 2021) has refined the channel. Our paper extends this literature from the extensive margin (did the city have a press?) to the intensive/compositional margin (what did the press print?). Market access is the natural explanatory variable for both, which keeps our framework comparable.

The theoretical anchor for our empirical analysis combines two well-developed literatures. The first, the monopolistic-competition trade theory of Dixit and Stiglitz (1977), Krugman (1979, 1980), and Helpman and Krugman (1985), delivers the home-market effect (larger markets support more firms producing more differentiated varieties). The second, the spatial product differentiation literature of Hotelling (1929), Lancaster (1979), Salop (1979), and Anderson et al. (1992), adds the structure that lets us speak about where in product space a firm chooses to position its variety. Appendix B develops the combined framework in full. Here we sketch only what is needed to interpret the empirical results in Section 8.

The framework rests on five identifiable assumptions, each of which can be contested separately on historical grounds. We label them A1–A5 and develop them formally in Appendix B. Informally, they are: A1, consumers value variety (CES aggregation across titles); A2, printing has increasing returns at the title level (fixed typesetting cost per title); A3, consumer preferences vary across cities in a positional way (Italian readers want different titles from Polish readers, in a sense that admits a distance metric in title space); A4, the early modern book trade had bilateral transport costs with a proximity advantage (closer markets weigh more in a city's effective demand); and A5, printers chose product positions strategically (they did not all print the same canon, and they responded to what other printers were producing). Each assumption contributes to a distinct empirical prediction, so the framework decomposes into three signed predictions whose required assumptions differ.

In what follows we denote the estimated within-city slope of composition metric m on $\log(\text{MA})$ as $\hat{\beta}_{\text{m}}$, with m ranging over diversity, distinctiveness, and specialization-entropy. The empirical specification that delivers these slopes is developed in Section 7. Here we simply state the signed predictions.

1. *Diversity rises with MA.* A1 + A2 are sufficient. Love of variety plus increasing returns at the title level delivers the home-market effect (more firms in larger markets, each producing

one differentiated variety). This prediction holds even with symmetric CES preferences and no positional differentiation. We expect $\hat{\beta}_{\text{diversity}} > 0$ in within-city panels.

2. *Distinctiveness from the European centroid falls with MA.* A3 + A4 + A5 are additionally required. Positional preferences (A3) and proximity-weighted demand (A4) make a high-MA city's effective demand tilt toward the European mainstream, and strategic positioning (A5) means firms in such cities locate near that demand-weighted centroid. We expect $\hat{\beta}_{\text{distinctiveness}} < 0$ in within-city panels.

3. *Topical entropy rises with MA.* A5 is the binding additional assumption. Within a high-MA city, firms strategically differentiate from each other to avoid head-to-head price competition, and with more firms in the city (via Prediction 1) the distribution of titles across topical clusters becomes more even. We expect $\hat{\beta}_{\text{spec-entropy}} > 0$ in within-city panels.

The decomposition is useful because the three predictions are not redundant. Each can fail on its own and the pattern of which fails diagnoses which assumption is binding for the empirical setting. A near-zero $\hat{\beta}_{\text{diversity}}$ would point to A1 or A2 failing (book demand is authority-seeking rather than variety-loving, or printing is constant-returns). A near-zero or wrong-signed $\hat{\beta}_{\text{distinctiveness}}$ points to A3 or A4 failing (homogeneous European preferences, or segmented bilateral trade). A near-zero $\hat{\beta}_{\text{spec-entropy}}$ points to A5 failing (printers did not strategically differentiate, perhaps regulation or guild structure pinned positions). Appendix B (Table 16) works through each conditional prediction explicitly.

These three signs are what we report in Section 8. The Smithian *Wealth of Nations* prediction, "the division of labor is limited by the extent of the market", would imply the opposite signs on all three. Specialization should tighten the within-city mix (negative diversity slope), push cities into distinctive local niches (positive distinctiveness slope), and concentrate output on a narrower set of topical clusters (negative entropy slope). Stigler (1951)'s own formalization of the extent-of-the-market doctrine in fact leaves room for increasing returns alongside finer specialization, so what we test here is the finer-specialization prediction specifically rather than Smithian thought as a whole. We are not simply contradicting Smith. We are reporting evidence consistent with the modern monopolistic-competition-plus-spatial-differentiation refinement of Smith, which differs in exactly these three predictions. Because our identifying within-city variation in $\log(\text{MA})$ is thin

(Section 6 reports 11–22% of total variance) and partly tracks the rising NW European complex over the panel window, we read the resulting within-city slopes as reduced-form associations consistent with the A1–A5 framework rather than as evidence that any particular assumption in the framework is separately identified. The Mokyr (2002, 2009, 2017) civilizational account of European knowledge integration is qualitatively consistent with the same pattern (mainstream-convergent variety expansion as the European republic of letters integrates). We treat Mokyr as the broader historical context within which the production-side mechanism plays out.

5 Constructing market access

We construct city-level log market access at the four 50-year anchor years 1450, 1500, 1550, and 1600 that bracket the four print windows used in the regression panel (1450–1499, 1500–1549, 1550–1599, 1600–1649). The empirical USTC corpus is cut at 1649 because years 1650–1656 each have fewer than five recorded titles.[4] The MA construction follows the standard procedure. Population data come from the Buringh European Urban Population dataset, which provides population estimates for approximately 2,260 European cities at century resolution before 1500 and half-century resolution thereafter (Buringh, 2021a,b), augmented with the Bosker et al. (2013) supplement for the MENA buffer ($\sim$70 cities). Pre-1500 anchors are not native to Buringh's half-century grid. We construct a 1450 anchor by linear interpolation between the 1400 and 1500 native anchors. Roman roads, medieval trade routes, navigable rivers, and seas are taken from standard digitized shapefiles (Bagnall et al., 2024; Talbert, ed, 2000; McCormick et al., 2013).

We rasterize the cost of travel on a 10 km $\times$ 10 km grid covering Europe and the MENA buffer. Each grid cell is assigned a least-cost-of-travel value based on the dominant transport feature in that cell, with the cost-per-cell schedule following Bairoch (1988) (porters $= 1.00$, roads $= 0.81$, rivers $= 0.21$, seas $= 0.08$). Three robustness specifications use travel cost parameterizations from Boerner and Severgnini (2014), Masschaele (1997), and Galloway et al. (1996). The four cost-per-cell schedules are collected in Appendix Table 7. We use the Bairoch parameterizations as the primary specification. The cost weights are relative. A unit of τ_{ij} is the cost (in Bairoch's combined

[4] We nevertheless retain population anchors out to 1650 for the growth regressions in Section 9, where the left-hand side is Δ_{50} log(pop) over 1500–1550, 1550–1600, and 1600–1650 windows.

time/effort/risk units) of moving a unit weight one cell on flat overland portage. Moving the same unit over a 10 km cell of sea costs 0.08 as much; over a navigable river, 0.21; over a road, 0.81. In ratio terms, moving a unit weight 100 km by navigable river carries the same τ-cost as moving it roughly 26 km overland by road, or 21 km by porterage.

For each cost specification we compute the least-cost path between every pair of the $\sim$2,330 cities using Dijkstra's algorithm on the rasterized cost surface.[5] The result is a $2{,}330 \times 2{,}330$ cost matrix. Figure 4 shows two illustrative paths on the Bairoch raster: Mainz $\rightarrow$ Venice (which follows the Rhine south, then crosses the Alps via the Brenner-area pass, then descends through the Po valley to the Adriatic) and Antwerp $\rightarrow$ Madrid (which heads down the Channel to the Atlantic, around the Iberian peninsula by sea, then overland to Madrid). The paths visibly hug rivers and coasts where available.

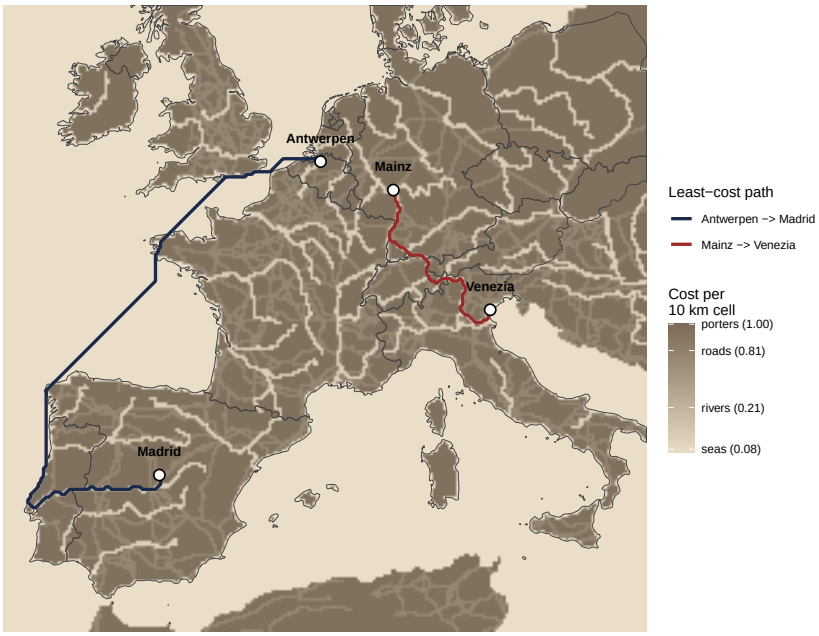


Figure 4: Two illustrative least-cost paths on the spec1_bairoch cost raster: Mainz $\rightarrow$ Venice (Rhine + Brenner Pass + Po valley), and Antwerp $\rightarrow$ Madrid (Channel + Atlantic coast + overland Iberia). The cost raster shades cells by per-cell cost; darker cells are more expensive portage, lighter cells are rivers/coasts.

[5]Implementation uses `gdistance::costDistance` in R. Full technical details—raster construction, edge weights, anchor interpolation, and the four cost specifications—are in Appendix A.

For each city i at anchor year y we aggregate to market access as

$$\mathrm{MA}_{iy} \;=\; \sum_{j \neq i} \mathrm{pop}_{jy}\, \tau_{ij}^{-\theta},$$

and the regression covariate is $\log \mathrm{MA}_{iy}$, with trade elasticity exponent $\theta = 1$ as the headline value. We exclude self-population in our headline measure. MA is conceptually access to trade outside the city, and own-city size enters the regressions separately—through the within-city demeaning of Section 7's Mundlak/BJ specification (which absorbs the city's own time-invariant scale) and through the explicit $\log(\mathrm{pop})$ convergence control in Section 9's Barro specification. Including own population in MA would conflate the trade-access quantity with own-city size, which is exactly what we want to condition on.[6]

Figure 5 shows MA at four anchor years for cities with recorded print activity (1500, 1550, 1600, and 1650). The Italian peninsula, the Low Countries, and the Rhine corridor have the highest MA throughout. By 1650 MA values are systematically higher across the entire panel as populations grow, but the spatial pattern is largely stable (the cross-sectional ranking changes little). Within-city variation in $\log(\mathrm{MA})$ over the 200-year window comes overwhelmingly from common population growth across Europe and from regional differential growth (Bruges declines while Antwerp, then Amsterdam rises). Section 6 quantifies this within-vs-between split.

[6] We construct two own-included variants (with $\tau_{ii} = 0.05$ and $\tau_{ii} = 0.10$) for diagnostic purposes only; signs on all three Section 8 metrics are preserved under both variants but magnitudes are not directly comparable to the no-own headline.

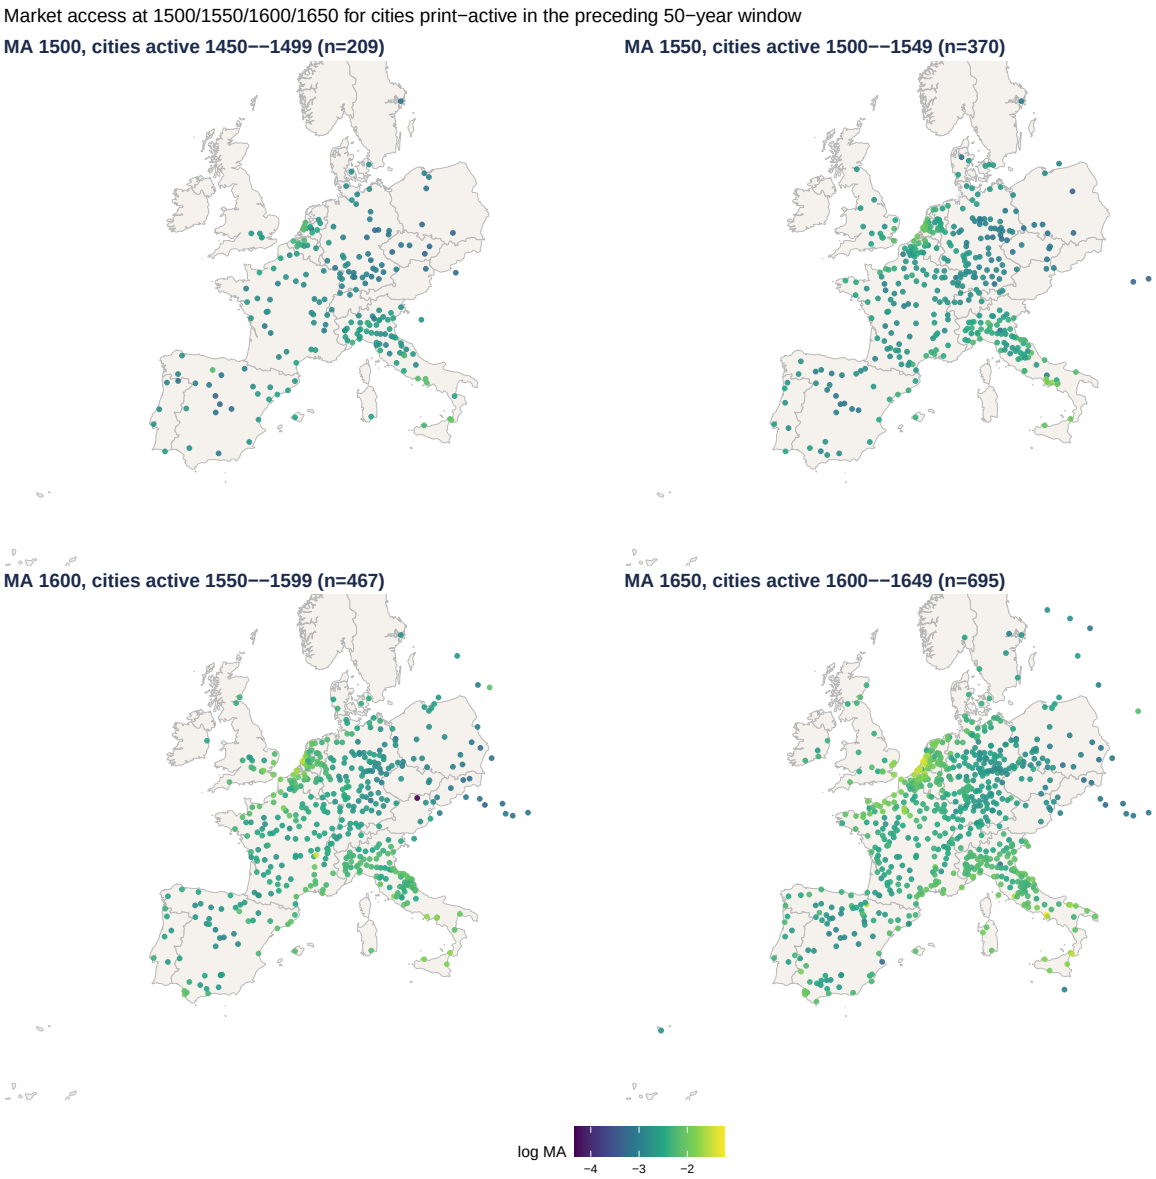


Figure 5: Market access at the four 50-year anchor years 1500, 1550, 1600, and 1650, plotted for cities with recorded print activity in the preceding 50-year window (so the 1500 panel covers cities active in 1450–1499, the 1550 panel covers 1500–1549, the 1600 panel covers 1550–1599, and the 1650 panel covers 1600–1649 — the last full window in our regression panel, given that USTC is cut at 1649). Italian peninsula, Low Countries, and Rhine corridor dominate throughout. By 1650 the entire distribution has shifted upward as European populations grew and as new printing centers entered the panel, but the cross-sectional ranking is largely preserved. Common log-MA color scale across the four panels.

6 The 50-year city-level panel

We construct our base panel with population anchors at 1500, 1550, 1600, 1650. We add 1450 by interpolation (Section 5). For each (city, language, period) cell, the panel records: the log(MA) at the period's opening anchor year; the city's diversity, distinctiveness, and specialization (entropy) metrics computed on the pooled set of titles printed in that 50-year window; and the title count n_{titles}. The three metrics (D, X, H) are formally defined in Section 2. The only choice to make at the panel-construction step is how to aggregate titles to the (city, language, period) cell. We do this at the title-set level rather than by averaging shorter-window statistics. For the (city, language, period) cell we collect every title published in that city in that language between y and $y + 49$, compute the cell-level city centroid $\bar{v}_{c,\ell,p}$ as the unweighted mean of the title embeddings in that set, and compute D, X, H directly from that pooled centroid.[7]

Restricted to the six panel-credible languages (Latin, German, Italian, French, Dutch, Spanish), the full panel has 1,878 city-period observations across 506 unique cities. Latin is the largest panel (795 obs, 385 unique cities) and Dutch is the smallest (119 obs, 66 cities). Appendix Table 10 reports per-language and pooled descriptive statistics for the four headline panel variables.[8]

The headline covariate log(MA) has means in the -2.3 to -2.8 band across languages and SDs 0.24–0.37.

Most variation in our metrics, and especially in log(MA), is between cities rather than within-city across periods (Table 2). The within-city share of log(MA) variance ranges from 4.5% (French) to 8.6% (Latin), so the identifying within-city movement is thin. Composition metrics carry more within-city variation than log(MA) does. Within-city shares for diversity range from 22% (French) to 35% (German), for distinctiveness from 19% (Spanish) to 53% (Latin), and for entropy from 25%

[7] The European centroid used for X is also a pooled title-set centroid, computed within the relevant language's embedding space. The Latin centroid is over all Latin titles in the period and the Italian centroid over all Italian titles. They are not mathematically combined into a multilingual centroid.

[8] Several decompositions reported in Appendix C use sub-samples of this 1,878-observation panel and therefore report smaller n. The useful-knowledge vs. ceremonial diversity decomposition (App. C.2) requires both useful and ceremonial subsets to clear the 5-title floor at each city-period cell, giving pooled $n = 1{,}136$. The share-of-useful test (App. C.3) requires the 10-title floor on useful-plus-ceremonial titles, giving pooled $n = 1{,}486$. The Latin-share decomposition (App. C.4) uses the broader total-diversity panel and incorporates Latin-share computed over the full USTC corpus at each (city, period), giving $n = 1{,}957$. The direction-of-distinctiveness analysis (App. C.5) reports per-subject panels whose n ranges between 737 and 1,882 depending on subject coverage by language. The headline panel of 1,878 is the maximum used in the paper. Sub-samples are always trimmed by data availability, never selected on outcome.

(Spanish) to 39% (French). The within-city variation is what identifies β in our panel regressions. It is meaningful but tight, and partly motivates our choice to use the Mundlak/Bell-Jones specification described in Section 7.

Table 2: Within-city share of total variance, by language, for log(MA) and each of the three composition metrics. Cells report the within-city share; between-city share is its complement. Derived from the six-language base panel (1,878 observations across 506 unique cities) at the (city, language, 50-year period) level.

Language	log(MA)	Diversity	Distinctiveness	Entropy
Dutch	5.8%	29.5%	32.6%	36.2%
French	4.5%	22.2%	22.1%	38.5%
German	7.2%	35.1%	32.6%	34.8%
Italian	8.3%	25.4%	38.2%	34.3%
Latin	8.6%	25.2%	52.6%	28.1%
Spanish	7.7%	26.4%	18.5%	24.6%

Figure 6 shows the per-decade weighted mean of each metric across cities by language. Each city-decade observation contributes to the per-language mean with weight equal to the number of titles printed in that city-decade. This prevents a small thinly-printing city from carrying the same weight as a large publishing center in the per-language mean and keeps the displayed series proportionate to the actual cross-corpus mass at each point in time. Diversity is roughly stable to slowly rising across the six languages, with the exception of the 1600–1649 period where German, and to a lesser extent Latin, diversity collapses, most likely because of the devastation of the Thirty Years' War. Distinctiveness mainly rises over time for all six languages over the 1450–1649 period. Cities increasingly diverge from the contemporaneous European centroid in raw level, a between-city differentiation pattern that is consistent with the deeper sample of new printing centers entering the panel after 1500. Specialization (entropy) displays no clear trend.

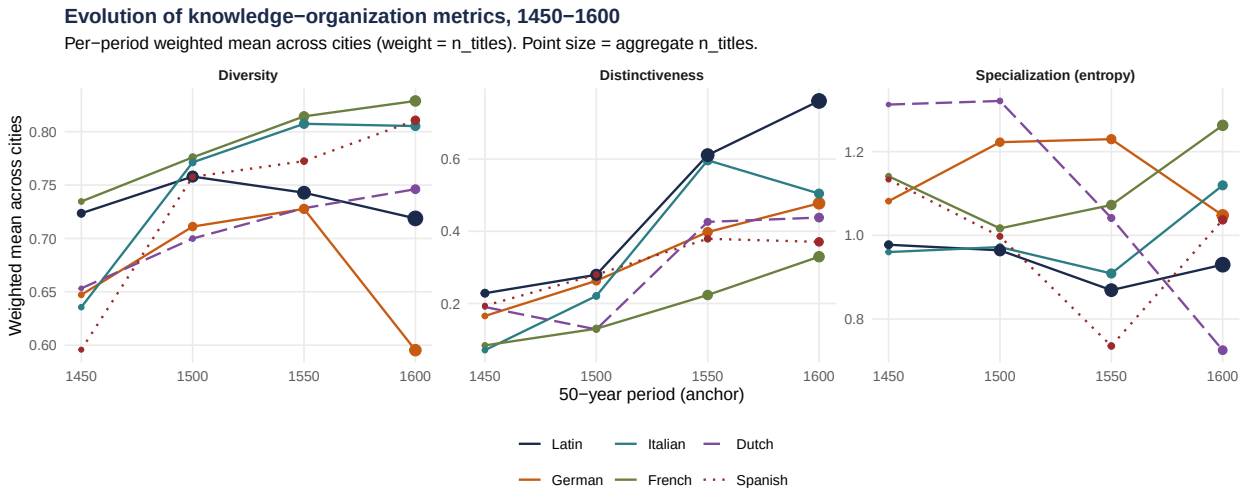


Figure 6: Per-decade weighted-mean of each metric across cities, by language. Weights are n_{titles} in each (city, decade) cell. Six panel-credible languages shown; colors follow the project colorblind-safe palette. Diversity is broadly flat with Latin slowly rising; distinctiveness rises across the panel for all six languages over 1450–1620; specialization-entropy rises for most languages except German, which declines sharply across the war-affected 1620–1649 window.

Our 1600 period spans titles published between 1600–1649, overlapping the Thirty Years' War (1618–1648), which devastated central Europe. The footprint shows up directly in the data within the war window itself. Magdeburg, the sixth-largest German-printing city in 1550–1599 (1,331 titles), disappears from the top-15 after the war begins (Tilly's 1631 sack killed roughly 25,000 of ∼30,000 inhabitants and burned the city). Leipzig under Saxon neutrality grew from 1,701 titles in 1550–1599 to 8,916 in 1600–1649 (5×), and Hamburg vaulted from outside the top-15 to second place (4,640 titles in 1600–1649) as refugee printers, capital, and trade flowed in. Within the 1600–1649 period, German diversity and entropy in the weighted means fall sharply in the post-1620 decades, a reversal of the upward trajectory seen in Latin, Italian, and French over the same decades. Latin-language print in German territories also dips notably in 1631–1638 (1,936 Latin titles in 1632, down from 3,438 in 1620), recovering by 1648.

We treat the war's contribution to the German 1600-period numbers as a sample-composition feature, not a flaw, and run sample-restriction sensitivity checks reported in Appendix D. The headline composition findings of Section 8 are attenuated by the war's German footprint. The pooled coefficients on diversity, distinctiveness, and entropy all strengthen when war-affected observations are excluded.

7 Empirical strategy

7.1 The Mundlak / Bell-Jones specification

We estimate the within-city effect of market access on knowledge composition using the Mundlak (1978) correlated random effects specification, popularized for panel applications by Bell and Jones (2015). Despite the "random effects" label that sometimes attaches to it, this is a fixed-effects estimator that identifies the within-cluster slope from within-cluster variation only, exactly as the canonical FE estimator does. Appendix C explains the equivalence in detail. On a balanced panel without singletons the two estimators deliver numerically identical within-city coefficients (Wooldridge, 2010, Ch. 10).

We use Mundlak rather than within-FE on our panel because the within-FE estimator drops singleton cities—those appearing in only one of the four 50-year periods. Singletons are 41–47% of cities in each of the six panel-credible languages (Appendix Table 17) and are systematically smaller, more intermittently printing, and more peripheral than the cities that survive. Dropping them biases the within-FE estimator toward the continuously-printing centers (Paris, Lyon, Frankfurt, Wittenberg, Venice, Rome, Antwerp) and away from the broader European urban print population that the headline claim is about. Mundlak retains the singletons in the sample at their own city mean (where they contribute exactly zero to identifying the within slope, no different from being dropped) so they can inform the between-city coefficient and the pooled regression's intercept structure without distorting the within identification.[9]

7.2 Specification

For each metric $y_{i\ell d} \in \{\text{diversity}, \text{distinctiveness}, \text{specialization-entropy}\}$:

$$y_{i\ell d} \;=\; \beta_W \cdot \widetilde{m}_{i\ell,d} \;+\; \beta_B \cdot \overline{m}_{i\ell} \;+\; \delta_d \;+\; \mu_\ell \;+\; \varepsilon_{i\ell d}$$

where:

[9] Appendix C documents the singleton mechanism, reports the within-FE and Mundlak coefficients side by side (Tables 18 and 19), shows that the headline distinctiveness sign flips between the two estimators precisely because of the singleton-driven sample-composition difference, and runs a leave-one-language-out robustness check on the Mundlak estimate (Table 20).

- $\widetilde{m}_{i\ell,d} = \log(\text{MA})_{id} - \overline{\log \text{MA}}_{i\ell}$ is the within-(city, language) deviation of $\log(\text{MA})$ at period d from the city's mean $\log(\text{MA})$ across the periods it appears in language ℓ.
- $\overline{m}_{i\ell}$ is that mean.
- δ_d is a period fixed effect, and μ_ℓ is a language fixed effect (pooled regressions only, per-language regressions drop μ_ℓ).

The coefficient β_W is the within-city longitudinal slope, measuring how much the metric moves when this city's MA moves above its own (city, language) typical level. The coefficient β_B on the city mean $\log(\text{MA})$ is the between-city cross-sectional slope, telling us whether high-MA cities have systematically higher (or lower) metric values on average.

The decomposition is computed at the (place, language) level rather than the place level. This is the appropriate granularity for our research question. We ask whether a city's $\log(\text{MA})$ trajectory in each language it prints in is associated with that language's metric trajectory at that city. Demeaning at the place level (across languages) would identify off cross-language differences within a city, which is a different (and less interpretable) contrast.

7.3 Standard errors

Standard errors are clustered at the city (place) level. Since the panel is short (four 50-year periods) and broad (around 506 unique cities across the six panel-credible languages), the asymptotic distribution of cluster-robust standard errors is close to nominal.

7.4 Weighting

Observations are weighted by $\sqrt{n_{\text{titles}}}$ to give greater influence to cells where the city-period-language metric is estimated from more titles. The square root choice is the standard inverse-variance approximation when sampling noise in the metric scales as $1/\sqrt{n_{\text{titles}}}$, applicable to the cell means and cosine-distance averages we estimate (Wooldridge, 2010, Ch. 19).[10] Robustness checks without weighting produce sign-stable results (Section 8).

[10] See also Solon et al. (2015) on weighting in panel regressions.

7.5 Pooled vs. per-language

The headline result is the pooled regression across the six panel-credible languages (Latin, German, Italian, French, Dutch, Spanish) with a language fixed effect μ_ℓ. We also report per-language regressions to expose the language-level heterogeneity that the pooled coefficient averages over. Where the pooled and per-language results disagree we discuss the heterogeneity directly rather than relying on the pooled coefficient alone.

7.6 Predictions

Section 4 and Appendix B develop the theoretical framework, monopolistic competition with spatial product differentiation, that motivates our empirical claims. The framework rests on five identifiable assumptions A1–A5 about how books were produced and consumed in 1450–1649, each contributing to a distinct empirical prediction. We test three signed within-city predictions and compare against the competing Smithian specialization prior (Table 3). Pooled estimates of the three coefficients are reported in Section 8. Appendix B reports the full conditional-prediction structure that lets us read the sign of each estimated coefficient as evidence about which of A1–A5 hold in the data. We emphasize that these are reduced-form associations rather than separately-identified structural parameters. The identifying within-city variation in $\log(\text{MA})$ is partly a proxy for proximity to the rising NW European urban complex, and we cannot rule out that part of the signed-coefficient pattern operates through channels A1–A5 do not exhaust.

Table 3: Headline predictions and the assumptions each requires.

Coefficient	Krugman-w-spatial-diff.	Smith	Required assumptions
$\beta_{\text{diversity}}$	> 0	< 0	A1, A2
$\beta_{\text{distinctiveness}}$	< 0	> 0	A1–A5
$\beta_{\text{specialization-entropy}}$	> 0	< 0	A1, A2, A3, A5

We treat the Mokyr (2002, 2009, 2017) account as the broader civilizational framing within which the production-side mechanism plays out. It is qualitatively consistent with the Krugman-plus-spatial-differentiation predictions but does not formalize signed predictions of the kind we test

in Section 8.

Sections 8 and 9 ask two sequential questions. Section 8 asks whether log(MA) shapes the composition of a city's print, by testing the three signed predictions above. Section 9 then asks whether the composition pattern Section 8 documents carries growth-relevant information beyond log(MA) itself, in the specifically Mokyr-style direction that the useful-knowledge framework predicts. The new-trade-theory baseline predicts no. Composition is in equilibrium an endogenous function of MA, and conditional on MA it should add no independent predictive power. Mokyr predicts yes, in the useful-knowledge slice of composition. Section 9 adjudicates between these two readings using a long-difference Barro growth specification with log(MA) and log(pop) controls. Section 9.1 develops the specification in detail.

8 Results

Our headline results are illustrated in Figure 7. Appendix Table 12 reports the regressions in full. The three signed predictions of Section 7.6 are all confirmed in the pooled specification and each finding survives a battery of robustness checks discussed at the end of this section.

Diversity rises with log(MA) within-city. The pooled slope is $\hat{\beta}_{\text{diversity}} = 0.467^{***}$ (SE 0.064, $n = 1{,}884$),[11] and the sign is positive in five of six per-language regressions, significant at the 5% level or better in four and at the 10% level in the fifth (Spanish): Latin 0.378^{***}, Italian 1.018^{***}, French 0.913^{***}, Dutch 0.420^{***}, Spanish 0.386^{*}. The pooled magnitude implies that a 1-log-unit increase in MA (about a 2–2.5 within-city standard-deviation movement) is associated with roughly a $0.47 \cdot \text{SD}_y$ increase in diversity, which is statistically robust and economically meaningful. The one exception is German whose within-city diversity coefficient is essentially zero (-0.069, SE 0.131). Two non-mutually-exclusive readings of the German null are plausible. First, German print from 1517 onward confessionalized sharply into Lutheran and Catholic strands. If rising market access in a German city pulled the city's print mix toward one confessional pole rather than broadening it across topics, the diversity metric would not respond. Second, the 1600–1649 German period is

[11] The 1,884 regression observations are slightly larger than the 1,878 (city, language, period) panel cells reported in Section 6. The 1,878 figure is the intersection count, cells with valid values on all three metrics simultaneously, whereas each metric's individual regression uses any cell with a valid value on that metric. A small number of cells contribute a valid D but not X or H (or vice versa), so the per-metric regression has marginally more usable observations than the intersection-of-all-metrics count. The gap is six observations for diversity.

mechanically dominated by the Thirty Years' War footprint described in Section 6. In war-affected windows German output narrows sharply and asymmetrically across cities and the within-city slope of diversity on MA is correspondingly muted. The war-restricted robustness result below provides indirect support for the second reading. The pooled diversity coefficient strengthens once war-affected observations are dropped, suggesting the German window is attenuating the headline.

Distinctiveness falls with log(MA) within-city. The pooled slope is $\hat{\beta}_{\text{distinctiveness}} = -0.709^{***}$ (SE 0.133, $n = 1{,}884$), negative in all six per-language regressions and significant in four: German -0.946^{***}, Italian -1.287^{***}, French -1.667^{***}, Dutch -1.060^{***}, Spanish -0.728 (expected sign, not significant). Latin is the remaining exception, with a slope essentially at zero (-0.058). The Latin null makes sense on its face. Latin was the period's pan-European lingua franca of learned discourse. Its readership was pan-European, and the topics Latin print addressed were almost by definition targeted at that audience rather than at any one city's local public. With Latin print's effective demand already centered on the pan-European mainstream regardless of where it was produced, increases in a Latin city's market access cannot pull the city's centroid closer to the European centroid than it already is. The high-amplitude vernacular distinctiveness slopes (French at -1.67, Italian at -1.29, Dutch at -1.06, German at -0.95) are consistent with the same logic in reverse. Vernacular markets are inherently more locally segmented so the room for MA-driven convergence toward the European centroid is correspondingly larger. The distinctiveness finding is the European print convergence result.

Specialization-entropy rises with log(MA) within-city. The pooled slope is $\hat{\beta}_{\text{spec-entropy}} = 0.751^{***}$ (SE 0.223, $n = 1{,}884$), positive in all six per-language regressions and significant in two: Italian 1.700^{***} and French 1.715^{***}. The remaining four coefficients carry the expected sign but are not significant individually: Latin 0.425, Dutch 1.345, Spanish 0.648, and German 0.093. Higher entropy means the city's titles are more evenly distributed across the ten topical clusters used in the metric, less concentrated in the descriptive sense. High-MA cities had less concentrated topical mixes. Taken with the diversity and distinctiveness results, this completes the picture. As market access rises, a city's print output becomes broader, more European-mainstream, and more evenly distributed across topics.

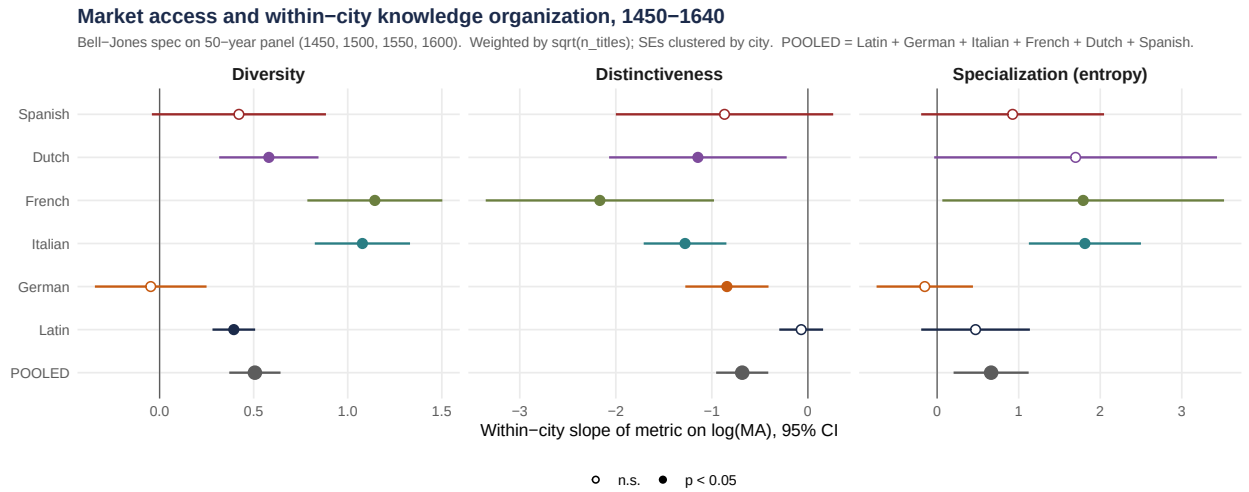


Figure 7: Pooled within-city Mundlak/BJ coefficients on log(MA) for the three composition metrics, in the six panel-credible languages and pooled. Bars show 95% confidence intervals; clustering at the city level. Full regression table in Appendix Table 12.

The three results are mutually consistent. In Latin, German, Italian, French, Dutch, and Spanish print between 1450 and 1649 cities with greater market access produced print output that was broader (more diverse), more European-mainstream (less distinctive), and less topically concentrated (more even). The broadening and the evening-out are the home-market effect that Krugman (1979, 1980) and the new economic geography literature predict, larger markets supporting more firms producing more differentiated varieties with broad demand. The convergence toward the European mainstream is the further prediction of the spatial product differentiation augmentation of Hotelling (1929), Salop (1979), and Anderson et al. (1992) developed in Section 4, under which firms in larger and more contested markets position their varieties nearer the dense center of product space. Both run against the finer-specialization prediction a naive reading of the *Wealth of Nations* would imply, though, as Section 4 notes, Stigler (1951)'s own formalization of the extent-of-the-market doctrine leaves room for increasing returns alongside finer specialization, so what we contradict is that finer-specialization prediction specifically rather than Smithian thought as a whole.

The patterns survive a series of standard checks. First, excluding the 1450 anchor period leaves the 3-period panel estimates essentially unchanged. Second, running the regressions unweighted produces sign-stable results for 24 of 30 (metric, language) pairs. Third, varying the cost specification across Bairoch, Boerner, Masschaele, and Campbell-Galloway-Murphy parametrizations

gives cross-spec correlations on $\log(\text{MA})$ of 0.81–0.97 (with the headline Bairoch spec correlating ≥ 0.91 with each of the other three) and signs that remain robust on all three pooled coefficients, all significant at conventional levels in all four parametrizations (Appendix Table 15). Fourth, varying the trade elasticity $\theta \in \{0.5, 1, 2, 3.8\}$ (the last being the structural value in Donaldson, 2018) preserves all three signed predictions at the 1% level (Appendix Table 14). Only at the extreme $\theta = 5$ do the distinctiveness and entropy signals attenuate to insignificance, the expected pattern when MA collapses toward a near-local-population proxy. Fifth, the headline coefficients are not driven by the eastern fringe of the panel whose MA estimates are most exposed to the sparser road/river shapefile coverage in eastern Europe. Dropping cities east of 25°E or 20°E moves all three pooled coefficients by less than a quarter of a standard error and leaves significance levels essentially unchanged (Appendix Table 11). Sixth, restricting each vernacular's market access to the cities that themselves print in the language, rather than weighting every neighbor by total population, preserves the pooled diversity and distinctiveness coefficients at all three inclusion thresholds, with only the specialization-entropy coefficient losing significance at the stricter thresholds (Appendix Table 33). Finally, dropping observations affected by the Thirty Years' War strengthens all three pooled coefficients rather than weakening them (Figure 8). The pattern is that the war attenuated the composition findings rather than driving them. Every coefficient gets larger in magnitude once war-affected observations are removed, and entropy nearly doubles. The full Thirty Years' War sensitivity battery is in Appendix E (Tables 26, 27, and 28).

A separate concern is that title length and cell size vary across (city, language, period) cells, and larger cells could mechanically post higher diversity by sampling a wider vocabulary. The length-restricted sub-sample of titles with at least seven words delivers essentially the same three signed coefficients as the full panel (Appendix Table 32). A more direct diagnostic that holds the per-cell title count fixed at 30 is reported in Appendix Figure 21. Across 100 permutation iterations the within-city diversity slope on $\log(\text{MA})$ is robustly positive (every iteration significant at the 5% level, 95% interval $[0.124, 0.165]$). The point estimate under this diagnostic is smaller than the headline of 0.467, consistent with part of the headline magnitude reflecting a mechanical vocabulary-coverage contribution alongside the composition response we are interested in. The sign-based prediction the framework delivers is what we test, and it survives.

The one-way city-clustered standard errors that produce the p-values above could understate

residual correlation across cities within a 50-year period. We therefore re-estimate the headline pooled specification with two-way clustering on (place, period) following Cameron and Miller (2015). Point estimates are unchanged by construction. Standard errors widen modestly for diversity (one-way SE $0.064 \rightarrow$ two-way 0.072, with p moving from 2×10^{-12} to 0.007) and for distinctiveness ($0.133 \rightarrow 0.159$, with p moving from 2×10^{-7} to 0.021). For entropy the two-way SE is actually tighter than the one-way SE ($0.223 \rightarrow 0.185$), reflecting positive within-period correlation in the residuals across cities. The entropy p-value moves from 0.001 to 0.027. All three headline coefficients remain significant under both clusterings. Diversity and distinctiveness remain significant at the 1% and 5% levels respectively under two-way clustering, and entropy at the 5% level.

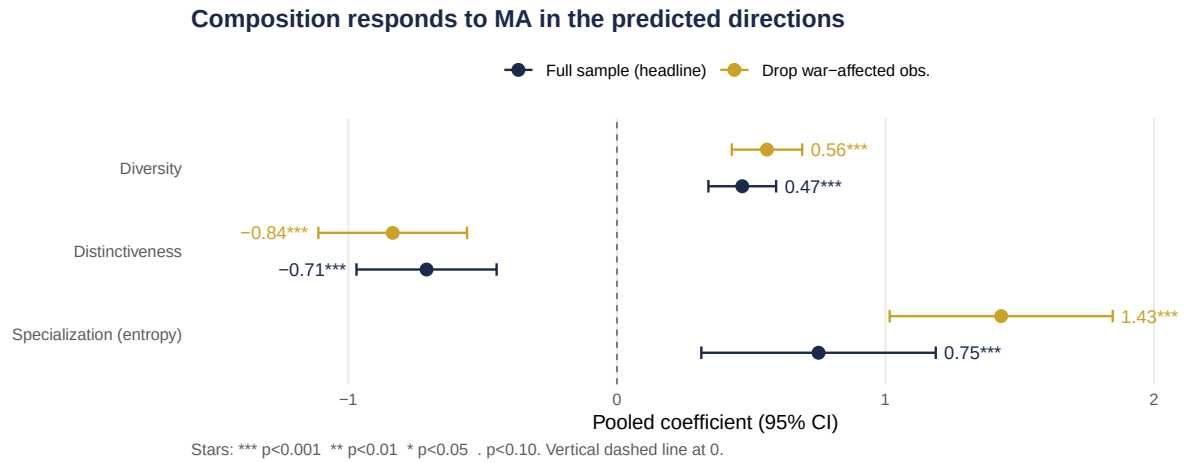


Figure 8: Pooled within-city Mundlak/BJ coefficients on $\log(\text{MA})$ for the three composition metrics, under the full sample and under the war-restricted sample (German 1600 period and Latin 1600 for German-territory cities dropped). 95% confidence intervals; dashed vertical line at zero.

Section 8 has established that market access reshapes the composition of a city's print in three signed predictions that all confirm new trade theory with spatial product differentiation and all reverse the Smithian intuition. The natural follow-up question is whether the composition pattern documented here carries growth-relevant information beyond market access itself. If new trade theory is the whole story composition is a sufficient statistic for MA's growth effect at this horizon and should add no independent predictive power conditional on MA and initial city size. If Mokyr's useful-knowledge framework is also at work, however, the slice of composition that tilts toward prescriptive instrumental content should carry an independent positive sign that survives condition-

ing on MA. Section 9 adjudicates between these two readings by asking whether the metrics that responded to MA in Section 8 predict subsequent 50-year city population growth, conditional on log(MA) and log(pop). The answer is informative on both counts. Distinctiveness and entropy are indistinguishable from zero in the growth regression (the new-trade prediction holds along those dimensions), but total diversity and the useful-knowledge share of titles carry positive, significant coefficients (the Mokyr prediction holds along the useful-knowledge slice). Section 9 develops this joint finding.

9 Knowledge composition and city growth

Section 8 established three robust within-city patterns in how print composition responds to market access. We now ask the obvious follow-up. Do those same composition metrics also predict subsequent economic outcomes, specifically city population growth? We relate our metrics to growth using the Buringh population panel at five half-century anchors — 1450, 1500, 1550, 1600, and 1650 — which gives us four 50-year growth windows: 1450–1500, 1500–1550, 1550–1600, and 1600–1650.[12]

9.1 Barro specification

Our headline specification is a Barro-style long-difference with an initial log(pop) convergence control:

$$\Delta_{50} \log(\text{pop})_{i,\, t \to t+50} \;=\; \beta\, y_{i,t} + \gamma \log(\text{MA})_{i,t} + \lambda \log(\text{pop})_{i,t} + \delta_t + \mu_\ell + \varepsilon_{i,t}.$$

Here $y_{i,t}$ is a composition metric for city i in period t. δ_t and μ_ℓ are period and language fixed effects. We do not include city FE.[13] The convergence coefficient λ is estimated rather than imposed.

[12] The 1450 population anchor is linearly interpolated by Buringh from native 1400 and 1500 values. The three later anchors (1500, 1550, 1600, 1650) are native. The 1450–1500 growth window therefore has a partly interpolated LHS. We confirmed before adopting the 4-window panel that the headline coefficients are robust to this addition. Dropping the 1450 window moves the pooled Spec A coefficients on diversity, distinctiveness, and entropy by less than a quarter of a standard error each.

[13] We considered a third design that regresses $\log(\text{pop})_{i,t}$ on a lagged composition metric with city and period fixed effects (within-city FE on the level, lagged metric on the RHS). We do not report it because the design is subject to severe Nickell (1981) bias on our $T = 4$ panel. Under the empirically observed persistence in log-population

Because we do not demean $\log(\text{pop})$ by city there is no within-cluster averaging of the lagged DV and therefore no Nickell bias on λ or β. With $T = 4$ and $N \approx 510$ cities, the Barro design fits our data structure. It permits convergence dynamics and it identifies β from cross-sectional variation in initial composition that predicts subsequent growth, controlling for initial size and MA.

Appendix D reports a complementary long-difference specification without the convergence control. The two specifications agree directionally on every coefficient we report below. The convergence-control choice shifts magnitudes modestly but not signs or significance patterns. The full regression table corresponding to Specification A is in Appendix Table 13.

9.2 Theoretical channels and the metric family

The new-trade-theory framework of Section 4 predicts that $\beta = 0$ for the three composition metrics of Section 2. In equilibrium composition is an endogenous function of MA, so conditional on MA it should carry no additional information about future growth. Several theoretical extensions, however, would deliver a positive composition–growth channel through different mechanisms (Table 4).

($\hat{\rho} \approx 1 - \hat{\lambda} \approx 0.96$), the Nickell bias on the within-city slope is of order $-(1+\rho)/(T-1) \approx -0.65$, which is large and signed against any positive composition-on-growth coefficient. The level-of-log-pop dependent variable also answers a subtly different question than the growth-rate DV we use here.

Table 4: Theoretical extensions to the new-trade-theory baseline that would predict a positive composition–growth channel, and the composition metric each one is most naturally tested against.

Extension to the framework	Predicted sign
Mokyr (2002, 2009) useful-knowledge channel: cities printing a higher share of instrumental, scholarly, scientific content drive growth via Republic-of-Letters knowledge transfer.	$\hat{\beta}_{\text{share-useful}} > 0$ on the extensive margin; $\hat{\beta}_{\text{within-useful diversity}} > 0$ on the intensive margin
Endogenous growth via knowledge spillovers (Romer, 1986; Grossman and Helpman, 1991). Jacobs externalities (Glaeser et al., 1992): diverse knowledge stocks generate ideas, and ideas drive growth.	$\hat{\beta}_{\text{diversity}} > 0$
Marshall externalities (Henderson, 2003): own-industry specialization generates within-industry productivity gains.	$\hat{\beta}_{\text{distinctiveness}} > 0$, $\hat{\beta}_{\text{diversity}} < 0$

These three extensions are not mutually exclusive. Mokyr's channel is about which subset of knowledge production matters (the prescriptive subset), Romer-Jacobs is about whether any breadth matters, and Marshall is about whether specialization matters. They can be tested separately because each maps to a distinct composition metric.

The three composition metrics of Section 2 — diversity, distinctiveness, and specialization-entropy — carry the Romer-Jacobs and Marshall tests directly. Testing Mokyr's useful-knowledge channel requires two additional metrics, both built on a partition of USTC titles into a useful (prescriptive, instrumental, scholarly, scientific) subset and a ceremonial (devotional, poetic, dramatic) subset. We adopt:[14]

- *Useful*: Classical Authors, Medical Texts, Jurisprudence, Philosophy and Morality, Academic Dissertations, Periodicals, Educational Books, History and Chronicles, Political Tracts.

[14] Each useful category captures a face of the propositional / prescriptive knowledge Mokyr (2002) identifies as the substantive carrier of the Industrial Enlightenment. Medical texts codify applied medicine, jurisprudence codifies institutional and commercial knowledge, philosophy and morality covers natural philosophy and ethics, academic dissertations are formal scholarly output, periodicals carry commercial and political information, educational books transmit instrumental literacy and numeracy, history and chronicles consolidate institutional and political knowledge, political tracts develop political-economy argument, and classical authors carry the recovered Greco-Roman corpus that fed humanist scholarship. Classical Authors is the most defensibly contested inclusion in useful, since it could be read as a humanist or rhetorical category rather than a propositional one, and we report the robustness of the share-of-useful coefficient to dropping it in Appendix D.3.

- *Ceremonial*: Religious, Funeral orations, Wedding pamphlets, Poetry, Drama.

On this partition we construct two metrics. The *share of useful content* in a (city, period, language) cell is the extensive-margin metric:

$$s^{\text{useful}}_{i,t,\ell} = \frac{n^{\text{useful}}_{i,t,\ell}}{n^{\text{useful}}_{i,t,\ell} + n^{\text{ceremonial}}_{i,t,\ell}},$$

the fraction of useful titles among useful-plus-ceremonial titles. *Within-useful diversity* is the intensive-margin metric: the diversity measure defined in Section 2, computed on the useful subset of titles only. The extensive metric asks what share of a city's print is useful-knowledge. The intensive metric asks how topically deep a city's useful-knowledge print is. Both enter the Barro specification of Section 9.1 identically. We restrict the sample to (city, period, language) cells with at least 10 useful-plus-ceremonial titles to keep the share well-measured.

We also test four further composition channels in Appendix D — USTC classification shares, ceremonial diversity, Latin share as a transnational-orientation proxy, and direction-of-distinctiveness in embedding space. None of these delivers a substantively positive growth-side coefficient at the pooled level. They corroborate the central finding here that the useful-knowledge channel is the channel that survives and they rule out several plausible alternative readings.

9.3 Headline city growth results

Table 5 reports the pooled Barro coefficient on each composition metric. Three patterns are immediate.

Table 5: Pooled Barro growth coefficients on each composition metric. Each row is a separate regression; standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$ (or $\sqrt{n_{\text{useful+ceremonial}}}$ for share-of-useful). Period and language fixed effects included; no city FE. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Composition metric	$\hat{\beta}_{\text{metric}}$	$\hat{\gamma}_{\text{MA}}$	$\hat{\lambda}_{\text{pop}}$	n
Diversity	0.143^{*}	0.073^{**}	-0.009	1,882
	(0.082)	(0.029)	(0.015)	
Distinctiveness	-0.058	0.072^{**}	-0.006	1,882
	(0.065)	(0.029)	(0.014)	
Specialization (entropy)	-0.036	0.070^{**}	-0.003	1,882
	(0.031)	(0.029)	(0.015)	
Share of useful content	0.123^{**}	0.068^{**}	-0.004	1,486
	(0.049)	(0.029)	(0.015)	

Distinctiveness and entropy do not predict growth. Conditional on $\log(\text{MA})$ and $\log(\text{pop})$, both coefficients are statistically indistinguishable from zero (distinctiveness -0.058, $p = 0.37$; entropy -0.036, $p = 0.25$). The Marshall externality prediction that specialized, distinctive cities grow faster finds no support. This is exactly the new-trade-theory prediction holding along these two dimensions. Composition responds to MA in Section 8 but carries no independent growth-relevant information beyond MA conditional on initial size.

$\log(\text{MA})$ itself predicts growth robustly. The pooled MA coefficient is $\hat{\gamma}_{\text{MA}} = 0.073^{**}$ (SE 0.029, $p \approx 0.013$) and it is stable across all four headline regressions. Cities with higher market access grew faster over the next 50 years conditional on initial size, consistent with the cumulative-causation prediction of Krugman (1991) at the 50-year horizon.[15]

Two compositional signals do predict growth above and beyond MA. Total diversity carries a marginally positive coefficient ($\hat{\beta} = 0.143^{*}$, $p = 0.084$). Cities with topically broader print output

[15] The pooled convergence coefficient $\hat{\lambda} = -0.009$ (SE 0.015) is not statistically distinguishable from zero. We do not detect mean reversion of city size at this horizon, conditional on metrics, MA, and period FE. Half-century convergence coefficients in the early modern European urban literature (de Vries, 1984; Bairoch, 1988) are typically small in absolute value, so the modest size of $\hat{\lambda}$ is consistent with that literature even though the coefficient is not individually significant. Per-language convergence coefficients (Table 13) carry the conventional negative sign in German (significant at the 1% level) and are positive in Italian (marginally significant) and French (not significant). The pooled null averages over this heterogeneity.

grew faster over the next 50 years, even after conditioning on log(MA) and log(pop). And the share of useful-knowledge titles carries a robustly positive coefficient ($\hat{\beta} = 0.123^{**}$, $p = 0.012$, $n = 1{,}486$). Both signals point in the same direction. The slice of composition that tilts toward broader, useful-knowledge, content predicts subsequent growth. The total-diversity signal is the Romer-Jacobs prediction (any breadth of knowledge drives growth) at marginal significance. The share-of-useful signal is the Mokyr prediction (the useful-knowledge slice specifically drives growth) at conventional significance.

Per-language detail is in Appendix Table 13. The full robustness battery for the share-of-useful coefficient (eight sample restrictions and specification choices) is in Appendix Table 23.

9.4 Two margins, two languages: the Latin asymmetry

The pooled headlines in Section 9.3 mask a sharper per-language pattern that, once it is on the table, organizes the substantive interpretation of every coefficient in Table 5. The *within-useful diversity* coefficient is large and significantly positive for Latin but null or negative for every vernacular language (Figure 9, right panel; per-language detail in Appendix Table 21), while the *share-of-useful* coefficient is the opposite, large and significantly positive for German and Italian but small and insignificant for Latin (Table 6 below).

Table 6: Share-of-useful and subsequent city growth (Barro long-difference).

	Pooled (1)	Latin (2)	German (3)	Italian (4)	French (5)	Dutch (6)	Spanish (7)
Share-of-useful$_{i,t}$	0.123**	0.084	0.247***	0.218**	0.167	0.066	0.270
	(0.049)	(0.076)	(0.081)	(0.105)	(0.134)	(0.235)	(0.207)
log(MA)$_{i,t}$	0.068**	0.084**	-0.027	0.059**	0.049	0.014	0.329***
	(0.029)	(0.034)	(0.033)	(0.025)	(0.051)	(0.090)	(0.106)
log(pop)$_{i,t}$	-0.004	-0.005	-0.071***	0.012	0.022	0.040	0.028
	(0.015)	(0.016)	(0.015)	(0.012)	(0.030)	(0.050)	(0.031)
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Language FE	Yes	No	No	No	No	No	No
Observations	1,486	666	286	148	185	93	108
Cities (clusters)	421	327	148	73	105	52	60

Notes: Dependent variable is $\Delta_{50}\log(\text{pop})_{i,t\to t+50}$. Share-of-useful is the fraction of useful titles in useful-plus-ceremonial titles in the cell. Sample restricted to (city, period, language) cells with ≥ 10 useful-plus-ceremonial titles. Standard errors clustered at the city in parentheses; the 'Cities (clusters)' row reports the number of unique cities contributing to each column and bounds the effective within-city degrees of freedom available for inference. Observations weighted by $\sqrt{n_{\text{useful+ceremonial}}}$. Period FE always included; language FE in the pooled column only. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

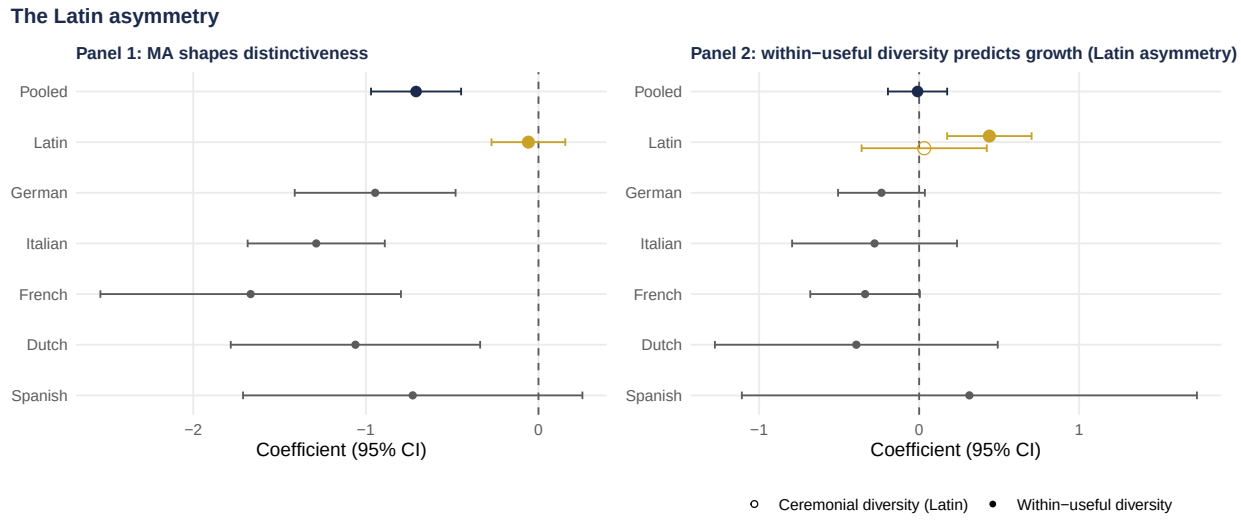


Figure 9: The Latin asymmetry. Panel 1 reports pooled and per-language within-city BJ coefficients of distinctiveness on log(MA) from Section 8's headline regression. Latin's slope is statistically null ($\hat{\beta} = -0.058$, $p = 0.60$). Every other language shows a sharply negative slope. Panel 2 reports pooled and per-language Barro growth coefficients on within-useful diversity (filled circles) from Appendix D.2's decomposition. Latin's coefficient (0.439, $p = 0.001$, gold filled circle) is the largest in the panel. The pooled coefficient and the other languages' coefficients are null or negative. Latin's ceremonial-diversity coefficient (open gold circle) sits at zero, confirming that the Latin growth signal comes through the useful subset rather than through any diversity within Latin print. Bars are 95% confidence intervals. Latin is highlighted in gold and the pooled coefficient in navy.

The two metrics are not capturing the same thing. They capture different margins of the same Mokyr useful-knowledge channel, and the Latin/vernacular split between them lines up cleanly with a core difference in the audiences that each medium served.

For a vernacular-printing city, the audience that read the city's output was overwhelmingly local, made up of vernacular literates within the city walls and the surrounding region. The compositional balance between useful content (legal handbooks, almanacs, commercial information, news pamphlets, basic medical guides, political tracts) and ceremonial content (devotional literature, hymnals, vernacular religious pamphlets, popular drama, broadside poetry) varies substantially across vernacular-printing cities and across periods, reflecting local religious-confessional and commercial-civic cultures. Reformation-era Wittenberg, for instance, printed overwhelmingly Lutheran-confessional material in German. Only 0.05 of its $1{,}658$ German titles in 1500–1549 fall in the useful-knowledge partition, a share that stays at 0.04 ($n = 1{,}486$) a century later. Augsburg and Hamburg in their seventeenth-century commercial peaks, by contrast, tilted further toward useful-knowledge content. Cities whose vernacular print supply skewed toward useful-knowledge content provided their local population with the kind of practical, instrumental literacy that translates into productive activity, and the extensive share-of-useful margin picks up that cross-city variation. Within the useful subset, vernacular content was comparatively homogeneous across cities (legal handbooks, almanacs, and news pamphlets look broadly similar across German, Italian, or French printing centers), so the intensive within-useful margin does not predict growth for vernacular print.

Latin print served a fundamentally different audience. Its readers were a trans-European scholarly community who could obtain Latin texts from any printing city in Europe. The compositional share of useful versus ceremonial in any single Latin-printing city's output mattered less than what the city contributed to the broader trans-European useful-knowledge stock. Cities whose Latin useful-knowledge print specialized deeply, like Padua in medicine and natural philosophy (useful share 0.83 in 1450–1499, 0.54 on 410 Latin titles in 1550–1599), Salamanca in canon law, Leiden in mathematics and natural science, Pavia in jurisprudence, or Bologna in law and medicine, built productive reputations that attracted skilled scholarly labor, sustained the university and professional infrastructure that supported their economies, and contributed to subsequent population growth. These same cities sit at the dense center of the European scholar-affiliation network that

de la Croix et al. (2025) reconstruct from the *Repertorium Eruditorum Totius Europae*, which their simulations show channelled ideas (botanic gardens, mathematical astronomy, Protestantism) preferentially toward high-exposure cities. Our Latin within-useful diversity coefficient plausibly captures the printed-output footprint of the same university-academy network. The natural concern is then that the Latin coefficient is identified off these few exemplar cities. Appendix Table 22 reports a leave-one-university-city-out robustness check. Dropping Bologna, Padua, Leiden, or Pavia one at a time leaves $\hat{\beta}_{\text{within-useful}}$ in the range 0.436–0.442, unchanged from the baseline 0.439 and significant at the 1% level in every case. The Latin within-useful diversity signal is not identified off any single university city; it is a feature of the broader Latin scholarly panel. The Latin margin also picks up dramatic within-city recompositions: Wittenberg's Latin useful share rises from 0.28 ($n = 780$) in 1500–1549, when Latin print was still dominated by Reformation theology, to 0.85 on $7{,}500$ Latin titles in 1600–1649 as the university shifted toward academic dissertations and natural-philosophy disputations. The intensive within-useful margin captures that depth-of-specialization variation across Latin-printing cities. The extensive share margin does not predict Latin growth because the relevant audience was trans-European rather than local. How much of any one city's Latin print was useful mattered less than whether that city was a distinctive, productive contributor to the European stock of scholarly useful knowledge.

The data confirm the asymmetry sharply. Table 6 shows the share-of-useful coefficient is positive and significant for German (0.247^{***}, $p = 0.003$) and Italian (0.218^{**}, $p = 0.040$) but small and insignificant for Latin (0.084, $p = 0.27$). Appendix Table 21 shows the within-useful diversity coefficient is large and significant for Latin (0.439^{***}, $p = 0.001$) but null or negative for German (-0.236), Italian (-0.279), and French (-0.338, marginally negative). The two metrics are essentially mutually exclusive in which languages they pick up. Same Mokyr channel, two margins, two language signatures.

This split also makes sense of Latin's role across both sections of the paper. In Section 8, Latin is the only language whose within-city distinctiveness coefficient on $\log(\text{MA})$ is statistically null, because Latin print's effective demand is already centered on the pan-European mainstream and cannot be pulled any closer to it by rising MA. In Section 9, Latin's within-useful diversity coefficient is the largest and most significant on the growth side, because Latin is the trans-European scholarly medium where depth of useful-knowledge content actually circulates. The same physical

fact about Latin print produces both findings. Figure 9 above plots the two coefficients that directly express this asymmetry.

9.5 Interpretation and caveats

The three predictions divide the four metrics cleanly. Distinctiveness and entropy carry no independent growth signal once MA is controlled for. Both move with MA in Section 8 and stop mattering on the growth side, which is what new trade theory predicts. Total diversity carries a marginally positive growth signal, consistent with the Romer-Jacobs view that breadth of knowledge drives growth. The two Mokyr-relevant metrics behave the way the Mokyr channel predicts. Share-of-useful predicts growth on the extensive margin, driven by vernacular composition shifts. Within-useful diversity predicts growth on the intensive margin, driven by Latin scholarly depth. The cross-language asymmetry is not noise. It is what we would expect if the channel works through different print media for different audiences.

Our findings refine the empirical print-and-growth literature on both margins. Dittmar (2011) and Boerner et al. (2021) have shown that whether a city had a printing press predicts subsequent population growth between 1500 and 1600. Our panel adds the compositional margin. What the press produced, particularly the share and depth of useful-knowledge content, predicts growth above and beyond whether a press existed. Our share-of-useful finding is the city-level, six-language, 1450–1650 analog of the volume-level, England-published, 1500–1900 corpus-textual finding of Almelhem et al. (2026). In their LDA-based analysis of 264,443 volumes published in England in the HathiTrust Digital Library, the volumes that became most progress-oriented during the Enlightenment were those using language at the nexus of science and political economy, with industrial-flavoured works at that same nexus carrying the strongest signal. The Mokyr channel finds further empirical support in Squicciarini and Voigtländer (2015), who use subscriber lists for Diderot's *Encyclopédie* as an upper-tail human-capital proxy and show that they predict French city industrialization a century later. Their proxy operates on the demand side, ours on the supply side. The timing of our growth signal also aligns with the productivity-growth chronology of Bouscasse et al. (2025), who estimate that English productivity growth accelerated around 1600, bracketed by our 1450 and 1650 panel endpoints.

The magnitudes are small. A one-standard-deviation increase in share-of-useful is associated with about a 1.5 log-point increase in 50-year city growth.

The growth channel may also operate at a longer horizon than the 1450–1650 window suggests. Our reduced-form panel cannot distinguish equilibrium-conditioned composition coefficients from causal ones, and the Krugman cumulative-causation dynamics in Krugman (1991) explicitly play out over centuries rather than single 50-year windows. A positive 50-year correlation could therefore overstate or understate the structural channel.

The Thirty Years' War (1618–1648) shapes the 1600–1650 growth window for German-territory cities. Appendix E reports the sensitivity, and both Section 8 composition signals and Section 9 diversity-on-growth strengthen once war-affected observations are removed.

We do not claim a single channel explains every coefficient. Italian print shows a significantly negative growth coefficient on diversity (-0.251^{**}, Table 13) that runs against the pooled positive, plausibly reflecting the outsized weight of Venice in the Italian-print panel (35% of all Italian titles, see Section 3) over a period in which Venice's own economic position weakened.

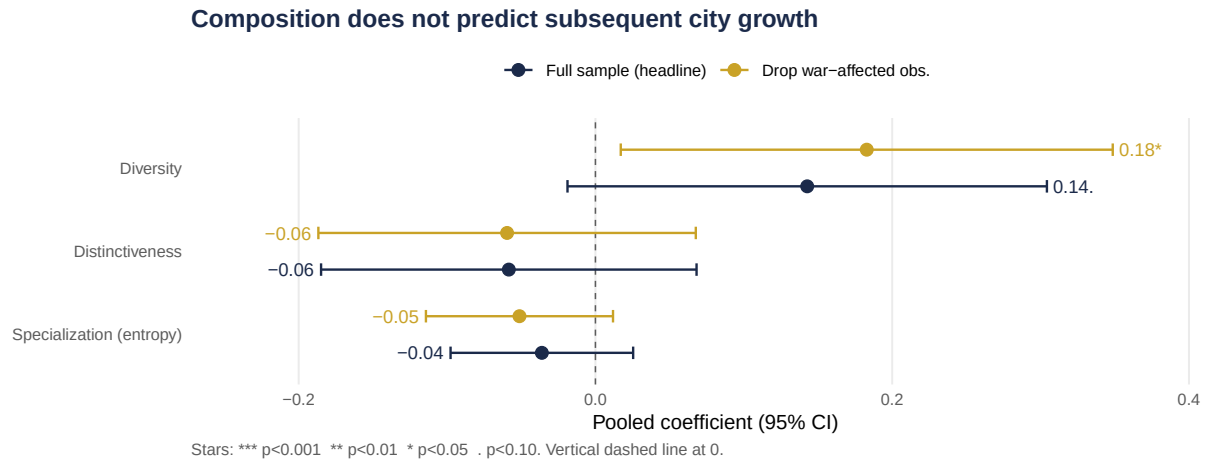


Figure 10: Pooled Barro long-difference coefficients of $\Delta_{50}\log(\text{pop})$ on each composition metric, conditional on $\log(\text{MA})$ and $\log(\text{pop})$, under the full sample and the war-restricted sample. 95% confidence intervals; dashed vertical line at zero. Full regression table in Appendix Table 13.

10 Conclusions

We document a joint MA → composition → growth relationship across six European languages between 1450 and 1650, with a specifically Mokyr-style middle node. The first link (Section 8) shows

that within-city increases in market access are associated with print output that is more topically diverse, less distinct from the contemporaneous European centroid, and more even across topical clusters. These three signed predictions all reverse the Smithian specialization prior and confirm the monopolistic-competition trade theory with spatial product differentiation laid out in Section 4 and developed in Appendix B. The second link (Section 9) shows that the useful-knowledge slice of that composition pattern, the fraction of titles in prescriptive, instrumental, scholarly subjects (0.123^{**}, $p = 0.012$ pooled) and within-useful diversity in the trans-European Latin medium (0.439^{***} in Latin), predicts subsequent 50-year city growth conditional on $\log(\text{MA})$ and initial $\log(\text{pop})$. Distinctiveness and entropy do not, as new trade theory alone would predict. The joint reading is that market access reshapes the composition of print at the city level (the new-trade prediction), and the useful-knowledge slice of that composition carries independent growth-relevant content above and beyond MA (the Mokyr prediction), with the trans-European Latin scholarly medium carrying the lion's share of the within-language signal.

The cleanest piece of evidence that the useful-knowledge channel is the active mechanism is the Latin asymmetry that runs across both sections. Latin is the only language with a statistically null distinctiveness slope in Section 8. Its print output is already centered on a pan-European audience, so increases in market access cannot pull its centroid any closer to a centroid it already sits on. Latin is simultaneously the language whose within-useful diversity coefficient is the largest and most significant in Section 9's growth regression. The same physical fact about Latin print as a trans-city scholarly medium produces both findings. It is invisible to a metric that measures distance from the European mainstream (Section 8), because Latin is the European mainstream of useful knowledge, and it carries the strongest growth signal of any language when restricted to the useful subset (Section 9), because that is the medium through which Mokyr's transnational stock of useful knowledge actually circulated.

Our two findings speak directly to the three literatures we introduced in Section 1. The composition pattern is the opposite of Smith's division-of-labor intuition in this setting. Rather than larger markets producing narrower local specialization, they produced broader, more mainstream-centered print output. That is not a refutation of Smith (1776) so much as a confirmation of the modern monopolistic-competition refinement of him (Dixit and Stiglitz, 1977; Krugman, 1980; Helpman and Krugman, 1985) once spatial product differentiation is added to the framework. Pro-

duction tilts toward the demand-weighted centroid in high-MA locations rather than into distinctive niches. The same composition pattern is qualitatively consistent with the broader European-integration dynamic that Mokyr (2002, 2009, 2017) emphasizes. What our data add is direct city-level measurement of the mainstream-convergence half of that integration, traceable in the embedding geometry rather than inferred from selected case studies. The empirical pattern also runs in a different direction from Eisenstein (1980)'s "broaden then deepen" framing. In our panel, as a city's MA rises, its print output becomes broader and more mainstream, not narrower and more specialized, and the seventeenth-century "deepening" she describes does not appear as a within-city tightening of topical concentration. Finally, our growth results refine the empirical print-and-growth literature initiated by Dittmar (2011) and continued by Dittmar and Seabold (2019); Boerner et al. (2021); Taylor and Hall (2026). The extensive margin (whether a city had a press) predicts growth in their work. We show that the compositional margin within printing cities adds incremental predictive power for a Mokyr-style useful-knowledge channel. Cities whose print output skewed toward useful-knowledge classifications grew systematically faster than comparably sized and comparably connected cities whose print output skewed toward ceremonial classifications.

The Latin within-useful diversity coefficient confirms the same channel through a different operationalization. The magnitude is modest (roughly twelve log-points of additional growth over the 50-year window when the useful share moves from zero to one), but the sign, robustness, and the Latin concentration of the signal are all consistent with Mokyr (2002, 2009, 2017)'s "Industrial Enlightenment" reading. Our findings are the cross-language, city-level, 1450–1650 complement to two recent volume-level textual analyses of English print. Grajzl and Murrell (2024) apply structural topic modeling to English print 1530–1700 and document the coevolution of ideas on religion, science, and institutions in the run-up to industrialization. Almelhem et al. (2026) apply latent Dirichlet allocation to volumes published in England 1500–1900 and document the mid-eighteenth-century secularization of scientific language and the progress-oriented turn at the science / political-economy nexus. Their analyses are national, English-focused, and centered on the Reformation and Enlightenment proper. Ours is European, six-language, and centered on the print revolution that immediately preceded them. All three papers identify the same Mokyr useful-knowledge channel through different corpora and different methods, and together they cover sequential phases of the same long-run cultural and economic transformation. Our period documents

the formation of the trans-European Latin scholarly medium and the early differentiation of useful-knowledge print across cities. Their period documents the seventeenth- and eighteenth-century turn in the language of that knowledge toward explicit belief in progress, in the language of England's industrialisers. The timing of our growth signal also aligns with the recent productivity-growth chronology of Bouscasse et al. (2025), who estimate that English productivity growth accelerated around 1600, a date bracketed by our 1450 and 1650 panel endpoints.

The broader speculative reading we draw from these findings is about the geography of intellectual life during a market-integration episode. One might have expected, as have many literatures across history, sociology, and cultural studies, that as European book markets thickened, local intellectual cultures would have developed more distinctively, with each city carving out its own niche. The opposite happened according to our data. As cities became more central in the European trade network their print output drifted toward a common European intellectual mainstream. The shared canon of humanist-scientific Latin we associate with the late Renaissance and the early Republic of Letters appears in our panel not as a feature of any one center but as a convergent equilibrium of the trade-network high-MA cities.

That convergence is, we suspect, part of why Mokyr (2017)'s integrated European knowledge stock became possible in the first place. Cumulative European intellectual integration required, as one necessary condition, that the highest-market-access producers tilt toward a shared mainstream rather than entrenching local idiosyncrasies. Our measurement captures the slope of that tilt directly. Whether the same logic governs intellectual convergence in modern market-integration episodes (in academic publishing, in digital media, in patent specializations across modern city networks) is a question we cannot answer with sixteenth-century data, but the mechanism is general enough to invite empirical work beyond the early modern period.

Data availability

The USTC corpus is distributed by the University of St Andrews under the license terms documented at the USTC project page (Universal Short Title Catalogue, 2022). The local snapshot used in this paper is dated 2022-12-13. The Buringh European Urban Population dataset is available from DANS (Buringh, 2021a) and described in Buringh (2021b). The Bosker et al. (2013) MENA

supplement is distributed with that paper. The R scripts that build the cost matrices, title embeddings, market-access panel, composition metrics, and regression tables, together with the figures and intermediate data files, will be released as a public replication package on publication and are available from the authors on request in the interim.

References

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, 2014, *129* (3), 1085–1140.

Almelhem, Ali, Murat Iyigun, Austin Kennedy, and Jared Rubin, "Enlightenment Ideals and Belief in Progress in the Run-up to the Industrial Revolution: A Textual Analysis," *Quarterly Journal of Economics*, 2026, *141* (1), 263–314.

Anderson, Simon P., André de Palma, and Jacques-François Thisse, *Discrete Choice Theory of Product Differentiation*, Cambridge, MA: MIT Press, 1992.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Arora, Sanjeev, Yingyu Liang, and Tengyu Ma, "A Simple but Tough-to-Beat Baseline for Sentence Embeddings," in "International Conference on Learning Representations (ICLR)" 2017.

Bagnall, Roger S., Richard J. A. Talbert, Sarah Bond, Tom Elliott, Sean Gillies et al., "Pleiades: A Gazetteer of Past Places," https://pleiades.stoa.org/ 2024. Open-access gazetteer of ancient and medieval places.

Bairoch, Paul, *Cities and Economic Development: From the Dawn of History to the Present*, Chicago, IL: University of Chicago Press, 1988.

Bell, Andrew and Kelvyn Jones, "Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data," *Political Science Research and Methods*, 2015, *3* (1), 133–153.

Boerner, Lars and Battista Severgnini, "Epidemic Trade," Technical Report, London School of Economics and Political Science, LSE Library 2014.

_ **, Jared Rubin, and Battista Severgnini**, "A Time to Print, a Time to Reform," *European Economic Review*, 2021, *138*, 103826.

Bosker, Maarten, Eltjo Buringh, and Jan Luiten van Zanden, "From Baghdad to London: Unraveling Urban Development in Europe, the Middle East, and North Africa, 800–1800," *Review of Economics and Statistics*, 2013, *95* (4), 1418–1437.

Bouscasse, Paul, Emi Nakamura, and Jón Steinsson, "When Did Growth Begin? New Estimates of Productivity Growth in England from 1250 to 1870," *Quarterly Journal of Economics*, 2025, *140* (2), 835–888.

Buringh, Eltjo, "European Urban Population, 700–2000," Dataset, version 1.0. 2021.

_ , "The Population of European Cities from 700 to 2000," *Research Data Journal for the Humanities and Social Sciences*, 2021, *6* (1), 1–18.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Church, Kenneth W. and Patrick Hanks, "Word Association Norms, Mutual Information, and Lexicography," *Computational Linguistics*, 1990, *16* (1), 22–29.

de la Croix, David, Matthias Doepke, and Joel Mokyr, "Clans, Guilds, and Markets: Apprenticeship Institutions and Growth in the Preindustrial Economy," *Quarterly Journal of Economics*, 2018, *133* (1), 1–70.

_ , Rossana Scebba, and Chiara Zanardello, "*Flora, Cosmos, Salvatio*: Pre-modern Academic Institutions and the Spread of Ideas," August 2025. Working paper, IRES/LIDAM, UCLouvain.

de Vries, Jan, *European Urbanization, 1500–1800*, Cambridge, MA: Harvard University Press, 1984.

Dittmar, Jeremiah E., "Information Technology and Economic Change: The Impact of the Printing Press," *Quarterly Journal of Economics*, 2011, *126* (3), 1133–1172.

_ and Skipper Seabold, "New Media and Competition: Printing and Europe's Transformation after Gutenberg," 2019. Conditionally accepted at the Journal of Political Economy. Working-paper version dated 2019; PDF copy at master_supporting_docs/Dittmar Seabald 2019 New Media.pdf.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 1977, *67* (3), 297–308.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 2018, *108* (4–5), 899–934.

_ and Richard Hornbeck, "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 2016, *131* (2), 799–858.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Eisenstein, Elizabeth L., *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*, Cambridge: Cambridge University Press, 1980. Two-volume work; abridged single-volume edition issued 1983.

Galloway, James A., Derek Keene, and Margaret Murphy, "Fuelling the City: Production and Distribution of Firewood and Fuel in London's Region, 1290–1400," *The Economic History Review*, 1996, *49* (3), 447–472.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, "Text as Data," *Journal of Economic Literature*, 2019, *57* (3), 535–574.

Glaeser, Edward L., Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer, "Growth in Cities," *Journal of Political Economy*, 1992, *100* (6), 1126–1152.

Grajzl, Peter and Peter Murrell, "A Macroscope of English Print Culture, 1530–1700, Applied to the Coevolution of Ideas on Religion, Science, and Institutions," *Social Science History*, 2024, *48* (3), 489–519.

Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.

Hamilton, William L., Jure Leskovec, and Dan Jurafsky, "Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change," in "Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)" 2016, pp. 1489–1501.

Helpman, Elhanan and Paul R. Krugman, *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*, Cambridge, MA: MIT Press, 1985.

Henderson, J. Vernon, "Marshall's Scale Economies," *Journal of Urban Economics*, 2003, *53* (1), 1–28.

Hoberg, Gerard and Gordon Phillips, "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis," *Review of Financial Studies*, 2010, *23* (10), 3773–3811.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39* (153), 41–57.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation Over the Long Run," *American Economic Review: Insights*, 2021, *3* (3), 303–320.

Kelly, Morgan, Joel Mokyr, and Cormac Ó Gráda, "The Mechanics of the Industrial Revolution," *Journal of Political Economy*, 2023, *131* (1), 59–94.

Krugman, Paul, "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, 1979, *9* (4), 469–479.

_ , "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 1980, *70* (5), 950–959.

_ , "Increasing Returns and Economic Geography," *Journal of Political Economy*, 1991, *99* (3), 483–499.

Lancaster, Kelvin, *Variety, Equity, and Efficiency: Product Variety in an Industrial Society*, New York: Columbia University Press, 1979.

Levy, Omer and Yoav Goldberg, "Neural Word Embedding as Implicit Matrix Factorization," in "Advances in Neural Information Processing Systems 27 (NeurIPS 2014)" 2014, pp. 2177–2185.

Masschaele, James, *Peasants, Merchants, and Markets: Inland Trade in Medieval England, 1150–1350*, New York: St. Martin's Press, 1997.

McCormick, Michael, Guoping Huang, Giovanni Zambotti, and Jessica Lavash, "Roman Road Network (Version 2008)," *Digital Atlas of Roman and Medieval Civilizations (DARMC)*, 2013. Harvard University. https://darmc.harvard.edu/.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean, "Distributed Representations of Words and Phrases and their Compositionality," in "Advances in Neural Information Processing Systems" 2013, pp. 3111–3119.

Mokyr, Joel, *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton, NJ: Princeton University Press, 2002.

_ , "The Intellectual Origins of Modern Economic Growth," *Journal of Economic History*, 2005, *65* (2), 285–351.

_ , *The Enlightened Economy: An Economic History of Britain, 1700–1850*, New Haven, CT: Yale University Press, 2009.

_ , *A Culture of Growth: The Origins of the Modern Economy*, Princeton, NJ: Princeton University Press, 2017.

Mundlak, Yair, "On the Pooling of Time Series and Cross Section Data," *Econometrica*, 1978, *46* (1), 69–85.

Nickell, Stephen, "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 1981, *49* (6), 1417–1426.

Pascali, Luigi, "The Wind of Change: Maritime Technology, Trade, and Economic Development," *American Economic Review*, 2017, *107* (9), 2821–2854.

Pennington, Jeffrey, Richard Socher, and Christopher D. Manning, "GloVe: Global Vectors for Word Representation," in "Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)" 2014, pp. 1532–1543.

Rashdall, Hastings, *The Universities of Europe in the Middle Ages*, Oxford: Clarendon Press, 1895. Reissued ed., F. M. Powicke and A. B. Emden, 3 vols., Oxford University Press, 1936.

Romer, Paul M., "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 1986, *94* (5), 1002–1037.

Rubin, Jared, "Printing and Protestants: An Empirical Test of the Role of Printing in the Reformation," *Review of Economics and Statistics*, 2014, *96* (2), 270–286.

Salop, Steven C., "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 1979, *10* (1), 141–156.

Smith, Adam, *An Inquiry into the Nature and Causes of the Wealth of Nations*, London: W. Strahan and T. Cadell, 1776.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge, "What Are We Weighting For?," *Journal of Human Resources*, 2015, *50* (2), 301–316.

Squicciarini, Mara P. and Nico Voigtländer, "Human Capital and Industrialization: Evidence from the Age of Enlightenment," *Quarterly Journal of Economics*, 2015, *130* (4), 1825–1883.

Stigler, George J., "The Division of Labor is Limited by the Extent of the Market," *Journal of Political Economy*, 1951, *59* (3), 185–193.

Talbert, Richard J. A., ed., *Barrington Atlas of the Greek and Roman World*, Princeton, NJ: Princeton University Press, 2000.

Taylor, Alexander N. and Jacob R. Hall, "The King's French: The Political Economy of Language and National Identity," 2026. Working paper, May 2026.

Universal Short Title Catalogue, "USTC Database," University of St Andrews, https://ustc.ac.uk 2022. Local snapshot dated 2022-12-13.

Verger, Jacques, *Les universités au Moyen Age*, 2nd ed., Paris: Presses Universitaires de France, 2007.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, MA: MIT Press, 2010.

Appendices (for online publication)

A Data sources and processing detail

This appendix documents the data inputs, parameter choices, and processing steps that produce the panel used in Sections 6–9. The goal is to let a reader replicate our construction without reading the scripts.

A.1 USTC corpus

Source. Universal Short Title Catalogue (USTC), University of St Andrews (Universal Short Title Catalogue, 2022). Local snapshot dated 2022-12-13.

Coverage extracted. All editions with `year` in $[1450, 1649]$ and a non-missing `place` field. The working corpus is 788,299 editions across 1,339 unique cities and seven principal languages (Latin, German, French, Italian, English, Dutch, Spanish).

Variables used. `ustc_sn` (catalog ID), `short_title` (used for embedding), `place` (printing city), `year`, `language1`, `classification1` (top-level USTC subject), and printer/author identifiers (used for diagnostics, not the headline regressions).

Cleaning. Place names are passed through a hand-curated Latin-imprint dictionary (e.g., `Lugduni` $\rightarrow$ `Lyon`, `Antverpiae` $\rightarrow$ `Antwerpen`, `Lutetiae` $\rightarrow$ `Paris`, `Wittembergae` $\rightarrow$ `Wittenberg`). Approximately 75 high-frequency Latin imprint variants are mapped to canonical Buringh spellings. Residual place-name normalization (German umlauts, accented characters, historical spelling variants) is handled separately.

A.2 Population panel — Buringh

Source. Buringh (2021a), "European Urban Population, 700–2000" (DANS, version 1.0; documented in Buringh (2021b)).

Coverage. $\sim$2,260 European and Mediterranean cities with population estimates at century anchors back to 700 CE and half-century anchors thereafter (1500, 1550, 1600, 1650, 1700, …,

2000).

Anchor years for our MA panel. We retain the five half-century anchors 1450, 1500, 1550, 1600, and 1650 for the regression panel. The 1450 anchor is constructed from Buringh by linear interpolation between the 1400 and 1500 native century anchors.

A.3 Population panel — Bosker–Buringh MENA supplement

Source. Bosker et al. (2013), "From Baghdad to London," *Review of Economics and Statistics*.

Coverage extracted. Cities in MENA buffer countries (Algeria, Egypt, Iraq, Israel, Lebanon, Libya, Morocco, Saudi Arabia, Syria, Tunisia, Yemen) at century anchors back to 800 CE. $\sim$70 cities total, of which 44 have positive population in the 1450–1650 window and so enter the MA gravity sum (Figure 11). The buffer is restricted to Mediterranean cities that contribute to European cities' market access through Levantine trade.

Why we use this. A sixteenth-century Mediterranean port like Cairo, Aleppo, or Istanbul is part of the economic environment of Venice, Marseille, and Barcelona. Excluding the MENA buffer would underestimate Mediterranean cities' market access by a non-trivial amount. The buffer enters MA only as origin/destination populations, not as a panel of cities to be measured for knowledge composition.

A.4 City-panel construction

Each source is passed through a linear-interpolation step at the target anchor years, then stacked. A unique-city catalog is built by rounding lat/lon to 0.01° and deduplicating; each unique city receives a `unified_city_id` code. Where Buringh and Bosker–Buringh report the same physical city, Buringh's row is kept. Figure 11 plots the 2,169 cities with positive population at any anchor year in the 1450–1650 window (the cities that actually contribute to the MA gravity sum), color-coded by source and sized by max population in that window.

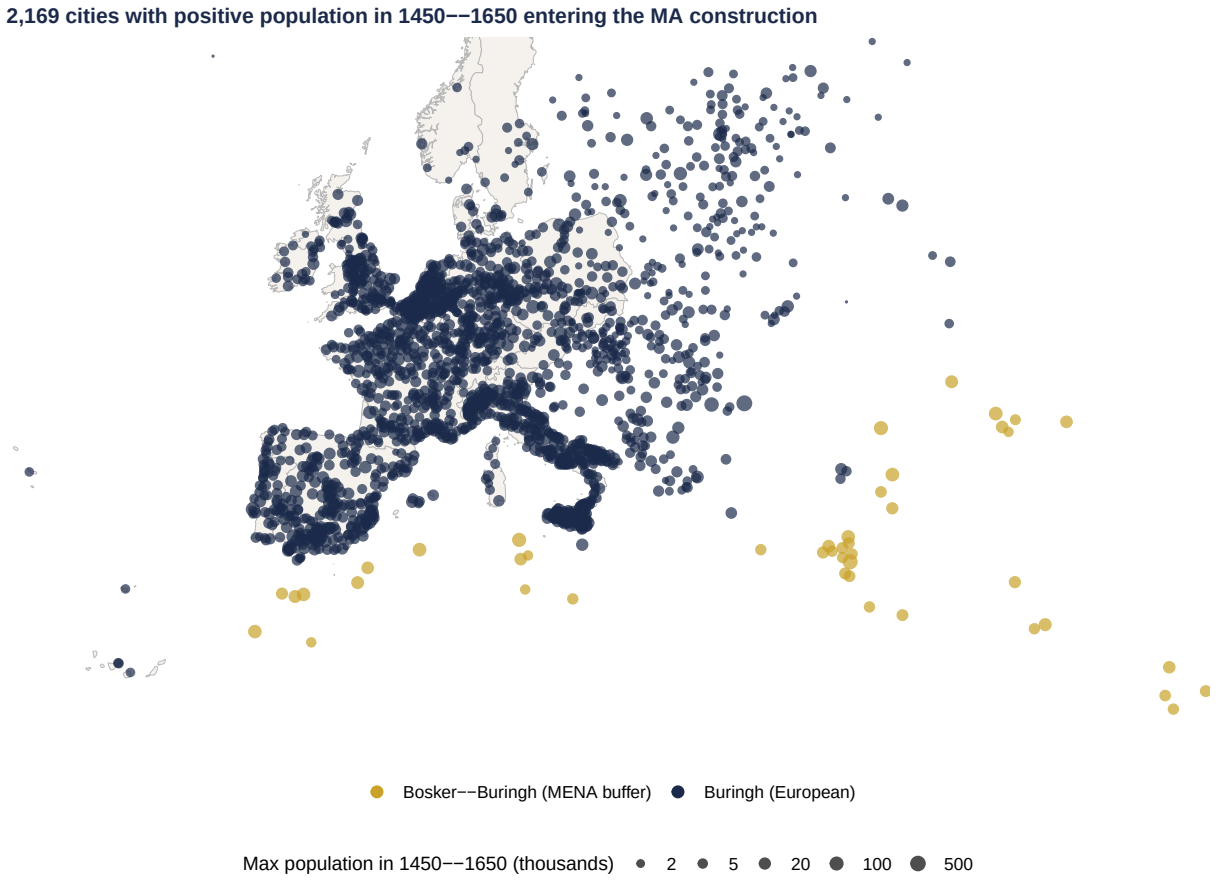


Figure 11: The 2,169 cities with positive population in 1450–1650 that feed the MA construction: the 2,125-city Buringh European set (navy) plus the 44-city Bosker–Buringh MENA supplement (gold). Dot size is proportional to the maximum population of each city over the five anchor years (1450, 1500, 1550, 1600, 1650) on a log scale; a handful of cities in the Buringh source dataset whose population is zero throughout the window are excluded (they contribute nothing to the gravity sum). The MENA buffer enters MA only as origin/destination populations contributing to the proximity-weighted sums for Mediterranean and southern European cities; it is not a panel of cities measured for knowledge composition.

A.5 Transport network shapefiles

All shapefiles are reprojected to the project's Europe Equidistant Conic projection (EEC) before rasterization.

Roman roads. 1,585 line features covering the Roman trunk network across Europe (McCormick et al. *Origins of the European Economy* derivative).

Medieval trade routes. 56 line features capturing major medieval long-distance routes (Hanseatic links, Alpine passes, Champagne-fair connectors).

Rivers. 184 navigable-river line features clipped to the European bounding box.

Seas. Single polygon covering the Mediterranean, Atlantic coast, Baltic, and North Sea.

A.6 Coordinate reference system

Europe Equidistant Conic projection with the following PROJ string:

```
+proj=eqdc +lat_0=0 +lon_0=0 +lat_1=43 +lat_2=62 +ellps=intl +units=m
```

Standard parallels at 43°N and 62°N minimize linear distance distortion across the European latitude range.

A.7 Cost-grid construction and cost specifications

A 10 km × 10 km grid covers the bounding box of all city coordinates. Each grid cell is assigned a least-cost-per-cell value based on the most efficient transport feature available there. The four cost specifications (relative costs per cell unit length) are listed in Table 7.

Table 7: Cost specifications. Lower values mean cheaper transport per unit length; portage is the numeraire at 1.00. Per-spec correlations on log(MA) at the city level are 0.88–0.98.

Specification	Porters	Roads	Rivers	Seas	Source
`spec1_bairoch` (primary)	1.00	0.81	0.21	0.08	Bairoch (1988)
`spec2_boerner_etal`	1.00	0.50	0.50	0.13	Boerner and Severgnini (2014)
`spec3_masschaele`	1.00	0.81	0.51	0.10	Masschaele (1997)
`spec4_campbell_etal`	1.00	0.50	0.24	0.18	Galloway et al. (1996)

A.8 Market access aggregation

For each city i at each anchor year y and each cost spec s:

$$\mathrm{MA}_{iy}^{(s)} = \sum_{j \neq i} \frac{\mathrm{pop}_{jy}}{\tau_{ij}^{(s)}},$$

with gravity exponent $\theta = 1$. The diagonal of the cost matrix is set to $+\infty$ so that $i = j$ drops from the sum (no-own-city variant). Two own-cost robustness variants replace the diagonal with $\tau_{ii} = 0.05$ and $\tau_{ii} = 0.10$ respectively before summing.

A.9 USTC place ↔ Buringh unified-city crosswalk

USTC `place` names are matched to Buringh `city` names through three stages: (1) exact match on cleaned name strings, (2) Stringdist match (Jaro–Winkler distance, threshold 0.85) for residuals with hand-review of borderline cases, (3) hand assignments for ∼150 remaining city names with unusual historical variants.

Coverage. 1,292 USTC place names map to a unified city ID; the residual ∼50 USTC places without a Buringh match are dropped from the regression panel. These are either very small print villages or non-European outposts.

A.10 Polity panel — Abramson shapefiles

Source. Abramson (2017), *International Organization*. Polity shapefiles distributed with the replication package, native projection Europe Albers Equal Area Conic, coverage every five years between 1100 and 1815.

Processing. Reprojected from Albers EAC to our EEC projection; point-in-polygon matching of each (city, decade) cell to the nearest available Abramson five-year file. The polity panel is included in the regression input as a candidate fixed-effect dimension; the headline specifications do not use polity FE.

A.11 Embedding pipeline parameters

All hyperparameters are set in `scripts/config/_project.R` (Table 8). Embeddings are trained *per language*: separate PPMI matrix and SVD for Latin, German, French, Italian, English, Dutch, Spanish. Within each language, the 200-dimensional title embedding for each title is computed by SIF-weighted averaging of the constituent word vectors and removing the first principal component from the resulting title-embedding matrix.

Table 8: Embedding-pipeline hyperparameters.

Parameter	Value	Notes
SEED	20251023	fixed for all stochastic steps
EMBED_DIM	200	SVD truncation rank
PMI_WINDOW	5	skipgram context window
PMI_MIN_N	20	minimum word frequency for PMI vocabulary
SIF_A	10^{-3}	SIF smoothing constant (Arora et al., 2017)
KMEANS_K	10	k for topical-entropy metric

A.12 Title-length descriptive statistics

The embedding pipeline trains on `short_title` fields from the USTC corpus. Because skip-gram window models and SIF averaging (Mikolov et al., 2013; Arora et al., 2017) are conventionally applied to longer text units, it is reasonable to ask whether USTC short titles supply enough words per document to populate either step. Table 9 reports per-language title-length statistics after a simple whitespace-and-punctuation tokenization of `short_title`. The pooled median is 17 words; only 12% of titles fall below the 5-word skip-gram window, and 80% have at least 7 words. The shortest distributions are Dutch (median 13) and French (median 12); the longest is English (median 34), reflecting the corpus's long descriptive subtitles in seventeenth-century English imprints. The vernacular distributions are sufficient for the PPMI co-occurrence step (5-word window, minimum word frequency 20) and for the SIF title aggregation that averages whatever word vectors are present.

Table 9: USTC title length, in words, by language. Cells report mean, quantiles, and shares of the per-title word count after a simple whitespace–punctuation split of `short_title`. 'Share < 5' is the share of titles below the 5-word skip-gram window; 'share ≥ 7' is the share of titles long enough to populate the long-title sub-sample used as a robustness check.

Language	n_{titles}	Mean	p10	p25	Median	p75	p90	Share < 5	Share ≥ 7
Latin	297,886	19.29	4.00	8.00	17.00	29.00	35.00	0.13	0.79
German	114,985	23.39	4.00	8.00	24.00	34.00	40.00	0.14	0.80
Italian	64,702	20.85	6.00	10.00	18.00	28.00	39.00	0.07	0.87
French	93,085	15.59	4.00	6.00	12.00	21.00	34.00	0.16	0.74
English	50,202	34.70	9.00	19.00	34.00	43.00	59.00	0.03	0.94
Dutch	38,911	14.39	3.00	6.00	13.00	21.00	29.00	0.16	0.75
Spanish	32,105	20.72	4.00	8.00	16.00	29.00	43.00	0.11	0.81
Pooled (7 langs)	691,876	20.53	4.00	8.00	17.00	30.00	39.00	0.12	0.80

A.13 Vernacular nearest-neighbor validation

Section 2 reports the nearest-neighbor figure for three Latin seed words as evidence that the embedding pipeline recovers substantively coherent intellectual neighborhoods. The same exercise carried out separately for the five vernacular languages confirms that the pipeline is not special to Latin: in each language the ten nearest words to a religious seed (church), a polity seed (war), and a learned-discourse seed (law) cluster in semantically coherent neighborhoods. Figures 12–16 report the result for each vernacular: German (*kirche*, *krieg*, *recht*), Italian (*chiesa*, *guerra*, *legge*), French (*eglise*, *guerre*, *loy*), Dutch (*kercke*, *oorloghen*, *recht*), and Spanish (*iglesia*, *guerra*, *ley*). The Dutch and French seeds use early-modern spellings (*kercke* not *kerk*; *oorloghen* not *oorlog*; *loy* not *loi*) because the USTC corpus pre-dates spelling standardization in those languages and the modern lemma forms have very low frequency in our vocabulary.

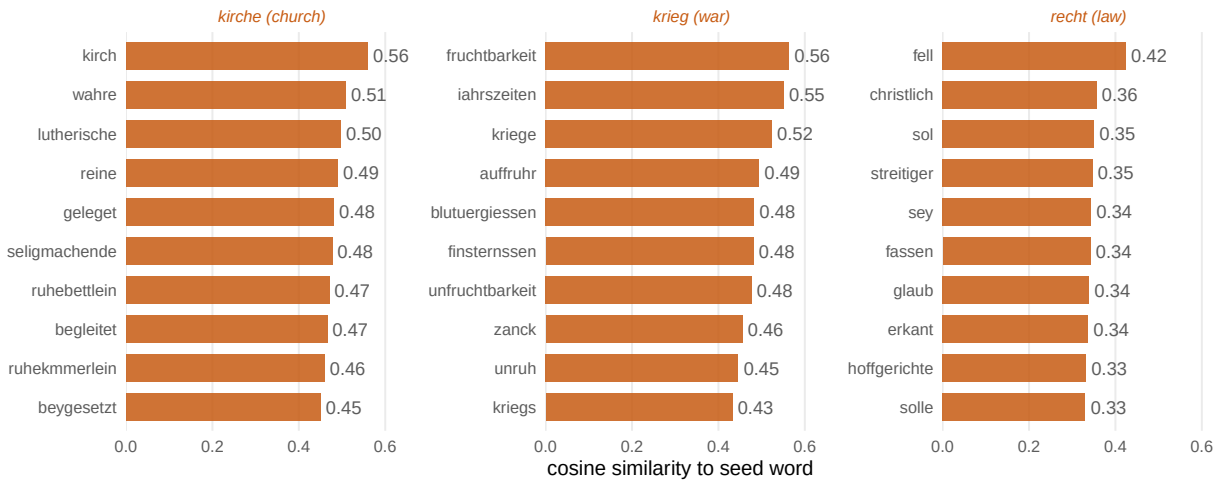


Figure 12: German: ten nearest words to each of *kirche* (church), *krieg* (war), and *recht* (law), ranked by cosine similarity in the German title-embedding space.

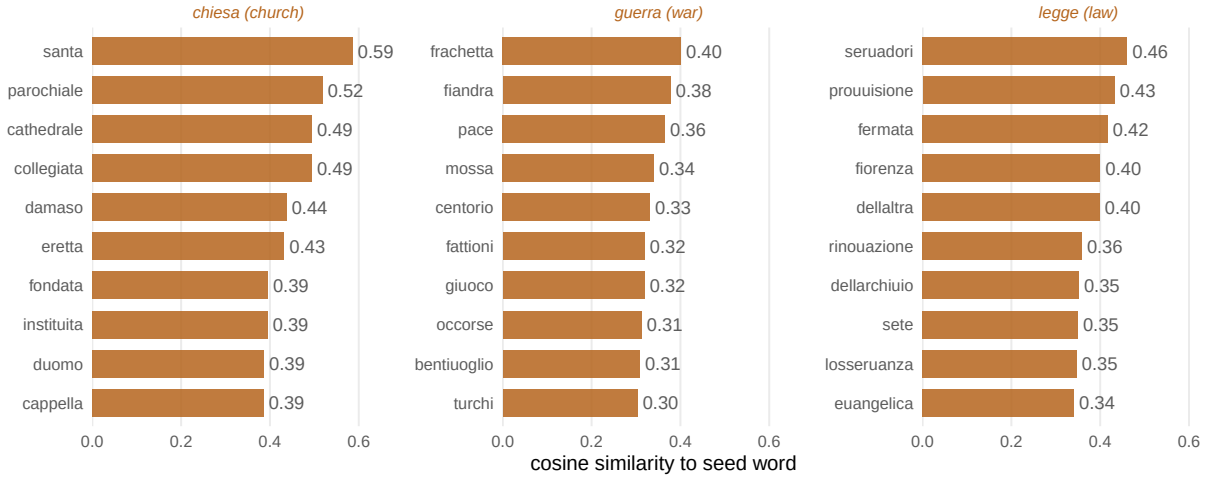


Figure 13: Italian: ten nearest words to each of *chiesa* (church), *guerra* (war), and *legge* (law).

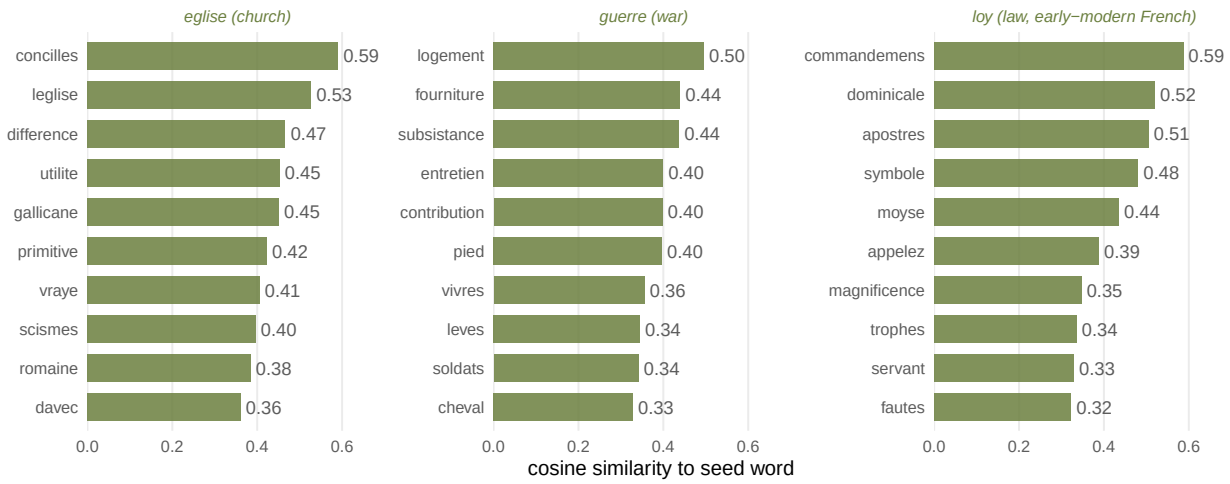


Figure 14: French: ten nearest words to each of *eglise* (church), *guerre* (war), and *loy* (law, early-modern French).

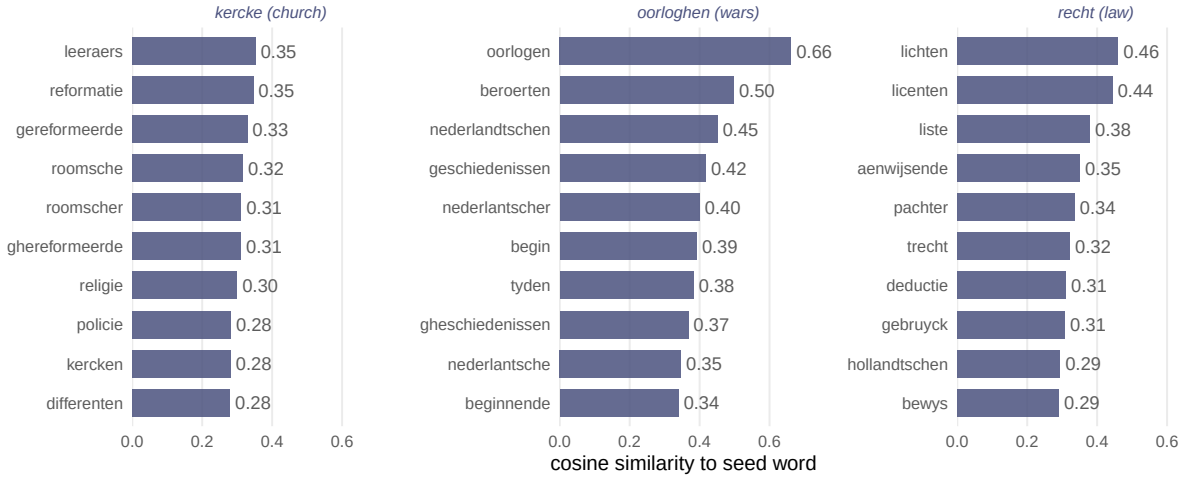


Figure 15: Dutch: ten nearest words to each of *kercke* (church), *oorloghen* (wars), and *recht* (law). Spellings reflect pre-standardization early-modern Dutch in the USTC corpus.

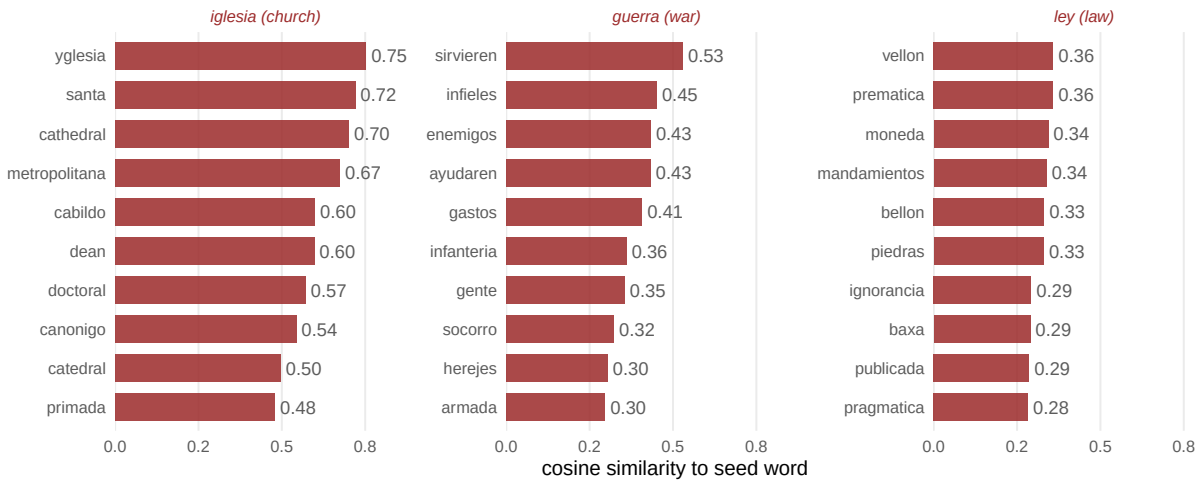


Figure 16: Spanish: ten nearest words to each of *iglesia* (church), *guerra* (war), and *ley* (law).

A.14 Embedding-validation figure: Latin titles in UMAP space

Figure 17 shows a stratified sample of 5,004 Latin titles projected from the 200-dimensional title-embedding space to two dimensions with UMAP (`uwot::umap`, `n_neighbors` = 15, `min_dist` = 0.10, cosine metric), colored by USTC `classification1`. The sample is stratified to the nine largest classifications by frequency plus "Other". Two observations support the embedding-vs-classification argument made in Section 2. *First*, classifications occupy partly-overlapping regions of the manifold rather than disjoint clusters—Academic Dissertations and Medical Texts form recognizable tight regions, but Religious titles scatter widely across most of the manifold and Jurisprudence/Poetry/Other overlap heavily. Discrete classification labels therefore flatten substantial intra-category variation. *Second*, even where classifications cluster (Academic Dissertations in the lower-left), the embedding-space neighborhoods include a non-trivial fraction of titles from other classifications—a Religious title in the same neighborhood as the Academic-Dissertations cluster is, by the embedding's logic, a Religious title written in the academic-disputation idiom, which the discrete `classification1` field cannot distinguish from a Religious title in another idiom. The UMAP is exploratory and is not used in any inferential procedure; the continuous metrics in Section 6 operate on the full 200-dimensional space, not the projection.

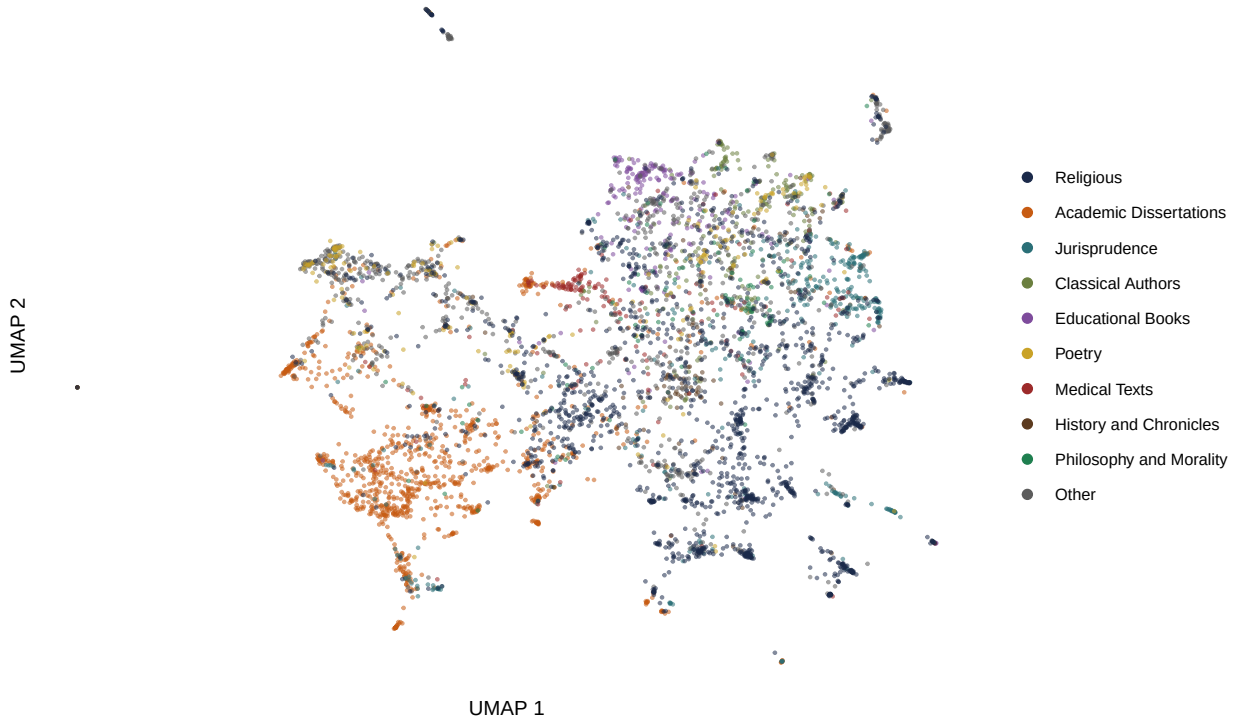


Figure 17: UMAP projection of 5,004 Latin titles from the 200-dimensional title-embedding space to two dimensions, colored by USTC `classification1`. Stratified sample from the nine largest classifications by frequency plus "Other". Academic Dissertations cluster tightly; Religious titles scatter across the manifold.

A.15 Software environment

R 4.5.1 on macOS 25.2. Key packages: `fixest` 0.14 (panel regressions), `sf` (vector spatial), `terra`/`raster` 3.6 (raster spatial), `gdistance` 1.6 (least-cost paths), `text2vec` (PPMI/SVD), `uwot` (UMAP), `fpc` and `cluster` (cluster diagnostics), `haven` (Stata `.dta` ingest), `here` (path resolution). Full lockfile in `renv.lock`.

A.16 Descriptive statistics for the six-language base panel

Table 10 reports descriptive statistics for the four headline variables of the six-language base panel used in Sections 8 and 9: diversity, distinctiveness, specialization entropy, and log MA. Each cell reports a sample mean above its standard deviation in brackets, computed over the (language, city, 50-year period) observations in the regression panel. The bottom rows report the total number of

titles in the underlying USTC corpus, the number of unique cities contributing at least one cell, and the number of (language, city, period) cells per language and pooled.

Table 10: Descriptive statistics, 6-language base panel

	Latin (1)	German (2)	Italian (3)	French (4)	Dutch (5)	Spanish (6)	Pooled (7)
Diversity	0.63	0.60	0.63	0.59	0.62	0.63	0.62
	[0.11]	[0.13]	[0.15]	[0.15]	[0.15]	[0.14]	[0.13]
Distinctiveness	0.74	0.59	0.70	0.71	0.62	0.59	0.69
	[0.22]	[0.22]	[0.21]	[0.19]	[0.19]	[0.21]	[0.22]
Specialisation (entropy)	0.59	0.92	0.72	0.70	0.62	0.67	0.69
	[0.39]	[0.37]	[0.35]	[0.39]	[0.38]	[0.38]	[0.40]
log(MA)	-2.68	-2.76	-2.54	-2.57	-2.48	-2.59	-2.64
	[0.45]	[0.44]	[0.40]	[0.47]	[0.55]	[0.41]	[0.45]
Total titles	294,603	115,199	62,884	89,805	34,069	29,063	625,623
Unique cities	385	181	108	144	66	77	506
City-period cells	795	355	209	258	119	142	1,878

Notes: Descriptive statistics for the 6-language base panel used in §8 and §9. Reported cell is one (language, city, 50-year period) observation. Each row lists the sample mean above the standard deviation in brackets. The diversity metric is the mean cosine distance from the city centroid to the titles printed in that city-period; distinctiveness is the cosine distance from the city centroid to the contemporaneous European centroid for that language; specialization is the Shannon entropy of titles across $K = 10$ topical clusters. log(MA) is the gravity-style log market access constructed in §5.

A.17 Eastern-cities robustness

Some print-active cities in our panel lie east of the bulk of European cities and may have noisier MA estimates, both because the bilateral-cost network (Roman roads, medieval routes, navigable rivers, seas) has sparser shapefile coverage east of roughly the Vistula (∼longitude 20°E), and because the surrounding city universe over which MA is summed is itself thinner in that region. To test whether the headline coefficients in Section 8 are driven by this eastern fringe, we re-run the pooled BJ within-city specification dropping cities east of two longitude cutoffs: 25°E (cuts the smallest set, mostly Vilnius and a handful of Polish peripherals) and 20°E (cuts a larger set additionally including Kraków and several other Polish/Hungarian peripherals). Longitude is the per-place modal coordinate from the USTC-derived geo lookup (`Data/ustc_europe_geo_coords_10-19-25.csv`).

Table 11 reports the result. The three signed within-city coefficients are essentially unchanged across both restrictions, with point estimates moving by no more than 0.04 (about a quarter of a

Table 11: Headline pooled BJ within-city coefficients on log(MA) under successive longitude restrictions. Dropping cities east of 25°E (mostly Vilnius and a few Polish peripherals) or east of 20°E (additionally Kraków, Tartú, several Polish/Hungarian peripherals) tests whether the headline is driven by the eastern fringe, whose MA values are most exposed to the sparser road/river shapefile coverage in that region.

	Full sample (1)	Drop lon > 25°E (2)	Drop lon > 20°E (3)
Diversity	0.456***	0.450***	0.448***
	(0.065)	(0.066)	(0.067)
Distinctiveness	-0.667***	-0.671***	-0.671***
	(0.128)	(0.130)	(0.132)
Specialisation (entropy)	0.712***	0.690***	0.666***
	(0.228)	(0.229)	(0.234)
Observations (any metric)	1884	1875	1843
Unique cities	506	501	486

Notes: Pooled BJ within-city coefficient $\hat{\beta}_W$ on the demeaned $\widetilde{\log \mathrm{MA}}_{i\ell,d}$ from the headline Mundlak/Bell-Jones specification of §7. Standard errors clustered at the city in parentheses. Observations weighted by $\sqrt{n_{\text{titles}}}$. Period and language FE included; the between-city coefficient $\hat{\delta}$ on $\overline{\log \mathrm{MA}}_{i\ell}$ is estimated but suppressed from the table. Eastern cutoffs use the per-place modal longitude from Data/ustc_europe_geo_coords_10-19-25.csv. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

standard error) and significance levels unaffected for diversity and distinctiveness. The entropy coefficient attenuates modestly under the more aggressive 20°E cutoff but remains positive at the 5% level. The eastern fringe is not what is driving the Section 8 results.

A.18 Headline regression tables: Sections 8 and 9

The figures in Sections 8 and 9 report coefficient plots; the full underlying regressions are collected here in standard economics-journal table format. Table 12 is the headline within-city composition regression visualized in Figure 7; Table 13 is the growth regression visualized in Section 9; Table 14 is the trade-elasticity robustness panel referenced in Section 5; Table 15 is the cost-specification robustness panel referenced in Section 8.

Table 12: Within-city market access and print composition (Mundlak/Bell-Jones)

	Pooled (1)	Latin (2)	German (3)	Italian (4)	French (5)	Dutch (6)	Spanish (7)
Panel A: Diversity							
Within $\widetilde{\log \mathrm{MA}}_{i\ell d}$	0.467***	0.378***	-0.069	1.018***	0.913***	0.420***	0.386*
	(0.064)	(0.055)	(0.131)	(0.126)	(0.148)	(0.126)	(0.230)
Between $\overline{\log \mathrm{MA}}_{i\ell}$	0.008	0.022	-0.060**	-0.008	0.051*	0.007	-0.055
	(0.013)	(0.014)	(0.029)	(0.034)	(0.028)	(0.032)	(0.035)
Panel B: Distinctiveness							
Within $\widetilde{\log \mathrm{MA}}_{i\ell d}$	-0.709***	-0.058	-0.946***	-1.287***	-1.667***	-1.060***	-0.728
	(0.133)	(0.109)	(0.238)	(0.203)	(0.444)	(0.369)	(0.502)
Between $\overline{\log \mathrm{MA}}_{i\ell}$	-0.020	0.008	0.029	-0.027	-0.143*	-0.129**	0.119
	(0.028)	(0.033)	(0.034)	(0.040)	(0.081)	(0.050)	(0.079)
Panel C: Specialisation (entropy)							
Within $\widetilde{\log \mathrm{MA}}_{i\ell d}$	0.751***	0.425	0.093	1.700***	1.715***	1.345	0.648
	(0.223)	(0.342)	(0.334)	(0.338)	(0.628)	(0.822)	(0.500)
Between $\overline{\log \mathrm{MA}}_{i\ell}$	-0.075*	-0.158**	-0.098	0.017	0.124	0.031	-0.229**
	(0.039)	(0.067)	(0.060)	(0.085)	(0.089)	(0.124)	(0.102)
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Language FE	Yes	No	No	No	No	No	No
R^2 (within)	0.098	0.110	0.038	0.353	0.329	0.085	0.138
Observations	1,884	796	357	209	258	122	142

Notes: Each entry reports the Mundlak/Bell-Jones decomposition of $\log(\mathrm{MA})$ into within-city deviations $\widetilde{\log \mathrm{MA}}_{i\ell d}$ and city means $\overline{\log \mathrm{MA}}_{i\ell}$. Standard errors clustered at the city are shown in parentheses below the coefficient. Observations are weighted by $\sqrt{n_{\text{titles}}}$. Period fixed effects always included; language fixed effects are included in the pooled column only. Sample: USTC, 1450–1640, four 50-year periods. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 13: Print composition and subsequent city growth (Barro long-difference)

	Pooled	Latin	German	Italian	French	Dutch	Spanish
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Diversity							
Diversity$_{i,t}$	0.143*	0.508***	-0.127	-0.251**	-0.123	0.643**	0.147
	(0.082)	(0.145)	(0.143)	(0.117)	(0.183)	(0.263)	(0.458)
log(MA)$_{i,t}$	0.073**	0.079**	-0.007	0.059**	0.055	0.010	0.303***
	(0.029)	(0.033)	(0.031)	(0.028)	(0.050)	(0.069)	(0.094)
log(pop)$_{i,t}$	-0.009	-0.022	-0.045***	0.025*	0.030	0.015	0.019
	(0.015)	(0.015)	(0.015)	(0.014)	(0.026)	(0.049)	(0.026)
Panel B: Distinctiveness							
Distinctiveness$_{i,t}$	-0.058	-0.091	0.066	0.130	-0.137	-0.303	-0.129
	(0.065)	(0.117)	(0.098)	(0.107)	(0.124)	(0.232)	(0.210)
log(MA)$_{i,t}$	0.072**	0.085**	-0.002	0.065**	0.037	-0.005	0.313***
	(0.029)	(0.034)	(0.030)	(0.028)	(0.049)	(0.080)	(0.098)
log(pop)$_{i,t}$	-0.006	-0.006	-0.042**	0.022	0.005	-0.017	0.015
	(0.014)	(0.015)	(0.016)	(0.017)	(0.021)	(0.044)	(0.027)
Panel C: Specialisation (entropy)							
Specialisation (entropy)$_{i,t}$	-0.036	-0.055	-0.014	-0.093*	0.007	0.007	0.028
	(0.031)	(0.041)	(0.055)	(0.047)	(0.073)	(0.110)	(0.087)
log(MA)$_{i,t}$	0.070**	0.078**	-0.001	0.065**	0.049	0.026	0.300***
	(0.029)	(0.034)	(0.030)	(0.028)	(0.050)	(0.077)	(0.104)
log(pop)$_{i,t}$	-0.003	-0.008	-0.045***	0.020	0.022	0.015	0.023
	(0.015)	(0.016)	(0.016)	(0.013)	(0.026)	(0.054)	(0.029)
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Language FE	Yes	No	No	No	No	No	No
R^2 (within)	0.017	0.044	0.036	0.069	0.036	0.059	0.108
Observations	1,882	794	357	209	258	122	142

Notes: Dependent variable is $\Delta_{50}\log(\text{pop})_{i,t\to t+50}$. Each panel regresses subsequent 50-year log-population growth on the contemporaneous composition metric, log(MA), and initial log(pop), with period and (for the pooled column) language fixed effects. Standard errors clustered at the city are shown in parentheses below the coefficient. Observations are weighted by $\sqrt{n_{\text{titles}}}$. Sample: USTC city-period cells with non-missing population at both t and $t+50$. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.10$.

Table 14: Robustness of pooled BJ coefficients to choice of trade elasticity θ

	Trade elasticity θ				
	0.5	1.0	2.0	3.8	5.0
	(1)	(2)	(3)	(4)	(5)
Diversity	0.515***	0.500***	0.377***	0.192***	0.150***
	(0.076)	(0.062)	(0.045)	(0.024)	(0.017)
Distinctiveness	-0.533***	-0.683***	-0.508***	-0.265***	-0.023
	(0.142)	(0.124)	(0.106)	(0.064)	(0.108)
Specialisation (entropy)	-0.012	0.566***	0.376**	0.158	-0.148
	(0.271)	(0.213)	(0.178)	(0.136)	(0.132)
Observations	1,884	1,884	1,884	1,884	1,884

Notes: Each cell reports the pooled within-city Mundlak/Bell-Jones coefficient on log(MA) for one of the three composition metrics (rows) under a different choice of trade elasticity θ (columns) in the MA aggregation $\text{MA}_{iy} = \sum_{j \neq i} \text{pop}_{jy} / \tau_{ij}^{\theta}$. Column $\theta = 1.0$ is the headline; $\theta = 3.8$ is the structurally estimated value in Donaldson (2018). The spec1_bairoch cost matrix is held fixed across columns; only the MA aggregation varies. Standard errors clustered at the city in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 15: Cross-specification robustness of the headline pooled BJ coefficients on log(MA). Each column re-estimates the pooled within-city BJ regression using a different cost-coefficient parametrisation; the headline panel uses Spec 1 (Bairoch). Cross-spec correlations on log(MA) range from 0.81 to 0.97, with the headline Spec 1 correlating ≥ 0.91 with each of the other three. SEs clustered at the city in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Bairoch	(2) Boerner	(3) Masschaele	(4) Campbell
Diversity	0.462***	0.465***	0.451***	0.488***
	(0.057)	(0.055)	(0.054)	(0.058)
Distinctiveness	-0.621***	-0.607***	-0.598***	-0.652***
	(0.115)	(0.110)	(0.109)	(0.121)
Specialisation (entropy)	0.532***	0.473**	0.463**	0.595***
	(0.196)	(0.200)	(0.192)	(0.209)
Observations	1,884	1,884	1,884	1,884

B Theoretical framework

This appendix develops the framework that motivates the three signed predictions tested in Section 8. The framework combines monopolistic competition with spatial product differentiation. It rests on five identifiable assumptions about latent features of early-modern book production and consumption; each assumption is needed for a specific subset of the predictions. The conditional-prediction structure below lets us read the sign of each empirical coefficient as evidence about which assumptions hold for European print in 1450–1649.

B.1 Setting and notation

Index cities by i (consumer location) and j (producer location). A "variety" is a single title, denoted by its position $x \in X$ in a product space X. We take X to be the 200-dimensional title-embedding space introduced in Section 2; the theory needs only that X is a metric space.

Let L_i denote the size of the book-reading population in city i; pop_i denotes Buringh population. Let $\tau_{ij} \geq 1$ denote the iceberg cost of shipping a book from j to i. Market access of city j is the

gravity-style sum

$$\text{MA}_j \;=\; \sum_{i \neq j} \frac{L_i}{\tau_{ij}^{\theta}},$$

where θ is the trade elasticity. Sections 5 and 6 explain how we operationalize L_i, τ_{ij}, and θ for the 1450–1649 period.

A central conceptual point is that our data are entirely on the production side. We observe what cities printed, not what individual readers consumed. The framework's predictions must therefore be cast as predictions about production. Section B.3 states each prediction in production-side terms; the consumer-side variety expansion familiar from Schiff (2015); Handbury and Weinstein (2015); Couture et al. (2021) is a downstream consequence of the same mechanism but requires consumer-side data we do not have.

B.2 Five building blocks

The full framework borrows from three literatures: Dixit and Stiglitz (1977), Krugman (1979, 1980), Helpman and Krugman (1985), and Fujita et al. (1999) for monopolistic competition with CES preferences and increasing returns; Hotelling (1929), Lancaster (1979), Salop (1979), and Anderson et al. (1992) for spatial product differentiation and strategic positioning; standard free-entry equilibrium for the remaining structure. Each block requires one assumption about 1450–1649 book markets. We label them A1–A5.

A1. Consumers value variety (CES aggregation). The representative consumer in city i aggregates titles via a CES function

$$U_i \;=\; \left[\sum_j \int_{x \in X_j} q_{ij}(x)^{(\sigma-1)/\sigma} \, dx \right]^{\sigma/(\sigma-1)},$$

where X_j is the set of titles produced in city j, $q_{ij}(x)$ is the consumer's quantity demanded of title x from j, and $\sigma > 1$ is the elasticity of substitution across titles. The latent claim about book demand is that early modern book consumers had at least some love-of-variety component in their aggregate preferences. The alternative not tested is authority-seeking demand (one canonical text per subject; $\sigma \to \infty$), confessional allegiance (variety bounded by religious identity), or pure status-good demand.

A2. Production has increasing returns at the title level. Each title is produced by a single firm with cost

$$C(q) \;=\; F + c\,q,$$

where F is the fixed cost of producing the title (typesetting, plate-making, editorial labor) and c is the marginal cost per copy (paper, ink, press labor). Equilibrium firm-level output is $q^* = F(\sigma - 1)/c$ and the number of titles produced in city j scales with city j's effective market size. The latent claim about printing technology is that fixed costs per title were non-trivial relative to per-copy costs. The alternative is constant returns ($F \to 0$), in which case no home-market effect operates. This claim is not directly attested by the historical literature: the fixed-cost-per-title question is rarely posed explicitly; the empirical confirmation of Prediction 1 below provides indirect evidence that A2 holds in the 1450–1649 European print industry.

A3. Spatially heterogeneous consumer preferences (positional differentiation). Each city's representative consumer has an ideal point $x_i^* \in X$ in product space, and utility from consuming title x depends on the distance $|x - x_i^*|$. Cities have heterogeneous ideal points: $x_i^* \neq x_j^*$ for $i \neq j$, at least on average. The latent claim is that cities had heterogeneous typical preferences over content. Wittenberg's reader had Lutheran-theological preferences; Madrid's had Catholic-courtly; Amsterdam's had commercial-cosmopolitan. The aggregate European reader has a centroid of all these positions.

Combined with A4 (below), A3 implies that the demand-weighted centroid for a producer in city j is

$$\bar{x}_j \;=\; \frac{\sum_i (L_i/\tau_{ij}^{\theta})\, x_i^*}{\sum_i (L_i/\tau_{ij}^{\theta})}.$$

High-MA cities have a $\bar{x}_j$ close to the European centroid (because they sum over many cities at low transport cost); low-MA cities have a $\bar{x}_j$ close to their own x_j^*. The variation in $\bar{x}_j$ across cities is the source of Prediction 2.

A4. Bilateral transport costs with proximity advantage. Books shipped from j to i are subject to iceberg cost $\tau_{ij} \geq 1$ that is decreasing in proximity. The cost surface is heterogeneous across city pairs. The latent claim is that book trade was costly enough that proximity mattered for trade volumes, but cheap enough that books did flow across cities (otherwise $\bar{x}_j$ collapses to x_j^* for every j and the distinctiveness prediction below has no bite).

A5. Firms choose product positions strategically. Each firm in city j chooses both whether to enter and what title to produce (its position $x \in X$). Profits depend on the density of demand at the chosen position, weighted by transport access. Firms within a city differentiate from each other Hotelling-style: a firm avoids head-to-head price competition by positioning away from its competitors. The latent claim is that printers in 1450–1649 had real choice over what to publish and were not constrained to a fixed topical template by local religious or political authorities. The historical literature has noted that strict censorship in post-Tridentine Spain or post-Reformation Wittenberg may have violated A5 in those specific contexts; the empirical question is whether A5 holds on average across the 1,300 USTC cities.

B.3 Predictions and their assumption requirements

Three signed predictions follow from the five assumptions. Each prediction requires only a specific subset.

Prediction 1: $\hat{\beta}_{diversity} > 0$ *(within-city diversity rises with MA).* From A1 + A2, the number of titles produced in city j is proportional to city j's effective market size divided by per-firm labor demand. Effective market size for a firm in city j is $\sum_i L_i \tau_{ij}^{-\theta} = \text{MA}_j$, so the number of titles $V_j \propto \text{MA}_j$ with the home-market multiplier from Krugman (1980): V_j rises more than proportionally with MA_j. More titles produced locally implies larger within-city dispersion in embedding space, which is our diversity metric.

Required assumptions. A1 (love of variety; otherwise no demand for differentiated titles), A2 (increasing returns; otherwise no home-market effect). A3, A4, A5 not required.

Prediction 2: $\hat{\beta}_{distinctiveness} < 0$ *(distinctiveness falls with MA).* From A3 + A4, the demand-weighted centroid $\bar{x}_j$ moves closer to the European centroid as MA_j rises. From A5, firms locate at or near $\bar{x}_j$. Therefore the average position of titles printed in j is closer to the European centroid for high-MA cities, and the city's distinctiveness (cosine distance from its centroid to the European centroid) is lower.

Required assumptions. A1, A2 (to ensure the variety count is high enough to characterize an average position), A3 (positional preferences; without it, no concept of distance from a centroid), A4 (proximity advantage; without it, $\bar{x}_j$ is the same for all cities), A5 (firms position strategically).

Prediction 3: $\hat{\beta}_{\textit{spec-entropy}} > 0$ *(topical entropy rises with MA).* From A5, firms within a city differentiate from each other. With more firms in city j (high MA via Prediction 1), the firms occupy more distinct positions in product space; the distribution of titles across topical clusters becomes more even; entropy rises.

Required assumptions. A1, A2 (to generate large V_j), A3 (positional preferences make differentiation meaningful), A5 (strategic differentiation). A4 not strictly required for the sign, but the magnitude of the effect depends on the demand-weighting structure.

B.4 Conditional predictions: signs if each assumption fails

The signed predictions above hold only if all the required assumptions hold. If a particular assumption fails, the corresponding prediction's sign weakens or reverses (Table 16).

Table 16: Conditional predictions: what each assumption's failure would imply for the corresponding coefficient sign.

Fails	Substantive interpretation	Predicted sign change
A1	Book demand was authority-seeking (one canonical text per subject) rather than variety-additive	$\hat{\beta}_{\text{diversity}} \to 0$ or < 0
A2	Printing was a constant-returns industry; fixed costs per title were trivial	$\hat{\beta}_{\text{diversity}} \to 0$
A3	European preferences were homogeneous across cities; no "mainstream" position concept	$\hat{\beta}_{\text{distinctiveness}} \to 0$; $\hat{\beta}_{\text{spec-entropy}}$ weakened
A4	Book trade was locally segmented; high transport costs across most city pairs	$\hat{\beta}_{\text{distinctiveness}} \to 0$ or > 0
A5	Printers did not strategically differentiate (e.g. under tight censorship or guild coordination)	$\hat{\beta}_{\text{spec-entropy}} \to 0$ or < 0

Reading the table in reverse, the empirical signs we estimate are informative about which specific assumptions hold:

- A positive $\hat{\beta}_{\text{diversity}}$ is consistent with both A1 and A2; a near-zero or negative slope is evidence that one of A1 or A2 fails (cannot tell which without further data).
- A negative $\hat{\beta}_{\text{distinctiveness}}$ is consistent with A3 and A4 holding; a near-zero or positive slope is evidence that one or both fail.
- A positive $\hat{\beta}_{\text{spec-entropy}}$ is consistent with A5; near-zero or negative is evidence that strategic differentiation does not operate.

B.5 Implications for Section 9 (city growth)

The framework also has implications for the city-growth analysis in Section 9, and they are clean enough to be worth stating explicitly.

The framework's headline growth prediction is about MA, not composition. Krugman (1991) and Fujita et al. (1999) extend the static Helpman-Krugman framework to dynamic agglomeration via cumulative causation: high-MA cities attract firms (forward linkages to demand) and workers (backward linkages to supply), and the resulting agglomeration further raises their MA. Cities at trade-integration nodes grow faster than peripheral cities in a way that compounds over time. The pooled $\hat{\gamma}_{\text{MA}} = 0.073^{**}$ ($p \approx 0.013$) in Section 9.3 is the reduced-form version of this prediction.

The framework predicts composition and growth are co-determined *by MA, not causally linked.* In the static Helpman-Krugman equilibrium, a city's MA determines both its production composition (number of titles, positions in product space, evenness across topical space — Section 8 outcomes) and its growth dynamics. Both are downstream consequences of MA. Conditional on MA, composition should not carry additional predictive power for subsequent growth. Our Barro specification of Section 9.1 —which controls for $\log(\text{MA})_{i,t}$ and $\log(\text{pop})_{i,t}$ when regressing on composition—tests exactly this conditional-on-MA prediction. The framework says the conditional coefficient on composition should be approximately zero. Our pooled nulls for all three composition metrics are therefore the model-consistent finding, not a failure of the framework.

Frameworks that would predict a non-zero composition-to-growth coefficient. Section 9.5 (and the extensions table therein) spells out three alternatives that would generate a positive $\hat{\beta}$ on composition net of MA: endogenous growth via knowledge spillovers (Romer, 1986; Grossman and Helpman, 1991), Jacobs externalities (Glaeser et al., 1992), and Marshall externalities (Hender-

son, 2003). Our pooled nulls reject all three at the 50-year horizon and conditional on MA. This is a substantive negative finding about each extension at this horizon, with the standard caveats: limited statistical power, the 50-year window may be too short for cumulative-causation dynamics that operate over centuries, and we cannot distinguish a Mokyr-style mechanism that is real but operating at continental rather than city-level resolution.

B.6 Inverse inference and what the historical literature does not already know

The paper's contribution is empirical, not theoretical. We do not derive a formal structural model. What our paper does is *inverse inference under the framework*: assuming the framework's broad structure holds, our reduced-form findings let us reverse-infer which of the five latent assumptions about 1450–1649 European book markets must hold (at least approximately). The empirical confirmation of all three signed predictions in Section 8 is jointly consistent with all five assumptions A1–A5 holding for this period and this corpus.

Some of these inverse inferences map to questions the historical literature on early-modern print has not directly resolved:

Increasing returns to printing at the title level (A2). The historical literature on the printing press has long discussed press-level setup costs (the press itself was a substantial capital investment) and print-run-level economies of scale (longer print runs spread the typesetting cost over more copies). What it has not directly settled is whether the per-title fixed cost was substantial relative to the per-copy variable cost in the 1450–1649 European industry. Our confirmation of Prediction 1 ($\hat{\beta}_{\text{diversity}} > 0$) is direct evidence that A2 holds: cities with larger effective markets produced disproportionately more distinct titles, exactly as the home-market multiplier predicts under non-trivial $F/(cq^*)$.

Spatially heterogeneous demand across European cities (A3). Historians of the Republic of Letters (Mokyr, 2002, 2009, 2017) emphasize the transnational character of early-modern European intellectual life—scholars in Paris read texts produced in Wittenberg, Leiden, Padua. A natural inference is that demand was relatively homogeneous across European cities (every educated European read a common canon). Our confirmation of Prediction 2 ($\hat{\beta}_{\text{distinctiveness}} < 0$) is evidence

against that homogeneity reading: cities had heterogeneous typical preferences over content, and high-MA cities' producers responded to the demand-weighted European average rather than to their own local average. The Republic of Letters operated at the elite-scholar layer; the bulk of book production served city-level demand that was spatially differentiated.

Integration of the early-modern book trade (A4). A long historiographical debate concerns whether the European book trade was effectively integrated across language, religious, and political boundaries, or whether segmentation (Catholic vs. Protestant book markets, Latin vs. vernacular markets, polity-by-polity censorship) was the operative force. Our confirmation of Prediction 2 is evidence that the trade was integrated enough that high-MA producers responded to aggregate European demand—not just to their own segmented local market. The empirical signed prediction confirms, indirectly, that A4 holds at least on average.

Strategic differentiation among printers (A5). The historical literature has both supported A5 (descriptions of printer-publishers choosing portfolios, competing on title positioning, courting specific author networks) and qualified it (descriptions of guild restrictions, royal printing privileges, post-Tridentine censorship). Our confirmation of Prediction 3 ($\hat{\beta}_{\text{spec-entropy}} > 0$) is consistent with A5 holding in the aggregate across the 1,300 USTC cities, even if it fails in specific tightly-controlled contexts (e.g., post-1561 Madrid). The language-specific anomalies in Section 8 (German diversity null; French Paris dominance) plausibly reflect A5 failures in those specific contexts.

We are not building a structural model and we do not estimate σ, F, c, or θ structurally. We are saying that given the standard monopolistic-competition-with-spatial-differentiation framework, our empirical findings let us reverse-infer that the five identifiable assumptions about latent features of the 1450–1649 European book market each hold at least approximately. The joint statement that all five hold is a substantive historical claim that direct evidence does not easily settle.

C Within-FE vs. Mundlak / Bell-Jones — the singleton-dropping diagnostic

Section 7.1 explains why we report the Mundlak / Bell-Jones specification as the headline rather than the within-FE estimator. The within-FE estimator drops fixed-effect singleton cities, and

the dropped cities are systematically smaller and more intermittently printing, producing sample-selection bias. This appendix first restates the Mundlak / Bell-Jones reparameterization of fixed effects and its equivalence to within-FE on a balanced panel, then documents the singleton-dropping mechanism on our panel.

C.1 Why Mundlak / Bell-Jones is a fixed-effects estimator

The mechanic. Decompose any covariate x_{it} into a cluster mean $\bar{x}_i$ and a within-cluster deviation $\widetilde{x}_{it} = x_{it} - \bar{x}_i$, and run

$$y_{it} \;=\; \beta_W \, \widetilde{x}_{it} \;+\; \beta_B \, \bar{x}_i \;+\; \delta_t \;+\; \alpha + \varepsilon_{it}$$

with δ_t a period FE. The two slopes have clean interpretations: β_W is the within-cluster slope (how much does y deviate from its cluster mean when x deviates from its cluster mean?) and β_B is the between-cluster slope (do clusters with higher x averages have higher y averages?).

Mundlak's theorem. On a balanced panel without singletons, $\hat{\beta}_W^{\text{Mundlak}}$ is numerically identical to the FE estimator $\hat{\beta}^{\text{FE}}$ from the model $y_{it} = \beta x_{it} + \alpha_i + \delta_t + \varepsilon_{it}$. This is a standard textbook result (Wooldridge, 2010, Ch. 10). The within-cluster identification of β_W is mechanically the same as the within-FE identification of β. There is no random-effects assumption being smuggled in.

Why the "random effects" label is misleading. The classical random-effects estimator (Swamy–Arora, Balestra–Nerlove) treats the cluster intercept α_i as drawn from a distribution and assumes it is uncorrelated with the regressors. This is the assumption that makes plain RE inconsistent when treatment is correlated with unobserved cluster heterogeneity—and it is the assumption the worry about "we shouldn't use RE because treatment is not randomized" is targeting. Mundlak (1978)'s contribution is to break this assumption: by explicitly including $\bar{x}_i$ as a regressor, any correlation between the cluster intercept and the time-invariant component of x is absorbed into β_B. After the augmentation, β_W is identified from within-cluster variation only and the random-effects inconsistency disappears. Mundlak / Bell-Jones is therefore not a random-effects estimator in the consistency-worry sense; it is a reparameterization of fixed effects that makes the within-cluster and between-cluster contributions to the data separately visible.

C.2 The singleton-dropping mechanism

When `fixest::feols(y ~ x | place + decade)` is run, observations from cities appearing in only one of the four 50-year periods are dropped as fixed-effect singletons. This is algebraically correct (a single-period city has no within-period variation in x, so it cannot inform the within-cluster slope). The result is that the within-FE estimator identifies β only from cities present in two or more periods.

Singleton rates in the six panel-credible languages (distinctiveness panel) are large (Table 17).

Table 17: Singleton-city rates by language in the 50-year panel.

Language	Cities	Singletons	% singleton
Latin	385	159	41.3%
German	181	76	42.0%
Italian	108	49	45.4%
French	144	68	47.2%
Dutch	66	29	43.9%
Spanish	77	34	44.2%

In every language, 41–47% of cities are dropped by within-FE. Pooled across all six languages, 142 of 506 unique cities (28%) are dropped as place-FE singletons in the standard within-FE specification. The dropped cities are not random. They are systematically smaller and more peripheral than the cities that survive. Continuously printing centers (Paris, Lyon, Frankfurt, Wittenberg, Venice, Rome, Antwerp) appear in all four periods; small intermittent cities appear in one. The within-FE estimator therefore identifies on a sample heavily weighted toward continuously printing centers, while the Mundlak estimator uses the full panel.

C.3 The pooled distinctiveness sign flip

The headline empirical comparison is the pooled distinctiveness regression. Four specifications using the same dependent variable, the same control variables, and the same panel are reported in Table 18.

Table 18: Pooled distinctiveness regression under four specifications. The Mundlak coefficient is sign-opposite to the within-FE coefficients because Mundlak retains the 142 singleton cities that within-FE drops.

Spec	Pooled $\hat{\beta}$	n obs	n cities used
Within-FE: `place + decade + language`	0.506^{**}	1,742	364
Within-FE: `place^language + decade + language`	0.453^{*}	1,461	(cells, not cities)
Within-FE: `place + decade` (no language FE)	0.414^{*}	1,742	364
Mundlak / BJ: `decade + language` (headline)	-0.709^{***}	1,884	506

The Mundlak coefficient is -0.709 on the full sample of 506 cities; the corresponding within-FE coefficient is 0.506 on the reduced sample of 364 cities. The 142 additional cities in the Mundlak sample are the singletons dropped by within-FE—small, peripheral, intermittently printing.

The three within-FE variants disagree among themselves but none of them match the Mundlak estimate. The gap between the within-FE variants and Mundlak (roughly 0.5 versus -0.7, a 1.2-unit swing) is driven by the singleton cities that Mundlak retains and within-FE drops. The fact that the within-city slope flips sign when these cities are reincluded is informative. It tells us the small, intermittently printing cities show a strongly negative MA-distinctiveness relationship that the large continuously printing centers do not.

C.4 Bell-Jones decomposition: within vs. between

The Bell-Jones decomposition separates the within-cluster and between-cluster slopes within the same regression (Table 19).

Table 19: Within-city vs. between-city Bell-Jones slopes for the pooled distinctiveness regression.

Coefficient	Estimate	SE	p
$\hat{\beta}_W$ on $\widetilde{\log \mathrm{MA}}_{i\ell,d}$ (within-city)	-0.709	0.133	< 0.001
$\hat{\beta}_B$ on $\overline{\log \mathrm{MA}}_{i\ell}$ (between-city)	-0.020	0.028	0.47

The within-city slope is significantly negative (cities with rising MA become less distinctive), while the between-city slope is not statistically distinguishable from zero. These are different sta-

tistical objects—different identification, different interpretation—and the within-city signal is concentrated entirely in the longitudinal within-city variation rather than in cross-sectional differences in city-level MA averages. The within-FE pooled coefficient picks up the within-city slope only on the multi-period subsample of cities, biasing the estimate toward continuously printing centers whose within-city MA variation is small.

C.5 Leave-one-language-out robustness for the Mundlak estimate

To confirm the Mundlak distinctiveness coefficient is not driven by any single language (Table 20):

Table 20: Leave-one-language-out robustness for the pooled Mundlak distinctiveness coefficient.

Sample	Pooled $\hat{\beta}$	SE	p	n
All 6 languages	-0.709^{***}	0.133	< 0.001	1,884
Drop Latin	-1.254^{***}	0.195	< 0.001	1,088
Drop German	-0.648^{***}	0.143	< 0.001	1,527
Drop Italian	-0.649^{***}	0.150	< 0.001	1,675
Drop French	-0.525^{***}	0.099	< 0.001	1,626
Drop Dutch	-0.674^{***}	0.145	< 0.001	1,762
Drop Spanish	-0.700^{***}	0.138	< 0.001	1,742

The pooled Mundlak coefficient remains negative and significant at the 1% level in every leave-one-out. The magnitude strengthens when Latin is dropped (Latin contributes the weakest per-language distinctiveness coefficient at -0.058, dragging the pool toward zero) and attenuates most when French is dropped (French contributes the strongest negative per-language coefficient at -1.67). No single language drives the headline.

D Composition channel decompositions

This appendix contains the decompositions that identify which slice of Section 8's composition pattern is the growth-relevant slice Section 9 documents. The Section 8 headline—MA pulls a city's print toward broader, more European-mainstream, more topically even content—averages over

many composition dimensions; the Section 9 headline—total diversity and the useful-knowledge share of titles predict subsequent city growth above and beyond MA—is concentrated in two of them. The headline substantive findings are share-of-useful (Section D.3) and Latin within-useful diversity (Section D.2); the other decompositions are null on the growth side and provide convergent evidence for the same Mokyr useful-knowledge channel rather than for any of the alternative readings we test.

All five exercises use the Barro framework of Section 9.1: long-difference DV ($\Delta_{50}\log(\text{pop})$), period and language fixed effects (no city FE), initial $\log(\text{pop})$ and $\log(\text{MA})$ as controls, weighted by $\sqrt{n_{\text{titles}}}$, SEs clustered at the city.

D.1 USTC classification share decomposition

Method. For each (city, period, language) cell, we computed the share of titles in each USTC top-level classification. We restrict to classifications with $\geq 5{,}000$ titles across the six panel-credible languages (19 classifications in the MA decomposition, 7 with enough non-zero shares to enter the growth regression cleanly). For each classification c, we ran (i) a within-city BJ regression of π_c on $\log(\text{MA})$ and (ii) a Barro growth regression of $\Delta_{50}\log(\text{pop})$ on π_c.

Result. The MA decomposition shows that within-city, as MA rises, cities print proportionally more *Periodicals*, *Jurisprudence*, *Medical Texts*, and *Classical Authors*, and proportionally fewer *Religious*, *Funeral Orations*, *Wedding Pamphlets*, and *Poetry*. Read through the lens of the economics of religion, this pattern implies that market integration is associated with the secularization of a city's print mix, with output shifting at the margin from devotional and occasional content toward useful-knowledge content. The growth decomposition is largely null. Of the seven classifications with enough data to enter both regressions, only Periodicals shows up positive on both axes ($\hat{\beta}_{\text{MA}} = 0.27$, $\hat{\beta}_{\text{growth}} = 0.22$, both $p < 0.01$). Religious is negative on both axes but the negative growth coefficient is plausibly mechanical (smaller, peripheral cities have higher Religious share as a residual category and grew slower).

Interpretation. The Periodicals signal is suspect because (a) the average share is only 0.88% of titles, (b) periodicals as a publishing format emerged in the early 1600s, so most of the identifying variation comes from Amsterdam, Antwerp, Frankfurt, and Cologne in the $1600 \rightarrow 1650$ growth

window, and (c) those cities' growth in this window was driven by the Dutch Golden Age and the wider commercial expansion that also gave them periodical publishing. The coefficient plausibly reflects reverse causation rather than a causal channel.

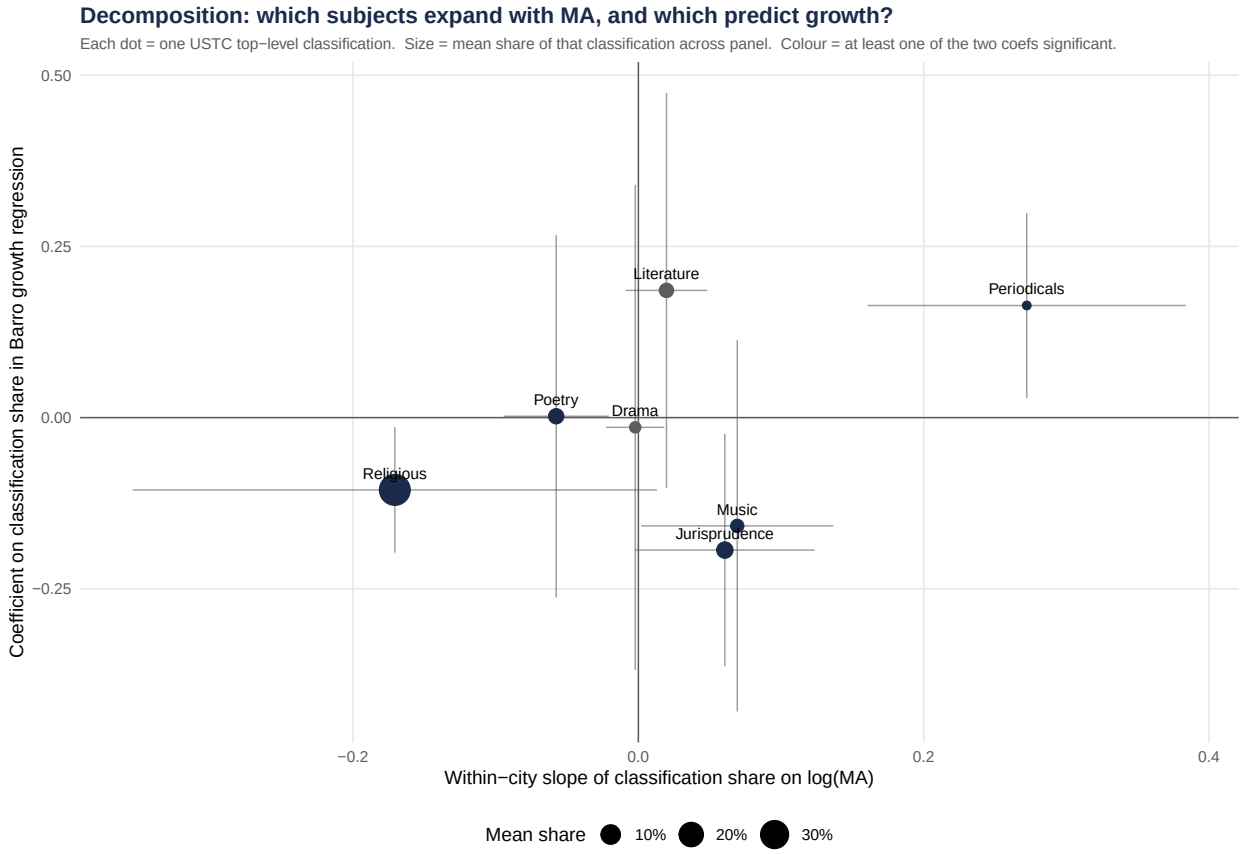


Figure 18: USTC classification share decomposition: within-city MA coefficient and Barro growth coefficient for each of seven classifications.

D.2 Useful-knowledge vs. ceremonial diversity decomposition

Method. We partition USTC classifications into two groups motivated by Mokyr's "useful knowledge" framework:

- *Useful*: Classical Authors, Medical Texts, Jurisprudence, Philosophy and Morality, Academic Dissertations, Periodicals, Educational Books, History and Chronicles, Political Tracts.
- *Ceremonial*: Religious, Funeral orations, Wedding pamphlets, Poetry, Drama.

For each (city, period, language) cell, we recompute within-city diversity (mean cosine distance to

centroid) restricted to titles in each group, yielding $D_{i,t}^{\text{useful}}$ and $D_{i,t}^{\text{ceremonial}}$ as separate metrics. The Barro growth regression then includes both diversity flavours simultaneously.

Result. Table 21.

Table 21: Useful-knowledge vs. ceremonial diversity decomposition: pooled and per-language Barro growth coefficients.

Sample	$\hat{\beta}_{\text{useful}}$ (p)	$\hat{\beta}_{\text{ceremonial}}$ (p)	n
Latin	0.439*** (0.001)	0.031 (0.88)	548
German	−0.236 (0.092)	0.404 (0.11)	202
Italian	−0.279 (0.29)	0.212 (0.35)	99
French	−0.338 (0.057)	−0.112 (0.71)	139
Dutch	−0.393 (0.39)	2.486 (0.065)	66
Spanish	0.314 (0.67)	0.862** (0.046)	82
POOLED	−0.010 (0.92)	0.282 (0.057)	1,136

The pooled useful-diversity coefficient is statistically null ($\hat{\beta} = -0.010$, $p = 0.92$); the pooled ceremonial coefficient is positive and marginally significant ($\hat{\beta} = 0.282$, $p = 0.057$). The per-language pattern, however, is sharply heterogeneous and informative. **Latin**—the cross-border language of European scholarship and the natural home of "useful-knowledge" content in this period—shows a large, highly significant positive useful-diversity coefficient ($\hat{\beta} = 0.439$, $p = 0.001$). German, French, and Italian useful coefficients are negative (insignificant except for marginal French at $p = 0.057$). The signs flip across languages and the pooled estimate averages them to zero.

Interpretation. The within-useful diversity channel supports the Mokyr-style hypothesis in the language that most plausibly carried trans-European useful knowledge—Latin—while showing the opposite sign in vernacular markets. This is consistent with a reading in which transnational scholarly print (Latin) is the medium of Mokyr's "Industrial Enlightenment," whereas vernacular useful-knowledge print sits closer to local consumer demand and is not the same channel. The pooled coefficient masks the heterogeneity.

Four of the five Latin exemplar cities named in §9 (Bologna, Padua, Leiden, Pavia) are in the USTC Latin panel and are also dense nodes of the de la Croix et al. (2025) scholar-affiliation network, which raises the concern that the Latin coefficient operates through university-city human-

capital channels rather than through Latin print specifically. The fifth, Salamanca, is a university printing centre whose Latin output falls below the USTC panel threshold, so it does not enter this check. Table 22 re-estimates the Latin within-useful diversity coefficient dropping each of the four in turn. The coefficient ranges from 0.436 to 0.442 across the four leave-one-out specifications, statistically indistinguishable from the baseline 0.439 and significant at the 1% level in every case.

Table 22: Leave-one-university-city-out robustness for the Latin within-useful diversity Barro coefficient. Each row re-estimates the §9.4 specification on Latin cells with one named university city excluded. Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Sample	$\hat{\beta}_{\text{within-useful}}$	SE	n_{obs}	Cities
Baseline (all Latin cities)	0.439***	(0.134)	548	270
Drop Bologna	0.437***	(0.134)	544	269
Drop Padova	0.442***	(0.135)	544	269
Drop Leiden	0.437***	(0.134)	545	269
Drop Pavia	0.436***	(0.135)	544	269

Notes: Salamanca, the fifth exemplar city named in §9.4, is not in the USTC Latin panel at the 5-title-per-cell threshold and so cannot be dropped.

A complementary test of the Mokyr channel that does not condition on within-subset diversity—the extensive-margin share of useful content—is reported in Section D.3 below; that test delivers a positive and robust pooled coefficient.

D.3 Share of useful-knowledge titles (extensive margin)

Method. The decomposition in Section D.2 measures the within-useful intensive margin: holding the partition fixed, how does diversity inside the useful subset relate to growth? A complementary Mokyr-style hypothesis operates at the extensive margin: cities whose share of printed titles falls in useful-knowledge classifications grow faster, regardless of how internally varied the useful titles are. We construct, for each (city, period, language) cell,

$$s^{\text{useful}}_{i,t,\ell} = \frac{n^{\text{useful}}_{i,t,\ell}}{n^{\text{useful}}_{i,t,\ell} + n^{\text{ceremonial}}_{i,t,\ell}},$$

the share of useful titles among useful-plus-ceremonial titles. We restrict to cells with at least 10 useful-plus-ceremonial titles to keep the share well-measured. The Barro regression is identical to Section D.2 otherwise (Spec A: log(pop) and log(MA) controls; period and language fixed effects; weighted by $\sqrt{n_{\text{useful+ceremonial}}}$; SEs clustered at city).

Result. The headline pooled and per-language coefficients are reported as Table 6 in Section 9, where the share-of-useful is one of the four headline composition channels for the growth regression. To summarize, the pooled coefficient is positive and significant at the 5% level ($\hat{\beta} = 0.123$, $p = 0.012$, $n = 1{,}486$); two per-language coefficients are individually significant (German 0.247^{***}, Italian 0.218^{**}); the others share the positive sign. Moving from no useful titles to all useful titles within the useful-ceremonial partition is associated with roughly 12 log-points of additional growth over the 50-year window, conditional on log(MA) and log(pop)—the same scale as Section 9.3's diversity coefficient.

Robustness. Eight robustness checks on the pooled coefficient are reported in Table 23. The coefficient survives dropping the German 1600 panel cell (Thirty Years' War onset), dropping all German-territory Latin 1600 cells in addition, dropping the 1450 window, dropping Classical Authors from the useful set, dropping cities east of 25E (and 20E), and switching from Specification A (Barro with log(pop) convergence control) to Specification B (long difference without log(pop)).

Table 23: Share-of-useful pooled coefficient under robustness checks.

Sample / specification	$\hat{\beta}$	SE	p	n
A: Barro (headline)	0.123^{**}	0.049	0.012	1,486
B: long-diff (no convergence)	0.122^{**}	0.049	0.013	1,486
Drop German 1600	0.112	0.060	0.063	1,355
Drop German + German-terr Latin 1600	0.121	0.066	0.069	1,273
Drop 1450 window	0.135^{***}	0.051	0.008	1,356
Drop Classical Authors from useful	0.116^{**}	0.049	0.019	1,478
Drop cities east of 25E	0.127^{***}	0.049	0.009	1,477
Drop cities east of 20E	0.130^{***}	0.049	0.009	1,452

Interpretation. Cities whose print output skewed toward useful-knowledge classifications grew systematically faster than cities of the same size and market access whose print output skewed to-

ward ceremonial classifications. This is the cleanest piece of evidence in the paper for a Mokyr-style "useful-knowledge → growth" channel: it is an extensive-margin test that does not condition on within-subset diversity, survives the standard robustness perturbations, and delivers the predicted positive sign. Combined with the Latin within-useful diversity coefficient in Section D.2 (0.439^{***}), the two tests provide converging evidence that the composition of useful versus ceremonial content carries growth-relevant information beyond market access and initial size.[16]

D.4 Latin share as a transnational-orientation proxy

Method. Latin print was by construction transnational in the period—a Latin title published in Lyon could be sold in Wittenberg without translation. Vernacular print was largely local. We compute, for each (city, period) over the full USTC corpus, the share of titles in Latin: share_Latin$_{i,t}$. This is then included in the Barro growth regression alongside total diversity.

Result. Table 24.

Table 24: Latin share and diversity in the Barro growth regression.

Sample	$\hat{\beta}_{\text{Latin-share}}$ (p)	$\hat{\beta}_{\text{diversity}}$ (p)	n
Latin	-0.104 (0.18)	0.560^{***} (< 0.01)	794
German	-0.047 (0.56)	-0.097 (0.53)	357
Italian	0.046 (0.62)	-0.251^{**} (0.034)	209
French	-0.073 (0.53)	-0.111 (0.55)	258
Dutch	-0.131 (0.58)	0.610^{**} (0.021)	122
Spanish	-0.407^{**} (0.025)	0.060 (0.89)	142
POOLED	-0.076 (0.29)	0.198^{***} (< 0.01)	1,957

The pooled Latin-share coefficient is statistically null ($\hat{\beta} = -0.076$, $p = 0.29$). The pooled diversity coefficient in this specification is positive and significant ($\hat{\beta} = 0.198$, $p < 0.01$), reflecting the same diversity signal documented in Section 9.3. Only Spanish shows a per-language

[16] One natural additional covariate is $\log(n_{\text{titles}})$, total print volume. Adding it as a control attenuates the share-of-useful coefficient by about half ($\hat{\beta} = 0.065$, $p = 0.19$). We do not report that specification in the main panel because total print volume is itself downstream of the same market-access and city-prosperity channels that drive growth (e.g. Dittmar, 2011), making it a bad control (a mediator). If a reader treats print volume as exogenous to the useful/ceremonial share, the controlled coefficient is the relevant test; we do not.

Latin-share coefficient that is significant (-0.407, $p = 0.025$), and the Spanish panel is among the smallest.

Interpretation. The simple "Republic of Letters → growth" proxy of Latin's share in a city's total print output is not a growth predictor at the city-period level for our six-language panel, once log(MA), log(pop), and total diversity are controlled. The Mokyr-style channel that does survive operates through the composition of useful versus ceremonial content within a city's print (Section D.3), and through within-useful diversity in the Latin subset itself (Section D.2). The Latin share—a much cruder proxy that mixes "useful" Latin with classical, devotional, and ceremonial Latin—does not carry the same signal.

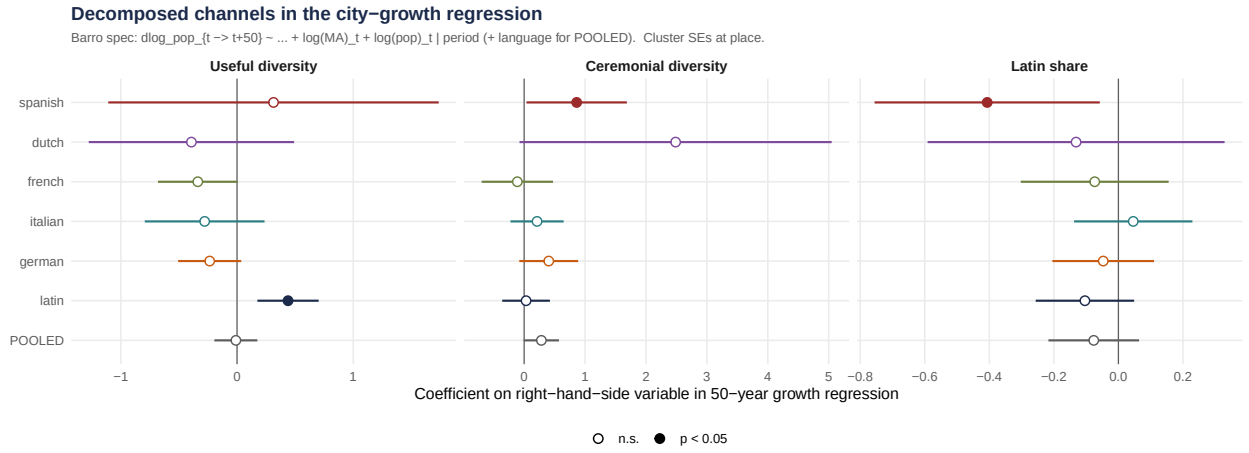


Figure 19: Useful vs. ceremonial diversity decomposition (left two panels) and Latin-share decomposition (right panel).

D.5 Direction-of-distinctiveness in embedding space

Method. Section D.1 decomposed cities by their count shares across USTC classifications. Here we decompose by direction in 200-dimensional embedding space—a continuous version that handles unclassified titles, captures within-classification heterogeneity, and avoids the small-city-residual concern of count shares.

For each language ℓ, we compute three sets of vectors in $\mathbb{R}^{200}$:

- *European centroid* $\mathbf{g}_t^\ell$—mean title embedding across all titles in language ℓ at period t.
- *Subject centroid* $\mathbf{C}_c^\ell$—mean title embedding across titles in language ℓ classified into USTC

top-level subject c, pooled across periods. We retain subjects with ≥ 200 titles in the language.

- *City centroid* $\mathbf{c}_{i,t}^{\ell}$—mean title embedding for city i at period t (the same centroid used in our distinctiveness metric).

For each (city, period, language, subject) cell we compute the projection of the city's deviation from the European centroid onto the subject's deviation from the European centroid:

$$\text{proj}_c(i,t,\ell) \;=\; \frac{(\mathbf{c}_{i,t}^{\ell} - \mathbf{g}_t^{\ell}) \cdot (\mathbf{C}_c^{\ell} - \mathbf{g}_t^{\ell})}{\|\mathbf{C}_c^{\ell} - \mathbf{g}_t^{\ell}\|}.$$

This is a scalar measuring *how much of a city's distinctiveness from the European centroid points in the direction of subject* c. Positive values mean the city's centroid leans toward subject c; negative means away from it.

For each subject c, we run the Barro growth regression with proj_c on the right-hand side, separately by language and pooled.

Result. Pooled coefficients across 12 subject directions (n between 737 and 1,882, depending on subject coverage by language) are reported in Table 25.

Table 25: Pooled growth coefficient on the projection of city-centroid deviation onto each subject direction, 12 subjects.

Subject direction	Pooled $\hat{\beta}$	p
Political Tracts	0.106	0.12
Periodicals	0.096	0.14
Music	−0.047	0.44
History and Chronicles	0.041	0.47
Religious	−0.024	0.49
Academic Dissertations	0.032	0.49
Classical Authors	0.024	0.61
Jurisprudence	−0.018	0.66
Poetry	0.008	0.80
Literature	−0.004	0.88
Medical Texts	−0.007	0.91
Philosophy and Morality	−0.001	0.99

None of the twelve pooled subject directions reaches conventional significance. Political Tracts and Periodicals—the two with the largest positive point estimates—are marginal ($p = 0.12$ and $p = 0.14$), with signs consistent with a commercial-information reading. Religious is no longer marginally significant once the panel is rebuilt on the v2 cost matrices.

Per-language standouts. Several individual language × direction coefficients are significant. In German, both Political Tracts (0.190, $p = 0.018$) and Periodicals (0.140, $p = 0.019$) directions are positive—German cities whose centroid leans toward those two "useful" content directions grew systematically more, consistent with the share-of-useful finding for German in Section D.3. In Spanish, Religious (-0.130, $p = 0.008$) and Poetry (-0.126, $p = 0.017$) directions are negative—Spanish cities whose centroid leaned toward ceremonial directions grew less. In Dutch, several directions (Political Tracts, Poetry, Literature, History) carry large positive coefficients ($\hat{\beta}$ between 0.71 and 1.10, all $p < 0.005$), driven by a small panel ($n = 122$) and dominated by the seventeenth-century expansion of Amsterdam.

Interpretation. The embedding-space decomposition adds a robustness check that the count-

based findings in Sections D.1 and D.3 are not simply artifacts of how USTC classifies titles. The pattern that survives—positive German coefficients on Political Tracts and Periodicals; the absence of any pooled significant negative direction—is consistent with the share-of-useful headline in Section D.3. The two methodologies converge on the same substantive conclusion: cities whose print leans toward useful-knowledge content (whether measured as a count share or as an embedding direction) grew systematically faster.

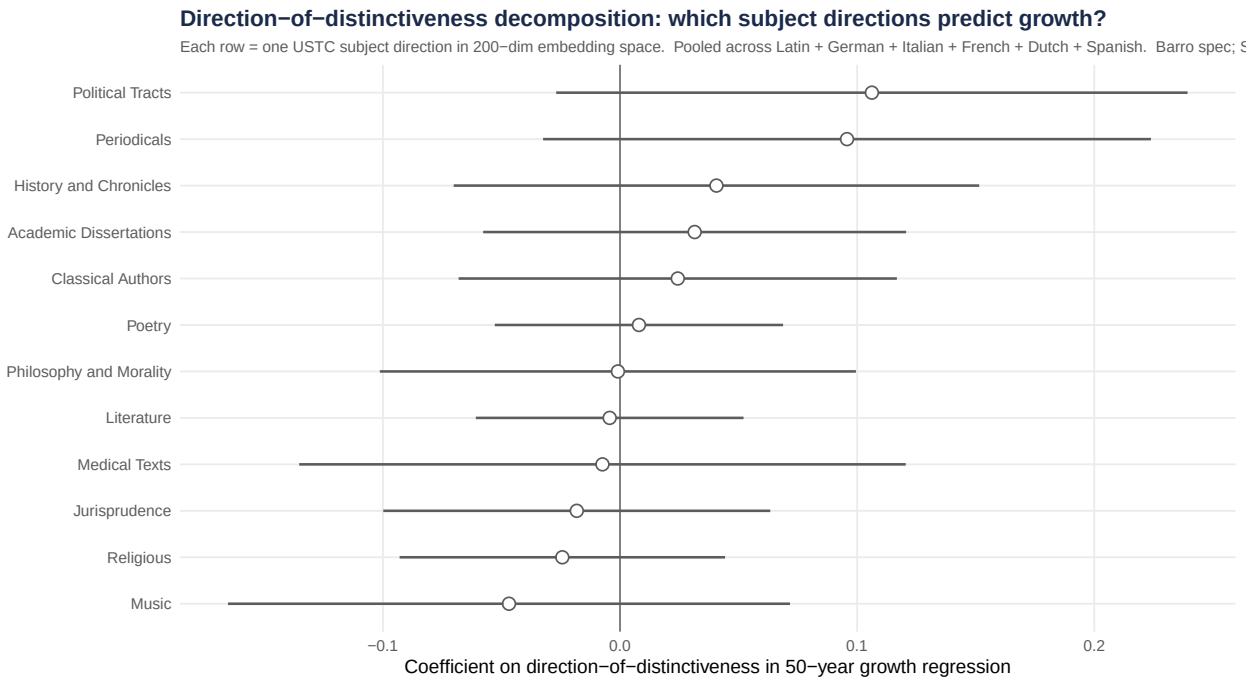


Figure 20: Forest plot of pooled coefficients across the 12 subject directions in the embedding-space decomposition.

E Thirty Years' War sample-restriction sensitivity

The Thirty Years' War (1618–1648) overlaps our 1600 period (titles 1600–1649) and devastated central Europe. This appendix documents the war's footprint in our panel and reports the sample-restriction sensitivity for both the headline composition findings (Section 8) and the city-growth regressions (Section 9).

E.1 Diagnostic: war footprint in the data

German metric reversal. The 1600 period is the only one in which German diversity falls (0.723 $\rightarrow$ 0.592) and German entropy falls (1.233 $\rightarrow$ 1.044). Latin, Italian, and French diversity / entropy are smooth across 1550 $\rightarrow$ 1600. The German reversal coincides exactly with the war years (1618–1648) which fall inside our 1600-period title window.

City-composition fingerprint. Comparing top-10 German-printing cities pre- and post-1600 (Table 26):

Table 26: Top-10 German-printing cities, pre-war (1550–1599) vs. war-era (1600–1649).

Pre-war (1550–1599)	Titles	War-era (1600–1649)	Titles
Nürnberg	3,128	Leipzig (5.2$\times$ pre-war)	8,916
Frankfurt	2,324	Hamburg (new entrant)	4,640
Leipzig	1,701	Frankfurt	4,213
Wittenberg	1,585	Nürnberg	2,409
Strasbourg	1,542	Strasbourg	2,104
Magdeburg	1,331	Erfurt	2,077
Köln	1,116	Köln	2,032
Augsburg	1,115	Munich	1,924
Erfurt	994	Wittenberg	1,724
Tübingen	777	Wolfenbüttel	1,717

Three war-driven shifts are immediate. *Magdeburg disappears* from the top-15 after 1600. Tilly's 1631 sack killed an estimated 20,000 of the city's roughly 25,000–30,000 inhabitants and burned the city; the press collapsed with the population. *Leipzig and Hamburg surge.* Leipzig (under Saxon neutrality) and Hamburg (free Hanseatic port) absorbed refugees, printers, and capital from war-affected regions. *Wolfenbüttel emerges* as Duke Augustus the Younger's library and academic center expanded under the protection of the Welf court.

Latin in the war years. Latin print volume dips notably during the war (2,111 Latin titles in 1632, down from 3,758 in 1620), and recovers to 3,427 by 1648. This is consistent with collateral damage to Wittenberg, Heidelberg, Tübingen, Strasbourg, and the German Jesuit colleges that produced

much of the Latin scholarly output.

E.2 Sensitivity samples

We construct three samples and re-run both the Section 8 headline regression and the Section 9 Barro growth regression on each:

- *(i) Full panel*—1,884 city × period × metric obs.
- *(ii) Drop the German 1600 period entirely*—excludes 155 obs (Germany at $t = 1600$).
- *(iii) Drop German 1600 + Latin 1600 for German-territory cities*—excludes 287 obs, where German-territory cities are USTC-coded Germany, Austria, Switzerland, Czechia, Slovakia, Hungary, or Poland (the Holy Roman Empire and immediate borderlands).

E.3 Section 8 sensitivity — composition findings strengthen

Pooled Mundlak/BJ coefficients on $\log(\text{MA})$, all three metrics (Table 27).

Table 27: Section 8 pooled BJ coefficients across the three sensitivity samples. Standard errors clustered at the city in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Metric	(i) Full	(ii) Drop Ger 1600	(iii) Drop Ger + GerTerr Lat 1600
Diversity	0.467^{***}	0.557^{***}	0.559^{***}
	(0.064)	(0.063)	(0.067)
Distinctiveness	-0.709^{***}	-0.652^{***}	-0.835^{***}
	(0.133)	(0.139)	(0.141)
Specialization (entropy)	0.751^{***}	0.895^{***}	1.431^{***}
	(0.223)	(0.237)	(0.212)
Observations	1,884	1,729	1,597

All three coefficients remain significant at the 1% level in every sample. Magnitudes strengthen once war-affected observations are removed—particularly entropy, which nearly doubles between

the full sample and the (iii) restriction. The war was attenuating the pooled composition story, not driving it.

E.4 Section 9 sensitivity — diversity signal strengthens

Pooled Barro growth coefficients (long-difference $\Delta_{50} \log(\text{pop})$ regressed on contemporaneous metric, controlling for $\log(\text{MA})$, $\log(\text{pop})$, period FE, language FE) are reported in Table 28.

Table 28: Section 9 pooled Barro growth coefficients on each composition metric across the three sensitivity samples. Pooled coefficients with city-clustered standard errors in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Sample	Diversity	Distinct.	Entropy
(i) Full sample	0.143^{*}	-0.058	-0.036
	(0.082)	(0.065)	(0.031)
	$n = 1{,}882$	$n = 1{,}882$	$n = 1{,}882$
(ii) Drop German 1600 period	0.208^{**}	-0.074	-0.039
	(0.085)	(0.066)	(0.032)
	$n = 1{,}727$	$n = 1{,}727$	$n = 1{,}727$
(iii) Drop German + German-territory Latin 1600	0.183^{**}	-0.060	-0.051
	(0.085)	(0.065)	(0.032)
	$n = 1{,}595$	$n = 1{,}595$	$n = 1{,}595$

The diversity coefficient on subsequent city growth, which is marginally positive in the full sample ($\hat{\beta}_{\text{diversity}} = 0.143^{*}$), strengthens to significant at the 5% level once war-affected observations are removed (0.208^{**} under sample (ii), 0.183^{**} under sample (iii)). The distinctiveness and entropy growth coefficients remain indistinguishable from zero in every sample, as in the headline specification. The Thirty Years' War was attenuating—not producing—the diversity-on-growth signal documented in Section 9.

E.5 Interpretation

The Thirty Years' War is a real feature of the 1600 period's data, with a clear footprint in German-language print and a smaller footprint in Latin print in German territories. Excluding war-affected observations:

- Strengthens the composition findings of Section 8 (diversity ↑, distinctiveness ↓, entropy ↑ with MA, all significant in every sample at the 1% level; entropy nearly doubles in magnitude).
- Strengthens the diversity-on-growth signal of Section 9: the diversity coefficient is marginally positive in the full sample (0.143^{*}) and becomes significant at the 5% level under both war-restricted samples (0.208^{**} and 0.183^{**}). The distinctiveness and entropy growth coefficients remain null in every sample.

We therefore retain the full sample as the primary specification for both sections and report this sensitivity here. A reviewer concerned that "your 1600 period is contaminated by war" can read this appendix and confirm that the contamination, if anything, works against our findings in both Section 8 and Section 9.

F Additional robustness checks

This appendix collects six robustness checks that target specific concerns about the identification of the headline coefficients in Sections 8 and 9. Each subsection addresses a discrete concern, reports the diagnostic, and compares the resulting coefficient to the headline. The qualitative pattern of the headline findings is preserved in every check.

F.1 Lagged market access

The Section 9 Barro specification uses contemporaneous $\log(\text{MA})_{i,t}$ as a control. Because $\log(\text{MA})_{i,t}$ is a population-weighted sum over neighbors at the same anchor year, common regional shocks that drive both neighbor populations at t and the focal city's growth from t to $t + 50$ could in principle contaminate the regressor. To address this we rebuild $\log(\text{MA})$ from neighbor populations at the

previous half-century anchor. For the 1500–1550 growth window we use neighbor pops at 1450 to weight MA. For 1550–1600 we use 1500 pops. For 1600–1650 we use 1550 pops. The 1450–1500 growth window drops from the sample because no native 1400 population anchor is available, so the lagged-MA panel runs to 878 city-period cells for the diversity / distinctiveness / entropy specifications and 800 cells for share-of-useful, against 1,882 and 1,486 in the headline. Table 29 reports the coefficients side by side.

Table 29: Lagged-MA robustness for the four headline Barro growth coefficients. Column (1) reproduces the §9 headline with contemporaneous $\log(\text{MA})_{i,t}$. Column (2) replaces it with $\log(\text{MA})_{i,t-50}$, built from neighbour populations at $t-50$ rather than t. The 1450–1500 growth window is dropped because no 1400 native anchor is available. Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$ (or $\sqrt{n_{\text{useful+ceremonial}}}$ for share-of-useful). Period and language FE included. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Metric	(1) Headline	(2) Lagged MA	$n_{(1)}$	$n_{(2)}$
Diversity	0.143*	0.216	1882	878
	(0.082)	(0.135)		
Distinctiveness	-0.058	-0.068	1882	878
	(0.065)	(0.081)		
Entropy	-0.036	-0.039	1882	878
	(0.031)	(0.042)		
Share-of-useful	0.123**	0.160**	1486	800
	(0.049)	(0.063)		

The diversity point estimate rises from 0.143 to 0.216 but loses significance with the reduced sample. The distinctiveness and entropy nulls hold. The share-of-useful coefficient strengthens from 0.123 to 0.160 and remains significant at the 5% level. Conditioning on a predetermined measure of trade exposure does not weaken any of the headline qualitative findings, and arguably strengthens the Mokyr-style useful-knowledge channel.

F.2 University-presence robustness for the Latin within-useful diversity coefficient

The Section 9 discussion of the Latin within-useful diversity coefficient names Bologna, Padua, Leiden, Pavia, and Salamanca as exemplar cities. These are also dense nodes of the European scholar-affiliation network reconstructed by de la Croix et al. (2025). The natural concern is that the Latin coefficient operates through university-city human-capital channels (cf. de la Croix et al. (2018)) rather than through Latin print specifically.

To test this directly we construct a (city, period) university-active indicator from a hand-curated list of ∼60 European university foundations between 1088 and 1650, mapped to USTC place names. 57 of the 517 panel cities are flagged. We then re-estimate the Latin within-useful diversity coefficient under four specifications that condition on university presence in different ways (Table 30).

Table 30: University-presence robustness for the Latin within-useful diversity Barro coefficient. Each row reports the coefficient on within-useful diversity in a Latin-only Barro growth regression (analogous to Table 21 Latin row) under a different treatment of university presence. The university-active indicator is from a hand-curated list of European university foundations 1088–1650; 57 of 517 panel cities are flagged. Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Specification	$\hat{\beta}_{\text{within-useful}}$	SE	n_{obs}	Cities
Baseline	0.439***	(0.134)	548	270
(i) + univ dummy	0.438***	(0.134)	548	270
(i') + univ x diversity interaction	0.396**	(0.166)	548	270
(ii) Non-university cities only	0.448***	(0.161)	419	236
(iii) University cities only	0.371	(0.257)	129	50

The university-active dummy itself is small and insignificant ($\hat{\beta} = 0.019$, $p = 0.63$). The within-useful × university-active interaction is also insignificant ($\hat{\beta} = 0.106$, $p = 0.77$). The cleanest sufficiency test is the non-university-cities-only specification, which gives $\hat{\beta}_{\text{within-useful}} = 0.448$ ($p = 0.006$) on the 236-city sub-panel. This is statistically indistinguishable from the full-sample baseline of 0.439. The Latin within-useful diversity signal is therefore not identified off the university exemplar cities. It is a feature of the broader Latin scholarly panel. The university-cities-only specification gives a slightly smaller and insignificant point estimate (0.371, $p = 0.16$),

reflecting the small ($n = 50$) sub-panel rather than evidence of a distinct university channel.

The university panel is approximate. It is curated from standard reference lists (Rashdall (1895), supplemented with sixteenth- and seventeenth-century foundations) rather than from a single authoritative source like Verger (2007), and short-lived or minor institutions are likely missed. The exemplar cities named in Section 9 are correctly flagged. The conclusion that the Latin coefficient survives on non-university cities is robust to plausible additions and deletions from the list.

F.3 Leave-one-classification-out for the share-of-useful coefficient

The Section 9 share-of-useful coefficient is identified from a partition of USTC `classification1` categories into a useful subset of 9 categories and a ceremonial subset of 5 categories. The Classical Authors inclusion in the useful side is acknowledged in the main text as the most contestable. On the ceremonial side the Religious category mixes vernacular Bibles, Calvinist catechisms, Lutheran scholastic disputations, and Greek-Hebrew biblical commentaries that could plausibly be coded differently. To test sensitivity we drop each of the 14 classifications in turn and re-estimate the pooled share-of-useful coefficient (Table 31).

Table 31: Leave-one-classification-out robustness for the pooled share-of-useful Barro coefficient. Each row reports the pooled coefficient on share-of-useful after dropping one USTC `classification1` from either the useful or the ceremonial side of the partition. Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{useful+ceremonial}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Sample	$\hat{\beta}_{\text{share-useful}}$	SE	n_{obs}
Baseline (full partition)	0.123**	(0.049)	1486
Drop one useful class:			
Drop Classical Authors	0.116**	(0.049)	1480
Drop Medical Texts	0.130***	(0.048)	1469
Drop Jurisprudence	0.153***	(0.052)	1444
Drop Philosophy and Morality	0.119**	(0.048)	1480
Drop Academic Dissertations	0.104**	(0.049)	1477
Drop Periodicals	0.093	(0.062)	1484
Drop Educational Books	0.119**	(0.049)	1457
Drop History and Chronicles	0.109**	(0.046)	1464
Drop Political Tracts	0.121**	(0.048)	1459
Drop one ceremonial class:			
Drop Religious	0.078	(0.053)	1080
Drop Funeral orations	0.119**	(0.048)	1473
Drop Wedding pamphlets	0.123**	(0.049)	1484
Drop Poetry	0.121**	(0.047)	1460
Drop Drama	0.123**	(0.048)	1475

The coefficient stays in the $[0.09, 0.15]$ band for 12 of the 14 single-class drops, with significance at the 5% level preserved in 11. Two drops are sensitivity points. The Periodicals drop on the useful side moves the coefficient to 0.092 ($p = 0.13$). The sample size barely changes ($n = 1{,}484$ versus $1{,}486$), so the attenuation reflects the absorption of news and current-information output into the residual rather than a sample-composition shift. The Religious drop on the ceremonial side moves the coefficient to 0.078 ($p = 0.14$). Here the sample size falls sharply from $1{,}486$ to $1{,}080$ because dropping the largest classification on either side leaves many cells without enough

useful-plus-ceremonial titles to clear the 10-title threshold the share-of-useful sample requires. The Religious drop is therefore partly a power statement rather than evidence that the headline depends on a particular treatment of religious print.

F.4 Length-restricted sub-sample for the Section 8 metrics

The three composition metrics that drive Section 8 are constructed from title embeddings on USTC `short_title` fields. `short_title` length varies across cells (Table 9), and one concern is that mechanical vocabulary-coverage variation could drive part of the headline pattern. To test this we restrict the title sample to titles with at least 7 words after a whitespace-and-punctuation tokenization (about 80% of titles across the seven languages), recompute the diversity, distinctiveness, and entropy metrics on the restricted set per (city, language, period) cell, and re-run the pooled Mundlak/Bell-Jones specification. Table 32 reports the headline and length-restricted coefficients side by side.

Table 32: Length-restricted sub-sample robustness for the three §8 pooled within-city Mundlak/Bell-Jones coefficients on log(MA). Column (1) reproduces the headline; column (2) recomputes diversity, distinctiveness, and entropy from the sub-sample of titles with at least seven words after a whitespace–punctuation tokenisation of `short_title` (about 80% of titles, per Table 9). Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Metric	(1) Headline	(2) $n_{\text{words}} \geq 7$	$n_{(1)}$	$n_{(2)}$
diversity	0.467***	0.455***	1884	1810
	(0.064)	(0.052)		
distinctiveness	-0.709***	-0.629***	1884	1810
	(0.133)	(0.125)		
entropy	0.751***	1.049***	1884	1810
	(0.223)	(0.201)		

The diversity coefficient is essentially unchanged ($0.467 \rightarrow 0.455$, both significant at the 0.1% level), as is distinctiveness ($-0.709 \rightarrow -0.629$, both at the 0.1% level). The entropy coefficient strengthens from 0.751 to 1.049. All three signed predictions of Section 8 survive the restriction to longer titles.

F.5 Permutation diagnostic for the diversity coefficient

The length restriction in F.F.4 addresses one form of vocabulary-coverage variation but not the most direct one. Larger (city, language, period) cells have more titles and therefore more opportunities to sample a wider vocabulary, which mechanically inflates the within-cell mean cosine distance that defines diversity. A cleaner diagnostic holds the number of titles per cell fixed.

For each (city, language, period) cell with at least 30 titles we draw a random sub-sample of exactly 30 titles, recompute diversity on the sub-sample, and re-estimate the pooled Bell-Jones within-city slope. Repeating this 100 times with different random seeds yields the distribution of the downsampled slope shown in Figure 21.

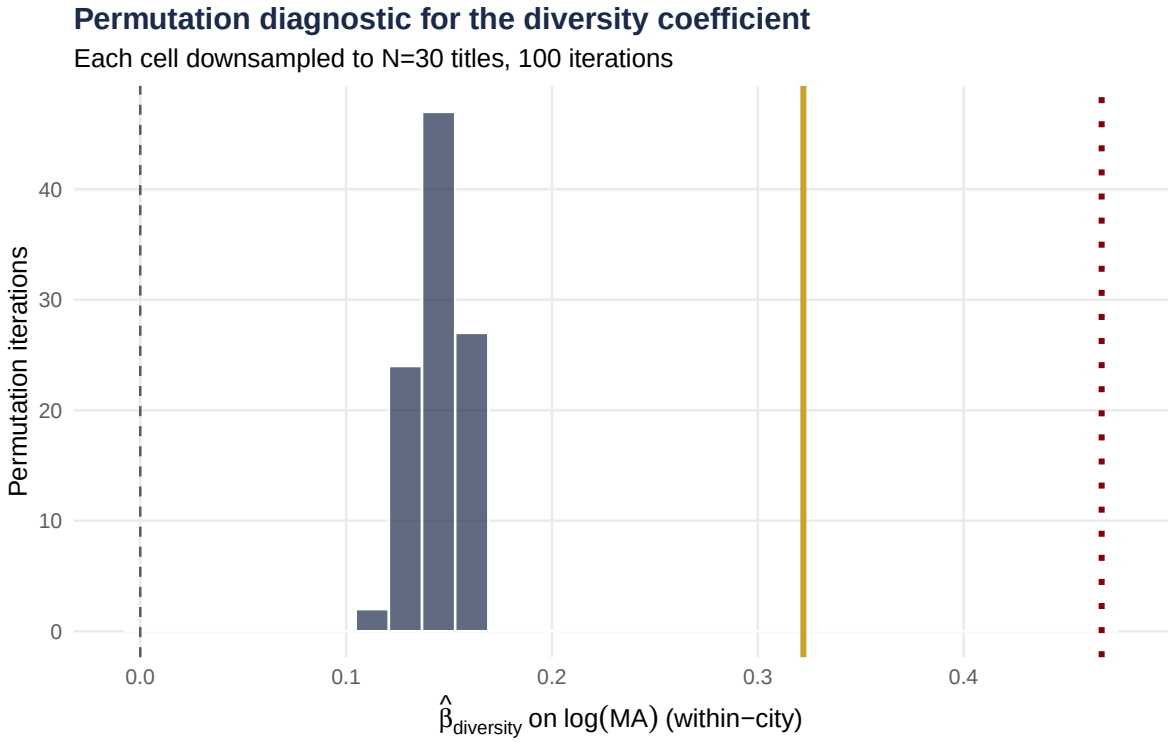


Figure 21: Distribution of the within-city Bell-Jones diversity slope on log(MA) across 100 permutation iterations, each downsampling every (city, language, period) cell to 30 titles before recomputing diversity. The gold line is the baseline on the restricted $N \geq 30$ panel with all titles in each cell used. The dotted red line is the full-panel headline of 0.467 reported in Section 8. The downsampled distribution is centered at 0.145 (SD 0.011, 95% interval $[0.124, 0.165]$). Every one of the 100 iterations is significant at the 5% level.

The headline slope of 0.467 separates into two components when the permutation diagnostic is applied. The portion that remains after fixing the per-cell title count at 30 is 0.145 on average. The

portion that is absorbed by removing the n_{titles} variation across cells is $0.467 - 0.145 = 0.322$. The qualitative finding that diversity rises with market access within-city is robust, since every one of the 100 permutation iterations delivers a significantly positive slope. The magnitude estimated under the permutation diagnostic is roughly one third the size of the headline. The headline estimate is best read as the composition response plus a mechanical vocabulary-coverage contribution, and the permutation estimate as the composition response alone. Both are positive, both are significantly different from zero, and both support the qualitative Section 8 conclusion.

F.6 Language-restricted market access

The market-access measure in Section 8 weights every neighboring city by its total population, regardless of the language its presses actually served. For the five vernaculars this raises a natural objection. The relevant market for French print is plausibly the set of cities that themselves print in French, not the full European population surface that also sustains Latin, German, and Italian output. To address this we rebuild a language-restricted market access for each vernacular. Let S_ℓ denote the set of cities printing at least N titles in language ℓ over 1450–1650. We zero out the population of every city outside S_ℓ, recompute $\log(\text{MA})$ on the restricted population surface, and re-estimate the Section 8 within-city Bell–Jones slopes on the restricted measure. Latin retains the unrestricted measure throughout, since the Latin market was genuinely pan-European. Table 33 reports the result.

Table 33: Language-restricted market access. Each vernacular's market is the set of cities printing at least N titles in that language (S_ℓ); population from cities outside S_ℓ is zeroed before recomputing log(MA). The $N \geq 1$ column is the headline (it keeps every city printing in the language); $N \geq 10$ and $N \geq 50$ are stricter stress tests. Panel A reports pooled within-city Mundlak/Bell-Jones slopes on restricted log(MA) for Spec C (the five vernaculars, each on its own restricted MA) and Spec D (Latin retained on unrestricted MA, the five vernaculars on restricted MA). Panel B reports the per-language within-city slopes at the headline $N \geq 1$. Standard errors clustered at the city in parentheses; observations weighted by $\sqrt{n_{\text{titles}}}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Panel A. Pooled within-city slope on restricted log(MA)

Metric	$N \geq 1$	$N \geq 10$	$N \geq 50$
Spec C: five vernaculars restricted ($n = 1,014$)			
diversity	0.462***	0.428***	0.421***
	(0.079)	(0.071)	(0.064)
distinctiveness	-0.937***	-0.828***	-0.754***
	(0.173)	(0.150)	(0.142)
entropy	0.623***	0.489**	0.429**
	(0.213)	(0.194)	(0.193)
Spec D: Latin unrestricted + five vernaculars restricted ($n = 1,810$)			
diversity	0.361***	0.343***	0.343***
	(0.048)	(0.044)	(0.042)
distinctiveness	-0.522***	-0.505***	-0.482***
	(0.106)	(0.097)	(0.095)
entropy	0.350*	0.291	0.269
	(0.195)	(0.185)	(0.182)

Panel B. Per-language within-city slope at $N \geq 1$

Language	diversity	distinctiveness	entropy
German	-0.064	-0.879***	0.229
Italian	0.940***	-1.218***	1.586***
French	0.784***	-1.305***	1.061
Dutch	0.409***	-0.710	0.989**
Spanish	0.443**	-1.050***	0.478

We lead with the $N \geq 1$ threshold, which keeps every city that printed even a single title in the language. This is the least arbitrary cutoff and it removes a specific concern with the absolute-count alternative. A count threshold is blind to within-city language share, so at $N \geq 10$ it discards a small town that printed almost entirely in the language alongside one that printed a few incidental titles in it. A composition diagnostic on the dropped cities confirms that the first type is real but immaterial for a population-weighted measure. The towns dropped at $N \geq 10$ that print overwhelmingly in the language, with an own-language output share above 0.8, account for under 0.2% of each language's titles and have a median population of around two to four thousand, the smallest cities in the panel. They cannot move a population-weighted neighbor's MA appreciably, so the $N \geq 10$ and $N \geq 50$ columns serve as stricter stress tests rather than as the preferred specification.

At $N \geq 1$ all three signed predictions of Section 8 survive in the substantively correct specification, Spec D, in which Latin keeps the unrestricted measure and the five vernaculars use the restricted one. The pooled within-city slopes are $\hat{\beta}_{\text{diversity}} = 0.361$ and $\hat{\beta}_{\text{distinctiveness}} = -0.522$, both significant at the 1% level, and $\hat{\beta}_{\text{spec-entropy}} = 0.350$, significant at the 10% level (Panel A). The per-language slopes at $N \geq 1$ (Panel B) track the headline pattern for the four high-volume vernaculars. Italian, French, Dutch, and Spanish all post positive diversity and negative distinctiveness slopes, and the German within-city diversity slope is essentially zero, exactly as in the Section 8 headline where German is the lone diversity exception.

The stricter thresholds attenuate the pooled slopes modestly and monotonically. Moving from $N \geq 1$ to $N \geq 50$ shrinks the Spec D diversity slope from 0.361 to 0.343 and the distinctiveness slope from -0.522 to -0.482, with significance unchanged. The one coefficient that crosses a significance threshold is the pooled specialization-entropy slope, which falls from 0.350 at $N \geq 1$ to 0.291 at $N \geq 10$ and 0.269 at $N \geq 50$ and is no longer significant at the stricter cutoffs. We read this attenuation as mild and do not interpret it as evidence that the discarded towns are decisive, for two reasons. First, the movement is small and never reverses a sign. Second, the attenuation is not universal. German runs the other way, with its distinctiveness slope strengthening from -0.879 to -1.120 and its specialization-entropy slope rising from 0.229 to 0.486 as the cutoff tightens, which is inconsistent with a story in which pruning low-volume towns uniformly removes genuine market signal. The honest summary is that the qualitative Section 8 findings hold under a language-restricted market, that the pooled diversity and distinctiveness results are robust

across every threshold, and that only the pooled specialization-entropy result is sensitive to how aggressively the market is pruned.

References (Appendices)

Abramson, Scott F., "The Economic Origins of the Territorial State," *International Organization*, 2017, *71* (1), 97–130.

Anderson, Simon P., André de Palma, and Jacques-François Thisse, *Discrete Choice Theory of Product Differentiation*, Cambridge, MA: MIT Press, 1992.

Arora, Sanjeev, Yingyu Liang, and Tengyu Ma, "A Simple but Tough-to-Beat Baseline for Sentence Embeddings," in "International Conference on Learning Representations (ICLR)" 2017.

Bairoch, Paul, *Cities and Economic Development: From the Dawn of History to the Present*, Chicago, IL: University of Chicago Press, 1988.

Boerner, Lars and Battista Severgnini, "Epidemic Trade," Technical Report, London School of Economics and Political Science, LSE Library 2014.

Bosker, Maarten, Eltjo Buringh, and Jan Luiten van Zanden, "From Baghdad to London: Unraveling Urban Development in Europe, the Middle East, and North Africa, 800–1800," *Review of Economics and Statistics*, 2013, *95* (4), 1418–1437.

Buringh, Eltjo, "European Urban Population, 700–2000," Dataset, version 1.0. 2021.

_ , "The Population of European Cities from 700 to 2000," *Research Data Journal for the Humanities and Social Sciences*, 2021, *6* (1), 1–18.

Couture, Victor, Benjamin Faber, Yizhen Gu, and Lizhi Liu, "Connecting the Countryside via E-Commerce: Evidence from China," *American Economic Review: Insights*, 2021, *3* (1), 35–50.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 1977, *67* (3), 297–308.

Fujita, Masahisa, Paul Krugman, and Anthony J. Venables, *The Spatial Economy: Cities, Regions, and International Trade*, Cambridge, MA: MIT Press, 1999.

Galloway, James A., Derek Keene, and Margaret Murphy, "Fuelling the City: Production and Distribution of Firewood and Fuel in London's Region, 1290–1400," *The Economic History Review*, 1996, *49* (3), 447–472.

Glaeser, Edward L., Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer, "Growth in Cities," *Journal of Political Economy*, 1992, *100* (6), 1126–1152.

Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.

Handbury, Jessie and David E. Weinstein, "Goods Prices and Availability in Cities," *Review of Economic Studies*, 2015, *82* (1), 258–296.

Helpman, Elhanan and Paul R. Krugman, *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*, Cambridge, MA: MIT Press, 1985.

Henderson, J. Vernon, "Marshall's Scale Economies," *Journal of Urban Economics*, 2003, *53* (1), 1–28.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39* (153), 41–57.

Krugman, Paul, "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, 1979, *9* (4), 469–479.

_ , "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 1980, *70* (5), 950–959.

_ , "Increasing Returns and Economic Geography," *Journal of Political Economy*, 1991, *99* (3), 483–499.

Lancaster, Kelvin, *Variety, Equity, and Efficiency: Product Variety in an Industrial Society*, New York: Columbia University Press, 1979.

Masschaele, James, *Peasants, Merchants, and Markets: Inland Trade in Medieval England, 1150–1350*, New York: St. Martin's Press, 1997.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean, "Distributed Representations of Words and Phrases and their Compositionality," in "Advances in Neural Information Processing Systems" 2013, pp. 3111–3119.

Mokyr, Joel, *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton, NJ: Princeton University Press, 2002.

_ , *The Enlightened Economy: An Economic History of Britain, 1700–1850*, New Haven, CT: Yale University Press, 2009.

_ , *A Culture of Growth: The Origins of the Modern Economy*, Princeton, NJ: Princeton University Press, 2017.

Mundlak, Yair, "On the Pooling of Time Series and Cross Section Data," *Econometrica*, 1978, *46* (1), 69–85.

Romer, Paul M., "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 1986, *94* (5), 1002–1037.

Salop, Steven C., "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 1979, *10* (1), 141–156.

Schiff, Nathan, "Cities and Product Variety: Evidence from Restaurants," *Journal of Economic Geography*, 2015, *15* (6), 1085–1123.

Universal Short Title Catalogue, "USTC Database," University of St Andrews, https://ustc.ac.uk 2022. Local snapshot dated 2022-12-13.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, MA: MIT Press, 2010.