Modern Library Holdings and Historic Economic Growth

Eric Chaney*

March 23, 2026

Abstract

In a broad sample of cities, population growth robustly correlates with changes in the number of observed scholarly deaths. I document this relationship using library data and develop a statistical framework to use proxy variables to improve noisy measures of economic activity. Combining existing city growth estimates with predictions based on author counts reduces measurement error. More generally, author counts provide a disaggregated, high-frequency economic proxy that opens new avenues for research in data scarce environments. I illustrate these opportunities by examining the impact of the Mongol invasions on Iraqi cities.

*University of Oxford. This project was inspired by a seminar comment by Robert J. Barro. Many individuals and seminar participants provided helpful discussions and comments. Emily Gilreath and Hayley Moreno provided assistance with the VIAF data and Yazan al-Karablieh able research assistance. Part of this research was carried out while I was Furer fellow at Harvard University, whose financial support is gratefully acknowledged. A previous version of this paper was circulated under the title "Modern Library Holdings and Historic City Growth." I am responsible for all remaining errors.

Estimates of historic economic activity are central to the study of long-run growth. While reliable indicators of past performance are scarce, urbanization has emerged as a widely-used economic proxy in recent years. This metric, however, has well-documented limitations. For example, city sizes are measured infrequently and imperfectly reflect economic conditions. As a result, existing data are often inadequate for the study of historic development.

This paper begins from the assumption that economic and intellectual activity were historically linked to derive a new economic measure from the number of deaths recorded in the union of the world's largest library collections.

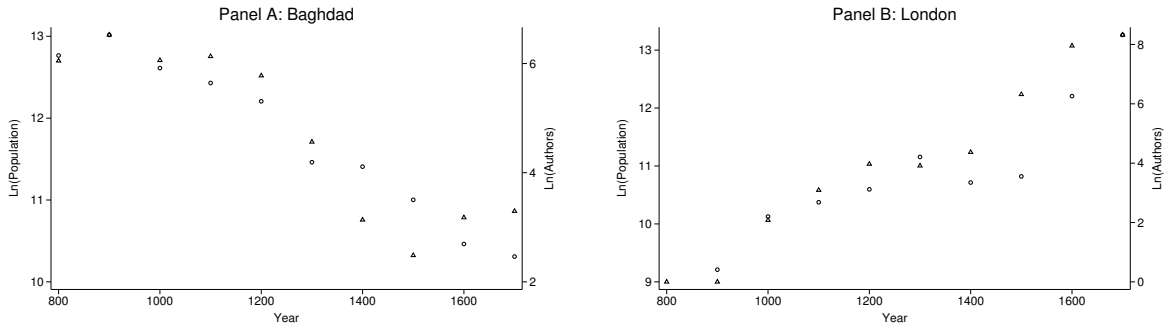


Figure 1: **Author Counts and Population Estimates of Baghdad and London 800-1700**

Notes: Panels A and B respectively provide the evolution of the total number of observed author deaths in Baghdad (Iraq) and London (UK). Circles denote the natural logarithm of existing population estimates (left-hand axis) and triangles the sum of author deaths (right-hand axis).

Figure 1 previews the empirical support for this paper's basic premise, illustrating a close relationship between author counts and population estimates —the only other regularly available disaggregated measure of historical development— in Baghdad and London over the millennium ending in 1800. This relationship, which holds in a broad sample of cities, is consistent with the idea that economic and scholarly activity were positively related (e.g. Kuznets, 1960; Kremer, 1993; Galor and Weil, 2000; Galor, 2023). It suggests that author counts, available at uniquely high levels of spatial disaggregation and temporal frequency, capture novel aspects of the development process.

Because the structural relationships linking scholarly activity to urban development are poorly understood, I develop a flexible machine learning framework to extract economic information from author counts without imposing strong parametric assumptions. Empirically, I show that changes in the observed number of author deaths robustly predict city-level population growth. The estimated relationship is roughly half as strong as that the link between luminosity and GDP in modern data (e.g. Henderson et al., 2012; Michalopoulos and Papaioannou, 2018)— a demanding benchmark given the noise in historical urbanization measures.

I then generalize the framework proposed by Henderson et al. (2012) by showing how these in-sample predictions can be used to improve existing urban population estimates. In other words, the paper extends proxy-based machine learning approaches —typically used to generate out-of-sample estimates— to eliminate measurement error within the observed sample. This yields a general framework for improving noisy measurements of an underlying variable using a vector of proxies for that variable. In this paper's context, the resulting algorithm shows how information created by ongoing advances in computation and digitization can be continuously incorporated into existing city size estimates (e.g. Glaeser et al., 2018; Combes et al., 2022; Hanlon and Heblich, 2022; Jedwab et al., 2022).

Conceptually, this framework can be viewed as wedding Bayesian shrinkage with proxy-based prediction and machine-learning techniques. It is rooted in the fact that conditioning on an exogenous set of proxies reduces the conditional variance of the latent city-growth signal relative to measurement error. This, in turn, provides a frequentist interpretation for familiar Bayesian shrinkage procedures, clarifying why this paper's approach improves upon existing data even when the underlying structural relationships are misspecified.

I also explore the use of author counts as a direct measure of differential development. This complementary strategy, which exploits high-frequency economic information uniquely captured by author counts, largely relies on existing econometric methods (Abadie et al., 2010, 2015; Abadie, 2021; Rambachan and Roth, 2023). It enables quantitative analysis in many environments for the first time.

The remainder of the paper is organized as follows. Section 1 describes

the data. Section 2 validates the data and develops the conceptual framework. A third section provides the empirical applications and a final section concludes.

1 Data

The idea that great libraries should gather all available intellectual output reaches back to antiquity (Heller-Roazen, 2002). Not until the nineteenth century, however, did the widespread emergence of national libraries and modern cataloging practices make it possible to systematically gather written material at scale (e.g. Strout, 1956; Willison, 1989; Krajewski, 2011).

During this period, librarians established cataloging rules to both organize and facilitate access to materials (e.g. Taylor and Joudrey, 2021, pp. 77-78). These efforts included the development of the authority file, which catalogers use to record the name variant that uniquely identifies authors in a given collection (Taylor, 1984; Wiederhold and Reeve, 2021). Although authority files revolutionized inter-library data sharing, language and coordination challenges hindered early attempts to fully exploit these new opportunities through the use of a standardized set of international authority files.

At the start of the twenty-first century, the German National Library, the Library of Congress and Online Computer Library Center proposed a Virtual International Authority File (VIAF) to address such difficulties. Rather than impose a single authoritative name for international use, the VIAF creates a unique number that links the authority files that represent the same author in different collections (e.g. Bennett et al., 2006). For example, the VIAF 120564917 links the files with authoritative names **Khafrī, Muḥammad ibn Aḥmad, active 1526** and **Ḵafrī, Šams al-Dīn Muḥammad b. Aḥmad**: those respectively used by the Library of Congress and the National Library of the Netherlands to denote the sixteenth-century Iranian astronomer.

Over the past two decades, the VIAF consortium has grown to contain many of the world's largest libraries. Updated monthly, the data provide a snapshot of all authors represented in the collections of participating institutions. In this paper, I use the 39,424,267 VIAFs —derived from over

100 million underlying authorities files—available in August 2024.[1]

These data begin to address the limitations of existing collections of academics, creative, notable or otherwise significant individuals (Schich et al., 2014; Serafinelli and Tabellini, 2022; Laouenan et al., 2022; de la Croix et al., 2024) which cover a small fraction of the total number of those who historically engaged in intellectual pursuits.[2] Moreover, they have unique quality guarantees rooted in the authority file's organizing function: cataloguers continuously revise and update the information in these files —including birth and death dates— to ensure catalogs correctly differentiate between authors with similar names (e.g. Robert E. Wolverton, 2005).

Despite these advantages, VIAF imperfectly captures the universe of interest: all individuals who historically engaged in sustained intellectual activity. While the data include the majority of history's most prominent inventors (e.g. 88% of those mentioned by Mokyr (1990)) they also contain sculptors, artists and other prominent figures less clearly linked to knowledge production. Moreover, because library collections are limited by budgetary constraints, scholarly expertise and the survival of historic documents (e.g. `https://www.loc.gov/acq/devpol/cps.html`), VIAF should be viewed as a non-random sample from the underlying population of authors.

These limitations create two sets of empirical challenges. The first, discussed in detail in section 2, stem from usual selection concerns. The second derive from VIAF's relatively sparse coverage of developing regions: a product, in part, of our imperfect understanding of historical knowledge production in these regions. To begin to address this gap, I supplement the VIAF data with authors from Islamic biographical dictionaries. This

[1]Thanks to Emily Gilreath of OCLC for providing the data necessary to approximate the number of underlying authority files.

[2]For example, the largest of these databases covers only 3% of VIAF authors and are largely subsumed by the VIAF collection: in a random sample of 100 individuals of the individuals from the Freebase dataset used by Schich et al. (2014); Serafinelli and Tabellini (2022), 76 had a VIAF. Although this falls to 59 in the larger Laouenan et al. (2022) dataset, 88% of those missing are classified as athletes or politicians. Although VIAF contains only half the historic academics in the de la Croix et al. (2024) sample, the overwhelming majority of these missing academics do not seem to have produced a single written work and in this sense seem less likely to have contributed to knowledge (e.g. de la Croix et al., 2024, p. 1552).

historical source —which has been described as providing an "astonishingly complete and exact" record of the region's scholarly population— extends European coverage rates to the MENA region (Humphreys, 1991, p. 187).

For the remainder of this paper I limit attention to these two high-coverage regions. In particular, I focus on the 1055 MENA and European cities contained in the augmented union of the Bairoch et al. (1988); Bosker et al. (2013); Chandler (1987) databases. As the Appendix details in greater depth, I assign MENA and European authors to these cities by place of death, using information contained in library catalogues, author names and Wikipedia. I geo-reference authors using the place of death —rather than place of birth or activity— for theoretical reasons described in the following section.

2 Validation and Conceptual Framework

At least since Kuznets (1960, p. 328), economists have argued that economic activity is positively related to the "number of geniuses, talented men, and generally gifted contributors to new knowledge." More recently, empirical and theoretical contributions (Kremer, 1993; Galor, 2023) have established an aggregate relationship between the number of exceptionally talented individuals and development as proxied by population size.

There are reasons to suspect this macro-level relationship held at higher levels of disaggregation. For example, Childe (1950) and Acemoglu et al. (2002) provide evidence that agricultural surplus —a basic prerequisite for sustained intellectual activity— historically tracked economic and urban development.

Nevertheless, this more disaggregate relationship between scholarly and economic activity was clearly complex: in some environments scholarly activity simply reflected economic activity while in others also driving economic development (e.g. Mokyr, 2005). It is also undoubtedly the case that institutional and cultural factors could weaken or strengthen the structural connection between economic and intellectual activity (e.g. Simon, 1992; Serafinelli and Tabellini, 2022).

Because these structural links are poorly understood, I focus on documenting —and developing a flexible framework to exploit— the empirical

relationship between these two variables. This alternative approach, while less demanding, continues to face a number of challenges. For example, author counts represent the number of historic deaths observed today, not the underlying scholarly populations of interest.

Variation in underlying scholarly populations can be inferred from observed author counts, however, if death and subsequent inclusion in the database occur roughly at random. In particular, let X_{ijt} denote the Bernoulli trial equal to one if the death of author i in scholarly community j at historic time t is observed today. The total number of deaths observed in a community of size n_{jt} is, in turn, $M_{jt} \equiv \sum_{i=1}^{n_{jt}} X_{ijt}$.[3] Under usual independence and regularity conditions, the delta method implies that the difference in observed deaths across two disjoint time windows of equal length $2c$ is equal to:

$$ln(\sum_{t=\tau-c}^{\tau+c-1} M_{jt}) - ln(\sum_{t=\tau-3c}^{\tau-c-1} M_{jt}) = ln(\bar{n}_{j\tau}) - ln(\bar{n}_{j\tau-1}) + ln(\bar{\gamma}_{j\tau}) - ln(\bar{\gamma}_{j\tau-1}) + \epsilon_{j\tau} \tag{1}$$

where $\bar{n}_{j\tau}$ and $\bar{\gamma}_{j\tau}$ denote, respectively, the average scholarly population and observation probability in community j across the time interval τ and $\epsilon_{j\tau} \sim N(0, \frac{1}{2c}[\frac{1-\bar{\gamma}_{j\tau}}{\bar{n}_{j\tau}\bar{\gamma}_{j\tau}} + \frac{1-\bar{\gamma}_{j\tau-1}}{\bar{n}_{j\tau-1}\bar{\gamma}_{j\tau-1}}])$ (e.g. Walsh, 1952, 1956).

Intuitively, if the relevant observation probabilities are constant over time, author growth measures the underlying scholarly changes with classical measurement error. More generally, regressions of measures of economic growth on author growth provide attenuated estimates of the underlying coefficient of interest —that relating economic and scholarly population growth— if the relevant observation probabilities evolve exogenously.

Columns 1-9 of Table 1 use existing city size estimates, the only consistently available disaggregated measure of historical development, to explore the bivariate relationship between economic and scholarly activity. The century-specific coefficients, the highest frequency at which city growth

[3] X_{ijt} can be further decomposed into indicators denoting death and subsequent inclusion in the sample conditional on death. I abstract from rare historical shocks (e.g. mass mortality events or systematic loss of records) although this framework can be easily modified to accommodate, and is often robust to, localized violations of the independence assumption.

Table 1: **Relationship Between City and Author Growth**

	Δln(Population)											
	900	1000	1100	1200	1300	1400	1500	1600	1700	All	<1700	<1700
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\Delta ln(Authors)$	0.10*	0.15**	0.09**	0.12***	0.13***	0.12***	0.12***	0.13***	0.29***	0.17***	0.12***	0.07***
	(0.06)	(0.06)	(0.04)	(0.03)	(0.02)	(0.04)	(0.03)	(0.02)	(0.03)	(0.01)	(0.01)	(0.02)
Constant	0.11**	0.19***	0.02	0.08***	0.14***	-0.09***	0.08***	0.15***	-0.05**	-	-	-
	(0.04)	(0.04)	(0.03)	(0.03)	(0.05)	(0.03)	(0.03)	(0.02)	(0.02)			
N	63	90	129	140	206	357	367	557	838	2747	1909	1909
R^2	0.07	0.09	0.05	0.04	0.05	0.03	0.04	0.04	0.14	0.10	0.09	0.43
Century Indicators?	-	-	-	-	-	-	-	-	-	Yes	Yes	Yes
City Indicators?	-	-	-	-	-	-	-	-	-	No	No	Yes

Notes: in columns (1)-(9) N denotes the number of observations in each century. In column 10 pools these observations, whereas columns 11 and 12 limit the analysis to periods prior to 1700.) ***, **and *denote significance at the 1%, 5% and 10% levels.

estimates are regularly available, are broadly stable and statistically significant. The strength of this relationship becomes particularly clear when these coefficients are variance averaged in column (10) using the regression:

$$z_{jt} = \alpha_t + \psi x_{jt} + \epsilon_{jt}$$

where z_{jt} denotes existing city growth estimates in city j and century t, x_{jt} author growth, α_t century indicators and ϵ_{jt} the error term.[4]

Columns 11 and 12 explore the robustness of this relationship. In particular, Column 11 limits the sample to years prior to 1700 and Column 12 additionally adds city indicators. In this more demanding specification, the reported elasticities are driven by variation of city and author growth around their city-specific growth paths.[5] These point estimates remain highly statistically significant: roughly one-third as strong as the relationship linking luminosity and GDP.

2.1 Improving Existing Data

These empirical patterns are consistent with the idea that scholarly and economic activity were structurally linked. They suggest that author counts contain novel economic information.

This section develops a framework to extract part of this information without strong structural assumptions. Suppose a sufficiently large number

[4]Throughout the paper I present heteroscedasticity-robust standard errors unless otherwise noted. This is rooted theoretically in the fact that Gibrat's law implies that city growth is i.i.d. (e.g. Eeckhout, 2004) although these standard errors are robust to alternative calculations.

[5]Results throughout this section are qualitatively similar if the sample is limited to the balanced panel of cities.

of estimates z_j measure an underlying variable y_j (e.g. city growth) with classical measurement error $z_j = y_j + \epsilon_j, E[\epsilon_j^2] = \sigma_z^2$.[6]

Although z_j is unbiased, researchers prioritizing group over individual performance can do better by exploiting information that the distribution of z_j is "too wide." Consider the shrinkage estimator

$$\hat{y}_j = (1 - \lambda)E[y] + \lambda z_j$$

with mean-squared error

$$E[(\hat{y}_j - y_j)^2] = \lambda^2\sigma_z^2 + (1 - \lambda)^2\sigma_y^2 \tag{2}$$

where $\sigma_y^2 = E[(y_j - E[y])^2]$.

Minimizing equation (2) yields familiar expressions: the minimum MSE $\frac{\sigma_y^2\sigma_z^2}{\sigma_y^2+\sigma_z^2} < \sigma_z^2$ is attained when $\lambda = \frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2}$ (e.g. Efron and Morris, 1973; Fay III and Herriot, 1979; Robbins, 1983; Kane and Staiger, 2008; Chetty et al., 2014; Angrist et al., 2017; Fessler and Kasy, 2019; Armstrong et al., 2022). In fact, because MSE loss penalizes squared bias and variance symmetrically, equation (2) is algebraically identical to the usual MSE of the weighted average of two unbiased estimators of y_j: z_j and a hypothetical second estimator with error variance σ_y^2. In other words, this paper's composite generalizes familiar gains from estimator averaging to settings with biased estimators: balancing the gains from noise reduction against the cost of introducing an exogenous source of estimation error from a second estimator —here, bias rather than measurement error.

Conditioning the analysis on a vector of proxies $\boldsymbol{x_j}$ (e.g. growth in author counts) orthogonal to the measurement error ϵ_j generates additional gains. Mechanically, it partitions observations into groups that are more homogenous in the underlying y_j which, in turn, increases the optimal weight on the conditional mean. In the maintained analogy, this is equivalent to an exogenous reduction in the variance of the second hypothetical unbiased estimator. Intuitively, in the limiting case where $\boldsymbol{x_j}$ perfectly predicts y_j, all the remaining variation in z_j is noise and the estimator fully shrinks z_j toward $E[y|\boldsymbol{x}]$, recovering the true y_j.

[6]For simplicity, I assume homoskedastic measurement error. Allowing for heteroskedasticity is straightforward and does not alter this section's basic insights.

More generally, shrinking towards a linear approximation $\boldsymbol{\psi}'\boldsymbol{x_j}$ to the conditional expectation function (CEF) yields the optimal shrinkage factor

$$\lambda_{\boldsymbol{\psi}'\boldsymbol{x_j}} = \frac{E[(y_j - \boldsymbol{\psi}'\boldsymbol{x_j})^2]}{E[(y_j - \boldsymbol{\psi}'\boldsymbol{x_j})^2] + \sigma_z^2}$$

with corresponding MSE

$$\frac{E[(y_j - \boldsymbol{\psi}'\boldsymbol{x_j})^2]\sigma_z^2}{E[(y_j - \boldsymbol{\psi}'\boldsymbol{x_j})^2] + \sigma_z^2}$$

that is strictly smaller than unconditional shrinkage whenever $\boldsymbol{x_j}$ has predictive content.[7] Letting R^2_{xz} denote the R^2 from regressing z_j on $\boldsymbol{x_j}$, $\lambda_{\boldsymbol{\psi}'\boldsymbol{x_j}}$ can be written as the share of the residual variance in z_j attributable to unexplained heterogeneity in y_j: $\frac{\frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2} - R^2_{xz}}{1-R^2_{xz}}$

In other words, conditional shrinkage weakly dominates unconditional shrinkage in MSE terms, even under misspecification of the conditional mean. However, the magnitude of these gains —the extent to which the predictive power of the proxy vector is effectively exploited— depends on a researcher's ability to approximate the true CEF and error variance.[8]

In some settings, theory provides —at least approximately— the appropriate CEF (e.g. Fessler and Kasy, 2019); when it does not, the CEF must be estimated. Machine learning methods provide a natural estimation strategy in this second case (e.g. Mullainathan and Spiess, 2017; Athey and Imbens, 2019) as illustrated by recent work that exploits economic proxies —such as night lights— to extend income measures to data-sparse regions (e.g. Donaldson and Storeygard, 2016; Jean et al., 2016; Khachiyan et al., 2022; Koch et al., 2024; Rossi-Hansberg and Zhang, 2025).

With the notable exception of Henderson et al. (2012), however, existing studies emphasize out-of-sample prediction and do not shrink the observed estimates toward their predicted values. While this practice is largely inconsequential when the original data are precisely measured, it

[7] I assume that the conditional variance of y_j is constant. If it is instead heteroskedastic, shrinkage using a constant factor based on the average conditional variance continues to yield MSE gains, though further improvements are possible by allowing the shrinkage factor to vary with the conditional variance.

[8] The composite is also robust to misspecification of the error variance. In particular, when the unexplained variance of the latent heterogeneity is smaller than σ_z^2, any convex combination that places positive weight on the fitted values reduces the MSE.

sacrifices potentially substantial MSE gains when the data are noisy. Henderson et al. (2012) make this point clearly: in the GDP-luminosity context they apply what can be seen as a nested version of the more general framework developed here to obtain in-sample improvements that machine learning approaches currently discard.[9] This paper's framework addresses this omission by combining shrinkage and prediction in a unified approach that systematically exploits information contained in the proxy vector — both within and out of sample.

2.2 Author Counts as a Direct Proxy for Differential Development

This section develops a complementary approach that uses author counts as a direct proxy for economic activity. Under minimal assumptions, it recovers differential changes in the scholarly populations which, in turn, often reflect differential economic development.

Equation (1) implies that the logarithm of city-level author counts recovers changes in average scholarly populations within a standard difference-in-differences (DiD) framework, provided the observation probabilities satisfy an analogue of the parallel trend assumption. Similarly, synthetic control methods can be applied to the logarithm of region-level aggregates, as long as the average observation probabilities do not vary systematically across regions (e.g. Abadie et al., 2010; Abadie, 2021).

The intuition here is straightforward. The logarithm of the sum of observed deaths in city (or region) j over a window of length $2c$ equals the logarithm of the average scholarly population, plus classical measurement error and a term driven by the logarithm of the observation probability. When the latter differences out —either across time or between treated and control units— the use of author counts introduces classical measurement error, yielding unbiased but less precise estimates.

[9]The composite estimator proposed by Henderson et al. (2012) is algebraically equivalent to the linear approximation to the optimal developed here when $\boldsymbol{x_j}$ is scalar. Although they abstract from nonlinear or higher-order dependencies, the analysis here implies that their estimator remains valid as a linear approximation to the theoretical optimum even when such dependencies are present. Under their maintained assumption of zero covariance between $\boldsymbol{x_j}$ and the error term, however, only this linear approximation —not the optimum itself— may be attainable.

Formally, equation (1) implies —letting D_j denotes binary treatment and D_τ window indicators— that β_τ in the canonical DiD specification

$$ln(\sum_{t=\tau-c}^{\tau+c-1} M_{jt}) = \alpha D_j + \delta_\tau + \sum_{\tau \neq \tau_0} \beta_\tau D_j \cdot D_\tau + \eta_\tau$$

consistently estimates the average treatment effect on the treated (ATET)

$$E[ln(\bar{n}_{j\tau}(1)) - ln(\bar{n}_{j\tau}(0))|D_j = 1, D_\tau = 1] \tag{3}$$

for $\tau > \tau_0$, under an augmented parallel trends assumption: the usual parallel trends assumption

$$\begin{aligned} E[ln(\bar{n}_{j\tau}(0))|D_j = 1, Post_\tau = 1] - E[ln(\bar{n}_{j\tau}(0))|D_j = 1, Post_\tau = 0] = \\ E[ln(\bar{n}_{j\tau}(0))|D_j = 0, Post_\tau = 1] - E[ln(\bar{n}_{j\tau}(0))|D_j = 0, Post_\tau = 0] \end{aligned} \tag{4}$$

plus the assumption the average observation probabilities also evolve in parallel

$$\begin{aligned} E[ln(\bar{\gamma}_{j\tau}(1))|D_j = 1, Post_\tau = 1] - E[ln(\bar{\gamma}_{j\tau}(0))|D_j = 1, Post_\tau = 0] = \\ E[ln(\bar{\gamma}_{j\tau}(0))|D_j = 0, Post_\tau = 1] - E[ln(\bar{\gamma}_{j\tau}(0))|D_j = 0, Post_\tau = 0] \end{aligned} \tag{5}$$

In other words, the logarithm of author counts used in the standard DiD framework will identify the differential response of scholarly population to a given treatment.[10]

In practice, small scholarly populations or low observation probabilities may generate zeros. In such environments, the discussion in Chen and Roth (2024) implies $exp(\delta_\tau) - 1$ in the Poisson model $E[\sum_{t=\tau-c}^{\tau+c-1} M_{jt}|D_j, D_\tau] = exp(\alpha D_j + \delta_\tau + \sum_{\tau \neq \tau_0} \beta_\tau D_j \cdot D_\tau)$ consistently estimates

$$\frac{E[\bar{n}_{j\tau}(1)|D_j = 1, D_\tau = 1] - E[\bar{n}_{j\tau}(0)|D_j = 1, D_\tau = 1]}{E[\bar{n}_{j\tau}(0)|D_j = 1, D_\tau = 1]}$$

[10]To be completely precise, these assumptions are too strong: the parameter of interest is also identified if the violation of the usual parallel trend assumption and that of the additional restriction on observation probabilities move in opposite directions but are of the same sign.

under the ratio version of equations (4) and (5).[11]

Alternatively, it is possible to apply synthetic control methods to aggregate counts across treated cities and a set of k control regions defined by the researcher. These approaches differ primarily in how untreated units are combined to construct the relevant counterfactual (e.g. Arkhangelsky et al., 2021; Roth et al., 2023).

To relate the scholarly ATET in equation (3) to economic outcomes, assume city GDP (Y_{jt}) and scholarly populations satisfy $n_{jt} = f(\phi_{jt}, Y_{jt})$ where f is continuously differentiable and ϕ_{jt} captures structural heterogeneity. For simplicity, consider a single-period window so that $\bar{n}_{j\tau} = n_{jt}$.[12] Dropping subscripts, a first-order Taylor approximation yields:

$$E[\frac{\phi(0)}{f(0)}\frac{\partial f(0)}{\partial \phi}(\frac{\phi(1)-\phi(0)}{\phi(0)}) + \frac{Y(0)}{f(0)}\frac{\partial f(0)}{\partial Y}(\frac{Y(1)-Y(0)}{Y(0)})|D=1] \qquad (6)$$

where $f(0) \equiv f(\phi(0), Y(0))$.

The first term in equation (6) captures the treatment's effect on scholarly population that operates through changes to the structural relationship between income and scholarly activity. The second term —this paper's focus— captures changes in scholarly activity driven by economic activity.

If a treatment does not alter the mapping between scholarly and economic outcomes and is independent of the income elasticities, the income ATET can be obtained by scaling the author-based ATET by the average income elasticity among treated cities. More generally, if scholarly activity is a normal good, the estimator identifies a weighted average of GDP treatment effects, with weights proportional to the income elasticity of scholarly activity. In this sense, using author counts in place of city-level GDP can only lead to spurious economic conclusions if a treatment sufficiently alters the structural mapping between income and scholarly production itself.

[11]For expositional simplicity, throughout this paper I generally abstract from the implications of interchanging expectations and logarithms. See Roth et al. (e.g. 2023) for a discussion of the sensitivity of the difference-in-differences framework to such variations in functional form. In addition, I extend for this discussion the assumption that the individual observation probabilities are relatively similar to all authors in, respectively, the control and treatment cities.

[12]The same basic intuition developed here, however, extends to multi-period windows with averages replacing single-period values.

Importantly, the relevant comparisons when evaluating this assumption is between the structural relationships in treated cities and their counterfactuals at time t. In other words, the validity of this assumption does not rest on assuming that the underlying structural relationships are stable over time. In many contexts, concerns surrounding this assumptions validity can be addressed using alternative economic proxies, such as existing city size estimates, as illustrated in the following section.

3 Empirical Applications

Here I revisit Baghdad's population estimates and assess the impact of the Mongol invasions on Iraqi cities. Both subjects remain contested, in part due to the absence of systematic quantitative evidence. For example, estimates of Baghdad's peak population range from two million (e.g. Hamdan, 1962) to 450,000 Bosker et al. (2013).

3.1 New Population Estimates for Baghdad, Iraq

Columns (2) and (3) of Table 2 report existing century-level population estimates from Bosker et al. (2013) along with the corresponding author counts. The two series broadly co-move, declining from medieval highs to lows towards the end of the sample.

This correspondence, however, often breaks down. For example, between 900 and 1000 the population estimate declines by 150,000 while author counts increase by 21.

Such discrepancies may reflect a host of factors, including time-varying observation probabilities or structural shifts in the underlying relationship between scholarly activity and urban population. Nevertheless, one of this paper's central insights is that author counts contain systematic economic information that can be used to improve city-size estimates in the aggregate. In other words, although individual revisions using the composite developed in section 2.1 may occasionally move in the wrong direction, the aggregate gains from correct adjustments are guaranteed to outweigh such losses.

I estimate the CEF underlying the composite estimator using Lasso

Table 2: **Comparison of Existing and New Growth Rates: Baghdad, Iraq**

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Year	Pop. Level	Author Level	Pop. Growth	Author Growth	Fitted	$\hat{\lambda}$	New Pop. Growth	95% Confidence Interval	New Pop. Level
800	350	276	-	-	-	-	-	-	222
900	450	564	0.25	0.71	0.27	0.74	0.26	[-0.16,0.68]	287
1000	300	585	-0.41	0.04	1.12	0.74	-0.01	[-0.43,0.41]	284
1100	250	373	-0.18	-0.45	-0.23	0.74	-0.19	[-0.61,0.23]	235
1200	200	501	-0.22	0.30	-0.27	0.74	-0.23	[-0.66,0.19]	187
1300	95	140	-0.74	-1.27	-0.55	0.74	-0.69	[-1.11,-0.27]	94
1400	90	52	-0.05	-0.99	-0.54	0.74	-0.18	[-0.60,0.24]	78
1500	60	12	-0.41	-1.47	-0.44	0.74	-0.41	[-0.83,0.01]	52
1600	35	23	-0.54	0.65	-0.08	0.74	-0.42	[-0.84,0.00]	34
1700	30	14	-0.15	-0.50	-0.00	0.85	-0.13	[-0.49,0.22]	30

Notes: column 2 provides the Bosker et al. (2013) population estimates and column 3 author counts. Columns 4 and 5 respectively report the population and author growth rates. Column 6 reports the fitted values obtained from OLS estimation of the Lasso specification whereas column 7 reports the optimal weight on the existing growth estimates. Column 8 provides the composite estimate corresponding to the optimal weighting and column 9 the Armstrong et al. (2022) 95% confidence interval. Column 10 provides the updated estimates of city levels.

for three reasons.[13] First, because the relationship between author and city growth is arguably approximately sparse, Lasso provides a near-oracle approximation to the true CEF (e.g. Belloni and Chernozhukov, 2013). Second, first-differencing removes the time-invariant components of city-specific observation probabilities. Third, because theory suggests city growth is i.i.d. (e.g. Eeckhout, 2004), estimating the model in growth rates instead of levels arguably mitigates overfitting concerns stemming from the likely non-stationary of city sizes and their proxies (Belloni et al., 2016; Babii et al., 2021; Brignoli et al., 2024).

I divide the sample at 1700 given evidence that existing data is more precisely measured that year. In both samples, I then reserve 40% of the data for the hold-out sample and estimate a linear Lasso specification in which I include century and geographic indicators, author growth, the interaction of author growth with both the century and geographic indicators, the initial population level and its square, the initial number of authors normalized by the baseline population and its interaction with author growth, the logarithm of the baseline number of observed deaths and its interaction

[13] Although recent work extends Lasso methods to panel environments, such extensions are less relevant in this paper's environment (e.g. Belloni et al., 2016, p. 591).

with author growth.[14]

In the 1700 subsample, Lasso achieves an in-sample R^2 of 0.26 and a hold-out R^2 of 0.25. In the pre-1700 sample, the in-sample and hold-out R^2 are respectively 0.15 and 0.14— roughly double the naive linear specification reported in Table 1. These R^2 are roughly half those obtained in modern settings: a demanding benchmark given the much higher quality of income and proxy data today (e.g. Khachiyan et al., 2022).

Column (6) of Table 2 reports the fitted growth rates. The optimal degree of shrinkage toward these fitted values depends on the (unobserved) signal-to-noise ratio in the original population estimates. To estimate this quantity, I use competing city size estimates from de Vries (1984) and Bosker et al. (2013), which imply signal-to-noise ratios of 0.89 for 1700 and 0.78 for 1600. For earlier periods, I assume that 1600 value applies. This strategy is conservative: it likely under-utilizes the predictive content and is consequently unlikely to under-perform the original estimates.

Column (7) reports the resulting weight on the original estimate, column (8) the composite growth estimate and column 9 the associated 95% confidence interval (Armstrong et al., 2022). Although these intervals are wide —reflecting the uncertainty surrounding Baghdad's historical population— they are approximately 12% shorter than the pre-shrinkage intervals.

Finally, Figure 2 evaluates the out-of-sample performance of the usual proxy-based estimates using newly constructed city growth estimates from Buringh (2021). These estimates explains approximately 10% of the variation in these independently constructed growth estimates. This, in turn, is consistent with both the in and hold-out sample performance, suggesting that author counts capture a stable fraction of the historical growth process.

[14] Although some of these predictors are unlikely to satisfy the exogeneity assumption, I include them for completeness as their exclusion does not qualitatively affect the analysis. The geographic regions are Southeastern Europe (Albania, Bosnia and Herzegovina, Bulgaria, Greece, Hungary, Kosovo, North Macedonia, Romania, Serbia and Slovakia), Central Europe (Austria, Czechia, Estonia, Germany, Latvia, Lithuania, Poland, Slovenia and Switzerland), Eastern Europe (Belarus, Russia and Ukraine), France (Belgium and France), Northern Europe (Denmark, Finland, Netherlands, Norway and Sweden), the UK (Ireland and the UK), Italy (Italy and Malta), Iberia (Portugal and Spain) and cities under Muslim rule. To select the tuning parameter I use cross validation with 10 folds.

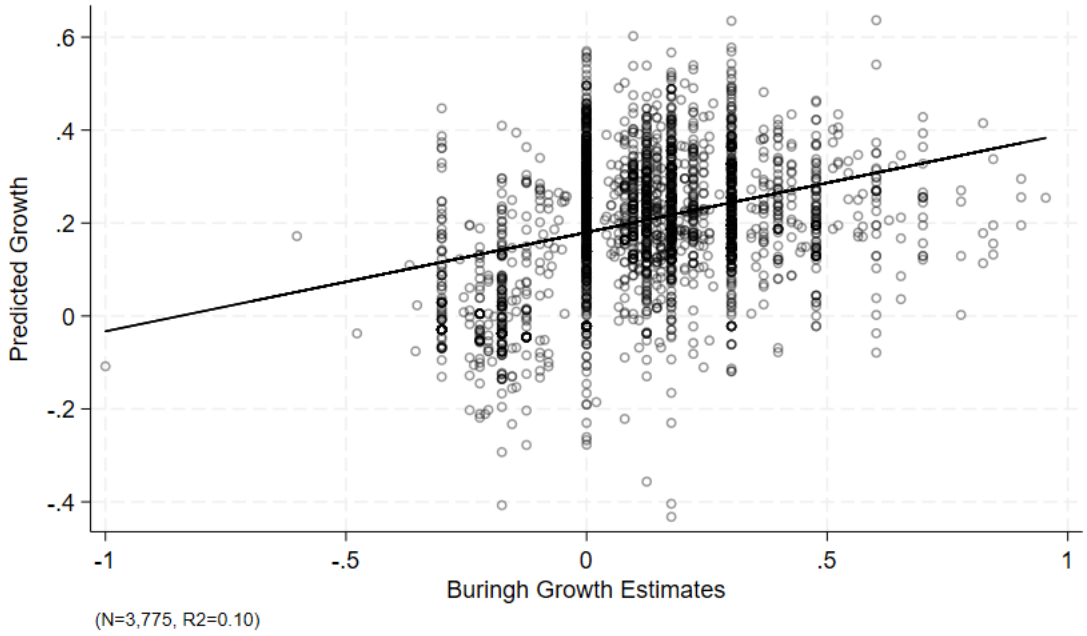


Figure 2: **Predicted and Estimated Growth prior to 1700**

3.2 The Impact of the Mongol Invasions on Iraq

Traditionally, historians have viewed the Mongol invasions of the thirteenth century as a critical juncture that "ended the great Muslim civilization of the Middle Ages" (e.g. Lewis, 1968, p. 49) in Iraq. More recent scholarship has reassessed this interpretation, arguing that economic recovery and renewed prosperity followed the conquest (e.g. Biran, 2022, p. 286). Because these contrasting conclusions rely, in part, on different readings of the same qualitative evidence the economic impact of the Mongol invasions remains a subject of ongoing scholarly debate (e.g., Jackson, 2017).

Quantitative analysis of this question has been limited by data constraints. Population estimates at the century level are available for only two Iraqi cities on the interval [800,1800), limiting DiD analysis. Alternative approaches face similar challenges: with only four pre-treatment observations, it is difficult to demonstrate that a synthetic control can reproduce Iraq's pre-invasion trends (e.g. Abadie, 2021).

Author counts enable systematic quantitative analysis in this environment for the first time. Panel A of Figure 3 plots the log average number of author deaths across the 13 Iraqi cities in the sample, with a vertical line

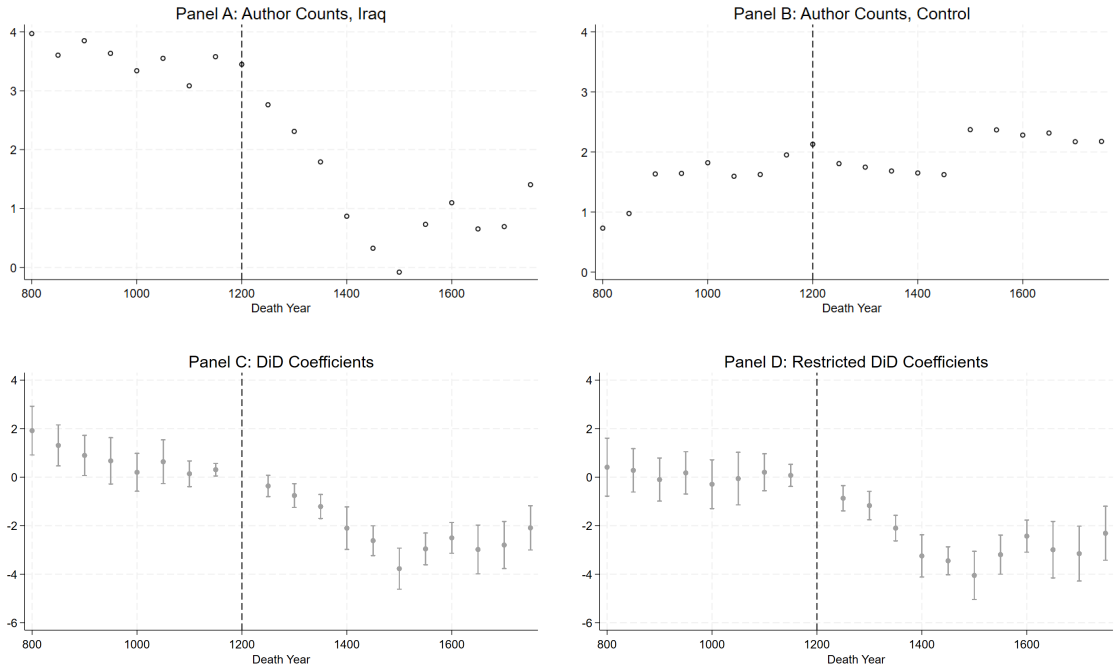


Figure 3: **DiD Estimates of the Effect of Mongol Invasions on Iraqi Cities**

marking [1200,1250), the final window prior to the Mongol invasion sack of Baghdad in 1258. This graph documents a sustained decline following the invasions. For example, total observed deaths fall from 409 in [1200,1250) —an average of 31.46 per city— to 1.38 per city in [1450,1500): a decline of 313 log points (96%). This decline is broad-based: in the second period I only observed deaths in a handful of cities.

These patterns could reflect the coverage of my sources rather than true scholarly decline. Panel B addresses this concern, plotting the average observed author deaths across 250 cities in the western Islamic world.[15] Although the average level of scholarly deaths are lower —consistent with Iraq's pre-Mongol prominence— the decline over time is more modest dropping from a mean of 8.412 to 5.072.

Under the ratio version of the parallel trend assumption (Chen and Roth, 2024) the implied counterfactual for Iraq over the interval [1450,1500)

[15]Here, I adopt a liberal definition of the western Islamic world, including cities in modern-day Algeria, Egypt, Israel, Lebanon, Libya, Morocco, Oman, Portugal, Saudi Arabia, Spain, Syria, Tunisia, Turkey and Yemen regardless of whether these cities were actually ruled by Islamic dynasties.

is 18.97, obtained by scaling the pre-invasion mean (31.46) by the ratio $\frac{5.072}{8.412}$. This yields a treatment effect of $ln(1.38) - ln(18.97)$ or -262 log points, the [1450,1500) coefficient in panel C.

Nevertheless, this panel provides reasons to question the validity of the parallel-trends assumption. Several pre-treatment coefficients are positive, providing evidence of differential pre-trends. This could reflect faster economic growth in the control regions, pre-existing decline in Iraq, or differential changes in the observation probabilities. In addition, parts of the control region were affected by a number of shocks during this period including increased European military activity in Iberia and western North Africa.

Panel D addresses these concerns by restricting the control group to 19 cities in the Arabian Peninsula and modern-day Libya and Tunisia. The pre-treatment fit improves in this modified specification, imply that the impact of the Mongol invasion led to a drop in Iraq's scholarly population of 344 log points by [1450,1500). In other words, the DiD analysis of the impact of the Mongol invasion suggests the invasions led to the near-elimination of Iraq's scholarly population over the long-run.

This empirical strategy, however, has limitations. For example, although historical reasoning suggests plausible control regions, it provides little guidance on which of these regions should be preferred. Moreover, although the 50 year windows reduce statistical noise, they limit the ability to pinpoint the exact timing of the shock. For example, did the inhabitants of Iraq anticipate the invasions after their start in the Eastern Islamic world 40 years earlier?

Figure 4 addresses these issues by estimating the ATET using synthetic control methods on the average number of scholarly deaths in Iraq and 11 regions by 20 year intervals.[16] Panel A shows the synthetic control closely tracks Iraq's scholarly trajectory prior to the invasion and a decline after the vertical line in 1200: the final 20 year interval prior to the start of the Mongol invasion of the Eastern Islamic world: suggesting Iraqis may have

[16]These include the cities of Cairo, Damascus, Istanbul, Mecca and Medina and the aggregate of remaining cities in the Arabian Peninsula and Israel, Egypt, Turkey, Tunisia, Lebanon-Syria and Algeria, Morocco, Portugal and Spain. Following Abadie (2021), I smooth the data using a 60-year moving average although the results are robust to using raw data.

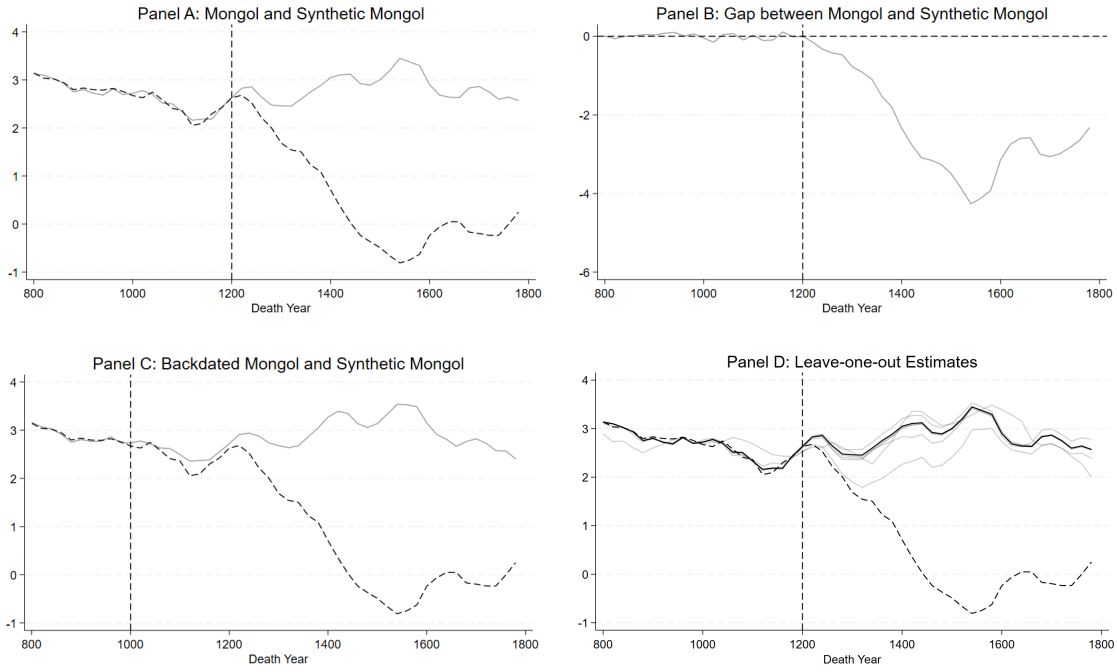


Figure 4: **Synthetic Control Estimates of the Effect of Mongol Invasions on Iraqi Cities**

anticipated the invasions. Panel B provides the corresponding estimate of the ATET, a decline of 337 log points by 1460 which is almost identical to the previously reported difference-in-differences estimate. These results are highly statistically significant: using the placebo method the average effect of the Mongol invasion is estimated to be a drop of 240 log points with a standard error of 89 log points.[17]

Panels C and D provide robustness checks. Panel C provides results backdating the Mongol invasions two hundred years to 1020. Reassuringly, the synthetic control tracks the average scholarly population of Iraq up until the invasions. Panel D provides the leave-one-out analysis, omitting from the sample one-at-a-time each of the regions that contribute to the synthetic control.[18] In general, the leave-one-out estimates track Iraq's average scholarly population closely prior to the invasion and the resulting

[17]This become a drop if 248 log points with a standard error of 76 log points if the regions containing cities in the Iberian peninsula are dropped.

[18]These regions with their corresponding weights in parentheses are Large Iberian Cities (.03826), Smaller Maghribi Cities (.0034), Tunisia (.0099), Smaller Turkish Cities (.0002), Mecca (.1992), Medina (.4262), Damascus (.0400) and Cairo (.2828).

estimates of the invasions over the long-run are all highly negative.[19] In other words, the results are robust to the omission of any particular region.

This analysis provides the first systematic quantitative evidence that the Mongol invasions contributed to Iraq's scholarly decline. Is this collapse in observed scholarly activity also reflective of economic decline? This paper's conceptual framework highlights two threats to this interpretation of the results. First, the invasions may have altered the structural relationship between scholarly production and economic activity. For instance, Iraq may have flourished economically even as subsequent regimes eliminated institutional support for scholarship. Second, the invasions may have affected the observation probabilities —for example, by shifting Iraq's cultural center away from the data sources on which this study relies.

Available population estimates weigh against of these explanations. As shown in Table 2, estimates for Baghdad indicate a sustained population decline following the conquest. Similar patterns are observed for other Iraqi cities with population data from both before and after the invasion. These declines are not consistent with the idea that Iraq did not suffer economically following the invasions.

Overall, these findings are consistent with aspects of both the revisionist and traditional interpretations of the Mongol invasion. On one hand, they suggest the revisionists are correct with they argue the Iraqi economy did not immediately collapse after the conquest. On the other, the results align with the traditional view that the invasion initiated institutional changes that gradually eroded Iraq's long-run economic vitality (e.g. Ashtor, 1976, p. 250). Further research is needed to better understand these empirical patterns, although they highlight how this paper's approach can facilitate systematic empirical analysis in data-scarce contexts.

[19]Specifications that show a decline in the synthetic control after the invasions place relatively high weight on cities in the Iberian Peninsula, which as noted above faced an unrelated large shock around the period of the Mongol invasions. Results are robust to omitting cities in this region, although I include it here for transparency.

4 Conclusion

This study shows that author counts proxy for historic economic activity. Changes in the number of authors affiliated with a city robustly correlate with existing population growth estimates, providing a relatively high-frequency measure of urban activity where traditional data are unavailable.

To illustrate how high-frequency variation in author counts deepens our understanding of historical economic activity, I examine the yearly impact of the Mongol invasions on Iraq.

As in Henderson et al. (2012), I shrink existing city growth estimates towards fitted values obtained by relating existing estimates to a proxy for these estimates. In so doing, I provide a Bayesian interpretation that generalizes their estimator to non-linear environments and those with more than one proxy while providing additional intuition for this result.

For cities without population estimates, author data enables previously impossible quantitative analysis. For others, it allows researchers to track the impact of events at a relatively high frequency for the first time.

The collection of authors used in this paper should be viewed as a first approximation to the population of scholars working in Europe and the MENA region. The results suggest that refining and extending these data will further our understanding of historic economic activity.

References

Abadie, A., "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects," *Journal of Economic Literature*, 2021, *59* (2), 391–425.

_ , A. Diamond, and J. Hainmueller, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 2010, *105* (6), 493–505.

_ , _ , and _ , "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 2015, *59* (2), 495–510.

Acemoglu, D., S. Johnson, and J.A. Robinson, "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," *Quarterly Journal of Economics*, 2002, *117* (4), 1231–1294.

Angrist, J., P. Hull, and P. Pathak, "Leveraging Lotteries for School Value-Added: Testing and Estimation," *Quarterly Journal of Economics*, 2017, *132* (2), 871–919.

Arkhangelsky, D., S. Athey, D. Hirshberg, G. Imbens, and S. Wager, "Synthetic Difference-in-Differences," *American Economic Review*, 2021, *111* (12), 4088–4118.

Armstrong, T., M. Kolesar, and M. Plagborg-Moller, "Robust Empirical Bayes Confidence Intervals," *Econometrica*, 2022, *90* (6), 2567–2602.

Ashtor, Eliyahu, *A Social and Economic History of the Near East in the Middle Ages*, Berkeley and Los Angeles: University of California Press, 1976.

Athey, S. and G.W. Imbens, "Machine Learning Methods that Economists Should Know About," *Annual Review of Economics*, 2019, *11*, 685–725.

Babii, A., E. Ghysels, and J. Striaukas, "Machine Learning Time Series Regressions With an Application to Nowcasting," *SSRN Working Paper*, 2021.

Bairoch, P., J. Batou, and P. Chevre, *La population des ville europeenees de 800 a 1850: Banque de donnees et analyse sommaire de resultats.*, Geneva: Librarie Droz, 1988.

Belloni, A. and V. Chernozhukov, "Least squares after model selection in high-dimensional sparse models," *Bernoulli*, 2013, *19* (2), 521–547.

_ , _ , C. Hansen, and D. Kozbur, "Inference in High-Dimensional Panel Models with an Application to Gun Control," *Journal of Business and Economic Statistics*, 2016, *34* (4), 590–605.

Bennett, Rick, Christina Hengel-Dittrich, Edward T. O'Neill, and Barbara B. Tillett, "VIAF (Virtual International Authority File): Linking Die Deutsche Bibliothek and Library of Congress Name Authority Files," in "World Library and Information Congress: 72nd IFLA General Conference and Council" Seoul, Korea August 2006. Meeting: 123 Cataloguing.

Biran, Michal, "Baghdād under Mongol Rule," in "Baghdād: From Its Beginnings to the 14th Century," Vol. 166 of *Handbook of Oriental Studies. Section 1: The Near and Middle East*, Leiden: Brill, 2022, pp. 285–315.

Bosker, M., E. Buringh, and J.L. 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.

Brignoli, Paolo Libenzio, Alessandro Varacca, Cornelis Gardebroek, and Paolo Sckokai, "Machine learning to predict grains futures prices," *Agricultural Economics*, 2024, *55* (3), 479–497.

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

Chandler, T., *Four thousand years of urban growth: an historical census*, Lewiston: Edwin Mellen Press, 1987.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, May 2024, *139* (2), 891–936.

Chetty, R., J.N. Friedman, and J.E. Rockoff, "Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates," *American Economic Review*, 2014, *104* (9), 2593–2632.

Childe, V. Gordon, "The Urban Revolution," *Town Planning Review*, 1950, *21* (1), 3–17.

Combes, P.-P., L. Gobillon, and Y. Zylberberg, "Urban economics in a historical perspective: Recovering data with machine learning," *Regional Science and Urban Economics*, 2022, *94*.

de la Croix, D., F. Docquier, A. Fabre, and R. Stelter, "The Academic Market and the Rise of Universities in Medieval and Early Modern Europe (1000-1800)," *Journal of the European Economic Association*, 2024, *22* (4), 1541–1589.

de Vries, J., *European Urbanization 1500-1800*, London: Methuen, 1984.

Donaldson, D. and A. Storeygard, "The View from Above: Applications of Sateliite Data in Economics," *Journal of Economic Perspectives*, 2016, *30* (4), 171–198.

Eeckhout, J., "Gibrat's Law for (All) Cities," *American Economic Review*, 2004, *94* (5).

Efron, B. and C. Morris, "Stein's Estimation Rule and Its Competitors–An Empirical Bayes Approach," *Journal of the American Statistical Association*, 1973, *68* (341), 117–130.

Fay III, R.E and R.A. Herriot, "Estimates of Income for Small Places: An Application of James-Stein Procedures to Census Data," *Journal of the American Statistical Assocation*, 1979, *74* (366), 269–277.

Fessler, P. and M. Kasy, "How to use Economic Theory to Improve Estimators: Shrinking toward Theoretical Restrictions," *Review of Economics and Statistics*, 2019, *101* (4), 681–698.

Galor, Oded, *The Journey of Humanity: A New History of Wealth and Inequality with Implications for Our Future*, New York: Dutton, 2023.

_ **and David N. Weil**, "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond," *American Economic Review*, 2000, *90* (4), 806–828.

Glaeser, E., S.D. Kominers, M. Luca, and N. Naik, "Big Data and Big Cities: the Promises and Limitations of Improved Measures of Urban Life," *Economic Inquiry*, 2018, *56* (1), 114–137.

Hamdan, G., "The Pattern of Medieval Urbanism in the Arab World," *Geography*, April 1962, *47* (2), 121–134.

Hanlon, W.W. and S. Heblich, "History and urban economics," *Regional Science and Urban Economics*, 2022, *94*.

Heller-Roazen, D., "Tradition's Destruction: On the Library of Alexandria," *Obsolescence*, 2002, *100*, 133–153.

Henderson, J.V., A. Storeygard, and D. Weil, "Measuring Economic Growth from Outer Space," *American Economic Review*, 2012, *102* (2), 994–1028.

Humphreys, R. Stephen, *Islamic History: A Framework for Inquiry*, revised ed., Princeton, NJ: Princeton University Press, 1991.

Jackson, Peter, *The Mongols and the Islamic World: From Conquest to Conversion*, New Haven: Yale University Press, 2017.

Jean, N., M. Burke, M. Xie, W.M. Davis, D.B. Lobell, and S. Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, 2016, *353* (6301), 790–794.

Jedwab, R., N. Johnson, and M. Koyama, "Medieval cities through the lens of urban economics," *Regional Science and Urban Economics*, 2022, *94*.

Kane, T. and D. Staiger, "Estimating Teacher Impacts on Student Achievement: An Experimental Evalution," *NBER Working Paper*, 2008, (14607).

Khachiyan, A., A. Thomas, H. Zhou, G. Hanson, A. Cloninger, T. Rosing, and A. Khandelwal, "Using Neural Networks to Predict Microspatial Economic Growth," *AER: Insights*, 2022, *4* (4), 491–506.

Koch, P., V. Stojkoski, and C. Hidalgo, "Augmenting the availability of historical GDP per capita estimates through machine learning," *PNAS*, 2024, *121* (39).

Krajewski, M., *Paper Machines about cards & catalogs, 1548-1929*, MIT Press, 2011.

Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *The Quarterly Journal of Economics*, 1993, *108* (3), 681–716.

Kuznets, Simon, "Population Change and Aggregate Output," in "Demographic and Economic Change in Developed Countries," Columbia University Press, 1960, pp. 324–351.

Laouenan, Morgane, Prashant Bhargava, Jean-Benoît Eyméoud, and et al., "A cross-verified database of notable people, 3500BC–2018AD," *Scientific Data*, 2022, *9*, 290.

Lewis, B., "The Mongols, the Turks and the Muslim Polity," *Transactions of the Royal Historical Society*, 1968, *18*, 49–68.

Michalopoulos, S. and E. Papaioannou, "Spatial Patterns of Development: a Meso Approach," *Annual Reviews of Economics*, 2018, *10*, 383–410.

Mokyr, Joel, *The Lever of Riches: Technological Creativity and Economic Progress*, New York: Oxford University Press, 1990.

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

Mullainathan, S. and J. Spiess, "Machine Learning: an Applied Econometric Approach," *Journal of Economic Perspectives*, 2017, *31* (2), 87–106.

Rambachan, A. and J. Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *Forthcoming.*

Robbins, H., "Some Thoughts on Empirical Bayes Estimation," *The Annals of Statistics*, 1983, *11* (3), 713–723.

Rossi-Hansberg, E. and J. Zhang, "Local GDP Estimates Around the World," *Working Paper*, 2025.

Roth, J., P.H.C. Sant'Anna, A. Bilinkski, and J. Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235*, 2218–2244.

Schich, M., C. Song, Y. Ahn, A. Mirsky, M. Martino, and A. Barabasi, "A network framework of cultural history," *Science*, 2014, *345*, 558–562.

Serafinelli, M. and G. Tabellini, "Creativity over Time and Space," *Journal of Economic Growth*, 2022, *27*, 1–43.

Simon, Julian Lincoln, *Population and Development in Poor Countries: Selected Essays*, Princeton, NJ: Princeton University Press, 1992.

Strout, R., "The Development of the Catalog and Cataloging Codes," *The Library Quarterly: Information, Community, Policy*, 1956, *26* (4), 254–275.

Taylor, Arlene, "Authority Files in Online Catalogs," *Cataloging & Classification Quarterly*, 1984, *4* (3), 1–17.

Taylor, Arlene G. and Daniel N. Joudrey, *The Organization of Information*, 4th ed., Santa Barbara, CA: Libraries Unlimited, 2021.

Walsh, J., "Large sample validity of the binomial distribution for lives with unequal mortality rates," *Scandinavian Actuarial Journal*, 1952, *1-2*, 11–15.

Walsh, J., "Actuarial validity of the binomial distribution for large numbers of lives with small mortality probabilities," *Scandinavian Actuarial Journal*, 1956, *1*, 39–46.

Wiederhold, R. and G. Reeve, "Authority Control Today: Principles, Practices, and Trends," *Cataloging & Classification Quarterly*, 2021, *59*, 129–158.

Willison, I., "The National Library in Historical Perspective," *Libraries & Culture*, 1989, *24* (1), 75–95.

Wolverton, Jr. Robert E., "Authority Control in Academic Libraries in the United States: A Survey," *Cataloging Classification Quarterly*, 2005, *41* (1), 111–131.