Modern Library Holdings and Historic City Growth

Eric Chaney*

September 26, 2022

Abstract

This paper provides a new proxy for preindustrial economic activity by combining records from the world's largest library collections with historical Islamic biographies. A two-pronged statistical framework exploits variation in the number of authors affiliated with a city to derive a high-frequency measure of economic development, improve existing city growth estimates and generate preliminary growth estimates where none exist. To illustrate how the data allow for economic analysis by year instead of century, the paper revisits growth during the rise of the Atlantic trade and documents the enduring impact of the Mongol invasions.

*University of Oxford and CEPR. This project was inspired by a seminar comment by Robert J. Barro. Many individuals and seminar participants provided helpful discussions and comments. 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. I am responsible for all remaining errors.

Estimates of historical city sizes underpin our understanding of long-run economic development (e.g. De Long and Shleifer, 1993; Acemoglu et al., 2002b, 2005; Dittmar, 2011; Nunn and Qian, 2011; Bosker et al., 2013; Cantoni, 2015). These estimates suffer, however, from a number of well-known limitations. For example, they are geographically skewed towards Europe, available only at low frequencies and are measured with significant error. Consequently, existing estimates are inadequate both for the study of the historical development of non-European regions and for the analysis of the economic impact of many European events.

In this paper, I construct a database of historical authors using over 30 million records from the world's largest library collections and 28,000 Islamic biographies to help address these limitations. Panel A of Figure 1 shows existing population estimates for Baghdad (Iraq) as well as the smoothed number of author deaths from the new database between 800 and 1800 CE.[1] Panel B provides the same quantities for London (UK). In both panels the number of authors broadly tracks existing population estimates. This pattern is not limited to large or author-intensive cities. Rather, city and author growth are robustly correlated across the sample of cities in Europe and the Middle Eastern and North African (MENA) region.

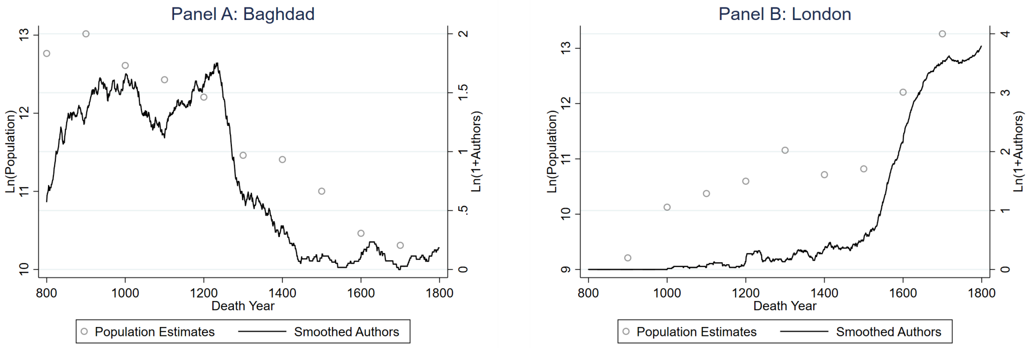

Figure 1: **Author Counts and Population Estimates of Baghdad and London 800-1800**

These empirical patterns suggest author counts can be used to help ad-

[1]Throughout, I use expressions such as "the number of author deaths" and "author counts" interchangeably. I do the same for the terms referring to city and economic growth.

dress the current lack of disaggregated high-frequency measures of historic economic activity. Nevertheless, exploiting author counts in this manner involves a few conceptual challenges. First, the number of observed author deaths is an imperfect measure of the underlying population of authors in a city. For example, observed author deaths may be contaminated by differences in the survival probabilities of works or in the age structure of scholarly populations. Second, the structural relationship between author and city populations may vary across cities or over time. Third, the distribution of author populations is skewed, leaving some cities with few author deaths.

To address these challenges, I develop a simple framework linking changes in author counts to city growth. In this model, I show that under a handful of arguably plausible assumptions difference-in-differences estimates using author deaths can be used to recover the sign of differential changes in economic activity across groups of cities by year.

In the second part of the conceptual framework, I explore how author deaths can be used to improve existing city growth data. Starting from Stein's Paradox, I derive a composite estimate of city growth from existing data and predicted growth based on growth in author counts. The validity of the composite rests on the assumptions of classical measurement error in existing city growth estimates and zero covariance between observed author count growth and this measurement error. I demonstrate that my composite estimate is numerically identical to that proposed by Henderson et al. (2012), linking their weighting to the broader Empirical Bayes/Machine Learning literatures and demonstrating how the analysis readily extends to more than one proxy (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).

Time-varying author survival or other sources of measurement error need not invalidate either framework. For example, although many historical authors are not in the data, the high-frequency analysis is not affected if the average observation probabilities evolve similarly across groups of cities. Similarly, because the validity of the composite estimate rests on the ability of author growth to predict city growth, selection, measurement error and variation in the underlying structural relationships will not invalidate the composite estimate inasmuch as they are not related to measurement error in existing city growth estimates.

The second part of the paper implements the conceptual framework. I first investigate the yearly evolution of author deaths in the Atlantic traders (Acemoglu et al., 2005) in the years after 1500, finding that these cities exhibit statistically faster growth by 1548. I then explore the development of author counts in the MENA region after the Mongol invasions of the 13th century. Cities conquered by the Mongols experience a differential drop in author counts following the invasions that becomes statistically significant in 1249 and persists until the end of the sample. To my knowledge, this represents the first systematic empirical analysis of this pivotal event in world history. Finally, I revise city growth estimates and derive preliminary city growth estimates where no estimates currently exist. This exercise is of particular value for the MENA region, where many historically important cities lack credible growth estimates.

The idea underpinning this paper is consistent with the work of a number of scholars who argue that collections of historical authors constitute one of the MENA region's best untapped sources of pre-industrial information (e.g. Bulliet, 1970; Romanov, 2017). While this literature demonstrates that variation in author counts can be used to shed light on past events, there are to my knowledge no other studies that document the correlation of author and city growth or that suggest the use of author counts as a proxy for historic economic activity.[2]

From a data construction standpoint, the paper adds to a growing list of studies using collections of individuals to enhance our understanding of the past (e.g. Schich et al., 2014; de la Croix and Licandro, 2015; Cummins, 2017; Gergaud et al., 2017; de la Croix et al., 2019; Mitchell, 2019; Serafinelli and Tabellini, 2022). Distinct from these studies, I construct a georeferenced database with significant non-European coverage that approximates the population of known authors working in Europe. These data are at least an order of magnitude larger than those previously used and open new avenues for research. For example, the economic impact of

[2]Gergaud et al. (2017) explore the relationship between city growth and notable people and do not find a robust relationship between the two although they do find a positive correlation between city growth and the number of entrepreneurs and artists.

events can be measured by year rather than century. Moreover, author counts are available in cities and time periods where alternative metrics of economic activity are unavailable. In this sense, the data assembled here provide a new proxy for historical development much as luminosity has done for the present (e.g. Henderson et al., 2012; Michalopoulos and Papaioannou, 2018).

As with any proxy, variation in author counts imperfectly captures the underlying growth process. For example, the data may reflect economic growth driven by trade or political considerations to a greater extent than other factors. Nonetheless, the robust correlation between city and author growth suggests that the relationship between author and economic growth is broad and that high-frequency variation is likely to be informative in many empirical environments.

Against this limitation, author counts have significant advantages over alternative proxies for historic development. For example, systematic information on historical authors is readily available across much of the world and far into the past. Competing proxies such as prices, wages, construction patterns or state finances (e.g. Allen et al., 2011; Buringh et al., 2020) are limited to relatively recent periods by documentary survival and/or are geographically skewed towards Western Europe. Author counts provide a widely available, regular measure of local economic activity that begins to address the current absence of such a metric.

The paper is organized as follows. Section 1 outlines the conceptual framework and section 2 describes the data. Section 3 presents the empirical applications. A final section concludes.

1 Author Counts and City Growth: a Conceptual Framework

Today, city sizes and input-based measures of innovation are positively related (e.g. Bettencourt et al., 2007; Carlino and Kerr, 2015). There are reasons to believe this relationship also held in the past. For example, sustained intellectual activity in preindustrial periods required access to agricultural surplus which flowed towards cities (e.g. Bairoch, 1988, p. 11).

Consequently, cities tended "to concentrate educational activities [...] and, if not research in the modern sense of the term, at least a certain systematic reflection" (Bairoch, 1988, p. 336). This rendered science, libraries and universities "urban phenomena" (Mokyr, 1993, p. 18).

Many models of intellectual and cultural production are consistent with the urban concentration of authors suggested by historical sources, modern scholarship (e.g. Gomez-Lievano et al., 2012) and empirically confirmed in section 2.1 of this paper. To maintain focus, however, I take this relationship as given to establish a framework that clarifies how variation in author counts can enhance our understanding of past economic activity. Questions such as the determinants of author mobility or of author productivity are consequently beyond this paper's scope.

The remainder of this section develops a simple theoretical framework that details conditions under which author counts can be used to enhance our understanding of historic economic development. First, I illustrate how the data can be used to provide a high-frequency measure of differential city growth and then derive the procedure to improve existing city growth estimates.

1.1 Deriving a High-Frequency Measure of City Growth from Author Counts

I begin by assuming that in each period authors independently die with probability p_i and enter this paper's dataset once dead with probability π_i. In other words, I assume the indicator X_i denoting the observed death of author i is a Bernoulli random variable with success probability $\gamma_i = p_i \pi_i$. The probabilities γ_i are assumed finite and relatively homogenous which implies, in turn, that as the number of authors grows, the distribution of observed author deaths from a population of C authors is well approximated by a normal distribution (e.g. Walsh, 1952) with mean $C\bar{\gamma}$, variance $C\bar{\gamma}(1-\bar{\gamma})$ and $\bar{\gamma} \equiv \frac{\sum_i \gamma_i}{C}$.

Let C_{jt} denote the number of authors working in city j at time t and d_j be an indicator variable indexing two groups of cities (with J_1 and J_0 constituent cities and $\sum_{d_j=1} C_{jt}$ and $\sum_{d_j=0} C_{jt}$ authors respectively).[3] Then

[3]For expositional simplicity I abstract from the fact that I measure an author's vital

$ln(\sum_{d_j=1} X_{ijt}) - ln(\sum_{d_j=1} X_{ijs}) - [ln(\sum_{d_j=0} X_{ijt}) - ln(\sum_{d_j=0} X_{ijs})]$ is also approximately normally distributed with mean

$$ln(\frac{\frac{1}{J_1}\sum_{d_j=1} C_{jt}}{\frac{1}{J_1}\sum_{d_j=1} C_{js}}) - ln(\frac{\frac{1}{J_0}\sum_{d_j=0} C_{jt}}{\frac{1}{J_0}\sum_{d_j=0} C_{js}}) + ln(\frac{\bar{\gamma}_{d_j=1,t}}{\bar{\gamma}_{d_j=1,s}}) - ln(\frac{\bar{\gamma}_{d_j=0,t}}{\bar{\gamma}_{d_j=0,s}}) \quad (1)$$

or the differential growth of the average number of authors working in cities in d_j between s and t plus the differential growth in the average observation probabilities over the same time period. For the remainder of this paper, I assume the second differential growth term is zero.[4]

If cities in d_j receive an absorbing and non-anticipated binary treatment at time g, equation 1 can be used to identify group-time average treatment effects (Callaway and Sant'Anna, 2021).[5] In particular, if $s = g-1$ equation 1 converges in probability to $ln\frac{E[C_{jt}(1)|d_j=1,t]}{E[C_{jt}(0)|d_j=1,t]} + \theta_t$ as the number of cities grows ($\theta_t \equiv ln\frac{E[C_{jt}(0)|d_j=1,t]}{E[C_{jt}(0)|d_j=1,g-1]} - ln\frac{E[C_{jt}(0)|d_j=0,t]}{E[C_{jt}(0)|d_j=0,g-1]}$). The assumption that $\theta_t = 0$ for $t \geq g$ leads to the convergence of equation 1 to the ATT expressed in percentage terms, adapting the parallel trends assumption based on a never-treated group in Callaway and Sant'Anna (2021) to this environment.

To relate this quantity to city growth, assume the relationship between the number of authors working in a city and its population P_{jt} is given by $C_{jt} = f_{jt}(P_{jt})$ where f_{jt} represents a continuous differentiable function that is allowed to vary by city and time with $f'_{jt}(P_{jt}) \geq 0\ \forall j, t$ (and $f'_{jt}(P_{jt}) \neq 0$ for some cities).[6] Then $ln\frac{E[C_{jt}(1)|d_j=1,t]}{E[C_{jt}(0)|d_j=1,t]}$ is approximately:

$$\frac{E[f'_{jt}(P_{jt}(0))(P_{jt}(1) - P_{jt}(0))|d_j = 1, t]}{E[f_{jt}(P_{jt}(0))|d_j = 1, t]} \quad (2)$$

statistics with error. Such measurement error, however, is easily incorporated in a standard manner and does not affect the conclusions.

[4]Note that $\bar{\gamma} = \sum_k \frac{c_k}{C}\gamma_k$, so this assumption would be violated if the average probability of death and/or the probability that an author enters the dataset conditional on death differentially changes in cities with $d_j = 1$. Sizeable and sustained changes in these quantities will be unlikely in many empirical environments.

[5]This setup is assumed for expositional ease and readily generalizes to a staggered treatment environment with multiple groups.

[6]Note that it is not necessary to assume that the relationship between author and city populations is stable over time as long as the parallel trends assumption $\theta_t = 0$ holds. For simplicity I abstract from situations in which the treatment affects the functional form of f_{jt} although such changes will not affect the conclusions as long as they are relatively small.

showing that it is possible to recover the sign of the city size ATT under relatively weak assumptions. For example, if $P_{jt}(1) - P_{jt}(0)$ is homogenous in sign or the treatment is unrelated to $f'_{jt}(P_{jt}(0))$, equation 2 and the city size ATT will have the same sign. More generally, author counts can be used to recover the sign of the city size ATT unless a sufficient number of author-intensive cities respond abnormally to a treatment.[7]

1.2 Improving Existing Growth Estimates: Stein's Paradox and Henderson et al. (2012)

Author counts can also be used to improve existing city growth estimates. To illustrate this point, begin by assuming that observed city growth z_j in city $j \in \{1, ..., k\}$ is measured with classical measurement error in the sense that $z_j = y_j + \epsilon_j$ where $\epsilon_j \overset{\text{i.i.d.}}{\sim} N(0, \sigma_z^2)$ and y_j, ϵ_j are independent.[8]

While the noisy z_j provide unbiased estimates of y_j, in a series of articles Stein (1955); James and Stein (1961); Stein (1962) showed that as long as $k \geq 4$ the James-Stein estimator $\hat{y_j} = \bar{z} + (1 - \frac{(k-3)\sigma_z^2}{\sum_j^k (z_j - \bar{z})^2})(z_j - \bar{z})$ is preferable to z_j in the sense that its mean-squared error (MSE) is strictly less than σ_z^2. This result, known as Stein's paradox, implies that a researcher can do better than using the individual z_j to estimate y_j as long as four or more city growth estimates exist.

To highlight the implications of this result, suppose noisy estimates of the growth rates of Algiers, Baghdad, Coimbra and Oxford between 1600 and 1700 exist that satisfy the James-Stein assumptions. Stein's paradox implies that new estimates obtained by "shrinking" the noisy estimates towards the average growth rate of the four cities will be more precise in an MSE sense than the original estimates. This result, called one of the most provocative in mathematical statistics in decades (Stigler, 1990), initially faced resistance. How, for example, could information on the growth rates of Algiers, Coimbra and Oxford be used to improve upon the unrelated

[7]Naturally, if one is willing to impose further structure it is possible to recover the relevant magnitudes. The general approach adopted in this paper, however, has the advantage of recovering the relevant sign of the treatment under assumptions that are arguably satisfied in many empirical environments.

[8]The normality assumption is for expositional ease and will be relaxed below. The assumptions on ϵ_j (minus normality) are what Schennach (2016) calls strongly classical measurement error.

problem of estimating the growth rate of Baghdad?

Efron and Morris (1973, 1977) provided greater intuition for this seemingly paradoxical result by noting the link between Stein's paradox and the Empirical Bayes approach. In particular, as k grows (I henceforth assume a large number of growth rates and that the mean of y_j, $E[y]$ and its variance σ_y^2 exist), $\hat{y_j}$ converges to Bayes's equation:

$$E[y] + \lambda_y(z_j - E[y]) \quad (3)$$

where $\lambda_y \equiv \frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2}$. Under the additional assumption that y_j is itself normally distributed, its distribution $N(E[y], \sigma_y^2)$ can be viewed as the prior distribution, $z_j|y_j \sim N(y_j, \sigma_z^2)$ the likelihood and equation 3 can be interpreted as the empirical Bayes estimator (or the mean of the posterior distribution). This result helped provide intuition for Stein's paradox giving a wide-sense extension of the James-Stein estimator to non-normal cases in which attention is restricted to composite estimates that are a linear function of z_j (Efron and Morris, 1973).

If a variable x_j exists that is independent of ϵ_j then this analysis applies conditional on x_j.[9] In particular, shrinking z_j towards $E[y|x]$ with $\lambda^*_{y|x} \equiv \frac{\sigma^2_{y|x}}{\sigma_z^2+\sigma^2_{y|x}}$ yields $MSE(\lambda^*_{y|x}|x) = \frac{\sigma^2_{y|x}\sigma_z^2}{\sigma^2_{y|x}+\sigma_z^2}$. If the conditional mean is instead approximated by the regression fitted values ψx_j, the conditional variance by $E[(y_j - \psi x_j)^2|x]$ and $\lambda^*_{y|x}$ by $\lambda_{y|x} \equiv \frac{E[(y_j-\psi x_j)^2|x]}{\sigma_z^2+E[(y_j-\psi x_j)^2|x]}$ then $MSE(\lambda_{y|x}|x) = \frac{E[(y_j-\psi x_j)^2|x]\sigma_z^2}{E[(y_j-\psi x_j)^2|x]+\sigma_z^2} \geq MSE(\lambda^*_{y|x}|x)$.

The total mean squared error associated with $\lambda_{y|x}$ is $E[MSE(\lambda_{y|x}|x)]$. Using Jensen's inequality it follows that:

$$\sigma_z^2 \geq MSE(\lambda_y) \geq MSE(\lambda) \geq E[MSE(\lambda_{y|x}|x)] \geq E[MSE(\lambda^*_{y|x}|x)] \quad (4)$$

where $MSE(\lambda) = \frac{E[(y_j-\psi x_j)^2]\sigma_z^2}{E[(y_j-\psi x_j)^2]+\sigma_z^2}$ or the MSE obtained by shrinking towards the fitted values ψx_j using the constant shrinkage factor $\lambda \equiv \frac{E[(y_j-\psi x_j)^2]}{E[(y_j-\psi x_j)^2]+\sigma_z^2}$.

Equation 4 suggests viewing the estimator based on λ as an approxima-

[9]Note that this analysis extends to a vector of variables $\boldsymbol{x_j}$ under the assumption that $\boldsymbol{x_j} \perp\!\!\!\perp \epsilon_j$, providing a framework in which multiple proxies can be combined to improve upon existing estimates (e.g. Fay III and Herriot, 1979).

tion to the conditional James-Stein estimator. From a Bayesian perspective this equation can be understood as providing an ordering of the quality of the initial guesses $E[y|x]$, ψx_j and $E[y]$. The z_j will be pulled further towards the initial prior as this guess improves (the (conditional) variance of y decreases) leading to greater MSE reduction. This reduction will also be greater if heteroscedasticity is correctly accounted for.

Equation 4 demonstrates that the shrinkage estimator based on λ does no worse than that based on λ_y. Because this inequality also holds when $cov(x_j, \epsilon_j) = 0$, the independence assumption can be weakened to one of zero covariance.[10]

It is straightforward to demonstrate that the optimal weighting $\frac{\frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2}-\rho_{xz}^2}{1-\rho_{xz}^2}$ in equation 11 of Henderson et al. (2012) is numerically identical to λ.[11] Interpreting this weighting as an approximation to the conditional James-Stein estimator, however, eliminates the need for the linearity assumption made by Henderson et al. (2012, p. 1005).[12] This nonparametric interpretation also provides greater intuition as to why the validity of their approach does not depend on the ability of the researcher to identify a group of related cities or to correctly specify the relationship between x_j and y_j. Finally, linking their weighting to the broader Empirical Bayes/Machine Learning literatures suggests a simple extension of the analysis to environments with multiple proxies (e.g. Fay III and Herriot, 1979).

2 Data

I use two types of data to estimate the number of author deaths by year (see the Appendix for an in-depth discussion of the data construction). The first derive from author records contained in the Virtual International

[10]When $cov(x_j, \epsilon_j) = 0$ but the two variables are not independent the composite estimate can be interpreted as an approximation to the infeasible conditional James-Stein estimator.

[11]Note that $E[(y_j - \psi x_j)^2] = \sigma_y^2(1 - \rho_{xy}^2)$ and $\rho_{xz}^2 = \frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2}\rho_{xy}^2$ which implies $\frac{E[(y_j-\psi x_j)^2]}{\sigma_z^2+E[(y_j-\psi x_j)^2]} = \frac{\frac{\sigma_y^2}{\sigma_y^2+\sigma_z^2}-\rho_{xz}^2}{1-\rho_{xz}^2}$.

[12]The assumptions made here are otherwise identical to those in Henderson et al. (2012) since Henderson et al. (2012, pp. 1005, 1007) assume classical measurement error for ϵ_j, that $cov(x_j, \epsilon_j) = 0$ and restrict consideration to composite estimates that are a linear function of z_j.

Authority File (VIAF) database (`www.viaf.org`). This database, which provides information on over 30 million authors, can be viewed as containing the union of authors whose works are held in the world's largest national libraries (e.g. The Library of Congress, The Bibliothèque Nationale de France or the National Library of Morocco). For Europe in the print era, I estimate that these data contain 81% of all authors whose works have survived until today. Such evidence, combined with the European focus and concentration of the participating libraries, suggest the VIAF data provide a reasonable approximation to the population of known European authors.

Authors working in the MENA region not covered by the VIAF database are drawn from 31 of the region's most well-known historical biographic dictionaries. These biographical dictionaries are of Islamic origin (e.g. Gibb, 1962; Makdisi, 1993) and were composed in many cities across the region. Although the reasons behind the emergence of this genre and its historical popularity remain topics of ongoing research, the information contained in these sources is both reliable and quantitatively significant (e.g. Petry, 1981, p. 8). When I combine the 28,477 authors in my sample of historical dictionaries with the VIAF data, the resulting database of 36,049 MENA authors represents to my knowledge the largest ever assembled. I estimate that these data contain approximately 30% of all historical MENA authors.

To georeference authors, I use a variety of sources such as online library catalogues, author names and Wikipedia. The Appendix demonstrates that my approach georeferences approximately 92% of all authors that could be referenced using online sources.

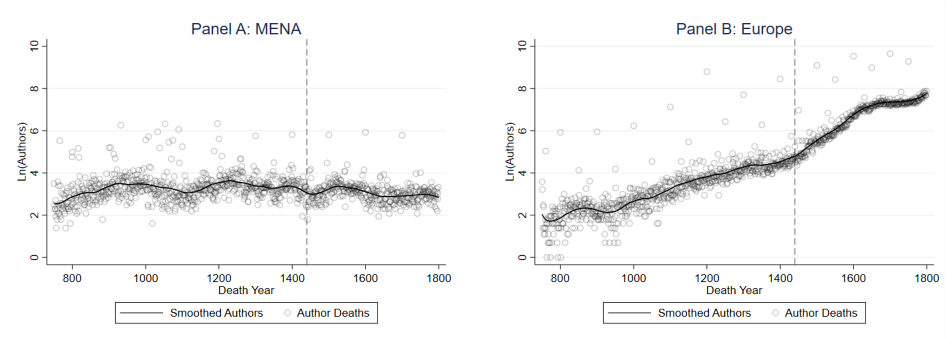

Figure 2: **Author counts 750-1800**

Figure 2 illustrates the evolution of the 453,487 European and 34,051 MENA author deaths that occurred on the interval [750,1800). Panel A details the logarithm of the number of author deaths in the MENA region and panel B presents the data for Europe.[13] The introduction of printing in Europe, marked by vertical lines in both panels, is one potential disruption to the relationship between city and author growth. This event, however, does not appear to have discretely affected the evolution of author counts over time. The jumps in the number of observations around round dates such as centuries reflect the approximate nature of many recorded death dates.

After constructing the data, I limit attention to the 196,945 georeferenced authors who died on the interval [750,1800) in 1,055 MENA and European cities contained in an augmented version of the union of the Bairoch et al. (1988), Bosker et al. (2013) and Chandler (1987) databases.[14] These include 982 cities with at least one estimated growth rate as well as 73 new MENA cities for which no growth estimates exist.[15] Table 1 reports the 10 largest cities as measured by author counts for each century, where MENA cities are highlighted. The shift from MENA to European leadership over time tracks well-known historical patterns (e.g. Bosker et al., 2013) despite the fact that author deaths imperfectly proxy for city size.

One implication of this paper's conceptual framework is that the evolution of author counts over time should reflect shocks to a city's population in the absence of abrupt changes in the probability of observing author deaths or in the structural relationship between author and city populations. Figure 3 provides graphical evidence consistent with this prediction for four selected cities.

In panel A, I plot the evolution of the logarithm of one plus author counts against time in Merv, Turkmenistan. In this panel the black line

[13]Note that the decrease in the variance of the logarithm of the total number of author deaths as the number of authors increase is predicted by the conceptual framework. In particular, this distribution is well approximated by $N(ln(C_t) + ln(\bar{\gamma_t}), \frac{(1-\bar{\gamma_t})}{C_t\bar{\gamma_t}})$.

[14]Competing city size estimates exist (e.g. de Vries, 1984; Buringh, 2021), although I focus on the Bairoch et al. (1988), Bosker et al. (2013) and Chandler (1987) data for comparability with previous studies.

[15]Note that the growth regressions are limited to the 982 cities with 2,750 existing growth estimates. The difference-in-differences analysis, however, draws from all 1,055 cities.

Table 1: **The 10 Largest Cities by Author Counts, 800-1800**

800		900		1000		1100		1200	
Kufa	400	**Cordoba**	439	**Cordoba**	471	**Baghdad**	342	**Baghdad**	466
Basra	231	**Baghdad**	431	**Baghdad**	408	**Nishapur**	155	**Damascus**	406
Baghdad	197	**Basra**	126	**Nishapur**	241	**Kufa**	149	**Cairo**	258
Medina	110	**Cairo**	111	**Cairo**	203	**Isfahan**	114	Paris	139
Cordoba	75	**Nishapur**	83	**Isfahan**	84	**Cordoba**	91	**Mosul**	119
Cairo	63	**Kufa**	78	**Rayy**	84	Constantinople	88	**Aleppo**	92
Mecca	61	**Mecca**	68	**Bukhara**	75	**Cairo**	82	**Isfahan**	86
Damascus	59	**Rayy**	64	**Damascus**	68	**Damascus**	70	Rome	83
Constantinople	56	**Damascus**	57	**Basra**	56	**Seville**	56	**Jerusalem**	81
Rayy	31	**Toledo**	54	**Qayrawan**	56	Rome	56	Palermo	76
1300		1400		1500		1600		1700	
Damascus	543	**Damascus**	369	Nuremburg	682	Paris	1802	Paris	4101
Cairo	315	**Cairo**	365	Venice	657	Nuremburg	1690	London	3484
Paris	214	Florence	267	Paris	590	London	1583	Leipzig	2121
Bologna	139	Nuremburg	215	Florence	579	Leipzig	1400	Nuremburg	2041
Florence	124	Venice	197	Rome	544	Rome	1293	Amsterdam	1828
Baghdad	124	Paris	180	Augsburg	324	Venice	1130	Rome	1485
Jerusalem	94	**Mecca**	169	**Damascus**	281	Wittenberg	826	Utrecht	1028
Aleppo	93	Rome	144	**Cairo**	270	Augsburg	684	Leiden	1014
Rome	92	Bologna	132	Milan	268	Antwerp	619	Hamburg	977
Constantinople	81	Prague	110	Bologna	248	Rostock	591	Jena	970

Notes: city names are followed by the number of authors who died on the interval $[t-50, t+50)$. MENA cities are in bold.

denotes smoothed deaths (using a rear moving average of 50 years) and the vertical line marks the Mongol sack of the city in 1221. The drop in author counts is consistent with historical evidence that the city did not recover from this event.[16] Panel B provides a similar graph for Magdeburg, Germany. In this panel the vertical line in 1631 marks both a drop in observed author deaths and the sack of the city (McLeod, 2001).

The impact of the Ottoman conquest of Tlemcen, Algeria in 1551 can clearly be seen in panel C, providing support for claims that the conquest negatively affected Tlemcen's economic prosperity (e.g. Lawless and Blake, 1976, p. 59).[17] Finally, panel D provides author counts for Avignon, France. In this panel, variation is consistent with historical evidence that the movement of the papacy from Rome to Avignon in 1309 quickly transformed the "village into a vibrant cosmopolitan city" (Rollo-Koster, 2015, p. 188). Similarly, author counts decline rapidly following the deposition of the Avignon pope in 1409.

[16]For a brief overview of the history of Merv, see `https://www.theguardian.com/cities/2016/aug/12/lost-cities-merv-worlds-biggest-city-razed-turkmenistan`.

[17]Tlemcen was one of the major northern 'ports' of the trans-Saharan trade (e.g. Niane, ed, 1997, p. 245).

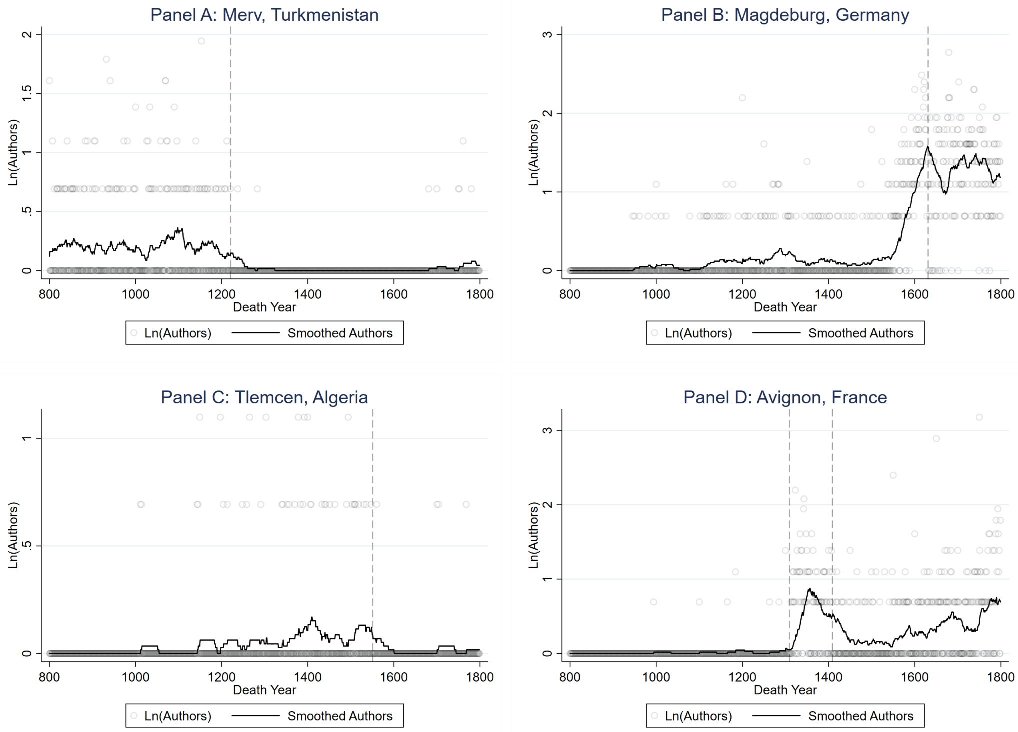

Figure 3: **Raw Author Counts, Smoothed Values and Selected Events**
Vertical lines denote Mongol Sack (panel A), Magdeburg Massacre (panel B), Ottoman Conquest (panel C), and Avignon Papacy (panel D)

2.1 Correlations with Existing City Growth Data

The link between city-specific shocks and high-frequency variation in author counts suggests city and author growth are correlated. Figure 4 confirms this by graphing the relationship between existing city growth estimates and growth in author counts net of century dummies.[18] In this figure, panels A and B present the relationship for MENA and European cities. The point estimates are highly significant in both regions despite the fact that the MENA sample size is limited by relatively few existing city growth estimates.

To explore how this relationship varies across time, I pool the data and

[18]Throughout, I use log differences to estimate the city growth rate. Author growth is the difference between the logarithm of one plus the sum of all authors who died on the 100 year interval [t-50,t+50) and the logarithm of one plus the sum of those who died on the interval [t-150,t-50). In the Appendix I justify this procedure and demonstrate that the results are robust to varying this window.

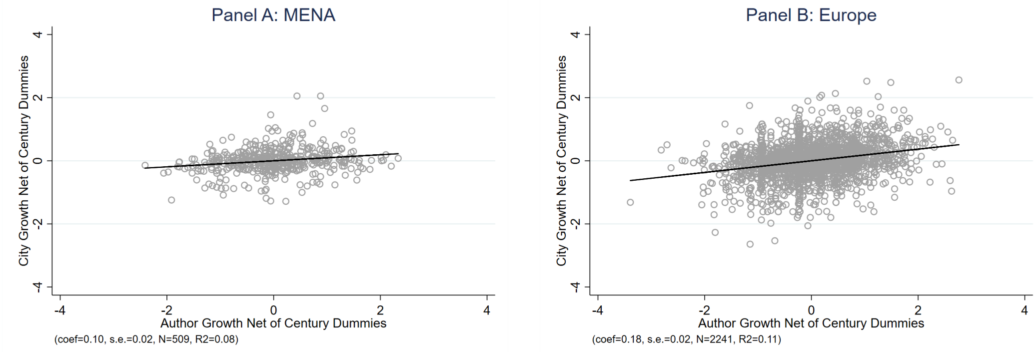

Figure 4: **Author Counts and Population Estimates of Baghdad and London 800-1800**

estimate the regression

$$z_{jt} = \alpha_t + \psi_t x_{jt} + \xi_{jt} \tag{5}$$

where z_{jt} denotes existing measures of city growth and x_{jt} denotes author growth in city j and century t.[19] The relationship between author and city growth, presented in columns (1)-(9) of Table 2, is always statistically significant after 900, with slope coefficients ranging between 0.09 and 0.29. The point estimates are also economically significant, becoming similar in magnitude to the relationship between GDP and night lights (Henderson et al., 2012) by the end of the sample.

Statistically, the point estimates are stable over time with the exception of $\hat{\psi}_{1700}$. This break, which is limited to European cities, entirely drives the difference between MENA and European slope coefficients in Figure 4. Prior to 1700, the $\hat{\psi}_t$ for Europe and the MENA regions are statistically indistinguishable. Because such a break does not necessarily invalidate the assumptions made in the conceptual framework, I leave a detailed exploration of this result—which may be related to intellectual developments during this period— to future research.

Is the relationship between author and city growth limited to large or author-intensive cities? Results in columns 10 and 11 examine this question by constraining the slope coefficient in equation 5 to be constant. In column

[19]Because Gibrat's law implies that city growth is i.i.d. (e.g. Eeckhout, 2004), there is little reason to worry about spatial correlation in the error term. For this reason, I present heteroscedasticity-robust standard errors throughout unless otherwise noted.

Table 2: **Relationship Between City and Author Growth**

	Δln(Population)										
	900	1000	1100	1200	1300	1400	1500	1600	1700	Small	Not Int.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
$\Delta ln(Authors)$	0.10	0.11*	0.09**	0.15***	0.14***	0.14***	0.12***	0.13***	0.29***	0.18***	0.13***
	(0.06)	(0.06)	(0.04)	(0.03)	(0.05)	(0.05)	(0.03)	(0.02)	(0.03)	(0.02)	(0.02)
Constant	0.11***	0.20***	0.02	0.03	0.13***	-0.10***	0.08***	0.15***	-0.05***	-	-
	(0.04)	(0.04)	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)		
N	64	91	130	141	207	357	366	556	838	1165	1259
Century Dummies?	-	-	-	-	-	-	-	-	-	Yes	Yes

Notes: the total number of observations in the regression reported in columns (1)-(9) is 2,750. In these columns, N denotes the number of observations in each century. In column 10 the regression is restricted to small cities (those whose population in t-1 is under 10,000) and in column 11 to non-author-intensive cities (those whose ratio of scholars to population in t-1 is under 0.5 authors per thousand inhabitants.) ***, **and *denote significance at the 1%, 5% and 10% levels.

10, I limit the sample to cities with baseline populations less than 10,000 and in column 11 to those cities with below-median baseline author-to-population ratios. In both columns, the estimated slope coefficient remains highly significant, suggesting the generality of the relationship between author and city growth.

3 Empirical Applications

Under a few assumptions (maintained throughout this section) author counts provide a high-frequency measure of historical economic development that can be used to improve existing city growth estimates. This section illustrates these points empirically. First, I examine the impact of the Atlantic trade and Mongol invasions by year. Second, I use author data to revise existing city growth data.

3.1 High-Frequency Analysis: Atlantic Traders and Mongol Invasions

3.1.1 The Atlantic Traders

Acemoglu et al. (2005) document that the Atlantic traders –Britain, France, the Netherlands, Portugal and Spain— grew more rapidly than the remainder of Western Europe after 1500 CE. Conceptually, these authors link this differential growth to the institutional activities of commercial interests strengthened by the Atlantic trade. Empirically, their analysis is limited by both the intermittent and low-frequency nature of the Bairoch et al. (1988) population estimates.

I revisit the evolution of economic activity in the Atlantic traders on the interval [1300,1800) using author counts from all Western European cities in my sample.[20] These include cities such as Nantes or Cádiz that Acemoglu et al. (2005) were forced to omit from their balanced panel due to missing data. Panel A of Figure 5 presents the estimated differential growth in observed author deaths $ln(1 + \sum_{j \in Atl} X_{ijt}) - ln(1 + \sum_{j \in Atl} X_{ij,1500}) - [ln(1+\sum_{j \notin Atl} X_{ijt}) - ln(1+\sum_{j \notin Atl} X_{ij,1500})]$ over time.[21] Panel B, in turn, graphs the 25 year rear moving average of these point estimates along with a 99% confidence interval obtained using a non-parametric bootstrap with 100 replications.

The moving average in panel B, which approximates the differential author growth of the Atlantic traders by year, is broadly consistent with the original study. In particular, the moving average becomes statistically different from zero beginning in 1548 and remains so through the end of the sample. Although the pre-trend is at times statistically significant, it is difficult to interpret given that the small number of author deaths in the Atlantic traders prior to 1500 violate the maintained theoretical assumptions.[22]

3.1.2 The Mongol Invasions

The Mongol invasions of the thirteenth century helped shape the economic trajectory of the eastern MENA region. Although the initial impact of these invasions is generally believed to have been large, their long-run impact remains a topic of scholarly debate. For example, while traditional scholarship attributed the region's economic decline to this event more recent

[20] I follow Acemoglu et al. (2002a, p. 2) and define Western European cities as those within the modern-day boundaries of Austria, Belgium, Britain, Denmark, Finland, France, Germany, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland.

[21] I add the average of the differential growth in author counts on the interval [1475,1500) to the series to correct for noise in the estimated values of $E[ln(X_{ij,1500})|Atl]$.

[22] In particular, the fact that the Atlantic traders are often at the zero lower bound of author deaths prior to 1500 works to create a spurious pre-trend. Relatedly, a recent literature (e.g. Roth, forthcoming) highlights problems involved with such pre-trend tests and suggests alternative approaches to avoid the pre-testing problem. For comparison with existing studies I present pre-trends while noting that the magnitude of the effects are such that violations of the parallel trends assumption are unlikely to entirely explain the results.

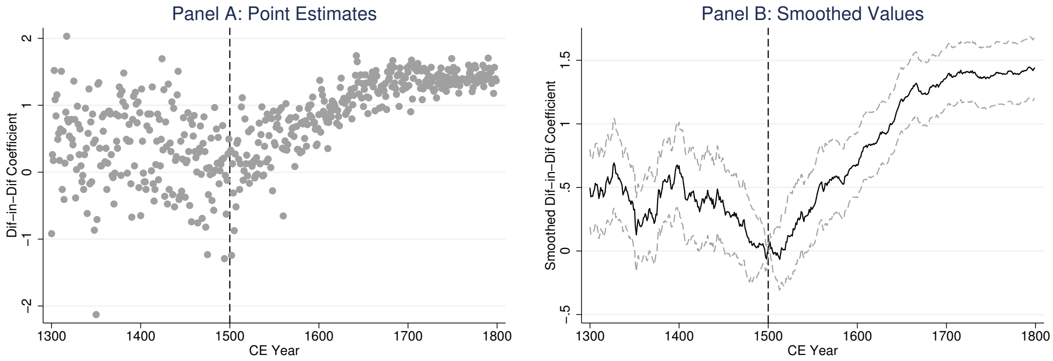

Figure 5: **Atlantic Traders**

studies have stressed "that the harmful effects of the Mongol conquests were not as great, as lasting, or even as extensive as was once thought" (Lewis, 1968, p. 52). The absence of systematic empirical evidence has made it difficult to conclusively differentiate between these conflicting views.

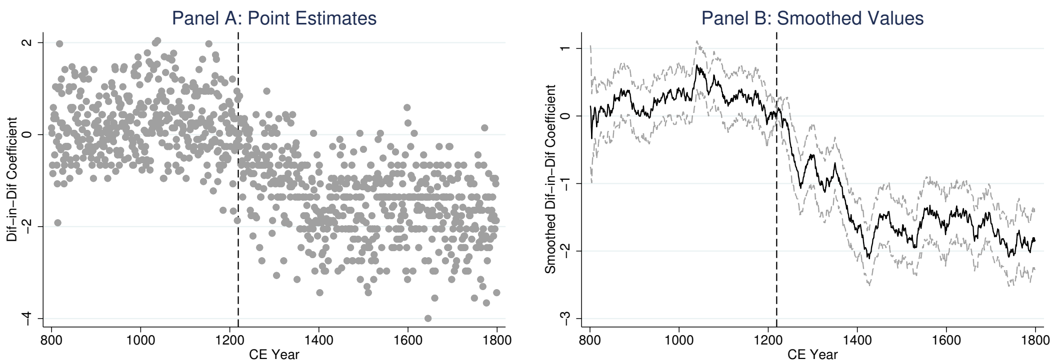

Figure 6: **Mongol Invasions**

Figure 6 uses author counts to provide what is to my knowledge the first systematic empirical investigation of the Mongol invasions.[23] As in Figure 5, the left-hand graph provides the yearly difference-in-differences estimates and the right-hand graph presents the associated moving average. In this figure, the vertical line corresponds to the start of the Mongol invasions of the eastern MENA region in 1219.[24]

[23]This figure compares cities conquered by the Mongols in Iraq and eastwards to never-treated cities in North Africa and the Arabian Peninsula. See the Appendix for details.

[24]Although cities were conquered at different point in time, the exact date a city is

The impact of the invasions appears immediately after 1219, becoming statistically significant by 1249 and continuing to deepen until approximately 1400. The estimated effects are both economically and statistically significant. For example, the point estimate of -1.84 in 1799 implies that the invasions led to a 84% decrease in the average number of authors working in Mongol-conquest cities relative to the counterfactual of no Mongol invasion. Although this result weighs against recent work minimizing the role of the Mongol conquests in the region's long-run economic decline (e.g. Toonen et al., 2020) and provides support for those continuing to espouse the traditional view (e.g. Kuru, 2019), it may mask significant heterogeneity. For example, Lewis (1968, p. 54) hypothesizes that the Mongol conquests had more long-lasting effects in modern-day Iraq than in other areas through the destruction of irrigation networks.[25] Other interpretations of this result are clearly possible and highlight the need for further research.

3.2 Improving City Growth Estimates

I apply the Henderson et al. (2012) weighting to the Bosker et al. (2013) estimates of Seville, Spain to illustrate how author counts can be used to improve existing growth estimates. Seville, under Islamic rule until 1248 CE, is one of a handful of Atlantic trader ports that appear in both the MENA and European data. In addition, there is well-documented uncertainty regarding its historical population.

Table 3 details the relevant calculations. Columns 2 and 3 provide the Bosker et al. (2013) estimates of Seville's population in thousands and the number of author deaths while columns 4 and 5 give the associated growth rates. As in previous examples, population and author counts broadly track one another, rising during the Islamic period and again in the century after 1500. The similar evolution of these two series, however, at times breaks down. For example, Bosker et al. (2013) estimate that between 1400 and 1500 Seville's population decreased by 44 log points whereas over

treated is often unclear since populations may have anticipated conquest. Given this complication, I assume a common treatment date of 1219 and note that this likely biases the estimated effects towards zero.

[25]See Heldring et al. (2020) for a discussion of the long-run evolution of irrigation infrastructure in Iraq.

Table 3: **Comparison of Existing and New Growth Rates: Seville, Spain**

(1) Year	(2) Pop. Level	(3) Author Level	(4) Pop. Growth	(5) Author Growth	(6) $\hat{\alpha} + \hat{\psi}x$	(7) $\hat{\lambda}$	(8) $\hat{y}$	(9) New Level	(10) de Vries
800	30	7	-	-	-	-	-	32	-
900	40	30	0.29	1.46	0.26	0.77	0.28	42	-
1000	90	44	0.81	0.38	0.24	0.77	0.68	83	-
1100	85	57	-0.06	0.26	0.04	0.77	-0.03	81	-
1200	80	64	-0.06	0.12	0.05	0.76	-0.03	78	-
1300	90	30	0.12	-0.76	0.02	0.77	0.10	86	-
1400	70	14	-0.25	-0.76	-0.21	0.77	-0.24	67	-
1500	45	86	-0.44	1.82	0.30	0.77	-0.27	51	25
1600	135	234	1.10	1.00	0.28	0.77	0.91	128	90
1700	72	146	-0.63	-0.47	-0.19	0.87	-0.57	72	96

Notes: column 2 provides the Bosker et al. (2013) population estimates and column 3 author counts. Columns 4 details the population growth rates and column 5 the growth rate of one plus the number of authors. Column 6 reports the fitted values obtained from author growth and the regression coefficients in Table 2 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 provides the city level time path assuming the Bosker et al. (2013) 1700 estimate is correct. Column 10 provides the de Vries (1984) alternative population estimates.

this period author counts increased by 182 log points.

The following three columns demonstrate how the composite estimate combines columns 4 and 5. Column 6 provides the fitted values obtained from the coefficients in Table 2 and the author growth estimates in column 5. The advantages of the Bayesian interpretation proposed in this paper become apparent at this point. As noted above, there is evidence that the underlying structural relationship between author and city growth is different for MENA cities in 1700. The structural relationship may also be non-linear or vary by city characteristics. The Bayesian view makes clear that while incorporating such information will lead to a better "prior" and larger MSE savings, ignoring it and shrinking towards the values in column 6 remains a valid empirical strategy.

Column 7 estimates the shrinkage parameter λ, or the weight that the Bosker et al. (2013) growth estimate receives in the composite. This was calculated using estimates of the signal-to-noise ratio obtained from the Bosker et al. (2013) and de Vries (1984) data where, to be conservative, I assume this ratio is constant prior to 1600.[26] Column 8 provides the new

[26] Regressing the Bosker et al. (2013) growth estimates on those given by de Vries (1984) yields reliability ratios of 0.89 for 1700 and 0.78 for 1600. These estimates of the

growth estimates.

For the year 1500, the composite estimate of -0.27 weights the existing growth estimate of -0.44 by 0.77 and the predicted growth rate of 0.30 by 0.23. Is this new estimate an improvement over the Bosker et al. (2013) estimate? Although in general it is impossible to determine if an individual composite estimate is superior to the original, in this case there is historical evidence to support this upward revision. The original Bosker et al. (2013) growth rate relies on population estimates for 1400 and 1500 provided by Chandler (1987) and Carande (1965). Over the past decades, new research has improved our understanding of Seville's population evolution. For example, Flores Varela (2005, p. 99) provides data suggesting that between 1405 and 1486 Seville's population grew by 30 log points. This is identical to the fitted value of 0.30 in column 6, suggesting that the revised estimate is preferable to the Bosker et al. (2013) estimate.

Column 9 provides revised estimates of Seville's population obtained by assuming the Bosker et al. (2013) population level for 1700 is correct and iterating backwards using column 8. Column (10) gives alternative Seville population estimates proposed by de Vries (1984). The data patterns weigh against adopting these revised estimates. For example, the population increase suggested by de Vries (1984) between 1600 and 1700 is at variance with the drop in author counts observed over this period.

I perform identical calculations to those detailed in Table 3 for the remaining growth estimates. In addition, I use the fitted values derived from author growth rates and the point estimates in Table 2 to derive preliminary estimates for the 6,088 missing growth rates in the 982 cities in the growth regressions. I do the same for the 657 missing growth rates in the additional 73 MENA cities where no growth estimates currently exist.[27] The resulting balanced panel is an improvement in both temporal and geographic scope over existing data.

signal to noise are likely too large given that Bosker et al. (2013) (which heavily relies on Bairoch et al. (1988)) and de Vries (1984) at times rely on similar information to derive their growth estimates. This error, however, will tend bias the new estimates towards existing growth estimates. ρ_{xz} is estimated from the data.

[27]In these calculations I abstract from the founding as well as the permanent abandoning of cities.

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 high-frequency measure of urban activity where traditional data are unavailable. I demonstrate that the statistical framework in Henderson et al. (2012) has a non-parametric interpretation rooted in Stein's Paradox, linking their approach to the broader Empirical Bayes and Machine Learning literatures.

To illustrate how high-frequency variation in author counts deepens our understanding of historical economic activity, I examine the yearly impact of two events. First, I document the differential growth of author counts in cities most exposed to the Atlantic trade after 1500. Then I provide new evidence that the Mongol invasions had a long-lasting impact on the Eastern MENA region.

As in Henderson et al. (2012), I shrink existing city growth estimates towards regression fitted values of city on author growth, generating a new set of city growth estimates. I also use the estimated coefficients to derive preliminary city growth estimates where none exist. I note that the Bayesian interpretation proposed in this paper readily generalizes to situations in which more than one proxy for city growth exists.

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 high frequency. These advances are general but are of particular import for MENA cities where many historically important cities lack credible population estimates.

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 to other non-European regions will further improve our understanding of historic economic activity, particularly where traditional data are unavailable.

References

Acemoglu, D., S. Johnson, and J. Robinson, "The Rise of Europe: Atlantic Trade, Institutional Change and Economic Growth," *NBER Working Paper 9378*, 2002.

__ , __ , and __ , "The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth," *American Economic Review*, 2005, *95* (3), 546–579.

__ , __ , 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.

Allen, R., J.-P. Bassino, D. Ma, C. Moll-Murata, and J.L. van Zanden, "Wages, prices, and living standards in China, 1738-1925: in comparison with Europe, Japan, and India," *Economic History Review*, 2011, *64* (S1), 8–38.

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.

Bairoch, P., *Cities and Economic Development*, Chicago: University of Chicago Press, 1988.

__ , 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.

Bettencourt, L., J. Lobo, D. Helbing, C. Kuhnert, and G. West, "Growth, Innovation, Scaling, and the Pace of Life in Cities," *Proceedings of the National Acadmey of Sciences*, 2007, *104* (17), 7301–7306.

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.

Bulliet, R., "A Quantitative Approach to Medieval Muslim Biographical Dictionaries," *Journal of the Economic and Social History of the Orient*, 1970, *13* (2), 195–211.

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.

_, B. Campbell, A. Rijpma, and J.L. van Zanden, "Church building and the economy during Europe's 'Age of the Cathedrals', 700-1500 CE," *Explorations in Economic History*, 2020, *76.*

Callaway, B. and P. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cantoni, D., "The Economic Effects of the Protestant Reformation: Testing the Weber Hypothesis in the German Lands," *Journal of the European Economic Association*, 2015, *13* (4), 561–598.

Carande, R., *Carlos V y sus Banqueros*, Madrid: Sociedad de Estudios y Publicaciones, 1965.

Carlino, G. and W.R. Kerr, "Agglomeration and innovation," in G. Duranton, J.V. Henderson, and W. Strange, eds., *Handbook of Regional and Urban Economics, Vol. 5*, Oxford: Elsevier, 2015, pp. 349–404.

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

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.

Cummins, N., "Lifespans of the European Elite, 800-1800," *Journal of Economic History*, 2017, *77* (2), 406–439.

de la Croix, D. and O. Licandro, "The longevity of famous people from Hammurabi to Einstein," *Journal of Economic Growth*, 2015, *20*, 263–303.

— , **F. Docquier, A. Fabre, and R. Stelter**, "The Academic Market and the Rise of Universities in Medieval and Early Modern Europe (1000-1800)," *Working Paper*, 2019.

De Long, J.B. and A. Shleifer, "Princes and Merchants: European City Growth before the Industrial Revolution," *The Journal of Law & Economics*, 1993, *36* (2), 671–702.

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

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

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.

— **and** — , "Stein's Paradox in Statistics," *Scientific American*, 1977, *236* (5), 119–127.

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.

Flores Varela, C., "La evolución de la población urbana de Andalucía en los siglos XV y XVI," *En la España Medieval*, 2005, *28*, 97–125.

Gergaud, O., M. Laouenan, and E. Wasmer, "A Brief History of Human Time," *Working Paper*, 2017.

Gibb, H., "Islamic Biographical Literature," in B. Lewis and P.M. Holt, eds., *Historians of the Middle East*, London: Oxford University Press, 1962, pp. 54–58.

Gomez-Lievano, A., H. Youn, and L.M.A Bettencourt, "The Statistics of Urban Scaling and Their Connection to Zipf's Law," *PLoS ONE*, 2012, *7* (7).

Heldring, L., R. Allen, and M. Bertazzini, "The Economic Origins of Government," *Working Paper*, 2020.

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

James, W. and C. Stein, "Estimation with Quadratic Loss," in "Pr oceedings of the Fourth Berkeley Symposium," University of California Press, 1961, pp. 361–379.

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

Kuru, A., *Islam, authoritarianism, and underdevelopment*, Cambridge: Cambridge University Press, 2019.

Lawless, R. and G. Blake, *Tlemcen: Continuity and Change in an Algerian Islamic Town*, London: Bowker, 1976.

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

Makdisi, G., ""Tabaqat"-Biography: Law and Orthodoxy in Classical Islam," *Islamic Studies*, 1993, *32* (4), 371–396.

McLeod, T., "Magdeburg, sack of (1631)," in R. Holmes, C. Singleton, and S. Jones, eds., *The Oxford Companion to Military History*, Oxford: Oxford University Press, 2001.

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

Mitchell, S., "London calling? Agglomeration economies in literature since 1700," *Journal of Urban Economics*, 2019, *112*, 16–32.

Mokyr, J., "Urbanization, Technological Progress, and Economic History," in H. Giersch, ed., *Urban agglomeration and economic growth*, Berlin: Springer-Verlag, 1993, pp. 3–38.

Niane, D., ed., in D. Niane, ed., *Africa from the Twelfth to the Sixteenth Century, General History of Africa*, Berkeley: Univerity of California Press, 1997.

Nunn, N. and N. Qian, "The Potato's Contribution to Population and Urbanization: Evidence from a Historical Experiment," *Quarterly Journal of Economics*, 2011, *126* (2), 593–650.

Petry, C.F., *The Civilian Elite of Cairo in the Later Middle Ages*, Princeton: Princeton University Press, 1981.

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

Rollo-Koster, J., *Avignon and Its Papacy, 1309-1417*, London: Rowman & Littlefield, 2015.

Romanov, M., "Algorithmic Analaysis of Medieval Arabic Biographical Collections," *Speculum*, 2017, *92* (S1), S226–S246.

Roth, J., "Pre-test with Caution: Even-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, forthcoming.

Schennach, S., "Recent Advances in the Measurement Error Literature," *Annual Review of Economics*, 2016, *8*, 341–377.

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.

Stein, C., "Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution," in "Proceedings of the Third Berkeley Symposium," University of California Press, 1955, pp. 197–206.

— , "Confidence Sets for the Mean of a Multivariate Normal Distribution," *Journal of the Royal Statistical Society*, 1962, *24* (2), 265–296.

Stigler, S., "The 1988 Neyman Memorical Lecture: A Galtonian Perspective on Shrinkage Etsimators," *Statistical Science*, 1990, *5* (1), 147–155.

Toonen, W., M. Macklin, G. Dawkes, J. Durcan, M. Leman, Y. Nikolayev, and A. Yegorov, "A hydromorphic reevaluation of the forgotten river civilizations of Central Asia," *Proceedings of the National Acadmey of Sciences*, 2020, *117* (52), 32982–32988.

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