Information Technology and Economic Change: The Impact of the Printing Press

Appendices

Jeremiah Dittmar*

January 10, 2011

*Contact: American University, Department of Economics, 4400 Massachusetts Avenue NW, Washington, DC 20016. Email: dittmar.academic@gmail.com. Web: www.jeremiahdittmar.com.

A Data Appendix

Appendix A is contained within the main body of the paper.

B Propensity Score Analysis

This section employs a propensity scoring approach developed in the program evaluation literature to examine the factors associated with adoption and the association between print technology and city growth.[1] Using propensity scores to control for the probability of adoption, I find that cities that adopted printing in the late 1400s grew 39 percent faster than similar cities that did not 1500-1600.

The propensity score is an index of the likelihood of adoption. In this context, it sheds light on potential endogeneity problems. Specifically, I find that while adoption of the printing press was associated with high growth, the likelihood of adoption was negatively associated with future growth. This analysis suggests that entrepreneurs established printing presses at cities that previously experienced relatively high growth, but that they did not accurately forecast future growth.

Let us denote the logarithm of gross city population growth over some period after 1500 by Y_i. Let us denote the binary adoption (or "treatment") variable by T_i:

$$T_i = \begin{cases} 1 & \text{if} \quad \text{city adopted printing press by 1500} \\ 0 & \text{if} \quad \text{city did not adopt printing press by 1500} \end{cases}$$

A vector X_i captures each city's pre-treatment population growth and other pre-treatment characteristics (e.g. the presence of a university, important religious site, or political capital; country indicators; location on a navigable river, port, or Roman site; and institutional variables). For every city i, we observe (T_i,Y_i,X_i). We posit:

$$Y_i \equiv Y_i(T_i) = (1 - T_i)Y_i(0) + (T_i)Y_i(1)$$

In a clean experiment, the average treatement effect (ATE) of adoption is:

$$ATE = \mathbb{E}_i \left[Y_i(1) - Y_i(0) \right].$$

But historical data are marked by an unobserved counterfactual. For any city we observe $Y_i(0)$ or $Y_i(1)$, not both. Hence to estimate the ATE we need to construct a comparison

[1] See Imbens and Wooldridge (2009), Imbens (2004), and Wooldridge (2002) for reviews.

of outcomes across similar treated and control observations.

The propensity score is the probability of technological adoption, conditional on city characteristics:

$$P(X_i) = \Pr(T_i = 1|X = X_i) = \mathbb{E}\left[T_i|X = X_i\right]$$

By accounting for this conditional probability, we can control for selection into technology adoption and examine the extent to which cities with printing presses grew faster (or slower) than otherwise similar cities that did not adopt the new information technology.

I estimate propensity scores using a logit model in which the binary variable capturing whether or not print technology was adopted by 1500 is a function of: city size, and variables capturing whether a city was on a port or navigable river or the site of Roman settlement and whether the city was historically the location of a university. I also include as controls country fixed effects, an extended version of DeLong and Shleifer's (1993) indicator for whether the prevailing regime was "Prince" or "Free", city latitude, longitude, and the interaction between latitude and longitude.[2] It is reasonable to inquire whether the establishment of printing presses in neighboring cities impacted adoption decisions elsewhere. However, I find no evidence of such effects once one controls for country fixed-effects and latitude and longitude and do not report these specifications.[3] Similarly, I find no effects of including controls for religious sites such as archbishoprics and episcopal sees. Specifications that include distance from Mainz as an additional control generate virtually identical results.

Table B1 presents parameter estimates from an OLS and logit regressions. It shows that adoption was significantly associated with city size and the presence of a university, and negatively associated with ports. The negative association with ports suggest that access to waterborne transport was not a significant positive determinant of adoption (it holds even if distance from Mainz is introduced as an additional regressor). City size in 1400 and city size in 1500 are included as regressors to capture the association between pre-treatment growth rates and adoption. The identifying assumption is that – although adoption occurred in the late the 15th century – the adoption decision did not impact

[2]The results I report below are not contingent on the inclusion of the extended DeLong-Shleifer freedom index. Including an indicator for political capitals does not substantively change the OLS results. Because all capitals adopted printing presses, these observations would be dropped from logit specfications. Because I include country fixed effects I conservatively restrict to those countries with at least one print city and one non-print city with population data. These countries are: England, Belgium, France, Switzerland, Germany, Hungary, the former Czechoslovakia, Italy, and Spain. The estimated effects are stronger when one drops country fixed effects and expands the sample to territories in which all cities with population data adopted the press (e.g. Portugal with Lisbon and Porto) or with no early printing cities (e.g. Russia and the former Yugoslavia).

[3]I examined the effect of neighbors' adoption within various distances and using different transformations of distance as weights.

city size in 1500.[4] The country fixed effects begin to capture and control for the regional aspect of diffusion, but should not be taken to suggest that national economies and were anything more than incipient.

I use the parameter estimates from Table B1 to compute propensity scores.[5] Figure B1 plots the densities of propensity scores for adopting and non-adopting cities. It shows that the propensity scores of cities that adopted the printing press in the late 1400s are typically high and that most non-adopting cities have low estimated propensity scores. However, a considerable number of cities are in the thin upper tail of the propensity score distribution for non-adopting cities. Non-adopting cities with estimated propensity scores over 0.5 include: Bordeaux, Reims, Braunschweig, Groningen, Lille, Maastricht, Cordoba, Arezzo, Aachen, Dublin, Tournai, Bourges, Montpellier, and Aix-en-Provence. Non-adopting cities with estimated propensity scores $\hat{P} \in (0.4, 0.5)$ include Bremen, Marseilles, Malaga, Beauvais, Dortmund, Rimini, Dordecht, Poznań, Salerno, Goslar, Mechelen, and Arras. Neither Amsterdam nor Berlin nor Madrid adopted the press in the 1400s.[6] More obviously, there is also substantial overlap in the distributions for adopting and non-adopting cities propensity scores for $\hat{P}(X_i) \in (0.20, 0.4)$. This overlap provides purchase for econometric identification.

The estimated propensity scores can be used to examine possible selection (endogeneity) effects in technology adoption. A selection problem would arise if (i) adoption was associated with above par growth in future years, and (ii) adoption was associated with the accurate expectation of above par growth – or, more broadly, with factors that augured well for city growth. If this were the case, the association between adoption and subsequent growth need not reflect the impact of the technology.

However, regression analysis of early technology adoption confirms that there was both a positive printing press effect and a *negative* association between the likelihood of adoption and future growth. In general, we expect an outcome Y_i to be some function of the treatment T_i and the propensity score $\hat{P}_i = \hat{P}(X_i)$ measuring the probability that a given observation receives treatment. Following an approach developed in the program evaluation literature, the estimated propensity score can be employed as a control

[4]Specifications that drop population in 1500 as a regressor yield similar estimates. If cities that adopted printing in the late 1400s already experienced a positive growth impact before 1500, this would conservatively bias the estimates down. Essentially, the regressions take the 1500-1600 as the first period in which treated units are consistently subject to treatment.

[5]A flexible logit specification in which adoption is a function of each of the variables in Table B1, their squares, and interactions yields similar propensity scores and does not substantively change the conclusions one draws about the association between print technology and city growth.

[6]This evidence contradicts Eisenstein's (1979: 440) claim that by 1500 there were printers' workshops in "every important municipal center."

Table B1

Regression Analysis of the Adoption of the Print Press

Independent Variable	Logit	OLS
(1)	(2)	(3)
Log Population 1500	0.60	0.06
	(0.49)	(0.05)
Log Population 1400	1.30 ***	0.17 ***
	(0.43)	(0.05)
University	2.57 ***	0.33 ***
	(0.75)	(0.07)
Roman Site	0.60	0.09
	(0.52)	(0.06)
Freedom 1500	-0.05	-0.04
	(1.83)	(0.19)
Freedom 1400	2.13 *	0.25 **
	(1.10)	(0.12)
Navigable River	0.80	0.10
	(0.57)	(0.07)
Port	-1.27 ***	-0.13 **
	(0.49)	(0.06)
Latitude & Longitude	Yes	Yes
Country FE	Yes	Yes
Observations	265	265
F Statistic		21.37 ***
Wald Chi Square	77.50 ***	
R Square	0.48	0.48

Note: The dependent variable is an indicator for cities that adopted the printing press 1450-1500. Regressions include country fixed effects and controls for city latitude, longitude, and the interaction between latitude and longitude. "Freedom" is the DeLong-Shleifer coding of political institutions. All variables described in the text of the paper and/or Appendix A. Heterskedasticity-robust standard errors in parentheses. Significance at the 90, 95, and 99 percent confidence indicated "*", "**", "***", respectively.

Figure B1
Propensity Score Densities

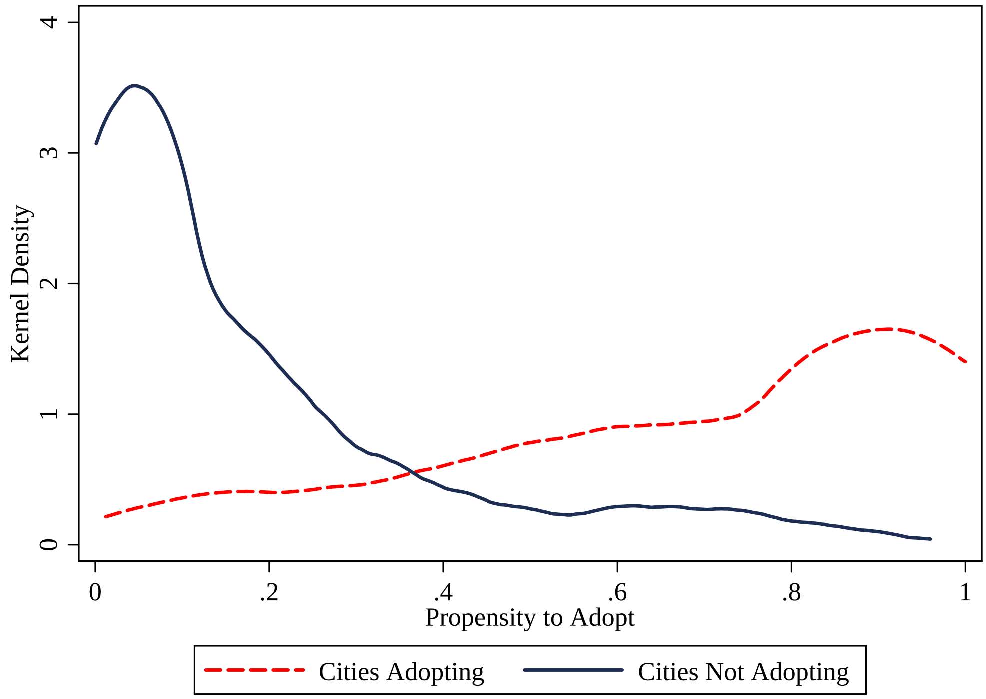

Note: This figure shows the distribution of propensity scores. See text for discussion of the upper-tail cities that did not adopt printing in the 1400s.

function and we can estimate the ATE in a model:

$$Y_i = \alpha_0 + \alpha_1 \hat{P}i + \alpha_2 T_i + \epsilon_i \tag{1}$$

Here the treatment effect is captured in α_2, the coefficient on technology adoption. The estimate of α_2 is consistent assuming (i) $\mathbb{E}\left[Y(1) - Y(0)|X_i\right]$ is uncorrelated with $\mathrm{Var}(T|X_i)$ and (ii) unconfoundedness (sometimes called "selection on observables").[7] Because $\mathrm{Var}(T|X_i)$ is a nonmonotonic quadratic in $P(X_i)$ and $\mathbb{E}\left[Y(1) - Y(0)|X_i\right]$ will likely be linear in several elements of X_i, zero correlation may hold approximately.[8] The key assumption is unconfoundedness.

Table B2 reports results estimating the model in equation (1) over several different periods. Panel A shows the baseline results. Three points are notable. First, the estimates show printing cities had no growth advantage prior to adoption. Second, the estimate of their growth advantage in the century after adoption is highly significant and very large: print cities grew an extra 0.33 log points (39 percentage points). For comparison, mean

[7]Formally, the unconfoundedness assumption is that $\mathbb{E}[Y(j)|T, X] = \mathbb{E}[Y(j)|X]$, for $j \in 0, 1$.
[8]See Wooldridge (2002: 617-618) for discussion.

Table B2
Printing and Log City Growth: Propensity Score Analysis

Variable	Pre-Adoption Log Growth 1400-1500	Post-Adoption Log Growth 1500-1600	Log Growth 1500-1700	Log Growth 1500-1800
(1)	(2)	(3)	(4)	(5)
Panel A: All Cities				
Propensity to Adopt	-0.20	-0.51 ***	-0.55 ***	-0.82 ***
	(0.13)	(0.13)	(0.19)	(0.17)
Adopt Printing by 1500	0.00	0.33 ***	0.21	0.20
	(0.09)	(0.10)	(0.15)	(0.13)
Observations	265	232	238	265
F Statistic	2.48 *	7.79 ***	5.75 ***	14.52 ***
Panel B: Non-German Cities				
Propensity to Adopt	-0.18	-0.57 ***	-0.66 ***	-0.91 ***
	(0.15)	(0.15)	(0.19)	(0.19)
Adopt Printing by 1500	0.01	0.37 ***	0.34 **	0.29 **
	(0.11)	(0.11)	(0.15)	(0.14)
Observations	217	189	201	217
F Statistic	1.51	7.75 ***	6.09 ***	14.42 ***

Note: The regressions are of the form: $Y_i = \alpha_0 + \alpha_1 P_i + \alpha_2 T_i + \epsilon_i$, where Y_i is city i's log population growth, T_i is an indicator capturing whether city i adopted the printing press by 1500, and P_i is the estimated propensity score. Significance at 90, 95, and 99 percent confidence denoted "*", "**", "***" respectively.

city growth for all cities was 0.27 log points (31 percentage points) both 1500-1600 and 1500-1700. Third, the estimated print growth advantage 1500-1800 is a more modest 0.2 log points and is not significant at conventional confidence levels.[9] These results reflect the massive demographic losses German print cities experienced during the 30 years war (1618-1648).[10]

In Table B2, Panel B shows the results excluding German cities. Outside Germany, cities that adopted the press in the late 1400s had no growth advantage 1400-1500 but a consistent, significant advantage of 0.3 log points after 1500. Essentially, these estimates control for the slow growth German print cities experienced during the 1600s. Many of the German print cities became sites of Protestant agitation. As Cantoni (2009) documents, Protestant cities were exposed to relatively severe negative shocks over the course of The Thirty Years War (1618-1648).

[9] Standard errors adjusted reflect presence of estimated regressors. In specifications of the propensity model where country fixed effects do not appear, the print effect retains significance through 1800.

[10] As noted above, the negative association between the probability of adoption and future growth reflects the fact that big cities were likely to adopt and to grow slowly. If one introduces city size as an additional regressor, the estimated impact of printing is unchanged while the negative association between probability of adoption and subsequent growth vanishes.

The negative association between the propensity to adopt and future growth suggests that adoption was not driven by correct expectations about future city growth. It is explained by the fact that (i) printing technology was typically adopted in cities that were already relatively large, and (ii) large cities grew relatively slowly 1500-1600 (and to some extent 1700-1750). In contemporary economies, random or size-independent growth is the norm.[11] However, as shown in Dittmar (2010), city growth in pre- and early modern Europe was non-random.

Where there is reason to suspect selection into treatment, and where we are willing to add the assumption that the expectation of the outcome is linear in the propensity score, we can further control for these effects by introducing a term that captures the association between the outcome and the interaction between treatment and the propensity score[12]:

$$Y_i = \alpha_0 + \alpha_1 \hat{P}_i + \alpha_2 T_i + \alpha_3 \left[T_i \cdot (\hat{P}_i - \mu_{\hat{P}}) \right] + \epsilon_i \tag{2}$$

Because it is natural to also be concerned about propensity scores $\hat{P}(X_i)$ close to 0 or 1, Imbens and Wooldridge (2009) propose a rule of thumb for trimming the data in order to improve overlap in covariate distributions. They suggest that researchers examine first the complete data and then observations with propensity scores $\hat{P} \in \mathbb{A} = [0.1, 0.9]$

Table B3 presents estimates of equation (2) that show no evidence of selection into treatment and indicate a print effect of over 0.3 log points. This holds for both the complete data and the smaller sample trimmed to exclude observations with propensity scores less than 0.1 or greater than 0.9.

These results suggest that cities that adopted the printing press in the late 1400s grew as much as 100 percent faster than cities that did not 1500-1600. These estimates may be conservative. Some cities that were not early adopters did subsequently adopt. This would likely have muted the advantage conferred by early adoption.

Propensity scores can also be exploited to test for cross-city spillovers to technology adoption. To test for cross-city spillovers, I examine whether the establishment of a printing press in neighboring cities was associated with variations in city growth. Specifically, I estimate:

$$Y_i = \alpha_0 + \alpha_1 P_i + \alpha_2 T_i + \alpha_3 P_i^* + \alpha_4 T_i^* + e_i \tag{3}$$

Here P_i and T_i are city i's propensity score and binary treatment. The terms P_i^* and

[11] See Gabaix (2008) and Gabaix and Ioannides (2004).

[12] Formally, the interaction term is the interaction between treatment and the deviation from the mean propensity score. The linearity assumption is $\mathbb{E}[Y(j)|\hat{P}]$ is linear in $\hat{P}$.

Table B3
Testing for Selection in Adoption: Log City Growth 1500-1600

Independent Variable	Complete Data	Trimmed Data
(1)	(2)	(3)
Propensity to Adopt	-0.50 ***	-0.18
	(0.18)	(0.26)
Adopt by 1500	0.34 ***	0.39 ***
	(0.11)	(0.13)
(Propensity) x (Adopt by 1500)	-0.04	-0.47
	(0.26)	(0.41)
Constant	0.32 ***	0.21 *
	(0.06)	(0.11)
Observations	232	121
R Square	0.06	0.10

Note: Parameter estimates for equation (2). Trimmed data restricts to cities with propensities $\hat{P} \in [0.1, 0.9]$. Heteroskedasticity-robust standard errors in parentheses. Significance at 90, 95, and 99 percent confidence denoted "*", "**", "***" respectively.

T_i^* capture the propensity scores and the technology adoption decisions in neighboring cities, and could be implemented as distance-weighted sums:

$$P_i^* = \sum_{j \neq i} \frac{P_j}{d_{ij}} \qquad \text{and} \qquad T_i^* = \sum_{j \neq i} \frac{T_j}{d_{ij}}$$

Here P_j and T_j are city j's propensity score and technology adoption decision, respectively, and d_{ij} is the distance between city i and city j. Distance is calculated using latitude and longitude as "great circle" distance.

Table B4 documents that there was no apparent association between the establishment of printing presses in neighboring cities 1450-1500 and city growth 1500-1600. Column (2) replicates the baseline propensity score results. Column (3) shows that the baseline results are robust to controlling for neighbors' adoption decisions. Column (4) shows that the baseline results are robust to controlling for neighbors' adoption decisions and neighbors' propensity to adopt.

Table B4
Testing for Cross-City Spillovers: Dependent Variable is Log City Growth 1500-1600

Variable	Model 1	Model 2	Model 3
(1)	(2)	(3)	(4)
Adopt Printing 1450-1500	0.33 ***	0.33 ***	0.31 ***
	(0.10)	(0.10)	(0.10)
Propensity to Adopt	-0.51 ***	-0.52 ***	-0.52 ***
	(0.13)	(0.14)	(0.14)
Neighbors Adopt Printing 1450-1500		0.02	0.71
		(0.24)	(0.54)
Neighbors Propensity to Adopt			-1.21
			(0.91)
Observations	232	232	232
F Statistic	7.79 ***	5.16 ***	4.1 ***

Note: Parameter estimates for equation (3). Heteroskedasticity-robust standard errors in parentheses. Significance at 90, 95, and 99 percent confidence denoted "*", "**", and "***".

C Difference-in-Differences and First-Differences

This section reports difference-in-differences and first-differences estimates of the effect of adopting printing. Using either difference-in-differences or first-differences, I find that the early adoption of printing technology was associated with a growth advantage of 0.17 log points 1500-1600. Mean city growth was 0.27 log points 1500-1600, implying that cities that adopted the technology in the late 1400s grew 60 percent faster than the average city over this period.

Difference-in-differences estimators account for the effects of unobserved confounding variables provided the latter are constant over time. The difference-in-differences estimator can be estimated:

$$Y_{it} = \alpha_0 + \alpha_1 T_i + \alpha_2 \mathit{YEAR1500}_t + \alpha_3 (T_i \cdot \mathit{YEAR1500}_t) + \beta' X_{it} + \epsilon_{it} \qquad (4)$$

As before, Y_{it} is log growth and T_i is an indicator capturing whether a city adopted printing technology in the late 1400s ("treated" observations). The variable $\mathit{YEAR1500}_t$ is an indicator for the post-treatment period.[13] X_{it} is a vector of additional city characteristics. The parameter of interest is α_3, which captures the average treatment effect of adopting print technology in the late 1400s.

[13] $\mathit{YEAR1500}_t = 1$ if $t = 1500$. As discussed above, the average city adopted the printing press in 1476. To the extent printing cities benefitted from technology adoption immediately (i.e. 1476-1500), the difference-in-difference estimates presented here will be conservative.

Table C1
Difference-in-Differences Estimates of Log City Growth

Variable	Baseline Model	Model with City FE
(1)	(2)	(3)
Year1500	0.02	0.02
	(0.07)	(0.10)
Print	-0.07	-0.40 ***
	(0.06)	(0.04)
Print x Year1500	0.17 ***	0.17 **
	(0.06)	(0.08)
Observations	516	516
R Squared	0.01	0.48

Note: Regression estimated for 258 cities on which populations are observed 1400, 1500, and 1600. Heteroskedasticity-robust standard errors clustered at country level. Significance at 90, 95, and 99 percent confidence denoted "*", "**", and "***" respectively.

Table C1 presents results from difference-in-difference regressions estimated over data for 1400-1600 (i.e. examining growth 1400-1500 and 1500-1600). It shows that the difference-in-difference estimate of the average treatment effect is $\hat{\alpha}_3 = 0.17$ and statistically significant in both the basic difference-in-differences model and a model that includes city fixed effects.[14]

We obtain similar estimates of the impact of technology adoption if we exploit the panel structure of the data and estimate an unobserved (fixed) effects model in a first-differenced equation. In this case, one examines the association between changes in growth rates and changes in a variable capturing the presence of a printing press at the start of each period. Formally, $\Delta Y_i \equiv Y_{i1500} - Y_{i1400}$, and T_i is equivalent to the change in an indicator capturing the presence of a printing press at time t ($T_i \equiv \Delta PRINT_i \equiv PRINT_{i1500} - PRINT_{i1400}$). The estimating equation is:

$$\Delta Y_i = \beta_0 + \beta_1 T_i + \nu_i \tag{5}$$

Estimating (5) over the balanced panel of 258 cities, one obtains $\hat{\beta}_1 = 0.17$ with heteroskedasticity-robust standard error of 0.1 and associated t-statistic of 1.75.

[14] Additional results (not shown here) indicate that (i) there is no association between universities and city growth and (ii) that there is also no association between university-print interactions and city growth.

D Regression Results with Balanced Panel

This section presents regression estimates of the association between city growth and the establishment of the printing press over the balanced panel of cities observed 1300-1800. In the unbalanced panel the estimated print effect is 0.19 log points 1500-1600 (reported in Table IV in the main body of the paper). Table D1 shows that in the balanced panel the estimated print effect is 0.29 log points 1500-1600 (see column 3).

Table D1
Regression Analysis of Print Media and Log City Growth in Balanced Panel

	Dependent Variable is Log City Growth			
	Pre-Adoption	Post-Adoption		
Independent Variable	Log Growth 1400-1500	Log Growth 1500-1600	Log Growth 1500-1700	Log Growth 1500-1800
(1)	(2)	(3)	(4)	(5)
Print Adoption 1450-1500	0.11	0.29 ***	0.24 *	0.31 **
	(0.10)	(0.10)	(0.14)	(0.14)
Editions Per Capita	0.11 **	0.02	0.04	0.06
	(0.05)	(0.05)	(0.07)	(0.07)
University	-0.18	0.06	0.17	0.17
	(0.12)	(0.09)	(0.12)	(0.14)
Roman Site	0.15 *	0.03	0.08	0.10
	(0.08)	(0.08)	(0.10)	(0.10)
Capital	0.29 **	1.03 ***	1.52 ***	1.91 ***
	(0.14)	(0.25)	(0.28)	(0.37)
Freedom Index	-0.26 *	0.66 **	0.47	0.27
	(0.16)	(0.29)	(0.30)	(0.37)
Atlantic Port	0.18	0.42 ***	0.68 ***	0.94 ***
	(0.20)	(0.15)	(0.18)	(0.18)
Mediterranean Port	0.17	0.10	0.60 ***	0.67 ***
	(0.14)	(0.17)	(0.20)	(0.22)
Baltic Port	-0.14	0.14	0.32	0.29
	(0.14)	(0.21)	(0.25)	(0.37)
Navigable River	0.13	0.12	0.16	0.22 *
	(0.09)	(0.08)	(0.11)	(0.12)
Log Population	-0.24 ***	-0.32 ***	-0.42 ***	-0.63 ***
	(0.06)	(0.06)	(0.08)	(0.09)
Country FE	Yes	Yes	Yes	Yes
Observations	202	202	202	202
R Square	0.37	0.41	0.45	0.56

Note: Significance at 90, 95, and 99 percent confidence denoted "*", "**", and "***" respectively.

E Printing and Urban Incomes

Given the association between printing and city growth, it is natural to wonder whether the printing press impacted incomes at the city level.

The historical data on wages is high frequency but limited to a small number of cities almost all of which adopted the press 1450-1500. For instance, Allen (2007) provides data on the wages of skilled craftsmen and unskilled laborers in 20 cities, 16 of which got printing presses 1450-1500. In the 4 cities without printing presses 1450-1500 (Madrid, Amsterdam, Lwow, and Warsaw), Allen's wage data are available only from the early to mid-1500s.[15] As a result, any inferences based on wage data must be extremely tentative.

Here I simply present one interesting finding: In cities for which data on skilled and unskilled wages are available, skill premia had been steadily declining from high levels reached following the Black Death (1348), but after the establishment of printing presses skill premia began to rise. Figure E1 documents this fact by plotting the evolution of the ratio of skilled (craftmen's) wages to unskilled (laborers') wages across eight cities with data available from the early 1400s.[16]

Figure E1: Printing and Skill Premia Across 8 Cities

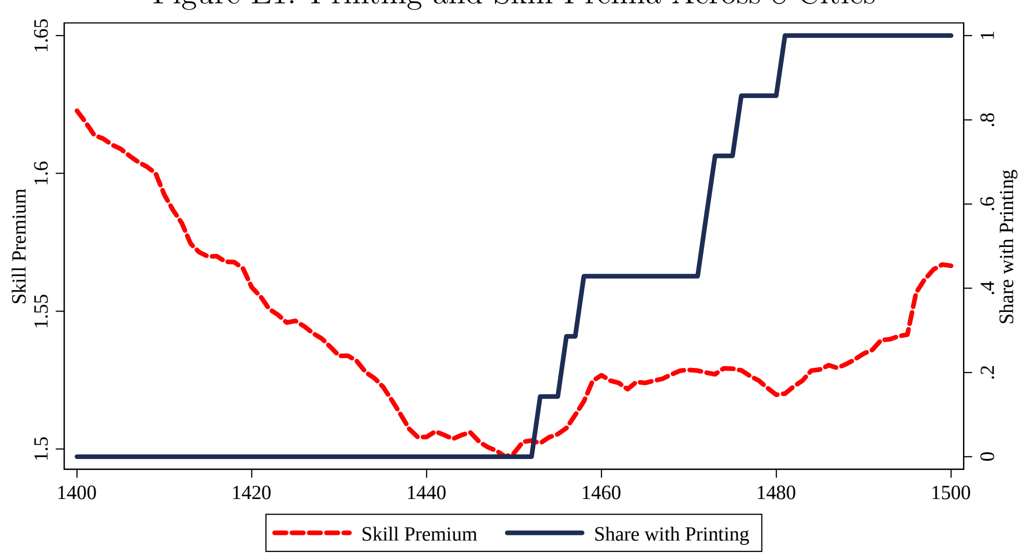

Note: This figure presents the 20-year moving average mean skill premium in cities with wage data available from the early 1400s: Antwerp, Valencia, Strasbourg, Paris, London, Oxford, Krakow, and Florence. Data from Allen (2007).

[15] Data compiled by the Global Price and Income History Group (UC Davis) and the International Institute of Social History are similarly limited.

[16] The data in Allen (2007) further show that wages for all workers in these cities rose 1500-1600. But while such increases are consistent with workers being compensated for the disamenities of larger cities, they also reflect the widespread price inflation of the 1500s.

References

[1] Acemoglu, Daron, Simon Johnson, and James Robinson, "The Rise of Europe: Atlantic Trade, Institutional Change and Economic Growth," *American Economic Review*, 95 (2005), 546-579.

[2] Acemoglu, Daron, Simon Johnson, and James Robinson, "The Rise of Europe," NBER Working Paper 9378 (2002).

[3] Allen, Robert (2007), *Wages, Prices & Living Standards: The World-Historical Perspective*, http://www.economics.ox.ac.uk/Members/robert.allen/WagesPrices.htm.

[4] Cantoni, Davide, "The Economic Effects of the Protestant Reformation: Testing the Weber Hypothesis in the German Lands," Harvard University Working Paper (2009).

[5] DeLong, J. Bradford and Andrei Shleifer, "Princes and Merchants: City Growth Before the Industrial Revolution," *Journal of Law and Economics*, 36 (1993), 671-702.

[6] Dittmar, Jeremiah, "Cities, Markets, and Growth: The Emergence of Zipf's Law," www.jeremiahdittmar.com (2010).

[7] Imbens, Guido, "Nonparametric Estimation of Average Treatment Effects Under Exogeneity," *The Review of Economics and Statistics*, 86 (2004), 4-29.

[8] Imbens, Guido and Jeffrey Wooldridge, "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 47 (2009), 5-86.

[9] Eisenstein, Elizabeth, *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe* (Cambridge: Cambridge University Press, 1979).

[10] Gabaix, Xavier, "Power Laws in Economics and Finance," NBER Working Paper 14299 (2008).

[11] Gabaix, Xavier and Yannis Ioannides, "The Evolution of City Size Distributions," in J. Henderson and J. Thisse eds., *The Handbook of Regional and Urban Economics*, Vol. 4 (2004).

[12] Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data* (Cambridge: MIT, 2002).