Laurier Centre for Economic Research & Policy Analysis

LCERPA Working paper No. 2025-8

July 2025

The Economic Effects of Long-Term Climate Change: Evidence from the Little Ice Age: Replication

Nikolai Cook
Wilfrid Laurier University

Hugo Cordeau
University of Toronto

Tongzhe Li
University of Guelph

Taylor Wright
Brock University

The Economic Effects of Long-Term Climate Change: Evidence from the Little Ice Age: Replication*

Nikolai Cook
Wilfrid Laurier University

Hugo Cordeau
University of Toronto

Tongzhe Li
University of Guelph

Taylor Wright
Brock University

June 27, 2025

Keywords: Replication in Economics, Climate Change, Environmental Economics, Economic History, City Population

JEL CODES: Q5, Q54, Q56, R23, N5

Abstract

Waldinger (2022) finds a positive relationship between temperature and city size during the climate change of 1600-1850. We show the main result differs by city size. Cities with less than 1000 inhabitants (which make up 23.5% of observations and are 49.6% of cities at some point) exhibit a strong and positive relationship between temperature and city size, whereas cities with always more than 1000 inhabitants exhibit a negative relationship. Further examination of the underlying city size data, which bins populations into coarse thousand-wide population intervals, finds the original analysis to be robust to a number of reasonable alternative researcher choices.

*Cook: (Corresponding) Wilfrid Laurier University, Lazaridis School of Business and Economics, 64 University Ave West, Waterloo, ON N2L 3C7 ncook@wlu.ca. Cordeau: University of Toronto, Department of Economics, 150 St. George St., Toronto, ON M5S 3G7, hugo.cordeau@mail.utoronto.ca. Li: University of Guelph, Department of Food, Agricultural and Resource Economics, 50 Stone Road E., Guelph, ON N1G 2W1, tongzhe@uoguelph.ca. Wright: Brock University, Department of Economics, 1812 Sir Isaac Brock Way, St. Catharines, ON L2S 3A1 twright3@brocku.ca. This replication is part of the 2023-02-22 Toronto Replication Games held by the Institute for Replication and hosted by the University of Toronto. The choice of article was made by the replicators from a list given to them by the Institute. We thank other participants of the Games for helpful conversations and the Institute for Replication for hosting them. We also thank the editor and three anonymous referees for their comments. Errors are ours.

1 Introduction

Waldinger (2022) combines annual paleo-climatic temperature data with a panel of European city populations every half-century during the Little Ice Age (1600-1850) to examine how cooler temperatures affected economic outcomes (proxied using city size) under adaptation; the gradual cooling over a long period should have offered time for economies to respond offering the researcher empirical evidence of a climate-economy effect accounting for adaptation.

In this replication, we show that the main results contain a nuance that should be highlighted: cities with less than 1000 inhabitants in the original Bairoch et al. (1988) data, which make up the single largest group of city size populations, exhibit a different temperature / city size relationship than the remainder of the sample. Accounting for heterogeneity while applying otherwise identical data and research methods provides evidence of a statistically significant effect that is of similar magnitude but of *opposite* sign for large cities than that of smaller cities.

We note that the unaccounted for heterogeneity we comment on here is confined to the main result; we do not evaluate the remainder of the evidence presented in the mechanism or adaptation analyses (which are based on novel but less geographically representative data). We wish to highlight that we consider the analysis of Waldinger (2022) to be well conducted and of great importance - understanding the effects of climate change under realized rather than potential adaptations and examining empirical evidence in place of relying on climate projections to answer these questions is a large contribution.

The remainder of this paper is structured as follows. First, we present our computational reproduction of the original article's results. Then we describe a data assumption made in the original article, and demonstrate the robustness of Waldinger (2022)'s results to any alternative that could have been applied. We then present our heterogeneity results using research methods otherwise identical to the original. We also include an analysis which attempts to alleviate some of the issues coarse binning of the original city size data may have led to in the analysis following lessons learned from Eeckhout (2004); we find the original article's results robust to our proposed solution.

2 Computational Reproduction of Waldinger (2022)

This section documents our computational reproduction of Waldinger (2022)[1] which posits that long-term climate change during the Little Ice Age (between 1600 and 1850) had significant negative economic effects (proxied by city size). The article then uses different datasets to demonstrate two potential mechanisms and two potential adaptations.

For context, Section II of Waldinger (2022) provides a description of the main dataset; a balanced panel of annual temperatures for Europe for each year since 1500 from Luterbacher et al. (2004) combined with data on city size (a proxy for economic activity) in 1600, 1700, 1750, 1800, and 1850 from Bairoch et al. (1988). Section III contains the main results. Section IV considers different outcome variables to explore mechanisms (agricultural productivity, mortality, and migration) while section V explores potential adaptations (trade, changing land use), all relying on less comprehensive data sets.[2]

[1]Published by the Journal of Political Economy in September 2022, and at the time of writing has an Altmetric Score of 67, placing it in the the top 5% of all research outputs scored by Altmetric and in the 96th percentile of research of the same age.

[2]Using data from the sixteenth to the nineteenth centuries for locations in 12 cities in the case of agricultural productivity, data from 1538-1838 for 404 English parishes in the case of mortality, and data from 1571-1871 for 7 English parishes in the case of migration. Using data from 1591-1857 for 750 European ports in the case of trade and a healthy 42 European countries for each decade since 1500 for land use.

We first confirm the computational reproducibility[3] of the article; all four of us replicators are able to independently use the provided Supplemental Material to *exactly* reproduce the published main and secondary results. For the main results, see our Table 1 (which is a replication of Waldinger (2022)'s Table 2 presented in their section "Empirical Strategy and Main Results"). There is one difference; in our version we have added p-values in square brackets.

Waldinger (2022) uses ordinary least squares to regress the natural logarithm of city size in a given year against the past 50-or-100 year-average temperature for that city.[4] Specifically, the following is estimated;

$$logCitySize_{it} = \beta + \gamma MeanTemperature_{it} + TimeFE_t + CityFE_i + \delta X_{it} + \epsilon_{it} \tag{1}$$

"where $logCitySize_{it}$ represents the natural logarithm of the size of city i in time period t, $Mean$ $Temperature_{it}$ represents the mean year temperature in city i and time period t over the past 100 and 50 years, $TimeFE_t$ represents time period fixed effects that control for variation in temperature and in city size over time that is common to all cities, and $CityFE_i$ represents city fixed effects. The city fixed effects control for time-invariant city characteristics–for example, distance to the ocean and waterways and permanent climatic or soil characteristics that may affect a city's access to trade or its agricultural productivity. The parameter X_{it} represents a number of control variables, each interacted with time period indicator variables. They are described in more detail when introduced into the analysis. The parameter ϵ_{it} is the error term." (Waldinger (2022), p. 2286)

The main result of (Waldinger, 2022, p.2289) is "a positive relationship between mean temperature and city size." This corresponds to column 1 and 2 of Table 1. There are three versions of standard errors included which we have computed in our tables. The coefficient of 0.532 represents a 70% increase in city size following a 1 degree Celsius warmer 50-year period ($e^{0.532} = 1.7023$). The next columns introduce different fixed effects specifications such as geographic controls in column 2. Columns 3 and 4 restrict the sample to non-capital cities and non-ocean-adjacent cities, respectively. Column 5 and 6 include a country specific linear time trend and country-by-time period fixed effects, respectively. As we move from left to right, the estimated effect of mean temperature increases from 0.532 to 1.170. The most statistically significant estimates have a p-value less than 0.001, while the least statistically significant estimate has a p-value of 0.064.

2.1 Robustness of Waldinger (2022)'s Results

This section demonstrates the robustness of the original article's results to what could be conceived as a critical assumption. Due to the coarse binning of the Bairoch et al. (1988) city size data, Waldinger (2022) necessarily assumes populations coded as 0 (which represent populations somewhere between 0 and 999) have a constant integer city size population of 500. Figure 1 shows that assuming *any* integer between 0 and 1000 for the censored data will return a statistically significant main result using identical research methods

[3]Following the definitions provided by the Institute for Replication, where **Computational Reproducibility** is the ability to duplicate the results of a prior study using the same data and procedures as were used by the original investigator.

[4]In a situation where data is systematically left-censored as in this case (the largest single bin of city size data codes all populations of 0 if the data-producers believed the population to be anywhere from 0-999 inhabitants), it is common to apply tobit regression models as ordinary least squares is biased. However, the replicators found that there is, at present, no widely adopted methodology that allows for fixed (rather than the available random) effects to be applied in a panel setting in conjunction with the tobit model. Because our aim is to remain as close as possible to the original research question (does within-city variation in temperature affect city growth) we present results which do not assume random effects or otherwise cross-sectional models. When the replicators estimated a tobit model with all else the same as in Waldinger (2022), the estimated effect of temperature was substantially lower, both in models with random effects and models that were cross-sectional only. Results available upon request.

to Waldinger (2022). In the left panel we present a histogram of the coefficients corresponding to column 2 of Table 1 where each coefficient comes from a regression where the censored city populations are assumed to be a constant integer increasing from 1 to 1000 in steps of 1 (we provide a dashed vertical line at Waldinger (2022)'s assumed 500). In the right panel we present the associated t-statistics. From this histogram the robustness of the positive and statistically significant main result to the *choice* of *any* constant number in the interval [1,1000] is apparent.[5]

3 Coarse Outcome Variable Data

We now turn to a description of the data and highlight two results of our own analysis. The data that Waldinger (2022) uses for their main results is a combination of annual temperatures provided by Luterbacher et al. (2004) and a panel of city populations in 1600, 1700, 1750, 1800, and 1850 provided by Bairoch et al. (1988). The result is 10,600 city-period observations representing 2,120 distinct cities during the 5 periods.

As noted in Waldinger (2022), every city-period in the original data provided by Bairoch et al. (1988) was coded as having 0 inhabitants if there were fewer than 1000 during that year. As the analysis ultimately estimates the effect of temperature on *percentage change* of city size, we noted the possibility that it would have been "easier" for smaller populations reported in this manner to show large population shifts - a contradiction of Gibrat's Law.[6]

Figure 2 presents a histogram of the populations corresponding to the 10,600 city-period observations used in the analyses; there is a striking proportion of city-period observations within the smallest population bin. As depicted in the leftmost bar in the histogram, 23.5% of observations require an assumed population, representing the single-largest category for city size. The remaining distribution of city sizes is depicted in blue. Above each bar is the percent of observations corresponding to a city size within that interval. Of the 2,120 distinct cities represented, 49.6% have an assumed population of 500 sometime in their panel.

Figure 3 provides a next step in examining the raw data. The horizontal axis is the change in temperature from one time period to the next. The vertical axis is the difference in the logarithm of city size from one period to the next. Red markers represent changes where either the previous or current period has an assumed population. Blue markers represent the remainder of observations in Bairoch et al. (1988). From the red markers in the figure, we can see a positive correlation between changes in temperature and changes in city size when population must be assumed. In particular, the largest changes in the data stem from observations with populations in the largest, leftmost, and very coarse population bin corresponding to the smallest city sizes. For the remaining markers (in blue) which represent always-large cities, there seems to be little correlation between changes in temperature and change in city size, although saturation of the graph (as well as the many variables included in the later extensive suite of fixed effects) may be concealing a correlation.

3.1 Accounting for Heterogeneity by City Size

In this section we provide results accounting for city size heterogeneity, using otherwise the same data and codes provided by Waldinger (2022). Small cities, which make up 23.5% of observations and 49.6% of cities

[5] This section can also be seen as an extension of Table 14 in the October 2015 working paper version of Waldinger (2022), which offers robustness of the results to assuming a city size of 1, 500, and 1000 whenever censoring occurs (not included in the final published version nor in its online appendices).

[6] Eeckhout (2004) refined Gibrat's Law for city populations into two facets: that city growth is proportionate and that the non-truncated distribution is lognormal rather than Pareto.

at some point in their panel, had a population between 0 and 999 according to data provided by Bairoch et al. (1988). Accounting for this 'left censoring' (in that the left hand of the city size tail is collapsed into a single value) offers additional nuance to the relationship between temperature and city size.

In Table 2, we use the same data and codes as in Table 1 with one exception. Here, we have created an indicator variable that takes a value of one if a city ever has a left-censored population. Effectively, this indicator variable divides cities into those in which Bairoch et al. (1988) provides more detailed interval data and those where a population must be assumed due to left-censoring. Because of the nature of the censoring, we can also note that these are 'large' and 'small' cities, respectively. We find that 'always large' cities exhibit a *negative*, large and statistically significant relationship with previous-century temperature, implying that warmer temperatures resulted in reduced large city populations. For small cities, the summation of the $Temp$ and $Temp \times Pop.Ever < 1000$ coefficients produces a positive, large, and statistically significant relationship with previous-century temperature.[7]

3.2 Accounting for Censorship Using Eeckhout (2004)

In this section we show that the main result's magnitude and statistical significance is not driven by the fact that much of the data requires an assumed population; even when the smallest populations are reconstructed using lessons learned from Eeckhout (2004) instead of assigned a constant integer, the main results of Waldinger (2022) remain robust.

In the top row of Figure 4, we begin by reproducing Figure 2 and Figure 3 while changing their horizontal axes to be common with the succeeding row.

In the second row, the first panel shows a histogram of ln(CitySize) when each city-year that required an assumed population is instead drawn from the [1,999] interval, weighted by probabilities informed by a fitted log-normal distribution (following the observation that human city sizes follow a log-normal distribution in Eeckhout (2004)) . With that observation, we recover the mean and variance of the uncensored distribution by applying an off-the-shelf tobit estimator to the left-censored ln(Citysize) data. This left-tail reconstruction allows us to examine whether it is the assumed constant number for censored cities that drives the original analysis' main results. The right panel displays the scatterplot of changes by city size status (in the spirit of Figure 3). When moving from the constant 500 assumption distribution to drawing from the now uncensored distribution, it seems that the red markers (small cities) still support an upward sloping regression line, while the blue markers support either one that is flat or negatively sloped. In the third row, the panels display the estimated regression coefficients and associated t-statistics following the specification from Table 1 column 2 following 1000 random draws from the now 'uncensored' distribution. Note that we do not require all cities in every regression to have the same draw - each city is given its own draw from the fitted distribution. Regressions following many of these draws then return a *higher* coefficient than the originally published estimate, leading us to consider that the large share of left-censored observations assigned a constant value is not driving the estimated effects. In the right panel, where we present the associated t-statistics, all of them are statistically significant at the 5% level.

[7] This unexplored heterogeneity may provide interesting corollary evidence that people migrated away from large cities when temperature is benign and toward large cities as temperatures fall, consistent with Dribe (2003) and the migration mechanism discussed in Waldinger (2022)'s Section IV, which uses data from English parishes between 1571 and 1871 to find that places with relatively more benign temperatures attracted more 'migrant marriages' than places with cooler temperatures.

4 Conclusion

The main result of Waldinger (2022) is "a positive relationship between mean temperature and city size" (p. 2289). The article provides empirical evidence of this by combining data from Luterbacher et al. (2004) for paleo-climatic temperatures and Bairoch et al. (1988) for city size. However, city size is not very detailed in the source data—for the 23.5% of observations where data is censored below 1000, Waldinger (2022) must assume the city has some number of inhabitants, ultimately choosing 500.

In this replication, we have found differing relationships between temperature and city size for large and small cities. For large cities (those with populations over 1000) the relationship between temperature and city size is negative and statistically significant while for small cities the relationship between temperature and city size is (probably) positive.

We also find that if the leftmost values (those representing populations between 0-999 in the source data) had instead been drawn from a fitted log-normal distribution (following the observation that human city sizes follow a log-normal distribution in Eeckhout (2004)), the originally estimated effect of temperature and its statistical significance seems robust.

We fully agree with Waldinger (2022) "that long-term temperature changes have important effects on economies." (p. 2310) and that such effects are important to understand. We hope more research will be conducted in this area, which excellently complements existing temperature-economy estimates based on forward-looking climate projections.

Figures

Figure 1: Results of Any Constant Assumption

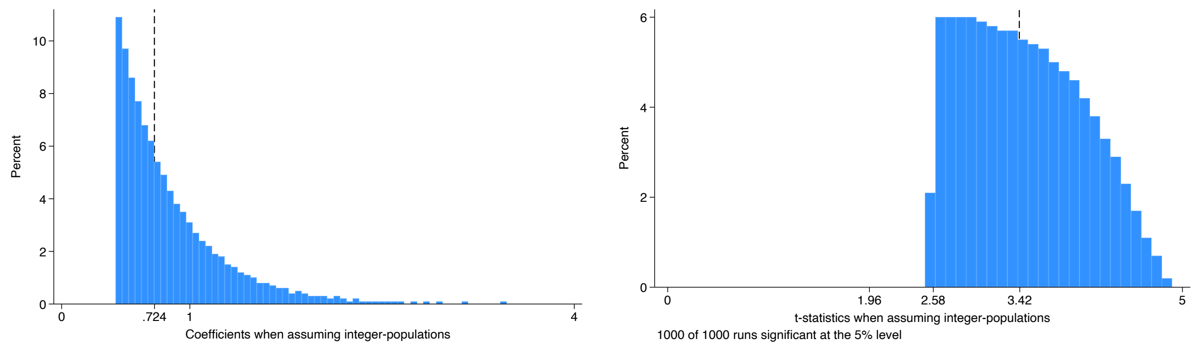

Notes: This figure presents the estimated coefficients (left panel) and associated t-statistics (right panel) of the $MeanTemperature$ regression term of column 2 in Table 1. Left panel includes a dashed line corresponding to the published coefficient 0.724. Right panel includes a dashed line corresponding to the published associated t-statistic of 3.42. Both represent 1000 regressions where the assumed population for the censored data is the integers 1, 2, 3, ... 1000. Assuming 500 corresponds to Waldinger (2022).

Figure 2: Distribution of City Size in Bairoch et al. (1988)

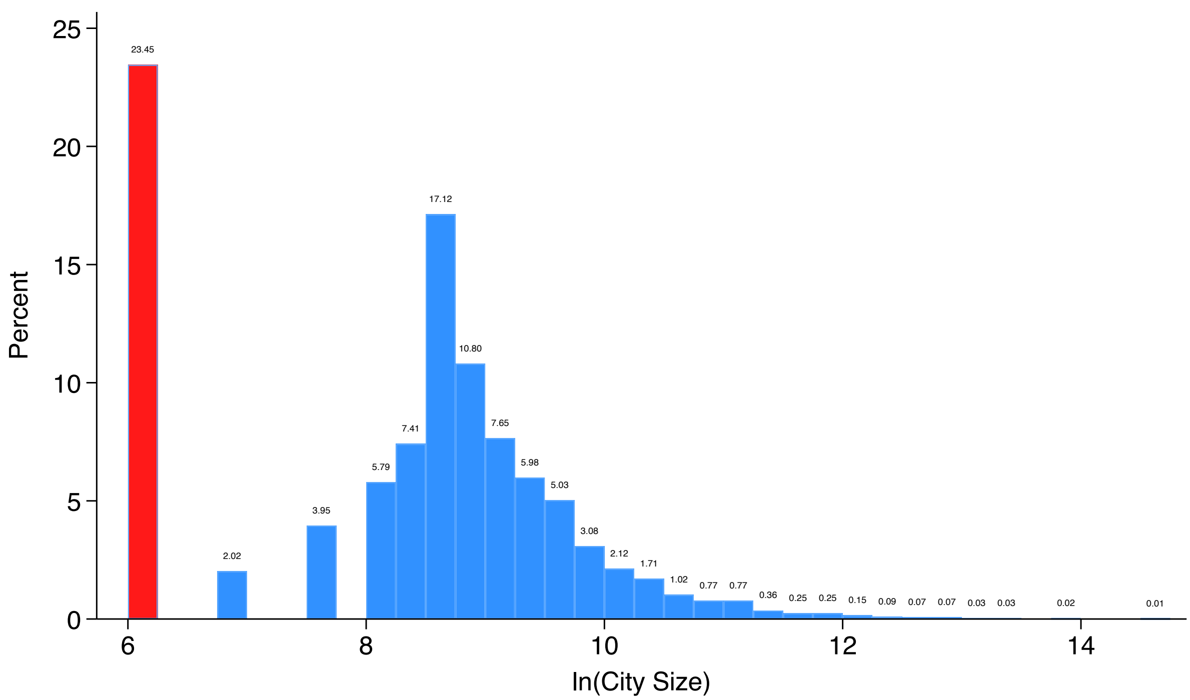

Notes: Best viewed in color. Red bar represents observations with assumed population. Blue bars represent remaining observations in Bairoch et al. (1988). Distribution of 10,600 observations of city size used in Table 1. The horizontal axis representing city size is provided in logarithmic scale. The vertical axis representing share of observations is provided in percent scale. Bars are assigned width of 0.25. Heights of bars in percent displayed on top of each. For example, 23.45% of observations are assigned ln(City Size)=6.21, consistent with an assumed population. A total of 2,120 distinct cities are used in the analysis. A total of 1,052 (49.6%) have an assumed population in at least one time period.

Figure 3: Scatter Plot of Temperature and City Size Changes, by City Size

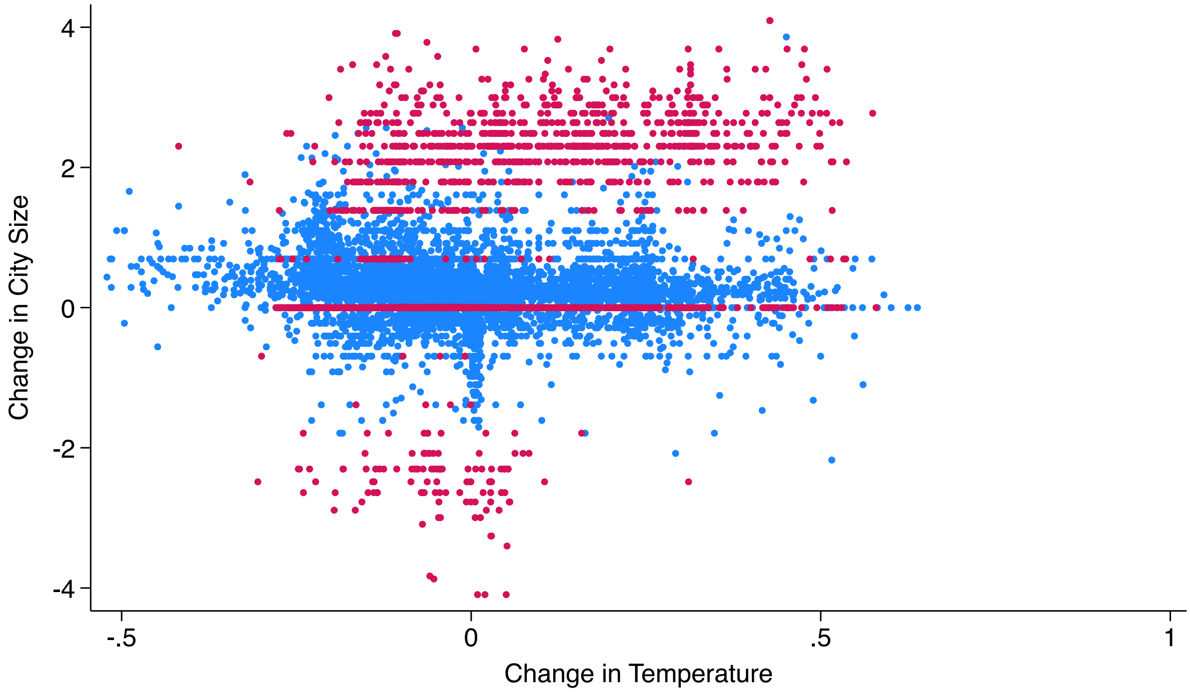

Notes: Best viewed in color. Red markers represent changes in population where either the current period or previous period had a population below 1000. Blue markers represent changes in population where both periods contained data from Bairoch et al. (1988). Horizontal axis represents the change in temperature between time periods. Vertical axis represents the change in ln(City Size) between time periods.

Figure 4: Re-Constructing the Uncensored Distributions

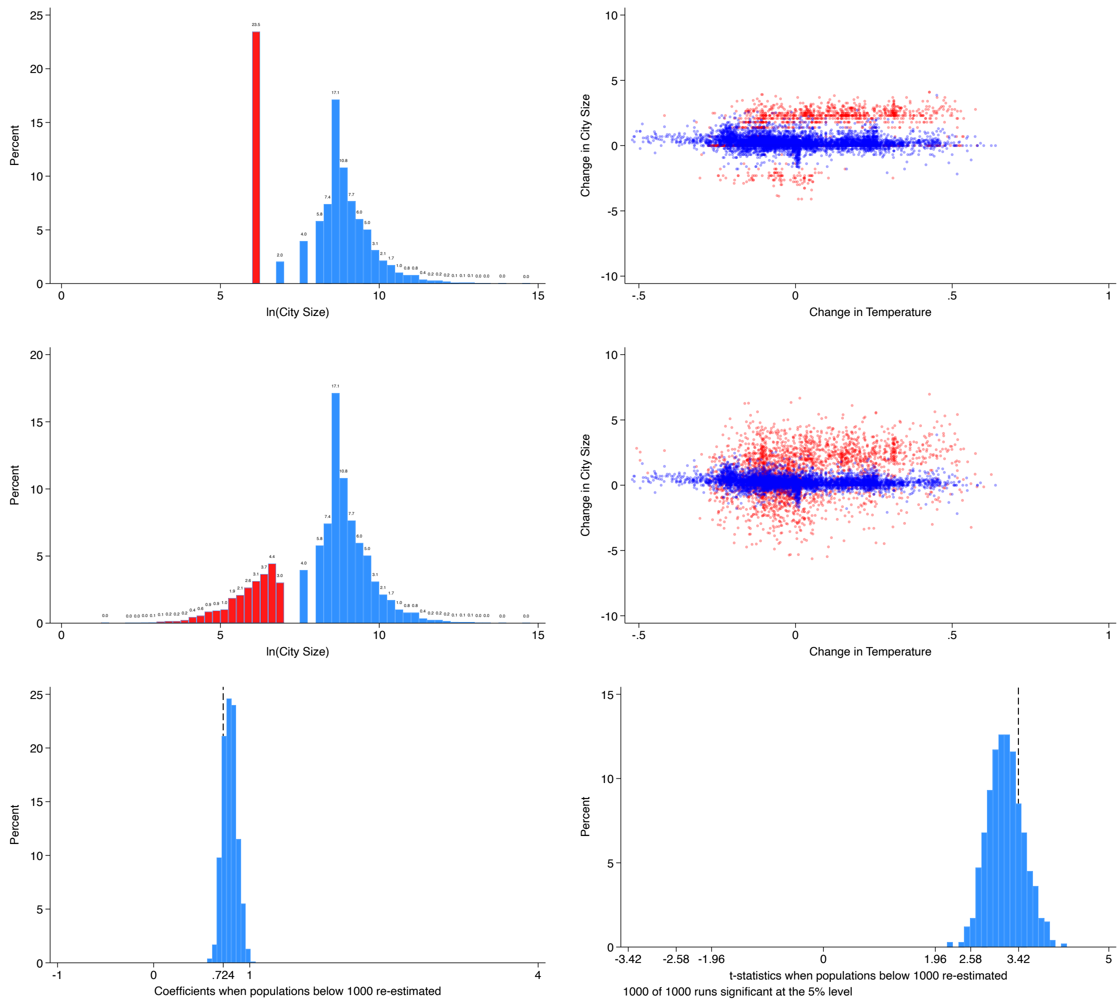

Notes: Best viewed in color. In the first row left panel, we reproduce ln(citysize) as in Waldinger (2022) and our Figure 2. Red represents city sizes set to 500 given data source's left censoring of populations between 1 and 1000 (Bairoch et al., 1988). In the first row right panel, we reproduce Figure 3. Red represents changes in city size where at least one of the two periods have a left-censored value. Blue represents changes in city size where neither periods are left-censored in Bairoch et al. (1988). In the second row left panel, we assign populations to left-censored city sizes using a fitted log-normal distribution following (Eeckhout, 2004), and repeat the first row left panel otherwise. In the second row right panel, we repeat the first row right panel, where red markers now indicate the new fitted city sizes. The third row left panel presents a histogram of estimated coefficients from 1000 regressions corresponding to the preferred specification of Waldinger (2022) using unique draws from the fitted log normal instead of a constant integer, with a vertical reference line provided at the published coefficient of 0.724 as in Table 1. The third row right panel presents a histogram of the associated t-statistics, with a vertical reference line provided at $t = 3.42$ as in Table 1.

5 Tables

Table 1: Temperature and City Size (Computational Reproduction of Waldinger (2022) Table 2)

	log CITY SIZE					
	(1)	(2)	(3)	(4)	(5)	(6)
Mean Temperature	0.532	0.724	0.749	0.931	0.842	1.170
Standard error clusters:						
Spatial and serial autocorrelation	(0.262)**	(0.268)***	(0.269)***	(0.328)***	(0.265)***	(0.429)***
	[0.042]	[0.007]	[0.005]	[0.005]	[0.001]	[0.006]
Two-way (city and region × time period)	(0.281)*	(0.282)**	(0.283)**	(0.323)***	(0.276)***	(0.465)**
	[0.064]	[0.013]	[0.011]	[0.006]	[0.003]	[0.015]
Temperature grid	(0.193)***	(0.212)***	(0.213)***	(0.274)***	(0.211)***	(0.386)***
	[0.006]	[0.001]	[0.000]	[0.001]	[0.000]	[0.002]
City fixed effects	Y	Y	Y	Y	Y	Y
Time period fixed effects	Y	Y	Y	Y	Y	Y
Country in 1600 linear time trend					Y	
Country in 1600 × time period f.e.						Y
Historical controls (× time period f.e.)	Y	Y	Y	Y	Y	Y
Geographic controls (× time period f.e.)		Y	Y	Y	Y	Y
Sample	All	All	Excl. Capitals	Excl. Ocean	All	All
Observations	10,600	10,600	10,510	8,395	10,600	10,600
R^2	0.767	0.769	0.759	0.766	0.779	0.783

Specifications exactly the same as in Waldinger (2022) Table 2. All four authors independently replicated this table. Standard errors in parentheses. p-values in square brackets. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 2: Temperature and City Size (Heterogeneity by City Size)

	log CITY SIZE					
	(1)	(2)	(3)	(4)	(5)	(6)
Mean Temperature	-1.510	-1.367	-1.393	-1.450	-1.150	-0.886
Standard error clusters:						
Assuming spatial and serial autocorrelation	(0.399)***	(0.326)***	(0.326)***	(0.349)***	(0.340)***	(0.429)**
	[0.000]	[0.000]	[0.000]	[0.000]	[0.001]	[0.039]
Two-way (city and region × time period)	(0.375)***	(0.352)***	(0.355)***	(0.390)***	(0.370)***	(0.521)*
	[0.000]	[0.000]	[0.000]	[0.000]	[0.003]	[0.095]
Temperature grid	(0.189)***	(0.208)***	(0.210)***	(0.263)***	(0.203)***	(0.354)**
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.012]
Mean Temperature × Population Ever < 1000	2.945	2.955	2.992	3.178	2.789	2.785
Standard error clusters:						
Assuming spatial and serial autocorrelation	(0.511)***	(0.399)***	(0.398)***	(0.403)***	(0.398)***	(0.417)***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Two-way (city and region × time period)	(0.454)***	(0.443)***	(0.444)***	(0.431)***	(0.433)***	(0.453)***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Temperature grid	(0.163)***	(0.162)***	(0.163)***	(0.168)***	(0.165)***	(0.173)***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Control variables:						
City fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time period fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Country in 1600 linear time trend					Yes	
Country in 1600 × time period f.e.						Yes
Historical controls (× time period f.e.)	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls (× time period f.e.)		Yes	Yes	Yes	Yes	Yes
Sample	All	All	Excl. Capitals	Excl. Ocean	All	All
Observations	10,600	10,600	10,510	8,395	10,600	10,600
R^2	0.781	0.782	0.773	0.782	0.791	0.795

Specifications exactly the same as in Waldinger (2022) Table 2 (also our Table 1), excepting our inclusion of an interaction term between temperature and an indicator variable that takes a value of one if a city ever had a population below 1000. Standard errors in parentheses. p-values in square brackets. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

References

Bairoch, P., Batou, J., and Chevre, P. (1988). The population of european cities. *Data bank and short summary of results*, 8001850.

Dribe, M. (2003). Dealing with economic stress through migration: Lessons from nineteenth century rural sweden. *European Review of Economic History*, 7(3):271–299.

Eeckhout, J. (2004). Gibrat's law for (all) cities. *American Economic Review*, 94(5):1429–1451.

Luterbacher, J., Dietrich, D., Xoplaki, E., Grosjean, M., and Wanner, H. (2004). European seasonal and annual temperature variability, trends, and extremes since 1500. *Science*, 303(5663):1499–1503.

Waldinger, M. (2022). The economic effects of long-term climate change: evidence from the little ice age. *Journal of Political Economy*, 130(9):2275–2314.