The Long Shadow of a Sudden River

Luke Heath Milsom*

20th March, 2025

Abstract

In 1134 a biblical storm created a river, the Zwin, that connected Bruges on the Belgian coast to the North Sea. However, by 1500 it had silted up. Using comprehensive city-level population and real wage data this paper first shows that, whilst navigable, this sudden river caused the population of Bruges to be some 400% higher. Then, using a dynamic quantitative spatial economics model I show that even hundreds of years after it had silted up the shadow of the Zwin can still be seen in Bruges and neighboring cities. Aggregate welfare remains on average 2% higher across Belgium as compared to the counterfactual world in which the Zwin had never existed.

JEL: R12, R13, N9, O18

*Heath Milsom: KU Leuven and The Institute For Fiscal Studies, luke.heathmilsom@kuleuven.be. With thanks to: Hannah Zillessen, Ferdinand Rauch, Eric Buyst, Frederic Vermeulen, Max Merczinek, and Shihang Hou. Heyi Zheng provided excellent research assistance on this project.

Cities are engines of productivity and beacons of opportunity [Glaeser, 2011]. Understanding how the prevailing urban structure was born, whether a higher-welfare spatial distribution of economic activity exists, and the role shocks and policy have to play are thus crucial economic questions. Place-based policies and an explicit policy goal of spatial equality are also becoming increasingly common. The long-run efficacy of such (expensive) programs will depend on their long-run impact and whether they fundamentally shift the prevailing spatial equilibrium of economic activity.

In this paper, I leverage a dramatic natural experiment to better understand the role temporary historical shocks can play in determining the long-run distribution of economic activity [Bleakley and Lin, 2012, Davis and Weinstein, 2002]. In the 12th century, a biblical storm caused a large sea inlet, the Zwin, to form [Glaeser, 2022, Ryckaert, 1989].[1] This new river connected the outskirts of Bruges[2] to the north sea [Trachet et al., 2015] and the lucrative trade that flowed through it [Glaeser, 2022]. However, some three to four hundred years later the Zwin had completely silted up and was no longer navigable. The sudden appearance and (not quite so sudden) disappearance of the Zwin has been attributed to Bruges' rise and subsequent fall [Dumolyn and Leloup, 2016, Van Houtte, 1966]. In this paper, I first empirically substantiate this claim before asking whether this temporary and sudden river had any immediate or long-lasting impact on the spatial distribution of economic activity across the whole Low Countries.

Using updated historic population data based on the Bairoch et al. [1988] dataset I show that during the period when the Zwin was navigable Bruges experienced a population boom. I first show the effect in the raw data before performing a difference-in-differences and event-study analysis relying on the plausibly exogenous timing and location of the Zwin by using other potential-Zwin locations as controls. I find that the Zwin increased Bruges population by on average 30,000 during the period it was navigable, on a population in 1100 of only 7,000. To allow for heterogeneity in the impact of the Zwin, and overcome valid concerns surrounding the possibility of contemporaneous shocks confounding my reduced-

[1] Technically a sea inlet is not a river, but for our purposes, the two formations are equivalent and so in this paper I will describe the Zwin as a river.

[2] For simplicity and clarity throughout this paper I will use the English spelling of cities in the low countries. Specifically, I will use "Bruges" and "Antwerp" rather than their Dutch spellings "Brugge" and "Antwerpen" or French spellings "Bruges" or "Anvers".

form estimates, I also take a market access approach. I calculate the market access of each of 79 cities in Europe using Djikstras algorithm over a cost surface that takes into account terrain ruggedness, river, and sea travel. Using exogenous variation in market access due only to changes in travel costs caused by the Zwin, I show that increases in market access cause increases in population.

To ask what impact the Zwin had on the entire urban structure of the low countries, I build a dynamic quantitative spatial economics model. I extend the framework of Allen and Donaldson [2020] to account for a key feature of the Medieval and early modern European economy: Merchants. Locations interact through costly migration and costly trade, and the past impacts today through historical amenity and productivity spillovers. Merchants choose where to locate to maximize profits and directly impact local productivity.

I estimate the model using city-level data on population, real wages, grain prices, and novel data on merchant numbers. I use plausibly exogenous variation in population, wages, and Merchants, stemming from three large-scale historical events: the Little Ice Age [Waldinger, 2022], the Black Death [Jedwab et al., 2019], and the Opening of the Atlantic to trade [Acemoglu et al., 2005a]. Additionally, I use plausibly exogenous variation in market access due to the Zwin. Quantitative information on merchant numbers is however not available in historical records. However, significant qualitative information in primary and secondary sources exists. I leverage this and text analysis from a large language model to estimate quantitative merchant numbers across time and space in various European cities. In sum this approach allows me to estimate model parameters and transport costs specific to my setting. Armed with parameter estimates I can then invert the model to back out rationalizing fundamentals and estimate in a reduced form manner the direct impact the Zwin may have had on local productivity in Bruges. In this way, I can flexibly allow the Zwin to impact multilateral travel costs, bilateral travel costs, and productivity in Bruges directly.

Using the estimated model I first find that the Zwin caused Bruges' population to be four times larger in 1400, confirming the reduced-form results. However, I also find that the Zwin also caused Bruges' population to still be 20% larger in 1800, than if the Zwin had never appeared. This is in spite of the fact that by 1800 the Zwin had been impassable for at least 300 years. The Zwin also impacted other cities in the Low Countries, although to a

lesser extent. These variations in population are mirrored in location-specific counterfactual welfare. In terms of aggregate welfare across the whole of Belgium, I find a fairly modest impact of the Zwin when it was navigable and thereafter, although its impact can still be felt in 1800 where welfare is around 2% higher than if the Zwin had never existed.[3]

In this paper, I contribute to the broad literature on understanding the short and long-run impact of shocks to urban systems and the spatial distribution of economic activity, and the intersection of history and urban economics more generally [Hanlon and Heblich, 2022, Nunn, 2020]. One strand of literature has considered the contrasting role fundamentals and path dependence have played in determining the observed distribution of economic activity [Lin and Rauch, 2022, Alesina et al., 2013]. Using exogenous variation in city size due to some disaster Davis and Weinstein [2002], Miguel and Roland [2011], Jedwab et al. [2019] for example, find that city-size rebounds quickly and interpret this as evidence for fundamentals. Turning to path-dependence, Bleakley and Lin [2012] and Jedwab et al. [2017] find evidence of persistence in the face of infrastructure obsoletism, which is evidence for path-dependence. However, Michaels and Rauch [2018], Redding and Sturm [2008], Gibbons et al. [2018] and Dell [2010] find evidence that large shocks can change long-run outcomes, providing evidence against path dependence. In this paper, I contribute to this debate and reconcile previously contrasting results by showing that shocks can have persistent effects over a long time frame but that ultimately (in my case) fundamentals determine the long-run spatial equilibrium. Secondly, I contribute to a nascent literature using quantitative economic geography models to understand the spatial impact of historical events [Nagy, 2022, Eckert and Peters, 2022, Ellingsen, 2021, Heblich et al., 2020]. I build on this literature by considering the dynamic impact of a temporary shock, combining counterfactual results with reduced-form evidence, and estimating setting-specific model parameters using plausibly exogenous variation.

The rest of this paper proceeds as follows: Section 1 describes the historical context and data, 2 performs the reduced form, and market access based analysis, section 3 sets up the dynamic quantitative spatial economics model, section 4 then describes the counterfactual

[3]A limitation of this approach is that I can trace the endogenous long-run impact of the Zwin through its impact on trade and migration, but not on other aspects such as political economy. The period I study is volatile, but the various conflicts are taken as exogenous and are not allowed to vary endogenously with local economic conditions.

results, finally section 5 concludes.

1 Historical context and data

1.1 Historical context

The hydrological, social, and economic history of the low countries and in particular the area between Bruges and the North Sea is a topic of long and still active academic debate (see Trachet et al. [2015] for a summary). Although this land appears stable and dry today, prior to modern diking, damming, and canals before 1800 or so it was a more volatile landscape. This was further exacerbated by a series of large tidal transgressions (known as Dunkirk I to IIIb) causing this low landscape to be crisscrossed with rivers, marshes, and routinely flooded areas [Soens et al., 2014]. In this historic landscape, it is possible to imagine a large storm creating a 15km inlet connecting the north sear to near Bruges — although this must still have been a remarkable event.

Although recent work has also emphasized the role of inclusive institutions in Bruges success, it remains clear that without access to the north sea, through the Zwin, Bruges would never have been able to succeed in the way it did between 1200 and 1500 [Charlier, 2011, Dumolyn and Leloup, 2016, Lambert, 2016, Ryckaert, 1989, Charlier, 2005, Dewilde et al., 2018, Glaeser, 2022]. Recent work has also highlighted the role of human activity directed by Bruge attempting to keep the Zwin navigable for as long as possible[4]. However, such efforts were completely in vain, as the necessary dredging technology simply didn't exist and 1500 the sandbanks at the entrance to the Zwin had become so large that only small ships could enter it Dewilde et al. [2018].

It's important to highlight that over the study period a lot more was happening around Bruges, and in the low countries in general, than just the advent and decline of the Zwin.

[4]For example Dumolyn and Leloup [2016] notes: "From about 1400 a pilotage service was organized to lead larger vessels around the sandbanks and in 1456 a signalization system was installed. During the fifteenth century, Bruges also deployed two dredge boats, one of which was aptly named the 'mole'. In the latter fifteenth century Bruges even invested in major hydraulic works to scourge the Zwin but to no advance; by 1500 the sandbanks in the Zwin estuary had become so large that only small ships could enter and navigate the stream. In 1486 only 73 ships called in at the port of Sluis while fifteen years later this number had dropped to merely 36".

Undoubtedly Bruges' institutions, diplomatic position, and outports (Sluis and Damme) amongst other features also factored into its success [Gelderblom, 2015, Van Houtte, 1966]. In this paper, I will not attempt to give a complete analysis of each of these factors but rather focus on the exogenous appearance and disappearance of the Zwin. One limitation of this work is that due to this I will bundle together factors influencing Bruges contemporaneously with the Zwin and can only estimate their combined effect. The advantage of using a dynamic quantitative spatial economics model is that I can estimate the impact of the Zwin conditioning both on economic shocks stemming from changes to the urban network, and on other events affecting other periods or geographies such as the 100 Years War, the 80 Years War, the advent of the Dukes of Burgundy, or the fall of Antwerp. This disadvantage is that such events are not allowed to endogenously occur in counterfactual scenarios.

1.2 Data

Europe from 700 to 1800 is an extremely data-scare environment. Systematic information at a geographically disaggregated level over this period is hard to come by. In this paper I rely on three main sources of data. The primary source, and by far the largest time series, is information on city-specific population estimates. I construct this dataset by combining the well-known Bairoch et al. [1988] data with updated and more credible estimates in key cities and periods. I combine this data with information for a subset of cities and time periods on real wages from Allen [2001] and estimated merchant numbers.

1.2.1 Population estimates

The main data for this paper are historical city-specific population estimates. I create a dataset of 652 cities across Europe (not including Russia) from three main sources. I use as my baseline the Bairoch et al. [1988] database which is well known as used, see for example Bosker et al. [2013], Acemoglu et al. [2005b], Glaeser [2022], Duranton and Puga [2020].[5] I combine this data with that from Chandler [1987] and Buringh [2021]. Where the three data sources overlap but disagree I take the average over the two which don't disagree. If

[5] Hanlon and Heblich [2022] show, using data from EconLit that the Bairoch et al. [1988] has been cited 172 times between 1992 and 2016 in economics papers (with 23 citations in "top 5" journals.

there is ambiguity I take the value from Buringh [2021] if it comes with a documented and reasonable source, and otherwise use as my baseline Bairoch et al. [1988]. I then make adjustments where there are known issues with the datasets described above and or more detailed research exists for specific locations which is contradictory. I revise the population of Italian cities following Malanima [1998], and Spanish cities following Bosker et al. [2013].

Particularly important in my setting it is well known that the Bairoch et al. [1988] estimate for Bruges in 1400 is a significant overestimate [Bosker et al., 2013]. Bruges' high-point is a particularly important estimate for this project, but one that is difficult to pin down due to rapid growth and periodic famines [Espeel, 2016]. I follow an authoritative modern source Stabel et al. [2018] who estimates a population of around 60,000 at the end of the 13th century in Bruges (around half of the estimate provided by Bairoch et al. [1988]). This estimate is based on earlier work by Prevenier [1975] who uses data on draft lists for the city militia. Such draft lists are informative as craft guilds and merchants had to deliver a certain proportion of their population to the city militia. Stabel et al. [2018] then add to this an estimate of the number of non-guild or otherwise exempt persons based on a conservative estimate of the proportion of the population such a group would normally consist of.

1.2.2 Real wages

I use information from Allen [2001] to estimate real wages in 18 European cities from 1500 to 1800. In each city these wages pertain to the purchasing power of building craftsmen, a common and fairly homogeneous group of workers who Allen [2001] argues is representative of broader wage patterns. Allen [2001] deflates nominal wages (expressed in grams of silver to allow international comparison) by a price index. Allen [2001] uses the Laspeyres index fixing quantities consumed but allowing prices to vary across cities and time. He uses information on twelve consumption goods (bread, beans/peas, meat, butter, cheese, eggs, beer, soap, linen, candles, lamp oil, and fuel). One notable omission is rent, justified by a lack of data and low importance in the pre-modern budget share (around 5% for the English working class in the eighteenth century). Using this information Allen [2001] reports welfare ratios (average earnings divided by the cost of poverty-line consumption), this is the data I use in this paper.

1.2.3 Grain prices

I use information from various sources [Federico et al., 2021, Shiue and Keller, 2007, Madsen et al., 2024] on grain prices at the city-year level. This information is crucial for understanding market integration across Europe over time and estimating how bilateral trade costs changed. It will allow us to estimate the impact of the Zwin on Bruges trade costs in a non-parametric fashion.

1.2.4 Weather, Black Death, and Atlantic traders

To credibly estimate model parameters I leverage sources of plausibly exogenous variation in local wages and population. For this I rely on three large scale historical events: the Little Ice Age, the Black Death, and the Opening of the Atlantic to trade.

First, I follow Waldinger [2022] and leverage temperature variation across time and cities due to the climatic event known as the Little Ice Age. Between about 1400 and 1900 temperatures in Europe dropped by as much as one degree on average relative to their previous baseline which the most extreme drops occurring in the 16th and 17th centuries. Second, I follow Acemoglu et al. [2005a] and use variation in the timing and impact of the Openning of the Atlantic for trade across European cities. Finally, following Jedwab et al. [2019] and others I use variation in deaths due to the black death and other plagues across cities and time.

1.2.5 Estimating merchant activity

There exists little to no precise quantitative information on merchant activity in the low countries and surrounding area from 1000 to 1800 [Puttevils, 2012] (due to the complete lack of merchant activity prior to 1000 here I focus on years post-1000). This is perhaps unsurprising as merchant populations were transitive and often seasonal. For example, in a rare peice of quantitative information in 1438 the mayor of Bruges, Nicolar Despars, describes the number of merchants taking place in a procession upon the arrival of the Duke of Burgundy (Murray [2005]). He mentions 136 Hansards, 150 Italians, and 48 Spanish. However, this is likely to be an incomplete list, and as the procession took place in December

at the low-point for merchant activity, a significant under estimate. Another example comes from the records of the money changer Collard de Marke who's books suggest over 400 Flemish and 500 foreign merchants were present in Bruges in the 1360's [Murray, 2005]. However, its unclear whether this information pertains to a single year or covers many, or whether these merchants merely passed through or were permanent residents. Nor is it know what proportion of the relevant merchant population one could expect to see on Marke's books. With this sparsity and unreliability of known quantitative information, I turn to less traditional sources to build a more comprehensive picture.

I leverage textual data and large language models to provide estimates every 100 years for the 20 biggest cities around Bruges in 1400 and Amsterdam. I use ChatGPT 4o Deep Research large language model[6]. In the appendix XX I provide a detailed account of the references used by ChatGPT and an abridged version of its qualitative discussion of the evolution of merchant numbers over this time. Leveraging a large language model in this way has the advantage of providing quantitative information on well known qualitative trends parsing and combining data from a vast number of sources. Additionally, using the output directly removes researcher degrees of freedom in attempting to compile such a dataset. However, this approach obviously comes with significant caveats, although where possible sources and facts have been checked, the processes generating results remains a total black box.

I take three approaches to validating the constructed merchant data. As discussed above direct validation is impossible, indeed if direct validation were possible it would not be necessary to take this approach. However, we can check whether the constructed data passes the proverbial "sniff-test". First, I show that the constructed estimates are close to the four data points found in historical sources for the number of merchants. Second, I show that the evolution of merchant numbers across cities confirms the qualitative patterns one would expect from the literature. Third, I follow the same procedure to generate population data, for which I have data, and compare the generated and historical sources and find a raw correlation of 0.77. Details can be found in sub-section E.3 in the appendix.

[6]You can find the exact prompt used here https://chatgpt.com/share/67a78014-6f84-800f-b86c-7d16d1196a5c.

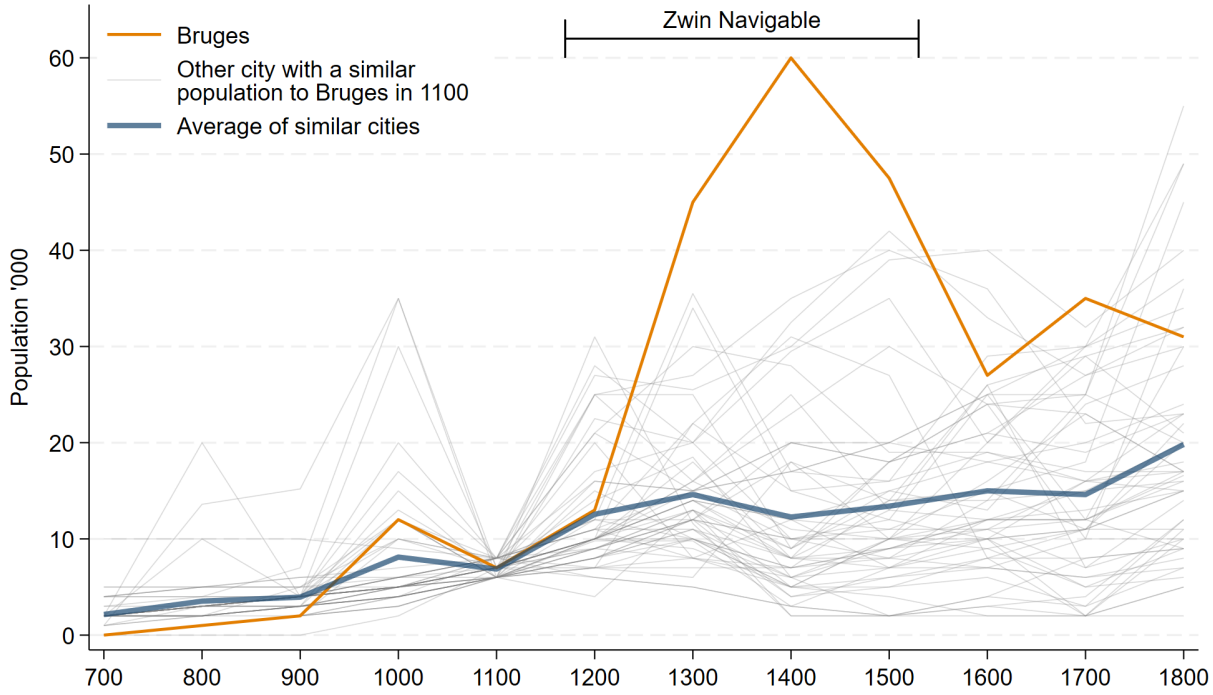
2 Reduced form impact of the Zwin

Figure 1 stays as close to the data as possible, and plots the population of Bruges against that of similar cities from 700 to 1800. Similar cities are defined simply as those having a similar population to Bruges in 1100, the year before the Zwin suddenly appeared. Figure 1 shows a large population spike for Bruges (plotted in orange) during the period that the Zwin is navigable, far in excess of any other similar city (other cities are plotted in gray, and the average over such cities in blue). Figure 1 clearly shows a large increase in Bruges population after the Zwin became navigable, as well as a large decrease after it was no longer navigable.[7] Indeed, out of all of the 643 cities in my sample the population of Bruges increased by more than any other city between 1200 and 1400, and decreased more than all but one city between 1400 and 1600[8]. Of course, figures such as 1 capture all changes in Bruges between 1200 and 1500, not just the impact of the Zwin. Thus, the impact of events such as the Zwin cannot be separately identified from other location-specific shocks that occurred in this time period. However, it does seem unlikely that any other large-event could explain a large proportion of the pattern of boom and bust shown in figure 1.

[7]The navigability of the Zwin was constantly changing in degrees, indeed almost immediately after appearing the Zwin began silting up Soens et al. [2014], however before 1500 it was still possible to get a large ship sufficiently far down the Zwin, whereas after 1500 this was no longer possible Dewilde et al. [2018] — and thus I have chosen 1500 as my cut off here.

[8]The one city who's population fell faster was Gent, popularly known as the factory of Bruges, where much of the Flemmish cloth was produced that was subsequently sold at Bruges.

Figure 1 Reduced form evidence



Notes: This figure plots the raw population data over time for Bruges (in orange) and each of 55 European cities which have a similar population to Bruges in 1100 (in thin gray). The (unweighted) average over these similar cities is given in thick blue.

To more formally assess the impact of the Zwin on Bruges whilst it was navigable, I perform a difference in differences analysis estimating equations of the form given in 1. In equation 1 Zwin_{it} takes the value one if the Zwin is navigable in location i in period t, and therefore only takes the value one between 1200 and 1500 (inclusive) in Bruges. City and year fixed effects are denoted by γ_i and τ_t, and idiosyncratic errors by ε_{it}.

$$\text{Pop}_{it} = \beta \cdot \text{Zwin}_{it} + \gamma_i + \tau_t + \varepsilon_{it} \tag{1}$$

This difference-in-differences specification relies on the perennial parallel trends assumption to make causal claims (as well as no contemporaneous shocks as discussed above). To increase the reasonableness of this assumption I consider two alternative control groups. First, I consider cities with a similar population to Bruges in 1100 — exactly those cities displayed in figure 1. Second, I consider 11 cities close to the coast near Bruges where plausibly the Zwin could have appeared instead of appearing in Bruges. This sample aims to leverage the randomness of the Zwin specifically, by comparing the trajectory of Bruges to

other cities that could plausibly been shocked.

In table 2 I show the estimated beta coefficient from equation 1 on the all sample and the two sub-samples discussed above. The impact of the Zwin when it was navigable is very similar across each sample/ specification. This methodology suggests that, during the period it was navigable, the Zwin increased the population of Bruges by between 24 and 33 thousand. On an 1100 population of 7,000, this is a large increase.

Table 1 Direct Impact of the Zwin on Population

	(1) All Sample	(2) All Sample	(3) Similar 1100 Population Sample	(4) Possible Alternative Rivers Sample
Zwin Navigable	33.20***	27.67***	23.52***	29.92***
	(0.486)	(0.379)	(0.896)	(8.266)
Year FE		X	X	X
City FE		X	X	X
# Cities	652	652	56	12
Observations	7824	7824	672	144
R2	0.00126	0.436	0.448	0.745

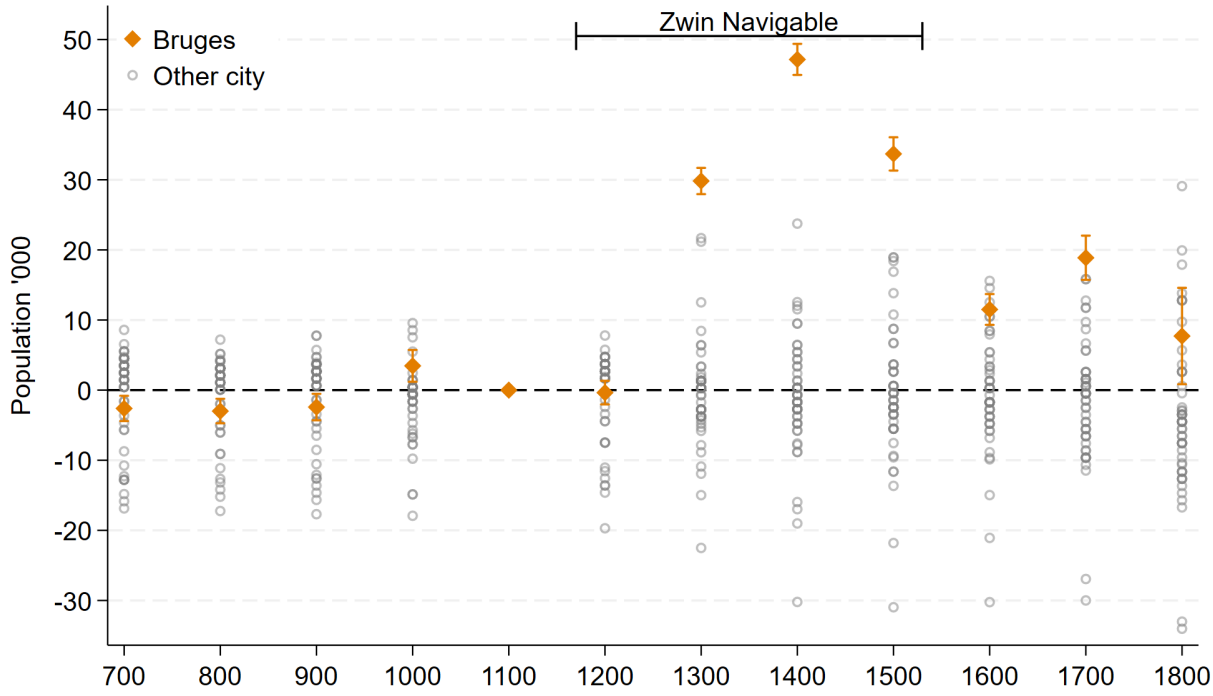
Notes: This table reports results from estimating the impact of the Zwin being navigable in Bruges. In column one I report raw estimates with no controls. In column two I include city and time fixed effects, therefore performing a two-way fixed effects analysis. In column three I restrict the sample to the 55 European cities which have a similar population to Bruges in 1100 (and Bruges itself). In column four I further restrict the sample to only include 11 cities which could plausibly have been shocked instead of Bruges. In all specifications standard errors are clustered at the city level.

To somewhat probe the parallel trends assumption, figure 2 shows the corresponding event study plot. In orange (with corresponding orange 95% confidence intervals) I plot the event study coefficients for Bruges relative to 1100. Figure 2 mirrors that of 1 showing no real pre-trends to speak of, and large impact of the Zwin. Figure 2 also gives some indicative evidence that the impact of the Zwin continued event long after it became impassible, with fairly large, positive, and significant coefficients in 1600, 1700, and 1800.

Figure 2 also shows placebo estimates from each of the 55 cities in my sample with a similar population to Bruges in 1100 (the same cities as shown in figure 1). To generate these estimates I estimate the same event study specification but suppose that each of the 55 alternative cities had been "treated" by the Zwin as opposed to Bruges. These estimates act as a placebo, as we would not expect any impact in these locations. Therefore because

the Bruges coefficients are much in excess of any of the placebo coefficients in 1300, 1400, and 1500 (as well as being fairly extreme in subsequent centuries), I take this as strong evidence that the impact of the Zwin as opposed to noise is being picked up. This approach will also account for spatial correlation, and thus provides a strong basis for inference. This exercise also suggests that no contemporaneous shock in any other city is anywhere near the same magnitude as that of the Zwin, implying once again that its unlikely such a shock affected Bruges. However we cannot rule it out, and its likely that some of the impact attributed to the Zwin is due to an alternative contemporaneous shock. In appendix C I perform a synthetic control analysis and find confirmative results.

Figure 2 Event study with placebos



Notes: This figure plots the event study coefficients for Bruges in orange (diamonds) with corresponding 95% confidence intervals indicated by vertical bars and calculated by clustering standard errors at the city level. In grey hollow circles, I plot the analogous event study estimated for each of 55 placebo cities which had a similar population to Bruges in 1100.

Figures 1 and 2 and table 2 show indicative evidence that the Zwin had a large impact on the population of Bruges. However, one should be cautious in causally interpreting the entire effect uncovered in this section. This is for three main reasons. First, contemporaneous shock cannot be ruled out, and although it seems improbable that they account for a significant

proportion of the estimated impact, it is eminently plausible that they account for some proportion. Second, over such a long time period with one treated unit the parallel (post) trends assumption becomes increasingly implausible. Evidence from using similar locations or alternative Zwin sites gives some reassurance by of course the assumption is impossible to test directly. Lastly, and crucially, the Zwin did not only impact Bruge. In historical sources it is attributed with causing Bruges's success, but also Antwerp, Bruges main rivals, decline. Similarly, cities well connected to Bruges (but not in competition for sea-born trade), gained by being able to trade with the increasingly rich and large hub of Bruges. These spillover effects will contaminate control cities in the analysis performed in this section. Again, it seems unlikely that the bias is sufficient to overturn the conclusions reached, but undoubtedly it should dampen them.

2.1 A market access approach

To overcome the issue of spillover effects contaminating control cities, estimate heterogeneous impacts, and crucially side-step the contemporaneous shocks critique, I turn now to a market access approach. The market access approach estimates the impact of the arrival (and silting up) of the Zwin on each location's market access which is then related to their population. There is a long and well-developed literature showing a strong empirical link in various settings between local economic activity and market access (see for example Donaldson and Hornbeck [2016], Redding and Sturm [2008], Jedwab and Storeygard [2022]) backed up by theory showing the generality of the market access approach Allen et al. [2020b].

The market access approach is intuitively attractive as it directly uses the plausibly exogenous variation in trade costs generated by Zwin. It is exactly for this reason that it can overcome the two main critiques of the event-study / difference-in-difference reduced form design: spillovers and contemporaneous shocks. By specifying the empirical nature of the shock, and linking it via theory to outcomes I trade off leveraging more precise variation caused by the Zwin with adding greater structure to the data.

The market access a given location i enjoys is given by $MA_i = \sum_j \tau_{ij}^{\phi} L_j MA_j^{-1}$. The cost of trading between i and j is given by τ_{ij}^{ϕ}, where τ_{ij} is the iceberg trade cost between i and j intuitively the proportion of the good that "melts" due to transit, therefore $\tau_{ij} \geq 1$. The

elasticity of trade with respect to the iceberg trade cost is given by ϕ. A location will have higher market access if it is close to (low τ_{ij}^{ϕ}) many large markets (large L_j) which themselves do not have many alternative trading partners (low MA_j).

I parameterise bilateral iceberg costs as depending on the travel time between each of the 652 cities in my sample. To do this I first construct a cost surface over the land, river, and sea routes of Europe by splitting the map of Europe into 1km pixels and calculating the cost of traversing each pixel as the inverse of travel time. Travel time across the sea is normalised to 1. Using Cogs (the most commonly available ship in the Middle Ages before Caravals) this is calibrated to about 15km/h on average. Travel time in rivers is calibrated to 1.25, i.e. it takes 25% longer to travel on a river than in the open sea. Travel on land is calibrated to 3 on flat land, corresponding to 5km/h. It increases with terrain ruggedness such that speed relative to ocean travel $= 2 + \exp(0.002 \times \text{Ruggedness})$. Over the most rugged terrain, this corresponds to a speed of around 500 meters an hour. Figure 3 displays the resulting cost surface.

Figure 3 Cost Surface



Notes: This figure shows the cost surface and location of cities in my sample. Darker colors indicate lower costs of travel.

The least cost path from/ to each of the 652 cities is then calculated over the cost surface using Djikstras algorithm. I calculate such travel costs in two scenarios, one without the Zwin and one with the Zwin. This process calculates two pair-wise travel time matrices $T^{NZ} = \{t_{ij}^{NW}\}$ and $T^{Z} = \{t_{ij}^{Z}\}$. I parameterise iceberg travel costs to be a log-linear function of travel times such that $\tau_{ij}^{\phi} = t_{ij}^{\kappa\phi}$. In the absence of data on bilateral trade flows in my setting, I turn to the literature for plausible estimates. Estimates of $\kappa \cdot \phi$ covering my historical setting are hard to come by, I use the estimate of -3.12 from Ellingsen [2021] who causally estimates this using plausibly exogenous variation in cross Atlantic trade in the 1800's. Given data on τ_{ij}^{ϕ} and L_j market access terms can then be calculated via a simple iterative algorithm which has a unique solution (up to scale) [Donaldson and Hornbeck,

2016]. In this manner, I can estimate market access terms for each city in each period where travel costs only vary with the Zwin, $\{MA_{it}\}$.

Using the calculated market access terms I can then relate changes in market access to changes in local population. I focus on the period from 1000 to 1400 to zoom in on the period where the Zwin had the largest impact, and there were no major transport innovations either over land or sea and therefore other than the Zwin travel costs remained fairly constant. I estimate equation 2.

$$\Delta_{1000,1400} \ln(L_i) = \alpha + \beta \cdot \Delta_{1000,1400} \ln(MA_i) + \varepsilon_i \tag{2}$$

Market access terms are however endogenous as they depend on contemporaneous populations. To overcome this I instrument changes in market access with those due only to changes in travel times caused by the Zwin, given in equation 3.

$$\Delta_{1000,1400} \ln\left(MA_i^{IV}\right) = \ln\left(MA_{i,1400}^{IV}\right) - \ln\left(MA_{i,1000}^{IV}\right) \tag{3}$$

$$= \ln\left(\sum_{j \neq i} \tau_{ij,1400}^{\phi} \frac{L_{j,1000}}{MA_{j,1400}^{IV}}\right) - \ln\left(\sum_{j \neq i} \tau_{ij,1000}^{\phi} \frac{L_{j,1000}}{MA_{j,1000}^{IV}}\right) \tag{4}$$

Armed with estimates of β, I can then calculate the overall impact of the Zwin for each location i assuming no change in population distribution $ZwinImpact_i = \hat{\beta} \Delta_{1000,1400} \ln\left(MA_i^{IV}\right)$. This measure of the impact of the Zwin will only capture the direct reduced-form effect. That is, it will only capture changes in market access due to the Zwin appearing not the subsequent changes that are likely to occur as individuals react through migrating as well as other spatial GE effects such as price changes. These secondary impacts are likely to significantly multiply the overall impact of the Zwin.

I estimate equation 2 using the instrument described in equation 3 on a subsample of 73 cities in the vicinity of the Zwin, these cities are shown on a map in figure 7 in the appendix. I choose a sub-sample of cities more likely to be impacted by the Zwin for reasons of numerical precision and computational constraints. Although the Zwin was a large event for Bruges and the surrounding cities, for locations far away it had little or no impact. When many such locations are included in my least-cost path calculations perturbations due

to issues around numerical precision can dominate real variation due to the Zwin. Increasing numerical precision to the degree to which this is no longer a problem makes the process with 653 cities computationally infeasible. That is, with so many locations almost completed unaffected by the Zwin noise in the estimation procedure dominates signal and the results become difficult to interpret. By focusing on a smaller subset of cities, although still quite some distance from the Zwin, I can overcome this challenge.

Table 2 shows the results. In column one, I perform OLS estimation. In column two, I use the instrument described above in a 2SLS procedure. As expected the coefficient attenuates — market access is higher in a location with unobserved better productivity or amenities and such a location would also attract a higher population. In columns three and four I perform pseudo-Poisson maximum likelihood estimation to account for the intensive margin which OLS drops as I take log of 0 population in various locations. PPML estimates the equation in levels and coefficients can be interpreted in the same way, as elasticities. In column four I take a control function approach to IV in PPML with a log-linear first stage. Again comparing columns three and four we see that naive estimation is biased up.

Coefficients can be interpreted as elasticities. Column four, my preferred specification, suggests that a one percent increase in market access causes a 22.6 percent increase in population. In standard deviation terms, a once standard deviation increase in market access is associated with around a 0.7 standard deviation increase in population.

Table 2 Market Access Regressions

	(1) OLS	(2) 2SLS	(3) PPML	(4) PPML IV
Log MA	21.53***	14.81***		
	(2.202)	(3.024)		
MA			32.35***	22.48***
			(4.453)	(3.487)
Control fn				8.356***
				(2.512)
Constant	11.55***		-18.07***	-11.74***
	(1.041)		(2.876)	(2.249)
# Cities	57	57	73	73

Notes: This table shows the results from estimating equation 2, on the sample of near-to-the-Zwin cities. In column I perform OLS esimtation. In column two 2SLS estimation, the first stage F-statistic is 15.21. In column three I perform PPML, and in column four I show the results from a two-stage control function approach estimated using PPML. Standard errors are robust.

3 A dynamic quantitative spatial economics model

I now extend my analysis to consider the impact the Zwin had on the spatial distribution of economic activity (and welfare) across the whole of the Low Countries, both when it was navigable and after it had become impassable. To do this I need to estimate the counterfactual spatial distribution of population and economic activity.

To do this I write down and estimate a quantitative spatial economics model building on Allen and Donaldson [2020] by including a key feature of the European medieval and early modern economy — merchants. Merchants facilitated trade and heralded the commercial revolution in Europe whereby a new rich merchant class was born. Nowhere (with the potential exception of the northern Italian cities) was this process more pronounced than in the Low countries. Merchants lowered trade frictions and were a highly mobile group. As a result, merchants could be found at a much greater spatial concentration than the population. Therefore the spatial distribution of merchants is more sensitive to shocks

inducing equilibrium switching, relative to that of the general population. As merchants bring trade and lower transaction costs within a city, and individuals exhibit higher costs to migration, the sudden movement of the merchant population away from a location could spell ruination for the remaining population.

Following Allen and Donaldson [2020] there are arbitrarily many locations $i \in N$ (cities) and $t \in \mathcal{T}$ time periods. Cities are connected and goods and individuals are allowed to move (with some cost) between them. History impacts the future through dynamic agglomeration effects in productivity and amenities. Intuitively infrastructure built some time ago might enhance (or decrease) productivity today. The advantage of leveraging this framework is that I can study the impact of the Zwin holding fixed all other shocks facing cities during this period, and trace out its impact across both time and space allowing for the endogenous and dynamic responses of individuals' location choice (and goods trade). The model also admits the potential for multiple long-run spatial equilibria whereby shocks can cause permanent changes to the distribution of economic activity. The disadvantage is that I have to impose considerable structure on the data, and shut down endogenous political economy responses.

Turning first to firms. Each location i emits a unique good in an Armington fashion. A continuum of firms ω in i produce this homogeneous good ($q_{it}(\omega)$) under perfect competition and CRTS using labor ($l_i(\omega)$) as the only factor of production.

$$q_{it}(\omega) = A_{it} l_{it}(\omega), \qquad A_{it} = \bar{A}_{it} L_{it}^{\alpha_1} L_{it-1}^{\alpha_2} \tag{5}$$

Where $\bar{A}_{it}$ is exogenous productivity and L_{it} is the total number of workers. α_1 captures aggregate contemporaneous spillovers, α_2 captures aggregate historical productivity spillovers. Intuitively α_1 captures what is more traditionally thought of as agglomeration forces, whereas α_2 captures factors like historical infrastructure which remain productive in the next period.

Individuals have CES preferences over differentiated location-specific goods with the elasticity of substitution σ, therefore consumption welfare is captured by local real wages (w_{it}/P_{it}). A location also generates utility for individuals in the form of local amenities (u_{it}),

and therefore location-time specific welfare is given by W_{it} in equation 6.

$$W_{it} = u_{it}\frac{w_{it}}{P_{it}}, \qquad u_{it} = \bar{u}_{it}L_{it}^{\beta_1}L_{it-1}^{\beta_2} \tag{6}$$

Where $\bar{u}_{it}$ is exogenous productivity and β_1, β_2 are analogous to α_1, α_2. For example β_1 captures contemporaneous congestion forces i.e. from non-tradeables or land and β_2 captures the impact of durable infrastructure on amenities, such as parks.

Bilateral trade from locations i to j incurs exogenous, symmetric, iceberg trade costs denoted by τ_{ijt}. Iceberg trade costs and CES demand generate the familiar gravity equation in trade [Allen et al., 2020b].

$$X_{ijt} = \tau_{ijt}^{1-\sigma}\left(\frac{w_{it}}{A_{it}}\right)^{1-\sigma}P_{jt}^{\sigma-1}w_{jt}L_{jt}, \qquad P_{it} = \left(\sum_{k=1}^{N}\tau_{kit}\left(\frac{w_{kt}}{A_{kt}}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{7}$$

Individuals decide where to move to maximize utility given in 6 subject to iceberg moving costs μ_{ijt} and some idiosyncratic preference draw ε_{jt} which is drawn from a Frechet distribution with dispersion parameter θ. Given this distributional assumption, migration will also follow a gravity structure, and the number of people moving from i to j in period t will be given by equation 8.

$$L_{ijt} = \mu_{ijt}^{-\theta}\Pi_{it}^{-\theta}W_{jt}^{\theta}L_{it-1} \tag{8}$$

Where $\Pi_{it} = \left(\sum_k \mu_{ikt}^{-\theta}W_{kt}^{\theta}\right)^{1/\theta}$, is a measure of labor market access.

I incorporate merchants within the dynamic quantitative spatial economics model of Allen and Donaldson [2020] by modeling them as a distinct set of agents. Merchants choose where to locate to maximise profits subject to idiosyncratic preferences which may reflect individual taste (merchants are after all also people with preferences), but also individual connections and knowledge. I model merchant income in a location i as some constant fraction α of the total trade T_i in that location.

$$\pi_i^M = \alpha A_i^M T_i \tag{9}$$

Merchant location-specific productivity, A_i^M, is a function of the number of merchants in

a location and some exogenous component $A_i^M = \bar{A}_i^M \cdot M_i^\lambda$. The parameter λ reflects the strength of Merchant-specific agglomeration economies. On the one hand, more merchants in a location mean that the "pie" needs to be split more ways, and so one may expect λ to be negative. On the other hand, merchants trade among themselves and so λ could be positive. Merchants have idiosyncratic location-specific preferences drawn from a type two extreme value distribution with dispersion parameter ϕ. From properties of the type two extreme value distribution we can therefore write the number of merchants in any given location as follows.

$$M_{it} = \left(A_{it}^M T_{it}\right)^\phi \left(\sum_j \left(A_{jt}^M T_{jt}\right)^\phi\right)^{-1} \tag{10}$$

Merchants also interact with the local community by increasing local productivity $A_{it} = \bar{A}_{it} M_{it}^\chi L_{it}^{\alpha_1} L_{it-1}^{\alpha_2}$.

Goods markets clear when total expenditure in a given location, $E_{it} = \sum_j X_{ijt}$, equals total income, which as labor is the only factor of production is the total wage bill $E_{it} = Y_{it} = w_{it} L_{it}$.

The dynamic equilibrium of this model is described in the appendix section A. The model can be solved via a simple iterative algorithm [Donaldson and Hornbeck, 2016]. For details and an in-depth discussion of the equilibrium properties of this model, and models in this class, see Allen and Donaldson [2020] and Allen et al. [2020a].

3.1 Estimating model parameters

Medieval and early modern Europe is not a data rich environment. However, leveraging all available data, I am able to estimate most model parameters.

To estimate model parameters, exogenous fundamentals, and the impact of the Zwin on travel costs and productivity, I use a five-step procedure. In step one, I use the cost surface and information on prices to estimate travel costs. In step two, I use estimated travel costs, population, and wage data to calculate market access terms. In step three, I use model-implied regressions and plausibly exogenous variation from large-scale historical events to estimate model parameters. In step four, I use the estimated model parameters and data to back out the rationalizing location fundamentals In step five, I estimate the effect of the

Zwin on fundamental productivity and travel costs.

3.1.1 Step 1: Calibrate migration and trade costs.

I parameterize the iceberg migration and trade costs as functions of the distance between each location, which I denote by t_{ijt}.

$$\tau_{ij}^{1-\sigma} = f_\tau\left(t_{ijt}\right) \qquad \mu_{ij}^{-\theta} = f_\mu\left(t_{ijt}\right) \tag{11}$$

In the absence of city-to-city trade or migration data I am unable to estimate f_τ or f_μ and so instead opt for a simple linear parameterization: $f_\xi(t_{ijt}) = t_{ijt}^{-\kappa_\xi}$ for $\xi = \{\tau, \mu\}$ where κ_ξ can then be calibrated. Over the time period studied cross-land transport technology changed little in the Low Countries (my sample ends in 1800 before the advent of railways in Belgium or the Netherlands, the first of which was inaugurated on May 5, 1835), and the main thoroughfares remain relatively unchanged.[9] Bilateral travel times are thus estimated using the least cost path across the generated cost surface as discussed in section 2.1. The only time variation comes from whether or not the Zwin is navigable. Estimates of κ_τ and κ_μ in a comparable setting are understandably hard to come by, I take as my baseline $\hat{\kappa_\tau} = 2.05$ and $\hat{\kappa_\mu} = 1.16$ from Ellingsen [2021] who uses migration and trade data from Latin America around 1800.

An alternative could be to estimate trade costs directly using information on prices. We have that $P_{it}^{1-\sigma} = \sum_j \tau_{ijt}^{1-\sigma} Y_{jt} P_{jt}^{\sigma-1}$. Parameterise iceberg trade costs as $\ln(\tau_{ijt}) = \kappa_\tau \cdot \ln(t_{ijt}) + \beta_{it} + \alpha_{jt} + \xi_{ijt}$, we can use data on prices, incomes, and travel costs to estimate $\kappa_\tau, \beta_{it}, \alpha_{it}$ by non-linear least squares. I can then estimate the direct reduced-form impact of the Zwin on $\hat{\beta}_{it}$ and $\hat{\alpha}_{it}$ therefore allowing the Zwin to affect multilateral as well as bilateral travel costs in a flexible manner. Operationalizing this approach remains a work in progress.

[9] There were of course some innovations on the margins, especially into the 17th centaury with the advent of stage coaches, turnpike-esque systems of improved roads, and increased canal activity. Although these late additions could have significant impacts on travel time they pale in comparison to the impact of the railroad and motorised transport. For example, although stage coaches in the mid to late 1700 could travel between Bruges and Brussels in a day — twice as fast as horseback travel 100 years before — even early trains in 1900 could do the same journey in a couple of hours.

3.1.2 Step 2: Find market access terms.

Having calculated transport costs, I can calculate labor market access and trade market access terms using available population and wealth data. We can solve the following system of non-linear equations via a simple iterative algorithm which has a unique up-to-scale solution [Donaldson and Hornbeck, 2016].

$$LMA_{it} = \sum_j t_{ijt}^{-\kappa_\mu} \frac{L_{jt-1}}{LMA_{jt}} \qquad TMA_{it} = \sum_j t_{ijt}^{-\kappa_\tau} \frac{Y_{jt}}{TMA_{jt}} \tag{12}$$

This is the same procedure implemented in section 2.1.

3.1.3 Step 3: Use model-implied regressions to estimate parameters

Using the trade gravity equation, migration gravity equation, and expression governing the spatial distribution of migrants along with standard market clearing conditions we can write derive the following series of log-linear relationships (see the appendix section A for details).

$$\begin{aligned}
(1-\beta_1\theta)\ln(L_{it}) &= (\beta_2\theta+1)\ln(L_{it-1}) + \theta\ln(w_{it}) + \frac{\theta}{1-\sigma}\ln(TMA_{it}) \\
&\quad + \ln(LMA_{it}) + e_{it} \\
\sigma\ln(w_{it}) &= \alpha_2(\sigma-1)\ln(L_{it-1}) + (\alpha_1(\sigma-1)-1)\ln(L_{it}) \\
&\quad + \chi(\sigma-1)\ln(M_{it}) + \ln(TMA_{it}) + v_{it} \\
\ln(M_{it}) &= \frac{\phi}{1-\lambda\phi}\ln(Y_{it}) + \tau_t + \xi_{it}
\end{aligned}$$

With data on $\{L_{it}, L_{it-1}, w_{it}, M_{it}, LMA_{it}, TMA_{it}\}$ I can estimate this series of equations (note $Y_{it} = w_{it}L_{it}$. This allows identification of all parameters with the exception of ϕ and λ which cannot be separately identified. λ captures the strength of merchant specific agglomeration forces and ϕ captures the dispersion of merchants location-specific idiosyncratic preferences. Intuitively, high agglomeration forces (large λ) or very dispersed preferences (low ϕ result in an observationally similar dispersion of merchants under a given set of location-specific outputs Y_{it}. To break this lack of identification I set $\lambda = 0$.

To estimate this series of simultaneous equations I take the following steps. I construct a dataset with information on wages, population, and merchants across 17 cities in 1500, 1600, 1700, and 1800 — this dataset is a subsample of the main population data, therefore I estimate parameters under the assumption that they would be consistent in other periods/ locations.[10] First, as I only have wage data for a sub-set of locations I cannot calculate TMA_{it} accurately and so I only use LMA_{it} to control for spatial factors affecting population and wages. Often TMA_{it} and LMA_{it} are very highly correlated so it is unlikely this step affects the results significantly. However, as a result of this I cannot credibly identify θ and σ and instead take estimates from the literature $\theta = 4, \sigma = 9$ following Donaldson and Hornbeck [2016]. Finally, this system is under-identified as there are three endogenous variables L_{it}, w_{it}, M_{it} and only one predetermined variable L_{it-1}. To identify the system I therefore need two further sources of plausibly exogenous variation. I leverage variation in temperature from the little ice age and the opening of trans-Atlantic trade routes as two such sources.

First, I follow Waldinger [2022] and leverage temperature variation across time and cities due to the climatic event known as the little ice age. Between about 1400 and 1900 temperatures in Europe dropped by as much as one degree on average relative to their previous baseline which the most extreme drops occurring in the 16th and 17th centauries. Waldinger [2022] shows that decreases in temperature negatively affected city size, a result I confirm here. I take the average temperature twenty years either side of each century in the grid cell (0.5° x 0.5°) a city resides in as my measure of temperature.[11] Second, I following Acemoglu et al. [2005a] I use variation in the timing and impact of of the Atlantic for trade across European cities. Acemoglu et al. [2005a] argue that cities which happened to have easy access to the Atlantic coast in Europe grew disproportionately following the opening of the Atlantic to trade around 1600. I use a dummy variable taking the value one if a city is in an Atlantic trading country after and including 1600 as my measure of Atlantic trading

[10]The cities are Amsterdam, Antwerp, Augsberg, Florence, Gdansk, Krakow, Leipzig, London, Madrid, Milan, Munich, Naples, Paris, Strasbourg, Valencia, Vienna, and Warsaw. This is obviously not a random sample of cities but rather those more likely to have available wage data in this period. Leipzig, Madrid, and Warsaw do not have wage data for 1500 leaving me with a total of 65 city-year observations in this dataset.

[11]I use gridded temperature data from Luterbacher et al. [2004] which starts in 1520, therefore for 1500 I take the average from 1520 to 1540.

potential.

Table 3 shows the resulting parameter estimates and table 4 in the appendix shows the first-stage results. The novel parameters χ and ϕ are estimated quite precisely as strongly positive. χ governs the degree to which merchants influence the local economy through local productivity. A value of 0.41 implies that a 10% increase in merchants in a location increases local productivity by 4%. ϕ disciplines the responsiveness of migration choices of merchants to the value of local trade. As merchants don't face the dispersion force of migration costs we would expect this parameter to be lower than that estimated for the general population, and indeed do find $\phi = 0.56$. The agglomeration parameters β_k, α_k are estimated with less precision which is perhaps unsurprising given the relative lack of data.

Table 3 Parameter Estimates

	Estimate	Standard Error
β_1	0.64	0.52
β_2	-0.50	0.32
α_1	-0.57	0.27
α_2	0.33	0.19
χ	0.41	0.12
ϕ	0.56	0.12

3.1.4 Step 4: Invert the model to back out location fundamentals

Armed with parameter estimates and data I can invert the model to back out underlying fundamentals. However, although wage and merchant data are useful to estimate model parameters they are not sufficient to invert the model. That is, I am only able to invert the model to identify a bundle of location fundamentals using available information on populations and market access. In particular, assuming that $TMA_{it} \approx LMA_{it}$ as is often found in the literature, I can solve model equations to find the following expression.

$$C_L \ln(L_{it}) + C_{L_{-1}} \ln(L_{it-1}) + C_{MA} \ln(MA_{it}) + C_\Omega \ln(\Omega_t) = C_u \ln(\bar{u}_{it}) + C_A \ln(\bar{A}_{it}) + C_{A^M} \ln(\bar{A}^M_{it}) \tag{13}$$

Where the coefficients are known combinations of structural parameters. Using data and parameter estimates I can recover the left hand side of this equation which we can denote by

D_{it}. As I cannot separately identify location fundamentals I make the simplifying assumption that amenities and merchant productivities do not vary over time. Then we can write equation 13 as $D_{it} = \lambda_i + \tau_t + \xi_{it}$ where $\xi_{it} = C_A \ln(\bar{A}_{it})$. This equation can be estimated allowing us to back out fundamental location productivity as $\hat{\bar{A}}_{it} = \exp\{\hat{\xi}_{it} C_A^{-1}\}$. Appendix section A.2 includes details on this procedure.

3.1.5 Step 5: Estimate the counterfactual path of fundamental productivities in the absence of the Zwin.

The Zwin directly effects transport costs to and from Bruges to all other locations. In addition, by potentially influencing merchant activity it can also indirectly impact local productivities. However, I also allow the Zwin to directly affect the productivity of Bruges. I estimate the Zwins impact using a dummy variable, including city and period fixed effects. This approach will attribute all within-city deviations from linear city-specific time trends and year fixed effects in fundamental productivities to the Zwin. That is, other shocks which are not modeled (not due to the changing economic geography of the Low Countries) and city-year specific, that coincide with the Zwin will be bundled into the estimated Zwin effects.

$$\hat{\bar{A}}_{it} = \beta \cdot \text{ZwinNavigable}_{it} + \gamma_i + \gamma_i \times t + \tau_t + \varepsilon_{it} \tag{14}$$

The variable $\text{ZwinNavigable}_{it}$ is a dummy variable that takes the value of one only in Bruges and during the period 1200 to 1500, subscripts on the remaining variables in equation 14 indicate the level of fixed effects being denoted. Estimating this equation I find that $\hat{\beta} = 6.06$ ($SE = 0.251$). To put this number into context the average year-on-year change in local productivities across the whole sample is around 1 (with a standard deviation of around 2), therefore the Zwin caused a shock around 6 times the average.

Figure 6 in the appendix graphically displays the estimated path of fundamental productivities taken in Bruges. In orange, I plot the estimated productivities that rationalize the observed population distribution. One can clearly see a significant spike in 1300 and 1400 as well as a significant dip in 1500. I then plot in a black dashed line the estimated path of

productivity in the absence of the Zwin which effectively smooths out the spikes and dips described before.

4 The long-run shadow of a temporary river

Figure 4 shows the the impact of the Zwin in Bruges and all of the cities in the "close to the Zwin" sample, on both population and welfare. On the left-hand side, I show the impact on Bruges itself. The upper left panel shows the impact on Bruges's population. One can see a clear and large effect, at its peak Bruges was about 400% more populous due to the Zwin. The reduced form results from section 2 suggested that the Zwin caused the population of Bruges to be between 340% and 470% larger whilst it was navigable — a comparable magnitude. However, the empirical results from the model allow me to go beyond those from the reduced form analysis in four key dimensions.

First, the model allows me to consider the heterogenous impact of the Zwin on cities other than Bruges. The top right hand panel of figure 4 shows the percentage change in population caused by the Zwin in each of the remaining 72 cities in my sample. Comparing the y-axis of this graph to that of the top left graph one can see a considerable change in the magnitude of the effect. However, the impact is still economically meaningful, with some locations being over 5% smaller in 1400, and a few per cent larger in 1600 due to the Zwin. The impact is also heterogeneous, with locations far away (or otherwise with few trading connections) to Bruges seeing the smallest impacts.

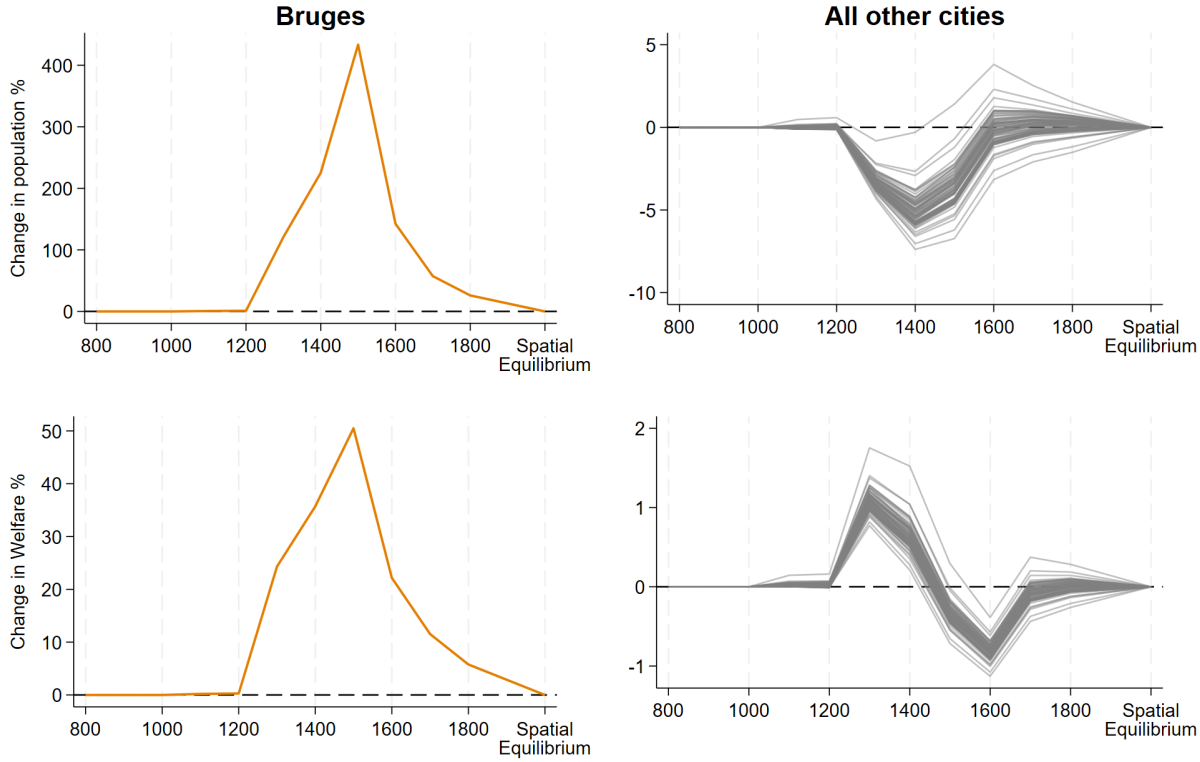
Second, using the model I can project the impact of the Zwin forwards, beyond 1500 after which it became impassable. This allows me to ask whether the Zwin cast any shadow on the spatial distribution of economic activity in Bruges and beyond. The top left hand panel of figure 4 shows clearly that in 1600, some 100 years after the Zwin had silted up completed, Bruges was still over twice the population it would have otherwise been if the Zwin had never appeared in the first place. This persistence continues to 1800, where Bruges population remains 20% or so larger due to the Zwin which disappeared now some 300 years ago.

Third, turning to the lower two graphs in figure 4, the model allows me to directly look

at the impact of the Zwin on local welfare. Population movements mitigate welfare changes, but due to costly migration do so imperfectly. As a result, local welfare in Bruges increased considerably during the period when the Zwin was navigable, and at its peak was around 50% higher than it would otherwise have been. This also persisted after the Zwin silted up, with welfare remaining almost 10% higher in 1800. The bottom right hand side figure then shows the welfare impact of the Zwin across cities not directly affected. Again, the magnitude of the impacts is considerably diminished relative to the effect on Bruges itself. However, the impact is still sizable. From 1200 to 1400 the impact is unambiguously positive due to two factors. First, other cities can now enjoy relatively cheaper goods from Bruges, and second, due to outmigration to Bruges those who remain have higher wages and face lower congestion. These forces are removed once the Zwin silts up causing a relative decline in welfare in 1500 and 1600 across the board.

Finally, I can leverage the model to analyse whether the Zwin caused a change in the long-run spatial equilibrium. In each of the panels of figure 4 I include an observation that corresponds to the spatial equilibrium. For each location, the economy returns to the spatial equilibrium. The sudden appearance and subsequent disappearance of the Zwin has no impact on the very long-run distribution of economic activity.

Figure 4 The model implied impact of the Zwin



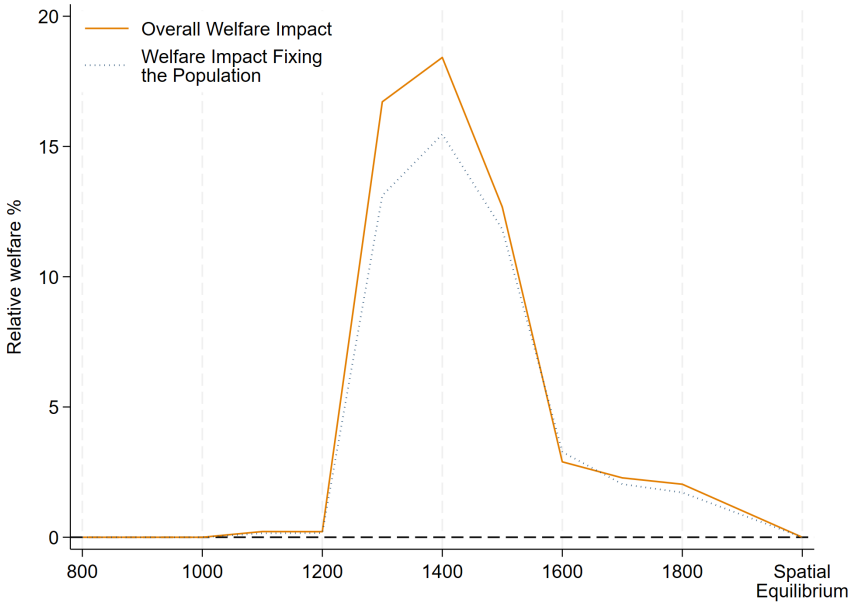
Notes: This figure shows the model-implied impact of the Zwin in Bruges and each of the other 72 cities in my "cloe to the Zwin" sample. In the top left figure, I plot the impact of the Zwin across time on Bruges' population measured as a percentage of its population in the absence of the Zwin. In the top left, I repeat this exercise for each of the other 72 cities. In the bottom two panels, I instead plot the model-implied change in welfare again measured as a percentage of a given location's welfare in the absence of the Zwin.

These results show that the Zwin did impact locations other than Bruges, and that this impact could be felt (particularly in Bruges itself) hundreds of years after the river ceased to exist. Additionally, as well as altering the spatial distribution of population, the Zwin had a material impact on local welfare both in Bruges and elsewhere — and this impact can also be detected centuries after the Zwin silted up. This sudden river cast a significant long-run shadow.

I can also use the model to consider how the Zwin affected aggregated Welfare across sub-sets of cities. In figure 5 I plot the population-weighted percentage change (relative to the never-Zwin baseline) in Welfare across all 17 Belgian cities in my sample. Figure 5 shows a large positive effect of an almost 20% increase for the period when the Zwin was navigable. It also shows that most of this effect does not come through within Belgium population redistribution. This effect persists, and Belgium as a whole remains over 2% better off some

300 years after the Zwin silted up than if the Zwin had never existed in the first place.

Figure 5 The impact of the Zwin on aggregate welfare in Belgium



Notes: This figure shows the aggregate impact of the Zwin on welfare across Belgium. In the solid orange line, I plot the population-weighted change in welfare due to the Zwin measured as a percentage of the never-Zwin counterfactual welfare. In the dotted line, I plot the same variable but fix the weighting as the distribution of population in the never-Zwin scenario.

5 Conclusion

In 1134 a large storm caused a sudden river to appear. This unique historical experiment allows us to look at the contemporaneous and long-run impact of an exogenous change to transport costs. In this paper, I show that this sudden river had a large contemporaneous impact on the spatial distribution of population and welfare. But more than this, I show that the sudden river still has a measurable impact some 300 years after it had silted up. That is, it cast a long shadow on the spatial economy of the low countries.

These results suggest that today's spatial distribution of economic activity is an amalgamation of past shocks. However, one should not overstate the length of the shadow. This paper suggests the Zwin caused Bruges population to be 400% higher at its height, 100% higher 100 years after it silted up, and only 10% high 300 years after it became impassable.

Appendix

A Model details

A.1 Existence and uniqueness of equilibrium

Recall from the main text three gravity-esque equations dictating migration, trade, and merchant activity across space.

$$L_{ijt} = \mu_{ijt}^{-\theta} W_{it}^{\theta} \left(\sum_{k} \mu_{ikt}^{-\theta} W_{kt}^{\theta} \right)^{-1} L_{jt-1} \tag{15}$$

$$X_{ijt} = \tau_{ijt}^{1-\sigma} w_{it}^{1-\sigma} A_{it}^{\sigma-1} \left(\sum_{k} \tau_{ikt}^{1-\sigma} w_{jt}^{1-\sigma} A_{kt}^{\sigma-1} \right)^{-1} E_{jt} \tag{16}$$

$$M_{it} = \left(A_{it}^{M} T_{it} \right)^{\phi} \left(\sum_{j} \left(A_{jt}^{M} T_{jt} \right)^{\phi} \right)^{-1} \tag{17}$$

Market clearing in the goods market implies that total expenditure in a given location is equal to total income which in turn equals the sum of all imports (and own consumption) which is total trade, $E_{it} = Y_{it} = \sum_j X_{ijt} = T_{it}$. As labor is the only factor of production local income is in turn given by the wage bill $Y_{it} = w_{it} L_{it}$. Labor market clearing in turn implies that the population of a given location is equal to the sum of in-migrants (including those who stay). Finally I impose symmetry in both trade and migration costs over space. Given this, we can find expressions for total population and income by summing respective gravity equations over j.

$$L_{it} = \left(\bar{u}_{it} L_{it}^{\beta_1} L_{it-1}^{\beta_2} w_{it} TMA_{it}^{1/(1-\sigma)} \right)^{\theta} L_{it-1} LMA_{it} \tag{18}$$

$$w_{it} L_{it} = w_{it}^{1-\sigma} \left(\bar{A}_{it} M_{it}^{\chi} L_{it}^{\alpha_1} L_{it-1}^{\alpha_2} \right)^{\sigma-1} TMA_{it} \tag{19}$$

Where $LMA_{it} = \sum_j \mu_{ijt}^{-\theta} L_{jt-1} LMA_{jt}^{-1}$ and $TMA_{it} = \sum_j \tau_{ijt}^{1-\sigma} Y_{jt} TMA_{jt}^{-1}$. From these two expressions and the equation determining the spatial distribution of merchants, we can derive two equations of the form described by Allen et al. [2020a]. To see this note we can write

the expression determining the spatial distribution of population as the following.

$$L_{it} = \bar{u}_{it}^{\frac{\theta}{1-\beta_1\theta}} L_{it-1}^{\frac{1+\beta_2\theta}{1-\beta_1\theta}} LMA_{it}^{\frac{1}{1-\beta_1\theta}} w_{it}^{\frac{\theta}{1-\beta_1\theta}} TMA_{it}^{\frac{\theta}{(1-\sigma)(1-\beta_1\theta)}} \tag{20}$$

Similarly, after some algebra, we find that the expression determining the spatial distribution of wages can be combined with that determining the distribution of merchants to find the following.

$$w_{it} = L_{it}^{\frac{(\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)}{1-\phi\lambda}x} L_{it-1}^{(\alpha_2(\sigma-1)-1)x} \bar{A}_{it}^{(\sigma-1)x} (\bar{A}_{it}^{M})^{\frac{\phi\chi(\sigma-1)}{1-\phi\lambda}x} \Omega_t^{\frac{\phi(\sigma-1)}{1-\phi\lambda}x} TMA_{it}^{x} \tag{21}$$

Where $x = \frac{1-\phi\lambda}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}$. Subbing equations 20 and 21 into each other we can find expressions for L_{it} and w_{it} in terms of market access terms and exogenous fundamentals or previous population. We can write $L_{it} = LMA_{it}^{a_1} \cdot TMA_{it}^{a_2} \cdot L_{it-1}^{a_3} \cdot F_{it}^{L}$ and $w_{it} = LMA_{it}^{b_1} \cdot TMA_{it}^{b_2} \cdot L_{it-1}^{b_3} \cdot F_{it}^{w}$ where F_{it}^{L} and F_{it}^{w} group together exogenous terms, and we have the following.

$$a_1 = \left(\frac{1}{1-\beta_1\theta}\right)\left(1 - \frac{\theta\left((\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)\right)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)^{-1}$$

$$a_2 = \left(\frac{\theta(1-\phi\lambda)}{(1-\beta_1\theta)(\phi\chi(1-\sigma)+\sigma(1-\sigma\lambda))} + \frac{\theta}{(1-\sigma)(1-\beta_1\theta)}\right)\left(1 - \frac{\theta\left((\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)\right)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)^{-1}$$

$$a_3 = \frac{1+\beta_2\theta}{1-\beta_1\theta} + \left(\frac{\theta}{1-\beta_1\theta}\right)\left(\frac{(\alpha_2(\sigma-1)-1)(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\chi)}\right)$$

$$b_1 = \left(\frac{1}{1-\beta_1\theta}\right)\left(\frac{(\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)\left(1 - \left(\frac{\theta}{1-\beta_1\theta}\right)\left(\frac{(\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)\right)^{-1}$$

$$b_2 = \left(\frac{\theta}{(1-\sigma)(1-\beta_1\theta)}\right)\left(\frac{(\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)\left(1 - \left(\frac{\theta}{1-\beta_1\theta}\right)\left(\frac{(\phi\chi+\alpha_1)(\sigma-1)-(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)\right)^{-1}$$

$$b_3 = \frac{(\alpha_2(\sigma-1)-1)(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)} + \left(\frac{\theta}{1-\beta_1\theta}\right)\left(\frac{(\alpha_2(\sigma-1)-1)(1-\phi\lambda)}{\phi\chi(1-\sigma)+\sigma(1-\phi\lambda)}\right)$$

Substituting these expressions into the market access terms we arrive at the following expressions, where $K_{jt} = F_{jt}^{L} \cdot F_{jt}^{w}$.

$$LMA_{it} = \sum_j \mu_{ijt}^{-\theta} \cdot LMA_{jt}^{-1} \cdot L_{jt-1} \tag{22}$$

$$TMA_{it} = \sum_j \tau_{ijt}^{1-\sigma} \cdot K_{jt} \cdot TMA_{jt}^{b_2+a_2-1} \cdot LMA_{jt}^{b_1+a_1} \cdot L_{jt-1}^{b_3+a_3} \tag{23}$$

These equations are of the form described in Allen and Donaldson [2020] with the sim-

plification that in this specification I do not allow forward looking behavior. As a result I can apply theorem one in Allen and Donaldson [2020] which states in the simplified case that, as in Allen et al. [2020a], the existence and uniqueness of the system depends on the spectral radius of the matrix of market access coefficients. Denote the matrix of coefficients by $E = \begin{pmatrix} |-1| & 0 \\ |b_1 + a_1| & |b_2 + a_2 - 1| \end{pmatrix}$. Then an equilibrium exists and is unique if the spectral radius of E is less than one, $\rho(E) < 1$, where the absolute value operator is taken element-wise. It is easy to show that $\rho(E) < 1$ iff $|b_2 + a_2 - 1| < 1$.

This describes conditions for existence and uniqueness of the equilibrium path. In this case pre-determined variables i.e. L_{it-1} do not impact the uniqueness and existence of period-by-period equilibrium. However, the long run equilibrium, where $L_{it} = L_{it-1} = L_i$ for a time-invariant set of exogenous fundamentals and transport costs will also depend on L_{it-1}. We can solve for the long-run spatial equilibrium to find the following series of equations.

$$LMA_i = \sum_j \mu_{ij}^{-\theta} LMA_j^{\frac{a_1}{1-a_3}-1} TMA_j^{\frac{a_2}{1-a_3}} (F_j^L)^{\frac{1}{1-a_3}} \tag{24}$$

$$TMA_i = \sum_j \tau_{ij}^{1-\sigma} LMA_j^{b_1+\frac{a_1(1+b_3)}{1-a_3}} TMA_j^{b_2-1+\frac{a_2(1+b_3)}{1-a_3}} (F_j^L)^{\frac{1+b_3}{1-a_3}} F_j^w \tag{25}$$

Therefore, the uniqueness of the long-run steady state will depend on the spectral radius of E_{LR} which is given by the following martrix.

$$E_{LR} = \begin{pmatrix} \left|\frac{a_1}{1-a_3} - 1\right| & \left|\frac{a_2}{1-a_3}\right| \\ \left|b_1 + a_1\left(\frac{1+b_3}{1-a_3}\right)\right| & \left|b_2 + a_2\left(\frac{1+b_3}{1-a_3}\right) - 1\right| \end{pmatrix} \tag{26}$$

A.2 Backing out location fundamentals

Using equations 18, 19, and 17 we can solve for fundamental amenities $\{\bar{u}_{it}\}$, goods production productivity $\{\bar{A}_{it}\}$, and merchant productivity $\{\bar{A}_{it}^M\}$ in each location in each period as

a function, data, market access terms, and parameters.

$$\bar{u}_{it} = L_{it}^{\frac{(1-\theta\beta_1)}{\theta}} L_{it-1}^{-\frac{1+\theta\beta_2}{\theta}} w_{it}^{-1} TMA_{it}^{\frac{1}{\sigma-1}} LMA_{it}^{-\theta} \tag{27}$$

$$\bar{A}_{it} = w_{it} L_{it}^{\frac{\alpha_1(\sigma-1)-1}{1-\sigma}} M_{it}^{-\chi} L_{it-1}^{-\alpha_2} TMA_{it}^{-\frac{1}{\sigma-1}} \tag{28}$$

$$\bar{A}_{it}^{M} = M_{it}^{\frac{1}{\phi}} w_{it}^{-1} L_{it}^{-1} \Omega_t^{-\frac{1}{\phi}} \tag{29}$$

Given data on $\{w_{it}, L_{it}, M_{it}\}$, parameters $\{\theta, \sigma, \phi, \chi, \alpha_1, \alpha_2, \beta_1, \beta_2\}$ and values for $\{LMA_{it}, TMA_{it}\}$ we can thus back out $\{\bar{u}_{it}, \bar{A}_{it}\}$. To operationalise this we make the following simplifying assumptions. First we assume market access terms are highly correlated and thus $LMA_{it} \approx TMA_{it}$. Next, as we do not have data on wages and merchants in each location in each period, we will not be able to separately identify the three location fundamentals. Instead, we can identifying a bundle. To see this we can solve the above system to find.

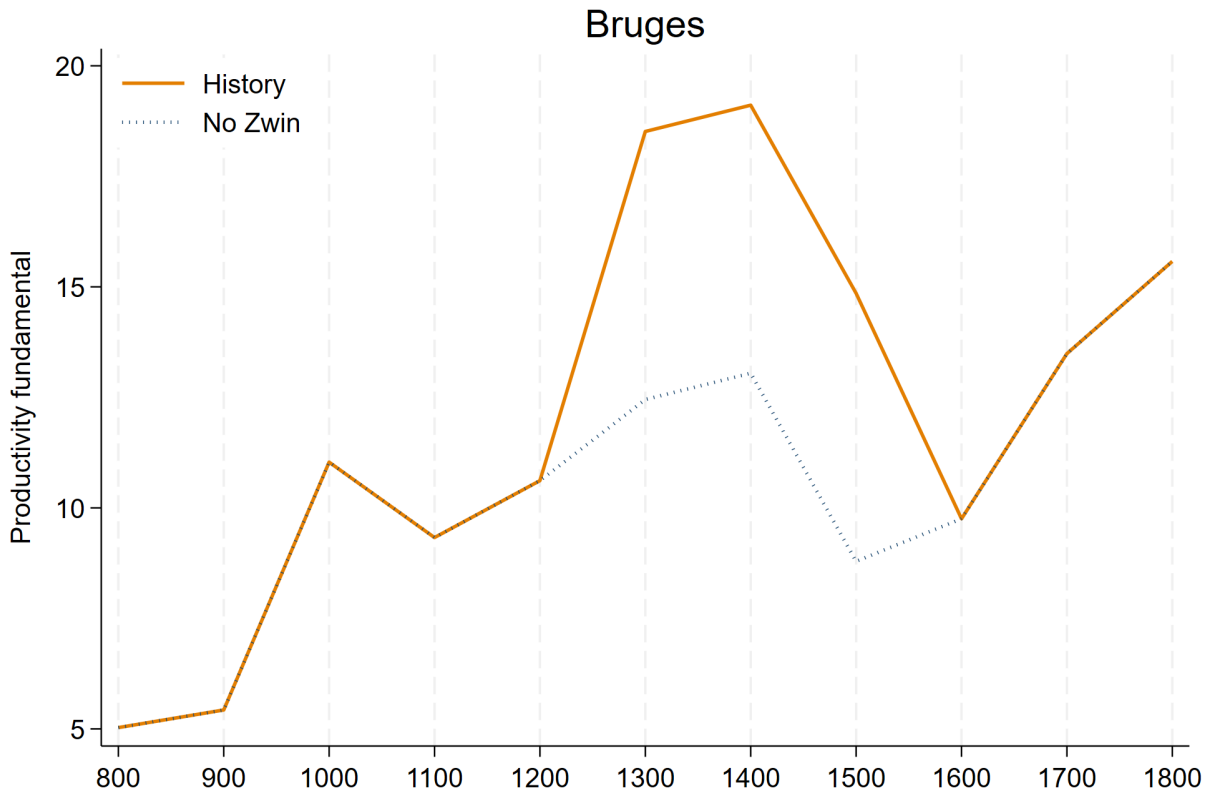
$$C_L \ln(L_{it}) + C_{L_{-1}} \ln(L_{it-1}) + C_{MA} \ln(MA_{it}) + C_\Omega \ln(\Omega_t) = C_u \ln(\bar{u}_{it}) + C_A \ln(\bar{A}_{it}) + C_{A^M} \ln(\bar{A}_{it}^M) \tag{30}$$

Where the necessary coefficients are given below.

$$\begin{aligned}
C_L &= 1 - \theta\beta_1 - \theta\frac{(\sigma-1)\chi\phi + (\sigma-1)\alpha_1 - 1}{\sigma - (\sigma-1)\chi\phi} \\
-C_{L_{-1}} &= \theta\frac{(\sigma-1)\alpha_1 - 1}{\sigma - (\sigma-1)\chi\phi} + \theta\beta_2 + 1 \\
-C_{MA} &= \frac{\theta}{\sigma - (\sigma-1)\chi\phi} + \frac{\theta}{1-\sigma} + 1 \\
C_A &= \theta\frac{\sigma-1}{\sigma - (\sigma-1)\chi\phi}
\end{aligned}$$

B Additional tables and figures

Figure 6 Counterfactual productivity fundamentals



Notes: This figure shows the path of the rationalizing productivity fundamentals in orange and the counterfactual no-Zwin path in dotted navy. The difference between these two lines is exactly the estimated Zwin shock.

Figure 7 Cities in my sample for the market access regressions



Notes: Cities in my "close to the Zwin" sample are shown in green.

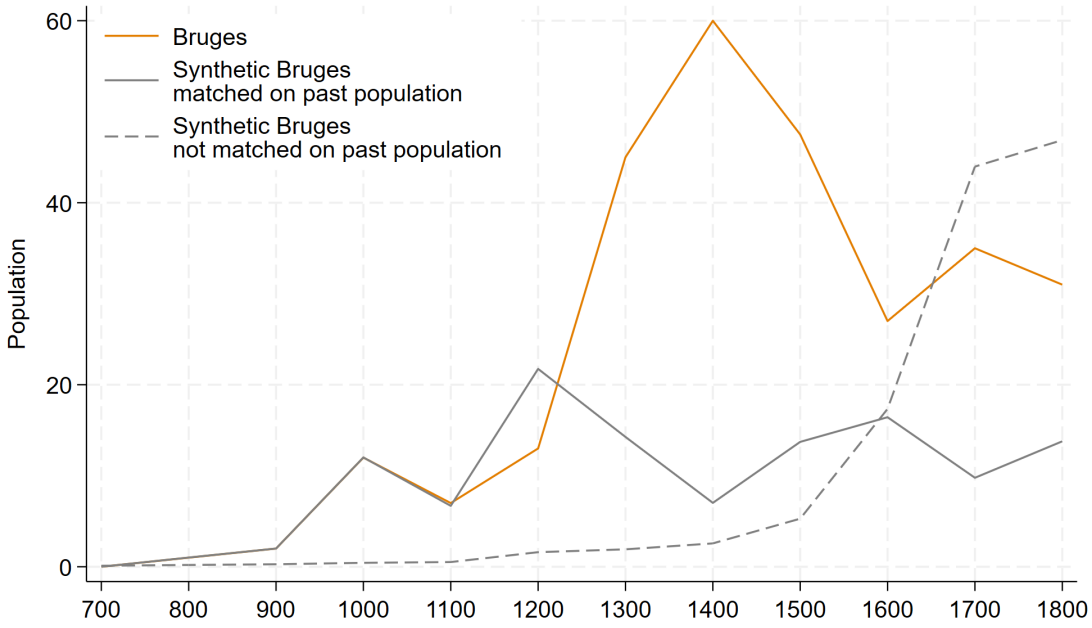
Table 4 Parameter estimates first stage

	(1) ln(wages)	(2) ln(pop)	(3) ln(Merchants)
Average Temperature	-0.0523**	-0.0190	-0.141***
	(0.0175)	(0.0288)	(0.0383)
Atlantic Trader	0.107	0.309*	0.729**
	(0.0898)	(0.130)	(0.236)
Lag ln(pop)	0.127*	0.967***	0.549***
	(0.0581)	(0.0679)	(0.131)
F	4.403	116.1	6.985
Observations	65	65	68

C Synthetic control results

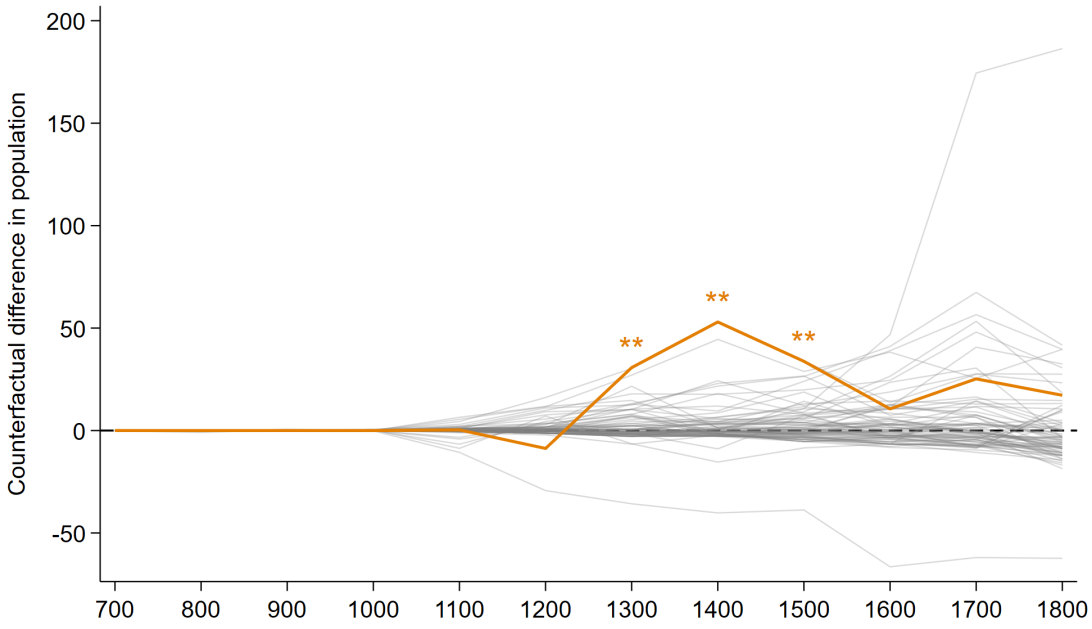
Synthetic control results. Population in 700, 800, 900, 1000, 1100, elevation, latitude, longitude, and population growth from 700 to 900 are used as matching variables. The weightings for Bruges are as such: Veliko Tarnovo (48.6%), Granada (4.4%), Burgos (5.7%), Badajoz (0.4%), and 409 cities are given the minimum 0.1% weighting. Figure 8 shows the results comparing historical Bruges to its synthetic control. This figure shows two synthetic controls, one that matches on past population, and one that does not. Figure 9 then shows inference in the synthetic control setting — permuting the synthetic control procedure over all 76 cities with a similar population to Bruges in 1100.

Figure 8 Synthetic control results: Bruges



Notes: This figure shows the population of Bruges and its estimated synthetic control. The solid gray line is the baseline synthetic control that matches on pre-1200 population, the dashed gray line does not match on past population.

Figure 9 Synthetic control results: Inference



Notes: This figure plots the difference between actual and counterfactual populations for each of the 76 cities in the sample of cities with a similar population to Bruges in 1100. Counterfactual populations are calculated via a synthetic control procedure for each city as described in the text. The population difference for Bruges in given in orange and stars indicate significance at the usual levels.

D Details on ChatGPTs sources and analysis

In this section I will provide an abridged version of the output from ChatGPT including all sources cited.

Paris. 1000: Paris was a regional market center under early Capetian rule with few merchants Kadens [2015]. 1100: King Louis VII granted a charter in 1170 to the Hanse parisienne des marchands de l'eau (Paris guild of river merchants), indicating an organized merchant group on the Seine. This guild monopolized river trade between Paris and Mantes, suggesting on the order of a few dozen active boat-trading merchants by the late 1100s [Picarda, 1901]. 1200: Paris is the largest city in Europe with powerful merchant guilds. Tax records indicate that the wealthy bourgeois of Paris (largely composed of merchants and financiers) made up about 140 families (2,000 individuals), who paid the bulk of city taxes Bove and Gauvard [2015]. 1300: Paris merchant class was well-established and powerful and Paris was the kingdom's principal market for textiles, wine, and luxury goods Bove

and Gauvard [2015], Étienne Boileau [1268]. 1400: In 1421 city records list 21 wholesale wine merchants in Paris [Fierro, 1996], showing the presence of specialized merchant groups. Despite turmoil during the Hundred Years' War (Paris was occupied 1418–1436), a modest community of perhaps 100 merchants continued to trade in goods like wine, grain, and cloth, supplying both the city and distant markets. 1500: The merchant guilds and confraternities (such as the Confrerie de Notre-Dame led by wealthy drapers and spice merchants) were thriving [Bove and Gauvard, 2015]. We estimate roughly 150 or more merchants were active in inter-city trade by 1500, dealing in textiles, spices, and metal goods. 1600: The Paris merchant community likely numbered a few hundred. (For example, in 1609, immigrants from Antwerp made up one-third of Amsterdam's 1500 merchants [Gelderblom, 2016]; Paris, as a larger city, had a significant but more domestically focused merchant class.) We estimate 200 international-scale merchants in Paris circa 1600. 1700: Parisian merchants (including prominent Huguenot, Italian, and Swiss families) invested in colonial companies and banking. The dissolution of guilds after 1791 makes direct counts difficult, but on the eve of the French Revolution, Paris had an extensive wholesale trading class. We estimate on the order of 300 merchants around 1700 engaged in large-scale trade (e.g. in textiles, colonial groceries, and luxury manufactures). Paris's dominance as a consumer market ensured a substantial merchant presence, even if many operated through agents in port cities. 1800: Many merchant-bankers and wholesale dealers are documented in the late 18th and early 19th centuries. A conservative estimate is that Paris had on the order of a few hundred (perhaps 300) significant merchants at 1800, though the nature of commerce was changing. (The official Chamber of Commerce of Paris, established in 1803, soon counted several hundred members, indicating the scale of the mercantile class.).

E Alternative shocks

Figure 4 shows that the Zwin did impact the population distribution and location of economic activity across the urban structure of the low countries, both contemporaneously and over a long horizon. However, there is no evidence of a fundamental shift in the spatial distribution to a new long-run equilibrium. This leads to the natural question: Could an alternative

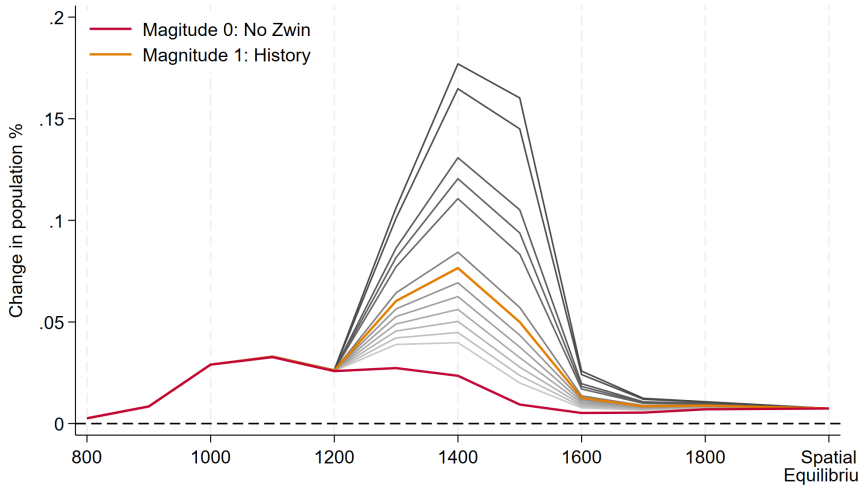
shock have led to different short or long-run effects, or have caused a shift in long-run spatial equilibrium?

I vary the "Zwin-shock" on a number of dimensions to get at this question. First, I ask what if the Zwin had reached all the way to Bruges, had stopped significantly further from the city walls, or had brought significantly more or less trade with it. That is, I vary the shock magnitude. Second, I ask what if the timing had been different. What if the Zwin had never silted up, or had taken an additional 100 before it had become impassable?

E.1 Varying shock magnitude

I vary the magnitude of the "Zwin shock" both on Bruges by multiplying the estimated impact of the Zwin by a factor γ which varies from 0 to 2. I multiply both the estimated impact on location productivity and that on transport costs. For each new counterfactual path of productivities, I estimate the dynamic quantitative spatial model described in section 3 and plot the results for Bruges in figure 10. In figure 10 darker lines depict more extreme shocks, the red line corresponds to the no-Zwin counterfactual (shock size of 0) and the orange to the actual history (shock size 1). Figure 10 shows that even relatively small shocks would have had a significant contemporaneous impact on the population of Bruges. However, in all cases, the same general pattern is followed, of a slow convergence back eventually to the same long-run spatial equilibrium.

Figure 10 Counterfactual productivity fundamentals

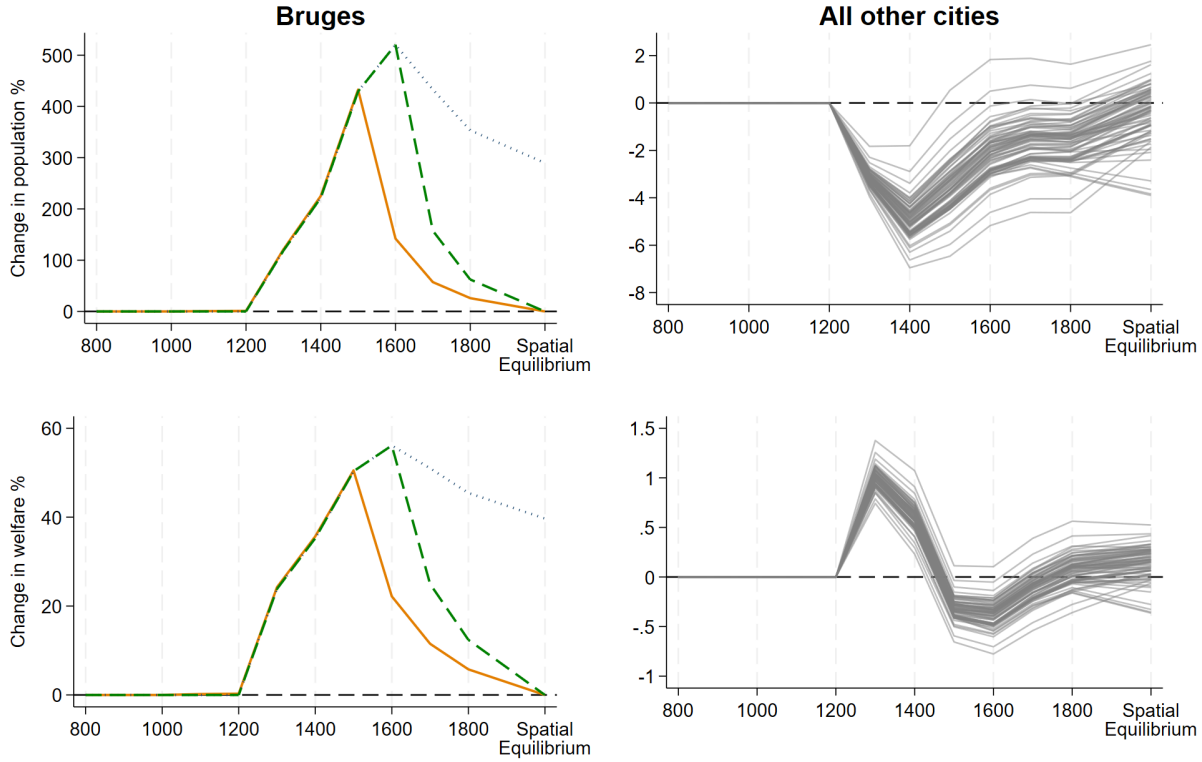


Notes: This figure replicates the counterfactual results of figure **??** varying the magnitude of the Zwin shock. I vary the magnitude from 0 (no Zwin counterfactual, depicted in red) to 1 (actual history, depicted in orange) to even more extreme shocks up to 2. Larger magnitudes are depicted in darker gray lines. Note that the y-axis scales vary across graphs for readability.

E.2 Varying shock length

Here I vary the shock length by considering what the impact of the Zwin would have been if it had silted up 100 years later or never silted up at all. Figure 11 shows unsurprisingly that maintaining the Zwin for an additional 100 years increases the population of Bruges over this period. The impacts of this longer-lasting Zwin are also more prominent in 1800. This counterfactual also doesn't alter the long-run spatial equilibrium, once again eventually all locations will revert to the same equilibrium distribution.

Figure 11 Counterfactual productivity fundamentals



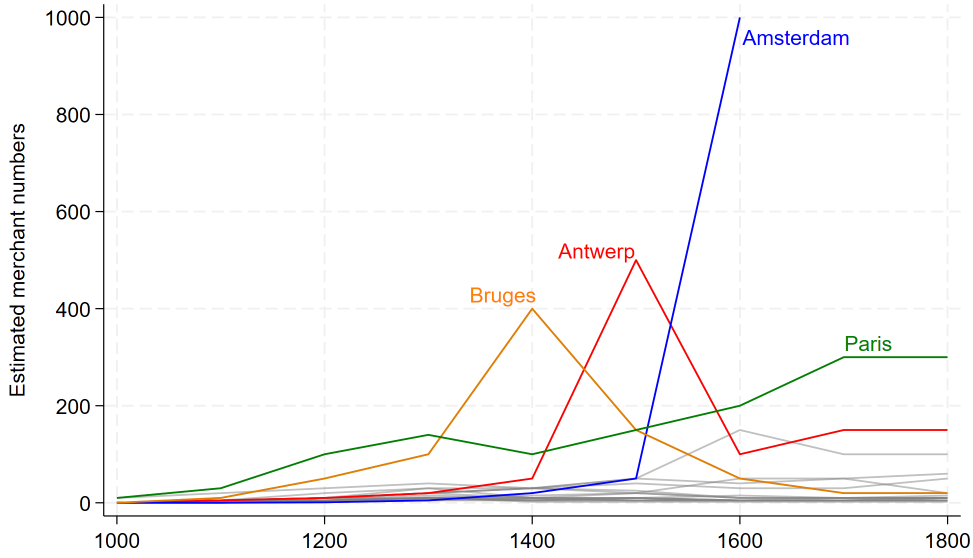
Notes: This figure replicates the counterfactual results of figure 4 varying the timing of the Zwin shock. In the solid line I plot the actual history, in the dotted line I plot the no-zwin counterfactual, in the dashed line I plot the counterfactual of the Zwin remaining navigable for 100 more years. In the right hand panels I show the effect of each of the remaining cities in the sample under the allways-Zwin counterfactual.

Contrastingly the "always Zwin" counterfactual manages to achieve a clearly different spatial equilibrium. In a world where the Zwin never silted up (and the positive productivity effect was maintained[12]), Bruges would have been a significantly bigger city by 1800. Additionally, other cities would have been smaller in 1800, as Bruges cannibalizes population from across the low countries. More strikingly, this counterfactual also causes a shift in the long-run spatial equilibrium. If the Zwin had never silted up Bruges would be larger today, at around three times its current size.

[12]This seems unlikely as the value of North Sea trade changes over time.

E.3 Validating generated Merchant data

Figure 12 Estimated Merchant Numbers



Notes: This figure plots estimated merchant numbers over time. I have not plotted merchant numbers for Amsterdam past 1600 which far exceed the current scale.

Table 5 Comparing generated and sourced merchant numbers

City	Year	Humans	ChatGPT
Bruges	1365	135	100-400
Bruges	1450	400	150-500
Antwerp	1560	1500	100-500
Amsterdam	1600	900	1000

Notes: This table compares generated merchant numbers from ChatGPT with those found in sources. "Human" merchant numbers are from Gelderblom [2013] and Murray [2005].

References

Daron Acemoglu, Simon Johnson, and James Robinson. The rise of europe: Atlantic trade, institutional change, and economic growth. *American economic review*, 95(3):546–579, 2005a.

Daron Acemoglu, Simon Johnson, and James A. Robinson. Chapter 6 institutions as a fundamental cause of long-run growth. volume 1 of *Handbook of Economic Growth*, pages 385–472. Elsevier, 2005b. doi: https://doi.org/10.1016/S1574-0684(05)01006-3. URL https://www.sciencedirect.com/science/article/pii/S1574068405010063.

Alberto Alesina, Paola Giuliano, and Nathan Nunn. On the origins of gender roles: Women and the plough. *The quarterly journal of economics*, 128(2):469–530, 2013.

Robert C Allen. The great divergence in european wages and prices from the middle ages to the first world war. *Explorations in economic history*, 38(4):411–447, 2001.

Treb Allen and Dave Donaldson. Persistence and path dependence in the spatial economy. Technical report, National Bureau of Economic Research, 2020.

Treb Allen, Costas Arkolakis, and Xiangliang Li. On the equilibrium properties of network models with heterogeneous agents. Technical report, National Bureau of Economic Research, 2020a.

Treb Allen, Costas Arkolakis, and Yuta Takahashi. Universal gravity. *Journal of Political Economy*, 128(2):393–433, 2020b.

Paul Bairoch, Jean Batou, and Chevre Pierre. *Population of European cities from 800 ò to 1850: database and summary analysis of the results (the)*, volume 4. Droz Library, 1988.

Hoyt Bleakley and Jeffrey Lin. Portage and path dependence. *The quarterly journal of economics*, 127(2):587–644, 2012.

Maarten Bosker, Eltjo Buringh, and Jan Luiten Van Zanden. From baghdad to london: Unraveling urban development in europe, the middle east, and north africa, 800–1800. *Review of Economics and Statistics*, 95(4):1418–1437, 2013.

Boris Bove and Claude Gauvard. *Le Paris du Moyen Âge*. Belin, 2015.

Eltjo Buringh. The population of european cities from 700 to 2000: Social and economic history. *Research Data Journal for the Humanities and Social Sciences*, 6(1):1–18, 2021.

Tertius Chandler. Four thousand years of urban growth: An historical census. *(No Title)*, 1987.

Roger H Charlier. Grandeur, decadence and renaissance. *Journal of Coastal Research*, pages 425–447, 2005.

Roger H Charlier. The zwin: from golden inlet to nature reserve. *Journal of Coastal Research*, 27(4):746–756, 2011.

Donald R Davis and David E Weinstein. Bones, bombs, and break points: the geography of economic activity. *American economic review*, 92(5):1269–1289, 2002.

Melissa Dell. The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903, 2010.

Brecht Dewilde, Jan Dumolyn, Bart Lambert, and Bram Vannieuwenhuyze. 'so one would notice the good navigability': economic decline and the cartographic conception of urban space in late fifteenth-and sixteenth-century bruges. *Urban History*, 45(1):2–25, 2018.

Dave Donaldson and Richard Hornbeck. Railroads and american economic growth: A "market access" approach. *The Quarterly Journal of Economics*, 131(2):799–858, 2016.

Jan Dumolyn and Ward Leloup. The zwin estuary: a medieval portuary network. In *Las sociedades portuarias de la Europa atlántica en la edad media*, volume 35, pages 197–212. Instituto de Estudios Riojanos, 2016.

Gilles Duranton and Diego Puga. The economics of urban density. *Journal of economic perspectives*, 34(3):3–26, 2020.

Fabian Eckert and Michael Peters. Spatial structural change. Technical report, National Bureau of Economic Research, 2022.

Sebastian Ellingsen. Long-distance trade and long-term persistence. *University of Pompeu Fabra Working Paper*, 2021.

Stef Espeel. De brugse stadsmagistraat in de late 14e eeuw. *Een prosopografische studie voor de*, 2016.

Giovanni Federico, Max-Stephan Schulze, and Oliver Volckart. European goods market integration in the very long run: from the black death to the first world war. *The Journal of Economic History*, 81(1):276–308, 2021.

Alfred Fierro. History and dictionary of paris. *(No Title)*, 1996.

Oscar Gelderblom. *Cities of Commerce: The Institutional Foundations of International Trade in the Low Countries, 1250–1650*. Princeton University Press, Princeton, 2013. Analyzes Bruges, Antwerp, and Amsterdam merchant communities; dataset of southern merchants to Amsterdam.

Oscar Gelderblom. *Cities of commerce: the institutional foundations of international trade in the Low Countries, 1250-1650*, volume 45. Princeton University Press, 2015.

Oscar Gelderblom. Merchants from the southern netherlands (1578–1630): Social and economic history. *Research Data Journal for the Humanities and Social Sciences*, 1(1):13–22, 2016.

Stephen Gibbons, Stephan Heblich, and Ted Pinchbeck. The spatial impacts of a massive rail disinvestment program: the beeching axe. 2018.

Edward Glaeser. *Triumph of the city: How urban spaces make us human*. Pan Macmillan, 2011.

Edward L Glaeser. Urban resilience. *Urban studies*, 59(1):3–35, 2022.

W Walker Hanlon and Stephan Heblich. History and urban economics. *Regional Science and Urban Economics*, 94:103751, 2022.

Stephan Heblich, Stephen J Redding, and Yanos Zylberberg. The distributional consequences of trade: evidence from the repeal of the corn laws. *Work in progress and incomplete*, 2020.

Remi Jedwab and Adam Storeygard. The average and heterogeneous effects of transportation investments: Evidence from sub-saharan africa 1960–2010. *Journal of the European Economic Association*, 20(1):1–38, 2022.

Remi Jedwab, Edward Kerby, and Alexander Moradi. History, path dependence and development: Evidence from colonial railways, settlers and cities in kenya. *The Economic Journal*, 127(603):1467–1494, 2017.

Remi Jedwab, Noel D Johnson, and Mark Koyama. Pandemics, places, and populations: Evidence from the black death. 2019.

Emily Kadens. Medieval law merchant: The tyranny of a construct. *Journal of Legal Analysis*, 7(2):251–289, 2015. Discusses Ypres merchants and medieval trade networks.

Bart Lambert. Merchants on the margins: fifteenth-century bruges and the informal market. *Journal of Medieval History*, 42(2):226–253, 2016.

Jeffrey Lin and Ferdinand Rauch. What future for history dependence in spatial economics? *Regional Science and Urban Economics*, 94:103628, 2022.

Jurg Luterbacher, Daniel Dietrich, Elena Xoplaki, Martin Grosjean, and Heinz Wanner. European seasonal and annual temperature variability, trends, and extremes since 1500. *Science*, 303(5663):1499–1503, 2004.

Jakob B Madsen, Peter E Robertson, and Longfeng Ye. Lives versus livelihoods in the middle ages: The impact of the plague on trade over 400 years. *European Economic Review*, 162: 104654, 2024.

Paolo Malanima. Italian cities 1300-1800. a quantitative approach. *Rivista di storia economica*, 14(2):91–126, 1998.

Guy Michaels and Ferdinand Rauch. Resetting the urban network: 117–2012. *The Economic Journal*, 128(608):378–412, 2018.

Edward Miguel and Gerard Roland. The long-run impact of bombing vietnam. *Journal of development Economics*, 96(1):1–15, 2011.

James M Murray. *Bruges, cradle of capitalism, 1280-1390*. Cambridge University Press, 2005.

Dávid Krisztián Nagy. Quantitative economic geography meets history: Questions, answers and challenges. *Regional Science and Urban Economics*, 94:103675, 2022.

Nathan Nunn. The historical roots of economic development. *Science*, 367(6485):eaaz9986, 2020.

Emile Picarda. *The Water Merchants: Parisian Hanseatic League and French Company*, volume 134. É. Bouillon, 1901.

Walter Prevenier. Bevolkingscijfers en professionele strukturen der bevolking van gent en brugge in de 14de eeuw. In *Album Charles Verlinden*, volume 46, pages 269–303. Belgisch centrum voor landelijke geschiedenis, 1975.

Jeroen Puttevils. *The ascent of merchants from the Southern Low Countries: from Antwerp to Europe, 1480-1585*. PhD thesis, University of Antwerp, 2012.

Stephen J Redding and Daniel M Sturm. The costs of remoteness: Evidence from german division and reunification. *American Economic Review*, 98(5):1766–1797, 2008.

M Ryckaert. De verzanding van het zwin. *Water. Tijdschrift over waterproblematiek*, 8: 204–207, 1989.

Carol H Shiue and Wolfgang Keller. Markets in china and europe on the eve of the industrial revolution. *American Economic Review*, 97(4):1189–1216, 2007.

Tim Soens, Dries Tys, and Erik Thoen. Landscape transformation and social change in the north sea polders, the example of flanders (1000-1800 ad). *Siedlungsforschung. Archäologie-Geschichte-Geographie*, 31:133–160, 2014.

Peter Stabel, Jeroen Puttevils, and Jan Dumolyn. Production, markets and socio-economic structures i, c. 1100-c. 1320. In *Medieval Bruges: c. 850-1550*, pages 86–123. Cambridge University Press, 2018.

Jan Trachet, Samuël Delefortrie, Kristof Dombrecht, Jan Dumolyn, Ward Leloup, Erik Thoen, Marc Van Meirvenne, and Wim De Clercq*. Turning back the tide: The zwin debate in perspective. a historiographical review of the medieval port system northeast of bruges. *Revue du Nord*, (5):305–321, 2015.

Jan A Van Houtte. The rise and decline of the market of bruges. *The Economic History Review*, 19(1):29–47, 1966.

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

Étienne Boileau. *Le Livre des Métiers de Paris*. 1268. Contains data on Paris trade guilds and merchant taxation.