KU LEUVEN

DEPARTMENT OF ECONOMICS

Persistent Spatial Equilibria. Evidence from a Sudden River.

Luke Heath Milsom

FACULTY OF ECONOMICS AND BUSINESS

DISCUSSION PAPER SERIES DPS 24.03

January 2024

Persistent Spatial Equilibria. Evidence from a Sudden River.

Luke Heath Milsom*

8th January, 2024

Abstract

This paper asks what impact a large, but temporary, productivity shock can have on the spatial distribution of economic activity across cities in the short and long run. To answer this question I use a dynamic quantitative spatial economics model and the natural experiment of a sudden river, the Zwin, that connected Bruges to the North Sea in the 12th century. I show that despite dramatic short-run impacts in Bruges and across the Low Countries during the period the Zwin was navigable as well as in the centuries after, this shock failed to alter the prevailing long-run spatial equilibrium. Simulating alternative shock magnitudes or locations also doesn't result in a change in spatial equilibrium, but a permanent shock would have. However, convergence is slow, in 1800 some 300 years after the Zwin became impassable aggregate welfare remains on average 2% higher across the Low Countries as compared to the counterfactual world where the Zwin had never existed.

JEL: R12, R13, N9, O18

*Heath Milsom: The Institute For Fiscal Studies and KU Leuven Department of Economics, luke.heath.milsom@ifs.org.uk. With thanks to: Hannah Zillessen, Ferdinand Rauch, Eric Buyst, Frederic Vermeulen, Max Merczinek, and Shihang Hou.

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 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. I then turn to ask what impact a different shock (what if the river had appeared elsewhere, or not silted up) would have had on the short-run and long-run urban structure of the Low Countries.

Using historic population data from Bairoch et al. [1988] I first show that during the period when the Zwin was navigable Bruges experienced a population boom and Antwerp a decline. To ask what impact this had on the entire urban structure of the low countries, I build a dynamic quantitative spatial economics model following Allen and Donaldson [2020]. Locations interact through costly migration and costly trade, and the past impacts today through historical amenity and productivity spillovers. I estimate and calibrate this model using available data on historic populations and existing parameter estimates from the literature. I then invert the model to back out rationalizing location-period specific fundamental productivities which I are used to estimate the impact the Zwin had on Bruges and Antwerp's productivity. Within this framework, the Zwin is modeled as a productivity

[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".

shock. Intuitively, sudden access to lucrative North Sea trade provided a more productive outlet for labor than that which was previously available.[3]

Using the estimated model I find that if the Zwin had never formed Burges's population in 1800 would have been similar to what it actually was. However, Antwerp, Bruges's main competitor for North Sea trade, would have been around 40% larger. 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. In 1400 Brussels would have been larger, but in 1500 smaller under the no-Zwin counterfactual. Even further afield, the impact of the Zwin can still be felt in Amsterdam in 1800 — in the no-Zwin counterfactual Amsterdam would have been 7.40% smaller. These variations in population are mirrored in location-specific counterfactual welfare. In terms of aggregate welfare across the whole Low Countries, 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.[4]

Although these changes are large and long-lived, and parameter estimates imply that multiple long-run spatial equilibria exist — I don't find evidence that the Zwin shock was sufficient to cause a change in the long-run urban equilibrium. Eventually, all cities converge back to the same equilibrium population distribution in the no-Zwin counterfactual as in actual history. This convergence however takes multiple centuries, and the impact of the Zwin can certainly still be felt in 1800.

To investigate what shocks could have changed the long-run equilibrium, I consider counterfactuals varying the magnitude of the Zwin shock, the affected city, or the length of time the Zwin was navigable. I find that changing the affected city or magnitude has no impact on the resulting long-run spatial equilibrium. These alternative counterfactuals do cause significantly different transition paths, with meaningful ramifications on population and welfare over a long horizon. However, if the Zwin had never silted up and was still navigable today, the spatial distribution of economic activity in the Low Countries would have looked quite different. Results from the model suggest that in this counterfactual the Low Countries converge to a new spatial equilibrium where Bruges consists of over 25% of the total (urban) population and Antwerp and Amsterdam are considerably smaller.[5]

[3] In this setting a "productivity shock" and a "trade shock" are isomorphic in the sense that the only direct way either interact with the model is to increase incomes in a given location.

[4] 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.

[5] For comparison, in 1800 Amsterdam consisted of 23% of the population of the 29 Low Country cities

Leveraging the full structure of the model I can also consider the city-specific and aggregate welfare effects of various counterfactuals. At the city level, changes to welfare mirror those of population as intuitively people move into more attractive i.e. high welfare areas. Turning to aggregate welfare, I find that in 1800 average welfare across the Low Countries was 2% higher as a result of this, by now, low obsolete river. I also consider the aggregate welfare effects under the counterfactual scenarios of a Zwin-like shock (in terms of productivity change) in each of the 29 cities in my sample. Shocks in different locations cause different changes in aggregate welfare (although in the long run, all converge to the same). Locations where the shock caused the largest gain in welfare in 1400 subsequently caused the largest falls in aggregate welfare in 1500 and beyond once the shock dissipated, relative to a no-Zwin world. Locations where shocks caused the largest gains in welfare in 1400 are in general more centrally located (higher market access), are more productive, and have larger populations. This reversal of fortunes is due to spatial misallocation. In "better" places the initial positive shock had a bigger impact, causing more people to move, when the shock is then removed more people are in the "wrong" place causing a reduction in welfare and an adjustment process as they slowly move back to their optimal location.

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,

that I consider.

Eckert and Peters, 2022, Ellingsen, 2021, Heblich et al., 2020]. I build on this literature by considering the dynamic impact of a temporary shock and combining counterfactual results with reduced-form evidence.

The rest of this paper proceeds as follows: Section 1 describes the historical context, data, and reduced form evidence, section 2 sets up the dynamic quantitative spatial economics model, section 3 then describes the main counterfactual results, section 4 considers alternative counterfactuals, and section 5 discusses welfare implications, finally section 6 concludes.

1 Historical context, data, and the reduced form impact of the Zwin on Bruges and Antwerp.

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[6]. However,

[6] 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".

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. 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

The main data for this paper are population estimates from Bairoch et al. [1988] which have been widely used in economics, see for example Bosker et al. [2013], Acemoglu et al. [2005], Glaeser [2022], Duranton and Puga [2020].[7] I use data on the population of 29 cities across the Low Countries every century from 1000 CE to 1800 CE. These 29 cities are chosen as those within modern-day Belgium, the Netherlands, and Luxembourg which appear in both the Bairoch et al. [1988] and the Chandler [1987] datasets. My sample ends in 1800 as after this date transportation technologies changed significantly. In this historical setting, systematic data on wages, prices, or housing is not available. As a result of this, I am forced to make various additional assumptions and simplifications, beyond those that are typically necessary when using quantitative spatial economics models.

Although I mainly use population data from Bairoch et al. [1988], this data source does have some known limitations and inaccuracies. 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,

[7] 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.

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.

1.3 Reduced form evidence

In figure 1 I plot the population of Antwerp and Bruges over time, indicating in gray the period when the Zwin was most likely navigable. Figure 1 descriptively shows a large spike in population in Bruges, and a corresponding dip in Antwerp, over the period when the Zwin is navigable[8]. Of course, figures such as 1 capture all changes in Bruges/ Antwerp 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.

[8] 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.

Figure 1 Reduced form evidence

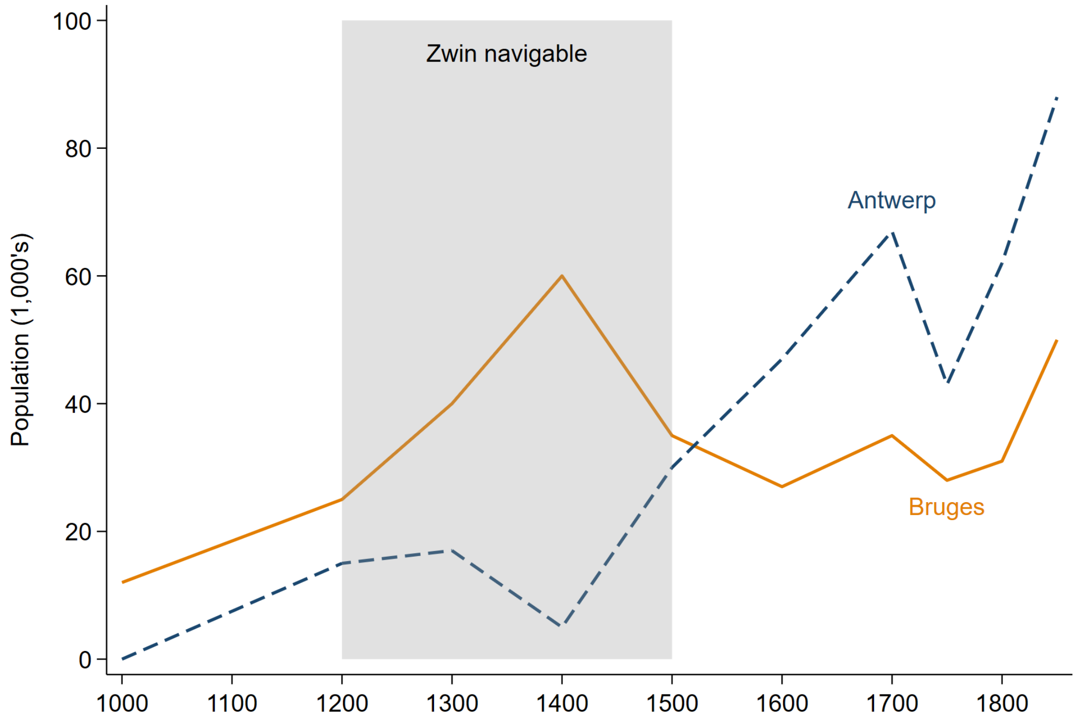

Notes: This figure shows the reduced form impact of the Zwin on Bruges and Antwerps population. Historical population data is mainly from Bairoch et al. [1988].

In appendix C I perform a synthetic control analysis to get a little closer to a plausibly causal effect and find confirmative results — Bruges is for a time above its counterfactual, and Antwerp below. In table 1 I quantify the impact shown visually in figure 1. To do this I construct a variable "$ZwinImpact$", which takes the value one for Bruges in 1300 and 1400 (when the Zwin was navigable), and minus one for Bruges in 1500, when the Zwin silted up and Antwerp enjoyed the trade that previously flowed through it. Analogously $ZwinImpact$ takes the value minus one for Antwerp in 1300 and 1400, and plus one in 1500. In table 1 I then report the estimated coefficient on this variable regressed against the local population. The results in table 1 show a robust and significant positive impact of the Zwin on population. In periods when the Zwin positively impacted a location, the population grew by around 10,000.

Table 1 Direct Impact of the Zwin on Population

	(1)	(2)	(3)	(4)
ZwinImpact	10.38^{*}	10.38	10.41^{*}	11.11^{**}
	(6.126)	(7.651)	(5.409)	(5.478)
Lat Lng controls			X	
Country FE			X	
Year FE		X	X	X
City FE				X
R2	0.00642	0.208	0.266	0.507
Observations	232	232	232	232

Notes: This table quantifies the reduced form impact of the Zwin on the populations of Bruges and Antwerp. The variable "*ZwinImpact*", takes the value one for Bruges in 1300 and 1400 (when the Zwin was navigable), and minus one for Bruges in 1500, when the Zwin silted up and Antwerp enjoyed the trade that previously flowed through it. Analogously *ZwinImpact* takes the value minus one for Antwerp in 1300 and 1400, and plus one in 1500. The estimated coefficient on this variable regressed against the local population is then reported in this table. In column one, I report the coefficient absent of controls. In column two I control for year fixed effects. In column three I control for a polynomial of degree two in latitude and longitude, country fixed effects, and year fixed effects. In column four I control for year and city fixed effects. Standard errors are robust.

Figure 1 and table 1 show indicative evidence that the Zwin had a large impact on the population of Bruges and Antwerp. To go beyond this, and ask how this sudden river impacted the urban structure of the entire area, or counterfactual questions such as "What if the Zwin never silted up", I turn next to developing a dynamic quantitative spatial economics model.

2 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 impose structure on the data and possible counterfactual path of each city by turning to a dynamic quantitative spatial economics model. Here I follow Allen and Donaldson [2020] closely, and present a slightly simplified version of their model. Within this framework, 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.

There are arbitrarily many locations $i \in N$ and $t \in \mathcal{T}$ time periods. 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.

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

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 2.

$$W_{it} = u_{it} \frac{w_{it}}{P_{it}}, \qquad u_{it} = \bar{u}_{it} L_{it}^{\beta_1} L_{it-1}^{\beta_2} \tag{2}$$

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].

$$X_{ijt} = \tau_{ijt}^{1-\sigma} \left(\frac{w_{it}}{P_{it}} \right)^{1-\sigma} P_{jt}^{\sigma-1} w_{jt} L_{jt}, \qquad P_{it} = \left(\sum_{k=1}^{N} \tau_{kit} \left(\frac{w_{kt}}{A_{kt}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{3}$$

Individuals decide where to move to maximize utility given in 2 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 4.

$$L_{ijt} = \mu_{ijt}^{-\theta}\Pi_{it}^{-\theta}W_{jt}^{\theta}L_{it-1} \tag{4}$$

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

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].

2.1 Persistence

The dynamic quantitative spatial economics model described above has the attractive property that, depending on the parameter values, it can exhibit within-period equilibrium uniqueness but long-run equilibrium multiplicity [Allen and Donaldson, 2020]. Intuitively α_1 and β_1 govern the within-period equilibrium properties of the model, if they are sufficiently small the dynamic equilibrium will exist and be unique. More specifically, if $\alpha_1 + \beta_1 < 1/\theta$, that is contemporaneous agglomeration forces are greater than the contemporaneous dispersion forces. If this condition holds the model will have a unique transition path.

Turning to the long-run equilibrium, that is the equilibrium to which the economy's transition path is converging. As discussed by Allen and Donaldson [2020], intuitively if agglomeration forces are strong enough location decisions may become self-reinforcing, and thus multiple equilibria could arise. The intuition here translates into a very similar condition as that discussed above, only now total agglomeration forces are what is important. That is multiple long-run equilibria will arise if $\alpha_1 + \alpha_2$ and $\beta_1 + \beta_2$ are sufficiently large. If this is the case the economy has the potential to exhibit path dependence where different initial population distributions may cause the economy to converge to different long-run equilibria. Indeed, each potential long-run equilibrium will be associated with a "basin of attraction", that is a set of possible initial population distributions that converge to the specific equilibrium. Therefore, although multiple equilibria may exist, it remains an empirical question to ask whether any one specific shock would be sufficient to cause a shift from one equilibrium to the next.

2.2 Estimation and calibration

The low countries from 1200 to 1800 are not a data-rich environment. Data on wages, intra-city migration, or intra-city trade is not available, but I do have data on populations every century. In this section, I describe how I overcome these data restrictions by leveraging fully the structure of the model and the data that is available. However, unavoidably, I am forced to make various somewhat unpalatable assumptions and calibrate almost all of the model parameters from the literature.

To identify the impact of the Zwin on the population of Belgian and Dutch cities I use a five-step procedure. The first three steps follow those of Allen and Donaldson [2020] in formulating migration and trade costs, solving for market access terms, and using a Rosen-Rosenback type regression to estimate model parameters. Step four is additionally necessary as without wage data the model is not invertible, and so I once again leverage population data and the structure of the model with a simplifying assumption to back out consistent fundamental productivities. Step five estimates the path of these fundamental productivities in the counterfactual world where the Zwin never appeared.

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}.

$$\tau_{ij}^{1-\sigma} = \exp\left(f_\tau(t_{ijt})\right) \qquad \mu_{ij}^{-\theta} = \exp\left(f_\mu(t_{ijt})\right) \tag{5}$$

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: $f_\xi(t_{ijt}) = \kappa_\xi \cdot t_{ijt}$ 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. Thus, I further simplify matters by estimating t_{ijt} in a time-invariant manner as the Haversine distance between each pair of cities based on their centroid latitude and longitude coordinates (as-the-crow flies distance).

Step 2: Find labor market access terms.

Having calculated transport costs, I can calculate labor market access terms using available population data. Denote exponent-inclusive iceberg costs as $M_{ij} = \mu_{ij}^{\theta}$ and noting that step one described above identifies M we can thus 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].

$$LMA_{it}^{\theta} = \Pi_{it}^{-\theta} = \sum_{j} M_{ij} \frac{L_{jt}}{\Pi_{jt}^{\theta}} \tag{6}$$

Step 3: Use a simple model-implied population regression to estimate β_2. Armed with the market access terms from step two, I can now turn to find the remaining parameters: $\{\sigma, \theta, \alpha_1, \alpha_2, \beta_1, \beta_2\}$. In this data-poor environment, it will not be possible to estimate all parameters credibly, instead, I focus on estimating the strength of the persistence effects in my setting. I focus on this one dimension because as discussed in subsection 2.1 it can crucially determine whether multiple equilibria exist. Thus, I take as my baseline the parameterization from Donaldson and Hornbeck [2016] with the exception of β_2 which I estimate. That is I take $\{\sigma = 9, \theta = 4, \alpha_1 = 0.19, \alpha_2 = -0.041, \beta_1 = -0.26\}$.

Intuitively the magnitude of β_2 should reflect the size of the persistence of city populations across time. If β_2 is large it implies that infrastructure built in the past has a large impact on utility today. One way to directly get at this would be to regress (log) population today on (log) population last period, where a larger coefficient would imply a larger β_2. Within the dynamic quantitative spatial economics model the same intuition holds with two adjustments. First, the model tells us that such a regression would suffer from endogeneity concerns as one needs to take into account the spatial properties of the network by including market access terms. Second, by solving for (log) population as a function of past population and market access terms we can determine how to correctly scale the coefficient on past population to find β_2.

To go from the theoretical identification approach discussed above to something empirically implementable I need to overcome three further challenges. First, I don't have data on wages, so cannot calculate trade market access. In this Malthusian period, one could proxy output with population and calculate trade market access as $MA_{it} = \sum_j \tau_{ij}^{1-\sigma} L_{jt}/MA_{jt}$, however, the resulting variable is very highly correlated with labor market access and so I opt not to include it in the subsequent regression. Second, as local productivity causes higher market access and shows up in the error term, market access will be correlated with the error. To overcome this I instrument market access with past market access not including i's population itself i.e. $MA_{it}^{IV} = \sum_{j \neq i} \mu_{ij}^{-\theta} L_{jt-1}/MA_{jt}^{IV}$. Third, as highlighted by Borusyak and Hull [2020] the IV for market access terms will suffer from the problem of endogenous exposure to exogenous shocks and thus need to be re-centered. As transport costs are fixed in my setting I take the average of market access terms over permuted population changes and use this "expected market access" to recenter my instruments.

Estimating the above-discussed regression using the IV procedure I find $\hat{\beta}_2 = 0.36$ which is close to the 0.31 estimated by Allen and Donaldson [2020].

The parameter values for $\alpha_1, \alpha_2, \beta_1, \beta_2$ imply, following the propositions described by Allen and Donaldson [2020], that the dynamic (short-run) equilibrium of the model exists and is unique, whereas the long-run equilibrium exhibits path-dependency. That is, multiple long-run equilibria exist, and different previous-period population distributions may lead to different long-run equilibria.

Step 4: Via solving and inverting the model back out fundamental productivities consistent with the observed population data.

As I don't have data on wages I cannot separately identify amenity and productivity fundamentals in each location. Instead, I make the simplifying assumption that amenities are constant over space, but allow productivity to vary. Thus, conditional on homogeneous amenities, I can solve for the vector of exogenous productivities in each time period that recovers the observed distribution of population.

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

Once I have recovered the fundamental productivities I turn to estimating the impact the appearance and disappearance of the Zwin had on them. As the historical account only suggests the Zwin impacted Bruges (positively) and Antwerp (negatively) I simply estimate its 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. This is undoubtedly a limitation of this approach, although, without access to better data, it is difficult to see how a more precise approach could be taken.

The specification is shown in the equation below. $\beta_{B,+}$ captures the positive impact the Zwin had on Bruges in 1300 and 1400, whereas $\beta_{B,-}$ captures the negative impact on Bruges in 1500. Analogously $\beta_{A,+}$ and $\beta_{A,-}$ capture the positive and negative impact the Zwin had on Antwerp in 1500 and 1300,1400 respectively. I further restrict the coefficients to be equal in absolute value, to reflect the fact that we consider the Zwin as location-specific productivity-enhancing, but not as something that creates productivity in and of itself. This is because the Zwin increases productivity through the opportunities trade provides, but does not in and of itself cause more trade in the aggregate.[9] Thus any benefit to Bruge will be reflected in a negative impact on Antwerp and vice-versa.

$$\begin{aligned}\bar{A}_{it} =& \beta_{B,+} \cdot \mathbb{1}[i = \text{Bruges}, t = 1300, 1400] - \beta_{B,-} \cdot \mathbb{1}[i = \text{Bruges}, t = 1500] \\ &+ \beta_{A,+} \cdot \mathbb{1}[i = \text{Antwerp}, t = 1500] - \beta_{A,-} \cdot \mathbb{1}[i = \text{Antwerp}, t = 1300, 1400] \\ &+ \gamma_i + \gamma_i \times t + \tau_t + \varepsilon_{it}\end{aligned}$$

I can then estimate the path of productivities in the absence of the Zwin as $\bar{A}^{NZ}_{it} = \bar{A}_{it} - \hat{\beta}_{B,+} \cdot \mathbb{1}[i = \text{Bruges}, t = 1300, 1400] - \hat{\beta}_{B,-} \cdot \mathbb{1}[i = \text{Bruges}, t = 1500] + \hat{\beta}_{A,+} \cdot \mathbb{1}[i = \text{Antwerp}, t = 1500] - \hat{\beta}_{A,-} \cdot \mathbb{1}[i = \text{Antwerp}, t = 1300, 1400]$. I estimate a statistically significant coefficient value of 9.375(4.405). The magnitude itself has no direct interpretation, but in a relative sense, this figure implies that the Zwin increased Bruges' productivity in 1400 by around two-thirds.

Figure 7 in the appendix graphically displays the estimated path of fundamental productivities taken in both Bruges and Antwerp. In orange, I plot the estimated productivities that rationalize the observed population distribution. One can clearly see in Bruges a significant spike in 1300 and 1400 as well as a significant dip in 1500. Similarly, Antwerp clearly displays a dip in 1300 and 1400 coupled with a spike in 1500. I then plot in a black dashed line the estimated path of productivity in the absence of the Zwin. In both cases, this path effectively smooths out the spikes and dips described before.

3 Results: The long-run shadow of a temporary river

Figure 2 shows the results from running my main counterfactual analysing the counterfactual world where the Zwin had never existed. In each panel, I plot the path of relative populations for each of Bruges, Antwerp, Brussels, and Amsterdam. Figure 8 in the Appendix shows the results for all 29 cities in my sample. In orange, I plot the historical path the population actually took, and in dotted blue the counterfactual path it would have taken if the Zwin

[9] Note that due to reduced aggregate transport costs, or increasing trade in an otherwise more productive location, the aggregate impact of the Zwin can still be to increase or decrease welfare.

had never appeared.

Figure 2 The impact of the Zwin on populations

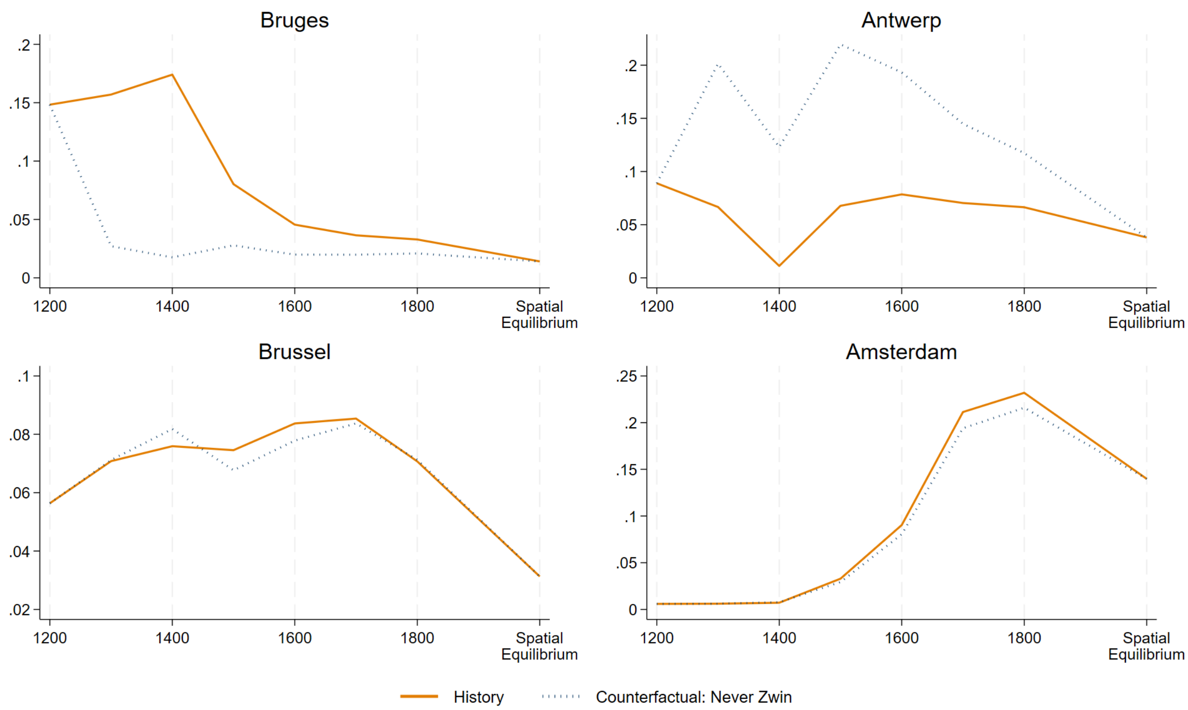

Notes: This figure shows the impact of the Zwin on the relative population of Bruges, Antwerp, Brussels, and Amsterdam. Populations are relative to the total population of the urban system, that is the whole Low Countries. In orange, I show the historical i.e. actual path of relative populations. The dashed line shows the counterfactual model-implied path that each city would have taken in the absence of the Zwin. Note that the y-axis scales vary across graphs for readability.

The impact of the Zwin whilst it was navigable (1200-1500) on Bruges and Antwerp is clear — it caused the population to increase significantly in Bruges and decrease significantly in Antwerp. The contemporaneous impact can also be seen as far afield as in Brussels, although there is little discernable impact on Amsterdam.

Turning to the longer-run transition path I find significant impacts of the Zwin many centuries after it became impassable. Whereas Bruges returned fairly rapidly to something close to its historical path after the Zwin silted up (around 1500), even in 1800 the counterfactual population of Antwerp is significantly different. If the Zwin had never appeared the model suggests that Antwerp would have been 11.73% of the total Low Countries urban population in 1800 whereas it actually accounted for only 6.64% of the total. Although there is no discernable impact on the population of Brussels by 1800, the Zwin does appear to leave a significant effect until 1700 or so. In Amsterdam, the impact of the Zwin appears to increase over time (at least until 1800), with the actual 1800 population 7.40% higher than that in the absence of the Zwin.

Projecting forwards from 1800 fixing amenities and productivities at their 1800 levels, I can also analyze whether the Zwin caused a fundamental change to the long-run spatial equilibrium of the economy. Although the model parameters allow multiple equilibria and short-run impacts appear large, I find that the economy eventually converges to the same long-run spatial equilibrium. However, the pace of convergence is extremely slow, for all practical purposes the Zwin did have an impact on the long-run spatial distribution of economic activity. In section 5 I discuss welfare implications.

4 Alternative shocks

Figure 2 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 same shock (in terms of productivity) had affected a different city. To do this I randomise which of the 29 Low Country cities in my sample get the positive Zwin shock that in reality was enjoyed by Bruges, and for simplicity don't alter Antwerps productivity fundamentals. Third, 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?

4.1 Varying shock magnitude

I vary the magnitude of the "Zwin shock" both on Bruges and Antwerp by multiplying the estimated impact of the Zwin by a factor γ which varies from 0 to 2. For each new counterfactual path of productivities, I estimate the dynamic quantitative spatial model described in section 2 and plot the results from Bruges, Antwerp, Brussels, and Amsterdam in figure 3. In figure 3 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 3 shows that even relatively small shocks would have had a significant contemporaneous impact on the populations of Bruges and Antwerp, as well as to a lesser extent on Brussels and Amsterdam. However, in all cases, the same general pattern is followed, of a slow convergence back eventually to the same long-run spatial equilibrium. Figure 3 also highlights the spatial interactions between cities. Larger shocks caused Bruges to grow more in 1400 at the expense of its neighbors including Brussels and Antwerp. However, in 1500 this shock instead benefits Antwerp and Bruges is hurt to a greater extent whereas Brussels gains as the influx of population from Bruges is less than the outpouring to Antwerp which remains relatively unproductive due to its diminished population.

Figure 3 Counterfactual productivity fundamentals

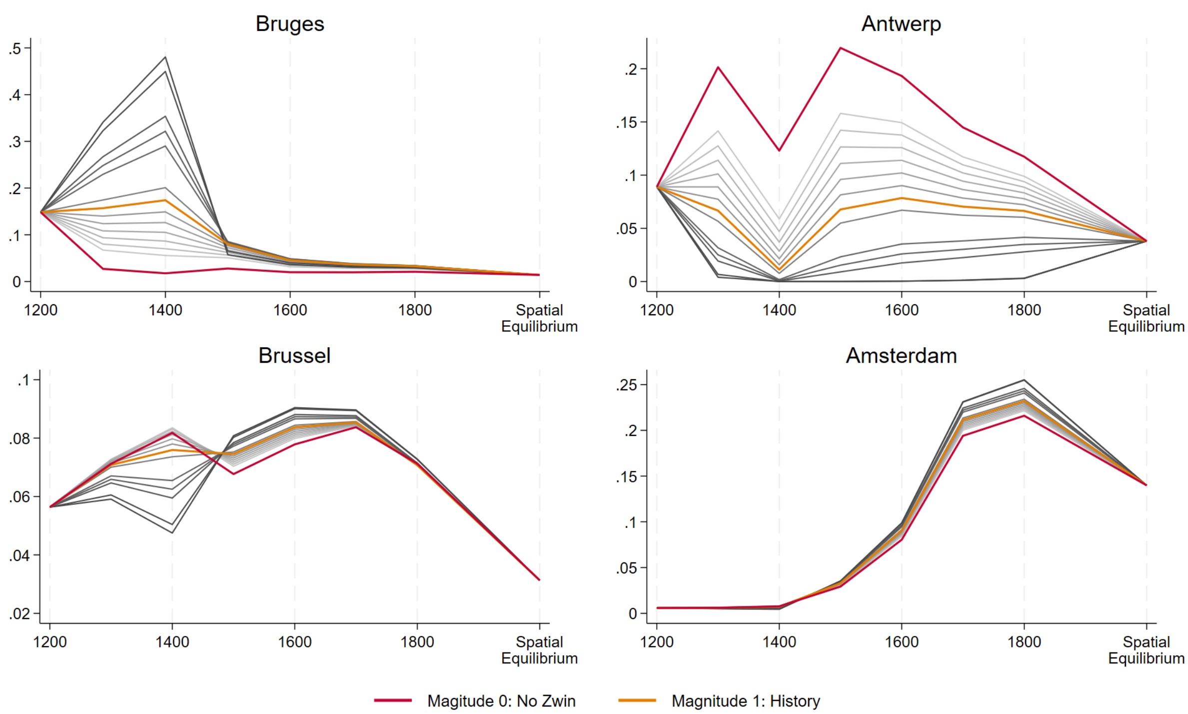

Notes: This figure replicates the counterfactual results of figure 2 varying the magnitude of the Zwin shock. I vary the magnitude from 0 (no Zwin counterfactual, depicted in red) to 1 (actual history, depicted in orange) to even more extreme shocks up to 2. Larger magnitudes are depicted in darker gray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

4.2 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. The estimated alternative paths for the fundamental productivities of Bruges and Antwerp can be found in figure 9 in the appendix.

Figure 4 shows unsurprisingly that maintaining the Zwin for an additional 100 years increases the population of Bruges over this period at the expense of Antwerp. The impacts of this longer-lasting Zwin are also more prominent in 1800 although most of the resulting "damage" to Antwerp appears to be achieved by the shorter-lived river. Amsterdam and Brussels similarly see minor changes in their populations. This counterfactual also doesn't alter the long-run spatial equilibrium, once again eventually all locations will revert to the same equilibrium distribution.

Figure 4 Counterfactual productivity fundamentals

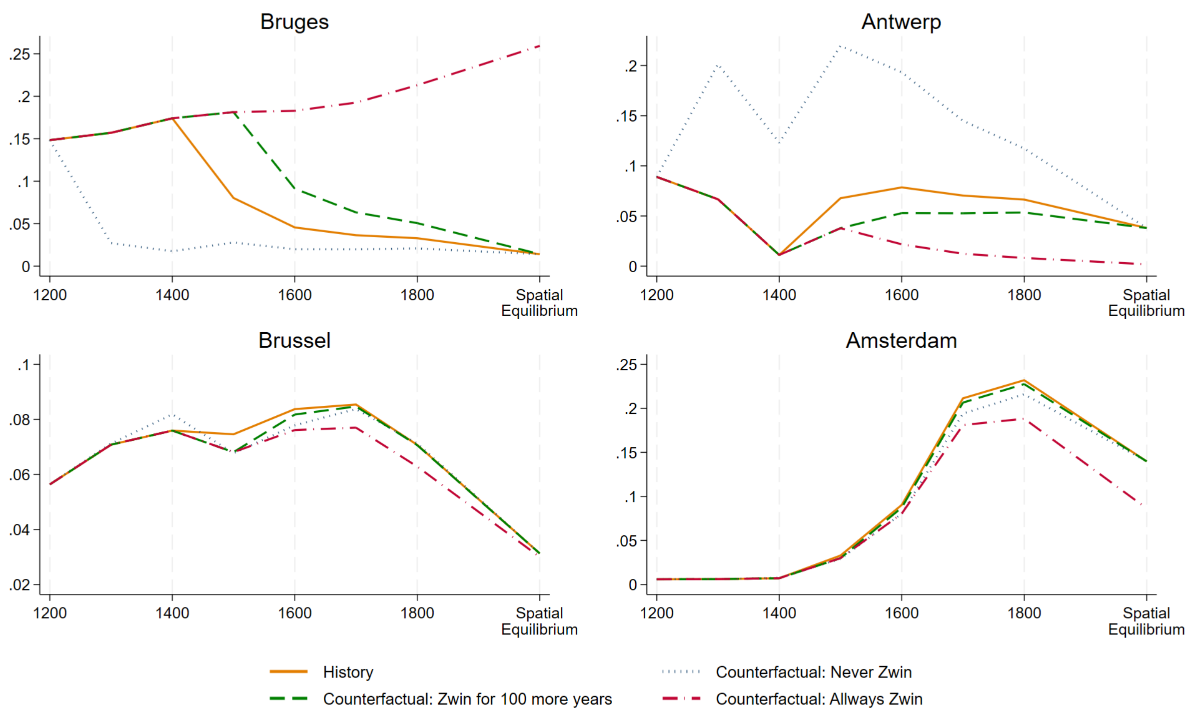

Notes: This figure replicates the counterfactual results of figure 2 varying the timing of the Zwin shock. In this figure, the y-axis is equalized across panels. In the solid line I plot the actual history, in the dotted line I plot the no-zwin counterfactual, in the dashed line I plot the counterfactual of the Zwin remaining navigable for 100 more years. Finally, in the dot-dashed line I plot the counterfactual depicting the scenario where the positive impact of the Zwin on Bruges, and the negative impact on Antwerp, never dissipated.

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[10]), Bruges would have been a significantly bigger city by 1800, whereas Antwerp is all but unpopulated. Additionally, Brussels and Amsterdam would have been smaller in 1800, Amsterdam more so, as Bruges cannibalizes population from across the low countries. More strikingly, this counterfactual also causes a shift in the long-run

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

spatial equilibrium. If the Zwin had never silted up Bruges would be larger today, accounting for around 25% of the total population of the 29 cities considered (for comparison in 1800 Amsterdam consisted of 23% of the total population). Interestingly, although Brussels eventually recovers to a similar spatial-equilibrium population, Amsterdam converges to a significantly smaller population.

5 Welfare

In the above discussion, I have focused on the tangible variable of local population. However, I can leverage the full structure of the model to back out welfare. City-period specific welfare in the model employed is given simply as $W_{it} = u_{it}\frac{w_{it}}{P_{it}}$, which intuitively is increasing in local real wages and amenities. Figures 10 and 11 in the appendix show the corresponding figures to 2 and 4 plotting welfare rather than population. Assuringly changes to welfare mirror those of population, intuitively if an area is more attractive (higher welfare) more people will move there, and the population will increase.

Figure 5 shows the population-weighted paths of aggregate welfare over time.[11] One can see that welfare is in general increasing as a result of technological progress, agglomeration economies, and the slow movement of people towards more productive locations. In the left-hand panel, I compare the aggregate welfare in the observed history with that from the no-Zwin counterfactual. In general welfare under both scenarios is similar with the exception of significantly lower welfare in 1500 compared to the no-Zwin counterfactual, but significantly higher in 1700. Again, this result highlights that although the spatial equilibrium does not change, aggregate welfare hundreds of years after the Zwin became impassable remains significantly higher. In 1800, three to four hundred years after the Zwin became obsolete, I still find that aggregate welfare is 2% higher as compared to the counterfactual world where the Zwin never existed at all.

In the right-hand panel of figure 5, I plot the path of aggregate (population-weighted) welfare for each of a set of alternative counterfactuals. In each of these counterfactuals, I vary the location of the Zwin shock. For each of the 29 cities in my sample I change their fundamental productivity path — adding to it the positive and subsequently negative estimated impact the Zwin had on Bruges. For all counterfactuals, I keep Antwerp's fundamental productivity path fixed at that which it would have experienced in the no-Zwin counterfactual. One should not think of this exercise as being one where a random river is

[11] Welfare is weighted by the counterfactual-specific population distribution.

constructed, especially considering many locations are quite far inland. Rather, it should be thought of more generally as a productivity shock occurring in a location.

For the same shock magnitude the right-hand panel of figure 5, shows considerable variation in resulting aggregate welfare across the whole Low Countries. This is again in spite of the fact that each counterfactual economy tends towards the same long-run spatial equilibrium. Additionally, this figure makes it clear that Bruges was not the best place for such a shock as the Zwin to happen. The aggregate welfare path under the Bruges as opposed to other location counterfactuals is amongst the lowest in the distribution. In particular, if Amsterdam had enjoyed a similarly positive shock akin to the Zwin-shock in Bruges, aggregate welfare across the low countries would have been significantly higher in 1800 relative to if the shock had occurred in any other location. This exercise highlights how susceptible long-run welfare is to past historical shocks.

Figure 5 Aggregate Welfare

Notes: These figures plot population-weighted aggregate welfare across time for various counterfactuals. In the left-hand panel in orange, I plot actual historical welfare and in dotted blue welfare under the counterfactual of the Zwin never forming. In the right-hand panel, I plot the population-weighted aggregate welfare under 29 counterfactuals. In each counterfactual, a different city is shocked to the same level as that which I estimate the Zwin impacted Bruges. The thick orange line in the right-hand panel shows Burges and therefore conforms to the orange "history" line in the right-hand panel.

In figure 12 in the appendix I focus on the heterogeneity highlighted above by considering changes in welfare under each scenario relative to changes under the no-Zwin baseline counterfactual. This figure highlights that on average locations that caused the largest aggregate increases in welfare in 1400 then caused the largest decreases in welfare in 1500 once the positive shock dissipated. Intuitively, areas that caused larger gains in 1300 and 1400 also caused a greater re-shuffling of the population, then in 1500 once the positive shock was removed, the population distribution if further from the optimum causing greater welfare losses. These losses then slowly dissipate over time. This mechanism also explains why almost all counterfactuals result in lower welfare relative to the no-Zwin counterfactual — temporary shocks cause inefficient allocations after the shock occurs which take time to undo.

In figure 6 I analyze what local characteristics empirically explain the observed heterogeneity. Locations with higher initial (1200) productivity, market access, and population cause larger initial increases in welfare and subsequently larger reductions. The first part of this story speaks to some multiplier effect in targeting "better" locations. Intuitively higher market access implies more people can move to the location to enjoy the higher productivity. Higher initial productivity means those that do move get more benefit from doing so, and similarly higher population implies more benefit through agglomeration economies. Targeting shocks at better-positioned locations can thus have larger positive effects, but if the shock turns out to be temporary this can backfire in the longer run. The higher aggregate welfare climbs when the shock is active, the higher it subsequently falls once it has been removed.

Figure 6 Explaining heterogeneity in Welfare effects

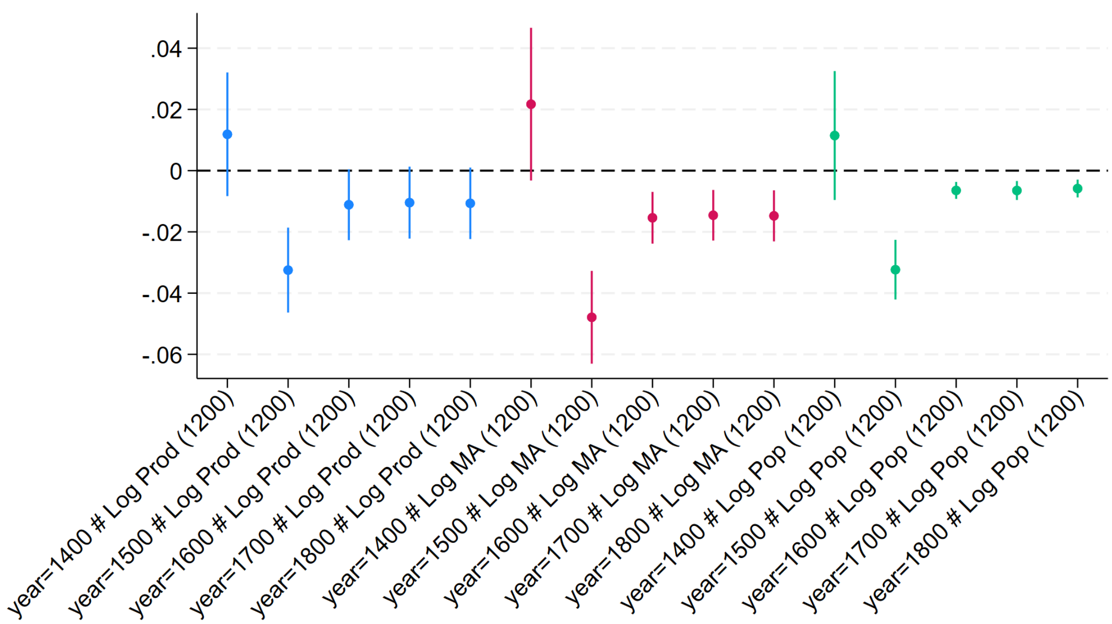

Notes: This figure plots the coefficients and corresponding confidence intervals of three regressions corresponding to each color. In each regression, I regress (log) the aggregate welfare change of each of the 29 city-specific counterfactuals relative to the no-zwin counterfactual against an explanatory variable at the shock-city level interacted by year dummies. I consider three explanatory variables capturing time-invariant characteristics of the shocked cities: log 1200 productivity, log 1200 market access, and log 1200 population. Standard errors are robust.

6 Conclusion

I show, using reduced form evidence and a dynamic quantitative spatial economics model, that a sudden river in the Low Countries, had a large effect on the population distribution of the Low Countries. The Zwin, although impassable by 1500 still impacted the population and welfare of cities across the low countries in 1800.

Despite these large impacts, I don't find evidence that this temporary shock fundamentally altered the long-run spatial equilibrium in the Low Countries. Through counterfactual exercises, I show that if the shock had been greater, or taken place elsewhere, it would also have been unlikely to change the long-run spatial equilibrium. However, if the Zwin had never silted up and was still navigable today, the urban system and spatial distribution of economic activity across Belgium and the Netherlands would look substantially different. This analysis shows that large and sustained shocks are needed to shift spatial equilibrium which is otherwise persistent. However, temporary shocks can have a large impact on the distribution of economic activity for centuries, even if eventually the system reverts.

I additionally find evidence of a "the higher you climb the harder you fall" effect. When shocks are temporary, larger positive contemporaneous effects can result in equally larger negative effects which may persist longer. This is due to the greater extent of misallocation resulting from the initial larger impact. This paper suggests that, at least in this context, moving the long-run spatial equilibrium is hard unless shocks are permanent in nature. Policymakers can still achieve long-run impacts through temporary shocks, but it's crucial that potential negative longer-run impacts are considered.

Appendix

A Dynamic quantitative spatial economics model equilibrium

For any initial population vector $\{L_{i0}\}$ and vectors of geographic fundementals $\{\bar{A}_{it}, \bar{u}_{it}, \tau_{ijt}, \mu_{ijt}\}$ such that for all i, t the following holds.

1. A locations income equals the value of purchases from it: $w_{it}L_{it} = \sum_j X_{ijt}$. Which implies that $w_{it}^{\sigma}L_{it}^{1-\alpha_1(\sigma-1)} = \sum_j K_{ijt}L_{jt}^{\beta_1(\sigma-1)}W_{jt}^{1-\sigma}w_{jt}^{\sigma}L_{jt}$. Where all exogenous and predetermined variables have been bundled together into the Kernel:
$K_{ijt} = (\tau_{ijt}(\bar{A}_{it}L_{it-1}^{\alpha_2}\bar{u}_{jt}L_{jt-1}^{\beta_2})^{-1})^{1-\sigma}$.

2. Trade is balanced. Income is fully spent $w_{it}L_{it} = \sum_j X_{jit}$. Which implies that: $w_{it}^{1-\sigma}L_{it}^{\beta_1(1-\sigma)}w_{it}^{\sigma-1} = \sum_j K_{ijt}L_{jt}^{\alpha_1(\sigma-1)}w_{jt}^{1-\sigma}$.

3. Total population equals the sum of those arriving. $L_{it} = \sum_j L_{ijt}$. This implies that: $L_{it}V_{it}^{-\theta} = \sum_j \mu_{ijt}^{-\theta}\Pi_{jt}^{-\theta}L_{jt-1}$.

4. Total population in the previous period equals the sum of those leaving. $L_{it-1} = \sum_j Lijt$. Which implies that: $\Pi_{it}^{\theta} = \sum_j \mu_{ijt}^{-\theta}W_{jt}^{\theta}$.

We can simplify this system by imposing symmetry in trade costs which implies that 1. and 2. can be combined. Thus we are left with a 3-equation model in each i, t with three unknowns $\{L_{it}, W_{it}, \Pi_{it}\}$, and unknown parameters $\{\alpha_1, \alpha_2, \beta_1, \beta_2, \sigma, \theta\}$.

B Additional tables and figures

Figure 8 in the appendix plots the difference between actual and no-Zwin counterfactual populations measured in percentage of the year-specific total Low Countries population for each city I study. Positive numbers mean populations are higher historically than in the no-Zwin counterfactual. In this figure, I omit Burges and Antwerp as they are significant outliers. Three things stand out from this figure. First, although the impacts appear small, a population increase of 0.5% of the total Low Countries Urban population in 1800 corresponds to 5,000 or so more people on a cross-city mean of 32,000 and so this is significant. Second, most cities appear to be converging to their historical populations over time with effects becoming more muted. Third, Amsterdam is an obvious exception to this rule.

Figure 7 Counterfactual productivity fundamentals

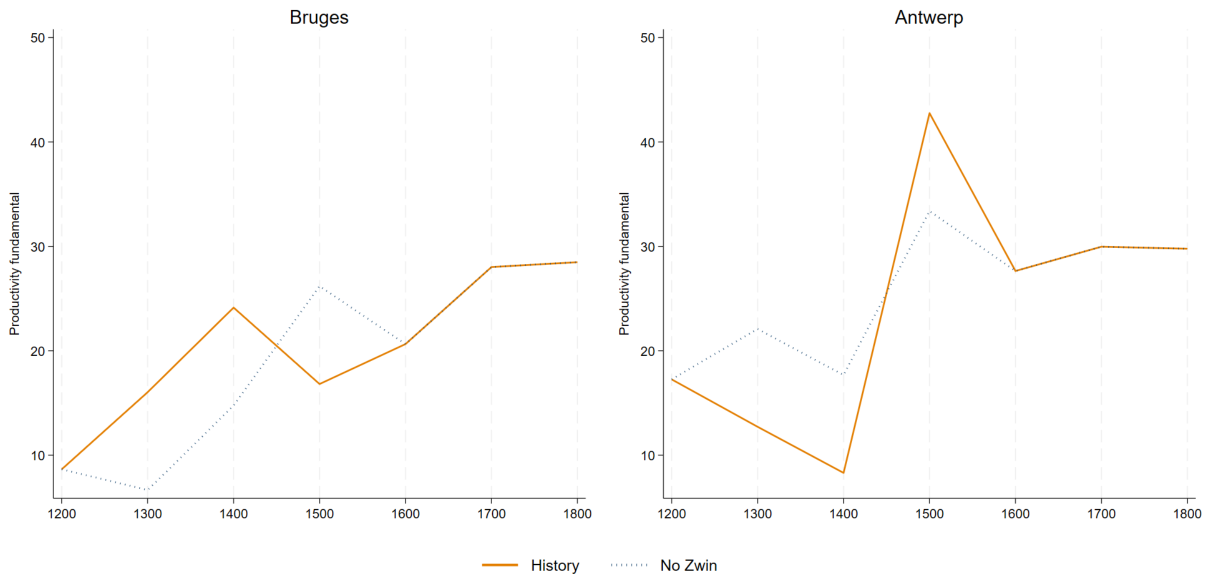

Notes: This figure shows the path of the rationalizing productivity fundamentals in orange and the counterfactual no-Zwin path in dotted navy. The difference between these two lines is exactly the estimated Zwin shock.

Figure 8 Difference in population between history and the no-Zwin counterfactual

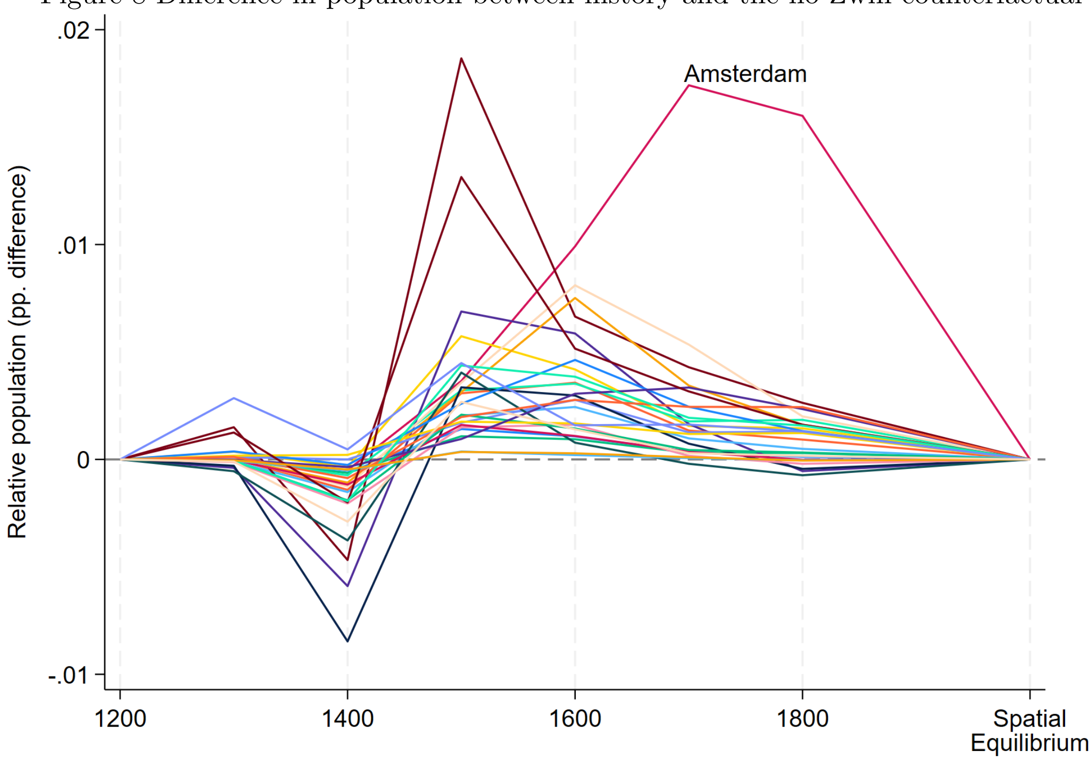

Notes: In this figure I plot the population difference (as a fraction of total population) between history and the no-Zwin counterfactual for each city in the Low Countries with the exception of Bruges and Antwerp.

Figure 9 Counterfactual paths for fundamental productivity

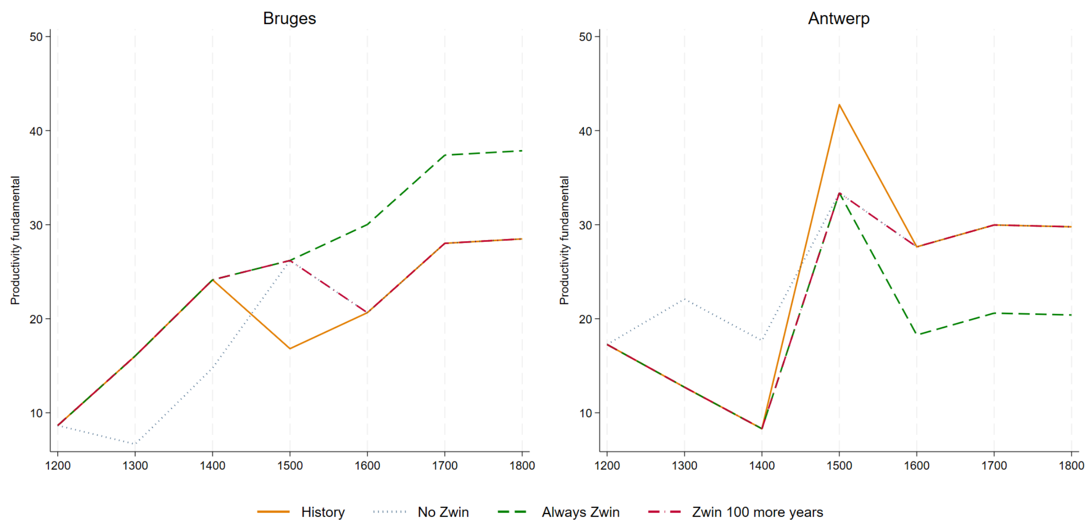

Notes: This figure shows the recovered counterfactual paths of productivity for Bruges and Antwerp under each of the following scenarios. First in a solid line, actual histroy. Second, in a dotted line history without the Zwin. Third, in a dashed line, history if the positive (and negative) effects of the Zwin never left. And finally in a dot-dashed line history if the Zwin had remained navigable for another 100 years.

Figure 10 Counterfactual welfare

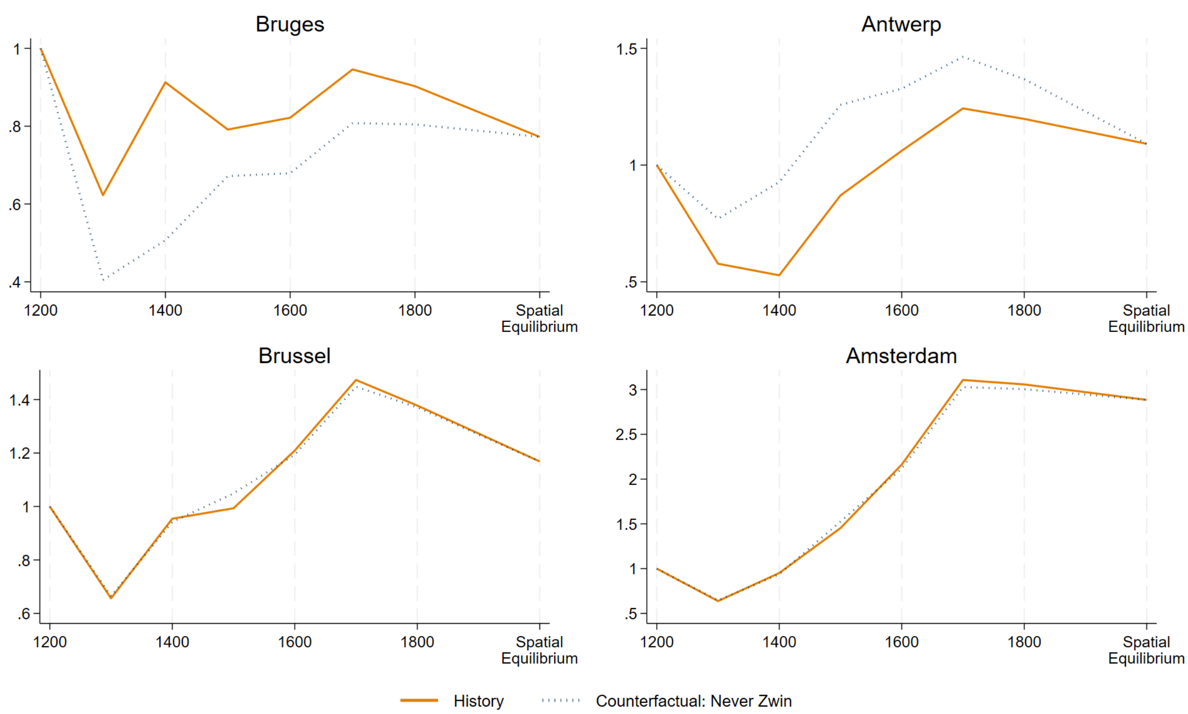

Notes: In this figure I show the counterfactual path of city-specific welfare for four cities in my sample Bruges, Antwerp, Brussels, and Amsterdam. In orange, I plot welfare as actually observed. In dotted blue, I plot welfare under the never-Zwin counterfactual world where the Zwin never appeared.

Figure 11 Counterfactual welfare

Notes: In this figure I plot the path of estimated city-specific welfare for four cities in my sample Bruges, Antwerp, Brussels, and Amsterdam. In solid orange, I plot the welfare actually observed. In dotted blue, I plot welfare under the never-Zwin counterfactual. In dashed green, I plot welfare under the counterfactual of the Zwin maintaining navigability for 100 more years. In dot-dashed red, I plot the counterfactual welfare if the Zwin had never become impassable.

Figure 12 Relative aggregate welfare for different location counterfactuals

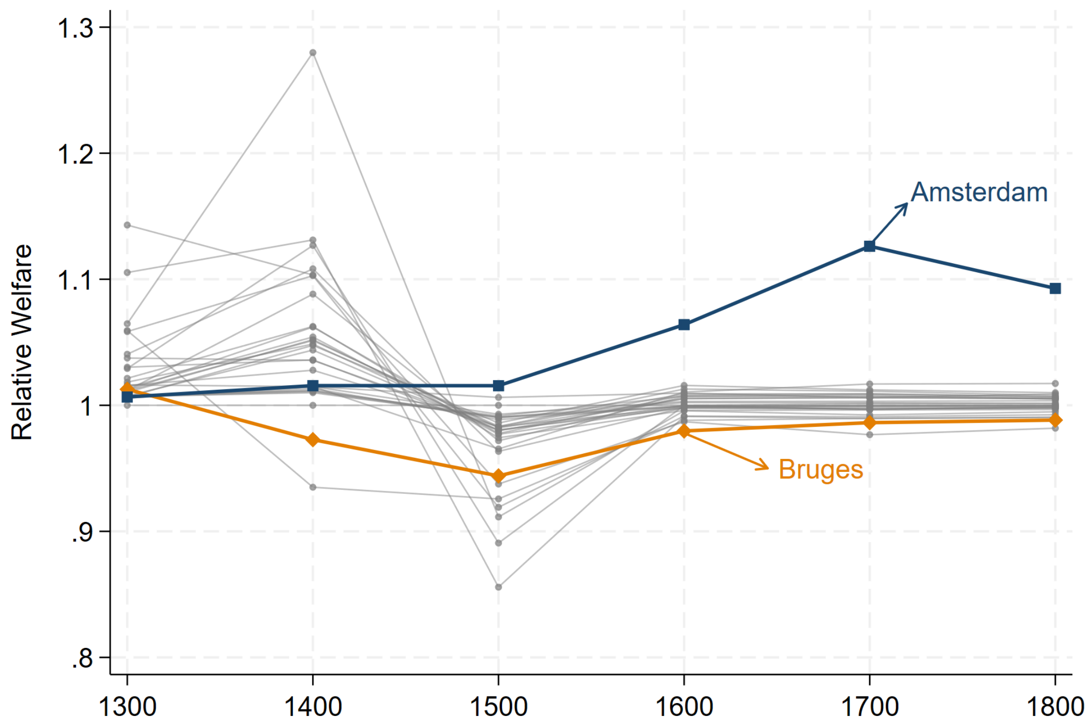

Notes: In this figure I plot the relative population-weighted aggregate welfare paths for the entire economy in each of 29 counterfactuals relative to the no-Zwin baseline counterfactual. Each counterfactual is represented as a line and corresponds to the positive Zwin shock occurring in that city. Amsterdam and Bruges are highlighted in orange and navy respectively.

C Synthetic control results

Synthetic control results. Population, labor market access, trade market access, latitude, and longitude are used as matching variables. I also match on 1800 population, as the theoretical results suggest little change in the long-run population. To reflect uncertainty around population estimates and thus to avoid overfitting on population I randomly perturb given populations by 10%. For Bruges, Liege gets 90.7% weighting and Gent 9.3%, for Antwerp the weightings are: Alost 6.3%, Amsterdam 7.2%, Brussels 50.8%, Dordrecht 16.2%, Mechelen 9.4%, and Middleburg 10.1%. The results are shown in figures 13 and 14 respectively. Figure 15 then shows

Figure 13 Synthetic control results: Bruges

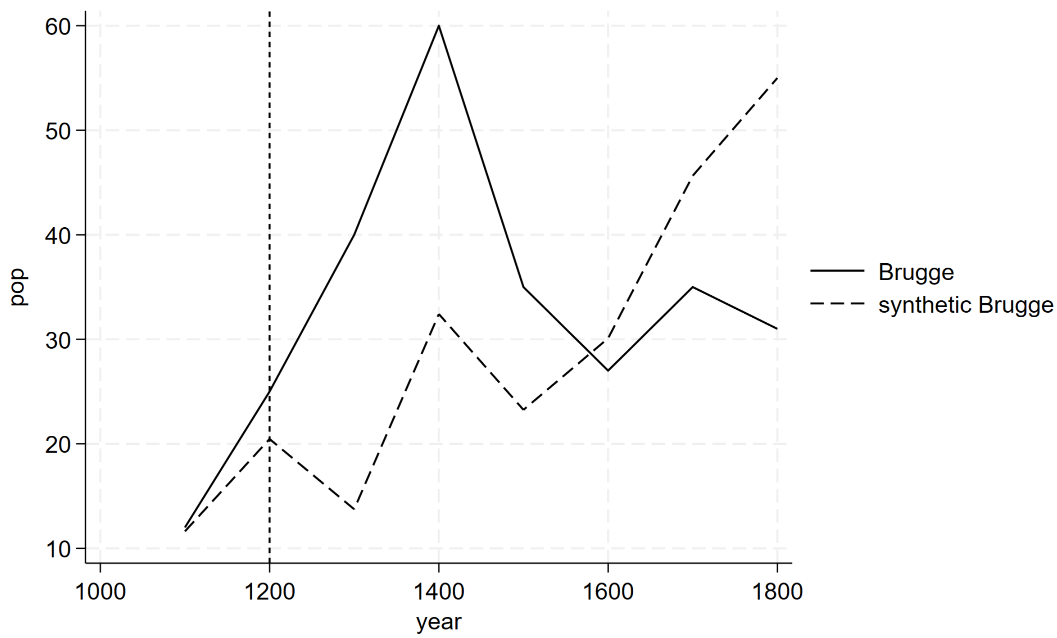

Notes: This figure shows the path of relative population for Bruges and its estimated synthetic control.

Figure 14 Synthetic control results: Antwerp

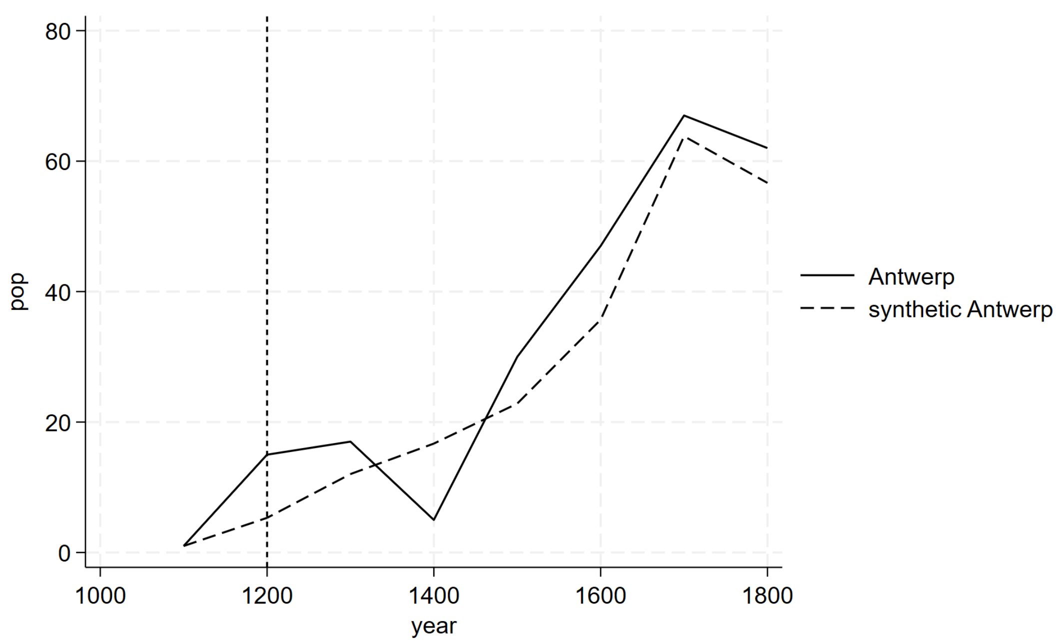

Notes: This figure shows the path of relative population for Antwerp and its estimated synthetic control.

Figure 15 Synthetic control results: Inference

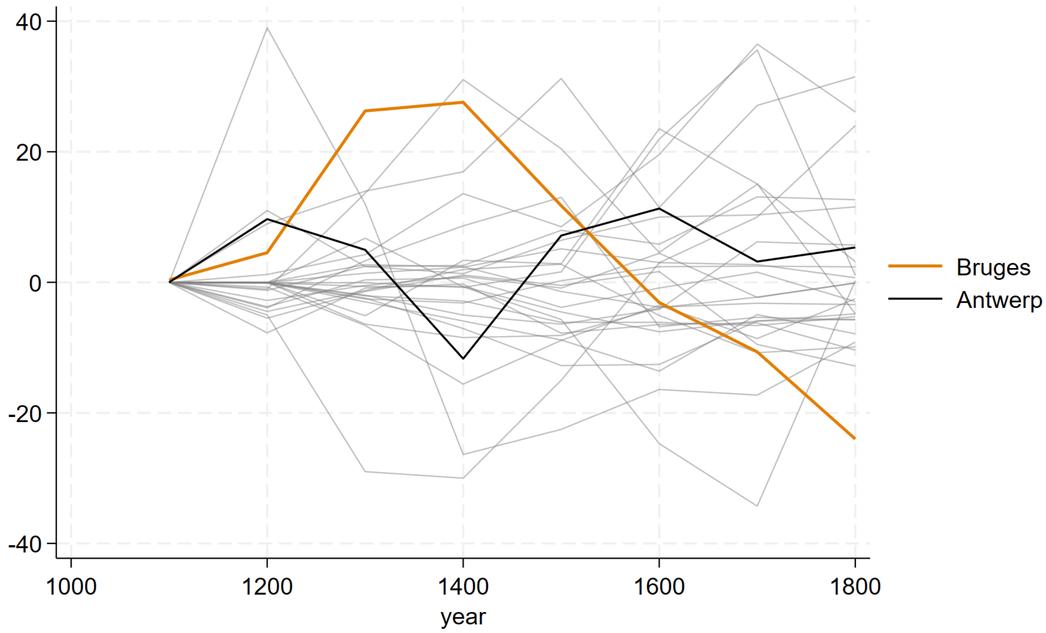

Notes: The figure shows the difference in relative populations between each city in the Low Countries and their estimated synthetic control. Bruges and Antwerp are highlighted in orange and black respectively.

References

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, 2005. 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.

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 t̀o 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.

Kirill Borusyak and Peter Hull. Non-random exposure to exogenous shocks: Theory and applications. Technical report, National Bureau of Economic Research, 2020.

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.

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.

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

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, 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.

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.

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.

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.

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.

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.

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.