

NBER WORKING PAPER SERIES

TRADE, MERCHANTS, AND THE LOST CITIES OF THE BRONZE AGE

Gojko Barjamovic
Thomas Chaney
Kerem A. Coşar
Ali Hortaçsu

Working Paper 23992
<http://www.nber.org/papers/w23992>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2017

This research is supported by the University of Chicago Neubauer Collegium for Culture and Society. Thomas Chaney acknowledges ERC grant N°337272–FiNet for financial support. Daniel Ehrlich, Simon Fuchs and Joonhwi Joo provided excellent research assistance. We are grateful to the Old Assyrian Text Project and its members for sharing the raw data (much of it still unpublished) behind this work. We thank Fikri Kulako lu for permission to use the photo of Kt 83/k 117. We thank Adam Anderson, Thomas Hertel, Michele Massa, Alessio Palmisano and Edward Stratford for valuable discussions and for sharing their research data, and to Dave Donaldson, Walker Hanlon, Sam Kortum, David Schloen and the participants at the Neubauer Collegium workshops, the Chicago Fed, the CEPR ERWIT 2016 conference, Harvard, Zurich, Brown, MIT, Yale, Georgia Tech and Georgetown, Hitotsubashi, CESifo, Dartmouth, Berkeley, the Princeton 2017 IES Summer conference, and the NBER Summer Institute 2017, for comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Gojko Barjamovic, Thomas Chaney, Kerem A. Coşar, and Ali Hortaçsu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Trade, Merchants, and the Lost Cities of the Bronze Age
Gojko Barjamovic, Thomas Chaney, Kerem A. Coşar, and Ali Hortaçsu
NBER Working Paper No. 23992
November 2017
JEL No. N15,N7,N75,R12

ABSTRACT

We analyze a large dataset of commercial records produced by Assyrian merchants in the 19th Century BCE. Using the information collected from these records, we estimate a structural gravity model of long-distance trade in the Bronze Age. We use our structural gravity model to locate lost ancient cities. In many instances, our structural estimates confirm the conjectures of historians who follow different methodologies. In some instances, our estimates confirm one conjecture against others. Confronting our structural estimates for ancient city sizes to modern data on population, income, and regional trade, we document persistent patterns in the distribution of city sizes across four millennia, even after controlling for time-invariant geographic attributes such as agricultural suitability. Finally, we offer evidence in support of the hypothesis that large cities tend to emerge at the intersections of natural transport routes, as dictated by topography.

Gojko Barjamovic
Harvard University
barjamovic@fas.harvard.edu

Kerem A. Coşar
University of Virginia
kerem.cosar@gmail.com

Thomas Chaney
Sciences Po
thomas.chaney@gmail.com

Ali Hortaçsu
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
hortacsu@uchicago.edu

This paper analyzes a large collection of commercial records from the earliest well-documented long-distance trade in world history: the Old Assyrian trade network connecting northern Iraq, Northern Syria and central Turkey during the Middle Bronze Age period (c. 2000-1650 BCE). The clay tablets on which the merchants recorded their shipment consignments, expenses, and contracts—excavated, translated and published by researchers for more than a century—paint a rich picture of an intra-regional exchange economy (Larsen, 2015).

Originating from the city of *Aššur* on the West bank of the River Tigris, some 100 km south of the modern-day Iraqi city of Mosul, several hundred Assyrian merchants settled in *Kaneš* (Kaneš) on a permanent or temporary basis. They maintained smaller expatriate trading settlements in a few dozen urban settlements on the central Anatolian Plateau and in the Trans-Taurus. *Kaneš* was the regional hub of the overland commodity trade involving the import of luxury fabrics and tin from *Aššur* to Anatolia (tin was originally sourced from Central Asia) in exchange for silver and gold bouillon (Barjamovic, in press). Assyrian merchants were also involved in a voluminous trade of copper and wool within Anatolia itself (Dercksen, 1996; Lassen, 2010).

The Assyrian texts depict a flourishing market economy, based on free enterprise and private initiative, profit-seeking and risk-taking merchants, backed by elaborate financial contracts and a well-functioning judicial system (Hertel, 2013). *Aššur* offered reliable legal procedures, a transparent system of taxation, and foreign policy that protected the Assyrian caravans and local investors involved in financing the risky long-distance trade. Assyrian merchants established trading colonies or “ports” among the small city-states of Anatolia. They negotiated with local Anatolian rulers, kings or ruling couples, the right to establish permanent trading settlements and maintain their own legal and financial institutions independent from the local community. Local rulers guaranteed the protection of passing merchant caravans against robbers and brigandage, and maintained roads and bridges, in exchange for tolls and taxes on transit trade.

Our first contribution is to extract systematic information on commercial linkages between cities from ancient texts. To do so, we leverage the fact that the ancient records we study can be transcribed into the Latin alphabet, allowing all texts to be digitized and parsed. We automatically isolate, across all records, the tablets which jointly mention at least two cities. We then systematically read those texts, which requires an intimate knowledge of the cuneiform script and Old Assyrian dialect of the ancient Akkadian language that the records are written in. Taking individual source context into account, this analysis relies exclusively upon a subset of records that explicitly refer to journeys between cities and distinguishes whether the specific journey was undertaken for the purpose of moving cargo, return journeys, or journeys undertaken for other reasons

(legal, private, etc.).

Our second contribution is to estimate a structural gravity model of ancient trade. We build a simple Ricardian model of trade. Further imposing that bilateral trade frictions can be summarized by a power function of geographic distance, our model makes predictions on the number of transactions between city pairs, which is observed in our data. The model can be estimated solely on bilateral trade flows and on the geographic location of at least some cities.

Our third contribution is to use our structural gravity model to estimate the geographic location of lost cities. While some cities in which the Assyrian merchants traded have been located and excavated by historians and archaeologists, other cities mentioned in the records can not be definitively associated with a place on the map and are now lost to us. Analyzing the records for descriptions of trade and routes connecting the cities and the landscapes surrounding them, historians have developed qualitative conjectures about potential locations of several of these lost cities. We propose an alternative, quantitative method based on maximizing the fit of the gravity equation. As long as we have data on trade between known and lost cities, with sufficiently many known compared to lost cities, a structural gravity model is able to estimate the likely geographic coordinates of lost cities. Our framework not only provides point estimates for the location of lost cities, but also confidence regions around those point estimates. For a majority of the lost cities, our quantitative estimates come remarkably close to the qualitative conjectures produced by historians, corroborating both such historical models and our purely quantitative method. Moreover, in some cases where historians disagree on the likely location of a lost city, our quantitative method supports the conjecture of some historians and rejects that of others.

Our fourth contribution is to test for the persistence of economic forces over a long horizon. Aside from allowing us to recover the location of lost cities, our gravity model yields a structural estimate for the fundamental economic size of ancient cities, when no reliable data on production and consumption, or even population size or density in the 19th century BCE survives. Instead, we structurally estimate fundamental city sizes in a general equilibrium trade model. Estimated ancient city sizes are strongly correlated with the economic size of those cities in the current era (based on trade between Turkish cities in 2014 CE). We argue that natural transportation networks shaped by the topography of the wider region—a factor usually overlooked by economists, but recognized by historians (Ramsay, 1890)—is critical in explaining the hierarchy of ancient cities and their modern counterparts.

Related literature. Our paper contributes to several literatures. First, we provide the earliest estimate of the gravity equation in trade, dating back to the 19th Century BCE, about four millennia earlier than existing estimates from the mid-19th century CE (Disdier and Head, 2008).

Second, we invert a structural gravity equation in order to locate lost cities, complementing qualitative approaches in history and archeology with a quantitative method rooted in economic theory. Our approach is loosely related to multidimensional scaling problems in other fields, where one searches for (unknown) coordinates of points such that the distances between those points are close to known distances. Multi-dimensional scaling has been applied for instance to locate eight parishes in Oxfordshire using data on marriages circa 1600-1850 CE (Kendall, 1971), and to match known archaeological sites to place names in Norway using night watchmen itineraries in the 13th century CE (Galloway, 1978). An earlier contribution (Tobler and Wineburg, 1971) uses a similar dataset as ours to locate Assyrian cities in Bronze Age Anatolia. Our method differs from and improves upon multidimensional scaling in that we use an explicit structural economic model. This allows us to infer not only the location of lost cities, but also the distance elasticity of trade, the size of cities (a theory-guided counterfactual measure), formal estimates of standard errors, and confidence intervals. Furthermore, compared to Tobler and Wineburg (1971), we use a much larger dataset that has become available for study in the meantime, systematically clean our data to identify meaningful economic exchanges, formally account for trade zeros, and confront our estimates to historical and contemporaneous evidence.

Finally, we provide novel evidence on the (very) long run determinants of the city size distribution. An important line of theoretical and empirical inquiry in economic geography involves attempts at explaining the distribution of economic and demographic size of cities over time. Locational fundamentals as dictated by geography is potentially an important factor (Davis and Weinstein, 2002). Agglomeration of economic activity for non-geographic reasons may magnify size differentials even across seemingly homogenous locations (Krugman, 1991). Path-dependence through lock-in effects could lead to the persistence of past factors—related to the fundamentals that may have been important once (Bleakley and Lin, 2012; Michaels and Rauch, 2016). Our results and historical setting suggest that centrality in transport routes dictated by topography may be an important geographic factor explaining the persistence of cities’ long-run economic fortunes.

The remainder of the paper is organized as follows. Section 1 describes our data. Section 2 derives our model and our estimation strategy. Section 3 discusses our estimates for the distance elasticity of trade, and the location of lost cities. Section 4 presents our estimates for city sizes, and explores the long-run determinants of the distribution of city sizes.

1 Ancient Trade Data

Our data comes from a collection of around 12,000 texts that constitute the hitherto deciphered and edited part of around 23,500 texts excavated primarily at the archaeological site of Kültepe, ancient *Kaneš*, located in Turkey’s central Anatolian province of Kayseri. These texts were inscribed on clay tablets in the Old Assyrian dialect of the Akkadian language in cuneiform script by ancient Assyrian merchants, their families and business partners.¹ The texts date back to a period between 1930 and 1775 BCE, with around 90% of the sample belonging to just one generation of traders, c. 1895 - 1865 BCE (Barjamovic et al., 2012).

Most texts under consideration are commercial: business letters, shipment documents, accounting records, seals and contracts. Fittingly, the tablets they were inscribed on were found in merchants’ houses and their archive rooms. In a typical shipment document or expense account, a merchant would inform partners about the cargo and related expenses:

(I paid) 6.5 shekels (of tin) from the Town of the Kanishites to Timelkiya. I paid 2 shekels of silver and 2 shekels of tin for the hire of a donkey from Timelkiya to Hurama. From Hurama to Kaneš I paid 4.5 shekels of silver and 4.5 shekels of tin for the hire of a donkey and a packer. AKT 8, 151 (lines 5-17)

Occasional business letters contain information about market and transport conditions:

Since there is a transporter and the roads are dangerous, I have not led the shipment to Hutka. When the road is free and the first caravan arrived safely here, I will send Hutka with silver. POAT 28 (lines 3-7)

Concerning the purchase of Akkadian textiles which you have written about, since you left the Akkadians have not entered the City; their land is in revolt, but should they arrive before winter, and if it is possible to make purchases profitable for you, we shall buy some for you. VS 26, 17 (lines 4-11)

While the actual cuneiform tablets are scattered all around the world in collections and museums, many of the texts have been transliterated into Latin alphabet, translated into modern language, published in various volumes, and recently digitized by assyriologists. In this draft, we use qualitative and quantitative information about cities and merchants mentioned in a sample of

¹Figure 1 shows a picture of a well preserved clay tablet.

9,728 digitized texts available to us and approximately 3,000 additional non-digitized texts.²

The version of the data we use, tabulated by Barjamovic (2011), mentions 39 unique settlements of varying size, ‘cities’ for short.³ Our measure of bilateral commercial interactions between cities is a count of all mentions of cargo shipments or individual merchants traveling from i to j ,

$$N_{ij}^{data} \equiv \text{number of mentions of travels from } i \text{ to } j.$$

We rarely have a description of the content of the shipments. So we are currently unable to identify the intensive margin of trade, the value of the wares being transported. N_{ij}^{data} measures instead the extensive margin of trade, a count of the number of shipments.

To construct this measure, we proceed in several steps. First, we automatically parse through all our 12,000 texts to identify any tablet which mentions at least two cities. We systematically isolate strings of characters corresponding to all possible spellings of city names.⁴⁵ We find 2,806 unique tablets containing at least two city names from this step.

Second, we systematically read all those 2,806 tablets, and identify all mentions of cargo shipments or individual travels. 198 unique tablets contain such mentions of itineraries. A typical business document will describe one or several itineraries of cargo shipments. The following is an excerpt from a memorandum on travel expenses describing cargo trips:

From Durhumit until Kaneš I incurred expenses of 5 minas of refined (copper), I spent 3 minas of copper until Wahšušana, I acquired and spent small wares for a value of 4 shekels of silver. AKT 8, 145 (lines 24-29)

From this sentence, we identify three shipments: from *Durhumit* to *Kaneš*, from *Kaneš* to *Wahšušana* and from *Durhumit* to *Wahšušana*. Note that for itineraries of this type, $A \rightarrow B \rightarrow C$, we count three trips, $A \rightarrow B$, $A \rightarrow C$ and $B \rightarrow C$, implicitly assuming some trade is going on along the way.

²We rely on data amassed through twenty years of collaborative effort of the Old Assyrian Text Project. The Old Assyrian Text Project website gives public access to a large part of the data (unfortunately, the site <http://oatp.net/> is currently down, due to lack of funding and human resources, but the digital archive can be accessed via www.web.archive.org). We are grateful to Thomas Hertel, Ed Stratford and all the members of the Old Assyrian Text Project for providing us with the underlying data files.

³79 city names appear in tablets. We restrict our analysis to 39 only, as too little is known about the other 40. The sample is further reduced from 39 to 26, as we need sufficient data on bilateral trade to identify our model.

⁴We exclude *Aššur*, the home city of the Assyrians, from our automated search for two reasons. First, the word *Aššur*, which occurs ca. 40,000 times, is also the name of the main Assyrian deity, and occurs very often as an element of personal names (cf., for instance the name Puzur-Aššur meaning "Refuge of Aššur" in Kt 92/k 313 cited above); second, the city of *Aššur* is often referred to as simply *ālum*—"the City" (comparable in use to references to the financial district of London), which appears ca. 10,000 times. Our automated search is not able to use a letter's context to distinguish between *Aššur* as a god, as part of a personal name or as a city; or the word for "city" as being *Aššur* or another city.

⁵The number of tablets where cities i and j appear jointly gives us an alternative measure of (non-directed) bilateral trade flows, N_{ij}^{joint} , which we use for robustness checks.

We add to these shipments additional information about merchants’ travels. While not necessarily about actual shipments of goods, references to individuals traveling inform us about a broader set of economic interactions. As our data comes from merchants’ letters and documents, we infer from a large number of merchants’ travels from i to j that trade flows from i to j are large.

We isolate 227 explicit cargo or merchants’ itineraries. From those 227 itineraries, we identify 391 directed travels (itineraries that involve more than two cities generate multiple travels).

In total, 26 cities appear as either origin or destination of those shipments. Of those 26 cities, 15 are ‘known’ and 11 are ‘lost’. ‘Known’ cities are either cities for which a place name has been unambiguously associated with an archaeological site, or cities for which a strong consensus among historians exists, such that different historians agree on a likely set of locations that are very close to one another. ‘Lost’ cities on the other hand are identified in the corpus of texts, but their location remains uncertain, with no definitive answer from archaeological evidence. From the analysis of textual evidence and the topography of the region, historians have developed competing hypotheses for the potential location of some of those. We propose to use data on bilateral trades between known and lost cities and a structural gravity model to inform the search for those lost cities.

Table 1 provides summary statistics. The mean number of travels across all city pairs is 0.60. As with modern international trade data, many city pairs do not trade: of all the 650 potential export-import relationships (directional ij and ji pairs of 26 cities), only 120 have a positive flow. The average N_{ij}^{data} for these trading pairs is 3.26, with a large dispersion (s.d. 4.23).

The top panel of figure 2 plots the map of cities, including a preview of the estimated locations of lost cities. The city of *Kaneš*, marked K , is geographically central to the system of cities under study. As discussed above, it was also the operational center of Assyrian merchants in central Anatolia. Trade flows, however, do not just display a hub and spoke structure around *Kaneš*, as seen by the rich pattern of bilateral ties between cities in the bottom panel of the figure.

2 Model and Estimation

We build a simple model of trade in which merchants arbitrage price differentials between cities. While stylized, this model captures key features of trade in the Middle Bronze Age. For instance, the model can accommodate a commodity produced outside of our network of trading cities, such as tin sourced from Central Asia, transported and traded between the Assyrian ports in Anatolia, and possibly exported by other commercial networks to distant places such as the Aegean and Egypt. We characterize equilibrium quantities in our model with direct counterparts in our dataset, such

as the count of transactions instead of their value.

Model. We follow [Eaton and Kortum \(2002\)](#) closely. There are $K + L$ cities, K of them known, and L of them lost. Tradable commodities (tin, copper, wool...) are indexed by ω . Merchants arbitrage price differentials between cities, subject to bilateral transaction costs. For simplicity, we assume iceberg ad valorem transaction costs, such that delivering one unit of a good from city i to city j requires shipping $\tau_{ij} \geq 1$ units of the good, while the remaining fraction $1 - 1/\tau_{ij}$ is lost in transit. We also explicitly assume a transaction cost for within city transactions, $\tau_{jj} \geq 1$, to capture the trade of cities with their hinterlands. If a merchant observes costs c_i and c_j in cities i and j such that $\tau_{ij}c_i < \tau_{jj}c_j$, she⁶ can exploit an arbitrage opportunity: buy τ_{ij} units of the good at a cheap cost $\tau_{ij}c_i$ in i , ship them to j , and sell at a high cost $\tau_{jj}c_j$ for a profit. We assume for tractability the cost of producing one unit of any commodity ω in city i , in any period, follows a Weibull distribution,

$$\Pr[c_i(\omega) \leq c] = 1 - \exp\left(-T_i w_i^{-\theta} c^\theta\right). \quad (1)$$

The cost $c_i(\omega)$ includes the marginal cost of production, any markup or distribution cost, but also w_i , a shifter to the cost of sourcing goods from city i reflecting the cost of local immobile factors, determined in equilibrium. The distribution of costs is i.i.d across commodities and over time, and costs are independent across cities. $\theta > 0$ is an inverse measure of the dispersion of costs. $T_i > 0$ controls the efficiency of sourcing goods from i .⁷

With merchants arbitraging away cost differences between cities,⁸ the equilibrium price for commodity ω in city j is the lowest cost among all possible sources, $p_j(\omega) = \min_k \{\tau_{kj}c_k(\omega)\}$.

⁶We use the conventional ‘she’ despite there being no documented instances of Assyrian female traders.

⁷This model can also accommodate cases where good ω is not produced locally, but instead is sourced from outside our network of $K + L$ trading cities and enters our system only through the gateway city i (e.g. tin mined in Central Asia, and shipped to *Aššur*). In that case, $T_i w_i^{-\theta}$ depends both on the fundamental efficiency and cost of local producers in i , T_i^{local} and w_i^{local} , on the efficiency and cost of outside producers sending goods to i , $T_i^{outside}$ and $w_i^{outside}$, and on the cost of sourcing goods from outside, $\tau_{outside,i}$,

$$T_i w_i^{-\theta} = \tau_{ii}^{-\theta} T_i^{local} \left(w_i^{local}\right)^{-\theta} + \tau_{outside,i}^{-\theta} T_i^{outside} \left(w_i^{outside}\right)^{-\theta}.$$

See section 4.2 for the example of *Qattara*, which acts as a gateway on the route to and from *Aššur*.

⁸As the merchants we consider are mobile, constantly traveling between cities, we do not consider the problem of repatriating the proceeds from this sale explicitly. In particular, we implicitly assume repatriation of profits is costless. If repatriating profits entails a cost, the τ_{ij} term would contain both the cost of shipping goods *and* of repatriating profits. Historical evidence suggests that some of the merchants profits were invested into real estate in *Aššur*, where house prices seemingly experienced a surge during the period ([Barjamovic et al., 2012](#)). In the absence of any systematic information on how profits are accrued and spent, we do not model profits explicitly. [Eaton et al. \(2012\)](#) show that if profits are redistributed using an outside good, their model’s predictions remain as in ours.

Assuming trade balance at the city level, total spending equals the amount paid to local factors,

$$X_i = \sum_k X_{ki} = w_i L_i, \quad (2)$$

where L_i is the size of city i 's population. With N i.i.d. draws for the costs $\{c_i(\omega)\}_{i=1 \dots K+L}$, the expected number of shipments going from i to j is

$$\mathbb{E}[N_{ij}] = N \frac{T_i w_i^{-\theta} \tau_{ij}^{-\theta} X_j}{\sum_k T_k w_k^{-\theta} \tau_{kj}^{-\theta}} = N \frac{X_i X_j}{X_{total}} \left(\frac{\tau_{ij}}{\Pi_i P_j} \right)^{-\theta}, \quad (3)$$

where we manipulate the model as in [Anderson and van Wincoop \(2003\)](#) to obtain the second expression, with $\Pi_i^{-\theta} = \sum_k \left(\frac{\tau_{ik}}{P_k} \right)^{-\theta} \frac{X_k}{X_{total}}$ a measure of outward resistance, $P_j^{-\theta} = \sum_k \left(\frac{\tau_{kj}}{\Pi_k} \right)^{-\theta} \frac{X_k}{X_{total}}$ a measure of inward resistance, and $X_{total} = \sum_k X_k$. We will rely on the result that if trade frictions are symmetric, $\tau_{ij} = \tau_{ji}, \forall i \neq j$, then $\Pi_i = P_j$ and expected trade is symmetric, $\mathbb{E}[N_{ij}] = \mathbb{E}[N_{ji}]$.

We use as our measure for the fundamental size of city i the counterfactual real value of its aggregate output if it were to move to complete autarky,⁹

$$Size_i \equiv \frac{w_i^{autarky} L_i}{P_i^{autarky}} \propto L_i T_i^{1/\theta}. \quad (4)$$

This measure for city size is convenient for two reasons. First, it only depends on exogenous parameters, L_i and $T_i^{1/\theta}$. If, for instance, trade frictions or the size of other cities were to change, this measure would remain invariant. Second, this measure can easily be computed using only trade data and an assumption for the trade elasticity θ , as shown below.

Estimation. The goal of our empirical strategy is to use the structural model (3) in order to estimate the structural parameters and the geographic location of lost cities.

For cities i and j with latitude-longitude (φ_i, λ_i) and (φ_j, λ_j) , we parametrize the symmetric trade cost function as

$$\tau_{ij}^{-\theta} = \delta (Distance_{ij}(\varphi_i, \lambda_i; \varphi_j, \lambda_j))^{-\zeta}. \quad (5)$$

δ , a simple scaling factor, controls for measurement units. The function $Distance_{ij}(\varphi_i, \lambda_i; \varphi_j, \lambda_j)$ maps geo-coordinates into geographic distances, in kms.¹⁰ ζ is the distance elasticity of trade.

⁹For a derivation of (4), see [Eaton and Kortum \(2002\)](#), equation (15) on page 1756.

¹⁰For latitudes (φ) and longitude (λ) measured in degrees, we use the Euclidean distance formula,

$$Distance_{ij}(\varphi_i, \lambda_i; \varphi_j, \lambda_j) = \frac{10,000}{90} \left(\sqrt{(\varphi_j - \varphi_i)^2 + \left(\cos \left(\frac{37.9}{180} \pi \right) (\lambda_j - \lambda_i) \right)^2} \right),$$

where 37.9 degrees North is the median latitude among known Assyrian cities. For locations in these latitudes, the difference between this Euclidean formula and the more precise Haversine formula is negligible. This approximation considerably speeds up the estimation. We will also need to know internal trade frictions. Since we do not observe internal trades, we cannot estimate within city transactions costs. We instead normalize internal distances, $Distance_{ii} = 30km$, capturing the economic hinterland of a city within the reach of a day's travel by foot or donkey.

We proceed to estimate the following vector of structural parameters

$$\{\zeta, (\varphi_{K+1}, \lambda_{K+1}) \cdots (\varphi_{K+L}, \lambda_{K+L}), Size_1 \cdots Size_{K+L}\}.$$

ζ is the distance elasticity of trade. (φ_l, λ_l) are the geo-coordinates of lost city l . $Size_i \propto L_i T_i^{1/\theta}$ is our measure of the fundamental size of city i . We use the structural model (3) and our trade cost assumption (5) to form an expression for trade shares,

$$s_{ij}^{model} \equiv \mathbb{E} \left[\frac{N_{ij}}{\sum_{k \neq j} N_{kj}} \right] = \frac{T_i w_i^{-\theta} \tau_{ij}^{-\theta}}{\sum_{k \neq j} T_k w_k^{-\theta} \tau_{kj}^{-\theta}} = \frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k \neq j} \alpha_k Distance_{kj}^{-\zeta}}, \quad (6)$$

with $\alpha_i = T_i w_i^{-\theta}$. Eaton et al. (2012) present the formal assumptions allowing to go from equation (3) in levels to (6) in shares. The empirical counterpart to trade shares is

$$s_{ij}^{data} \equiv \frac{N_{ij}^{data}}{\sum_{k \neq j} N_{kj}^{data}}. \quad (7)$$

Under the identifying assumption that predicted expected trade shares (6) and observed trade shares (7) differ by a mean-zero i.i.d. error term under the true parameters, we jointly estimate the distance elasticity of trade, ζ , the geo-coordinates of lost cities, and the α_i 's, by minimizing the sum of squared differences between observed and predicted trade shares

$$\beta \equiv (\zeta; \cdots (\varphi_l, \lambda_l) \cdots ; \cdots \alpha_i \cdots) = \arg \min_{\beta \in \mathbf{B}} \sum_j \sum_{i \neq j} \left(s_{ij}^{data} - s_{ij}^{model} \right)^2, \quad (8)$$

where \mathbf{B} contains a set of constraints we impose on the location of lost cities, derived from historical evidence.¹¹ We only impose very general constraints for which there is a broad consensus among historians. None of these constraints are binding in the neighborhood of our point estimates. We use our structural model (3)-(4) to recover fundamental city sizes from parameters estimated in (8),

$$Size_i \propto L_i T_i^{1/\theta} \propto \hat{\alpha}_i^{1+1/\theta} \sum_k \widehat{Distance_{ki}}^{-\hat{\zeta}} \hat{\alpha}_k, \quad (9)$$

where we use $\theta = 4$ from Simonovska and Waugh (2014).¹² As the absolute level of sizes cannot be identified, we arbitrarily normalize $Size_{Kaneš} \equiv 100$, so city sizes are all relative to that of Kaneš.

¹¹We list this set of constraints in Appendix A.

¹²To recover $L_i T_i^{1/\theta}$ we need to know the trade elasticity parameter θ . In the absence of consistent information on differences in commodity prices between Anatolian market places, our data does not allow us to directly estimate θ . We therefore choose $\theta = 4$ from the literature. Since the parameter θ only affects the absolute scale of our estimates of city sizes, but not relative city sizes (in logs), this choice is of little consequence.

To derive $L_i T_i^{1/\theta}$ in (9), we use the structural equation (6) to get $\hat{\alpha}_i \propto T_i w_i^{-\theta}$ so $L_i T_i^{1/\theta} \propto \hat{\alpha}_i^{1/\theta} w_i L_i$. From market clearing, $w_i L_i = X_i$. From (3) and (6), $\hat{\alpha}_i \propto T_i w_i^{-\theta} = X_i / \Pi_i^{-\theta}$ so $X_i \propto \hat{\alpha}_i \Pi_i^{-\theta}$. From the definition of $\Pi_i^{-\theta}$ and symmetry, $\Pi_i^{-\theta} \propto \sum_k \tau_{ik}^{-\theta} X_k / P_k^{-\theta} = \sum_k \tau_{ik}^{-\theta} X_k / \Pi_k^{-\theta} \propto \sum_k \tau_{ik}^{-\theta} \hat{\alpha}_k$. Combining $\tau_{ik}^{-\theta} \propto Distance_{ik}^{-\zeta}$ and the above, we get the proposed formula. City sizes are only identified up to a multiplicative constant, hence the \propto sign.

Interestingly, equation (9) shows how to recover the fundamental size of a city, in a counterfactual autarky state, using only observable trade flows. In this simple gravity setting, the term $\alpha_i^{1+1/\theta}$ corresponds to an exporter fixed effect, the propensity of a city to trade after controlling for distance. The extra term $\sum_k Distance_{ki}^{-\zeta} \alpha_k$ adjusts for the endogenous response of factor prices in general equilibrium: if city i is either centrally located and/or located near large trading partners ($Distance_{ki}^{-\zeta} \alpha_k$ large for some k 's), it faces an upward pressure on the price of local fixed factors. This depresses its exports by eroding its competitiveness. In autarky, this depressing effect of trade on factor prices disappears. Equation (9) formally adjusts for this endogenous factor price response.

Robust (White) standard errors are computed analytically and account for heteroskedasticity.¹³ To gauge visually the precision of estimates for the location of lost cities, we draw maps with confidence areas around our point estimates. Those confidence areas represent the set of likely locations of those lost cities, taking into account measurement error and model misspecification. The confidence area for lost city l corresponds to a contour plot for the 2-dimensional distribution of $B = 20,000$ Monte Carlo draws from $(\zeta; \dots (\varphi_l, \lambda_l) \dots; \dots \alpha_i \dots) \sim \mathcal{N}(\hat{\beta}, \hat{\Sigma})$, where $\hat{\beta}$ and $\hat{\Sigma}$ are the estimated parameters and co-variance matrix from solving the minimization problem (8). This procedure accounts not only for the precision of the latitude and longitude of city l , but also for the co-variance of those geo-coordinates with all other estimated parameters.

We also compute a measure of the precision of our location estimates akin to a standard error,

$$precision(l) = \sqrt{\frac{1}{B} \sum_{b=1}^B (Distance_{l,b})^2}, \quad (10)$$

where $Distance_{l,b}$ is the distance between the estimated location for l and its Monte Carlo draw b . This measure of precision is expressed in kms. Heuristically, it corresponds to the radius of a circle around our point estimate, such that the true location lies inside this circle with probability 75%.

Our non-linear estimation is closely related to [Silva and Tenreyro \(2006\)](#) and to [Eaton et al. \(2012\)](#). In particular, we use information contained in trade zeros to inform our estimation. There are, however, two key differences imposed on us by the data. The first obvious difference is that, unlike with modern trade data, we do not know the location of some cities. We use instead our

¹³For the parameter vector $\beta = (\zeta; \dots (\varphi_l, \lambda_l) \dots; \dots \alpha_i \dots)$, define the error term $e_{ij}(\beta) = s_{ij}^{data} - s_{ij}^{model}(\beta)$ evaluated at β using (6) and (7). The variance-covariance matrix of the estimated parameter vector $\hat{\beta}$ is given by

$$\hat{\Sigma} = \left(\sum_j \sum_{i \neq j} \nabla e_{ij}(\hat{\beta})' \nabla e_{ij}(\hat{\beta}) \right)^{-1} \left(\sum_j \sum_{i \neq j} e_{ij}(\hat{\beta})^2 \nabla e_{ij}(\hat{\beta})' \nabla e_{ij}(\hat{\beta}) \right) \left(\sum_j \sum_{i \neq j} \nabla e_{ij}(\hat{\beta})' \nabla e_{ij}(\hat{\beta}) \right)^{-1},$$

where $\hat{\beta}$ is the solution to the minimization problem (8). We compute the standard errors for city sizes, $Size_i$, by applying the Delta method to (9), using the above covariance matrix $\hat{\Sigma}$.

model to estimate those locations. Heuristically, the distance elasticity ‘translates’ our data on bilateral trade flows into geographic distances. Our structural model also gives us guidance on how to properly control for city sizes given their observed trade. Given the distance elasticity and city sizes, a simple triangulation-type technique can recover the location of lost cities. Our estimating equation (8) essentially does all that at once, finding parameters such that the gravity model fits the data as closely as possible. It further provides estimates of standard errors and confidence areas around our point estimates. The second difference is that we do not observe aggregate production. So we define trade shares in equations (6) and (7) excluding internal trade.

The next two sections present the estimation results in detail. Section 3 focuses on the location of lost cities, and section 4 on city sizes.

3 The Lost Cities of the Bronze Age

We present our results for the distance elasticity of trade and the estimated location of lost cities, and we confront our results to historical evidence in section 3.1. To further gauge the plausibility of our estimates, we suggest a quantitative method to systematically use the qualitative information contained in our ancient texts to construct admissible regions for the lost cities in section 3.2. Finally, as a proof of concept, we fictitiously “lose” the location of some known cities, and compare the known locations to our recovered gravity estimates in section 3.3.

3.1 Using Gravity to Recover the Location of Lost Cities

Panel A of table 2 presents our estimates for the distance elasticity of trade, $\zeta = 3.825$ with a standard error of 0.582. This suggests that trade around 1880 BCE was falling off much faster with distance than is the case today, with modern elasticity estimates typically near unity (Disdier and Head, 2008). However, two observations are in order. First, our estimate for ζ using non-directional data from joint attestation of city names is substantially smaller, $\zeta = 1.970$ with a standard error of 0.162.¹⁴ Second, modern estimates of the trade elasticity for shipments transported by road are substantially larger than unity —see Cosar and Demir (2016) for an estimate around 2 based on overland transit of exports from Turkish cities.

Figures 3-6 show our point estimates and confidence regions for each lost city separately.¹⁵ Each figure depicts a map of the region containing the Anatolian cities where Assyrian merchants operate.

¹⁴Appendix F presents robustness checks where we replicate all our results using our alternative non-directional measure of trade flows from joint mentions, N_{ij}^{joint} .

¹⁵See Appendix F for the geo-coordinates of lost cities, along with robust standard errors.

A shaded area shows the set of locations ruled out by assumption as inadmissible from historical sources. A “+” sign depicts the estimated location from our structural estimation (8), surrounded by contours representing the confidence area for that city. Those confidence areas offer a visual sense of the precision of our estimates. For most cities, our estimates are very tight, in the sense that the confidence area is at most 100km wide, and often much smaller. This visual message is confirmed by the measure of the precision of our estimated locations in panel B of table 2: all measures of precision are smaller than 50km (30 miles), and less than 30km in 7 of 11 cases. This is to be compared to the average distance of 281km between known cities.

We add to those maps two other locations. The “F” sign corresponds to the site suggested by historian Massimo Forlanini (Forlanini, 2008); the “B” sign corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011). This allows us to compare our estimates, obtained by a purely quantitative method—a structural gravity estimation, to those obtained by historians from a purely qualitative method—based on ancient itineraries, topographical studies, surviving toponyms, etc. We consider this comparison to be an informal external validity test.

In four cases, *Durhumit*, *Ninassa*, *Sinahuttum*, and *Washaniya*, our gravity estimates for the location of lost cities are extremely close to the conjecture of at least one of the two historians. In two of those cases, both historians either agree on the same site (*Sinahuttum*), or their conjectures are very close to each other (*Ninassa*). In the other two cases however (*Durhumit* and *Washaniya*), the historians strongly disagree; in those two cases, our gravity estimates are closer to the proposals made by Barjamovic than Forlanini. We view these cases where our structural gravity estimates agree with at least one historian’s proposal as an endorsement that the true location of those cities is indeed at or very near those sites. Again, as we do not use the historians conjectures as input in our estimation, those converging views are unlikely to be coincidental.

In the case of *Tuhpiya*, Forlanini and Barjamovic agree on the likely location of that city, but our gravity estimate is far from theirs. However, our estimate near the modern city of Sorgun corresponds to an earlier proposal by Cornelius (1963).

Two cases offer mixed results, *Kuburnat* and *Suppiluliyia*. In both cases, Forlanini and Barjamovic disagree by about 70km, and our gravity estimate also disagrees with both conjectures by about 70km. In those cases, we cannot convincingly draw any strong inference. In the case of *Suppiluliyia* however, the confidence area is elongated along a Northwest-Southeast axis (precision of 65km), with one end touching upon the conjecture by Barjamovic. In other words, while our model rejects Forlanini’s conjecture, we cannot reject that of Barjamovic.

The cases of *Hahhum*, *Mamma*, and *Zalpa* are interesting for another reason. In all three cases,

Forlanini’s and Barjamovic’s conjectures are either the same, or near each other. Our gravity estimates however are consistently shifted towards the North and West of theirs. The gravity estimates for those three cities happen to lie in the Taurus mountain range, a rugged high altitude and snow covered area. As we do not impose our gravity estimates to be in hospitable locations, nothing prevents this from happening. Forlanini and Barjamovic on the other hand impose the realistic constraint that cities are in accessible and suitable places and draw in historical information about their location near the Euphrates River, which the gravity estimates ignore.

Finally, in the case of the lost city of *Purushaddum*, our estimate is vastly different from both historians’. It is also a city for which both historians disagree strongly. Even though we cannot make any informed recommendation as to which historian is more likely to be right, we believe that in this case our gravity estimate is unlikely to be near the true site, for two reasons. First, our point estimate and its entire confidence region seems to be incompatible with detailed descriptions in our texts of the itineraries followed by caravans (see section 3.2 below). Second, again judging from qualitative historical sources, it seems that *Purushaddum* was a peripheral city, to the West of our network of trading cities. In such a case of a peripheral city, the triangulation exercise embedded in our estimator is unlikely to give reliable results (see section 3.3 below).

To conclude, we often find a remarkable agreement between our quantitative method for locating the lost cities of the Bronze Age and the qualitative method of historians using soft information. We view our results as plausible, with the exceptions of *Hahhum*, *Mamma*, *Zalpa*, and *Purushaddum*, with the first three being systematic and explainable by reference to actual topography, which ‘shifts’ the system to a position South and East, and the latter being in a peripheral position in the Assyrian trade network. Furthermore, in the cases of *Durhumit*, *Washaniya*, and *Suppiluliya*, the gravity estimates favor the proposals offered by Barjamovic over those given by Forlanini.¹⁶

3.2 Imposing Further Constraints on Lost Cities from Merchants Itineraries

To further assess the validity of our gravity based estimates for the location of lost cities, we use the qualitative information in the tablets on detailed itineraries of merchants to define admissible regions for the location of lost cities. This methodology is a mathematical counterpart to the contextual analysis of merchant itineraries by historians (Barjamovic, 2011).

In order to construct those admissible regions, we extract from our corpus of texts systematic

¹⁶For robustness, we replicate all our results using joint attestations of city names, N_{ij}^{joint} , as a non-directional measure of trade, instead of our preferred directional measure N_{ij}^{travel} . Appendix E shows estimation maps for this alternative measure, while Appendix F lists the estimated geo-coordinates of lost cities and replicates our estimation results. Our conclusions are robust, but the results with this noisier measure of bilateral trade are less precise.

information describing the routes followed by merchants as they travel between multiple cities. A typical multi-stop itinerary, which documents travels between both known and lost cities is found in the following excerpt from tablet Kt 83/k 117. That letter, sent to the Assyrian port authorities at *Kaneš* from its emissaries at the Assyrian port in *Wahšušana* describes how missives sent from *Wahšušana* to *Purušhaddum* will travel by two different routes, presumably during a conflict, so as to ensure safe arrival:

To the Port Authorities of Kaneš from your envoys and the Port Authorities of Wahšušana. We have heard the tablets that the Station(s) in Ulama and Šalatuwar have brought us, and we have sealed them and (hereby) convey them on to you. On the day we heard the tablets, we sent two messengers by way of Ulama and two messengers by way of Šalatuwar to Purušhaddum to clear the order. We will send you the earlier message that they brought us so as to keep you informed. The Secretary Ikūn-pīya is our messenger.
Kt 83/k 117 (lines 1-24 [Günbattı \(1998\)](#))

This letter contains two itineraries: *Wahšušana* → *Ulama* → *Purušhaddum*, and *Wahšušana* → *Šalatuwar* → *Purušhaddum*. For both of these itineraries, two cities are known (*Wahšušana* and *Ulama* for the first, *Wahšušana* and *Šalatuwar* for the second¹⁷), and one is lost (*Purušhaddum*). These are two examples of the type $A \rightarrow B \rightarrow X$ where city X is lost.

Using all such mentions of multi-stop itineraries, we impose two sets of constraints on the admissible location of lost cities: a set of “short detour” constraints, and a set of “pit stop” constraints.

The “short detour” constraint assumes that when deciding which itinerary to follow, merchants do not deviate too much from a direct route. For any segment of an itinerary with three stops, involving at least one lost city X , we assume that the intermediate stop does not represent too much of a detour compared to a direct trip without the intermediate stop. Formally we impose

$$||AX|| + ||XB|| \leq (1 + \lambda) ||AB||, \quad (\text{“short detour”})$$

where $||AB||$ represents the duration, in hours, of the fastest route going from A to B .¹⁸ This constraint means that going from A to B via X does not represent more than a $\lambda\%$ detour compared to going straight from A to B . In other words, we impose that merchants follow more or less direct routes when going over a multi-stop trip.

¹⁷ *Wahšušana* and *Ulama* are “known” in the sense that Forlanini and Barjamovic agree on locations within 30km.

¹⁸ To compute this measure of distance, we collect systematic information on elevation on a fine grid. We use [Langmuir \(1984\)](#)’s formula for calculating the time it would take for a normal human being to travel between any two contiguous grid-points. We further allow maritime travel along the coast, and we prevent crossing large impassable rivers except in a few locations (fording). We then use [Dijkstra \(1959\)](#)’s algorithm to compute the optimal travel route between any two grid-points. See Appendix B for details.

The “pit stop” constraint assumes that caravans are required to make frequent stops, in order to rest, replenish supplies, feed their pack animals (donkeys were subjected to harsh treatments by their caravan leaders), and possibly do side trades. For any lost city X , we formally impose

$$\|AX\| \leq \|\text{average segment}\| + \mu \|\text{s.d. segment}\|, \quad (\text{“pit stop”})$$

where $\|\text{average segment}\|$ is the duration, in hours, of the average segment between two known cities, and $\|\text{s.d. segment}\|$ its standard deviation. This constraint means that any segment involving at least one lost city is no more than μ standard deviations longer than the average known segment.

Figure 7 depicts a graphical example of how to construct such an admissible region by combining constraints from different itineraries. In this fictitious example, we consider one lost city X , which appears in two different itineraries, $A \rightarrow X \rightarrow B$, and $C \rightarrow D \rightarrow X$. The figure also shows how raising the parameters λ and μ widens the size of the admissible region.

Those two sets of constraints, “short detour” and “pit stop”, seem reasonable. In essence, we assume that when deciding the itinerary for a caravan or a trip, merchants do not deviate too much from an optimal route. By varying how much of a departure from the most direct route we allow, changing the parameters λ and μ , we can make those constraints arbitrarily loose. Historical evidence suggests that Assyrian merchants were indeed following close to optimal routes (Palmisano, 2013; Palmisano and Altaweel, 2015; Palmisano, 2017).

We systematically collect all mentions of multi-stop itineraries from our 12,000 texts. Jointly imposing the “Short detour” and “Pit stop” constraints corresponding to any mention of a lost city, we construct admissible regions for all lost cities. Note that this a joint estimation exercise, as many itineraries involve multiple lost cities, so the constraints for one lost city have to be compatible with the constraints for other lost cities.¹⁹

We present our results in a series of maps in figures 8-11. Each map depicts the admissible region for a given city, using the above procedure to analyze information from merchants’ itineraries, for two sets of parameters: a “tight” constraint (dotted red line, $\lambda = 2.6$ and $\mu = 1.3$) and a “loose” constraint (dashed blue line, $\lambda = 4$ and $\mu = 2$). For comparison, we also show on the same map our point estimate and confidence region from estimating our gravity model (8), as well as the location proposed by historians Forlanini (2008) and Barjamovic (2011). We reach two conclusions.

The first is that constraints from merchants’ itineraries do not impose very strict constraints on the location of lost cities. Even with the parameters corresponding to a “tight” set of constraints (red dotted lines), the admissible regions are fairly wide for most cities. The reason we cannot

¹⁹Appendix C describes our method in details.

impose a stricter set of constraints is that in order to satisfy all constraints jointly, we are bound to have a relatively loose interpretation of our constraints. In practice, with stricter constraints, the admissible region for the lost cities of *Purushaddum* and *Kuburnat* would be empty sets. Having non-empty sets for those two cities forces us to have wide admissible regions for the other cities.

The second conclusion is that our gravity estimates for the location of lost cities are compatible with the qualitative information about the itineraries followed by merchants, with the exception of *Purushaddum*. In the case of *Purushaddum*, we infer that our estimate using solely trade data and the gravity model (8) is likely to be inaccurate. For all other cities, as we never use this qualitative information to guide our estimation, we see this compatibility as an encouraging sign that our estimates based on a structural gravity model are valid.

3.3 Proof of Concept: What If We Fictitiously “Lose” Some Known Cities?

To evaluate our inverse-gravity method for estimating the coordinates of lost cities, we propose a proof of concept exercise: using trade data only, we fictitiously “lose” some known cities, use our structural gravity model to recover their locations, and compare those to their true location.

We perform this exercise using two datasets. First, we use our ancient trade dataset, setting the locations of lost cities to their GMM estimates from section 3.1. Second, we use data on intra-national trade flows between Turkish provinces in 2014 CE.²⁰ For comparability, we restrict our sample to the 21 Turkish provinces, out of the total set of 81, that comprise the area within which our ancient cities are situated. We assign the location of each province to its central city.

For each dataset, we perform three separate exercises: we pick one known city/province at a time, and estimate its coordinates, as if it had been lost. For the ancient dataset, we “lose” *Kaneš* (a large, central city), *Hanaknak* (a small central city), and *Salatuwar* (a large peripheral city). For the modern dataset, we “lose” Kayseri (a large central city, located near the site of ancient *Kaneš*), Kirşehir (a small central city), and Eskişehir (a large peripheral city).

The results of our within sample predictions are presented in figure 12. For both ancient and modern data, our estimates for central cities are very near their true locations. For the ancient dataset, the distance between the true location and the estimated location is 12km for *Kaneš*, and 7.5km for *Hanaknak*. In both cases, the true location lies well within the confidence region for the estimated location. For the modern dataset, the distance between the true location and the estimated location is 35 km for Kayseri and 23 km for Kirşehir. Those distances are to be compared with an average distance of 281km for known ancient cities, and 180km for modern provinces. For

²⁰The data is publicly available at the website of the [Turkish Ministry for Science, Industry and Technology](#).

peripheral cities, our gravity estimation is less precise. The true location of ancient *Salatuwar* is 133 km away from its estimated location; the true location of modern Eskişehir is 63 km from its estimated location; and both are outside the confidence region for the estimated location.

We conclude from this proof of concept exercise that our proposed inverse gravity estimation of the location of lost cities is reliable for central cities, but less precise for peripheral cities.

4 Persistence and Determinants of City Sizes

We now turn to a systematic discussion of our estimates of ancient city sizes, and of the long-run determinants of the city size distribution. Since reliable historical or archaeological evidence on the size of the ancient cities in question does not exist, we are unable to offer a direct external validity check. Instead, we confront our estimates of ancient city sizes to measures from the modern data, and explore the geographic and topographic determinants of city sizes.

4.1 City Size Estimates

Our estimates of the fundamental size of ancient cities ($L_i T_i^{1/\theta}$) are presented in panels B and C of table 2. A few observations are in order. First, our estimates do not achieve the level of statistical significance typical of modern econometric estimations. This is to be expected given the relative sparsity of our dataset, dating back four millennia.

Second, there is a wide dispersion in estimated city sizes. This is partly due to our choice of the theory-grounded measure $L_i T_i^{1/\theta}$ for city size, instead of the exporter fixed effect in (6), α_i : the differences between the two measures are magnified by two orders of magnitude.

Finally, one outlier stands out: the point estimate for *Qattara* is orders of magnitude larger than all other size estimates. There is a simple explanation for this extreme estimate. *Qattara* is precisely located by historical sources at the modern site of Tell al-Rimah in northern Iraq, known to have been a small independent polity during the period in question. Being the first major station on the route from *Aššur* to Anatolia (or the last stop on the return journey), virtually all trade between *Aššur* and Anatolia passed through it, even though trade in *Qattara* itself was out of bounds for the Assyrian merchants, who only held a right to pass through its territory (Barjamovic, 2008). We attribute many mentions of caravans passing by on their way to and from *Aššur* as trades with *Qattara*.²¹ Our estimator, not knowing the true destination of those trades, infers a very large size for *Qattara*. This is an example of a gateway city (see footnote 7), where *Qattara* serves as the sole

²¹Remember that we are currently unable to distinguish between the word *Aššur* used for the city, versus for a deity or as element of personal names, nor cases where merchants refer to *Aššur* simply as *ālum*, “the city”.

gateway to *Aššur* and the world beyond. In other words, the inaccurate size estimate for *Qattara* is consistent both with historical evidence, and with the logic of the gravity equation.

4.2 The Distribution of City Sizes over Four Millenia

In order to confront our ancient size estimates to modern size measures, we match the locations of ancient sites with corresponding modern urban settlements. We drop *Qattara* because it is an outlier and lies in a different country, Iraq. We then project three alternative measures of modern city sizes on our ancient size estimates and a control for geographic amenities.

Our first measure of modern size, $Population_i$, measures the total urban population living within 30km of ancient city i .²² For this measure, we use the 2012 urban population of districts (LAU-1 level, which are subdivisions of NUTS-3 level Turkish provinces). Our second measure, $NightLight_i$, is the total nighttime luminosity of the area within 30km of ancient city i . In the absence of modern-day city-level income data, nighttime luminosity is a strong correlate of local incomes (Hodler and Raschky, 2014). Our third measure, $L_i T_i^{1/\theta}|_{modern}$, mimics our estimate of ancient city sizes but uses modern trade data instead. Using data on intra-national trade flows between Turkish provinces in 2014, we solve (8)-(9) to recover size estimates from a gravity model.²³

Table 3 presents the results from the estimation of various specifications of

$$\ln(Size_i|_{modern}) = a + b \ln(L_i T_i^{1/\theta}|_{ancient}) + c \ln(CropYield_i) + u_i, \quad (11)$$

where $Size_i|_{modern}$ is either population, night lights or $L_i T_i^{1/\theta}|_{modern}$, depending on specification. $CropYield_i$, a measure of agricultural suitability around i , controls for local amenities.²⁴ All regressions are estimated via weighted least squares. The weights are the total count of importing and exporting relationships in the data, i.e., the number of itineraries in which a city is mentioned.

Columns 1, 4, and 7 of table 3 show the results of simple specifications without any geographic controls, for each of the three measures of modern size. The correlation between ancient and corresponding modern city sizes is high and significant, ranging from 35% for *NightLights* to 65% for *Population* and $LT^{1/\theta}|_{modern}$. Figure 13 plots the modern population against ancient size, as

²²For lost cities, we use our own estimates from section 3.1. The results are robust to using the locations proposed by Barjamovic (2011) instead, to restricting the sample to known cities only, or to using alternative procedures for matching ancient and modern towns. See the tables in appendix F for those robustness checks.

²³Since the intra-national trade data is at the province-level, we assign ancient cities to modern provinces based on the distance between province population centroids and ancient sites. See appendix A for details.

²⁴To approximate agricultural conditions in the distant past, we use the low-input level rain-fed cereal suitability index of IIASA/FAO (2012). We average this measure within an area of 30km radius. In unreported regressions, we also experimented with other geographic controls: elevation, ruggedness, distance to the nearest river, and distance to modern mineral deposits of gold, silver and copper. None of those controls were significant, nor were their estimated impact on modern sizes robust. We therefore exclude those controls from our regression.

in column 1 of table 3, and shows the surprising persistence of the city size from the the 20th century BCE to the 21st century CE. Columns 2, 5, and 8 document that agricultural suitability, *CropYield*, is also statistically correlated with modern city sizes. However, comparing column 2 (respectively 5 and 8) to column 3 (respectively 6 and 9), once we control for ancient city sizes, the coefficient on agricultural suitability becomes insignificant, while the relationship between ancient and modern city sizes remains positive and significant.

The strong and robust correlation of city sizes over four millennia is unlikely to be a mere coincidence. While this gives us confidence that our estimates for ancient city sizes are plausible, it also begets the question of how one can explain this surprising persistence of the city size distribution? Over the course of history, sufficiently many large shocks would have reset any initial advantage of a particular site, with a series of states rising and collapsing, radical changes in institutions and political boundaries, migrations and shifts in population for the region, climate change, large earthquakes, the rise and fall of religions, etc. In the remainder of this section, we test in detail two alternative hypotheses to explain this persistence: time-invariant local amenities, versus the advantageous location as a natural trading hub conferred by the topography of the land.

4.3 Determinants of City Sizes: the Road-knot Hypothesis

To probe the determinants of the persistence of city sizes, we project our ancient size estimate ($L_i T_i^{1/\theta}$) onto geographic observables: local agricultural suitability,²⁵ and a measure of ‘global’ advantage.

Our concept of ‘global’ advantage is novel: we define, for each site, a measure of its proximity to intersections of roadways. In developing this measure, we build upon the early work of Ramsay (1890), who proposed a topographical approach to the study of the historical geography of the region. Based on his reading of early Greek and Roman authors and his own exploratory travels, Ramsay suggested that a key to understanding Anatolian urban history is the realization that only a limited number of routes cross the area as dictated by the topography. He observed that the intersections of these routes—what he called “road-knots”—predict the location of major urban centers throughout history, in spite of a number of major political and social upsets. The exact position of the settlement could vary from period to period, but would remain in the immediate vicinity as dictated by the intersection. Ramsay’s basic hypothesis, that the existence of road-knots

²⁵In unreported regressions, we experimented with alternative measures of local amenities: elevation, ruggedness, distance to the nearest river, and distance to the nearest known copper deposit documented in the Early Bronze Age, using for this measure the list of Anatolian mines known at the time, compiled in Massa (2016). As none of those measures were either significant or robust, we exclude them from our regressions.

may be causally related to the presence of major administrative and trading centers, was further elaborated and advanced by French (1993).²⁶

The location of ancient *Kaneš* is a case in point: it is located at the northwestern end of Taurus crossings connecting the central Anatolian plateau to the upper Mesopotamian plain. The main settlement in the Bronze and parts of the Iron Age was at *Kaneš* itself, but in late Hellenistic times it moved to its current location, the regional capital of Kayseri 20km to the west. Several other large ancient and corresponding modern cities, such as *Hurama*-Karahöyük/Elbistan, *Mama*-Kale/Maraş, and *Samuha*-Kayalıpınar/Sivas, are also placed on road-knots (Barjamovic, 2011).

For our first measure, *RomanRoads_i*, we locate the intersections of roads from detailed maps of the Roman transportation network in Anatolia (French, 2016), and record the number of roads radiating from each intersection (3 for a T-crossing, 4 for an X-crossing etc). The variable *RomanRoads_i* assigns the number of Roman roads intersecting at points within 20km of city *i*, which varies between two and five.²⁷ While capturing the location of cities vis-a-vis the actual historical road network, this measure has two shortcomings. First, there is about a two thousand years gap between the Middle Bronze Age and the Roman period. This concern is partially alleviated by the fact that Roman roads themselves follow older trails (Ramsay, 1890; French, 1993). The other shortcoming concerns the potential endogeneity of the road network itself: it is plausible that roads endogenously connect large cities, so that large cities “cause” roads, rather than the reverse.

Our second measure, *NaturalRoads_i*, addresses this concern by constructing natural roadways across space.²⁸ We use detailed data on the topography of the entire region surrounding Anatolia and implement Langmuir (1984)’s formula to compute travel times for a normal human being walking on a rugged terrain. We complement this formula by collecting information on impassable rivers and river crossings (fords), and allowing for maritime travel near the coast. We use Dijkstra (1959)’s algorithm to compute the optimal route between any two points. Our approach in defining natural routes finds support in Palmisano (2013), Palmisano and Altaweel (2015) and Palmisano (2017), who argue that ancient routes followed least-effort paths closely. Armed with this measure of optimal travel routes, we consider a very large number of routes between origin-destination pairs, varying the length of these routes, from very short distances (less than a hundred km), to long

²⁶A similar analysis by Cronon (2009) emphasizes Chicago’s location at the intersection point of overland and water transportation routes as a key factor in its growth. For topographical and historical determinants of city sizes, see also Bleakley and Lin (2012) for mid-Atlantic and southern U.S. cities that were once portage sites, and Michaels and Rauch (2016) for French cities originating from Roman towns.

²⁷If there is no intersection within a 20km distance, a city assumes a score of two since each city is necessarily on a road itself. 16 out of 25 cities have a road score of two. Two cities, *Kaneš* and *Ulama*, have road scores of five.

²⁸See appendices B and D for details.

distances (several thousands km), in order to account for both local and global exchanges. We record all intersections or overlaps of those routes and define a natural road score for each location.

Figure 15 shows a heat map of our natural roads measure for Turkey. Major modern urban settlements and transportation arteries, not included on this map, overlap with our road measure, although neither were used as input. Note that some maritime areas near the shore have a high road score (e.g. North of Istanbul). This happens because our fictitious traveler often finds it optimal to travel by sea near the coast, so that many maritime optimal paths intersect or overlap at sea.

Our variable $NaturalRoads_i$ is the number of intersections or overlaps of optimal paths within 20 km of city i .²⁹ In essence, it measures the propensity of a given site to be connected to a natural routes network. This measure is arguably time invariant as it uses topographical data as input only.

Table 4 presents the results from the estimation of various specifications of

$$\ln \left(L_i T_i^{1/\theta} |_{ancient} \right) = d + e \ln (CropYield_i) + f \ln (Roads_i) + v_i. \quad (12)$$

All regressions are estimated via weighted least squares. The weights are the total count of importing and exporting relationships in the data, i.e., the number of itineraries in which a city is mentioned. We control for $CropYield_i$ in columns 1, 4 and 5. $Roads_i$ corresponds either to the number of roman roads intersecting near city i , $RomanRoads_i$, in columns 2 and 4, or to the number of natural paths intersecting or overlapping near city i , $NaturalRoads_i$, in columns 3 and 5.

Table 4 presents robust and significant evidence in support of the road-knot hypothesis. Cities with a high road score, located near the intersection of many routes, tend to be systematically larger. The *RomanRoad* variable has a p -value of 0.107 in column 2, and 0.146 when controlling jointly for *CropYield*. Our a priori measure of the connectedness of a city to the natural road network, *NaturalRoad*, is on the other hand strongly significant, with a p -value below 0.01 both on it own in column 3, and when controlling for *CropYield* in column 4. Our natural road score accounts for more than half (53%) of the variation of ancient city sizes. Figure 14 presents visual evidence of this strong correlation, and shows it is not driven by outliers.

Agricultural suitability, *CropYield*, is also correlated our estimates of ancient city sizes, accounting for about a quarter of the variation in column 1 of tale 4. But this correlation becomes insignificant when we control for our natural road score. Among all measures of local amenities, *CropYield* is the only variable that is significantly correlated with city size. Elevation, ruggedness,

²⁹We use 20km as a benchmark for the natural drift of city locations over time. As explained above, 20 km is the distance between the modern city of Kayseri and ancient *Kaneš*. French (1993) describes another instance in which a modern city, Aksaray, is 18 km away from the ancient site of Acemhöyük due to local relocations of towns throughout history.

distance to the nearest river, and distance to mineral deposits known to have been exploited in the Early Bronze Age are all either insignificant, or driven by one or two outliers.

To recap, along with agricultural suitability, the proximity of a city to natural roads is a time-invariant attribute that can account for the persistence of city sizes. Topography dictates the path followed by optimal roads, today and in the past. Access to this natural road network confers an economic advantage which can explain the emergence and sustainability of large urban centers.

Conclusion

Business documents dating back to the Bronze Age—inscribed into clay tablets and unearthed from ancient sites in Anatolia—give us a window to analyze economic interactions between Assyrian merchants and Anatolian cities 4000 years ago. The data allows us to construct proxies for trade between ancient cities and estimate a structural gravity model. Two main results emerge.

First, more cities are named in ancient texts than can be located unambiguously by archaeological and historical evidence. Assyriologists develop proposals on potential sites based on qualitative evidence (Forlanini, 2008; Barjamovic, 2011). In a rare example of collaboration across disciplines, we use a theory-based quantitative method from economics to inform this quest in the field of history. The structural gravity model delivers estimates for the coordinates of the lost cities. For a majority of cases, our quantitative estimates are remarkably close to qualitative proposals made by historians. In some cases where historians disagree on the likely site of lost cities, our quantitative method supports the suggestions of some historians and rejects that of others.

Second, we analyze the correlation between the estimated economic size of ancient cities and their time-invariant geographic attributes, as well as their future economic outcomes. Despite a gap of 4000 years, we find that ancient economic size predicts the income and population of corresponding regions in present-day Turkey. We argue that the persistence of cities' fortunes in the very long run can best be explained by their strategic position in the network of natural trade routes, as proposed by Ramsay (1890). While access to mineral deposits may have played a role in the early emergence of some cities, such as the mines in the Early Bronze Age near *Kaneš* and *Durhumit* (Massa, 2016), it seems that key to the emergence and persistence of the urban network in Anatolia is the ability of cities to access natural routes, and integrate into the broader trading network. We hope to further explore this hypothesis in other historical settings and regions in future research.