Trade, Merchants, and the Lost Cities of the Bronze Age*

Gojko Barjamovic
Thomas Chaney
Kerem Coşar
Ali Hortaçsu

January 23, 2019

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. Having structurally estimated ancient city sizes, 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. We also document persistent patterns in the distribution of city sizes across four millennia, find a distance elasticity of trade in the Bronze Age close to modern estimates, and show suggestive evidence that the distribution of ancient city sizes, inferred from trade data, is well approximated by Zipf's law.

*Corresponding author: Thomas Chaney, Department of Economics, Sciences Po, 28 rue des Saints Peres, 75007 Paris, France, thomas.chaney@gmail.com. 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. Francisco Del Villar, Daniel Ehrlich, Simon Fuchs, Joonhwi Joo and Stefan Pauly provided excellent research assistance. We are grateful to the Old Assyrian Text Project and its members for sharing the data underlying this work. We thank Fikri Kulakoğlu 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 Dave Donaldson, Walker Hanlon, Sam Kortum, David Schloen and the participants at various seminars and conferences for comments and suggestions.

I Introduction

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 on a permanent or temporary basis in *Kaneš* (Kanesh) near modern-day Kayseri in Turkey. They maintained smaller expatriate trading settlements in a few dozen urban centers 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 originally sourced from Central Asia) in exchange for silver and gold bullion (Barjamovic, 2018). 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, and digitized. We parse all texts and automatically isolate all 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 identifies exclusively a subset of records that explicitly refer to trades between cities.

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 known cities. We estimate a distance elasticity of trade in the Bronze Age equal to 1.9, surprisingly close to modern estimates.

Our third contribution is to use the 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 descriptions of trade routes connecting the cities and the landscapes surrounding them, historians have developed qualitative conjectures about potential locations 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. In some cases where historians disagree on the location of a lost city, our quantitative method supports the conjecture of some historians against others.

Our fourth contribution is to propose an explanation for the size distribution of ancient cities: cities which are centrally located in the transportation network, determined solely by the topography of the wider region, tend to be large. Our general equilibrium 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. We show that natural transportation networks—a factor usually overlooked by economists, but recognized by historians (Ramsay, 1890)—is critical in explaining the hierarchy of ancient city size estimates. We also provide evidence that the city size distribution is persistent over millennia, with estimated ancient city sizes strongly correlated with the economic size of those cities in the current era. Finally, we find suggestive evidence that the distribution of population of ancient urban settlements is closely approximated by Zipf's law, much like the distribution of modern city sizes.

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, and with a distance elasticity of trade close to modern estimates (Disdier and Head, 2008; Cosar and Demir, 2016).

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 regions. 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. We also show that our structural estimates yield more plausible estimates than a naive multi-dimensional scaling approach.

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 are 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 the transportation network, shaped by the topography of the land, is an important geographic factor explaining the hierarchy of city sizes.

The remainder of the paper is organized as follows. Section II describes our data. Section III derives our model and our estimation strategy. Section IV discusses our estimates for the distance elasticity of trade, and the location of lost cities. Section V presents our estimates for city sizes, and explores the determinants of the distribution of ancient city sizes. Section VI compares our structural gravity model to estimates from a naive gravity model.

II 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. Figure I shows a picture of a well preserved clay tablet.[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, ca. 1895 - 1865 BCE (Barjamovic, Hertel, and Larsen, 2012).

Since *Kaneš* was home to the main expatriate court adjudicating on disputes within the Assyrian commercial activities in Anatolia during that time, main Assyrian merchants all maintained houses and commercial storage in the city. The merchants settled at *Kaneš* typically acted both as agents of larger trading houses in the mother city of *Aššur*, as well as partners in local trade ventures. This required them to keep records on trade endeavors throughout their commercial circuit, regardless of whether it involved *Kaneš* or not. Such records were to a large extent archived at *Kaneš* alongside dossiers of legal and commercial records coming from elsewhere within the network, and archival copies of texts going out to other cities in Anatolia. This to some degree alleviates any geographical bias of the sources and the commercial geography that they reflect.

The city of *Kaneš* experienced a major conflagration that destroyed all Assyrian merchant houses and sealed off and preserved many of the commercial archives *in situ* ca. 1840 BCE (Manning, Barjamovic, and Lorenzen, 2017). This is the main reason why the material, which represents the world's oldest consistent archive of trade data, survives to this day. Unlike papyrus, paper or parchment, clay is ubiquitous, inexpensive and preserves well in the ground, so the *Kaneš* archives survived where most other materials would have perished. The closest comparable corpora of ancient

[1] The figure shows a photograph of tablet Kt 83-k 117. The transliteration of the cuneiform script into the latin alphabet reads: *a-na kà-ri-im Kà-ne-eš*$_6$ *qí-bi-ma um-ma ší-ip-ru-ku-nu ù kà-ru-um Wa-ah-šu-ša-na-ma tup-pè-e wa-bar-tum ša Ú-lá-ma ù Ša-lá-tù-ar ú-šé-bi*$_4$*-lu-nim-ma ni-iš-ta-me-ma ni-ik-nu-uk-/ma na-áš-ú-ni-ku-nu-tí i-ša-am-ší tup-pè-e ni-iš-ta-me-ù 2 šì-ip-ri ha-ra-an Ú-là-ma-ma 2 šì-ip-ri ha-ra-an Ša-là-tù-ar-ma a-na Pu-ru-uš-ha-tim a-na a-wa-tim za-ku-im ni-iš-ta-pàr a-wa-tàm pà-ni-tàm-ma ša ù-bu-lu-ni-ni ni-ša-pà-ra-ku-nu-tì-ma ù-za-ku-nu : ni-pà-tì I-ku-pì-a DUB.SAR šì-pàr-ni.* City names have been underlined, giving an example of how an automated search for strings of characters can identify mentions of city names. The English translation of part of the text is on page 17. We thank Fikri Kulakoğlu for permission to use the photo of this tablet.

trade data are almost 3000 years later, coming e.g., from the medieval Italian merchant archives and the Cairo Genizah.

Most texts under consideration, found in merchants' houses, are commercial: business letters, shipment documents, accounting records, seals and contracts. 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.* [Tablet AKT 8/151, lines 5-17]

> *In accordance with your message about the 300 kg of copper, we hired some Kaneshites here and they will bring it to you in a wagon...Pay in all 21 shekels of silver to the Kaneshite transporters. 3 bags of copper are under your seal...Here, Puzur-Aššur spent 5 minas of copper for their food. We paid 5 2/3 minas of copper for the wagon.* [Tablet Kt 92/k 313, lines 4-8, 14-22]

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.* [Tablet 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.* [Tablet 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. We use qualitative and quantitative information about cities and merchants mentioned in a sample of 9,728 digitized texts available to us and approximately 2,000 additional non-digitized texts.[2]

[2] We rely on data amassed through twenty years of collaborative effort of the Old Assyrian Text Project. The project's website gives public access to a large part of the data (sadly, due to insufficient funding the site http://oatp.net/ is no longer active, 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.

The version of the data we use, tabulated by Barjamovic (2011), mentions 79 unique settlements, 'cities' for short. Out of these 79 cities distributed across modern-day Iraq, Syria and Turkey, we restrict our analysis to 25 Anatolian cities in Turkey (appendix B explains in detail the sample selection criteria). Our directed measure of bilateral commercial interactions between cities is a count of all mentions of cargo shipments or individual merchants traveling from i to j,

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

Since we rarely have a description of the content of the shipments, we are unable to identify the intensive margin of trade, i.e., 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. To do so, we systematically isolate strings of characters corresponding to all possible spellings of city names.[3] We find 2,806 unique tablets containing at least two city names from this step.

Second, we systematically read all those 2,806 tablets, identify all mentions of cargo shipments or individual merchants' travels, and discard coincidental mentions of cities (see appendix B for an example of a coincidental joint attestation). 198 unique tablets contain such mentions of cargo and merchants' itineraries. A typical business document will describe one or several itineraries. The following example is an excerpt from a memorandum on travel expenses describing cargo trips:

> *From <u>Durhumit</u> until <u>Kaneš</u> I incurred expenses of 5 minas of refined (copper), I spent 3 minas of copper until <u>Wahšušana</u>, I acquired and spent small wares for a value of 4 shekels of silver.* [Tablet 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, $A \to B \to C$, we count three trips, $A \to B$, $B \to C$ and $A \to C$, implicitly assuming some trade is going on along the way. In the rare cases where an itinerary loops back, we do not count the return trip. This procedure

[3] We exclude *Aššur*, the home city of the Assyrians, from our automated search for three 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. Third, in order to analyze the long-run determinants of city sizes in a consistent manner, we limit our attention to cities within the boundaries of modern-day Turkey so as to eliminate confounding institutional factors. Being situated in northern Iraq, *Aššur* does not satisfy this criterion.

isolates 227 explicit cargo or merchants' itineraries, from which we identify 391 directed travels between city pairs (itineraries with more than two cities generate multiple travels).

Of the 25 cities in our sample, 15 are 'known' and 10 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, archaeology 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 I provides summary statistics. The mean number of travels across all city pairs is 0.63. As with modern international trade data, many city pairs do not trade: of all the 600 potential export-import relationships (directed ij and ji pairs out of 25 cities), only 114 have a positive flow. The average N_{ij}^{data} for these trading pairs is 3.33, with a large dispersion (s.d. 4.31).

Figure II plots all cities on a map, including a preview of the estimated locations of lost cities, and the bilateral trade flows between them.[4] The city of *Kaneš* 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*, with rich patterns of bilateral ties between cities. This further reassures us that we are not over sampling *Kaneš*-related trades.

III Model and Estimation

Model. We adapt Eaton, Kortum, and Sotelo (2012)'s finite sample version of the Eaton and Kortum (2002) gravity model of trade to our setting. The Eaton and Kortum model is particularly well suited for two reasons. First, it is a model of arbitrage pricing which plausibly describes Assyrian merchants' trading strategy. Second, this model makes an explicit prediction about the count of shipments, which we observe, rather than the value of shipments, about which we have almost no information. When bringing the model to the data, we depart from Eaton, Kortum,

[4]In the top panel, known cities are in grey (*Hanaknak*, *Hattuš*, *Hurama*, *Kaneš*, *Karahna*, *Malitta*, *Mamma*, *Šalatuwar*, *Šamuha*, *Tapaggaš*, *Timelkiya*, *Ulama*, *Unipsum*, *Wahšušana*, *Zimišhuna*), and lost cities are in black (*Durhumit*, *Hahhum*, *Kuburnat*, *Ninašša*, *Purušhaddum*, *Šinahuttum*, *Šuppiluliya*, *Tuhpiya*, *Wašhaniya*, *Zalpa*). In the bottom panel, thin lines indicate $0 < N_{ij}^{data} \leq 3$, and thick lines $N_{ij}^{data} > 3$. We also include *Aššur*, the home city of the Assyrians, in order to give a full picture of the related geography, even though it is not included in our sample (see footnote 3). See appendix B for a frequency plot of directed shipment counts (Appendix Figure I).

and Sotelo (2012) and other modern gravity estimates such as Silva and Tenreyro (2006): unlike with modern trade data, we do not know the location of some cities. We use instead our model to estimate those locations. In other words, we treat some distances as unknowns instead of data.

There are $K+L$ cities, K of them known, and L of them lost. A finite number of 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. 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 $c_i(\omega)$ and $c_j(\omega)$ for good ω in cities i and j such that $\tau_{ij} c_i(\omega) < \tau_{jj} c_j(\omega)$, she[5] can exploit an arbitrage opportunity: buy τ_{ij} units of the good at a cheap cost $\tau_{ij} c_i(\omega)$ in i, ship those τ_{ij} units to deliver one unit in j, sell at a high price $\tau_{jj} c_j(\omega)$ for a profit, without the threat of being undercut by local sellers who could at best buy τ_{jj} units locally at cost $c_j(\omega)$ to sell at cost $\tau_{jj} c_j(\omega)$ per unit.

We assume for tractability that the local cost of one unit of any commodity ω in city i, at any time, is drawn from a Weibull distribution,

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

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 below. 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 and $T_i > 0$ controls the efficiency of sourcing goods from i.

Denote by $c_{ij}(\omega) = \tau_{ij} c_i(\omega)$ the marginal cost of delivering good ω from origin city i to destination city j. Under the assumption that merchants can freely exploit any arbitrage opportunity,[6] the probability that a shipment sourced by destination j originates from i is equal to

$$\Pr\left[c_{ij}(\omega) \leq \min_k \left\{c_{kj}(\omega)\right\}\right] = \frac{T_i \left(\tau_{ij} w_i\right)^{-\theta}}{\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}}. \tag{2}$$

[5]We use the conventional 'she'. Although there are no documented instances of female Assyrian traders, women did occasionally participate in the trade as partners.

[6]As the merchants we consider are mobile, constantly traveling between cities, we do not consider the problem of repatriating the proceeds from this sale explicitly. We implicitly assume repatriation 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 new shipments and real estate in Aššur, where house prices seemingly experienced a surge during the period (Barjamovic, Hertel, and Larsen, 2012, p.72). In the absence of any systematic information on how profits are accrued and spent, we do not model profits explicitly. Eaton, Kortum, and Sotelo (2012) show that if profits are redistributed using an outside good, the predictions remain.

We define two additional conditional probabilities (see mathematical appendix A for formal derivations): the probability that, conditional on not sourcing locally, destination j sources good ω from origin i; and the probability that conditional on not sourcing a good either locally of from a lost city, destination j sources good ω from known origin i,

$$\Pr\left[c_{ij}(\omega) \le \min_{k\neq j}\{c_{kj}(\omega)\} \middle| c_{jj}(\omega) > \min_{k\neq j}\{c_{kj}(\omega)\}\right] = \frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\neq j} T_k(\tau_{kj}w_k)^{-\theta}}, \tag{3}$$

$$\Pr\left[c_{ij}(\omega) \le \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\} \middle| \min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\}\right] = \frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\in\mathcal{K}\setminus\{j\}} T_k(\tau_{kj}w_k)^{-\theta}}, \tag{4}$$

where we denote the set of K known cities by $\mathcal{K}$, and the set of L lost cities by $\mathcal{L}$. We form conditional probability (3) because, in our dataset, unlike in modern trade data, we do not observe internal transactions, a purchase in city j of a good sourced locally in j. We also form conditional probability (4) to estimate the distance elasticity of trade using known distances only.

Equations (2), (3), and (4) will form the basis of our estimation. It is important to note that the Eaton and Kortum (2002) model makes explicit predictions about the probability of a shipment occurring, equation (2). The empirical counterpart to this probability can be formed using only data on the count of shipments, and does not require knowledge of the value of shipments. This property is crucial to us, as our dataset contains information on the count of shipments, but not on the value of shipments. Among modern trade models, this feature is unique to the Eaton and Kortum (2002) model, and one of our main motivation for using it. Note that this model also predicts that trade shares in value are equal to trade shares in counts.[7] We will use this property to close our model in general equilibrium in order to derive counterfactual measures of city sizes.

Estimation. Our estimation proceeds in three steps. First, we parametrize trade costs as a function of distance only. Using data on shipments among known cities only, we can estimate the distance elasticity of trade. Second, imposing the estimated trade cost function, we jointly estimate city sizes for all cities and the geographic location of lost cities. Estimating unknown locations for lost cities, and therefore distances involving lost cities, is novel compared to conventional estimates of the gravity equation in trade. Third, we combine our estimates to compute a measure of the

[7] In the Eaton and Kortum (2002) model, the fraction of shipments imported by j originating from i in count, $N_{ij}/\sum_k N_{kj}$, is equal to the fraction of j's spending on imports from i, $X_{ij}/\sum_k X_{kj}$, in expectation,

$$\mathbb{E}\left[\frac{N_{ij}}{\sum_k N_{kj}}\right] = \mathbb{E}\left[\frac{X_{ij}}{\sum_k X_{kj}}\right] = \frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_k T_k(\tau_{kj}w_k)^{-\theta}}.$$

This holds because the distribution of prices of goods delivered in destination j is independent of the goods' origin. See property (b) on page 1748 and equations (8) and (10) in Eaton and Kortum (2002).

fundamental size of cities, solving a full general equilibrium version of our model. Heuristically, the distance elasticity 'translates' our data on bilateral trade flows into geographic distances. A simple triangulation-type technique can then recover the location of lost cities. Our three-steps procedure formally estimates parameters such that the gravity model fits the data as closely as possible, and further provides estimates of standard errors and confidence regions around our point estimates.

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

$$\tau_{ij}^{-\theta} \;=\; \mu \cdot \left(Distance_{ij}\left(\varphi_i, \lambda_i; \varphi_j, \lambda_j\right)\right)^{-\zeta}. \tag{5}$$

A scaling factor, μ, controls for measurement units, and ζ is the distance elasticity of trade. The function $Distance_{ij}\left(\varphi_i, \lambda_i; \varphi_j, \lambda_j\right)$ maps geo-coordinates into geographic distances, in kms.[8]

We use Euclidean distances, i.e. as the crow flies distances, instead of least effort distances that would account for the topography of the local terrain. There are two reasons motivating this choice. First, when estimating our gravity model, we need to solve a complex non linear minimization program—see problem (8) below. With an explicit Euclidean formula for distance, we can take the first order conditions of this program with respect to the latitudes and longitudes of lost cities. Had we used least effort distances instead, we would have had to compute all possible least effort distances for pairs of points on a discrete grid, and solved our minimization program by brute force. This is computationally infeasible.[9] Second, we use information on the topography of the local terrain as an external validity check on our estimates (see sections IV.B and V.B below). Not bringing topographical data into our estimation gives credence to those validity checks.

Our model being based upon Eaton, Kortum, and Sotelo (2012)'s finite sample version of the Eaton and Kortum model, predicted trade shares between city pairs are random. In particular they can be zero, as often happens in the data, if the lowest realized cost to deliver a good from i to j is higher that the lowest realized cost from all other origins. Beyond this finite sample randomness, we can easily add a multiplicative disturbance term to the trade cost function (5), without altering

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

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

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.

[9]Such an analysis would require us to consider all possible locations for our lost cities—all combinations of 10 sites chosen from millions of grid points—compute least effort paths for each guess, calculate our objective function, and iterate many times over. This requires computational power beyond current capabilities.

our estimation strategy.[10]

Step 1 of our estimation estimates the distance elasticity of trade ζ using trade among known cities. Under the parametrization (5) for the trade cost function, the following moment condition equates the expected fraction of goods from i to j among shipments from known cities, with the probability (4) of a good being sourced from i to j conditional on being sourced from a known city,

$$\begin{aligned}\mathbb{E}&\left[\frac{N_{ij}}{\sum_{k\in\mathcal{K}\setminus\{j\}} N_{kj}}\right]\\ &= \Pr\left[c_{ij}(\omega) \leq \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\} \middle| \min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\}\right]\\ &= \alpha_i \beta_j Distance_{ij}^{-\zeta}, \quad (6)\end{aligned}$$

with $\alpha_i = w_i^{-\theta} T_i$ and $\beta_j = \mu / \sum_{k\in\mathcal{K}\setminus\{j\}} w_k^{-\theta} T_k (\tau_{kj})^{-\theta}$. Under this moment condition (6), we estimate the distance elasticity $\hat{\zeta}$ by Poisson Pseudo Maximum Likelihood. This estimation uses only trade shares among known cities, $N_{ij} / \sum_{k\in\mathcal{K}\setminus\{j\}} N_{kj}$, bilateral distances between known cities, $Distance_{ij}$, and origin and destination fixed effects to control for α_i and β_j. This follows closely the procedure in Eaton, Kortum, and Sotelo (2012), with the only difference that we derive conditional probabilities (non-internal trade among the subset of known cities only) and not unconditional ones.

Step 2 of our estimation uses the distance elasticity $\hat{\zeta}$ from step 1, and our dataset on all trade flows between known and lost cities. We estimate exporter fixed effects $\alpha_i = w_i^{-\theta} T_i$ for all cities, and the latitudes φ_l and longitudes λ_l of lost cities, collected in the vector of parameters $\boldsymbol{\beta}$

$$\boldsymbol{\beta} = (\alpha_1, \cdots, \alpha_{Kane\check{s}-1}, \alpha_{Kane\check{s}+1}, \cdots, \alpha_{K+L}, \varphi_{K+1}, \lambda_{K+1}, \cdots, \varphi_{K+L}, \lambda_{K+L})',$$

where we arbitrarily normalize $\alpha_{Kane\check{s}} \equiv 100$. We use (3) to form the moment condition,

$$\mathbb{E}\left[\frac{N_{ij}}{\sum_{k\neq j} N_{kj}}\right] = \Pr\left[c_{ij}(\omega) \leq \min_{k\neq j}\{c_{kj}(\omega)\} \middle| c_{jj}(\omega) > \min_{k\neq j}\{c_{kj}(\omega)\}\right] = \frac{\alpha_i Distance_{ij}^{-\hat{\zeta}}}{\sum_{k\neq j} \alpha_k Distance_{kj}^{-\hat{\zeta}}}. \quad (7)$$

[10] To account for departures from the simplest gravity model where only distance matters, we can add a multiplicative disturbance term drawn from a joint Gamma distribution as in Eaton, Kortum, and Sotelo (2012),

$$\tau_{ij}^{-\theta} = \mu Distance_{ij}^{-\zeta} \nu_{ij}, \text{ with } \nu_{ij} \sim Gamma\left(\frac{1}{\eta^2} \frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j} \alpha_k Distance_{kj}^{-\zeta}}, \frac{\eta^2}{\alpha_i Distance_{ij}^{-\zeta}}\right).$$

Under this distributional assumption, treating the ν's as realizations from a random variable, the moment condition (7) remains the same (see mathematical appendix A for a formal derivation),

$$\mathbb{E}\left[\frac{N_{ij}^{data}}{\sum_{k\neq j} N_{kj}^{data}}\right] = \mathbb{E}\left[\frac{\alpha_i Distance_{ij}^{-\zeta} \nu_{ij}}{\sum_{k\neq j} \alpha_k Distance_{kj}^{-\zeta} \nu_{kj}}\right] = \frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j} \alpha_k Distance_{kj}^{-\zeta}}.$$

Reassuringly, both realized (random) and expected trade shares are all in $(0, 1)$ and add up to one.

It simply states that the expected share of shipments from i to j equals the probability of a shipment being sourced from i to j. Under this condition moment (7), we estimate the parameters $\boldsymbol{\beta}$ by a method of moments, solving the following non linear least squares minimization problem,[11]

$$\boldsymbol{\beta} \;=\; \arg\min_{\boldsymbol{\beta}} \sum_{j} \sum_{i \neq j} \left(\frac{N_{ij}}{\sum_{k \neq j} N_{kj}} - \frac{\alpha_i Distance_{ij}^{-\hat{\zeta}}}{\sum_{k \neq j} \alpha_k Distance_{kj}^{-\hat{\zeta}}} \right)^2 . \tag{8}$$

Note that our NLLS estimator (8) uses data contained in trade zeros explicitly. For instance, consider a case where the observed trade share from i to j_1 is zero, but the trade share from i to j_2 is positive. Our estimator trades off a low origin i fixed effect α_i to match the observed zero share from i to j_1, and a higher α_i to match the positive share from i to j_2. Furthermore, if i's location is lost, so that both $Distance_{ij_1}$ and $Distance_{ij_2}$ are unknowns to be estimated, our estimator also trades off a large $Distance_{ij_1}$ to match the observed zero share from i to j_1 and a lower $Distance_{ij_2}$ to match the positive share from i to j_2, where the Euclidean geometry of our space imposes a mechanical link between $Distance_{ij_1}$ and $Distance_{ij_2}$.

Step 3 of our estimation collects all estimated parameters to compute a measure of fundamental city sizes. To derive this measure, we need to fully solve our model in general equilibrium. This requires a number of additional assumptions. First, we solve this model under the continuous limit of Eaton and Kortum (2002), i.e., we assume away the randomness due to finitely many shipments. Second, we assume perfect competition for simplicity, so that prices equal marginal costs.[12] With no arbitrage, the equilibrium price for commodity ω in city j is the lowest cost among all possible sources, $p_j(\omega) = \min_{i \in \mathcal{K} \cup \mathcal{L}} \{c_{ij}(\omega)\}$. Third, we assume trade balance at the city level, so that total spending equals the amount paid to local factors,

$$X_i = \sum_k X_{ki} = w_i Pop_i, \tag{9}$$

where Pop_i is the size of city i's population. 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,[13]

$$Size_i \;\equiv\; \frac{w_i^{autarky} Pop_i}{P_i^{autarky}} \;\propto\; Pop_i T_i^{1/\theta}, \tag{10}$$

with $P_i^{autarky}$ the ideal price index in city i under autarky. This measure for city size is convenient because it only depends on exogenous parameters, Pop_i and $T_i^{1/\theta}$. If, for instance, trade frictions

[11] See appendix A for details on how we address the issue of local minima in our non-linear minimization problem.

[12] Imperfect (Bertrand) competition as in Bernard et al. (2003) would give identical results, with all aggregate variables simply shifted by a multiplicative constant.

[13] For a derivation of (10), see Eaton and Kortum (2002), equation (15) on page 1756.

or the size of other cities were to change, this measure would remain invariant. This measure can be computed using our parameter estimates and an assumption for the trade elasticity θ only,[14]

$$Size_i \;\propto\; Pop_i T_i^{1/\theta} \;\propto\; \hat{\alpha}_i^{1+1/\theta} \sum_k \widehat{Distance}_{ki}^{-\hat{\zeta}} \hat{\alpha}_k, \tag{11}$$

where we use $\theta = 4$ from Simonovska and Waugh (2014). 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*. Equation (11) shows how to recover the fundamental size of a city, in a counterfactual autarky state, using only observable trade data. 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 ($Distance_{ki}$ small and/or α_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 (11) formally adjusts for this endogenous factor price response.

Standard errors. Robust (White) standard errors are computed analytically and account for heteroskedasticity and for the two-step nature of our estimation (PPML then NLLS). To gauge visually the precision of estimates for the location of lost cities, we draw maps with confidence regions around our point estimates. For each lost city l, we draw four contours such that the true location lies inside with respectively 50%, 75%, 90% and 99% probability. These elliptical contours are computed using analytical solutions for the iso-density curves of the estimated distribution of the geo-coordinates of lost city l, $\mathcal{N}(\hat{\boldsymbol{\beta}}_l, \hat{\boldsymbol{\Sigma}}_l)$ with mean $\hat{\boldsymbol{\beta}}_l' = (\hat{\varphi}_l, \hat{\lambda}_l)$ and covariance matrix $\hat{\boldsymbol{\Sigma}}_l$. They account not only for the precision of the latitude and longitude of city l, but also for the co-variance of those geo-coordinates.

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

$$precision\,(l) \;=\; \sqrt{\mathbb{E}_{(\varphi,\lambda)\sim\mathcal{N}(\hat{\boldsymbol{\beta}}_l,\hat{\boldsymbol{\Sigma}}_l)}\left[\left(Distance\left(\hat{\varphi}_l,\hat{\lambda}_l;\varphi,\lambda\right)\right)^2\right]}, \tag{12}$$

where $Distance(\hat{\varphi}_l, \hat{\lambda}_l; \varphi, \lambda)$ is the distance between the estimated location for l and a location drawn from our estimated $\mathcal{N}(\hat{\boldsymbol{\beta}}_l, \hat{\boldsymbol{\Sigma}}_l)$. Heuristically, it means that our point estimates are within

[14] See mathematical appendix A for a formal derivation. To recover $Pop_i T_i^{1/\theta}$ 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 choose $\theta = 4$ from the literature instead. Since the parameter θ mostly affects the absolute scale of our estimates of city sizes, but not relative city sizes (in logs), this choice is of little consequence.

this distance, $precision\,(l)$ expressed in km, from the true location with probability 75%.[15]

IV 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 IV.A. 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 IV.B. As a proof of concept, we fictitiously "lose" the location of known cities, and compare their known locations to our recovered gravity estimates in section IV.C. Finally, we propose to use our gravity model to evaluate the validity of potential unnamed archaeological sites in section IV.D.

IV.A Using Gravity to Recover the Location of Lost Cities

Table II presents the estimated geo-coordinates of lost cities, along with robust standard errors. Panel A of table III presents our estimates for the distance elasticity of trade, $\zeta = 1.912$ with a standard error of 0.189. This suggests that the impact of distance on trade around 1880 BCE was surprisingly similar to what it is today, with modern elasticity estimates typically near unity (Disdier and Head, 2008; Chaney, 2018), and estimates for shipments transported by road above unity —see Cosar and Demir (2016) for a distance elasticity around 2 based on overland transit of exports from Turkish cities, almost equal to our ancient estimate.

Figures III and IV show maps with our point estimates and confidence regions for each lost city separately. A "•" sign depicts the estimated location from our structural estimation (8), surrounded by contours representing the confidence regions for that city (50th, 75th, 90th and 99th percentiles).[16] For most cities, our estimates are tight, in the sense that the confidence area is at most 100 km 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 III: all but three of our measures of precision are smaller than 100 km (60 miles), and less than 50 km in four cases. This to be compared to the average distance of 223 km 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). Those historians base they proposals on a

[15]See mathematical appendix A for analytical formulas of standard errors, confidence regions, $precision\,(l)$ and and their derivations.

[16]The confidence regions for *Šuppiluliya* are not shown as they are too wide.

careful analysis of ancient texts, ancient itineraries, topographical studies, surviving toponyms, etc.[17] This comparison allows us to confront our estimates, obtained by a purely quantitative method—a structural gravity estimation, to those obtained by historians from a purely qualitative method. We consider this comparison to be an informal external validity test.

In five cases, *Durhumit*, *Kuburnat*, *Ninašša*, *Šinahuttum*, and *Wašhaniya*, our gravity estimates for the location of lost cities are close to the conjecture of at least one of the two historians (less than 70 km - 45 miles). In one case, *Ninašša*, our estimate favors the proposal made by Forlanini over that from Barjamovic. In two cases, *Durhumit* and *Wašhaniya*, our gravity estimates are remarkably close to the proposals made by Barjamovic (48 km and 13 km respectively), and favor Barjamovic over Forlanini. For *Šinahuttum*, both historians agree on the same location, and our estimate is extremely close (24 km). For *Kuburnat*, Forlanini and Barjamovic disagree by about 70 km, and our gravity estimate is about 70 km from both proposals. We view these cases where our structural gravity estimates are close to at least one historian's proposal as an endorsement that the true locations of those cities are indeed at or very near those sites. 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 *Hahhum*, our estimate is not nearly as close to the historians' proposals, but the distance between their proposal and our estimate is of the same order of magnitude as the precision of our point estimate (100 km distance versus 60 km precision). Our gravity estimate is shifted towards the North and West of their proposal. It lies 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 its location on the Euphrates River, which the gravity estimate ignores.

In the case of *Tuhpiya*, Forlanini' and Barjamovic's proposals are near each other, but our gravity estimate is far from theirs (130 and 110 km respectively). However, our estimate near the modern town of Sorgun-Yozgat (22 km) corresponds to an earlier proposal by Cornelius (1963).

Finally, in three cases, *Purušhaddum*, *Šuppiluliya* and *Zalpa*, our estimates are statistically too imprecise to draw any definitive conclusion. For *Purušhaddum* and *Zalpa*, our estimates are also far from both historians' proposals. For *Šuppiluliya*, our estimate is not very far from the historians' proposals (about 90 km from both), but it is so imprecise that we cannot draw any inference: the precision for *Šuppiluliya* (90,000 km) is more than twice the circumference of the earth.

To conclude, we often find a remarkable agreement between our quantitative method for locating

[17] We describe in appendix B a few of the steps those historians use to infer the likely location of lost cities.

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 *Purušhaddum*, *Šuppiluliya*, and *Zalpa*, which are imprecisely estimated. Furthermore, in the case of *Ninašša*, our gravity estimates favor the location proposed by Forlanini over this given by Barjamovic, while in the cases of *Durhumit* and *Wašhaniya*, they favor Barjamovic over Forlanini.

IV.B Gravity Estimates of Lost Cities versus 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). It is also reminiscent of the pioneering work of Gardin and Garelli (1961) and their use of computer programming to aggregate information contained in the Assyrian texts in the 1960's.

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

> *To the Port Authorities of <u>Kaneš</u> from your envoys and the Port Authorities of <u>Wahšušana</u>. We have heard the tablets that the Station(s) in <u>Ulama</u> and <u>Šalatuwar</u> 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 <u>Ulama</u> and two messengers by way of <u>Šalatuwar</u> to <u>Purušhaddum</u> 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.*
> [Tablet Kt 83/k 117 (Günbattı, 1998), lines 1-24]

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

$$||AB|| + ||BC|| \leq (1+\delta)\,||AC||\,, \qquad \text{(“short detour”)}$$

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 C via B does not represent more than a δ% detour compared to going straight from A to C.

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

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

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 σ standard deviations longer than the average known segment.

Figure V 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, $A \to X \to B$, and $C \to D \to X$. The figure also shows how raising the parameters δ and σ widens the size of the admissible region.[19]

Those two sets of constraints, "short detour" and "pit stop", seem reasonable, and historical evidence suggests that Assyrian merchants were indeed following close to optimal routes (Palmisano, 2013; Palmisano and Altaweel, 2015; Palmisano, 2017).

[18] To compute this measure of distance, we collect systematic information on elevation on a fine grid. We use Langmuir (1984) formula for calculating the time for a normal human being to travel between any two contiguous grid-points. We prevent crossing large impassable rivers except in a few locations (fording). We then use Dijkstra (1959) algorithm to compute the optimal travel route between any two grid-points. See appendix C for details.

[19] The top row of figures only imposes the "short detour" constraint, while the bottom row of figures further imposes the "pit stop" constraint. The left figures show the example of an itinerary of the type $A \to X \to B$, where A and B are know, and X is lost. For example, points X_1 and X_2 are two possible candidates such that going from A to B via X_1 (or X_2) represents only a 5% detour compared to going straight from A to B ("short detour" constraint). But only point X_1 also satisfies the constraint that each leg of the trip (A to X and X to B) are no more than 0.4 standard deviations longer than the average trip ("pit stop" constraint). The middle figures show similar exercises for an itinerary of the type $C \to D \to X$, with C and D known and X lost. The right figures jointly impose constraints from both itineraries. Darker shades of grey correspond to shorter detours.

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. Appendix D provides further details.

We present our results in a series of maps in figures VI and VII. Each map depicts the admissible region for a given city (dashed line), using the above procedure to code information from merchants' itineraries, with parameters $\delta = 2.6$ for the "short detour" constraint and $\sigma = 1.3$ for the "pit stop" constraint. For comparison, we also show on the same map our point estimate and 90th percentile confidence region from estimating our gravity model (8) (solid ellipse), as well as the locations proposed by historians Forlanini (2008) and Barjamovic (2011) ("▲ F" and "▲ B" signs).

Our gravity estimates for the location of lost cities all lie within their admissible regions,[20] and are therefore compatible with the qualitative information from merchants' itineraries, with the exception of *Purušhaddum*. *Purušhaddum* was the main Assyrian market in Anatolia after *Kaneš*. It was located where the Assyrian zone of trade intersected with a regional network further to the west (Barjamovic, 2008; Erol, 2013), but its location remains debated (Forlanini, 2017, p.242f). Unfortunately, its peripheral position in the Assyrian network makes it difficult for the gravity model to suggest a location (see section IV.C below), with large and imprecise confidence regions (figure III). The addition of the constraints imposed merchants' itineraries suggests that the actual position is more likely to be sought in the overlap between the confidence region and the admissible region (the intersection of the solid line ellipse and the dashed line region on figure VI). At the Southwest corner of the region of overlap is a possible candidate that has not yet been surveyed by archaeologists, Akşehir Karahöyük (Barjamovic, 2017, p.314).

As we never use the information contained in merchants' itineraries to estimate our structural gravity model, we see this compatibility as an encouraging sign that our estimates are consistent with historians' qualitative methodology. Our procedure for extracting information from ancient text in an automated and systematic manner is also complementary to that qualitative methodology.

IV.C Proof of Concept: What If We Fictitiously "Lose" Some Known Cities?

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

[20] A few admissible sets are wide, and do not impose a strict constraint on the location of lost cities, e.g. *Tuhpiya*. The reason we cannot impose a stricter set of parameters 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 parameters, the admissible region for some lost cities would be empty sets, e.g. *Kuburnat*.

We pick one known city at a time, and estimate its coordinates, as if it had been lost. We perform this exercise separately for each of the 15 known cities. Each time, we set the distance elasticity to our estimated $\hat{\zeta} = 1.9$, we fix the other known cities to their true location, and the lost cities to their estimated location from table II, but re-estimate all other parameters solving the same non-linear least squares problem as (8).

Figure VIII presents the results of our within-sample predictions.[21] For each known city, an arrow goes from the true site, denoted by a "♦" sign, towards the estimated location, denoted by a "•" sign. The results suggest that our estimates are very precise for central cities (*Hattuš*, *Hanaknak*, *Kaneš*, *Mamma*, *Tappagaš*, *Šamuha*, *Unipsum*, *Zimišhuna*), but less so for peripheral cities (*Hurama*, *Malitta*, *Šalatuwar*, *Timelkiya*, *Ulama*, *Wahšušana*). One exception is *Karahna*, a poorly estimated central city; the likely reason is that trade with *Karahna* was extremely limited and therefore its size (the second smallest of our 25 cities) and location are imprecisely estimated. Among the nine centrally located cities, the average distance between the true and estimated locations is 40 km (median 33 km), and often substantially lower (*Zimišhuna*=1 km, *Hattuš*=3 km, *Tapaggaš*=17 km, *Hanaknak*=19 km, *Kaneš*=33 km). We conclude that our proposed inverse gravity estimation of the location of lost cities is reliable for central cities, but less precise for peripheral cities.

IV.D Using Gravity to Evaluate the Validity of Potential Archaeological Sites

Having estimated the parameters of a structural gravity model, we can in theory evaluate whether any potential archaeological site is a good candidate for a given lost city. Our estimation delivers a probability distribution, over the two-dimensional geographic space, for the likely location of each lost city. For any potential site for lost city l with latitude φ and longitude λ, $\hat{f}_l(\varphi, \lambda)$ assigns a probability (density) that this is the true site for l,

$$\hat{f}_l(\varphi, \lambda) = \frac{\exp\left(-\frac{1}{2(1-\hat{\rho}^2_{\varphi_l,\lambda_l})}\left[\frac{(\varphi-\hat{\varphi}_l)^2}{\hat{\sigma}^2_{\varphi_l}} + \frac{(\lambda-\hat{\lambda}_l)^2}{\hat{\sigma}^2_{\lambda_l}} - \frac{2\hat{\rho}_{\varphi_l,\lambda_l}(\varphi-\hat{\varphi}_l)(\lambda-\hat{\lambda}_l)}{\hat{\sigma}_{\varphi_l}\hat{\sigma}_{\lambda_l}}\right]\right)}{2\pi\hat{\sigma}_{\varphi_l}\hat{\sigma}_{\lambda_l}\sqrt{1-\hat{\rho}^2_{\varphi_l,\lambda_l}}}, \tag{13}$$

where the estimated geo-coordinates for city l $(\hat{\varphi}_l, \hat{\lambda}_l)$, and their estimated variances $(\hat{\sigma}^2_{\varphi_l}, \hat{\sigma}^2_{\lambda_l})$ and correlation $(\hat{\rho}_{\varphi_l,\lambda_l})$, are given in table II. Of course, this formula only summarizes the best possible estimate for a given location according to our structural gravity model, and it should be complemented with additional historical and archaeological evidence.

[21] Appendix Table I lists the geo-coordinates for all known city, both true and estimated, as well as the distances, in km, between the true and estimated locations. As a robustness check, we also run an alternative proof of concept exercise, estimating not just the location of one fictitiously lost city, but re-estimating the locations of all ten truly lost cities as well. The results are presented in Appendix Table II. The results are similar, although less accurate, as we are compounding measurement error for one fictitiously lost city with that of the ten truly lost cities.

We propose to apply formula (13) on a list of 87 unnamed archaeological sites in Anatolia, known to have been occupied during the Middle Bronze Age period when Assyrian merchants were active in Anatolia (those sites are tabulated in Barjamovic (2011), Appendix 2.1 and 2.2. on pp. 72-74). This allows us to point to some of the strengths and limits of our quantitative method. For brevity, we focus our discussion here on two lost cities, *Durhumit* and *Wašhaniya*.[22]

The location of the city of *Durhumit* is controversial and has been the topic of intense debate among historians in recent years. The city was a central market of copper during the period of Assyrian trade and is mentioned over 200 times in the trade records from *Kaneš*. It reappears in documentation of the Hittite state in the 14th-13th century BCE as a fortified imperial border province. Assyrian and Hittite sources seem to favor a location in different directions (Forlanini, 2008; Matthews and Glatz, 2009; Barjamovic, 2011, p.261-265; Cammarosano and Marizza, 2015, p.180f; Kryszeń, 2016, p.343ff; Corti, 2017, p.232). Scholarly disagreement follows the same east-west axis as our structural gravity estimate, with its east-west confidence region (figure III). The analysis constrained by itineraries in turn seems to favor a central northern position of the city (figure VI). Our gravity estimates ranks the unexcavated archaeological site of Ayvalıpınar as the most likely candidate for *Durhumit*. An Assyrian seal carved in a workshop at *Kaneš* (Lassen, 2014, p.118) was found on the surface of the site in strong suggestion that it formed part of the Assyrian trade network, and it has previously been thought to belong to the region of *Durhumit* based on qualitative analyses of the data (Barjamovic, 2011, p.386; Dönmez, 2017, p.88). However, it is probably located too far south and inside the core area of the Hittite state to be the city itself. The second most likely candidate, near Oluz, has been under archaeological excavation since 2007 revealing an occupational hiatus between the Early Bronze Age and the period of the Hittite Empire (Dönmez, 2017). This effectively eliminates it as candidate site for *Durhumit*. The third candidate (Ferzant) is not an urban site but a cemetery. The fourth candidate (Doğantepe) is a large site that has not been subject to systematic excavation. It is a viable candidate for *Durhumit*, although there may be other, possibly better proposals (Dönmez, 2017; Corti, 2017, p.222).

The city of *Wašhaniya* is known to have been located as the first major stop on a route leading west from *Kaneš* to *Wahšušana*. The gravity estimate corresponds fairly well to the conjecture proposed by historians (Forlanini, 2008; Barjamovic, 2011). The most likely candidate is Yassıhöyük, which has come under excavation within the last decade and revealed findings dated to the period of Assyrian trade (Omura, 2016). Excavations at a number of sites located along the historically

[22] The full set of results are presented on Appendix Table III. For each lost city l, the table lists the five most likely unnamed archaeological sites, along with their distance from our gravity estimate for the location of l, and the (logged) probability density that this is the correct site according to (13).

important route leading southeast from Yassıhöyük to Kayseri have revealed remains from the period (Weeden and Matsumura, 2017, p.108), including Suluca, Zank and Topaklı on the list of likely candidates. The site of Kırşehir Kalehöyük is also located close to the predicted location of our gravity model, but does not figure on the list of candidates because it lacks clear remains dated to the Bronze Age. The main mound now has the Alaaddin Mosque (built 1230 CE) and a high school built on it (Adıbelli, 2013). Dense later occupation of its surroundings makes it difficult to ascertain whether the city was occupied during the period of Assyrian trade.

We draw two lessons from this analysis. First, our structural gravity model should prove useful in selecting the most likely among a list of candidate archaeological sites to locate lost cities, but this selection ought to always be complemented by historical evidence. Second, it is likely that any list of candidate archaeological sites will be incomplete, as important ancient cities may lay buried under modern settlements, inaccessible to archaeologists, and may never be found.

V Determinants and Persistence of City Sizes

We now turn to a systematic discussion of our estimates of ancient city sizes, and of the determinants of the city size distribution. With no reliable historical or archaeological evidence on the size of those ancient cities to use as external validity, we explore the geographic and topographic determinants of city sizes, confront our estimates of ancient city sizes to measures from modern data, and characterize the distribution of ancient city sizes.

V.A City Size Estimates

Our estimates of the fundamental size of ancient cities ($Pop_i T_i^{1/\theta}$), presented in panels B and C of table III, do not achieve the conventional levels of statistical significance. This is to be expected given the sparsity of our four millennia old data.

We should also note that there does not seem to be any systematic bias for larger cities to be more or less likely to have been unambiguously located by historians.[23] We offer a potential rationale for this finding below. We provide evidence that city sizes are persistent, so that large ancient cities tend to be located at or near large modern cities. As archaeologists are rarely able to survey and excavate densely populated urban areas, this suggests that at least some large ancient cities may never be discovered, as they lay buried underneath modern cities.

[23] A Welch's t-test of equality between the sizes of known and lost cities gives a p-value of 0.29.

V.B Determinants of City Sizes: the Road-knot Hypothesis

To probe the determinants of ancient city sizes, we project our ancient size estimate ($Pop_i T_i^{1/\theta}$) on two geographic observables: terrain ruggedness, a local measure of the defensive advantage of a site,[24] 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 in Asia Minor, Ramsay suggested that key to understanding urban geography in the classical antiquity is the realization that the local terrain only allows a limited number of routes to cross the area. He observed that the zones where such routes intersect formed what he called "road-knots," which tend to 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 within the zone of intersection could vary from period to period, but would remain in its immediate vicinity. Ramsay's basic hypothesis, that the existence of road-knots may be causally related to the presence of major administrative and trading centers, was further elaborated and advanced by French (1993).[25]

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 20 km to the west. Several other large ancient and corresponding modern cities, such as *Hurama*-Karahöyük/Elbistan, *Mamma*-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 20 km of city i, which varies be-

[24] In unreported regressions, we experimented with alternative measures of local amenities: crop yields under primitive technologies, elevation, 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.

[25] A 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 modern French cities originating from ancient Roman towns. Dalgaard et al. (2018) find that the density of two-millennia old Roman roads is correlated with the current road network and contemporary economic activity in Europe.

tween 2 and 5.[26] 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, $NaturalRoads_i$, is immune to this reverse causality concern. 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 allow for maritime travel near the coast. We use Dijkstra (1959)'s algorithm to compute the optimal route between any two points (see appendices C and E for details). 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. We weight each route of duration d by a weight proportional to $d^{-\zeta}$ using our estimated ancient distance elasticity of trade $\zeta = 1.9$, and record all intersections or overlaps of those routes. This measure corresponds to the notion of betweenness centrality in the network of optimal routes. Implicitly, we are assuming in the background that a gravity model with distance elasticity $\zeta = 1.9$ governs the movement of a population uniformly distributed over space. We define a road-knot score equal to the number of intersections or overlaps for each location. Our variable $NaturalRoads_i$ is the simple average of this road-knot score within 20 km of city i. In essence, it measures the propensity of a given site to connect to the natural routes network. This measure is arguably exogenous as it only uses topographical data as input.

Figure IX shows a heat map of our road-knot scores for Turkey and the surrounding region. Major modern urban settlements and transportation arteries, not included on this map, overlap with our road measure, although neither were used as input.

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

$$\ln\left(Pop_i T_i^{1/\theta}|_{ancient}\right) = a + b \cdot \ln\left(Roads_i\right) + c \cdot \ln\left(Ruggedness_i\right) + u_i. \tag{14}$$

[26]We use 20 km 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. If there is no intersection within 20 km, a city assumes a score of 2 since each city is necessarily on a road itself. 16 out of 25 cities have a road score of 2. Two cities, *Kaneš* and *Ulama*, have road scores of 5.

$Roads_i$ corresponds either to the number of natural paths intersecting or overlapping near city i, $NaturalRoads_i$, in columns 1 and 4, or to the number of roman roads intersecting near city i, $RomanRoads_i$, in columns 2 and 5. In columns 3, 4, and 5, we control for $Ruggedness_i$, a measure of how rugged the terrain is around i (Riley, DeGloria, and Elliot, 1999).

We find robust and significant evidence in support of the road-knot hypothesis. The more road intersections near in a city, the larger it is. While the *RomanRoad* variable has a positive but non significant effect (columns 2 and 5), our a priori measure of the connectedness of a city to the natural road network, $NaturalRoad$, is strongly significant, with a p-value below 0.015 both on it own in column 1, and when controlling for *Ruggedness* in column 4. Our natural road score accounts for more than a fifth (22%) of the variation of ancient city sizes.[27] Figure X presents visual evidence of this strong correlation, and shows it is not driven by outliers.

Two observations are in order. First, our method explains which among the existing ancient cities are large, which are small. We do not attempt to explain *where* cities are located, only how large they are given their location. Second, our measure of connectedness to the natural road network, $NaturalRoads$, is particularly relevant in this central part of modern Turkey, a high plateau with many smaller mountains. Had we applied our method on a flat plain, such as lower Mesopotamia, Eastern China, Northwestern Europe, or the U.S. Midwest, the topography would presumably have offered little guidance on natural road access of a particular location, and access to waterways instead might play a larger role. Anatolia, with its clearly defined mountain ranges and valleys, is a particularly well suited laboratory to test our road-knot hypothesis.

Ruggedness, is also correlated with our estimates of ancient city sizes, accounting for 17% of the variation in column 3 of table IV, and it remains significant when we control for road-knots (columns 4 and 5). It suggests that the defensive value of a site contributed to the emergence of larger cities. Among all measures of local amenities, *Ruggedness* is the only variable significantly correlated with city size. Crop yield, elevation, distance to the nearest river, and distance to mineral deposits exploited in the Early Bronze Age are all either insignificant, or driven by outliers.

V.C The Distribution of City Sizes over Four Millenia

Next we confront our ancient size estimates to modern size measures, and document a strong persistence of the distribution of city sizes over four millennia. To do so, we match the locations of ancient sites with corresponding modern urban settlements. We then project two alternative

[27] The results are robust to using the locations proposed by Barjamovic (2011) instead of using our estimated locations for lost cities, and to restricting the sample to known cities only. See Appendix Table IV.

measures of modern city sizes on our ancient size estimates and a control for geographic amenities.

We drop one outlier from our sample, *Purušhaddum*, which is matched with the modern city of Ankara. A minor provincial town at the turn of the 20th century, Ankara was chosen as the capital of the Turkish Republic by Mustafa Kemal Atatürk in 1923. As a result, it rapidly grew to be the second largest city of the country. It is therefore now much larger than any other city in our sample, primarily due to the idiosyncratic positive effect of assuming a political role in recent history.[28]

Our first measure of modern size, $Population_i$, measures the total urban population living within 20 km of ancient city i.[29] 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 20 km of ancient city i. In the absence of modern city-level income data, nighttime luminosity is a strong correlate of local incomes (Hodler and Raschky, 2014).

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

$$\ln\left(Size_i|_{modern}\right) \;=\; d \;+\; e\cdot\ln\left(Pop_i T_i^{1/\theta}|_{ancient}\right) \;+\; f\cdot\ln\left(CropYield_i\right) \;+\; v_i, \qquad (15)$$

where $Size_i|_{modern}$ is either population or night lights depending on specification. $CropYield_i$, a measure of agricultural suitability around i, controls for local amenities.[30]

Columns 1 and 4 of table V show the results of simple specifications without any geographic controls for each measure of modern size, population and night lights. These results are also plotted in the two panels of figure XI. The correlation between ancient and modern sizes is high, 0.38 for both measures of modern size. This surprising level of persistence in city sizes from the the 20th century BCE to the 21st century CE is robust to controlling for modern local crop yields (columns 3 and 6), while crop yields are not statistically significant on their own (columns 2 and 5).

The strong and robust correlation of city sizes over four millennia is unlikely to be a mere coincidence, which gives us confidence that our estimates for ancient city sizes are plausible. While our results do not offer a definitive explanation for this persistence, two mechanisms highlighted in the

[28]We note the role of Ancyra as a Roman, Byzantine and Ottoman provincial center of varying importance, but nothing in history warrants its extreme current size based on political events in the 20th century CE.

[29]For lost cities, we use our own estimates from section IV.A. For two ancient cities, Tuhpiya and Samuha, there are no modern-day urban population centers within 20 kms of their coordinates. We winsorize their populations to the smallest population within the sample. The results are robust to dropping them and estimating with a sample of 22 cities. The results are also robust to using Barjamovic (2011) instead of our gravity estimates for the location of lost cities, to restricting the sample to known cities only, or to using alternative procedures for matching ancient and modern towns. See Appendix Table V for those robustness checks.

[30]To construct $CropYield_i$, we use the low-input level rain-fed cereal suitability index of IIASA/FAO (2012). We average this measure within an area of 20 km radius around the coordinates of ancient city i. In unreported regressions, we also experimented with other geographic controls: elevation, 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.

literature are potentially at play. The first mechanism is path dependence. Despite a series of large shocks, with 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, people seem to have come back to the same locations to restart cities. The second mechanism is the effect of time invariant fundamental characteristics. We have shown that the advantageous location as a natural trading hub conferred by the topography of the land is a key determinant of ancient city sizes. To the extent that transport routes are shaped by similar constraint throughout history, topography may have continued to affect the relative size of cities. We hope to further explore this mechanism in other historical settings and regions in future research.

V.D Did Zipf's Law for Cities hold in the Bronze Age?

In the absence of any reliable historical evidence on the population sizes of ancient cities, can we use our structural gravity model to evaluate whether the distribution of ancient cities was governed by Zipf's law? Formally, our structural estimates of city sizes do not inform us directly about population sizes, as they confound population and efficiency, $Size_i \propto Pop_i T_i^{1/\theta}$. We can however use our findings in the previous section V.C to get suggestive evidence on the distribution of population sizes. The regression of modern population sizes against our structural measure of ancient city sizes suggests a correspondence between structural size and population: $\ln(Population_i) \approx constant + 0.23 \ln(Pop_i T_i^{1/\theta})$ (see column 1 in table V). Using this correspondence naively, we then test whether the distribution of ancient population sizes follows Zipf's law, applying the methodology advocated by Gabaix and Ibragimov (2011), by estimating

$$\ln\left(Rank_i - 1/2\right) \;=\; g \;-\; h \cdot \ln\left(Population_i\right) \;+\; w_i, \tag{16}$$

where $Rank_i$ is city i's population rank, starting from the largest city, and h is the Zipf exponent.

Figure XII presents the results of estimating (16). The data suggests that the distribution of city population sizes in the Bronze Age is very well approximated by Zipf's law, with a Zipf exponent of $h = 1.08$ (robust $s.e. = 0.211$ and $R^2 = 0.719$), very close to modern estimates—Rosen and Resnick (1980) find an average Zipf exponent of 1.13 for 44 countries in 1970. This finding would suggest that Zipf's law is a stable empirical regularity over four millenia. It should of course be interpreted with due caution. First, our variable *Population* is estimated with error, so that our estimate for the Zipf exponent may be biased. And second, we have no direct evidence on actual population sizes, and rely instead on trade data and a structural gravity model to infer population sizes in the Bronze Age. It is however an intriguing finding, worthy of further investigations.

VI Structural versus Naive Gravity

We conclude with a brief comparison between our structural gravity approach and a 'naive' gravity similar to Tobler and Wineburg (1971). In order to provide a meaningful comparison, we use the exact same measure of bilateral trade flows as in our main analysis, but perform an estimation similar to theirs.[31] First, we define an undirected measure of interactions between cities i and j , I_{ij} in their notation, by adding the number of shipments going from i to j and from j to i: $I_{ij} = N_{ij} + N_{ji}$ where N_{ij} is our directed measure of shipment counts. Second, we impose that city size, P_i in their notation, is proportional to the total number of shipments traded by city i, $P_i = \sum_{j \neq i} I_{ij}$. Finally, we postulate a 'naive' gravity structure linking sizes and distances to interactions,

$$I_{ij} \ = \ k \cdot \frac{P_i P_j}{Distance_{ij}^2} \Leftrightarrow \frac{P_i P_j}{I_{ij}} \ = \ \beta \left((\varphi_i - \varphi_j)^2 + \cos^2 \left(\frac{37.9}{180} \pi \right) (\lambda_i - \lambda_j)^2 \right),$$

where β is a simple multiplicative constant, and the φ's are the latitudes and the λ's the longitudes of cities i and j. Collecting β and all the geo-coordinates of lost cities in the vector $\boldsymbol{\theta} = (\beta, (\varphi_{K+1}, \lambda_{K+1}), \cdots, (\varphi_{K+L}, \lambda_{K+L}))$, we estimate this model by non linear least squares,

$$\boldsymbol{\theta} \ = \ \arg\min_{\boldsymbol{\theta}} \sum_j \sum_{i \neq j} \left(\frac{P_i P_j}{I_{ij}} - \beta \left((\varphi_i - \varphi_j)^2 + \cos^2 \left(\frac{37.9}{180} \pi \right) (\lambda_i - \lambda_j)^2 \right) \right)^2 . \qquad (17)$$

We then compare our results using structural versus naive estimates.[32]

First, the estimated locations of lost cities from both models are far apart—123.3 km on average. So the modeling choices, structural versus naive gravity, have a substantial impact on our estimates. Our structural estimates are also substantially closer to the proposals from historian Barjamovic (2011)—87.1 km on average, than the naive estimates—154.4 km on average. Our structural gravity model seems better at identifying the location of lost cities than a simpler naive gravity model.

Second, our structural city-size estimates are only weakly correlated with the naive size measure, the total trade originating from a city: the correlation between $\ln(Pop_i T_i^{1/\theta})$ and $\ln(P_i)$ is 0.4, significant at the 5% level. Controlling for distance, and correcting for general equilibrium forces does have a sizable impact on city size estimates.[33] Moreover, our structural estimates for city sizes are significantly related to modern city sizes, while naive estimates are not. For instance, using size

[31] Tobler and Wineburg (1971) define trade flows between i and j as the total number of joint attestations of i and j in Assyrian letters. Using that definition instead of shipment counts would confound the difference between structural and 'naive' gravity with the difference between those two alternative definitions of bilateral trade flows

[32] Further details of this comparison are presented in Appendix Table VI for comparing location of lost cities estimates, and Appendix Table VII for comparing city size estimates.

[33] In practice, most of the difference between naive and structural gravity is accounted for by our control for distance. The correlation between a simple exporter fixed effect, $\ln(\alpha_i)$ which controls for distance but not for general equilibrium, and our size measure, $\ln(Pop_i T_i^{1/\theta})$ which controls for distance and general equilibrium, is 0.98.

estimates from naive gravity in the estimating equation (15) to test for the persistence of economic activity over 4000 years gives an insignificant coefficient of logged modern population on logged ancient size (0.313 with a p-value of 0.376 to be compared to 0.230 with a p-value of 0.035 for our structural estimate), and a poor fit ($R^2 = 0.035$ to be compared to $R^2 = 0.145$ for our structural estimate). Our structural estimates for size are also significantly related to measures of access to natural roads, both for all cities together, and for the subset of lost cities only. Naive size estimates are related to access to natural roads only when considering all cities, but not for the subset of lost cities only. Our structural estimates for city sizes seem more plausible than naive estimates.

To recap, estimating a structural rather than a naive gravity model, delivers not only different, but also more reliable estimates for the location of lost cities and the sizes of ancient cities.

VII Conclusion

Business documents dating back to the Bronze Age—inscribed on 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 a measure 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 show that the relative sizes of ancient cities are explained by their 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 hierarchy of the urban system in Anatolia is the ability of cities to access natural routes, and integrate into the broader trading network. We also document a strong correlation between the estimated economic size of ancient cities and modern size measures, controlling for geographic attributes. Despite a gap of 4000 years, ancient economic size predicts the income and population of corresponding regions in present-day Turkey.

HARVARD UNIVERSITY

SCIENCES PO

UNIVERSITY OF VIRGINIA, AND CEPR

UNIVERSITY OF CHICAGO, AND NBER

References

ADIBELLI, I. A. (2013): "2012 Kırşehir Kale Hoyük Kazısı," *Kazı sonuçları toplantısı*, 35, 296–310.

BARJAMOVIC, G. (2008): "The Geography of Trade. Assyrian Colonies in Anatolia c. 1975-1725 BC and the Study of Early Interregional Networks of Exchange," in *Anatolia and the Jazira during the Old Assyrian period*, ed. by J. G. Dercksen, Leiden: Nederlands Instituut voor het Nabije Oosten, 87–100.

——— (2011): *A Historical Geography of Anatolia in the Old Assyrian Colony Period*, Copenhagen: Museum Tusculanum Press.

——— (2017): "A Commercial Geography of Anatolia: Integrating Hittite and Assyrian Texts, Archaeology and Topography," in *Hittite Landscape and Geography*, ed. by L. Ullmann and M. Weeden, Leiden: Brill., 311–318.

——— (2018): "Interlocking Commercial Networks and the Infrastructure of Trade in Western Asia during the Bronze Age," in *Trade and Civilisation: Economic Networks and Cultural Ties from Prehistory to the Early Modern Era*, ed. by K. Kristiansen, T. Lindkvist, and J. Myrdal, Cambridge: Cambridge University Press, 113–142.

BARJAMOVIC, G. J., T. K. HERTEL, AND M. T. LARSEN (2012): *Ups and Downs at Kanesh: chronology, history and society in the Old Assyrian period*, Leiden: Nederlands Instituut voor het Nabije Oosten.

BERNARD, A. B., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and productivity in international trade," *American economic review*, 93, 1268–1290.

BLEAKLEY, H. AND J. LIN (2012): "Portage and path dependence," *The Quarterly Journal of Economics*, 127, 587–644.

CAMMAROSANO, M. AND M. MARIZZA (2015): "Das Land Tumanna und sein König in den hethitischen Quellen," *Die Welt des Orients*, 45, 158–192.

CHANEY, T. (2018): "The gravity equation in international trade: An explanation," *Journal of Political Economy*, 126, 150–177.

CORNELIUS, F. (1963): "Neue Aufschlüsse zur hethitischen Geographie," *Orientalia*, 32, 233–245.

CORTI, C. (2017): "The North: Hanhana, Hattena, Ištahara, Hakpiš, Nerik, Zalpuwa, Tummana, Pala and the Hulana River Land," in *Hittite Landscape and Geography*, ed. by L. Ullmann and M. Weeden, Leiden: Brill, 219–238.

Cosar, A. K. and B. Demir (2016): "Domestic road infrastructure and international trade: Evidence from Turkey," *Journal of Development Economics*, 118, 232 – 244.

Cronon, W. (2009): *Nature's Metropolis: Chicago and the Great West*, WW Norton & Company.

Dalgaard, C.-J., N. Kaarsen, O. Olsson, and P. Selaya (2018): "Roman Roads to Prosperity: Persistence and Non-Persistence of Public Goods Provision," *CEPR Discussion Paper DP12745*.

Davis, D. R. and D. E. Weinstein (2002): "Bones, bombs, and break points: the geography of economic activity," *American Economic Review*, 1269–1289.

Dercksen, J. G. (1996): *The Old Assyrian Copper Trade in Anatolia*, Leiden: Nederlands Instituut voor het Nabije Oosten.

Dijkstra, E. W. (1959): "A note on two problems in connexion with graphs," *Numerische mathematik*, 1, 269–271.

Disdier, A.-C. and K. Head (2008): "The puzzling persistence of the distance effect on bilateral trade," *The Review of Economics and statistics*, 90, 37–48.

Dönmez, c. (2017): "The Contribution of New Research to Hittite Historical Geography of Amasya Province," in *Places and Spaces in Hittite Anatolia I: Hatti and the East. Proceedings of an International Workshop on Hittite Historical Geography in Istanbul, 25th-26th October 2013*, ed. by M. Alparslan, Istanbul: Türk Eskiçağ Bilimleri Enstitüsü, 75–90.

Eaton, J. and S. Kortum (2002): "Technology, Geography and Trade," *Econometrica*, 70, 1741–79.

Eaton, J., S. Kortum, and S. Sotelo (2012): "International Trade: Linking Micro and Macro," *NBER Working Paper No.17864*.

Erol, H. (2013): "ina maqāt ubrī lu kaspam lu hurāṣam ašamma ... 'Yabancilar geldiğinde ya gümüş ya da altın satın alacağim ...'," in *Cahit Günbattı'ya Armağan. Studies in Honour of Cahit Günbattı*, ed. by I. Albayrak, H. Erol, and M. Çayır, 87–98.

Forlanini, M. (2008): "The Central Provinces of Hatti. An Updating," in *New Perspectives on the Historical Geography and Topography of Anatolia in the II and I Millennium BC*, ed. by K. Strobel, (EOTHEN 16) Firenze: LoGisma Editore, 1, 145–188.

——— (2017): "South Central: The Lower Land and Tarhuntašša," in *Hittite Landscape and Geography*, ed. by L. Ullmann and M. Weeden, Leiden: Brill., 239–252.

French, D. (1993): "Colonia Archelais and Road-Knots," in *Aspects of Art and Iconography: Anatolia and Its Neighbors Studies in Honor of Nimet Özgüç*, ed. by M. J. Mellink, E. Porada, and T. Özgüç, Ankara: Türk Tarih Kurumu Basimevi, 201–7.

——— (2016): *Roman Roads and Milestones of Asia Minor. Vol. 3 Milestones, Fasc. 9. An Album of Maps, Electronic Monograph*, Ankara: British Institute at Ankara.

Gabaix, X. and R. Ibragimov (2011): "Rank- 1/2: a simple way to improve the OLS estimation of tail exponents," *Journal of Business & Economic Statistics*, 29, 24–39.

Galloway, P. (1978): "Restoring the map of medieval Trondheim: a computer-aided investigation into the Nightwatchmen's itinerary," *Journal of Archaeological Science*, 5, 153–165.

Gardin, J.-C. and P. Garelli (1961): "Étude des Établissements assyriens en Cappadoce par ordinateurs," *Annales. Histoire, Sciences Sociales*, 16, 837–876.

Günbatti, C. (1998): "Karumlar Arasındaki Mektuplaşmalardan Yeni Örnekler," in *XXXIV. Uluslararası Assiriyoloji Kongresi, 6-10/VII/1987-Istanbul*, ed. by H. Erkanal, V. Donbaz, and A. Uğuroğlu, Ankara: Türk Tarih Kurumu Basımevi.

Hertel, T. K. (2013): *Old Assyrian Legal Practices: Law and Dispute in the Ancient Near East*, Leiden: Nederlands Instituut voor het Nabije Oosten.

Hodler, R. and P. A. Raschky (2014): "Regional Favoritism," *The Quarterly Journal of Economics*, 129, 995–1033.

IIASA/FAO (2012): *Global Agro-Ecological Zones (GAEZ v3.0)*, IIASA, Laxenburg, Austria and FAO, Rome.

Kendall, D. G. (1971): "Maps from marriages: an application of non-metric multi-dimensional scaling to parish register data," *Mathematics in the archaeological and historical sciences*, 303–318.

Krugman, P. (1991): "Increasing returns and economic geography," *Journal of political economy*, 99, 483–499.

Kryszeń, A. (2016): *A Historical Geography of the Hittite Hinterland. (Alter Orient und Alter Testament 437)*, Münster: Ugarit-Verlag.

Langmuir, E. (1984): *Mountaincraft and leadership: a handbook for mountaineers and hillwalking leaders in the British Isles*, Edinburgh: Scottish Sports Council.

Larsen, M. T. (2015): *Ancient Kanesh: A Merchant Colony in Bronze Age Anatolia*, Cambridge University Press.

Lassen, A. W. (2010): "The Trade in Wool in Old Assyrian Anatolia," *Jaarbericht Ex Oriente Lux*, 42, 159–179.

——— (2014): "The Old Assyrian Glyptic Style: An Investigation of a Seal Style, Its Owners and Place of Production," in *Current Research in Kültepe/Kanesh: An Interdisciplinary and Integrative Approach to Trade Networks, Internationalism, and Identity during the Middle Bronze Age (Journal of Cuneiform Studies, Supplemental Series 4)*, ed. by L. Atici, F. Kulakoğlu, G. Barjamovic, and A. Fairbairn, The American Schools of Oriental Research: Bristol CT, 107–121.

Manning, S. W., G. Barjamovic, and B. Lorenzen (2017): "The Course of C14 Dating Does Not Run Smooth: Tree-Rings, Radiocarbon, and Potential Impacts of a Calibration Curve Wiggle on Dating Mesopotamian Chronology," *Journal of Ancient Egyptian Interconnections*.

Massa, M. (2016): *Networks before Empires: cultural transfers in west and central Anatolia during the Early Bronze Age*, Unpublished PhD dissertation, University College London.

Matthews, R. and C. Glatz (2009): "The historical geography of north-central Anatolia in the Hittite period: texts and archaeology in concert," *Anatolian Studies*, 59, 51–72.

Michaels, G. and F. Rauch (2016): “Resetting the Urban Network: 117-2012,” *The Economic Journal*.

Omura, M. (2016): “2015 Yassıhöyük kazıları,” *Kazı sonuçları toplantısı*, 38, 343–352.

Palmisano, A. (2013): “Computational and Spatial Approaches to the Commercial Landscapes and Political Geography of the Old Assyrian Colony Period.” in *Time and History in the Ancient Near East. Proceedings of the 56th Rencontre Assyriologique Internationale, Barcelona, July 26-30, 2010.*, ed. by L. Feliu, J. Llop, A. M. Albà, and W. Lake., Eisenbrauns, 767–783.

——— (2017): “Drawing Pathways from the Past: the Trade Routes of the Old Assyrian Caravans Across Upper Mesopotamia and Central Anatolia,” in *Movement, Resources, Interaction. Proceedings of the 2st Kültepe International Meeting. Kültepe, July 26-30, 2015. Studies Dedicated to Klaas Veenhof. Kültepe International Meetings 2 (SUBARTU 39)*, ed. by F. Kulakoğlu and G. Barjamovic, Turnhout: Brepols, 29-48.

Palmisano, A. and M. Altaweel (2015): “Landscapes of interaction and conflict in the Middle Bronze Age: From the open plain of the Khabur Triangle to the mountainous inland of Central Anatolia,” *Journal of Archaeological Science: Reports*, 3, 216–236.

Ramsay, W. M. (1890): *The Historical Geography of Asia Minor*, London: John Murray.

Riley, S. J., S. D. DeGloria, and R. Elliot (1999): “Index that quantifies topographic heterogeneity,” *Intermountain Journal of sciences*, 5, 23–27.

Rosen, K. T. and M. Resnick (1980): “The size distribution of cities: an examination of the Pareto law and primacy,” *Journal of Urban Economics*, 8, 165–186.

Silva, J. S. and S. Tenreyro (2006): “The Log of Gravity,” *The Review of Economics and statistics*, 88, 641–658.

Simonovska, I. and M. E. Waugh (2014): “The Elasticity of Trade: Estimates and Evidence,” *Journal of International Economics*, 92, 34 – 50.

Tobler, W. and S. Wineburg (1971): “A Cappadocian speculation,” *Nature*, 231, 39–41.

Weeden, M. and K. Matsumura (2017): “Central West: Archaeology,” in *Hittite Landscape and Geography*, ed. by L. Ullmann and M. Weeden, Leiden: Brill, 106–118.

Tables

Table I: **Descriptive Statistics**

	Mean	St. Dev.	Min	Max	Observations
Known cities					15
Lost cities					10
Number of unique tablets					198
Number of itineraries					227
Number of travels					391
N_{ij}^{data} (all $i \to j$ pairs for $i \neq j$)	0.63	2.29	0	23	600
$N_{ij}^{data} > 0$ only	3.33	4.31	1	23	114
$Distance_{ij}$ in km (i and j known)	223	113	17	576	105

Notes: The ancient data comes from a textual analysis of clay tablets inscribed in the cuneiform script, written by Assyrian merchants in the 2nd millennium BCE. Most texts are digitized and will be available as tagged and searchable files through the OARE-project, currently being built as part of the Neubauer Project.

Table II: **Lost Cities' Geo-coordinates**

	Latitude	(s.e.)	Longitude	(s.e.)	Correlation
Durhumit	40.47	(0.025)	35.65	(0.445)	-0.952
Hahhum	38.429	(0.274)	38.04	(0.517)	0.68
Kuburnat	40.712	(0.582)	36.52	(0.512)	-0.06
Ninassa	38.977	(0.778)	34.614	(0.482)	0.86
Purushaddum	39.71	(1.54)	32.872	(0.669)	0.774
Sinahuttum	39.956	(0.333)	34.866	(0.165)	0.863
Suppiluliya	40.021	(1022.82)	34.618	(58.796)	1.0
Tuhpiya	39.611	(0.18)	35.199	(0.307)	0.528
Washaniya	39.157	(0.219)	34.311	(0.265)	-0.01
Zalpa	38.805	(0.648)	37.862	(1.199)	0.878

Notes: This table presents the estimated geo-coordinates, latitudes and longitudes, from solving our structural gravity model (8). All latitudes are North, and all longitudes are East. Robust (White) standard errors are in parentheses, as well as the estimated correlation between latitude and longitude.

Table III: **Gravity Estimation Results**

Panel A: Distance elasticity and statistics

ζ (dist. elast.)	1.912
	(0.189)
Observations	600

Panel B: Sizes and locations of lost cities

				Distance to historians' proposals, in km			
	$Pop_i T_i^{1/\theta}$	(s.e.)	Precision	Forlanini	(s.e.)	Barjamovic	(s.e.)
Durhumit	0.174	(0.409)	49	220	(46)	48	(26)
Hahhum	64.556	(85.929)	62	102	(27)	102	(27)
Kuburnat	11.22	(24.316)	76	72	(46)	70	(41)
Ninassa	0.21	(0.45)	87	71	(47)	93	(49)
Purushaddum	0.076	(0.1)	154	168	(75)	193	(112)
Sinahuttum	1.515	(2.021)	34	24	(20)	24	(20)
Suppiluliya	0.012	(5.246)	89914	89	(54307)	85	(54240)
Tuhpiya	0.579	(0.912)	38	128	(35)	112	(33)
Washaniya	5.413	(6.462)	35	68	(25)	13	(19)
Zalpa	28.695	(60.689)	145	103	(75)	131	(70)

Panel C: Sizes of known cities

	$Pop_i T_i^{1/\theta}$	(s.e.)
Hanaknak	2.062	(3.679)
Hattus	2.967	(4.183)
Hurama	5.091	(9.564)
Kanes	100.0	(70.902)
Karahna	0.008	(0.084)
Malitta	0.057	(0.096)
Mamma	114.683	(189.19)
Salatuwar	1.513	(4.836)
Samuha	3.508	(4.847)
Tapaggas	0.091	(0.218)
Timelkiya	101.922	(129.238)
Ulama	0.099	(0.498)
Unipsum	44.294	(64.469)
Wahsusana	1.416	(3.53)
Zimishuna	0.0	(0.001)

Notes: This table presents the results from estimating our structural gravity model from equations (6), (8) and (11) using directional data, N_{ij}^{data}. Our measure of fundamental city size, $Pop_i T_i^{1/\theta}$, defined in (10), is the counterfactual real output of city i if it were to move to complete autarky. Precision, measured in km, is defined in (12). Distance to historians' proposals measures the distance, in km, between our point estimate and the conjecture by historians Forlanini (2008) and Barjamovic (2011). Robust (White) standard errors in parentheses.

Table IV: **Determinants of Ancient City Sizes**

	$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$				
	(1)	(2)	(3)	(4)	(5)
$\log(NaturalRoads)$	1.404**			1.783***	
	(0.013)			(0.002)	
$\log(RomanRoads)$		1.990			2.387
		(0.378)			(0.220)
$\log(Ruggedness)$			2.371**	3.189***	2.495***
			(0.012)	(0.000)	(0.003)
N	25	25	25	25	25
R^2	0.224	0.038	0.166	0.508	0.220

Notes: This table presents the estimation from various specifications of (14). Each observation is an ancient city. The dependent variable $PopT^{1/\theta}|_{ancient}$ is the ancient city size estimate. Explanatory variables start $NaturalRoads$, the average natural road score of the area within 20 km of the ancient city as defined in subsection V.B. *RomanRoads* is the number of Roman roads radiating from an intersection within 20 km of the ancient city (French, 2016). *Ruggedness* is the Terrain Ruggedness Index (Riley, DeGloria, and Elliot, 1999). In unreported regressions, we experimented with alternative measures of local amenities (elevation, crop yield, distance to the nearest river, and distance to the nearest known copper, gold or silver deposit documented in the Early Bronze Age), but none of those measures were either significant or robust. Robust p-values are in parentheses.

Table V: **Persistence of Economic Activity across 4000 Years**

	$\log(Population)$			$\log(NightLights)$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.230**		0.297**	0.124**		0.178***
	(0.035)		(0.015)	(0.036)		(0.008)
$\log(CropYield)$		0.727	1.781*		0.777	1.407***
		(0.507)	(0.079)		(0.143)	(0.003)
N	24	24	24	24	24	24
R^2	0.145	0.015	0.226	0.143	0.059	0.312

Notes: This table presents the estimation from various specifications of (15). Each observation is an ancient city after dropping *Purušhaddum-Ankara* from the sample. Dependent variables are modern-day size measures: *Population* and $NightLights$ are total urban population and night luminosity within 20 km of the ancient city, respectively. Explanatory variables are $PopT^{1/\theta}|_{ancient}$, the ancient city size estimate and $CropYield$, the average rain-fed low-input cereal suitability index of the area within 20 km of the ancient city. In unreported regressions, we also experimented with other geographic controls (elevation, distance to the nearest river, and distance to modern mineral deposits of gold, silver and copper), but none of those measures were either significant or robust. Robust p-values are in parentheses.

Figures

Figure I: **Tablet Kt 83-k 117**

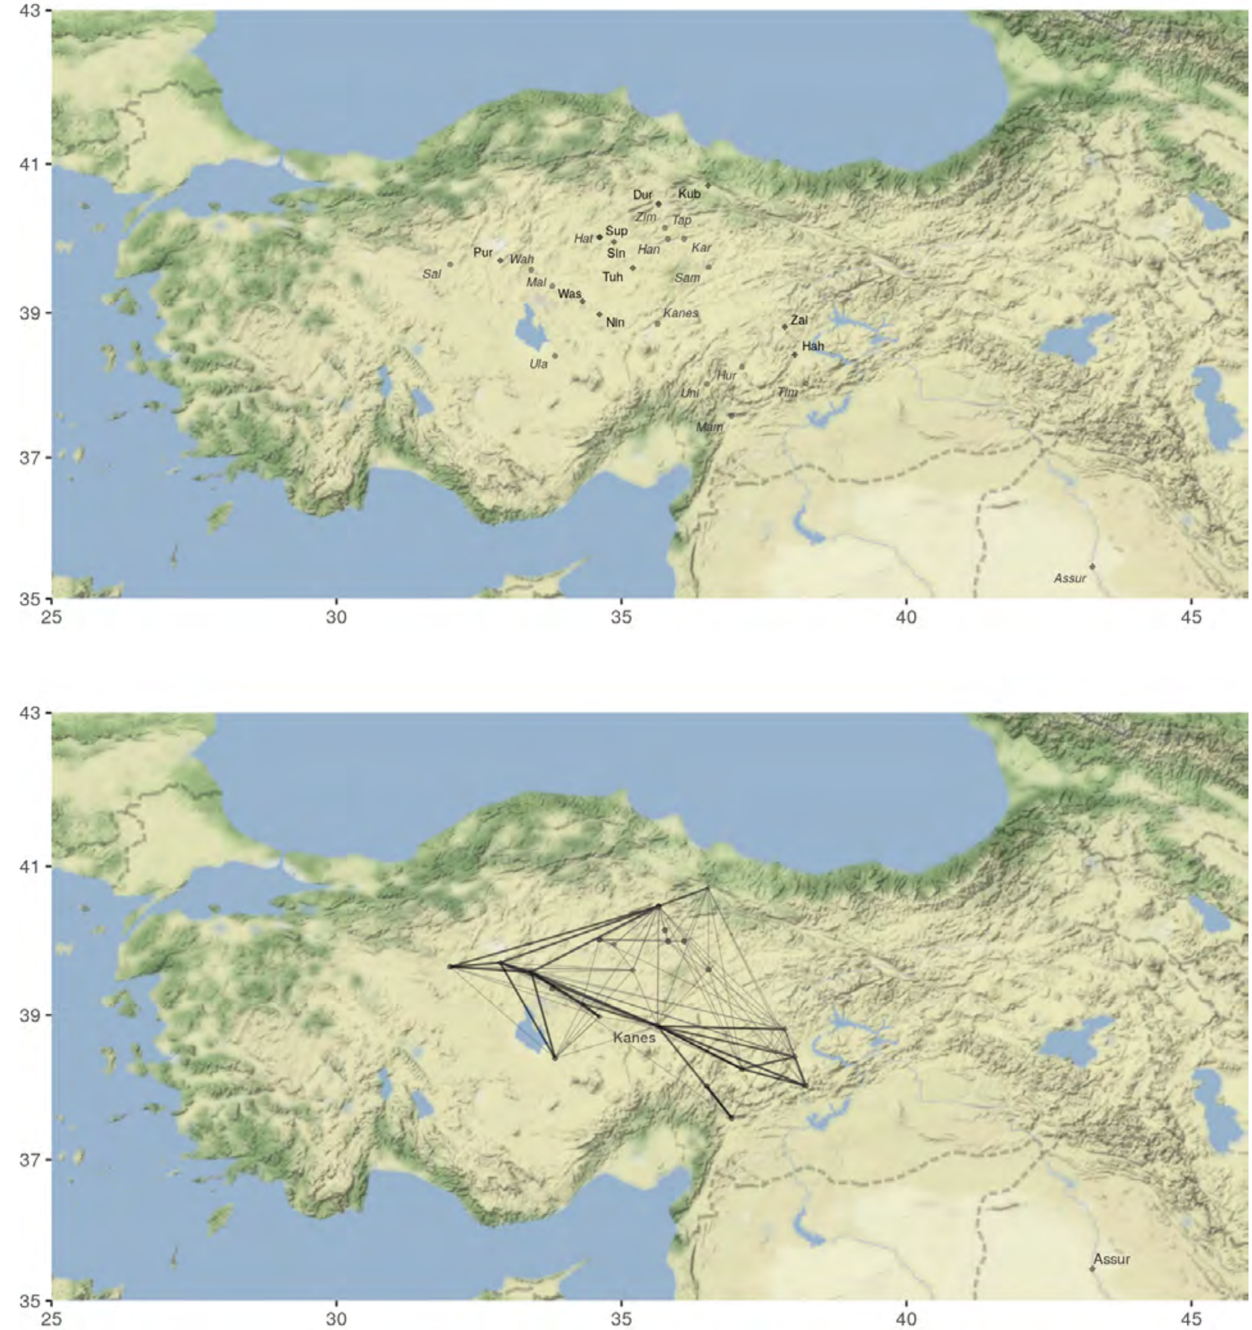

Figure II: **Cities and Trade in Anatolia in the Bronze Age.**

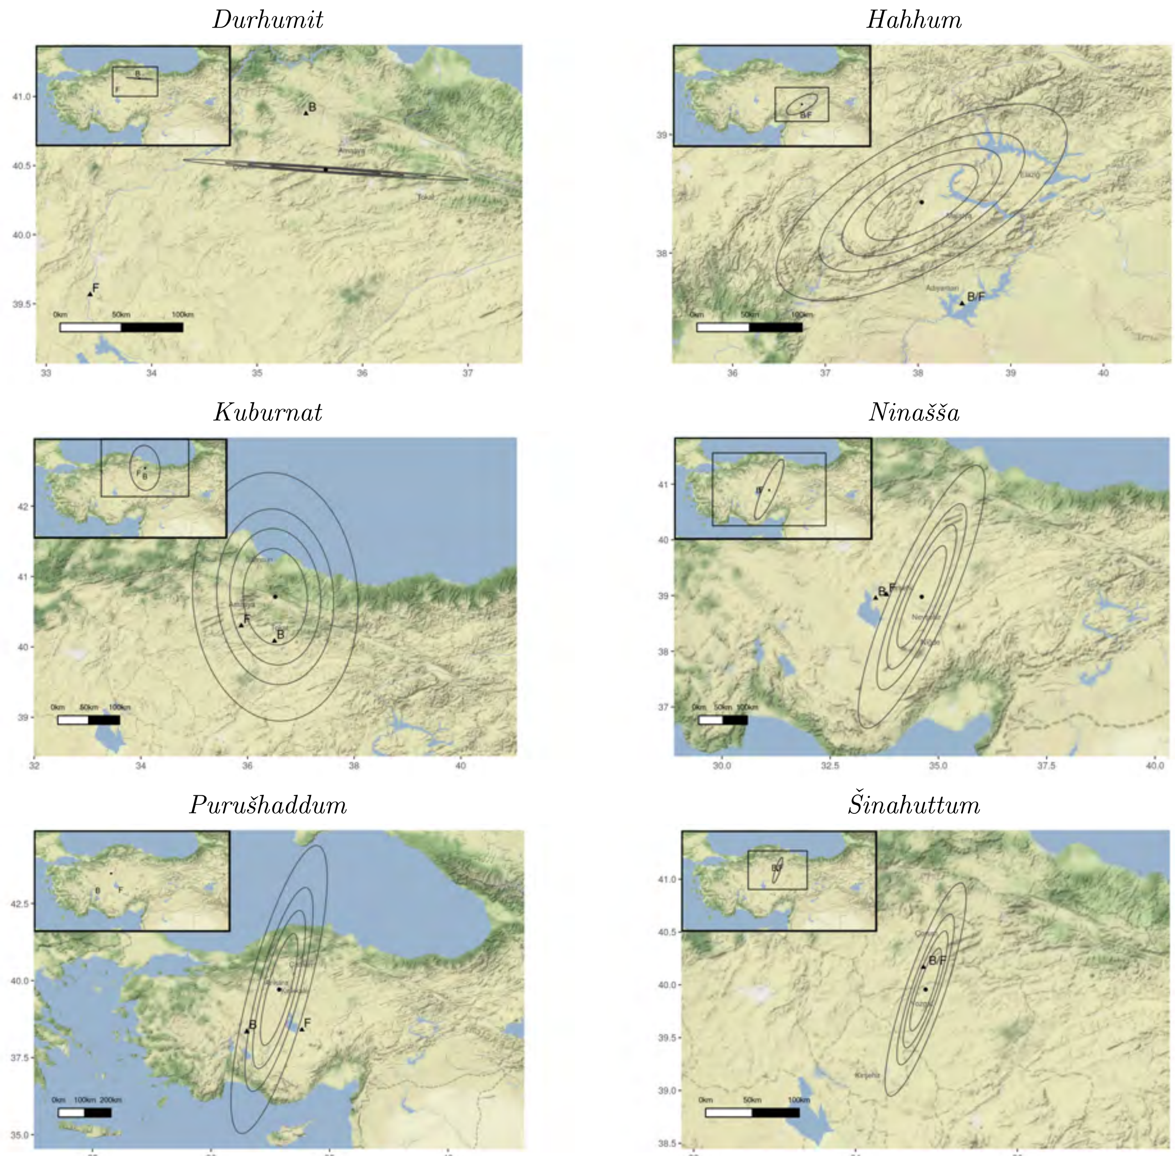

Figure III: **Locating Lost Cities (I).**

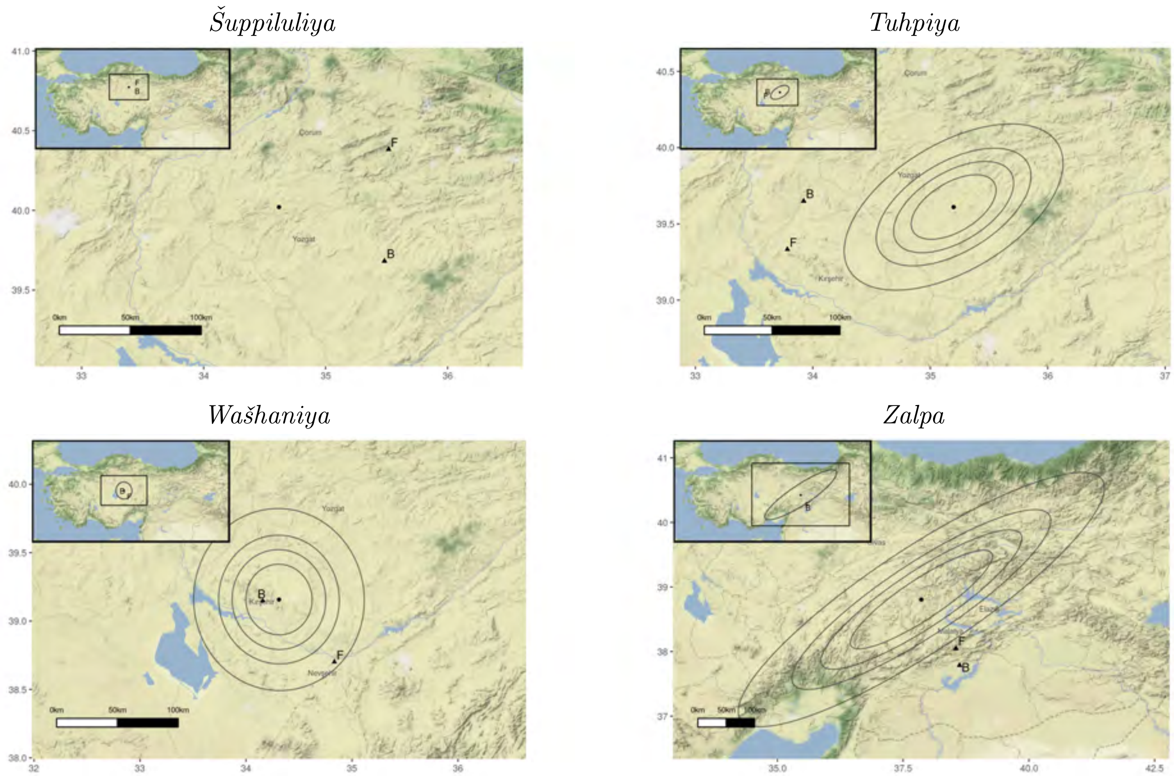

Figure IV: **Locating Lost Cities (II).**

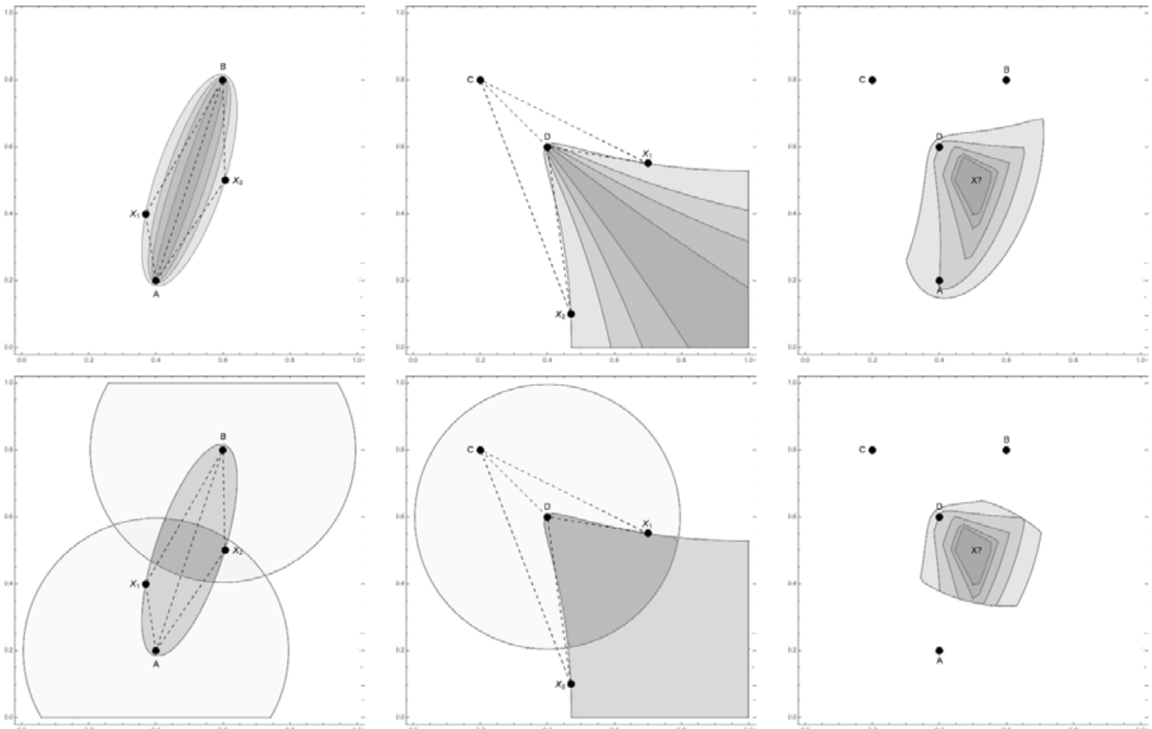

Itinerary $A \to X \to B$ | *Itinerary* $C \to D \to X$ | *Multiple Itineraries*

Figure V: **Constraints on Lost Cities from Merchants' Itineraries, Example.**

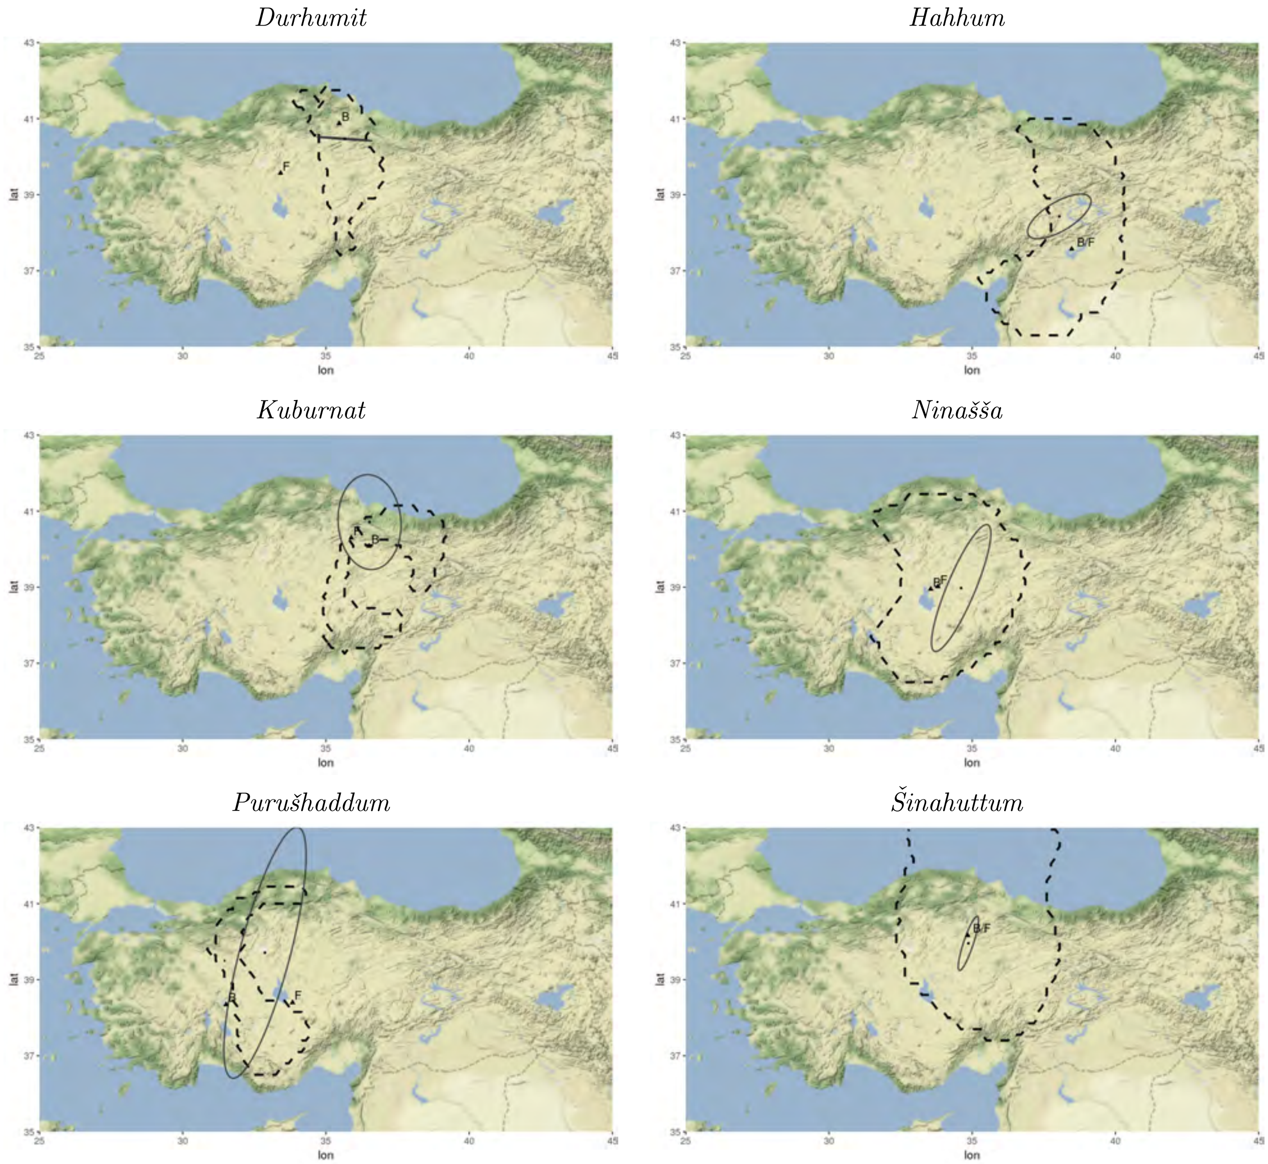

Figure VI: **Constraints on Lost Cities from Merchants' Itineraries (I).**

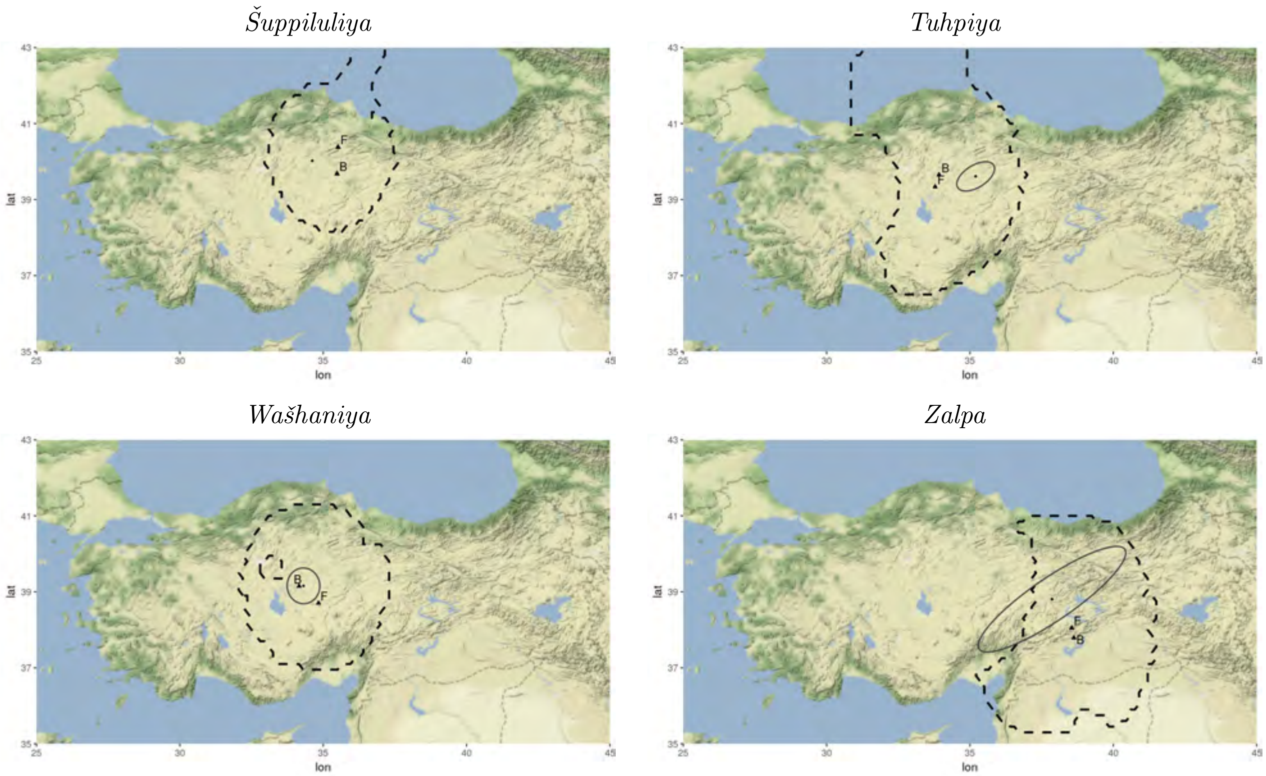

Figure VII: **Constraints on Lost Cities from Merchants' Itineraries (II).**

Figure VIII: **Proof of Concept: Recovering Fictitiously Lost Cities.**

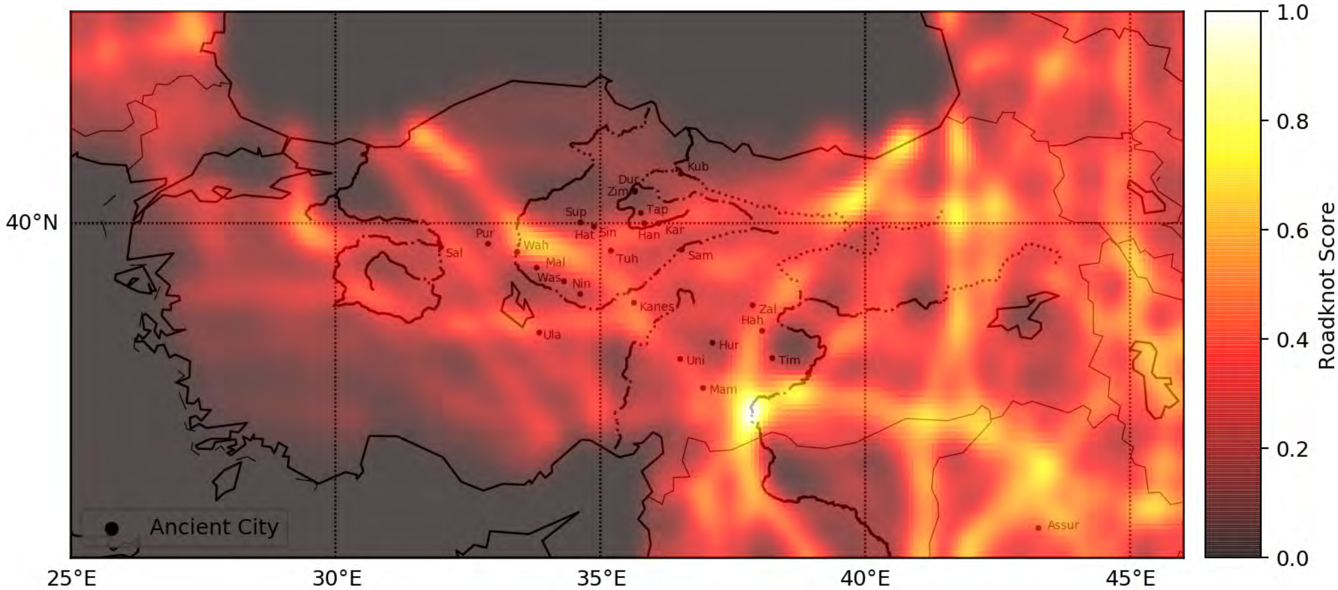

Figure IX: **Natural Roads Scores.**

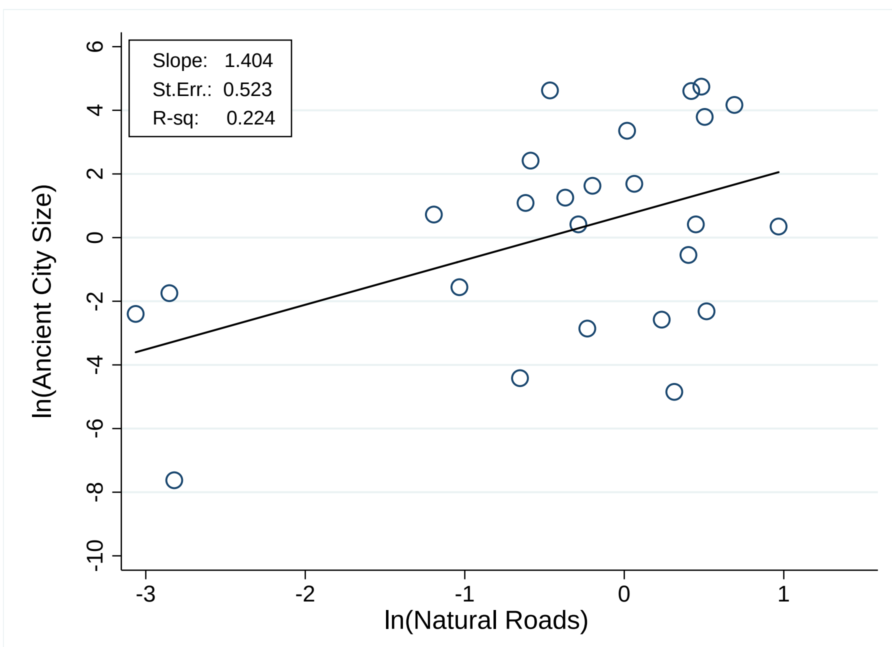

Figure X: **The Topographical Determinants of Ancient City Sizes**

Panel A: Population

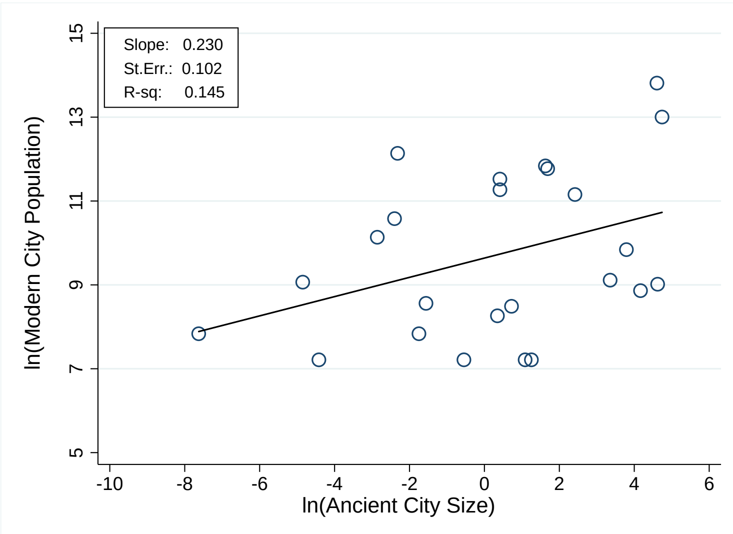

Panel B: Night Lights

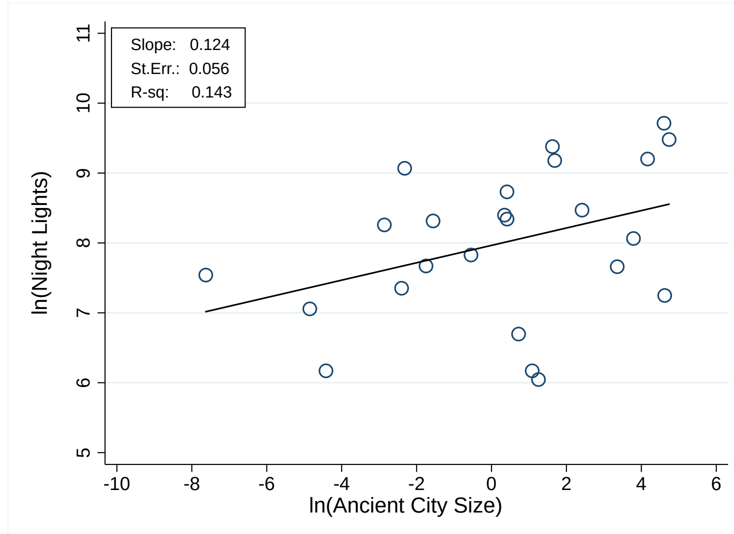

Figure XI: **Ancient and Modern City Sizes**

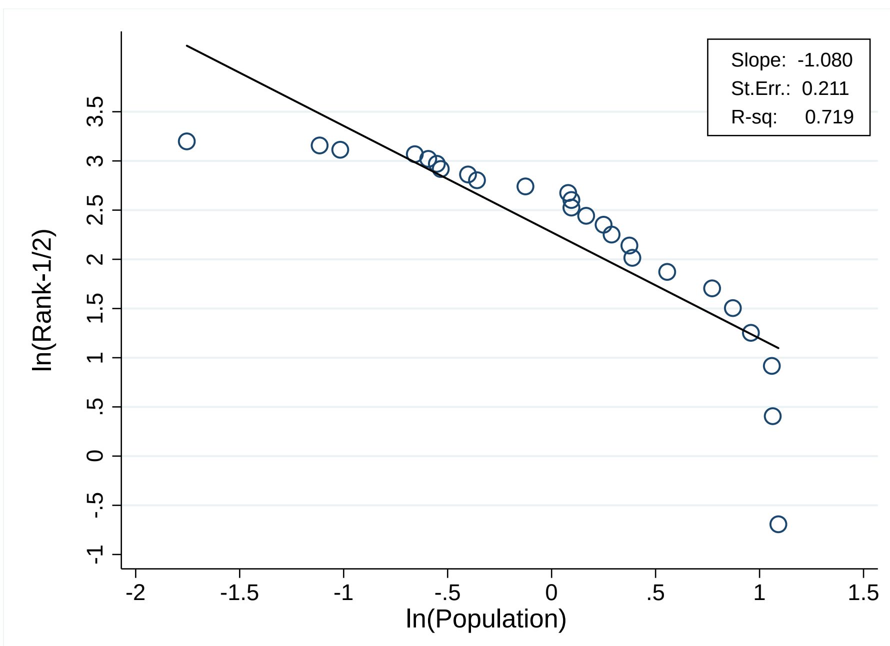

Figure XII: **Zipf's Law for Cities in the Bronze Age.**

Online Appendix for:

Trade, Merchants, and the Lost Cities of the Bronze Age

(not for publication)

by Gojko Barjamovic, Thomas Chaney, Kerem Coşar, and Ali Hortaçsu

This appendix contains additional information and robustness checks. It is organized as follows. Section A contains all mathematical proofs. Section B presents detailed information about our data sources and optimization procedure. Section C contains detailed instructions for the construction of optimal travel routes using only topographical data. Section D offers detailed instructions for replicating our analysis of the information contained in merchants' multi-stops itineraries. Section E gives a detailed description of how we construct our *NaturalRoads* measure. Section F contains additional tables.

A Mathematical Proofs and Optimization

Derivation of shipment probabilities, equations (2), (3), and (4). Equation (2) is the unconditional probability that origin city i is the cheapest source for good ω in destination city j. Given the assumption of Weibull distributed costs in (1), the probability distribution for the cost of delivering good ω from origin i to destination j is also Weibull,

$$\begin{aligned} G_{ij}(c) &= \Pr\left[c_{ij}(\omega) \leq c\right] \\ &= \Pr\left[\tau_{ij} c_i(\omega) \leq c\right] \\ &= 1 - \exp\left(-T_i \left(w_i \tau_{ij}\right)^{-\theta} c^{\theta}\right). \end{aligned}$$

Equation (2) is then derived exactly as in Eaton and Kortum (2002),

$$\begin{aligned} \Pr\left[c_{ij}(\omega) \leq \min_k \{c_{kj}(\omega)\}\right] &= \Pr\left[c_{ij}(\omega) \leq \min_{k \neq i} \{c_{kj}(\omega)\}\right] \\ &= \int_0^{\infty} \Pi_{k \neq i} \left(1 - G_{kj}(c)\right) dG_{ij}(c). \end{aligned}$$

We use the c.d.f. $G_{kj}(c) = 1 - \exp\left(-T_k \left(\tau_{kj} w_k\right)^{-\theta} c^{\theta}\right)$ and the corresponding p.d.f. $dG_{ij}(c) = \theta T_i \left(\tau_{ij} w_i\right)^{-\theta} c^{\theta-1} \exp\left(-T_i \left(\tau_{ij} w_i\right)^{-\theta} c^{\theta}\right) dc$ to get,

$$\begin{aligned} \Pr\left[c_{ij}(\omega) \leq \min_k \{c_{kj}(\omega)\}\right] &= T_i \left(\tau_{ij} w_i\right)^{-\theta} \int_0^{\infty} \Pi_k \exp\left(-T_k \left(\tau_{kj} w_k\right)^{-\theta} c^{\theta}\right) \theta c^{\theta-1} dc \\ &= T_i \left(\tau_{ij} w_i\right)^{-\theta} \int_0^{\infty} \exp\left(-\left(\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}\right) c^{\theta}\right) \theta c^{\theta-1} dc \\ &= T_i \left(\tau_{ij} w_i\right)^{-\theta} \left[\frac{-\exp\left(-\left(\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}\right) c^{\theta}\right)}{\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}}\right]_0^{\infty} \\ &= \frac{T_i \left(\tau_{ij} w_i\right)^{-\theta}}{\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}}. \square \end{aligned}$$

Equation (3) is the probability that origin i is the cheapest source for good ω in destination j, conditional on j not sourcing good ω internally,

$$\Pr\left[c_{ij}(\omega) \leq \min_{k \neq j} \{c_{kj}(\omega)\} \,|\, c_{jj}(\omega) > \min_{k \neq j} \{c_{kj}(\omega)\}\right] = \frac{\Pr\left[c_{ij}(\omega) \leq \min_k \{c_{kj}(\omega)\}\right]}{\Pr\left[c_{jj}(\omega) > \min_{k \neq j} \{c_{kj}(\omega)\}\right]}.$$

The numerator is given above (derivation of (2)). For the denominator, following similar derivations,

$$\begin{aligned} \Pr\left[c_{jj}(\omega) > \min_{k \neq j} \{c_{kj}(\omega)\}\right] &= 1 - \Pr\left[c_{jj}(\omega) \leq \min_{k \neq j} \{c_{kj}(\omega)\}\right] \\ &= 1 - \frac{T_j \left(\tau_{jj} w_j\right)^{-\theta}}{\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}} \\ &= \frac{\sum_{k \neq j} T_k \left(\tau_{kj} w_k\right)^{-\theta}}{\sum_k T_k \left(\tau_{kj} w_k\right)^{-\theta}}. \end{aligned}$$

Taking the ratio to form a conditional probability, we have the proposed,

$$\Pr\left[c_{ij}(\omega) \leq \min_{k\neq j}\{c_{kj}(\omega)\} | c_{jj}(\omega) > \min_{k\neq j}\{c_{kj}(\omega)\}\right] = \frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\neq j} T_k(\tau_{kj}w_k)^{-\theta}}. \square$$

Equation (4) is the probability that destination j imports a given good ω from origin i conditional on (a) j not importing from any lost city and (b) j not purchasing good ω internally. Conditions (a) and (b) are satisfied if and only if $\min_{l\in\mathcal{L}\cup\{j\}}\{c_{lj}(\omega)\} > \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\}$. We first characterize the distribution of $\min_{l\in\mathcal{L}\cup\{j\}}\{c_{lj}(\omega)\}$,

$$\begin{aligned}
\Pr\left[\min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) \leq c\right] &= 1 - \Pr\left[\min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > c\right] \\
&= 1 - \Pi_{l\in\mathcal{L}\cup\{j\}} \Pr[c_{lj}(\omega) > c] \\
&= 1 - \Pi_{l\in\mathcal{L}\cup\{j\}} (1 - \Pr[c_{lj}(\omega) \leq c]) \\
&= 1 - \Pi_{l\in\mathcal{L}\cup\{j\}} \exp\left(-T_l(w_l\tau_{lj})^{-\theta} c^\theta\right) \\
&= 1 - \exp\left(-\left(\sum_{l\in\mathcal{L}\cup\{j\}} T_l(w_l\tau_{lj})^{-\theta}\right) c^\theta\right),
\end{aligned}$$

i.e. a Weibull distribution with shape parameter $\sum_{l\in\mathcal{L}\cup\{j\}} T_l(w_l\tau_{lj})^{-\theta}$. Given this distribution, we can easily form the conditional probability (4), following the same steps as above,

$$\begin{aligned}
&\Pr\left[c_{ij}(\omega) \leq \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\} \middle| \min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\}\right] \\
&= \frac{\Pr[c_{ij}(\omega) \leq \min_{k\in\mathcal{K}\cup\mathcal{L}}\{c_{kj}(\omega)\}]}{\Pr\left[\min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > \min_{k\in\mathcal{K}\setminus\{j\}}\{c_{kj}(\omega)\}\right]} = \frac{\Pr[c_{ij}(\omega) \leq \min_{k\in\mathcal{K}\cup\mathcal{L}}\{c_{kj}(\omega)\}]}{\Pr\left[\min_{l\in\mathcal{L}\cup\{j\}} c_{lj}(\omega) > \min_{k\in\mathcal{K}\cup\mathcal{L}}\{c_{kj}(\omega)\}\right]} \\
&= \frac{\frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\in\mathcal{K}\cup\mathcal{L}} T_k(\tau_{kj}w_k)^{-\theta}}}{1 - \frac{\sum_{l\in\mathcal{L}\cup\{j\}} T_l(\tau_{lj}w_l)^{-\theta}}{\sum_{k\in\mathcal{K}\cup\mathcal{L}} T_k(\tau_{kj}w_k)^{-\theta}}} = \frac{\frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\in\mathcal{K}\cup\mathcal{L}} T_k(\tau_{kj}w_k)^{-\theta}}}{\frac{\sum_{k\in\mathcal{K}\setminus\{j\}} T_k(\tau_{kj}w_k)^{-\theta}}{\sum_{k\in\mathcal{K}\cup\mathcal{L}} T_k(\tau_{kj}w_k)^{-\theta}}} = \frac{T_i(\tau_{ij}w_i)^{-\theta}}{\sum_{k\in\mathcal{K}\setminus\{j\}} T_k(\tau_{kj}w_k)^{-\theta}}. \square
\end{aligned}$$

Derivation of moment condition (7) with multiplicative disturbance term, footnote 10. We follow a lightly edited version of Eaton, Kortum, and Sotelo (2012), and add to the trade cost function a multiplicative disturbance term drawn from a Gamma distribution,

$$\tau_{ij}^{-\theta} = \mu Distance_{ij}^{-\zeta}\nu_{ij}, \text{ with } \nu_{ij} \sim Gamma\left(\frac{1}{\eta^2}\frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}, \frac{\eta^2}{\alpha_i Distance_{ij}^{-\zeta}}\right).$$

Treating the ν's as realizations from a random variable, we rely on the scaling property of the Gamma distribution to obtain

$$\alpha_i Distance_{ij}^{-\zeta}\nu_{ij} \sim Gamma\left(\frac{1}{\eta^2}\frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}, \eta^2\right).$$

The joint distribution of n Gamma distributed variables normalized by their sum is Dirichlet,

$$\left(\cdots, \frac{\alpha_i Distance_{ij}^{-\zeta}\nu_{ij}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}\nu_{kj}}, \cdots\right)_{i\neq j} \sim Dirichlet\left(\cdots, \frac{1}{\eta^2}\frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}, \cdots\right).$$

Note that by definition of the Dirichlet joint distribution, all shares are in $(0,1)$ and they add up to one. Using the definition for the mean of a Dirichlet distribution, we recover our proposed moment condition (7),

$$\mathbb{E}\left[\frac{\alpha_i Distance_{ij}^{-\zeta}\nu_{ij}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}\nu_{kj}}\right] = \frac{\frac{1}{\eta^2}\frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}}{\sum_{l\neq j}\frac{1}{\eta^2}\frac{\alpha_l Distance_{lj}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}} = \frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j}\alpha_k Distance_{kj}^{-\zeta}}. \square$$

Derivation of our measure of size, equation (10). To express $Size_i \propto Pop_i T_i^{1/\theta}$ as a function of observables and model parameters only, we first use the definition our exporter fixed effect α_i estimated from (8),

$$\alpha_i \propto T_i w_i^{-\theta} \Rightarrow Pop_i T_i^{1/\theta} \propto \alpha_i^{1/\theta} w_i Pop_i.$$

From market clearing,

$$w_i Pop_i = X_i \Rightarrow Pop_i T_i^{1/\theta} \propto \alpha_i^{1/\theta} X_i.$$

The volume of trade from i to j is simply equal to total expenditure in j multiplied by the probability of sourcing a good from origin i.

$$X_{ij} = \frac{T_i w_i^{-\theta} Distance_{ij}^{-\theta} X_j}{\sum_k T_k w_k^{-\theta} Distance_{kj}^{-\theta}}.$$

We then manipulate this expression as in Anderson and van Wincoop (2003) to obtain

$$X_{ij} = \frac{T_i w_i^{-\theta}\tau_{ij}^{-\theta} X_j}{\sum_k T_k w_k^{-\theta}\tau_{kj}^{-\theta}} = \frac{X_i X_j}{X_{total}}\left(\frac{\tau_{ij}}{\Pi_i P_j}\right)^{-\theta},$$

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, $X_{ij} = X_{ji}$. Using the equivalence between trade shares in value and in count in the Eaton and Kortum (2002) model,

$$\frac{X_{ij}}{X_j} = \frac{T_i w_i^{-\theta}\tau_{ij}^{-\theta}}{\sum_k T_k w_k^{-\theta}\tau_{kj}^{-\theta}} = \mathbb{E}\left[\frac{N_{ij}}{\sum_k N_{kj}}\right] = \frac{\alpha_i\tau_{ij}^{-\theta}}{\sum_k \alpha_k\tau_{kj}^{-\theta}}.$$

Combining this with the above expression for bilateral trade, we get

$$\frac{X_{ij}}{X_j} = \frac{X_i}{X_{total}}\left(\frac{\tau_{ij}}{\Pi_i P_j}\right)^{-\theta} = \frac{\alpha_i\tau_{ij}^{-\theta}}{\sum_k \alpha_k\tau_{kj}^{-\theta}}, \forall i \neq j \Rightarrow X_i \propto \alpha_i \Pi_i^{-\theta}.$$

From the above, the definition of $\Pi_i^{-\theta}$, and symmetry, $P_k = \Pi_k$, we derive

$$\Pi_i^{-\theta} \propto \sum_k \tau_{ik}^{-\theta} X_k / P_k^{-\theta} = \sum_k \tau_{ik}^{-\theta} X_k / \Pi_k^{-\theta} \propto \sum_k \tau_{ik}^{-\theta} \alpha_k.$$

Combining $\tau_{ik}^{-\theta} \propto Distance_{ik}^{-\zeta}$ and the above, we get the proposed formula,

$$Size_i \ \propto \ Pop_i T_i^{1/\theta} \ \propto \ \alpha_i^{1+1/\theta} \sum_k Distance_{ki}^{-\zeta} \alpha_k . \square$$

Analytical formulas for standard errors. We follow Cameron and Trivedi (2005) throughout. The notation is also borrowed from that book, specifically pages 200-202. Let $\boldsymbol{\theta} = (\zeta, \cdots)'$ be the parameters of the first step (PPML), collecting the distance elasticity of trade ζ and the exporter and importer fixed effects. $\boldsymbol{\beta}$ is the vector of parameters of the second step (NLLS). Let K_1 be the number of city pairs for step 1, and K_2 the number of city pairs for step 2. K_1 includes all known cities in our sample that import from other known cities that import from other known cities ... etc. K_2 includes all cities (known or lost) in our sample, so that $K_1 < K_2$.

The variance-covariance matrix of $(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\beta}})$ is

$$\hat{\boldsymbol{\Sigma}} = \widehat{var}(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\beta}}) = \frac{1}{K_2} \begin{pmatrix} G_{11} & 0 \\ G_{21} & G_{22} \end{pmatrix}^{-1} \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \begin{pmatrix} G_{11} & 0 \\ G_{21} & G_{22} \end{pmatrix}^{-1\top} .$$

We now explain what the components of these block matrices are.

Consider first the Poisson Pseudo MLE step. Let y_{ij} be the outcome variable, trade share in counts, and x_{ij} be the vector of covariates. The first component of $\boldsymbol{\theta}$ is the distance elasticity of trade. The other components are coefficients associated with the exporter and importer dummies. The log-likelihood is then:

$$\log \mathcal{L}(\boldsymbol{\theta}) = \sum_{i,j} \left(-\exp\{x_{ij}'\boldsymbol{\theta}\} + y_{ij} x_{ij}'\boldsymbol{\theta} - \log y_{ij}! \right) .$$

In m-Estimators terminology (Cameron and Trivedi, p.118), we have:

$$Q_{K_1}(\boldsymbol{\theta}) = \frac{1}{K_1} \sum_{i,j} \left(-\exp\{x_{ij}'\boldsymbol{\theta}\} + y_{ij} x_{ij}'\boldsymbol{\theta} - \log y_{ij}! \right) = \frac{1}{K_1} \sum_{i,j} q_{i,j}(\boldsymbol{\theta}).$$

In the notation of two-step m-Estimation (p.200), we have that $h_1(w_{ij}, \boldsymbol{\theta}) = \frac{d}{d\boldsymbol{\theta}} q_{ij}(\boldsymbol{\theta})$. Thus,

$$G_{11} = \frac{1}{K_1} \sum_{i,j} \frac{dh_1(w_{ij}, \hat{\boldsymbol{\theta}})}{d\boldsymbol{\theta}'} = \frac{1}{K_1} \sum_{i,j} \frac{d^2 q_{ij}(\hat{\boldsymbol{\theta}})}{d\boldsymbol{\theta} d\boldsymbol{\theta}'} = \frac{1}{K_1} \sum_{i,j} \exp\{x_{ij}'\boldsymbol{\theta}\} x_{ij} x_{ij}'$$

and

$$S_{11} = \frac{1}{K_1} \sum_{ij} h_1(w_{ij}, \hat{\boldsymbol{\theta}}) h_1(w_{ij}, \hat{\boldsymbol{\theta}})' = \frac{1}{K_1} \sum_{ij} \frac{dq_{ij}(\hat{\boldsymbol{\theta}})}{d\theta_1} \frac{dq_{ij}(\hat{\boldsymbol{\theta}})}{d\theta_1'} = \frac{1}{K_1} \sum_{ij} (y_{ij} - \exp\{x_{ij}'\boldsymbol{\theta}\})^2 x_{ij} x_{ij}'.$$

This coincides with the BHHH estimate (p.138).

Consider now the standard non-linear least squares step. Let e_{ij} be the difference between the model and data trade shares. Given the first stage estimate $\hat{\boldsymbol{\theta}}$, the first order condition of the least squares problem is:

$$\sum_{i,j} e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})=0.$$

In the notation of the two-step m-Estimation (p.200), we have that $h_2(w_{ij},\hat{\boldsymbol{\theta}},\hat{\boldsymbol{\beta}})=e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})$. The nonlinear least squares simplifications from p.153 apply, so that

$$G_{22}=\frac{1}{K_2}\sum_{i,j}\frac{\partial}{\partial\boldsymbol{\beta}}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}'}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}}).$$

Following p.201, we have

$$S_{22}=\frac{1}{K_2}\sum_{i,j}h_{2i}h_{2i}'=\frac{1}{K_2}\sum_{i,j}\left[e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\right]\left[e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}'}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\right]$$

Finally, we compute the interactions terms, G_{21}, S_{12} and S_{21}. Given that we know h_1 and h_2, we use the following variation on p.202:

$$\begin{aligned}
G_{21}&=\frac{1}{K_2}\sum_{i,j}\frac{\partial h_2}{\partial\boldsymbol{\theta}'}=\frac{1}{K_2}\sum_{i,j}\frac{\partial}{\partial\boldsymbol{\theta}'}\left[e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\right]\\
S_{12}&=\frac{1}{K_1}\sum_{i,j}h_{1i}h_{2i}'=\frac{1}{K_1}\sum_{i,j}\left(\frac{d}{d\boldsymbol{\theta}}q_{ij}(\hat{\boldsymbol{\theta}})\right)e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\frac{\partial}{\partial\boldsymbol{\beta}'}e_{ij}(\hat{\boldsymbol{\beta}};\hat{\boldsymbol{\theta}})\\
S_{21}&=S_{12}^{\top}.
\end{aligned}$$

To derive the standard errors for city sizes, $Size_i$, we apply the Delta method to (10), using the above covariance matrix $\hat{\boldsymbol{\Sigma}}$.

Analytical formulas for iso-density contours and *precision* (l) **in (12).** Iso-density contours are points with latitude-longitude (φ,λ) such that $f_l(\varphi,\lambda)=c$ where f_l is the p.d.f. of the bi-variate normal distribution with estimated mean $(\hat{\varphi}_l,\hat{\lambda}_l)$, variance $(\hat{\sigma}^2_{\varphi_l},\hat{\sigma}^2_{\lambda_l})$, and correlation $\hat{\rho}_{\varphi_l,\lambda_l}$,

$$c=\frac{1}{2\pi\hat{\sigma}_{\varphi_l}\hat{\sigma}_{\lambda_l}\sqrt{1-\hat{\rho}^2_{\varphi_l,\lambda_l}}}\exp\left(-\frac{1}{2(1-\hat{\rho}^2_{\varphi_l,\lambda_l})}\left[\frac{(\varphi-\hat{\varphi}_l)^2}{\hat{\sigma}^2_{\varphi_l}}+\frac{(\lambda-\hat{\lambda}_l)^2}{\hat{\sigma}^2_{\lambda_l}}-\frac{2\hat{\rho}_{\varphi_l,\lambda_l}(\varphi-\hat{\varphi}_l)(\lambda-\hat{\lambda}_l)}{\hat{\sigma}_{\varphi_l}\hat{\sigma}_{\lambda_l}}\right]\right).$$

This is the formula for an ellipse. We use four values for c corresponding to 50%, 75%, 90%, and 99% confidence regions. For instance, for the 75th percentile, we use c_{75} defined as,

$$\Pr\left[(\varphi,\lambda)\,|\text{ s.t. } f_l(\varphi,\lambda)\geq c_{75}\right]=0.75.$$

To derive an analytical formula for $precision(l)$, we start from the definition,

$$precision\,(l) = \sqrt{\mathbb{E}_{(\varphi,\lambda)\sim\mathcal{N}(\hat{\boldsymbol{\beta}}_l,\hat{\boldsymbol{\Sigma}}_l)}\left[\left(Distance\left(\hat{\varphi}_l,\hat{\lambda}_l;\varphi,\lambda\right)\right)^2\right]}.$$

Using the Euclidean formula for distance and linearity of the expectation operator, we get,

$$\begin{aligned}\mathbb{E}_{(\varphi,\lambda)\sim\mathcal{N}(\hat{\boldsymbol{\beta}}_l,\hat{\boldsymbol{\Sigma}}_l)}\left[\left(Distance\left(\hat{\varphi}_l,\hat{\lambda}_l;\varphi,\lambda\right)\right)^2\right] &= \left(\frac{10000}{90}\right)^2\left(\mathbb{E}\left[(\varphi-\hat{\varphi}_l)^2\right]+\cos^2\left(\frac{37.9}{180}\pi\right)\mathbb{E}\left[(\lambda-\hat{\lambda}_l)^2\right]\right)\\ &= \left(\frac{10000}{90}\right)^2\left[\hat{\sigma}^2_{\varphi_l}+\cos^2\left(\frac{37.9}{180}\pi\right)\hat{\sigma}^2_{\lambda_l}\right]\end{aligned}$$

Thus, our formula for the geographic precision (in kms) for lost city l is,

$$precision\,(l) = \frac{10000}{90}\sqrt{\hat{\sigma}^2_{\varphi_l}+\cos^2\left(\frac{37.9}{180}\pi\right)\hat{\sigma}^2_{\lambda_l}},$$

where $(\hat{\sigma}^2_{\varphi_l},\hat{\sigma}^2_{\lambda_l})$ are the variances for the estimated latitudes and longitudes $(\hat{\varphi}_l,\hat{\lambda}_l)$ of city l.

Summary of our Optimization Procedure. Our estimates of $(\zeta;\cdots(\varphi_l,\lambda_l)\cdots;\cdots\alpha_i\cdots)$ are computed in two steps. First, the distance elasticity of trade – ζ – is estimated by Poisson Pseudo Maximum Likelihood in the subsample with known location data. Specifically, the independent variable in the Poisson regression model is the observed trade shares, and the dependent variables are the log of distance between cities, destination city dummies and origin city dummies. The estimate for ζ is the resulting coefficient for the log of distance between cities. Our estimation uses the `ppml` STATA command written by Silva and Tenreyro (2006).

Second, the geo-coordinates of lost cities and the α_i's are estimated by minimizing the sum of squared differences between observed and predicted trade shares given our estimated of ζ. This function of our parameters was coded in Python-Numpy, and the optimization performed using IPOPT. In particular, we ran 260 jobs in parallel. Each job executed the minimization process 20 times, starting from random initial values. The geo-coordinates' initial values were uniformly drawn between 36 and 42 degrees of latitude and 27 and 45 degrees of longitude, whereas the α_i's initial values were uniformly drawn between 0 and 200. The IPOPT specifications include: $100,000$ maximum iterations, `MA57` as the linear solver, a tolerance level of $1.0e^{-08}$, and an acceptable tolerance level of $1.0e^{-07}$. Issues of local minima are dealt away upon inspection of all minimization results. For our main specification, 242 out of $5,200$ executions yield the same lowest sum of squared differences and have the same parameter estimates, up to 6 decimal points.

B Data Sources and Optimization Procedure

B.1 Data Construction and Selection of Cities

The database of ancient texts that we have access to contains references to 79 unique cities in Anatolia (Barjamovic, 2011). Decades of scholarship has successfully disambiguated place names from terms previously mistaken as cities while they actually refer to types of textiles or people.[34] We drop 40 cities with only a single mention in the entire corpus: their parameters are not identified by our structural gravity estimation. Other than the remaining 39 Anatolian cities that appear more than once in the most up to date corpus of Assyrian texts, there are two cities *outside of Anatolia* with known locations: *Aššur* and *Qattara*. As explained in the text (see footnote 3), we exclude *Aššur* from our analysis because the word for *Aššur* is ambiguous, appearing in persons' names and as the name of the main Assyrian deity. This makes it impossible to run an automated search for the city in the digitized text. We also drop the city of *Qattara* because it was known to be a small independent polity in proximity to *Aššur*, where Assyrian merchants were allowed to pass by, but not to trade. *Qattara* being also the first stop on the way from *Aššur* to Anatolia (last stop on return trips), any mention of *Qattara* can only mean a shipment to or from *Aššur*, not a shipment to or from *Qattara* itself. Similar to *Aššur*, this city is also outside of the modern day Turkish geography we focus on.

The second criterion is that, out of the set of 39 Anatolian cities, we drop 9 lost cities about which professional Assyriologists have only speculative conjectures in terms of locations. This is due to both the relatively low number of mentions to these cities in the texts as well as the lack of precise descriptions of their geography—see Barjamovic (2011), Table 39 in page 411.

The final criterion is mechanical: any city that is disconnected from the rest of the network of trading cities cannot be identified in a gravity model. If there is no trade at all between city i and any other city, our estimator will assign a size zero to that city, and/or an infinite distance if this city is lost. Note that network connectedness is an iterative notion. We start by dropping all disconnected cities. Having dropped those cities, some cities which were only trading with the dropped cities are dropped in turn, which further eliminates the cities that were only trading with those, and so forth. This iterative process leaves us with 25 cities.

The next figure shows the distribution of shipment counts, for all (directed) city pairs with a positive number of shipments, $N_{ij}^{data} > 0$.

[34] For more details, see Bilgiç (1951) and Michel and Veenhof (2010).

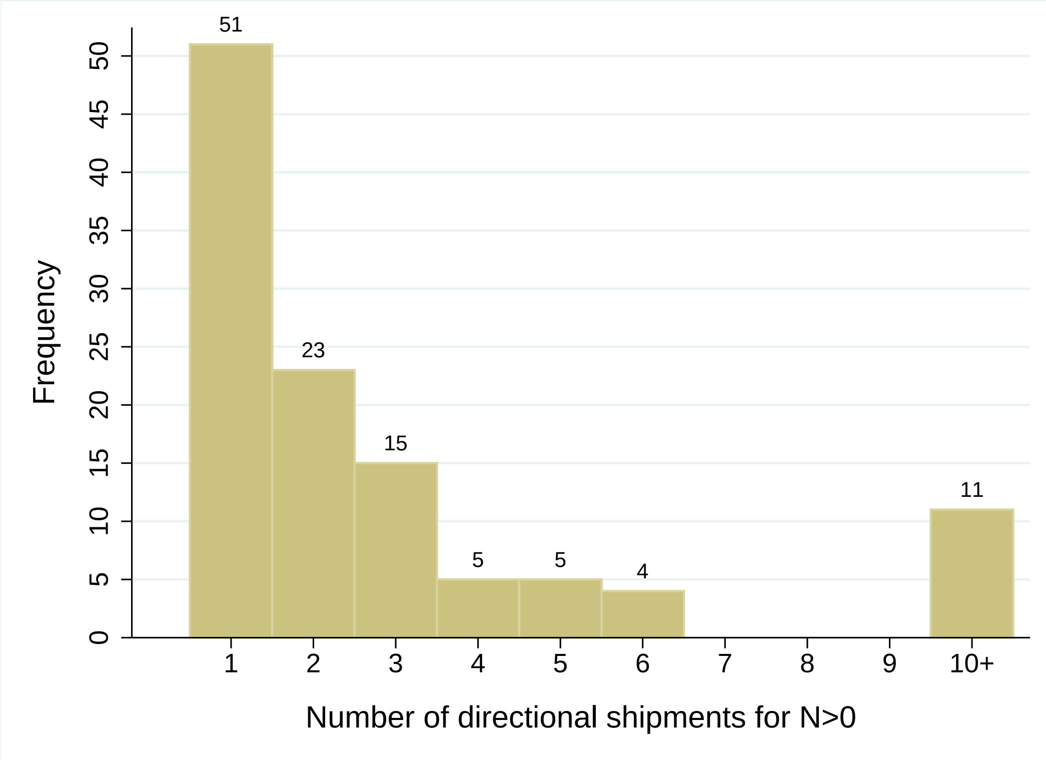

Appendix Figure I: **Frequency of shipment counts.**

B.2 An Example of a Coincidental Joint Attestation of Two Cities

After reading the 2,806 tablets which mention at least two cities, we discard any case where the two cities are mentioned in the same tablet for reasons unrelated to any trade relationship between those cities. Below is an example of a purely coincidental joint attestation of two city names, *Hahhum* and *Wahšušana*, underlined in the text for clarity,

> *From Enlil-bani to Aššur-idi with Cc to Ṭab-ṣilli-Aššur and Aššur-[xx] . We had half a pound of silver transferred in Hahhum, and I gave it to you in the year of Aššur-malik. You said: "After I make a transport to the City, I will either give you the proceeds or I will take goods on commission for you." I paid you in refined silver and we saw each other about five times when you cheated me. While I was staying in Wahšušana, my representatives seized you, but you gave them nothing. My representatives said ...*
>
> [Tablet BIN 6, 38 (NBC 3808) lines 1-18]

In this text, the merchant Enlil-bani accuses another merchant, Ṭab-ṣilli-Aššur, of having cheated him. Enlil-bani mentions the city of *Hahhum* in the context of a financial transaction between both merchants (second line). The city of *Wahšušana* is mentioned in passing in the same letter (penultimate line) to describe Enlil-bani's whereabouts while his representatives were trying to recover his due. It has nothing to do with the initial financial transaction. There is no direct economic connection between both cities that can be drawn from this letter.

B.3 A Partial Example of how Historians Locate Lost Cities

Historians Forlanini (2008) and Barjamovic (2011) use a series of references to ancients cities during the Middle Bronze Age period, complemented with references from later periods, to make informed proposals for the location of lost cities. A detailed list of references to such proposals was collected by Nashef (1992) with updates in Ullmann and Weeden (2017).

An example of a few snippets of information that guide the reasoning of historians would be the following references related to the city of *Hahhum*. The following text locates *Hahhum* on a river:

> *I met Elali in Hahhum while I was staying there at the bank of the river in Habnuk.*
> [Tablet I 469]

Another text, in this case a later Hittite royal annal, makes it explicit that the river in question was the Euphrates:

> *I the Great King Tabarna took away from Hahhu and presented it to the Sun-God. The Great King Tabarna removed the hands of its slave girls from the grindstone and its slaves? hands he removed ... he released their belts and he put them in the temple of the Sun-Goddess of Arinna ... The great river Euphrates, no one had crossed it. [The Great King] Tabarna crosses it on foot, and his troops crossed it [on] foot after him*
>
> [Text KBo 10.1]

A large number of additional references eventually allows historians to formulate with confidence the hypothesis that *Hahhum* lies to the South and East of *Kaneš*, in a position that allowed it to control an important river crossing. Its neighbor on the opposite river bank was *Badna*, and *Timelkiya* was the next state on the route north-west to *Kaneš*. Gradually, the positions of the cities can be established in relation to one another to form a network that can then be placed on a map. In rare cases, place names may survive to help guide this process, if their location in later times is known.

B.4 Data on Modern-day Trade, Local Resources, and Topography

Modern-day population: Data on the 2012 urban population of Turkish districts is obtained from the website of the Turkish Statistical Agency (https://biruni.tuik.gov.tr/EdUygulamaDis/zul/

loginEN.zul?lang=en).

Night time luminosity: Data on nighttime light emissions intensity is used as a proxy for local GDP at the very granular level, as in Hodler and Raschky (2014). The data for 2003 is from the National Oceanic and Atmospheric Administration (NOAA), available at http://ngdc.noaa.gov. Weather satellites from the U.S. Air Force measure light intensity between 8:30PM and 10:00PM, removing observations affected by cloud conditions, and correcting for likely ephemeral lights or background noise.

Crop yields: We use the low-input level rain-fed cereal suitability index of IIASA/FAO (2012) available at http://www.fao.org/nr/gaez/en/.

Ruggedness: We use the Terrain Ruggedness Index of Riley, DeGloria, and Elliot (1999), defined as the square root of the sum of squared elevation changes in 8 directions.

Elevation: Elevation data to calculate the natural road scores is obtained from IIASA/FAO (2012) and is available at http://www.fao.org/nr/gaez/en/. See appendix C for details.

Rivers and lakes: The shapefile of rivers and lakes in Turkey, used in calculating the natural road scores, has been downloaded from http://www.naturalearthdata.com/. See appendix C for details.

C Optimal Travel Routes

To define optimal travel routes, we use topographical data, and Dijkstra's shortest path algorithm.

Formally, we collect data on elevation on a fine grid (each pixel's side is 5 arc minutes, or about 10 km). The elevation data is downloaded from FAO-GAEZ, which itself is based on NASA's shuttle radar topography mission (IIASA/FAO, 2012). We collect elevation data on a large area around central Anatolia, in order to avoid a core-periphery bias –the tendency to have more road-crossings in locations in the center of the map. The total area is contained between 20 and 50 degrees of longitude East, and 30 and 43 degrees of latitude North. It corresponds to a wide area between Hungary in the Northwest, Kazakhstan in the Northeast, Kuwait in the Southeast, and Libya in the Southwest, which extends well beyond central Anatolia. We remove from this map the Arabian desert, assuming implicitly that traveller were not crossing it. We do include any maritime area where the sea is less than 500 m deep, allowing maritime travel along the coasts, but preventing high sea travel. The elevation map for the entire region we consider is depicted below.

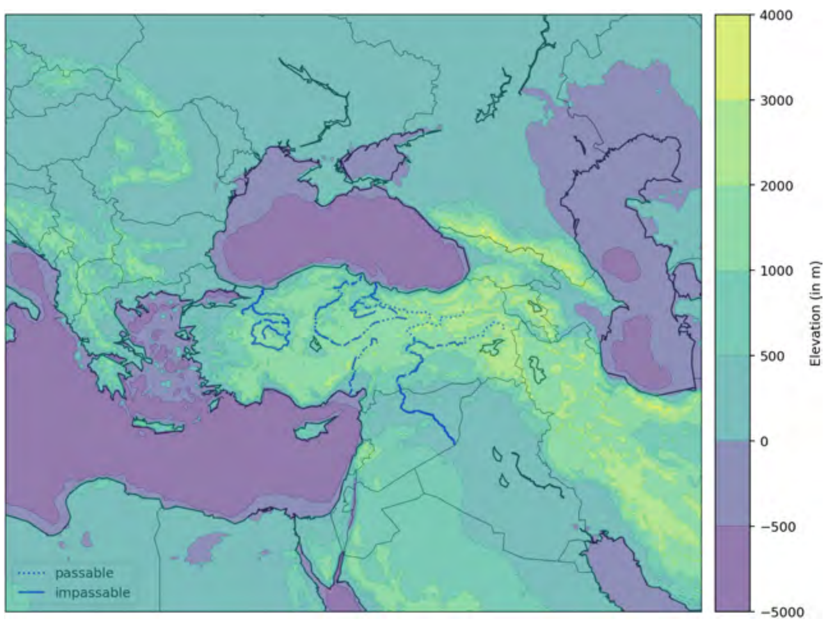

Appendix Figure II: **Elevation.**

Given this fine elevation grid, we compute travel times between any pixel and its eight neighbors (North-South, East-West, and diagonally). First, we compute the horizontal distance between any two contiguous pixels (in meters), and the signed elevation difference between them (in meters). We then apply the formula from Langmuir (1984) to translate distances and slope into travel times (in seconds). The parametrization of Langmuir's formula we apply for overland travel is as follows.

We assume it takes 0.72 seconds to travel 1 meter horizontally; it takes an additional 6 seconds for each vertical meter uphill; going downhill 1 vertical meter on a gentle slope (less than or equal to 21.25%) saves 2 seconds per vertical meter; going downhill on a steep slope (more than 21.25%) adds an additional 2 seconds per vertical meter. For maritime travel, we assume traveling by boat is 10% faster than traveling over a featureless plain overland. However, in order to avoid many very short trips by sea, we assume it takes 1 hour to embark on a boat (from land to sea), and 1 hour to disembark (from sea to land). This assumption of a fixed cost of loading/unloading boats creates a natural tendency for sea ports to emerge where natural terrestrial routes (e.g. a valley) connect to the sea. Finally, we manually code major lakes, and three main rivers, the Euphrates, the Red river, and the Green river in Turkey. For the three rivers, we collect information on impassable segments of the river (e.g. deep and steep canyon), as well as easy crossings/fords known to have been used in the Bronze Age, using data from Palmisano (2013). We impose a prohibitive penalty for crossing major lakes and those three rivers over segments where they are deemed impassable, and allow crossing as if it were on dry land for the river crossings.

Having defined travel times between any pixel and its eight neighbors, we apply Dijkstra's algorithm to compute the optimal travel paths between any two pixels (Dijkstra, 1959). We use those travel paths and travel times when coding the information contained in merchants itineraries in section IV.B (see the details in appendix D), and when defining natural roadways to compute the variable $NaturalRoads$ in section V.B (see the details in appendix E).

D Constraints from Merchants Itineraries

To impose constraints on the location of lost cities, using information contained in multi-stop merchant itineraries, we proceed as follows.

First, we collect systematic information on multiple itineraries of either merchants or caravans in our corpus of texts. We keep only itineraries with at least one stop in a lost city.

Second, we compute the following two statistics for all segments of those itineraries from one known city to another known city: the average travel length $||\text{average segment}||$, and the standard deviation of the segment lengths $||\text{s.d. segment}||$. Our measure of length is the optimal travel time between the two ends of each segment, defined in appendix C.

Third, we jointly impose on all lost cities the "short detour" and "pit stop" constraints defined in section IV.B, using all mentions of itineraries. In other words, for all lost cities jointly, we search for all grid points such that both constraints are satisfied. To solve for this multi-dimensional search, we proceed sequentially. We start by imposing all "pit stop" constraints where one city is known, and one city is unknown. This gives us an admissible region for each lost city mentioned at least once alongside a known city. We then impose all the "short detour" constraints involving two known cities and one lost city, searching only within the admissible regions of the previous step. This further restricts the size of the admissible regions for each lost city mentioned at least once alongside known cities. We finally solve a minimization problem using all the remaining "pit stop" and "short detour" constraints, imposing a penalty for a violation of the constraints.

E Constructing Natural Road Scores

To compute the $NaturalRoads$ variable, we proceed as follows.

We use the large region depicted in figure II, which extends widely beyond central Anatolia. For any two pixels on that map, we know the route of the shortest path from one to the other, using the procedure described in appendix C. Our purpose is to compute, for each pixel on the map, the number of optimal paths that intersect on that pixel. This corresponds to the notion of between-ness centrality in the network of optimal routes.

In order to distinguish between short distance routes (arguably travelled frequently) versus long distance routes (arguably travelled infrequently), we implicitly assume a gravity model of migrations. For any optimal route from pixel i to pixel j with duration d_{ij}, we assume this route is travelled with a probability proportional to $d_{ij}^{-\hat{\zeta}}$, where $\hat{\zeta} = 1.9$ is our estimate for the distance elasticity of trade in the Middle Bronze Age, and where we use the algorithm described in appendix C to compute shortest paths and their durations. This probability weighting corresponds to an implicit gravity model of migrations, where humans are uniformly distributed over space, and travel between locations according to a gravity model with distance elasticity $\hat{\zeta}$.

For any pair of pairs of points on Figure II, (A,B) and (C,D), with shortest paths of durations d_{AB} and d_{CD}, we draw the pair (A,B) with probability proportional to $d_{AB}^{-\hat{\zeta}}$ and the pair (CD) with probability proportional to $d_{CD}^{-\hat{\zeta}}$. If the two paths either intersect or overlap on a given pixel, we record their intersection for that pixel point. We repeat this procedure 1 million times. Each pixel receives a "road-knot score" equal to the number of intersections or overlaps recorded on that pixel.

For each ancient city i, either known or lost (we use our structural gravity estimates for the location of lost cities as our main specification, but also experiment with alternative locations as robustness), our variable of interest, $NaturalRoads_i$, simply adds up this "road-knot score" from all pixels which are within 20 km of city i.

References

ANDERSON, J. E. AND E. VAN WINCOOP (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–92.

BARJAMOVIC, G. (2011): *A Historical Geography of Anatolia in the Old Assyrian Colony Period*, Copenhagen: Museum Tusculanum Press.

BILGIÇ, E. (1951): "Die Ortsnamen der kappadokischen Urkunden im Rahmen der alten Sprachen Anatoliens," *Archiv für Orientforschung*, 15, 3–37.

CAMERON, A. C. AND P. K. TRIVEDI (2005): *Microeconometrics: methods and applications*, Cambridge university press.

DIJKSTRA, E. W. (1959): "A note on two problems in connexion with graphs," *Numerische mathematik*, 1, 269–271.

EATON, J. AND S. KORTUM (2002): "Technology, Geography and Trade," *Econometrica*, 70, 1741–79.

EATON, J., S. KORTUM, AND S. SOTELO (2012): "International Trade: Linking Micro and Macro," *NBER Working Paper No.17864*.

FORLANINI, M. (2008): "The Central Provinces of Hatti. An Updating," in *New Perspectives on the Historical Geography and Topography of Anatolia in the II and I Millennium BC*, ed. by K. Strobel, (EOTHEN 16) Firenze: LoGisma Editore, 1, 145–188.

HODLER, R. AND P. A. RASCHKY (2014): "Regional Favoritism," *The Quarterly Journal of Economics*, 129, 995–1033.

IIASA/FAO (2012): *Global Agro-Ecological Zones (GAEZ v3.0)*, IIASA, Laxenburg, Austria and FAO, Rome.

LANGMUIR, E. (1984): *Mountaincraft and leadership: a handbook for mountaineers and hillwalking leaders in the British Isles*, Edinburgh: Scottish Sports Council.

MICHEL, C. AND K. VEENHOF (2010): "The Textiles Traded by the Assyrians in Anatolia (19th-18th centuries BC)," in *Textile Terminologies in the Ancient Near East and Mediterranean from the Third to the First Millennia BC*, ed. by C. Michel and M.-L. Nosch, Ancient Textiles Series Vol. 8. Oxford and Oakville: Oxbow Books, 210–271.

NASHEF (1992): *Die Orts- und Gewassernamen der altassyrischen Zeit*, Wiesbaden: Ludwig Reichert.

PALMISANO, A. (2013): "Computational and Spatial Approaches to the Commercial Landscapes and Political Geography of the Old Assyrian Colony Period." in *Time and History in the Ancient Near East. Proceedings of the 56th Rencontre Assyriologique Internationale, Barcelona, July 26-30, 2010.*, ed. by L. Feliu, J. Llop, A. M. Albà, and W. Lake., Eisenbrauns, 767–783.

RILEY, S. J., S. D. DEGLORIA, AND R. ELLIOT (1999): "Index that quantifies topographic heterogeneity," *Intermountain Journal of sciences*, 5, 23–27.

SILVA, J. S. AND S. TENREYRO (2006): "The Log of Gravity," *The Review of Economics and statistics*, 88, 641–658.

Tobler, W. and S. Wineburg (1971): "A Cappadocian speculation," *Nature*, 231, 39–41.

Ullmann, L. and M. Weeden (2017): *Hittite Landscape and Geography*, Leiden: Brill.

F Additional Tables

Appendix Table I: **Proof of Concept, Recovering Fictitiously Lost Cities**

	True Coordinates		Estimated Coordinates		
	Latitude	Longitude	Latitude	Longitude	Distance in km
Hattus	40.021	34.61	40.044	34.635	3
			(0.53)	(0.58)	
Kanes	38.85	35.633	38.955	35.277	33
			(0.407)	(0.375)	
Karahna	40	36.1	40.021	34.618	130
			(2964.306)	(170.666)	
Tapaggas	40.148	35.762	40	35.818	17
			(317.12)	(97.762)	
Hanaknak	40	35.817	40.164	35.764	19
			(0.015)	(0.24)	
Hurama	38.261	37.114	39.643	37.327	155
			(0.715)	(0.457)	
Malitta	39.363	33.787	38.996	34.644	86
			(0.659)	(0.442)	
Mamma	37.583	36.933	38.02	36.5	62
			(1.603)	(0.513)	
Salatuwar	39.655	31.994	39.64	33.192	105
			(0.895)	(0.413)	
Samuha	39.619	36.528	39.399	36.11	44
			(0.196)	(0.383)	
Timelkiya	38.027	38.234	38.403	37.488	78
			(0.451)	(0.411)	
Ulama	38.411	33.834	39.852	33.196	170
			(0.402)	(0.966)	
Unipsum	38.021	36.503	37.583	36.933	61
			(1462.246)	(1105.701)	
Wahsusana	39.584	33.418	38.457	32.077	172
			(17.831)	(27.255)	
Zimishuna	40.461	35.65	40.47	35.65	1
			(307617.802)	(461223.631)	
Mean					76
Median					62

Notes: This table presents the results from our 'proof-of-concept' exercise in section IV.C. For all known cities, we list the true geo-coordinates of each city, their estimated geo-coordinates from estimating a model similar to (8), and the distance, in kms, between the true and estimated locations. The results for all cities are shown on single map on figure VIII. All latitudes are North, and all longitudes are East. Robust (White) standard errors in parentheses.

Appendix Table II: **Proof of Concept (robustness), Recovering One Fictitiously Lost City and Ten Lost Cities Jointly**

	True Coordinates		Estimated Coordinates		
	Latitude	Longitude	Latitude	Longitude	Distance in km
Hattus	40.021	34.61	39.997	36.131	133
			(0.322)	(0.073)	
Kanes	38.85	35.633	39.313	33.918	159
			(0.252)	(0.156)	
Karahna	40	36.1	40.046	34.547	136
			(0.754)	(0.475)	
Tapaggas	40.148	35.762	40	35.817	17
			(167.04)	(115.323)	
Hanaknak	40	35.817	40.15	35.761	17
			(0.25)	(0.522)	
Hurama	38.261	37.114	39.139	38.226	138
			(1.221)	(4.488)	
Malitta	39.363	33.787	38.888	35.282	141
			(0.515)	(0.141)	
Mamma	37.583	36.933	38.02	36.503	61
			(1.37)	(0.782)	
Salatuwar	39.655	31.994	39.561	33.356	120
			(0.959)	(0.604)	
Samuha	39.619	36.528	38.3	37.118	155
			(0.09)	(0.563)	
Timelkiya	38.027	38.234	38.261	37.114	102
			(0.019)	(0.045)	
Ulama	38.411	33.834	39.835	33.234	167
			(0.4)	(1.277)	
Unipsum	38.021	36.503	37.583	36.933	61
			(389261.721)	(646044.057)	
Wahsusana	39.584	33.418	39.003	31.926	146
			(10.185)	(50.561)	
Zimishuna	40.461	35.65	39.234	34.213	186
			(23.126)	(50.515)	
Mean					116
Median					136

Notes: This table presents the results from a robustness check of our 'proof-of-concept' exercise in appendix table I. For each line, we set the distance elasticity at $\zeta = 1.9$, and using (8), we estimate the geo-coordinates of one fictitiously lost city alongside the ten truly lost cities. We also re-estimate all other parameters of the model. For *Karahna* and *Zimišhuna*, our minimization algorithm hits the non-negativity constraint on their α's. For all known cities, we list the true geo-coordinates of each city, their estimated geo-coordinates from estimating (8), and the distance, in kms, between the true and estimated locations. All latitudes are North, and all longitudes are East. Robust (White) standard errors in parentheses.

Appendix Table III: **Assigning Lost Cities to Archaeological Sites**

Lost city gravity estimate	Candidate site	Distance to gravity estimate (in km)	Log(p.d.f.)
Durhumit 40.47, 35.65			
	Ayvalıpınar 40.46, 35.65	0.97	1.68
	Oluz Höyük 40.55, 35.63	8.62	-45.66
	Ferzant 40.6, 35.38	27.53	-114.33
	Doğantepe 40.6, 35.6	14.71	-130.09
	Boyalı 40.31, 34.26	123.58	-463.26
Hahhum 38.43, 38.04			
	Imikuşağı 38.52, 38.46	37.9	0.32
	Değirmentepe 38.48, 38.45	36.18	0.26
	Imamoğlu 38.48, 38.48	38.92	0.18
	Arslantepe 38.38, 38.36	28.69	0.17
	Yassıhöyük (Tanır) 38.39, 36.91	98.88	-3.38
Kuburnat 40.71, 36.52			
	Tekkeköy (Samsun) 41.2, 36.45	54.57	-1.01
	Dündartepe 41.25, 36.35	61.74	-1.12
	Kaledoruğu (Kavak) 41.08, 36.04	58.61	-1.26
	Kayapınar Höyüğü 40.16, 36.25	65.42	-1.27
	Bolus (Aktepe) 40.07, 36.5	71.72	-1.28

Continued on the next page

Appendix Table III: **Assigning Lost Cities to Archaeological Sites (continued)**

Lost city gravity estimate	Candidate site	Distance to gravity estimate (in km)	Log(p.d.f.)
Ninassa 38.98, 34.61			
	Suluca Karahöyük (Hacibektaş) 38.93, 34.55	7.42	0.11
	Topakhöyük 38.61, 34.29	49.37	-0.12
	Zank 38.95, 34.79	16.05	-0.19
	Topaklı 39.01, 34.83	18.93	-0.2
	Uşaklı/Kuşaklı Höyük 39.8, 35.1	100.87	-0.45
Purushaddum 39.71, 32.87			
	Karaoğlan 39.73, 32.83	4.3	-1.13
	Külhöyük (Haymana) 39.48, 32.67	31.43	-1.18
	Ballıkuyumcu 39.77, 32.52	31.77	-1.51
	Çomaklı / İlmez 37.72, 32.5	223.91	-2.21
	Ortakaraviran II 37.38, 32.09	267.98	-2.27
Sinahuttum 39.96, 34.87			
	Yassıhöyük (Yozgat) 39.99, 34.88	3.97	2.05
	Çengeltepe 39.84, 34.87	12.99	1.79
	Eskiyapar 40.16, 34.77	23.91	-0.44
	Mercimektepe 40.88, 35.34	110.6	-2.25
	Suluca Karahöyük (Hacibektaş) 38.93, 34.55	116.96	-3.73

Continued on the next page

Appendix Table III: **Assigning Lost Cities to Archaeological Sites (continued)**

Lost city gravity estimate	Candidate site	Distance to gravity estimate (in km)	Log(p.d.f.)
Suppiluliya 40.02, 34.62			
	Suluca Karahöyük (Hacibektaş) 38.93, 34.55	120.98	-3.65
	Alaca Höyük 40.23, 34.68	24.3	-51.97
	Büyüknefes (Bronze Age Site) 39.85, 34.5	21.61	-201.29
	Eskiyapar 40.16, 34.77	20.46	-373.9
	Arslantepe 38.38, 38.36	375.37	-Inf
Tuhpiya 39.61, 35.2			
	Çadır (Sorgun) 39.68, 35.14	9.27	1.25
	Uşaklı/Kuşaklı Höyük 39.8, 35.1	22.76	0.34
	Çengeltepe 39.84, 34.87	38.43	-1.48
	Boğazlıyan / Yoğunhisar 39.17, 35.23	49.42	-3.01
	Üyük 40.15, 35.85	82.78	-3.27
Washaniya 39.16, 34.31			
	Yassıhöyük (Çoğun / Kırşehir) 39.32, 34.08	27.42	0.34
	Suluca Karahöyük (Hacibektaş) 38.93, 34.55	32.51	0.08
	Harmandalı 38.95, 33.95	39.14	-0.38
	Zank 38.95, 34.79	48.08	-1.08
	Topaklı 39.01, 34.83	48.17	-1.11

Continued on the next page

Appendix Table III: **Assigning Lost Cities to Archaeological Sites (continued)**

Lost city gravity estimate	Candidate site	Distance to gravity estimate (in km)	Log(p.d.f.)
Zalpa 38.81, 37.86			
	Yassıhöyük (Tanır) 38.39, 36.91	95.1	-0.86
	Yalak (Boz Höyük) 38.3, 36.44	137.32	-1.39
	Sarız 38.47, 36.5	125.22	-1.67
	Imikuşağı 38.52, 38.46	60.94	-2.32
	Değirmentepe 38.48, 38.45	62.95	-2.56

Notes: This table gives a list of the five most likely potential archaeological sites for each lost city. The table lists each lost city, with its the geo-coordinates (first latitude, North, then longitude, East) derived from estimating our gravity model (for instance, the estimated coordinates for Durhumit are 40.47 degrees North and 35.65 degrees East). It then lists the top five candidate sites, with their corresponding geo-coordinates. The sites are ordered in decreasing order of probability density, where we evaluate our estimated probability density, $\hat{f}_l(\varphi, \lambda)$ for each lost city l at each potential site's coordinates, using equation (13). For each potential site, we display the distance (in km) from its corresponding lost city gravity estimate, and the (log) probability density for that site. For instance, Ayvalıpınar is the most likely location for Durhumit, at a distance of 0.97km, and with a log(p.d.f.) of 1.68.

Appendix Table IV: **Determinants of Ancient City Sizes, Robustness**

	$\log\left(PopT^{1/\theta}\|_{ancient}\right)$				
	(1)	(2)	(3)	(4)	(5)
Panel A: Barjamovic (2011) locations					
$\log\left(NaturalRoads\right)$	0.141			0.163	
	(0.281)			(0.234)	
$\log\left(RomanRoads\right)$		0.776*			0.783*
		(0.068)			(0.076)
$\log\left(Ruggedness\right)$			0.067	0.135	0.080
			(0.766)	(0.557)	(0.679)
N	25	25	25	25	25
R^2	0.062	0.145	0.006	0.084	0.153
Panel B: known cities only					
$\log\left(NaturalRoads\right)$	2.105*			2.526*	
	(0.086)			(0.057)	
$\log\left(RomanRoads\right)$		4.540			4.771
		(0.207)			(0.125)
$\log\left(Ruggedness\right)$			2.722	3.970***	2.878*
			(0.123)	(0.001)	(0.059)
N	15	15	15	15	15
R^2	0.311	0.126	0.111	0.536	0.251

Notes: This table replicates the results in table IV. Panel A uses the locations of lost cities proposed by Barjamovic (2011) instead of our structural gravity estimates. Panel B uses only the subsample of cities with known locations. Robust p-values are in parentheses.

Appendix Table V: **Persistence of Economic Activity across 4000 Years, Robustness**

	log(*Population*)			log(*NightLights*)		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Barjamovic (2011) locations						
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.691		0.899	0.675**		0.889***
	(0.194)		(0.143)	(0.014)		(0.002)
$\log\left(CropYield\right)$		0.518	1.123		0.558	1.158*
		(0.681)	(0.390)		(0.408)	(0.086)
N	25	25	25	25	25	25
R^2	0.045	0.007	0.073	0.134	0.025	0.229
Panel B: known cities only						
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.171*		0.209**	0.066		0.091
	(0.062)		(0.026)	(0.203)		(0.120)
$\log\left(CropYield\right)$		1.138	1.866		0.935	1.253*
		(0.458)	(0.169)		(0.254)	(0.088)
N	15	15	15	15	15	15
R^2	0.135	0.035	0.223	0.062	0.073	0.184
Panel C: ancient site matched to the largest modern settlement						
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.235**		0.303**			
	(0.031)		(0.013)			
$\log\left(CropYield\right)$		0.716	1.791*			
		(0.516)	(0.078)			
N	24	24	24			
R^2	0.154	0.015	0.236			
Panel D: ancient site matched to the closest modern settlement						
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.252**		0.311***			
	(0.019)		(0.010)			
$\log\left(CropYield\right)$		0.478	1.582			
		(0.648)	(0.112)			
N	24	24	24			
R^2	0.207	0.008	0.283			

Notes: This table replicates the results in table V. Panel A uses the locations of lost cities proposed by Barjamovic (2011) instead of our structural gravity estimates. Since none of these locations correspond to the modern city of Ankara, this estimation features the full set of 25 cities. See text for details. Panel B uses only the subsample of 15 cities with known locations, both in estimation and in backing out $PopT^{1/\theta}\vert_{ancient}$ values. Instead of matching ancient settlements to the *sum of all* modern settlements within 20 km, Panel C and D match them to the largest modern urban settlement within 20 km and to the closest one, respectively. Since the dependent variables in Panels C and D follow administrative boundaries, *NighLights* variable does not apply to these specifications. Robust p-values are in parentheses.

Appendix Table VI: **Structural versus Naive Gravity: Locating Lost Cities**

	(1)	(2)	(3)
Durhumit	127	48	99
Hahhum	193	102	256
Kuburnat	62	70	70
Ninassa	81	93	149
Purushaddum	241	193	432
Sinahuttum	34	24	35
Suppiluliya	67	85	37
Tuhpiya	200	112	99
Washaniya	132	13	143
Zalpa	96	131	224
Mean	123.3	87.1	154.4

Notes: This table compares the estimated locations of lost cities using either our structural gravity model (8) or a naive gravity model (17) similar to that used by Tobler and Wineburg (1971). Column 1 reports the distance between the structural and naive estimates for lost city locations. Column 2 gives the distance between the location proposed by Barjamovic (2011) and the structural estimates. Column 3 gives the distance between the location proposed by Barjamovic (2011) and the naive estimates. All distances are in km.

Appendix Table VII: **Structural versus Naive Gravity: Ancient City Sizes**

	Structural estimates				Naive estimates			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Persistence of economic activity across 4000 years								
	$\log(Population)$							
$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$	0.230**	0.297**	0.387**	0.533**	0.313	0.307	-0.139	-0.140
	(0.035)	(0.015)	(0.035)	(0.063)	(0.376)	(0.367)	(0.717)	(0.677)
$\log(CropYield)$		1.781*		2.238*		0.921		-0.008
		(0.079)		(0.093)		(0.439)		(0.997)
N	24	24	10	10	25	25	10	10
R^2	0.145	0.226	0.362	0.487	0.035	0.059	0.010	0.010
Sample of cities	all	all	lost	lost	all	all	lost	lost
Panel B: Determinants of ancient city sizes								
	$\log\left(PopT^{1/\theta}\vert_{ancient}\right)$							
$\log(NaturalRoads)$	1.404**	1.783***	1.163*	1.505**	0.630***	0.678***	0.682	1.097*
	(0.013)	(0.002)	(0.092)	(0.038)	(0.000)	(0.000)	(0.301)	(0.053)
$\log(Ruggedness)$		3.189***		2.147**		0.449		1.022**
		(0.000)		(0.012)		(0.145)		(0.025)
N	25	25	10	10	25	25	10	10
R^2	0.224	0.508	0.178	0.378	0.296	0.348	0.124	0.508
Sample of cities	all	all	lost	lost	all	all	lost	lost

Notes: This table compares the estimated sizes of ancient cities using either our structural gravity model (8) and (11) or a naive gravity model (17) similar to that used by Tobler and Wineburg (1971). Panel A replicates the results in table V, using either our structural estimates (columns 1-4) or naive estimates (column 5-8). Columns 1 and 2 simply reproduce columns 1 and 3 of table V for comparison. Columns 3 and 4 replicate the specifications of columns 1 and 2 on the subset of lost cities only. Column 5 to 8 replicate the specifications of columns 1 to 4 using naive estimates instead of structural ones. To offer a meaningful comparison between structural and naive estimates, we do not drop *Purušhaddum* in columns 3-4 and 7-8, as it is an outlier (near modern Ankara) for the structural estimates (columns 3-4), but not for the naive ones (columns 7-8). Panel B replicates the results in table IV, using either our structural estimates (columns 1-4) or naive estimates (column 5-8). Columns 1 and 2 simply reproduce columns 1 and 4 of table IV for comparison. Columns 3 and 4 replicate the specifications of columns 1 and 2 on the subset of lost cities only. Column 5 to 8 replicate the specifications of columns 1 to 4 using naive estimates instead of structural ones. Robust p-values are in parentheses.