

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

GOJKO BARJAMOVIC
 THOMAS CHANEY
 KEREM COŞAR
 ALI HORTAÇSU

We analyze a large data set of commercial records produced by Assyrian merchants in the nineteenth century BCE. Using the information 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 cases, our estimates confirm the conjectures of historians who follow different methodologies. In some instances, our estimates confirm one conjecture against others. We also structurally estimate ancient city sizes and offer evidence in support of the hypothesis that large cities tend to emerge at the intersections of natural transport routes, as dictated by topography. Finally, we 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. *JEL* Codes: F1, N7, N9, R12.

I. INTRODUCTION

This article 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 (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

*This research is supported by the University of Chicago Neubauer Collegium for Culture and Society. Thomas Chaney acknowledges European Research Council grant No. 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.

© The Author(s) 2019. Published by Oxford University Press on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2019), 1455–1503. doi:10.1093/qje/qjz009.
 Advance Access publication on March 27, 2019.

a century paint a rich picture of an intraregional 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 (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 legal and financial institutions independent from the local community. Local rulers guaranteed protection for passing merchant caravans against robbers and brigands, 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 these texts and automatically isolate all tablets that jointly mention at least two cities. We systematically read those texts, which requires an intimate knowledge of the cuneiform script and Old Assyrian dialect of the ancient Akkadian language in which the records are written. 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. Our model makes predictions on the number of transactions between city pairs, which is observed in our data. Further imposing that bilateral trade frictions can be summarized by a power function of geographic distance, the model can be estimated solely on bilateral trade flows and 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 cannot 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 for 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 provides not only 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 that 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 nineteenth century BCE survives. We show that natural transportation networks—a factor usually overlooked by economists, but recognized by historians (Ramsay 1890)—are critical in explaining the hierarchy of ancient city size estimates. We 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.

I.A. Related Literature

Our article contributes to several literatures. First, we provide the earliest estimate of the gravity equation in trade, dating back to the nineteenth century BCE, about four millennia earlier than existing estimates from the mid-nineteenth 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 to locate lost cities, complementing qualitative approaches in history and archaeology 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. Multidimensional 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 thirteenth century CE (Galloway 1978). An earlier contribution (Tobler and Wineburg 1971) uses a similar data set as ours to locate Assyrian cities in Bronze Age Anatolia. Our method differs from and improves on 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 two-dimensional confidence regions. Furthermore, compared to Tobler and Wineburg (1971), we use a much larger data set that has become available for study in the meantime, systematically clean our data to identify meaningful economic exchanges, formally account for trade zeros, and compare our estimates to historical and contemporaneous evidence. We also show that our structural estimates yield more plausible estimates than multidimensional scaling, even using the same data.

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 homogeneous 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 2018). 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 article is organized as follows. Section II describes our data. Section III derives our model and estimation strategy. Section IV discusses 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 the structural gravity model to estimates from a naive gravity model.

II. ANCIENT TRADE DATA

Our data come 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

1. The transliteration of the cuneiform script in Figure I reads: *a-na kà-ri-im Kà-ne-eš₆ qí-bi-ma um-ma šit-ip-ru-ku-nu ù kà-ru-um Wa-ah-šú-ša-na-ma tup-pè-e wa-bar-tum ša Ū-là-ma ù Ša-là-tù-ar ú-šé-bi₄-lu-nim-ma ni-iš-ta-me-ma ni-ik-nu-uk-lma na-áš-ú-ni-ku-nu-tí i-ša-am-šit tup-pè-e ni-iš-ta-me-ù 2 ši-ip-ri ha-ra-an Ū-là-ma-ma 2 ši-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-ti-ma ù-za-ku-nu : ni-pà-ti I-ku-pi-a DUB.SAR ši-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 in Section IV.*



FIGURE I

Tablet Kt 83-k 117

We thank Fikri Kulakoğlu for permission to use the photo of this tablet.

BCE, with around 90% of the sample belonging to just one generation of traders, c. 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, major Assyrian merchants 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. Such records were largely 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. To some degree, this 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 c. 1840 BCE ([Manning,](#)

Barjamovic, and Lorenzen 2017). This is the main reason 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 trade data are almost 3,000 years later, coming, for example, 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 $\frac{2}{3}$ 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 the Latin alphabet, translated into modern language, published in various volumes, and recently digitized.

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 nondigitized texts.²

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 (Online 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 ,

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

Because we rarely have a description of the content of the shipments, we are unable to identify the intensive margin of trade, that is, 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 our 12,000 texts to identify any tablet that 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.

2. We rely on data amassed through 20 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.

3. We exclude Aššur, the home city of the Assyrians, from our automated search for three reasons. First, the word Aššur, which occurs c. 40,000 times, is also the name of the main Assyrian deity and occurs very often as an element of personal names (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 about 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, 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.

Second, we systematically read those 2,806 tablets, identify all mentions of cargo shipments or individual merchants' travels, and discard coincidental mentions of cities (see [Online Appendix B](#) for an example of a coincidental joint attestation). One hundred and ninety-eight 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 Durhumit until Kaneš I incurred expenses of 5 minas of refined (copper), I spent 3 minas of copper until Wahšušana, I acquired and spent small wares for a value of 4 shekels of silver. (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, $A \rightarrow B \rightarrow C$, we count three trips, $A \rightarrow B$, $B \rightarrow C$, and $A \rightarrow C$, 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 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 the lost cities. We propose using 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 (std. dev. 4.31).

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 $i \rightarrow j$ 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 come from a textual analysis of clay tablets inscribed in the cuneiform script, written by Assyrian merchants in the second 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.

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, 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 oversampling Kaneš-related trades.

III. MODEL AND ESTIMATION

III.A. 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 that plausibly describes Assyrian merchants' trading strategy. Second, it 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, and Sotelo

4. 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 note 3). See Online Appendix B for a frequency plot of directed shipment counts (Online Appendix Figure 1).

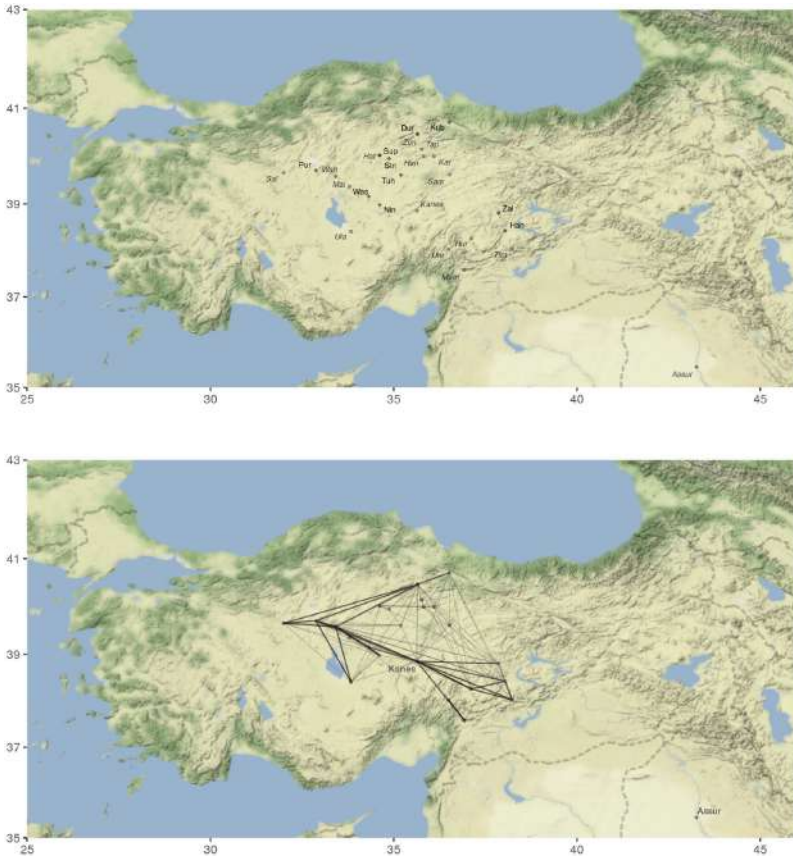


FIGURE II

Cities and Trade in Anatolia in the Bronze Age

In the top panel, known cities are gray diamonds (Hanaknak, Hattuš, Hurama, Kaneš, Karahna, Malitta, Mamma, Šalatuwar, Šamuha, Tapaggaš, Timelkiya, Ulama, Unipsum, Wahšušana, Zimišhuna), and lost cities are black circles (Durhumi, 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 $N_{ij}^{data} > 3$.

(2012) and other modern gravity estimates such as [Silva and Tenreiro \(2006\)](#): unlike with modern trade data, we do not know the location of some cities. Instead we use 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 tradeable 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 $\tau_{ij}c_i(\omega) < \tau_{ij}c_j(\omega)$, she⁵ 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 $\tau_{ij}c_j(\omega)$ for a profit, without the threat of being undercut by local sellers who could at best buy τ_{ij} units locally at cost $c_j(\omega)$ to sell at cost $\tau_{ij}c_j(\omega)$ 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,

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

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

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. Because 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, 72). In the absence of any systematic information on how profits are accrued and spent, we do not model profits explicitly.

probability that a shipment sourced by destination j originates from i is equal to

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

We define two more conditional probabilities (see [Online 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 or from a lost city, destination j sources good ω from known origin i ,

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

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

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 data set, 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 \(4\)](#) 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 shipment counts and does not require knowledge of the value of shipments. This property is crucial to us, as our data set contains information on shipment

[Eaton, Kortum, and Sotelo \(2012\)](#) show that if profits are redistributed using an outside good, the predictions remain.

counts but not on the value of shipments. Among modern trade models, this feature is unique to the [Eaton and Kortum \(2002\)](#) model and is one of our main motivations for using it. Note that this model also predicts that trade shares in value are equal to trade shares in counts.⁷ We use this property to close our model in general equilibrium to derive counterfactual measures of city sizes.

III.B. Estimation

Our estimation proceeds in three steps. First, we parameterize 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 with conventional estimates of the gravity equation in trade. Third, we combine our estimates to compute a measure of the 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-step procedure formally estimates parameters such that the gravity model fits the data as closely as possible and 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 parameterize the symmetric trade cost function as

$$(5) \quad \tau_{ij}^{-\theta} = \mu \cdot (\text{Distance}_{ij}(\varphi_i, \lambda_i; \varphi_j, \lambda_j))^{-\zeta}.$$

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

A scaling factor, μ , controls for measurement units, and ζ is the distance elasticity of trade. The function $Distance_{ij}(\varphi_i, \lambda_i; \varphi_j, \lambda_j)$ maps geocoordinates into geographic distances, in km.⁸

We use Euclidean distances, that is, “as the crow flies” distances, instead of least-effort distances that would account for the topography of the local terrain. There are two reasons for this choice. First, when estimating our gravity model, we need to solve a complex nonlinear minimization program—see [equation \(8\)](#) 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 [Section IV.B](#) and [V.B](#)). Not bringing topographical data into our estimation gives credence to those validity checks.

Because our model is based on [Eaton et al. \(2012\)](#)’s finite sample version of the Eaton and Kortum model, predicted trade shares between city pairs are random. In particular, they can be 0, as often happens in the data, if the lowest realized cost to deliver a good from i to j is higher than the lowest realized cost from all other origins. Beyond this finite sample randomness, we

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

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

where 37.9 degrees north is the median latitude among known Assyrian cities. For locations in these latitudes, the difference between this Euclidean formula and the more precise Haversine formula is negligible. This approximation considerably speeds up the estimation. We also need to know internal trade frictions. Because we do not observe internal trades, we cannot estimate within-city transactions costs. We instead normalize internal distances, $Distance_{ii} = 30$ km, 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.

can easily add a multiplicative disturbance term to the trade cost function (5), without altering our estimation strategy.¹⁰

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

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

with $\alpha_i = w_i^{-\theta} T_i$ and $\beta_j = \frac{\mu}{\sum_{k \in \mathcal{K} \setminus \{j\}} w_k^{-\theta} T_k (\tau_{kj})^{-\theta}}$. 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, $\frac{N_{ij}^{data}}{\sum_{k \in \mathcal{K} \setminus \{j\}} N_{kj}^{data}}$; bilateral distances between known cities, Distance_{ij} ; and origin and destination fixed effects to control for α_i and β_j . This closely follows the procedure in Eaton, Kortum, and Sotelo (2012), with the only difference that we derive conditional probabilities (noninternal trade among the subset of known cities only) and not unconditional ones.

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),

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

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

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

By construction, both realized (random) and expected trade shares lie in $[0, 1]$ and add up to 1.

Step 2 of our estimation uses the distance elasticity $\hat{\zeta}$ from step 1, and our data set 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 β

$$\beta = (\alpha_1, \dots, \alpha_{Kaneš-1}, \alpha_{Kaneš+1}, \dots, \alpha_{K+L}, \varphi_{K+1}, \lambda_{K+1}, \dots, \varphi_{K+L}, \lambda_{K+L})'$$

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

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

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 moment condition (7), we estimate the parameters β by a method of moments, solving the following nonlinear least squares minimization problem,¹¹

$$(8) \quad \beta = \arg \min_{\beta} \sum_j \sum_{i \neq j} \left(\frac{N_{ij}^{data}}{\sum_{k \neq j} N_{kj}^{data}} - \frac{\alpha_i \text{Distance}_{ij}^{-\hat{\zeta}}}{\sum_{k \neq j} \alpha_k \text{Distance}_{kj}^{-\hat{\zeta}}} \right)^2$$

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 0, 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 , under the constraint that shares add up to 1. Furthermore, if i 's location is lost, so that both $\text{Distance}_{i,j_1}$ and $\text{Distance}_{i,j_2}$ are unknowns to be estimated,

11. See Online Appendix A for details on how we address the issue of local minima in our nonlinear minimization problem.

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), that is, 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,

$$(9) \quad X_i = \sum_k X_{ki} = w_i Pop_i,$$

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

$$(10) \quad Size_i \equiv \frac{w_i^{autarky} Pop_i}{P_i^{autarky}} \propto Pop_i T_i^{\frac{1}{\theta}},$$

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 $T_i^{\frac{1}{\theta}}$. If, for instance, trade frictions or the size of other cities were to change, this measure would remain invariant. This measure can be computed using our

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 equation (10), see Eaton and Kortum (2002), equation (15).

parameter estimates and an assumption for the trade elasticity θ only,¹⁴

$$(11) \quad Size_i \propto Pop_i T_i^{\frac{1}{\theta}} \propto \hat{\alpha}_i^{1+\frac{1}{\theta}} \sum_k \widehat{Distance}_{ki}^{-\hat{\zeta}} \hat{\alpha}_k,$$

where we use $\theta = 4$ from [Simonovska and Waugh \(2014\)](#). Because 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 $\hat{\alpha}_i^{1+\frac{1}{\theta}}$ corresponds to an exporter fixed effect, the propensity of a city to trade after controlling for distance. The extra term $\sum_k \widehat{Distance}_{ki}^{-\hat{\zeta}} \hat{\alpha}_k$ adjusts for the endogenous response of factor prices in general equilibrium: if city i is either centrally located or located near large trading partners ($\widehat{Distance}_{ki}$ small or $\hat{\alpha}_k$ large for some k 's), it faces an upward pressure on the price of local fixed factors. This depresses its exports by eroding its competitiveness. In autarky, this depressing effect of trade on factor prices disappears. [Equation \(11\)](#) formally adjusts for this endogenous factor price response.

III.C. Standard Errors

Robust (White) standard errors are computed analytically and account for heteroskedasticity and the two-step nature of our estimation (PPML then NLLS). To visually gauge 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 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 geocoordinates of lost city l , $\mathcal{N}(\hat{\beta}'_l, \hat{\Sigma}'_l)$ with mean $\hat{\beta}'_l = (\hat{\varphi}_l, \hat{\lambda}_l)$ and covariance matrix $\hat{\Sigma}'_l$. They

14. See [Online Appendix A](#) for a formal derivation. To recover $Pop_i T_i^{\frac{1}{\theta}}$ 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 in practice.

account not only for the precision of the latitude and longitude of city l but also for the covariance of those geocoordinates.

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

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

where $\textit{Distance}(\hat{\varphi}_l, \hat{\lambda}_l; \varphi, \lambda)$ is the distance between the estimated location for l and a location drawn from our estimated $\mathcal{N}(\hat{\beta}_l, \hat{\Sigma}_l)$. Heuristically, it means that our point estimates are within this distance, $\textit{precision}(l)$ 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 with 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. [Table III](#), Panel A 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, nearly equal to our ancient estimate.

15. See [Online Appendix A](#) for analytical formulas of standard errors, confidence regions, $\textit{precision}(l)$ and their derivations.

TABLE II
LOST CITIES' GEOCOORDINATES

	Latitude	(Std. err.)	Longitude	(Std. err.)	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
Ninašša	38.977	(0.778)	34.614	(0.482)	0.86
Purušhaddum	39.71	(1.54)	32.872	(0.669)	0.774
Šinahuttum	39.956	(0.333)	34.866	(0.165)	0.863
Šuppiluliyā	40.021	(1,022.82)	34.618	(58.796)	1.0
Tuhpiya	39.611	(0.18)	35.199	(0.307)	0.528
Wašhaniya	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 geocoordinates, 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. The last column gives the estimated correlation between latitude and longitude, used to compute confidence regions.

Figure III shows maps with our point estimates and confidence regions for each lost city separately. A closed circle 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 Table III, panel B: all but three of our measures of precision are smaller than 100 km (60 miles), and less than 50 km in four cases. This can 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 (2008); the “▲ B” sign corresponds to the site suggested by historian Gojko Barjamovic (2011). Those historians base their proposals on a careful analysis of ancient texts, ancient itineraries, topographical studies, surviving toponyms, and so on.¹⁷ This comparison allows us to confront our estimates, obtained by a purely quantitative method—a structural gravity estimation, to

16. The confidence regions for Šuppiluliyā are not shown because they are too wide.

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

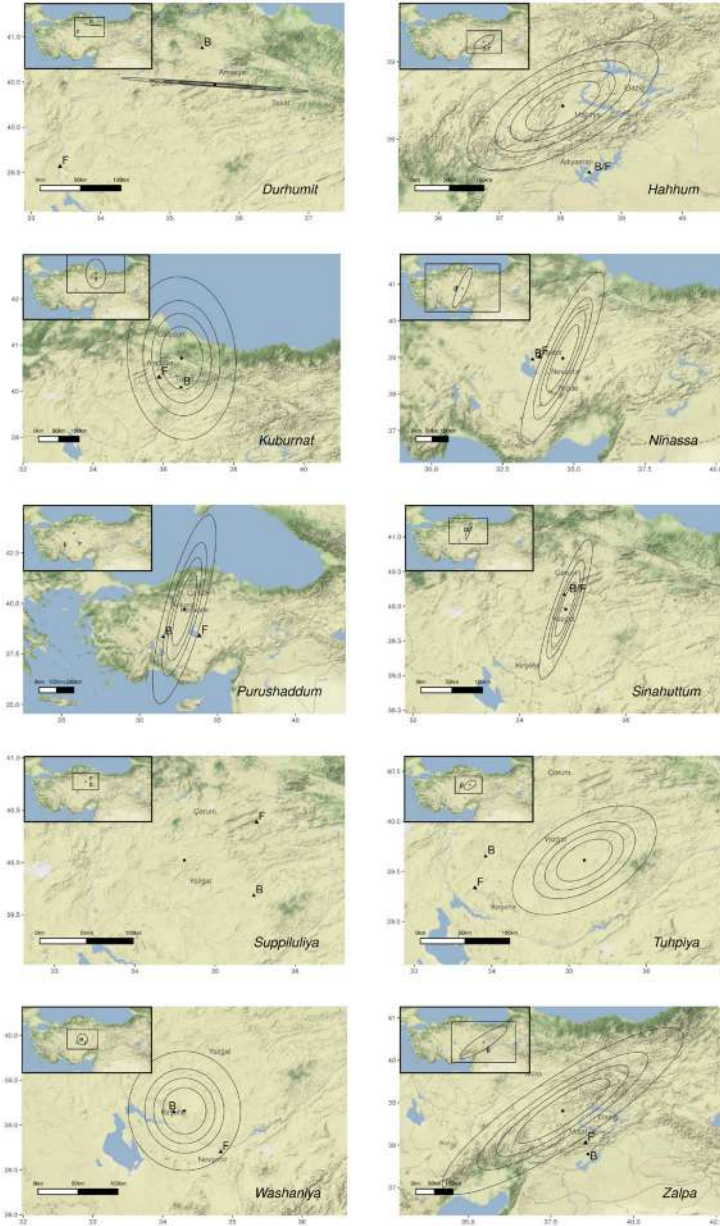


FIGURE III
Locating Lost Cities

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							
	$Pop_i T_i^{\frac{1}{\theta}}$	(Std. err.)	Precision	Distance to historians' proposals, in km			
				Forlanini	(Std. err.)	Barjamovic	(Std. err.)
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)
Ninašša	0.21	(0.45)	87	71	(47)	93	(49)
Purušhaddum	0.076	(0.1)	154	168	(75)	193	(112)
Šinahuttum	1.515	(2.021)	34	24	(20)	24	(20)
Šuppiluliyā	0.012	(5.246)	89,914	89	(54,307)	85	(54,240)
Tuhpiya	0.579	(0.912)	38	128	(35)	112	(33)
Wašhaniya	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							
	$Pop_i T_i^{\frac{1}{\theta}}$	(Std. err.)					
Hanaknak	2.062	(3.679)					
Hattuš	2.967	(4.183)					
Hurama	5.091	(9.564)					
Kaneš	100.0	(70.902)					
Karahna	0.008	(0.084)					
Malitta	0.057	(0.096)					
Mamma	114.683	(189.19)					
Šalatuwar	1.513	(4.836)					
Šamuha	3.508	(4.847)					
Tapaggaš	0.091	(0.218)					
Timelkiya	101.922	(129.238)					
Ulama	0.099	(0.498)					
Unipsum	44.294	(64.469)					
Wahšušana	1.416	(3.53)					
Zimišhuna	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, $Pop_i T_i^{\frac{1}{\theta}}$, defined in equation (10), is the counterfactual real output of city i if it were to move to complete autarky. Precision, measured in km, is defined in equation (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 are in parentheses.

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 or 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. Because we do not use the historians' conjectures as an 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 toward the north and west of their proposal. It lies in the Taurus mountain range, a rugged, high-altitude, snow-covered area. As we do not impose our gravity estimates to be in hospitable locations, nothing prevents this from being the case. Forlanini and Barjamovic impose the realistic constraint that cities are in accessible and suitable places and draw on historical information about its location on the Euphrates River, which the gravity estimate ignores.

In the case of Tuhpiya, Forlanini's 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, Šuppiluliyā 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 Šuppiluliyā, our estimate is not very far from their historians' proposals (about 90 km from both), but it is so imprecise that we cannot draw any inference: the precision for Šuppiluliyā (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 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, Šuppiluliyā, and Zalpa, which are imprecisely estimated. Furthermore, for Ninašša, our gravity estimates favor the location proposed by Forlanini over that given by Barjamovic, whereas 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 1960s.

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 multistop 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 Kaneš from your envoys and the Port Authorities of Wahšušana. We have heard the tablets that the Station(s) in Ulama and Šalatuwar have brought us, and we have sealed them and (hereby) convey them on to you. On the day we heard the tablets, we sent two messengers by way of Ulama and two messengers by way of Šalatuwar to Purušhaddum to clear the order. We will send you the earlier message that they brought us so as to keep you informed. The Secretary Ikūn-pīya is our messenger. (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 → Ulama → Purušhaddum, and Wahšušana → Šalatuwar → Purušhaddum. For both of these itineraries, two cities are known

(Wahšušana and Ulama for the first, Wahšušana and Šalatuwar for the 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 multistop 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 with a direct trip without the intermediate stop. Formally we impose

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

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 $\delta\%$ detour compared to going straight from A to C .

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

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

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.

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

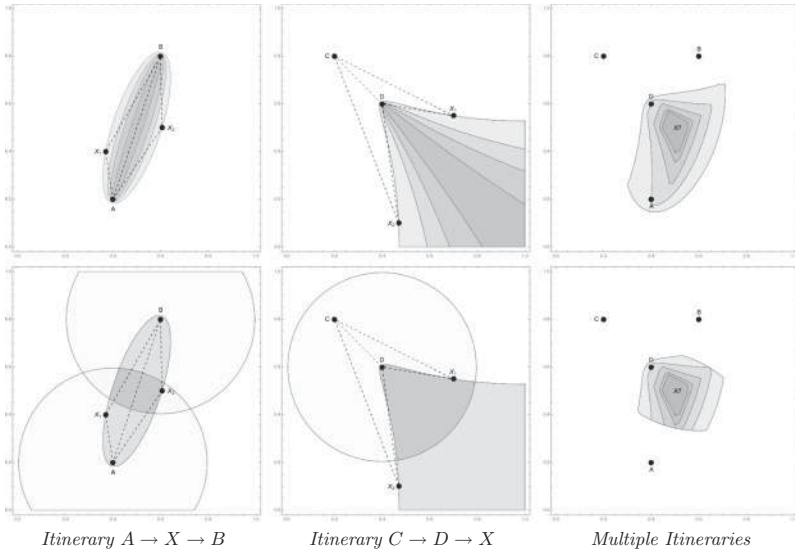


FIGURE IV

Constraints on Lost Cities from Merchants' Itineraries

The top row only imposes the short detour constraint; the bottom row further imposes the pit stop constraint. The left figures show the example of an itinerary of the type $A \rightarrow X \rightarrow B$, where A and B are known, 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 with 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 $C \rightarrow D \rightarrow X$, with C and D known and X lost. The right figures jointly impose constraints from both itineraries. Darker shades of gray correspond to shorter detours and shorter segments between pit stops.

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

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, 2017; Palmisano and Altaweel 2015).

We systematically collect all mentions of multistop 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. [Online Appendix D](#) provides further details.

We present our results in a series of maps in [Figure V](#). Each map depicts the admissible region for a given city (dashed line), using the 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 our point estimate and 90th percentile confidence region from estimating our gravity model (8) (solid line ellipse), as well as the locations proposed by [Forlanini \(2008\)](#) and [Barjamovic \(2011\)](#) ("▲ F" and "▲ B" signs).

Our gravity estimates for the location of lost cities all lie within their admissible regions¹⁹ and are therefore compatible with the qualitative information from merchants' itineraries, with the exception of Puruṣhaddum. Puruṣhaddum 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](#), 242f). Unfortunately, its peripheral position in the Assyrian network makes it difficult for the gravity model to suggest a location (see [Section IV.C](#)), with large and imprecise confidence regions ([Figure III](#)). The addition of the constraints imposed on 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 V](#)). 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](#), 314).

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

19. A few admissible sets are wide and do not impose a strict constraint on the location of lost cities, for example, for Tuḫpiya. The reason we cannot impose a stricter set of parameters is that 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, for example, for Kuburnat.

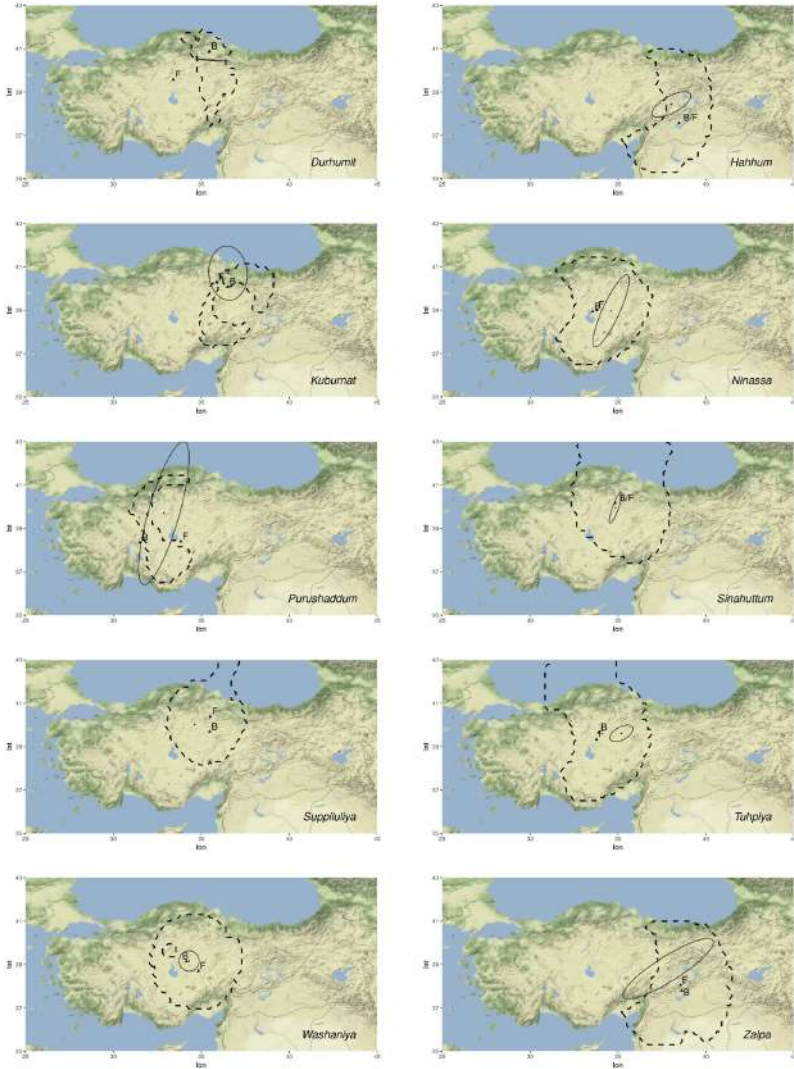


FIGURE V

Constraints on Lost Cities from Merchants' Itineraries

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.



FIGURE VI

Proof of Concept: Recovering Fictitiously Lost Cities

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.

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 reestimate all other parameters solving the same nonlinear least squares problem as equation (8).

Figure VI presents the results of our within-sample predictions.²⁰ For each known city, an arrow goes from the true site,

20. Online Appendix Table 1 lists the geocoordinates for all known cities, both true and estimated, as well as the distances, in km, between the true and estimated locations. As a robustness check, we run an alternative proof of concept exercise, estimating not just the location of one fictitiously lost city but reestimating the locations of all 10 truly lost cities as well. The results are presented in Online Appendix Table 2. The results are similar, although less accurate, as we are compounding measurement error for one fictitiously lost city with that of the 10 truly lost cities.

denoted by a diamond, toward the estimated location, denoted by a closed circle. 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 ,

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

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

We propose to apply formula (13) to 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). 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.²¹

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 marketplace for copper during the period of Assyrian trade and is mentioned more than 200 times in the trade records from Kaneš. It reappears in documentation of the Hittite state in the fourteenth–thirteenth 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, 261–265; Cammarosano and Marizza 2015, 180f; Kryszewski 2016, 343ff; Corti 2017, 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 V). Our gravity estimation ranks the unexcavated archaeological site of Ayvalpınar as the most likely candidate for Durhumit. An Assyrian seal carved in a workshop at Kaneš (Lassen 2014, 118) was found on the surface of the site, a strong indication 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, 386; Dönmez 2017, 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 a 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 subjected to systematic excavation. It is a viable candidate for Durhumit, although there may be other, possibly better proposals (Dönmez 2017; Corti 2017, 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.

21. The full set of results are presented in Online Appendix Table 3. 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 equation (13).

The gravity estimate corresponds fairly well to the conjecture proposed by historians (Forlanini 2008; Barjamovic 2011). The most likely candidate is Yassihöyük, which has come under excavation within the past decade and revealed findings dated to the period of Assyrian trade (Omura 2016). Excavations at a number of sites along the historically important route leading southeast from Yassihöyük to Kayseri have revealed remains from the period (Weeden and Matsumura 2017, 108), including Suluca, Zank, and Topaklı on the list of likely candidates. The site of Kırşehir Kalehöyük is also 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 (erected in 1230 CE) and a high school built on it (Adbelli 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 always ought to 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 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, compare 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 ($Pop_i T_i^{\frac{1}{\beta}}$) are presented in Table III, Panels B and C. We note they 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.²² 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. Because 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 under modern cities.

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

To probe the determinants of ancient city sizes, we project our ancient size estimate ($Pop_i T_i^{\frac{1}{\theta}}$) on two geographic observables: terrain ruggedness, a local measure of the defensive advantage of a site,²³ and a measure of “global” advantage.

Our concept of global advantage is novel. We define, for each site, a measure of its centrality in the transportation network. In developing this measure, we build on 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 exploratory travels in Asia Minor, Ramsay suggested that the key to understanding urban geography in classical antiquity is the realization that the local terrain allows only 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

22. A Welch’s *t*-test of equality between the sizes of known and lost cities gives a *p*-value of .29.

23. 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.

trading centers, was further elaborated and advanced by French (1993).²⁴

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 (three for a T-crossing, four 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 between two and five.²⁵ While capturing the location of cities vis-à-vis the actual historical road network, this measure has two shortcomings. First, there is about a 2,000-year gap between the Middle Bronze Age and the Roman period. This concern is partially alleviated by the fact that Roman roads follow older trails (Ramsay 1890; French 1993). The other shortcoming concerns the potential endogeneity of the road network: it is plausible that roads endogenously connect large cities, so that large cities “cause” roads, rather than the reverse.

24. 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 (2018) 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.

25. 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. Sixteen out of 25 cities have a road score of 2. Two cities, Kaneš and Ulama, have road scores of 5.

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's \(1984\)](#) formula to compute travel times for a normal human 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's \(1959\)](#) algorithm to compute the optimal route between any two points (see [Online Appendices C and E](#) for details). Our approach in defining natural routes finds support in [Palmisano \(2013, 2017\)](#) and [Palmisano and Altaweel \(2015\)](#), 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 assume 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 VII](#) 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 an input.

[Table IV](#) presents the results from the estimation of various specifications of

$$(14) \quad \ln\left(Pop_i T_i^{\frac{1}{\sigma}}\right) = a + b \cdot \ln(Roads_i) + c \cdot \ln(Ruggedness_i) + u_i.$$

$Roads_i$ 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)–(5),




FIGURE VII
Natural Roads Scores

TABLE IV
DETERMINANTS OF ANCIENT CITY SIZES

	$\log \left(PopT^{\frac{1}{\theta}} \right)$				
	(1)	(2)	(3)	(4)	(5)
$\log(NaturalRoads)$	1.404** (.013)			1.783*** (.002)	
$\log(RomanRoads)$		1.990 (.378)			2.387 (.220)
$\log(Ruggedness)$			2.371** (.012)	3.189*** (.000)	2.495*** (.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 equation (14). Each observation is an ancient city. The dependent variable $PopT^{\frac{1}{\theta}}$ is the ancient city size estimate. Explanatory variables are: *NaturalRoads*, the average natural road score of the area within 20 km of the ancient city as defined in Subsection VB; *RomanRoads*, the number of Roman roads radiating from an intersection within 20 km of the ancient city (French 2016); and *Ruggedness*, the Terrain *Ruggedness* Index (Riley, DeGloria, and Elliott 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. ** and *** indicate significance at the 5%, and 1% levels, respectively.

we control for *Ruggedness_i*, a measure of how rugged the terrain is around i (Riley, DeGloria, and Elliott 1999).

We find robust and significant evidence in support of the road-knot hypothesis. The more road intersections near a city, the larger it is. While the *RomanRoads* variable has a positive but nonsignificant effect (columns (2) and (5)), our a priori measure of the connectedness of a city to the natural road network,




FIGURE VIII

The Topographical Determinants of Ancient City Sizes

The horizontal axis is the logged index for road-knots, *NaturalRoads* for ancient cities. The vertical axis is the logged estimate of $PopT^{\frac{1}{5}}$ for ancient cities. The regression line corresponds to [Table IV](#), column (1).

NaturalRoads, is strongly significant, with a p -value below .015 both on its 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.²⁶ [Figure VIII](#) 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

26. The results are robust to using the locations proposed by [Barjamovic \(2011\)](#) instead of our estimated locations for lost cities and to restricting the sample to known cities only. See [Online Appendix Table 4](#).

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 Table IV, column (3) 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 Millennia*

Next we confront our ancient size estimates with 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 project two alternative measures of modern city sizes on our ancient size estimates and a control for geographic amenities.

We drop one outlier from our sample, Puruṣhaddum, which is matched with the modern city of Ankara. A minor provincial town at the turn of the twentieth 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 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.²⁷

Our first measure of modern size, $Population_i$, measures the total urban population living within 20 km of ancient city i .²⁸ We

27. 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 twentieth century CE.

28. For lost cities, we use our own estimates from Section IV.A. For two ancient cities, Tuḥpiya and Samuha, there are no modern-day urban population centers within 20 km of their coordinates. We winsorize their populations to the smallest population within the sample. The results are robust to dropping them and

TABLE V
PERSISTENCE OF ECONOMIC ACTIVITY ACROSS 4,000 YEARS

	log (<i>Population</i>)			log (<i>NightLights</i>)		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log \left(PopT^{\frac{1}{\theta}} _{ancient} \right)$	0.230** (.035)		0.297** (.015)	0.124** (.036)		0.178*** (.008)
$\log (CropYield)$		0.727 (.507)	1.781* (.079)		0.777 (.143)	1.407*** (.003)
<i>N</i>	24	24	24	24	24	24
<i>R</i> ²	0.145	0.015	0.226	0.143	0.059	0.312

Notes. This table presents the estimation from various specifications of equation (15). Each observation is an ancient city after dropping Puruṣhaddum-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 $PopT^{\frac{1}{\theta}} |_{ancient}$, the ancient city size estimate from Table III, 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 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 significant or robust. Robust *p*-values are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

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

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

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.²⁹

estimating with a sample of 22 cities. The results are also robust to using Bar-jamovic (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 Online Appendix Table 5 for those robustness checks.

29. 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

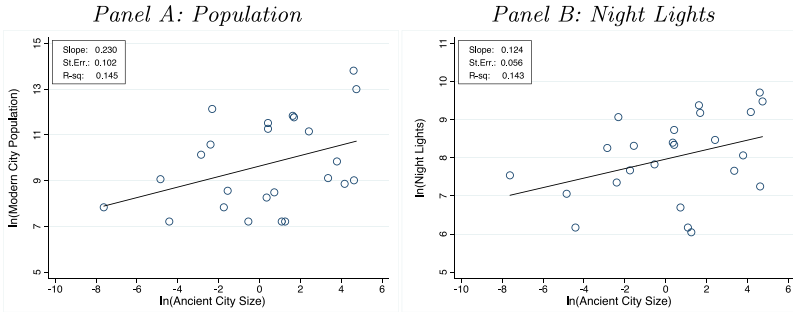


FIGURE IX
Ancient and Modern City Sizes

The horizontal axis is the logged estimate of ancient city sizes, $PopT^{\frac{1}{5}}$. The vertical axis is the logged 2012 urban population (Panel A) and night luminosity (Panel B) within 20 km of the ancient city's location. The regression lines correspond to Table V, column (1) in Panel A, and column (4) in Panel B.

Table V, columns (1) and (4) 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 IX. 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 twentieth century BCE to the twenty-first 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. Although our results do not offer a definitive explanation for this persistence, two mechanisms highlighted in the 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, and so on, people seem to

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 effects on modern sizes robust. We therefore exclude those controls from our regression.

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 constraints throughout history, topography may have continued to affect the relative size of cities. We hope to 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, $Size_i \propto Pop_i T_i^{\frac{1}{\sigma}}$. However, we can use our findings in [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: $\ln(Population_i) \approx constant + 0.23 \ln(Pop_i T_i^{\frac{1}{\sigma}})$ (see [Table V](#), column (1)). Using this correspondence naively, we test whether the distribution of ancient population sizes follows Zipf's law, applying the methodology advocated by [Gabaix and Ibragimov \(2011\)](#), by estimating

$$(16) \quad \ln(Rank_i - 0.5) = g - h \cdot \ln(Population_i) + w_i,$$

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

[Figure X](#) presents the results of estimating [equation \(16\)](#). The data suggest 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 std. err. = 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 millennia. It should of course be interpreted with due caution. First, our variable *Population* is estimated with error, so our estimate for the Zipf exponent may be biased. Second, we have




FIGURE X

Zipf's Law for Cities in the Bronze Age

The horizontal axis is the log of the population of ancient cities, where we apply the transformation $\ln(\text{Population}_i) = 0.23 \cdot \ln(\text{Pop}_i T_i^{\frac{1}{\beta}})$ to our structural estimate of ancient city sizes using [Table V](#), column (1). The vertical axis is the log of rank minus one-half. The regression line corresponds to the estimate of [equation \(16\)](#).

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 the multidimensional scaling exercise in [Tobler and Wineburg \(1971\)](#). To provide a meaningful comparison, we use the same measure of bilateral trade flows as in our main analysis but perform an estimation similar to theirs.³⁰ First, we define an undirected measure of

30. [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

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 : $I_{ij} = N_{ij}^{data} + N_{ji}^{data}$ where N_{ij}^{data} 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,

$$\begin{aligned} I_{ij} &= k \cdot \frac{P_i P_j}{\text{Distance}_{ij}^2} \Leftrightarrow \frac{P_i P_j}{I_{ij}} \\ &= \beta \left((\varphi_i - \varphi_j)^2 + \cos^2 \left(\frac{37.9}{180} \pi \right) (\lambda_i - \lambda_j)^2 \right), \end{aligned}$$

where k and β are simple multiplicative constants, the φ 's are latitudes, and the λ 's are longitudes. Collecting β and the coordinates of lost cities in the vector $\theta = (\beta, (\varphi_{K+1}, \lambda_{K+1}), \dots, (\varphi_{K+L}, \lambda_{K+L}))$, we estimate this model by nonlinear least squares,

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

We compare our results using structural versus naive estimates.³¹

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

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.

31. Further details are presented in [Online Appendix Table 6](#) for comparisons of the estimates of lost cities locations and in [Online Appendix Table 7](#) for comparisons of the estimates of city sizes.

from a city: the correlation between $\ln(\text{Pop}_i T_i^{\frac{1}{\theta}})$ 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.³² Moreover, our structural estimates for city sizes are significantly related to modern city sizes, whereas naive estimates are not. For instance, using size estimates from naive gravity in estimating equation (15) to test for the persistence of economic activity over 4,000 years gives an insignificant coefficient of logged modern population on logged ancient size (0.313 with a p -value of .376 to be compared with 0.230 with a p -value of .035 for our structural estimate), and a poor fit ($R^2 = 0.035$ compared with $R^2 = 0.145$ for our structural estimate). Our structural estimates for size are also significantly related to measures of access to natural roads, 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 4,000 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

32. 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, $\ln(\text{Pop}_i T_i^{\frac{1}{\theta}})$, which controls for distance and general equilibrium, is 0.98.

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). Although 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 4,000 years, ancient economic size predicts the income and population of corresponding regions in present-day Turkey.

HARVARD UNIVERSITY
SCIENCES PO

UNIVERSITY OF VIRGINIA AND CENTER FOR ECONOMIC POLICY
AND RESEARCH

UNIVERSITY OF CHICAGO AND NATIONAL BUREAU OF ECONOMIC
RESEARCH

SUPPLEMENTARY MATERIAL

An [Online Appendix](#) for this article can be found at *The Quarterly Journal of Economics* online. Data and code replicating tables and figures in this article can be found in [Barjamovic et al. \(2019\)](#), in the Harvard Dataverse, doi: [10.7910/DVN/P81WBZ](https://doi.org/10.7910/DVN/P81WBZ).

REFERENCES

- Adıbelli, I.A., “2012 Kırşehir Kale Hoyük Kazısı,” *Kazı sonuçları toplantısı*, 35 (2013), 296–310.
- Barjamovic, G., “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*, K. Kristiansen, T. Lindkvist and J. Myrdal, eds. (Cambridge: Cambridge University Press, 2018), 113–142.
- Barjamovic, Gojko, “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*, J. G. Dercksen, ed. (Leiden: Nederlands Instituut voor het Nabije Oosten, 2008), 87–100.
- , *A Historical Geography of Anatolia in the Old Assyrian Colony Period* (Copenhagen: Museum Tusulanum Press, 2011).
- , “A Commercial Geography of Anatolia: Integrating Hittite and Assyrian Texts, Archaeology and Topography,” in *Hittite Landscape and Geography*, L. Ullmann and M. Weeden, eds. (Leiden: Brill, 2017), 311–318.
- Barjamovic, Gojko, Thomas Chaney, Kerem Coşar, and Ali Hortaçsu, “Replication Data for: ‘Trade, Merchants, and the Lost Cities of the Bronze Age,’” Harvard Dataverse (2019), doi:10.7910/DVN/P81WBZ.
- Barjamovic, Gojko, Thomas Klitgaard Hertel, and Mogens Trolle Larsen, *Ups and Downs at Kanesh: Chronology, History and Society in the Old Assyrian Period* (Leiden: Nederlands Instituut voor het Nabije Oosten, 2012).
- Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, “Plants and Productivity in International Trade,” *American Economic Review*, 93 (2003), 1268–1290.
- Bleakley, Hoyt, and Jeffrey Lin, “Portage and Path Dependence,” *Quarterly Journal of Economics*, 127 (2012), 587–644.
- Cammarosano, M., and M. Marizza, “Das Land Tumanna und sein König in den hethitischen Quellen,” *Die Welt des Orients*, 45 (2015), 158–192.
- Chaney, Thomas, “The Gravity Equation in International Trade: An Explanation,” *Journal of Political Economy*, 126 (2018), 150–177.
- Cornelius, Friedrich, “Neue Aufschlüsse zur hethitischen Geographie,” *Orientalia*, 32 (1963), 233–245.
- Corti, C., “The North: Hanhana, Hattena, İstahara, Hakpiš, Nerik, Zalpuwa, Tummana, Pala and the Hulana River Land,” in *Hittite Landscape and Geography*, L. Ullmann and M. Weeden, eds. (Leiden: Brill, 2017), 219–238.
- Cosar, A. Kerem, and Banu Demir, “Domestic Road Infrastructure and International Trade: Evidence from Turkey,” *Journal of Development Economics*, 118 (2016), 232–244.
- Cronon, William, *Nature’s Metropolis: Chicago and the Great West* (New York: Norton, 2009).
- Dalgaard, Carl-Johan, Nicolai Kaarsen, Ola Olsson, and Pablo Selaya, “Roman Roads to Prosperity: Persistence and Non-Persistence of Public Goods Provision,” CEPR Discussion Paper DP12745, 2018.
- Davis, Donald R., and David E. Weinstein, “Bones, Bombs, and Break Points: The Geography of Economic Activity,” *American Economic Review* 92, (2002), 1269–1289.
- Dercksen, Jan Gerrit, *The Old Assyrian Copper Trade in Anatolia* (Leiden: Nederlands Instituut voor het Nabije Oosten, 1996).
- Dijkstra, Edsger W., “A Note on Two Problems in Connexion with Graphs,” *Numerische mathematik*, 1 (1959), 269–271.
- Disdier, Anne-Célia, and Keith Head, “The Puzzling Persistence of the Distance Effect on Bilateral Trade,” *Review of Economics and Statistics*, 90 (2008), 37–48.
- Dönmez, Ş., “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*, M. Alparslan, ed. (Istanbul: Türk Eskiçağ Bilimleri Enstitüsü, 2017), 75–90.
- Eaton, Jonathan, and Samuel Kortum, “Technology, Geography and Trade,” *Econometrica*, 70 (2002), 1741–1779.
- Eaton, Jonathan, Samuel Kortum, and Sebastian Sotelo, “International Trade: Linking Micro and Macro,” NBER Working Paper no. 17864, 2012.

- Erol, H., "İna maqāt ubri lu kaspam lu hurāsam ašamma ... 'Yabancılar geldiğinde ya gümüş ya da altın satın alacağım ...'," in *Cahit Günbattı'ya Armağan. Studies in Honour of Cahit Günbattı*, I. Albayrak, H. Erol and M. Çayır, eds. (Ankara: Ankara Üniversitesi Dil ve Tarih Coğrafya Fakültesi, 2013), 87–98.
- Forlanini, Massimo, "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*, Karl Strobel, ed. (Firenze: LoGisma, 2008), 145–188.
- , "South Central: The Lower Land and Tarhuntašša," in *Hittite Landscape and Geography*, L. Ullmann and M. Weeden, eds. (Leiden: Brill, 2017), 239–252.
- French, D. "Colonia Archelais and Road-Knots," in *Aspects of Art and Iconography: Anatolia and Its Neighbors: Studies in Honor of Nimet Özgüç*, M. J. Mellink, E. Porada and T. Özgüç, eds. (Ankara: Türk Tarih Kurumu Basımevi, 1993), 201–207.
- French, David, *Roman Roads and Milestones of Asia Minor. Vol. 3 Milestones, Fasc. 9. An Album of Maps, Electronic Monograph* (Ankara: British Institute at Ankara, 2016).
- Gabaix, Xavier, and Rustam Ibragimov, "Rank- 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business & Economic Statistics*, 29 (2011), 24–39.
- Galloway, Patricia, "Restoring the Map of Medieval Trondheim: A Computer-Aided Investigation into the Nightwatchmen's Itinerary," *Journal of Archaeological Science*, 5 (1978), 153–165.
- Gardin, Jean-Claude, and Paul Garelli, "Étude des Établissements assyriens en Cappadoce par ordinateurs," *Annales. Histoire, Sciences Sociales*, 16 (1961), 837–876.
- Günbattı, Cahit, "Karumlar Arasındaki Mektuplaşmalardan Yeni Örnekler," in *XXXIV. Uluslararası Assirioloji Kongresi, 6-10/VII/1987-Istanbul*, H. Erkanal, V. Donbaz and A. Uğuroğlu, eds. (Ankara: Türk Tarih Kurumu Basımevi, 1998).
- Hertel, Thomas Klitgaard, *Old Assyrian Legal Practices: Law and Dispute in the Ancient Near East* (Leiden: Nederlands Instituut voor het Nabije Oosten, 2013).
- Hodler, Roland, and Paul A. Raschky, "Regional Favoritism," *Quarterly Journal of Economics*, 129 (2014), 995–1033.
- IIASA/FAO, *Global Agro-Ecological Zones (GAEZ v3.0)* (Laxenburg, Austria: IIASA and Rome: FAO, 2012).
- Kendall, David G., "Maps from Marriages: An Application of Non-metric Multi-dimensional Scaling to Parish Register Data," *Mathematics in the Archaeological and Historical Sciences* (1971), 303–318.
- Krugman, Paul, "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99 (1991), 483–499.
- Kryszeń, A., *A Historical Geography of the Hittite Hinterland (Alter Orient und Alter Testament 437)* (Münster: Ugarit, 2016).
- Langmuir, Eric, *Mountaineering and Leadership: A Handbook for Mountaineers and Hillwalking Leaders in the British Isles* (Edinburgh: Scottish Sports Council, 1984).
- Larsen, Mogens Trolle, *Ancient Kanesh: A Merchant Colony in Bronze Age Anatolia* (Cambridge: Cambridge University Press, 2015).
- Lassen, Agnete Wisti, "The Trade in Wool in Old Assyrian Anatolia," *Jaarbericht Ex Oriente Lux*, 42 (2010), 159–179.
- , "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)*, L. Atici, F. Kulakoğlu, G. Barjamovic and A. Fairbairn, eds. (Bristol, CT: The American Schools of Oriental Research, 2014), 107–121.

- Manning, S.W., Gojko Barjamovic, and B. Lorenzen, "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*, 13 (2017), 70–81.
- Massa, Michele, *Networks before Empires: Cultural Transfers in West and Central Anatolia during the Early Bronze Age*. PhD diss., University College London, 2016.
- Matthews, R., and C. Glatz, "The Historical Geography of North-Central Anatolia in the Hittite Period: Texts and Archaeology in Concert," *Anatolian Studies*, 59 (2009), 51–72.
- Michaels, Guy, and Ferdinand Rauch, "Resetting the Urban Network: 117–2012," *Economic Journal*, 128 (2018), 378–412.
- Omura, M., "2015 Yassihöyük kazıları," *Kazı sonuçları toplantısı*, 38 (2016), 343–352.
- Palmisano, Alessio, "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*, L. Feliu, J. Llop and A. M. Albà, eds. (Winona Lake, IN: Eisenbrauns, 2013), 767–783.
- , "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)*, F. Kulakoğlu and G. Barjamovic, eds. (Turnhout: Brepols, 2017), 29–48.
- Palmisano, Alessio, and Mark Altaweel, "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 (2015), 216–236.
- Ramsay, William Mitchell, *The Historical Geography of Asia Minor* (London: John Murray, 1890).
- Riley, Shawn J., Stephen D. DeGloria, and Robert Elliot, "Index that Quantifies Topographic Heterogeneity," *Intermountain Journal of Sciences*, 5 (1999), 23–27.
- Rosen, Kenneth T., and Mitchel Resnick, "The Size Distribution of Cities: An Examination of the Pareto Law and Primacy," *Journal of Urban Economics*, 8 (1980), 165–186.
- Silva, J. M. C. Santos, and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 88 (2006), 641–658.
- Simonovska, Ina, and Michael E. Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92 (2014), 34–50.
- Tobler, Waldo, and S. Wineburg, "A Cappadocian Speculation," *Nature*, 231 (1971), 39–41.
- Weeden, M., and K. Matsumura, "Central West: Archaeology," in *Hittite Landscape and Geography* L. Ullmann and M. Weeden, eds. (Leiden: Brill, 2017), 106–118.