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

Gojko Barjamovic † Thomas Chaney‡ Kerem Coşar§ Ali Hortaçsu¶

June 27, 2017
PRELIMINARY AND INCOMPLETE: DO NOT CIRCULATE.

Abstract

We analyze a large dataset of commercial records produced by Assyrian merchants in the 19th Century BCE. Using the information collected from those records, we estimate a structural gravity model of long-distance trade in the Bronze Age. We find a distance elasticity for ancient trade close to modern estimates. We also use our structural gravity model to locate lost ancient cities. In many instances, our structural estimates confirm the conjecture of historians who followed a radically different method. In some instances, our estimates confirm one conjecture against others. Finally, confronting our structural estimates for ancient city sizes to historical data on political centers, as well as modern data on regional trade and incomes, we document persistent patterns in the distribution of city sizes across four millennia even after controlling for time-invariant geographic attributes such as agricultural suitability, mineral resources and defensive ability. Discussing related quantitative and qualitative evidence from historical geography, we conjecture that the locational advantage brought by natural transport routes dictated by topography is a key factor in explaining the persistence of the city size distribution.

*This research is supported by the University of Chicago Neubauer Collegium for Culture and Society. Thomas Chaney acknowledges ERC grant N°337272–FiNet for financial support. Daniel Ehrlich, Simon Fuchs and Joonhwi Joo provided excellent research assistance. We are grateful to Adam Grant Anderson, Thomas Hertel, Alessio Palmisano for valuable discussions and sharing their data, and to participants at Neubauer Collegium workshops, CEPR ERWIT 2016 conference and various presentations for comments and suggestions.

†Harvard University, barjamovic@fas.harvard.edu.
‡Sciences Po and CEPR, thomas.chaney@gmail.com.
§University of Virginia and CEPR, keremcosar@gmail.com.
¶University of Chicago and NBER, hortacsu@gmail.com.

This paper analyzes a vast collection of commercial records from the earliest documented long-distance trade in world history: The Old Assyrian trade network connecting northern Iraq, northern Syria and central Turkey during the Middle Bronze Age period (c. 1945-1730 BCE). The clay tablets on which the merchants inscribed their shipment consignments, expenses, and contracts —excavated, translated and published by archaeologists and historians for more than a century— paint a rich picture of an exchange economy.

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, a few hundred Assyrian merchants settled in Kanesh (Kaneš) on a permanent or temporary basis. They maintained smaller expatriate trading settlements in some 40 urban centers on the central Anatolian Plateau and Northern Syria. Kaneš was the regional hub of the overland commodity trade involving the import of luxury fabrics and tin from Aššur to copper rich Anatolia in exchange for silver and gold. Assyrian merchants were also involved in a voluminous trade of copper and wool within Anatolia itself.

Our first contribution is to extract systematic information on commercial linkages between cities from ancient texts. To do so, we use two complementary approaches. This first method leverages the fact that the ancient records we study can be transcribed into the Latin alphabet, allowing all texts to be digitized and parsed. We automatically search for joint mentions of city pairs across all records. Since those records all come from merchants' archives, and primarily deal with business matters, we take a joint mention of two cities in a text as evidence of some economic interaction between them. The second method involves systematic reading of texts, which requires an intimate knowledge of the Assyrian dialect of the ancient Akkadian language that the records are written in. Taking individual source context into account, this analysis relies exclusively upon a smaller set of records that explicitly refer to journeys between cities and distinguishes whether the specific journey was undertaken for the purpose of moving cargo, return journeys, or journeys undertaken for other reasons (legal, private, etc.).

Our second contribution is to estimate a structural gravity model of ancient trade. We build a simple Ricardian model of trade in precious commodities. Further imposing that bilateral trade frictions can be summarized by a power function of geographic distance, our model makes predictions on the number of transactions between city pairs, as in our data. Our model can be estimated solely on bilateral trade flows data and on the geographic location of at least some cities. Our structural estimate for the distance elasticity of trade, -1.47, is surprisingly close to modern estimates.

Our third contribution is to use our structural gravity model to estimate the geographic location of lost cities. While some of the cities in which the Assyrian merchants settled have been located

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

Our fourth contribution is to test for the persistence of economic forces over a long horizon. Aside from allowing us to recover the location of lost cities, our gravity models yields a structural estimate for the fundamental economic size of ancient cities, when no reliable data on production and consumption, or even population size or density in the 19th Century BC survives. We infer city sizes instead from the propensity of those cities to trade with others in a structural equilibrium trade model. Projecting these ancient city sizes on a set of time-invariant, observable local amenities and resources, such as crop yields, elevation, ruggedness, proximity to waterways or natural resources, we do not find any robust relationship. Estimated ancient city sizes are, however, strongly correlated with the political importance of those cities in the subsequent four millennia, as well as the economic size of those cities in the current era (based on trade between Turkish cities in 2014), even after controlling for these geographic attributes. At face value, these findings support the hypothesis that the spatial distribution of economic activity may be driven by random factors or early advantages that persists over very long periods through path-dependence and lock-in effects, while local fixed factors may play a secondary role. In light of quantitative and qualitative evidence from historical geography, however, we argue that a factor usually overlooked by economists, natural transportation networks shaped by the topography of the wider region, is a critical factor in explaining the hierarchy of ancient cities and their modern counterparts.

Related literature. Our paper contributes to several literatures. First, we provide the oldest estimate for the distance elasticity of trade, dating back to 19th century BCE. This precedes the available estimates going back to the mid-19th century CE by about four millennia (Disdier and Head, 2008). Second, we invert the structural gravity framework in order to locate lost cities, complementing qualitative approaches in history and archeology with a quantitative method rooted in economic theory. We improve on an earlier contribution by Tobler and Wineburg (1971) in that we use a much larger dataset on bilateral economic interactions between cities, and flexibly estimate the distance elasticity jointly with the coordinates of lost cities, while they imposed a quadratic distance elasticity. Finally, we provide novel evidence on the determinants of the city size distribution. An important line of theoretical and empirical inquiry in economic geography involves attempts at explaining the discrepancies in the economic and demographic size of cities observed at any point in time. First-nature forces, i.e., locational fundamentals as dictated by geography, is potentially an important factor (Davis and Weinstein, 2002). Second-nature forces, i.e., agglomeration of economic activity for non-geographic reasons, may magnify these discrepancies or generate size differentials even across seemingly homogenous locations (Krugman, 1991). Finally, path-dependence through lock-in effects could lead to the persistence of past factors—related to the fundamentals that may have been important once (Bleakley and Lin, 2012; Michaels and Rauch, 2016), or to random events such as the quality of governance at some point in history. Our results and historical setting suggest that centrality in transport routes dictated by topography may be an important geographic factor in explaining persistence of cities' long-run economic fortunes.

The remainder of the paper is organized as follows. Section 1 describes our data. Section 2 derives our model and our estimation strategy. Section 3 presents our structural estimates for the distance elasticity of trade, and the location of lost cities. Section 4 assesses the role of local geographic factors in explaining city sizes, and tests for the long-run persistence of the distribution of city sizes.

1 Ancient Trade Data

Our data comes from a collection of around 12,000 texts that constitute the conserved 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 Akkadian language in cuneiform script by ancient Assyrian merchants, their families and business partners. The texts date back to a period between 1945 and 1730 BCE, with around 90%

of the sample belonging to just one generation of traders, c. 1910 - 1880 BCE.

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

> *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-Assur spent 5 minas of copper for their food. We paid 5 2/3 minas of copper for the wagon.*
>
> 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.*
>
> POAT 28 (lines 3-7)

While the actual cuneiform tablets are scattered all around the world in collections and museums, many of the texts have been transliterated into Latin alphabet, published in various volumes, and recently digitized by historians. In this draft, we use qualitative and quantitative information about cities and merchants mentioned in a sub-sample of 9,000 digitized texts available to us and an additional 3,000 non-digitized texts.[1]

The version of the data we use, tabulated by Barjamovic (2011), mentions 51 unique cities. We focus our attention on 29 of those cities, as the others are either not mentioned jointly with other cities, so that we cannot construct any measure of trade flows for them, or too little is known about them from our data. Of the 29 cities of interest, 17 are known cities, and 12 are lost cities. 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 each other. Lost cities on

[1]The Old Assyrian Text Project website (http://oatp.net/) gives public access to the data. We are grateful to Thomas Hertel, Ed Stratford and all the members of the Old Assyrian Text Project for providing us with the underlying data files.

the other hand are those identified in the corpus of texts, but their location remains uncertain, with no definitive answer from archaeological evidence. From the analysis of textual evidence and the topography of the region, historians have developed competing theories about potential locations for some of those lost cities. We propose to use data on bilateral commercial interactions between known and lost cities and a structural gravity model to inform the search for those lost cities.

We construct three measures of bilateral commercial interactions between cities. The texts also contain qualitative information about prices, financial contracts and resolution of legal disputes, which we do not use but hope to analyze in future work.

The first measure is a count of all mentions of actual cargo shipments from i to j in tablets,

$$N_{ij}^{cargo} \equiv \#\text{ of mentions of cargo traveling from } i \text{ to } j.$$

To construct this measure, we systematically read and translated all tablets. All mentions of cargo shipments in the text were identified. A typical letter will describe one or several itineraries of cargo shipments. The following is an excerpt from a memorandum on travel expenses describing several cargo trips. City names are underlined:

> *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.*
>
> Kt 91/k 424 (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 the type $A \to B \to C$, we count three trips, $A \to B$, $A \to C$ and $B \to C$. As we only very rarely have a description of the content of the caravans, we are unable to identify the intensive margin of shipments, i.e., the value of the wares being transported. Instead, we measure the extensive margin, simply counting the number of shipments. In total, from reading through about 12,000 texts, we extract 253 unique tablets with explicit shipping details, from which we identify 318 shipments.

The total number of shipments we can identify is unfortunately too small, and bilateral flows between cities contains too many zeros to identify our model. As a consequence, we add to these shipments additional information about merchants' travels as our second measure. We count all mentions of actual travels of individuals from i to j in tablets,

$$N_{ij}^{travel} \equiv \#\text{ of mentions of persons traveling from } i \text{ to } j.$$

This includes not only caravans transporting cargo, as in our first measure N^{cargo}, but also travels of individuals which may not be directly involved in shipping goods. So by construction, $N^{cargo} \subset N^{travel}$. The following letter sent to the Assyrian port authorities at *Kaneš* from its emissaries at the Assyrian port in *Wahšušana* describes how missives sent from *Wahšušana* to *Purušhaddum* will travel by two different routes, presumably during a conflict, so as to ensure safe arrival:

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

From this letter, we identify 9 trips: from *Ulama* to *Wahšušana* and from *Šalatuwar* to *Wahšušana* (the letters received by the emissaries); from *Wahšušana* to *Ulama*, from *Ulama* to *Purušhaddum* and from *Wahšušana Purušhaddum* (the first messengers); from *Wahšušana* to *Šalatuwar*, from *Šalatuwar* to *Purušhaddum* and from *Wahšušana* to *Purušhaddum* (the second messengers); from *Wahšušana* to *Kanesh* (the message expected to be forwarded in the near future). We do not count the first mention of a trip from *Wahšušana* to *Kaneš* (the actual and forwarded letters sent by the emissaries): Most of our data comes from letters found in merchants' archives in the city of *Kaneš*. By construction, all those letters involve a trip to *Kaneš* as the letter is being sent there. Counting those trips would systematically bias our measures in favor of finding large inflows into *Kanesh.* As with N_{ij}^{cargo}, for itineraries of the type $A \to B \to C$, we count three trips involving all downstream pairs.

While the information in N^{travel} is not necessarily about actual shipment of goods, it informs us about a broader set of economic interactions. As our data comes from letters between merchants, if N_{ij}^{travel} is large, we infer that flows of trade, services and people from i to j are large. Mentions of persons' travels adds 87 trips to the 318 shipments of X^{cargo}, for a total of 299 itineraries and 405 trips (itineraries that involve more than two cities generate multiple trips).

The third measure is a simple automated count of all joint attestations of cities i and j in our 9,000 digitized tablets,

$$N_{ij}^{joint} \equiv \#\text{ of joint attestations of cities } i \text{ and } j.$$

If cities A, B and C are mentioned jointly in a given tablet, no matter how many times each, we count one joint attestation between A and B, one between A and C, and one between B and C. We use all possible spellings of city names.[2] We extract a total of 944 joint attestations of city pairs. These joint attestations may be about a specific shipment from one city to the another, about a merchant traveling between cities, about a financial contract involving parties in different cities, or simply about a letter mentioning those cities for different reasons. While this measure may pick up economically irrelevant linkages between cities, the benefit from using this measure is two-fold. First, this measure is automated, allowing us to quickly extract large amounts of information at little cost, without the need to painstakingly read through thousands of ancient texts. Measurement error is traded off by the mere quantity of information and the efficacy of extracting it. Second, this measure complements our other measures, adding information about potentially meaningful economic interactions that a direct human read of the texts would not systematically pick up. For instance, a financial contract which involves counterparts in cities i and j would not count into our other measures N^{cargo} and N^{travel}, although it informs us about a direct economic interaction between agents in i and j.

While the data for N^{cargo} and N^{travel} are collected using a qualitative method —reading through the texts and understanding their meaning— very different from the method used for N^{joint} —automatically searching for co-occurrences— those measures are correlated. For instance, we find a 62% correlation between N^{travel} and N^{joint}.

Table 1 provides summary statistics for the ancient data. Note that by only counting meaningful economic ties, the human-read data eliminates one known and one lost city from the analysis. The mean number of travels across all city pairs is 0.56. As in the modern international trade data, many city pairs do not trade: of all the 702 potential export-import relationships (directional ij and ji pairs of 27 cities), only 122 have a positive flow. Average N^{travel} for these trading pairs is 3.27. As expected, the machine-read non-directional data N^{joint} contains a higher number of bilateral links since it does not distinguish coincidental co-occurrences of cities in texts from business-related travels between them.

Figure 1 plots the map of cities, including a preview of estimated locations from both data sets. The city of Kaneš, marked K, is geographically central to the system of cities under study. As discussed above, it was also the operational center of Assyrian merchants in central Anatolia. Trade

[2]We exclude *Aššur*, the capital city of the Assyrians, from our automated search for two reasons: First, *Aššur*, is also the name of the main Assyrian deity, and is used as the word for a calendar month; second, the city of *Aššur* is often referred to as simply "the city." Our automated search is not able to use a letter's context to distinguish between *Aššur* as a goddess, as a month, or as a city; or the word for "city" as being *Aššur* or another city.

flows, however, do not just display a hub-spoke structure around Kaneš, as seen by the rich pattern of bilateral ties between cities in figure B.1.

2 Model and Estimation

We build a simple model of trade in which merchants arbitrage price differentials between cities. While stylized, this model captures key features of trade in the Bronze Age. For instance, the model can accommodate a commodity produced outside of our network of trading cities, such as tin sourced from Central Asia, traded locally among Assyrian colonies, and exported to distant places such as Egypt. We also make every effort to characterize equilibrium quantities in our model with direct counterparts in our dataset, such as the count of transactions instead of their value.

Model. We follow Eaton and Kortum (2002) closely. There are $K + L$ cities, K of them known, and L of them lost. Tradable commodities (tin, copper, wool...) are indexed by ω. Merchants arbitrage price differentials between cities, subject to bilateral transaction costs. For simplicity, we assume iceberg ad valorem transaction costs, such that delivering one unit of a good from city i to city j requires shipping $\tau_{ij} \geq 1$ units of the good, while the remaining fraction $1 - 1/\tau_{ij}$ is lost in transit. If a merchant observes a cost c_i in a city i, and a cost c_j in city j, such that

$$\tau_{ij} c_i < \tau_{jj} c_j,$$

she can exploit an arbitrage opportunity: Buy τ_{ij} units of the good at a cheap cost $\tau_{ij} c_i$ in i, ship them to j, and sell at a high cost $\tau_{jj} c_j$ for a profit.[3] We assume the cost of producing one unit of any commodity ω in city i, in any period, follows a Weibull distribution,

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

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

[3] As the merchants we consider are mobile, constantly traveling between cities themselves, we do not consider the problem of repatriating the proceeds from this sale explicitly. In particular, we implicitly assume repatriation of profits is costless. If repatriating profits entails a cost, the τ_{ij} term would contain both the cost of shipping goods *and* of repatriating profits. We also explicitly assume a transaction cost even for within city transactions, $\tau_{jj} \geq 1$, to capture local distribution costs.

controls the efficiency of sourcing goods from i.[4]

Similar to Eaton and Kortum (2002), the equilibrium price for commodity ω in city j—with merchants arbitraging away cost differences between cities—is the lowest cost among all possible sources, $p_j(\omega) = \min_k \{\tau_{kj} c_k(\omega)\}$. The distribution of prices in city j is then Weibull with shape parameter θ and scale parameter $\sum_k T_k w_k^{-\theta} \tau_{kj}^{-\theta}$. Assuming trade balance at the city level, total spending, including spending on locally produced goods, equals the amount paid to local factors,

$$X_i = \sum_k X_{ki} = w_i L_i, \tag{2}$$

where L_i is the size of city i's population. Since the productivity parameter T_i and population size L_i play isomorphic roles, and since we have no direct measure of the population of ancient cities, we normalize $L_i = 1, \forall i$, so that T_i is the sole determinant of city i's fundamental size. With N i.i.d. draws for the costs $\{c_i(\omega)\}_{i=1\cdots K+L}$, the expected number of shipments going from i to j is

$$\mathbb{E}[N_{ij}] = N \frac{T_i w_i^{-\theta} \tau_{ij}^{-\theta} X_j}{\sum_k T_k w_k^{-\theta} \tau_{kj}^{-\theta}}. \tag{3}$$

It will prove useful to manipulate the model as in Anderson and van Wincoop (2003) to obtain the following expression for the expected number of shipments from i to j,

$$\mathbb{E}[N_{ij}] = N \frac{X_i X_j}{X_{total}} \left(\frac{\tau_{ij}}{\Pi_i P_j} \right)^{-\theta}, \tag{4}$$

where $\Pi_i^{-\theta} = \sum_k \left(\frac{\tau_{ik}}{P_k} \right)^{-\theta} \frac{X_k}{X_{total}}$ is a measure of outward resistance, $P_j^{-\theta} = \sum_k \left(\frac{\tau_{kj}}{\Pi_k} \right)^{-\theta} \frac{X_k}{X_{total}}$ a measure of inward resistance, and $X_{total} = \sum_k X_k$. In particular, we will rely on the result that if trade frictions are symmetric, $\tau_{ij} = \tau_{ji}, \forall i \neq j$, then the inward and outward resistance terms are equal, $\Pi_i = P_i$, and trade is symmetric, $\mathbb{E}[N_{ij}] = \mathbb{E}[N_{ji}]$.

[4] For instance, following Kortum (1997), if the distribution of efficiency for producing commodity ω locally is Pareto with shape parameter θ, and if the probability that any local worker/entrepreneur tries her luck operating a local facility producing ω is exponentially distributed, then the cost of producing good ω in city i (the inverse of the efficiency) is distributed according to the Weibull distribution (1), where T_i is exactly proportional to the number of local households and their efficiency. With Weibull distributed costs, larger and/or more efficient cities, in the sense of cities with a higher T_i, will also be cities with lower costs on average:

$$\mathbb{E}_i[c] \propto w_i T_i^{-1/\theta}.$$

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

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

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

We parametrize the symmetric trade cost function as

$$\tau_{ij}^{-\theta} = \delta Distance_{ij}^{-\zeta}, \forall i \neq j. \tag{5}$$

The parameter ζ is the distance elasticity of trade, and δ a simple scaling factor which controls the units of measurement of distances. For cities i and j with latitude-longitude (φ_i, λ_i) and (φ_j, λ_j), $Distance_{ij} = H(\varphi_i, \lambda_i; \varphi_j, \lambda_j)$, where the function H maps geo-coordinates into geographic distances, measured in kms.[5]

We proceed to estimate the following vector of structural parameters

$$\left(\zeta, (\varphi_{K+1}, \lambda_{K+1}) \cdots (\varphi_{K+L}, \lambda_{K+L}), T_1^{1/\theta} \cdots T_{K+L}^{1/\theta}\right).$$

ζ is the distance elasticity of trade. (φ_l, λ_l) are the geo-coordinates of lost city l. $T_i^{1/\theta}$ is our measure of the fundamental size of city i, the counterfactual aggregate welfare of city i if it were to move to full autarky (Eaton and Kortum, 2002). We use the structural model (3) and our trade cost assumption (5) to form an expression for trade shares,

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

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

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

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

$$(\zeta; \cdots (\varphi_l, \lambda_l) \cdots; \cdots \alpha_i \cdots) = \arg\min \sum_j \sum_{i\neq j} \left(s_{ij}^{data} - \frac{\alpha_i Distance_{ij}^{-\zeta}}{\sum_{k\neq j} \alpha_k Distance_{kj}^{-\zeta}}\right)^2. \tag{8}$$

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

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

where 37.9 degrees North is the median latitude among known Assyrian cities. For locations in the Near East, the difference between this Euclidean formula and the more precise Haversine formula is negligible. This approximation considerably speeds up the estimation.

As the absolute level of the α_i's cannot be identified, we arbitrarily normalize $\alpha_{Kaneš} \equiv 1$.

We use our structural model, in particular equations (3) and (4), and parameter estimates from (8), to recover an estimate of fundamental city sizes, $T_i^{1/\theta}$,

$$T_i^{1/\theta} \propto \hat{\alpha}_i^{1+1/\theta} \sum_k Distance_{ki}^{-\hat{\zeta}} \hat{\alpha}_k, \tag{9}$$

where we use $\theta = 4$ from Simonovska and Waugh (2014).[6] We arbitrarily normalize $T_{Kaneš}^{1/\theta} \equiv 100$, so city sizes are all relative to that of Kaneš.

The minimization problem (8) is highly non-linear, primarily because we have to estimate the geographic coordinates of lost cities. As a result, our algorithm for solving (8) does not converge, for any initial value, when we use data on cargo shipments, N^{cargo}, as this dataset is too sparse. It does converge for our two other measures of bilateral linkages, the number of (cargo and non-cargo) travels between cities, N^{travel}, and the number of joint attestations of city names, N^{joint}. In the case of travels between cities, N^{travel}, we use 702 moments (all directional trade flows between 16 known cities and 11 lost cities) to estimate 49 parameters (the distance elasticity, 22 coordinates and 26 city sizes after the normalization of $T_{Kaneš}^{1/\theta} \equiv 100$). In the case of joint attestations of city names, N^{joint}, we use 406 moments (all undirected trade flows between 17 known cities and 12 lost cities) to estimate 53 parameters.

Our non-linear estimation is closely related to Silva and Tenreyro (2006) and to Eaton et al. (2012). In particular, we use information contained in trade zeros to inform our estimation. There are, however, three key differences imposed on us by the data. The first obvious difference is that, unlike with modern trade data, we do not know the actual geographic location of some cities. We use instead our model to estimate those locations. The second difference is that we do not observe internal trade, i.e., within-city transactions. We define trade shares in equations (6) and (7) excluding internal trades. The third difference is that in cases where we are unable to identify the direction of trade, we use our structural model to generate predictions for undirected trade flows.

[6] To get $T_i^{1/\theta}$ we need to know the trade elasticity parameter θ. In the absence of any data on differences in commodity prices, our data does not allow us to directly estimate θ. We therefore choose $\theta = 4$ from the literature . Since the parameter θ only affects the absolute level of our estimates of city sizes, but not relative city sizes (in logs), this choice is of little consequences.

We also need to know internal trade frictions. Since we do not observe internal trades, we cannot estimate city transactions costs. We instead normalize internal distances, $Distance_{ii} = 30km$, capturing the economic hinterland of a city within the reach of a day's travel.

To derive $T_i^{1/\theta}$ in (9), we use the structural equation (6) to get $\hat{\alpha}_i \propto T_i w_i^{-\theta}$ so $T_i^{1/\theta} \propto \hat{\alpha}_i^{1/\theta} w_i$. From our normalization $L_i = 1$ and market clearing, $w_i = X_i$. From (3), (4), and (6), $\hat{\alpha}_i \propto T_i w_i^{-\theta} = X_i / \Pi_i^{-\theta}$ so $X_i \propto \hat{\alpha}_i \Pi_i^{-\theta}$. From the definition of $\Pi_i^{-\theta}$ and symmetry, $\Pi_i^{-\theta} \propto \sum_k \tau_{ik}^{-\theta} X_k / P_k^{-\theta} = \sum_k \tau_{ik}^{-\theta} X_k / \Pi_k^{-\theta} = \sum_k \tau_{ik}^{-\theta} \hat{\alpha}_k$. Combining $\tau_{ik}^{-\theta} \propto Distance_{ik}^{-\zeta}$ and the above, we get the proposed formula. City sizes are only identified up to a multiplicative constant, hence the $\propto$ sign.

Standard errors. Standard errors are calculated by bootstrapping and account for sampling error. For each bootstrap $b = 1 \cdots B$, we generate two datasets, one for directional travels $\left(N^{travel}\right)$ and one for non-directional joint attestations $\left(N^{joint}\right)$ by sampling with replacement as many tablets as there are in our data and solve (8) and (9). For each parameter $\beta \in \left\{\zeta, T_1^{1/\theta} \cdots T_{K+L}^{1/\theta}\right\}$, and the corresponding bootstrapped estimates $\left\{\beta^b\right\}_{b=1\cdots B}$, the standard error is

$$s.e.\left(\beta\right) = \sqrt{\frac{1}{B}\sum_{b=1}^{B}\left(\beta - \beta^b\right)^2}.$$

The confidence area for city l corresponds to a contour plot for the 2-dimensional distribution of bootstrapped locations in the set $\left\{\left(\varphi_l^b, \lambda_l^b\right)\right\}_{b=1\cdots B}$. We also compute a measure of the precision of our location estimates akin to a standard error,

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

where $Distance_{l,l^b}$ is the distance between the estimated location for l in the true data and that in bootstrapped dataset b. This measure of precision is expressed in kms.

We present the results from our estimation in the next two sections. In section 3, we focus on the distance elasticity of trade and on the location of lost cities. In section 4, we focus on our estimates of city sizes.

3 Distance Elasticity and the Lost Cities of the Bronze Age

Table 2 presents our GMM estimates of the distance elasticity of trade. The distance elasticity estimated from data dating back four millennia are eerily close to modern estimates, typically around -1. When using joint attestations as our measure of bilateral commercial links between cities $\left(N^{joint}\right)$, we find a precisely estimated distance elasticity of -1.47 with a standard error of 0.23, near modern estimates. When using instead the total number of persons' travels $\left(N^{travel}\right)$, our estimate for the distance elasticity is higher in absolute term and less precise, -2.40 with a standard error of 1.26. It is however close to modern estimates of the trade elasticity for shipments transported by road —see Cosar and Demir (2016) for an estimate around -2 based on overland transit of exports from Turkish cities.

Our estimation also allows to recover the likely location of ancient lost cities. Not only can we estimate the geo-coordinates for each lost city, but our bootstrap method for computing standard errors also gives us entire confidence areas around our point estimates.

Figure 1 shows the location of all cities, both known and lost. The lost cities are denoted by a "+" sign, the lost cities by a "o" sign, and central city of *Kanesh* with a "K" sign. Panel A corresponds to estimates using joint attestations $\left(N^{joint}\right)$ as our measure of bilateral economic exchanges between cities; panel B corresponds to estimates using data on persons' travels $\left(N^{travel}\right)$. Our estimates when using joint attestations $\left(N^{joint}\right)$ are more concentrated in the center of the network of trading cities than when using explicit mentions of persons' travels $\left(N^{travel}\right)$. This apparent discrepancy between point estimates is partly misleading. Once measurement error is taken into account (we use a bootstrap method to correct for sampling error), confidence areas most often overlap.

We show our point estimates and confidence areas for four important cities in figures 2–5 (our estimates for the other cities are in appendix figures B.2–B.8). Each figure depicts a map of the region containing the Assyrian colonies in Anatolia (central Turkey). A green dot depicts the estimated location of one lost city from our structural estimation. A series of contours around this dot show the estimated confidence region for that lost city. We add to this map two other locations. The first, depicted with a "B" corresponds to site suggested by historian Gojko Barjamovic (Barjamovic, 2011); the second, depicted with an "F" corresponds to the site suggested by historian Massimo Forlanini (Forlanini, 2008). For each city, we show two maps, each corresponding to a different dataset: Panel A corresponds to estimates using joint attestations $\left(N^{joint}\right)$ as our measure of bilateral economic exchanges between cities; panel B corresponds to estimates using data on persons' travels $\left(N^{travel}\right)$.

Those maps visually deliver several messages. First, the size and shape of the confidence regions around our location estimates give a visual sense of the precision of those estimated locations. For most cities, our estimates are tight, in the sense that the confidence area is at most 100km wide (to be compared to distances within our network of trading cities of more than a thousand kms). Second, we can compare our estimated results from two datasets, collected with intrinsically different methods: joint attestations $\left(N^{joint}\right)$ come from a purely automated string search, while persons' travels $\left(N^{travel}\right)$ come from a careful reading and understanding of ancient texts. For all precisely estimated city locations, the estimates from both datasets are surprisingly close. Third, we can compare our estimates, obtained by a purely quantitative method—a structural gravity estimation, to those obtained by historians from a purely qualitative method—reading through texts and isolating contextual information about likely sites. In several cases where historians' conjectures do not drastically diverge, our estimates are very close to these conjectures, giving credence to both our quantitative method and their qualitative approach. In some cases, our estimates are close to one historian's conjecture but not the other's, giving a strong endorsement to the former.

Figure 2 shows the estimated location of *Durhumit*. It is an important city for several reasons. First, it is an important market place, sometimes acting as a substitute for the capital city of *Kaneš*, with smugglers bringing tin and textiles from *Aššur* along the illegal "Narrow Track", in exchange for locally sourced copper, avoiding to pay the entry toll in *Kaneš*. Some goods brought to *Durhumit* were further shipped to *Purušhaddum*. Second, this is one of two lost Assyrian cities, along with *Purušhaddum*, for which there is the largest disagreement between historians Barjamovic and Forlanini. Our gravity-based estimate for the location of *Durhumit* strongly favors the conjecture of Barjamovic, against that of Forlanini. Our estimated location using either directional or non directional data is within 50kms of the conjecture of Barjamovic. For both datasets, our confidence areas are fairly tight, and contain Barjamovic's conjecture with an 80% probability when using directional data $\left(N^{travel}\right)$.

Figure 3 shows the estimated location of *Wašhaniya*. Our estimate is closer to the conjecture of Barjamovic. When using non-directional data $\left(N^{joint}\right)$, Barjamovic's conjecture lies within our confidence region, while Forlanini's conjecture does not. When using directional data $\left(N^{travel}\right)$, both conjectures lie within our confidence region, but our point estimate is extremely close to Barjamovic (within 10kms).

Figure 4 shows the estimated location of *Šinahuttum*. This is a case where Forlanini's and Barjamovic's conjectures are very close, and where our structural gravity-based estimation is almost identical to those historians' conjecture. Our point estimates for both datasets are within 15kms of the historians' conjecture, with a tight confidence region which contains those conjectures. We take this finding as evidence that the actual site for the city of *Šinahuttum* is to be found in either of the historians' suggested sites with a very high degree of confidence.

Finally, figure 5 shows the estimated location for the city of *Purušhaddum*. As for *Durhumit*, this is a case where Forlanini and Barjamovic strongly disagree on the likely location of the city. However, unlike for *Durhumit*, our location is extremely imprecisely estimated, with extremely wide confidence areas for both our datasets. For both datasets, Forlanini's conjecture lies within our confidence region. But given that those confidence regions are 500 kms wide in the north-south direction, our estimate does not allow us to accept Forlanini's conjecture and reject Barjamovic's conjecture with confidence. This is a case where our estimates are too imprecise to decide among competing historians' suggestions.

Table 3 lists the distances between our point estimates for the location of lost cities and the conjectured locations from historians Forlanini and Barjamovic. Our directional estimates allow us to say with a high degree of confidence that the likely locations of *Šinahuttum* (figure 4), *Mamma*

(figure B.2), *Šuppiluliya* (figure B.4) and *Hahhum* (figure B.5) are at or very near the historians suggested sites (our estimates are within 50kms). For *Durhumit* (figure 2) and *Wašhaniya* (figure 3), we can with a high degree of confidence confirm Barjamovic's conjecture against Forlanini's. For *Purušhaddum* (figure 5), *Kuburnat* (figure B.6), *Tuhpiya* (figure B.7) and *Zalpa* (figure B.8), our estimates are very far (more than 100kms) from the historians conjectures, and we are unable to make any recommendation as to the likely site of those lost cities.

4 The Distribution of City Sizes: from 2000 BCE to 2014 AD

Table 4 reports the estimates of $T_i^{1/\theta}$ using both the directional and non-directional datasets. The correlation between the natural logarithms of the two results is 0.57. We use these estimates to explore the drivers of the city size distribution. We present three main findings. First, we find no evidence that city sizes are driven by locational fundamentals (section 4.1). Second, we show that large cities are more likely to become political and administrative centers (section 4.2). Third, the distribution of city sizes is surprisingly stable over four millennia (section 4.3).

4.1 Productivity and Locational Fundamentals

To test the hypothesis that city sizes are determined by locational fundamentals, we project our estimated $T_i^{1/\theta}$'s onto various geographic observables such as elevation, ruggedness, crop yield, as well as distance to rivers and to deposits of minerals that were economically valuable in the ancient economy. Some of these geographic characteristics, such as elevation, ruggedness and river access, are truly time invariant. Some others, such as crop yields and availability of minerals, may change over time due to variation in local climate and soil conditions, improvements in irrigation, arrival of new crops that affect regions differentially, and advances in mineral sciences enabling the discovery of deposits previously unknown to the locals. Still, present-day measures of these variables are likely to be highly correlated with their counterparts 4000 years ago.

Table 5 reports the results. Across all columns, we regress the natural logarithm of $T_i^{1/\theta}$ estimated from the directional data—shown at columns 4 and 6 of table 4—on the natural logarithms of elevation, ruggedness, cereal yield and distance to the nearest major river. In columns 2, 4 and 6, we add the (log) distance to the nearest source of economically relevant metals. To check robustness, we report specifications using all cities in columns 1-2, only the cities with estimated locations agreeing with the conjectures of Barjamovic (2011) in columns 3-4, and only the cities the locations of which are precisely known in columns 5-6. All regressions are estimated via weighted

least squares. The weights are the total count of importing and exporting relationships in the data, i.e., the number of itineraries in which a city is mentioned.[7]

The key message of table 5 is the absence of a significant and robust relationship between geographic fundamentals and estimated city sizes. Across all specifications, average cereal yield has a consistent and plausible positive sign, significant at the 10 percent level in one specification (column 4). Elevation and ruggedness flip signs across samples but are insignificant across all columns. City sizes are positively correlated with distance to the nearest river, but the coefficient is significant at the 10 percent level only if we restrict the sample to known cities (columns 5-6). The counter-intuitive result that proximity to a river does not confer an advantage to a city is consistent with the fact that due to the steep rapids caused by topography, rivers in central Anatolia are not suitable for transportation. We also don't find any meaningful association between city sizes and distance to copper, gold or silver deposits. This result, however, should be interpreted with care: among the explanatory variables, location of metal deposits as currently known is plausibly the least informative one for the ancient conditions.[8]

Next, we test whether ancient prominence persists into the future. The fact that we found no robust relationship between ancient sizes and geographic attributes makes this exercise informative about the potential role of time-invariant fundamental advantages versus path-dependence in accounting for long-term economic outcomes of cities and regions.

In what follows, we use all the cities in our sample and weight observations by the total count of cities' appearances in the itinerary data. The results are robust to dropping unknown cities if their locations are estimated to be inconsistent with qualitative evidence, and using known cities only.

4.2 Future Capital Cities

Six cities in our data served as capitals of a state at some point in their history after the era under study—see Appendix A.1 for detailed information including the list of cities. To check whether these cities differ from the rest, we first note that average $\ln(T_i^{1/\theta})$ equals 7.45 and 5.67 for capital and non-capital cities, respectively. Similarly, median size among capital cities is 1.45 log points higher than among non-capital cities.

For a more detailed analysis, we regress a binary variable that takes the value one for capitals, and zero otherwise, on $\ln(T_i^{1/\theta})$ while controlling for the set of geographic attributes introduced above. Table 6 reports the results. The first two columns are results from a linear probability

[7] In constructing elevation, ruggedness and cereal yields, we use the average within an area of 30 km radius—the reach of a day's travel at the time—around the city to account for its economic hinterland.

[8] In fact, the minimum distance is above 30 km for all three metals.

model. The third column reports average marginal effects from a probit regression. Across all columns, ancient size estimates are positively correlated with the future capital status of cities with varying degrees of significance while geographic attributes are estimated to be insignificant.

The result that size in the past is positively correlated with the propensity of a city being a capital in the future could be driven by various channels that are not necessarily mutually exclusive: it could be that a city's prominence persists and increases its probability of being selected as the center of a state that takes the region under its control. Similarly, such cities may simply have a higher chance of surviving into the future. Alternatively, their resource base and prominence could support a military and political expansion, as a result of which they become the center of a state commanding a wider territory. This latter possibility, however, is not supported by the result that the resource-based geographic fundamentals used as controls are insignificant.

4.3 Present-day Size Estimates, GDP and Population

Moving further into the future, we now ask whether ancient sizes are correlated with present-day size estimates of the same locations. To address this question, we find the set of present-day cities in Turkey corresponding to our system of ancient cities. There are 14 such cities since some contain multiple ancient cities within their boundary. Using 2014 trade flows within this system of 14 cities, we estimate $T_i^{1/\theta}|_{modern}$ by the same econometric procedure except that we do not need to estimate locations in this case.[9] Denoting the original estimates of ancient productivities by $T_i^{1/\theta}|_{ancient}$, figure 6 and table 7 present the results. Despite being estimated using trade flows that are 4000 years apart, the correlation is remarkably high and significant. Ancient sizes capture around 60 percent of the variation in modern sizes.

Continuing the same line of inquiry, we replace the dependent variable with the (log) total income and population of present-day cities. Since the Turkish Statistical Institute does not publish city-level GDP data for the period after 2001, we use estimates by the Economic Policy Research Foundation for 2013, the latest year for which data available.[10] Figures 7 and 8, and table 8 report the results. Even after controlling for time invariant geographic attributes, size estimated from ancient data is correlated with present-day income and population.[11]

[9]The only ancient city that is not within the boundaries of present-day Turkey is Qattara, which is near Tal Afar/Iraq. Since 2014 trade data is only for Turkish cities, this ancient city drops from the subsequent analysis.

[10]The data is accessible at the website http://www.tepav.org.tr/tr/haberler/s/4054. Economic Policy Research Foundation uses nighttime luminosity to estimate city-level incomes as in Hodler and Raschky (2014).

[11]Figures 6, 7, and 8 leave out Ankara, which is an outlier with high present-day levels due to its status as Turkey's capital city.

4.4 Discussion

Our exploration yields an intriguing result: a variety of geographic attributes—agricultural potential, access to mineral resources, distance to rivers and terrain-related defensive capability—are not correlated with estimated economic size of ancient cities, which are in turn correlated with future prominence and economic size of the same localities. We now discuss whether this strong persistence in regional outcomes lends support to models of city growth through path-dependence and lock-in effects, as in Bleakley and Lin (2012) for the case of mid-Atlantic and southern U.S. cities that were once portage sites at fall lines, and in Michaels and Rauch (2016) for French cities originating from Roman towns.

A potential caveat is that the geographic attributes we used in our analysis do not adequately capture locational advantages that matter throughout history in a persistent manner. One such factor is the geography of natural transport networks. In a rich account of historical geography of Anatolia, Ramsay (1890) argues that the topography of the region imposes strict restrictions on main passages in the east-west and north-south directions. As a result, cities located on these routes—especially those at the intersection points, the so-called road-knots—enjoyed an advantage of being ideal distribution centers for regional trade.[12] While we do not have a readily available quantitative measure of this geographic attribute, we reckon that it plays an important role in explaining the long-run persistence of economic size across Anatolian cities. The very location of Kanesh and the present-day city of Kayseri—whose center is only 20 km away from Kanesh—is a case in point: it lies at the western side of Taurus crossings connecting the central Anatolia plateau to the upper Mesopotamian plain. Several other ancient and corresponding present-day cities with high productivity estimates, such as Hurama-Elbistan, Zalpa-Gaziantep, Salatuwar-Eskisehir and Samuha-Sivas, are also placed on road-knots (Barjamovic, 2011; French, 1993). The main transportation arteries in Turkey overlap with the known Roman roads, which themselves may have presumably followed the ancient routes from the Bronze Age.[13] Recent GIS analysis by Palmisano (2013) and Palmisano and Altaweel (2015) confirms that ancient routes were indeed very close to be the least-cost pathways, which increases our confidence that the modern-transportation routes are not merely overlaid on ancient networks through a lock-in effect.

[12] 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.

[13] A similar long-run continuity in the Inca road network has been used by Martincus et al. (2017) to instrument for the development of the modern road system in Peru.

Conclusion

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

Three main results emerge. First, the estimated distance elasticity of trade is remarkably close to estimates from present-day datasets. This is consistent with Disdier and Head (2008) who document a puzzling persistence of the distance effect in international trade during the modern-era. In light of the apparent improvements in transportation technologies, recent theories offer explanations for this pattern based on trade frictions that are very persistent themselves, such as the way traders form networks and interact within them (Chaney, 2014). The ancient setting under study is one in which such relationships are of primary importance to conduct long-distance trade under severe enforcement problems. In future work, we hope to utilize the information on individual merchants and their networks available in the ancient data to shed further light on this subject.

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

Finally, we analyze the correlation between the estimated economic size and observable time-invariant geographic attributes of ancient cities, as well as their future economic outcomes. This exercise brings fresh evidence on a seminal question in economic geography: what explains the size of cities? We do not find any significant explanatory power for ancient economic size in agricultural potential, mineral resources and terrain-related defensive capability. We do, however, find that ancient economic size predicts the propensity of these cities to become a regional center in the future, as well as the income and population of corresponding regions in present-day Turkey. We argue that the persistence of cities' fortunes across 4000 years can best be explained by their time-invariant locational advantage in the topography of trade routes, as proposed by Ramsay (1890).

References

Anderson, J. E. and E. van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–92.

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

Bleakley, H. and J. Lin (2012): "Portage and path dependence," *The quarterly journal of economics*, qjs011.

Chaney, T. (2014): "The network structure of international trade," *The American Economic Review*, 104, 3600–3634.

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

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

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

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

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

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

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

French, D. (1993): "Colonia Archelais and Road-Knots," *Aspects of Art and Iconography: Anatolia and Its Neighbors Studies in Honor of Nimet Özgüç*, 201–7.

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

Kortum, S. S. (1997): "Research, patenting, and technological change," *Econometrica*, 65, 1389–1419.

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

Martincus, C. V., J. Carballo, and A. Cusolito (2017): "Roads, exports and employment: Evidence from a developing country," *Journal of Development Economics*, 125, 21–39.

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

Palmisano, A. (2013): "Computational and Spatial Approaches to the Commercial Landscapes and Political Geography of the Old Assyrian Colony Period." in *Time and History in the Ancient Near East. Proceedings of the 56th Rencontre Assyriologique Internationale, Barcelona, July 26-30, 2010.*, Eisenbrauns., 767–783.

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

Ramsay, W. M. (1890): *The historical geography of Asia Minor*, vol. 4, John Murray.

Silva, J. S. and S. Tenreyro (2006): "The log of gravity," *The Review of Economics and statistics*, 88, 641–658.

Simonovska, I. and M. E. Waugh (2014): "The elasticity of trade: Estimates and evidence," *Journal of International Economics*, 92, 34 – 50.

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

Figures

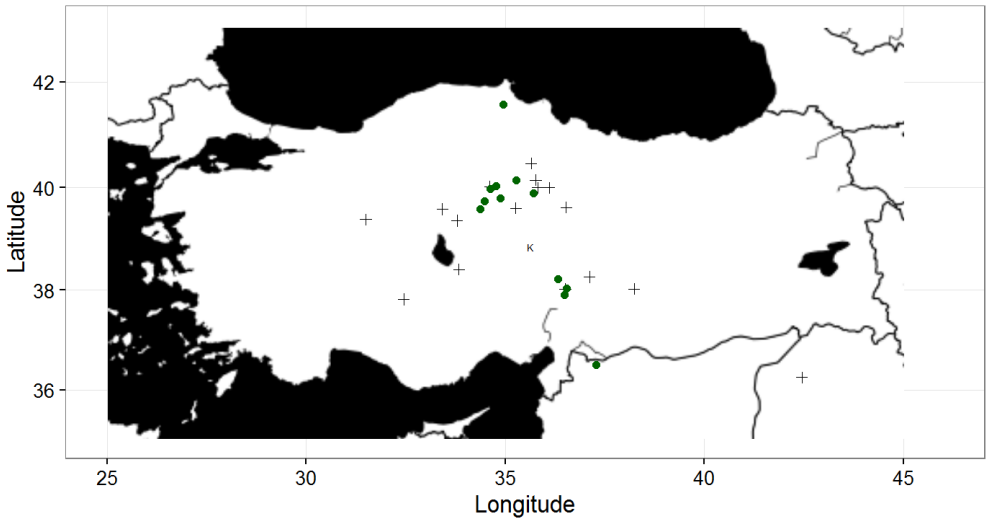

Panel A: Non-directional trade (N_{ij}^{joint})

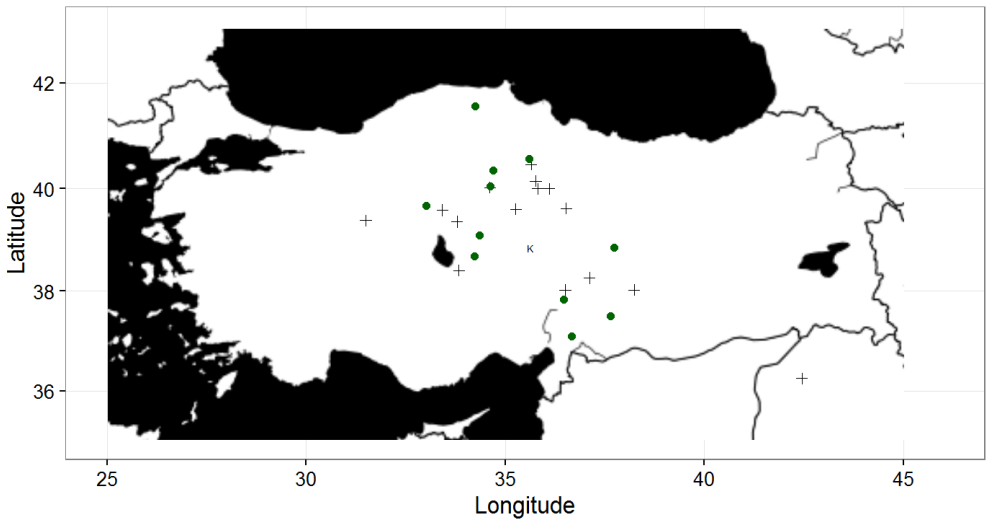

Panel B: Directional trade (N_{ij}^{travel})

Figure 1: Known and lost cities.

Notes: The maps show the of all Assyrian cities. The "+" signs correspond to the location of known cities. The "K" sign denotes *Kanesh*, the capital city of Assyrian colonies. The "o" signs denote the estimated location of lost cities, from the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. *Sources*: Old Assyrian Text Project.

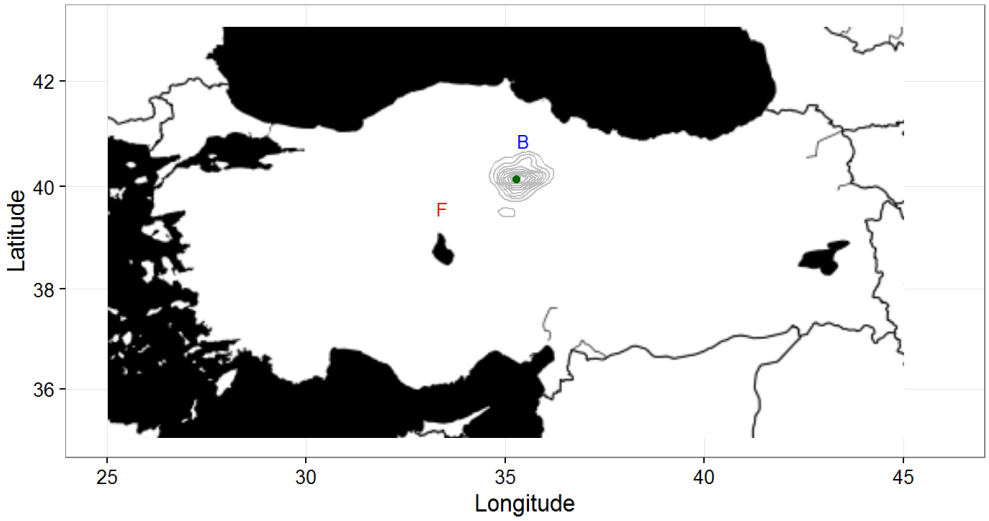

Panel A: Non-directional trade (N_{ij}^{joint})

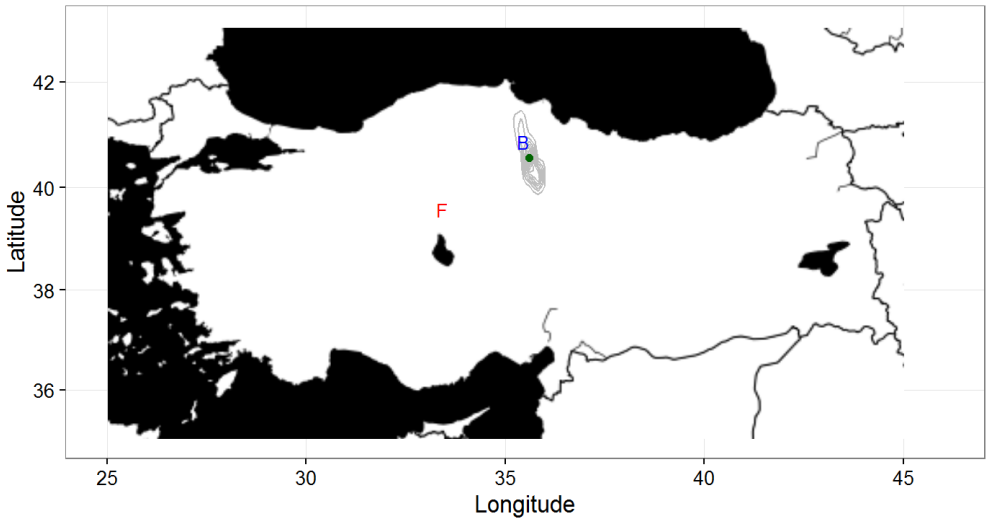

Panel B: Directional trade (N_{ij}^{travel})

Figure 2: Finding lost cities: *Durhumit*.

Notes: The maps show the estimated location for the ancient city of *Durhumit*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

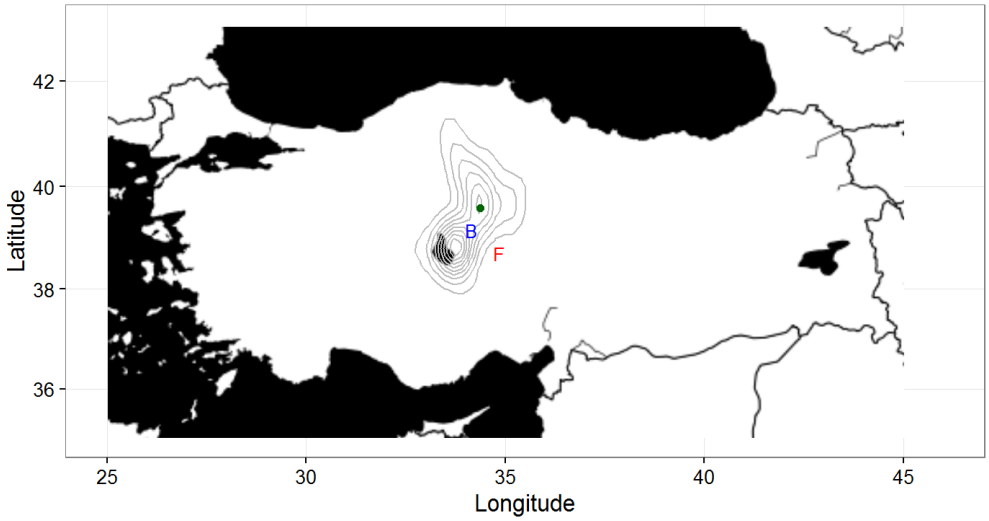

Panel A: Non-directional trade (N_{ij}^{joint})

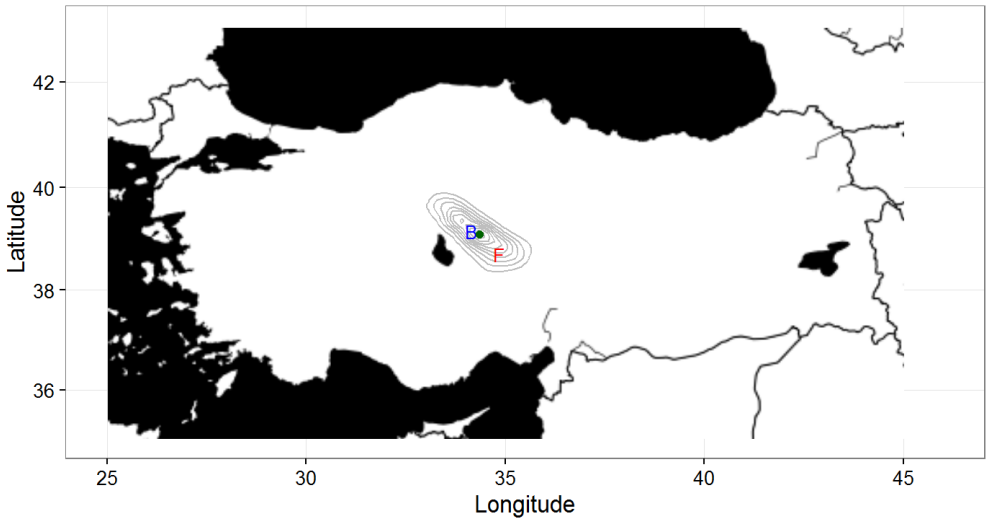

Panel B: Directional trade (N_{ij}^{travel})

Figure 3: Finding lost cities: *Washaniya*.

Notes: The maps show the estimated location for the ancient city of *Washaniya*. IIn both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

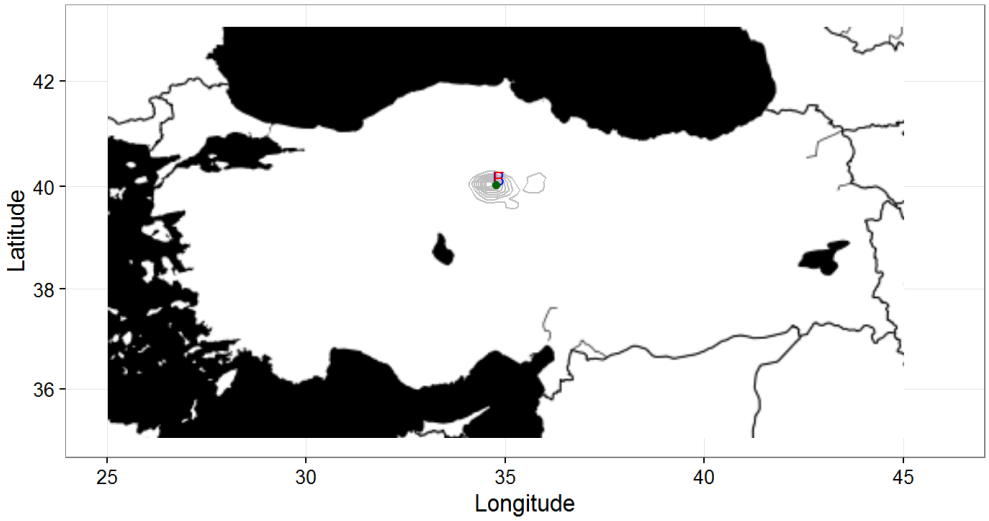

Panel A: Non-directional trade (N_{ij}^{joint})

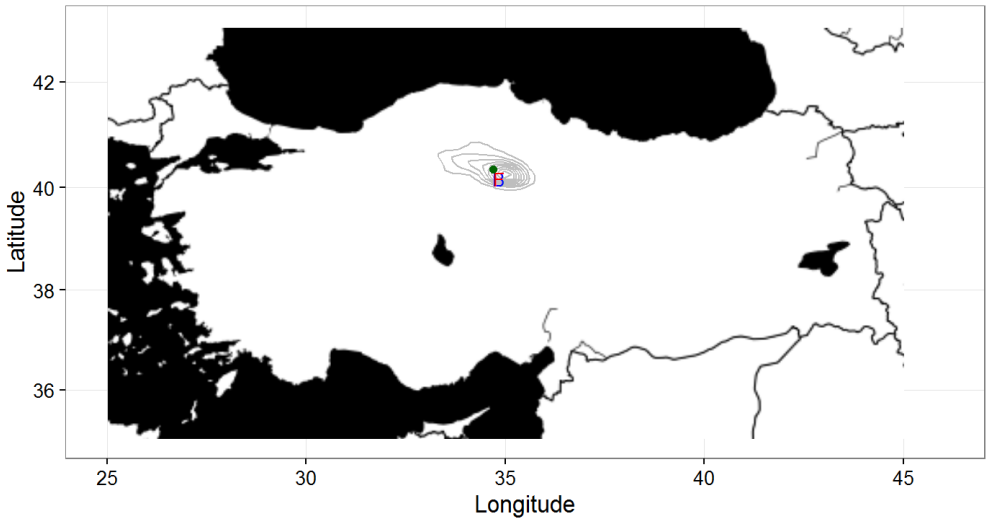

Panel B: Directional trade (N_{ij}^{travel})

Figure 4: Finding lost cities: *Sinahuttum*.

Notes: The maps show the estimated location for the ancient city of *Sinahuttum*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

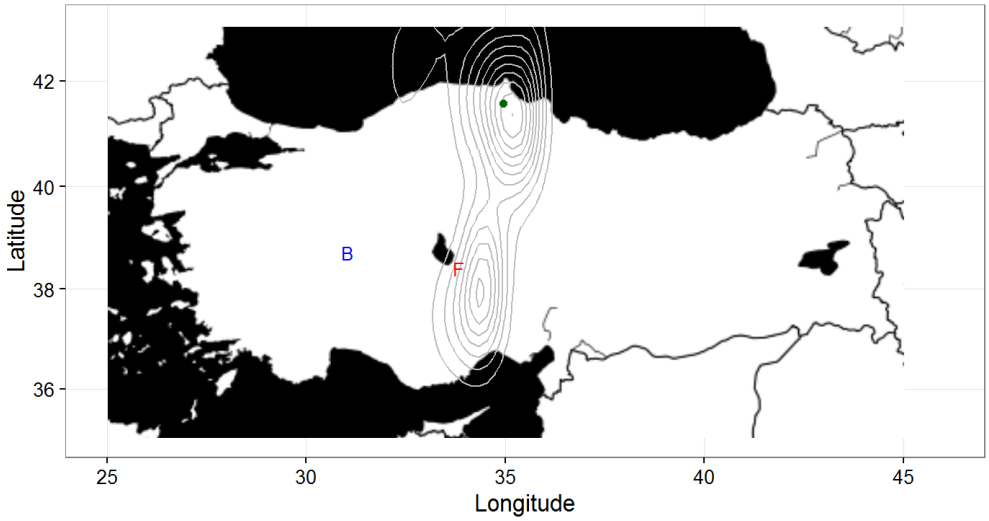

Panel A: Non-directional trade (N_{ij}^{joint})

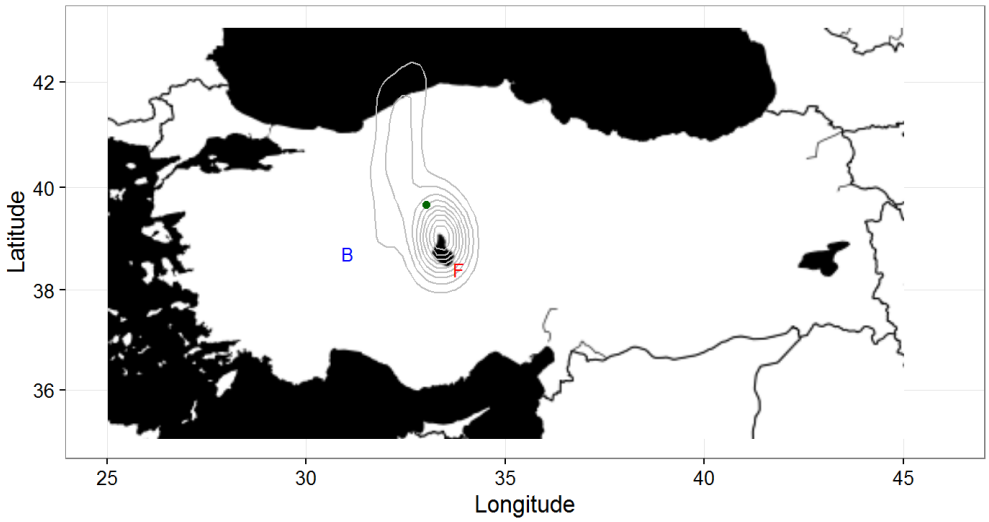

Panel B: Directional trade (N_{ij}^{travel})

Figure 5: Finding lost cities: *Purushaddum*.

Notes: The maps show the estimated location for the ancient city of *Purushaddum*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

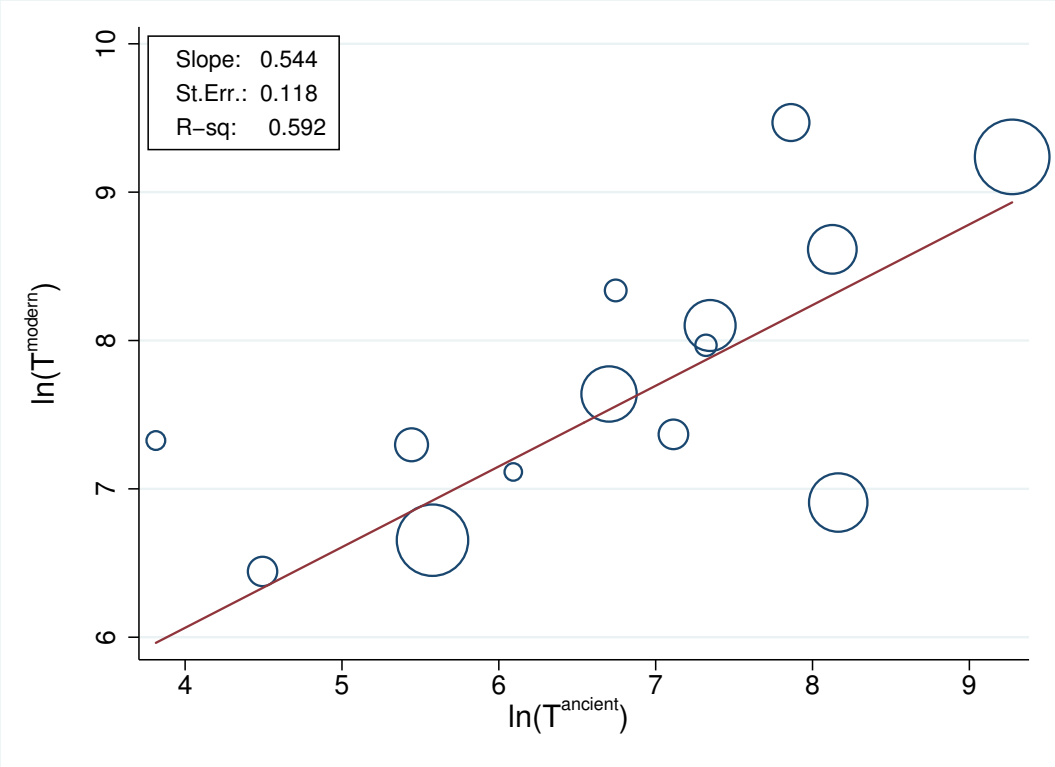

Figure 6: **Ancient and Modern City Sizes**

Notes: Each observation represents a present-day city in Turkey. Marker symbols are proportional to regression weights, which is number of mentions for ancient cities. $\ln(T^{1/\theta}|_{modern})$ in the y-axis is the size estimated using 2014 trade flows within the system of present-day cities corresponding to the ancient ones. $\ln(T^{1/\theta}|_{ancient})$ in the x-axis is estimated size of ancient cities. When multiple ancient cities fall within the same present-city boundary, we take their average size weighted by their number of mentions and then take natural logarithm. The regression line corresponds to the results in column 1 of table 7.

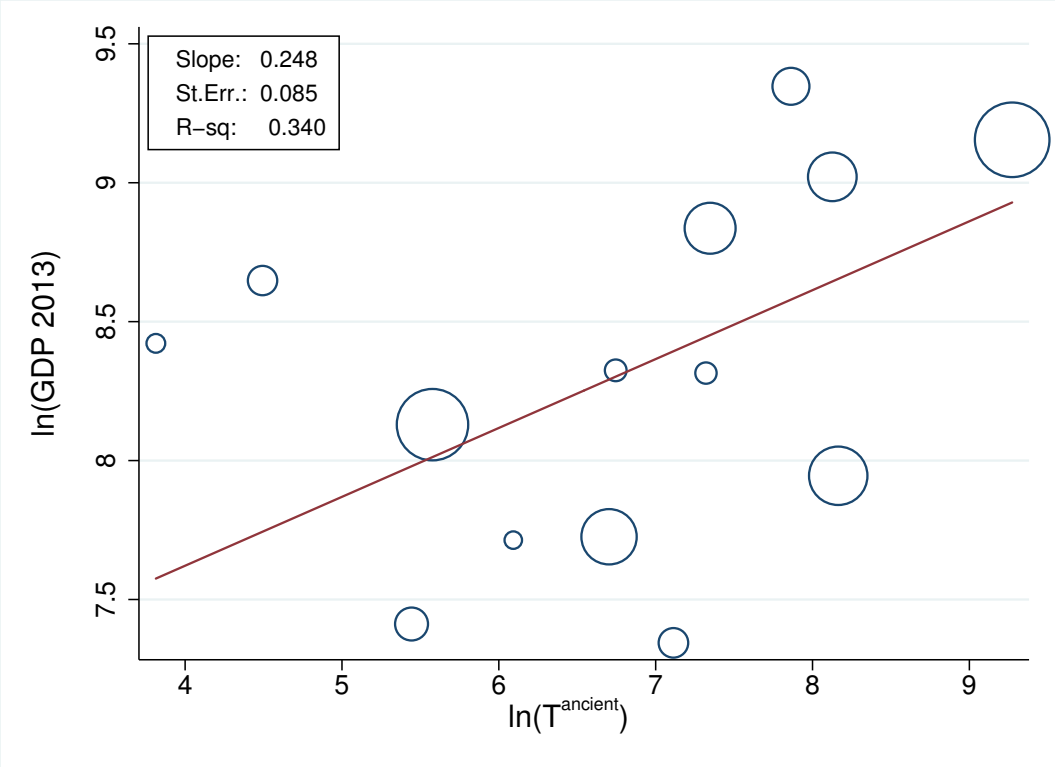

Figure 7: **Ancient City Sizes and Present-day GDP**

Notes: Each observation represents a present-day city in Turkey. Marker symbols are proportional to regression weights, which is number of mentions for ancient cities. $\ln(T^{1/\theta}|_{ancient})$ in the x-axis is estimated size of ancient cities. When multiple ancient cities fall within the same present-city boundary, we take their average size weighted by their number of mentions and then take natural logarithm. The regression line corresponds to the results in column 1 of table 8.

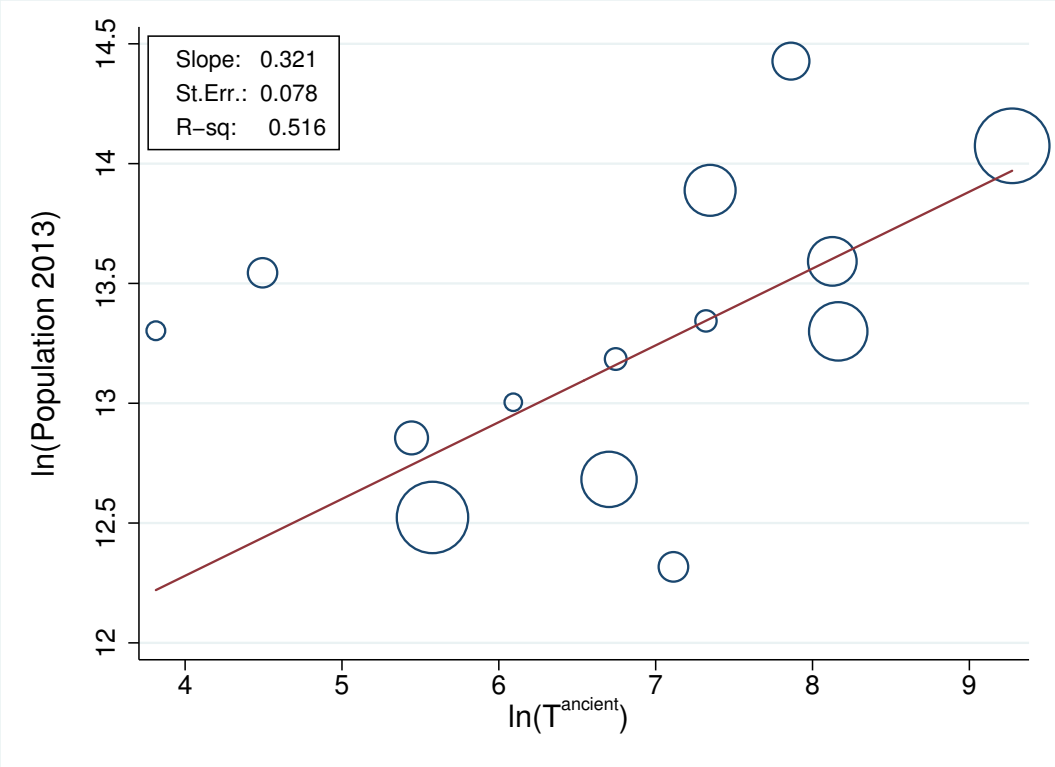

Figure 8: **Ancient City Sizes and Present-day Population**

Notes: Each observation represents a city in Turkey. Marker symbols are proportional to regression weights, which is number of mentions for ancient cities. $\ln(T^{1/\theta}|_{ancient})$ in the x-axis is estimated sizes of ancient cities. When multiple ancient cities fall within the same present-city boundary, we take their average size weighted by their number of mentions and then take natural logarithm. The regression line corresponds to the results in column 3 of table 8.

Tables

Table 1: **Descriptive Statistics**

	Mean	St. Dev.	Observations
Directional, human-read data			
Known cities			16
Lost cities			11
N_{ij}^{cargo} | all	0.49	1.98	650
N_{ij}^{cargo} | $N > 0$	3.24	4.15	98
N_{ij}^{travel} | all	0.56	2.17	702
N_{ij}^{travel} | $N > 0$	3.27	4.31	122
Non-directional, machine-read data			
Known cities			17
Lost cities			12
N_{ij}^{joint} | all	2.32	7.57	406
N_{ij}^{joint} | $N > 0$	5.62	10.97	168

Notes: The ancient data comes from a textual analysis of clay tablets inscribed in the cuneiform alphabet, sent by Assyrian merchants in the 2nd millennium BCE. Most texts are digitized and available through the Old Assyrian Text Project.

Table 2: **Distance Elasticity** (ζ)

	Non-directional Data (1)	Directional Data (2)
Distance elasticity: ζ	1.47	2.40
	(0.23)	(1.26)
Number of observations	406	702
Number of observations $\neq 0$	168	122
Number of trade flows	944	405
Number of starting values	2000	2000
Achieved minimum	1.540	4.384
Number of bootstraps (trials)	100	100
Number of bootstraps (successes)	83	65

Notes: This table presents the results of the estimates of the distance elasticity of trade (ζ) from (8). Collumn 1 uses non directional data on joint attestations of city names $\left(N^{joint}\right)$; column 2 uses directional data on persons' travels $\left(N^{travel}\right)$. Bootstrapped standard errors in parentheses. The lower panel of the table describes key statistics from the estimation.

Table 3: **Lost Cities Locations: Structural Estimates versus Historians Conjectures**

	Non-directional data		Directional data	
	Forlanini	Barjamovic	Forlanini	Barjamovic
	(1)	(2)	(3)	(4)
Durhumit	106	52	135	23
Hahhum	98	99	45	46
Kuburnat	30	44	141	108
Mamma	33	33	31	31
Ninassa	61	73	33	41
Purushaddum	226	285	98	125
Sinahuttum	10	10	15	15
Suppiluliya	55	50	53	52
Tuhpiya	66	52	156	133
Ursu	75	72	–	–
Washaniya	66	32	38	11
Zalpa	121	129	123	118
$\mathbb{E}[Distance]$	79	78	79	64
$\min[Distance]$	10	10	15	11
$\max[Distance]$	226	285	156	133

Notes: This table presents the distances (in kms) between our estimates of lost cities coordinates and the conjectured coordinates from historians Massimo Forlanini (Forlanini, 2008) (columns 1 and 3) and Gojko Barjamovic (Barjamovic, 2011) (columns 2 and 4). Columns 1 and 2 use non directional data on joint attestations of city names $\left(N^{joint}\right)$, while columns 3 and 4 use directional data on persons' travels $\left(N^{travel}\right)$. The lower panel presents simple statistics (mean, minimum and maximum).

Table 4: **City Sizes** $\left(T_i^{1/\theta}\right)$

Known cities	Non-directional Data (1)	Directional Data (2)	*Lost cities*	Non-directional Data (3)	Directional Data (4)
Amkuwa	45.92	105.37	Durhumit	100	89.23
	(59.66)	(144.58)		(0)	(432.50)
Hattus	29.64	100.00	Hahhum	248.58	352.63
	(65.52)	(0)		(139.38)	(1246.54)
Kanes	911.62	1064.49	Kuburnat	23.80	8.94
	(692.71)	(1761.75)		(28.88)	(105.22)
Karahna	4.04	4.58	Mamma	31.08	35.66
	(7.05)	(9.34)		(45.58)	(73.87)
Qattara	129.60	5484.75	Ninassa	16.30	20.79
	(303.71)	(57582.51)		(39.02)	(36.30)
Tapaggas	6.54	4.32	Purushaddum	574.98	4.05
	(4.36)	(8.85)		(657.41)	(77.53)
Hanaknak	24.51	50.30	Sinahuttum	11.43	62.50
	(21.56)	(66.41)		(22.22)	(41.89)
Hurama	59.83	210.32	Suppiluliya	3.14	0.01
	(53.97)	(358.59)		(42.38)	(11.54)
Malitta	0.00	6.09	Tuhpiya	23.10	401.38
	(3.32)	(9.00)		(24.22)	(898.34)
Salatuwar	104.34	338.15	Ursu	8.58	–
	(85.89)	(529.62)		(13.22)	–
Samuha	13.66	151.12	Washaniya	31.20	144.01
	(11.00)	(239.94)		(26.62)	(146.85)
Timelkiya	201.06	350.07	Zalpa	66.46	259.79
	(126.73)	(1198.05)		(119.51)	(1519.60)
Ulama	37.07	24.54			
	(32.81)	(67.39)			
Unipsum	2.32	33.88			
	(9.74)	(57.50)			
Ussa	8.72	–			
	(12.64)	–			
Wahsusana	332.79	26.43			
	(246.62)	(203.01)			
Zimishuna	6.25	0.43			
	(6.81)	(5.20)			

Notes: This table presents the results of the structural estimates of city sizes $\left(T_i^{1/\theta}\right)$ from (8) and (9), for both known cities (columns 1 and 2) and lost cities (columns 3 and 4). Columns 1 and 3 use non directional data on joint attestations of city names $\left(N^{joint}\right)$; columns 2 and 4 use directional data on persons' travels $\left(N^{travel}\right)$. Bootstrapped standard errors in parentheses.

Table 5: **City Sizes and Locational Fundamentals**

Dependent variable: $\ln(T^{1/\theta})$	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(elevation)$	-0.878	-0.754	-0.593	0.615	2.243	2.280
	(0.660)	(0.720)	(0.772)	(0.797)	(0.535)	(0.656)
$\ln(slope)$	0.489	0.335	0.523	-0.184	-1.129	-1.054
	(0.586)	(0.741)	(0.568)	(0.871)	(0.482)	(0.653)
$\ln(cereal)$	1.487	2.373	1.680	2.627*	0.250	1.793
	(0.419)	(0.126)	(0.367)	(0.063)	(0.899)	(0.519)
$\ln(river)$	0.505	0.631	0.515	0.673	0.619*	0.475*
	(0.194)	(0.110)	(0.201)	(0.106)	(0.065)	(0.068)
Copper		0.150		-1.467		1.340
		(0.949)		(0.595)		(0.695)
Gold		0.786		0.659		1.400
		(0.444)		(0.548)		(0.210)
Silver		0.555		0.512		-1.385
		(0.476)		(0.550)		(0.138)
N	27	27	24	24	16	16
R^2	0.197	0.378	0.223	0.458	0.412	0.540
Cities	All	All	Located	Located	Known	Known

Notes: Robust p-values are in parentheses. Dependent variable is the estimated sizes of ancient cities. *elevation* and *slope* are average altitude and slope of terrain within a 30 km radius of city coordinates, respectively. *cereal* is average the crop suitability value for low input level rain-fed cereals within a 30 km radius of city coordinates. *river* is the distance to the nearest river. Copper, gold and silver are distances to the nearest deposit of the particular metal. Columns 1-2 use all cities. Columns 3-4 exclude lost cities for which estimated locations significantly diverge from qualitative historical evidence. Last two columns use only cities with known locations.

Table 6: **City Sizes and Future Capitals**

Dependent variable: Future capital (binary)	(1)	(2)	(3)
$\ln(T^{1/\theta})$	0.118**	0.112*	0.147**
	(0.013)	(0.054)	(0.028)
$\ln(elevation)$		-0.143	-0.214
		(0.781)	(0.645)
$\ln(slope)$		0.0440	0.154
		(0.757)	(0.473)
$\ln(cereal)$		0.390	0.386
		(0.411)	(0.329)
$\ln(river)$		-0.0186	-0.044
		(0.863)	(0.595)
N	27	27	27
R^2	0.211	0.270	
pseudo-R^2			0.253
Method	LPM	LPM	Probit

Notes: Robust p-values are in parentheses. Dependent binary variable is equal to one for cities that have been a capital at some point of their history, and zero for others. *elevation* and *slope* are average altitude and slope of terrain within a 30 km radius of city coordinates, respectively. *cereal* is average the crop suitability value for low input level rain-fed cereals within a 30 km radius of city coordinates. *river* is the distance to the nearest river.

Table 7: **Ancient and Present-day City Sizes**

Dependent variable: $\ln(T^{1/\theta}\|_{modern})$	(1)	(2)
$\ln(T^{1/\theta}\|_{ancient})$	0.544***	0.397*
	(0.001)	(0.094)
$\ln(elevation)$		-0.404
		(0.769)
$\ln(slope)$		-1.395
		(0.157)
$\ln(cereal)$		-0.703
		(0.641)
$\ln(river)$		0.00466
		(0.986)
N	14	14
R^2	0.592	0.764

Notes: Robust p-values are in parentheses. Dependent variable is the estimated city sizes of present-day Turkish cities corresponding to the ancient locations. *elevation* and *slope* are average altitude and slope of terrain within a 30 km radius of city coordinates, respectively. *cereal* is average the crop suitability value for low input level rain-fed cereals within a 30 km radius of city coordinates. *river* is the distance to the nearest river.

Table 8: **Ancient City Sizes and Present-day City Characteristics**

	(1)	(2)	(3)	(4)
Dependent variable:	$ln(GDP)$	$ln(GDP)$	$ln(Population)$	$ln(Population)$
$\ln(T^{1/\theta}\|_{ancient})$	0.248**	0.373*	0.321***	0.275
	(0.013)	(0.062)	(0.001)	(0.137)
$\ln(elevation)$		-0.163		0.633
		(0.882)		(0.546)
$\ln(slope)$		-0.455		-0.824
		(0.486)		(0.225)
$\ln(cereal)$		-0.407		-1.582
		(0.704)		(0.151)
$\ln(river)$		-0.301		-0.0429
		(0.156)		(0.821)
N	14	14	14	14
R^2	0.340	0.569	0.516	0.663

Notes: Robust p-values are in parentheses. Dependent variables are total output (columns 1-2) and population of present-day Turkish cities corresponding to the ancient locations. *elevation* and *slope* are average altitude and slope of terrain within a 30 km radius of city coordinates, respectively. *cereal* is average the crop suitability value for low input level rain-fed cereals within a 30 km radius of city coordinates. *river* is the distance to the nearest river.

Technical Appendix for:
TRADE, MERCHANTS, AND THE LOST CITIES OF THE BRONZE AGE
(not for publication)

by Gojko BARJAMOVIC, Thomas CHANEY, Kerem COŞAR, and Ali HORTAÇSU

A Additional Data

A.1 Future Political Centers

- Kanes (known): Kayseri. The capital of the Kingdom of Cappadocia.
- Hattus (known): Corum, Bogazkale. The capital of the Hittite Kingdom.
- Hurama (known): Kahramanmaras, Elbistan. The capital of Dulkadir Emirate.
- Hahhum (lost): Adiyaman, Samsat. The capital of the Kingdom of Commagene.
- Mamma (lost): Kahramanmaras. The capital of neo-Hitite state Gurgum and also of Dulkadir Emirate after the sacking of Elbistan by Mongols.
- Durhumit (lost): Amasya. The capital of the Pontus Kingdom.

B Additional Figures

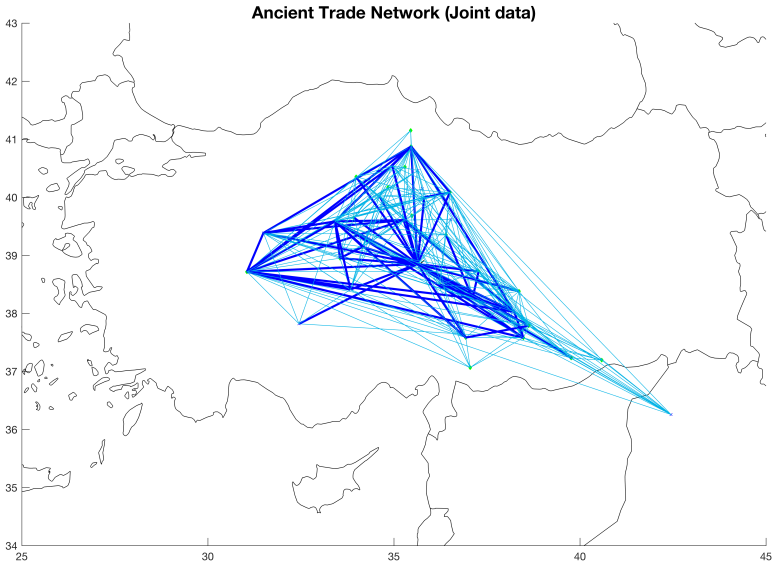

Figure B.1: Ancient trade network

Notes: Thick lines indicate $N_{ij}^{joint} > 4$ *in the non-directional data, and thin lines* $4 \geq N_{ij}^{joint} > 0$.

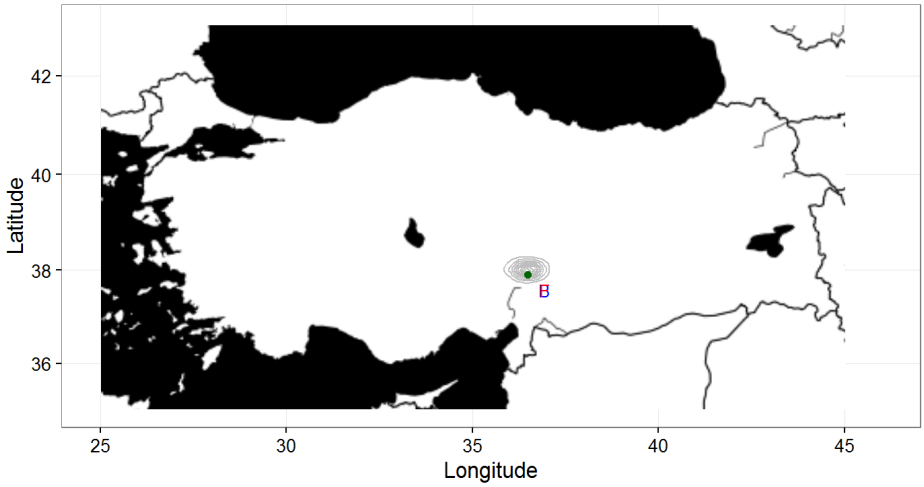

Panel A: Non-directional trade (N_{ij}^{joint})

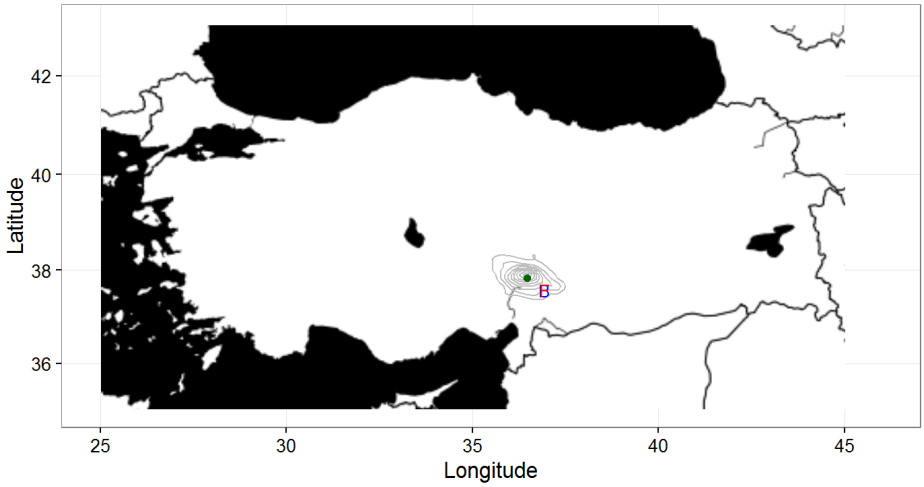

Panel B: Directional trade (N_{ij}^{travel})

Figure B.2: Finding lost cities: *Mamma*.

Notes: The maps show the estimated location for the ancient city of *Mamma*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

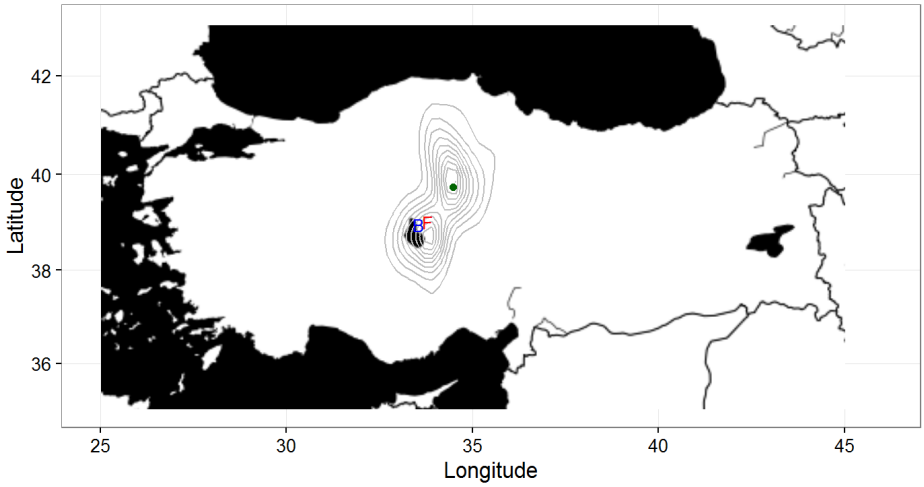

Panel A: Non-directional trade (N_{ij}^{joint})

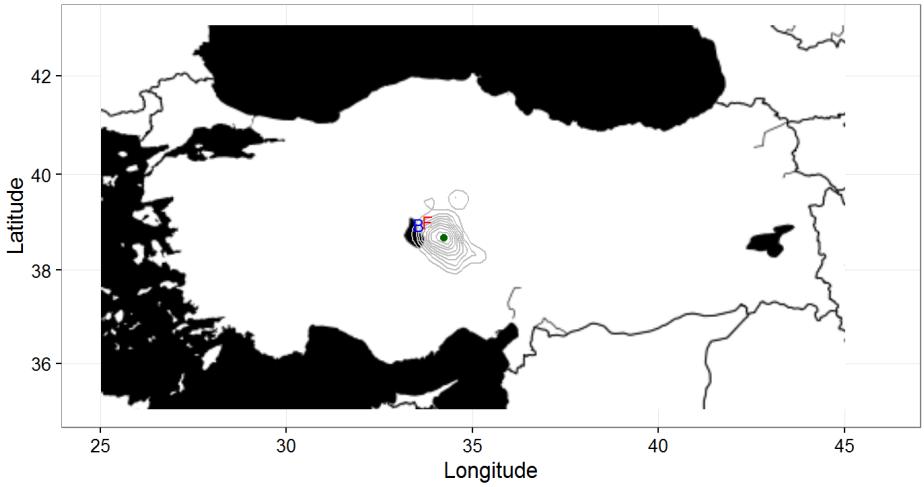

Panel B: Directional trade (N_{ij}^{travel})

Figure B.3: Finding lost cities: *Ninassa*.

Notes: The maps show the estimated location for the ancient city of *Ninassa*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

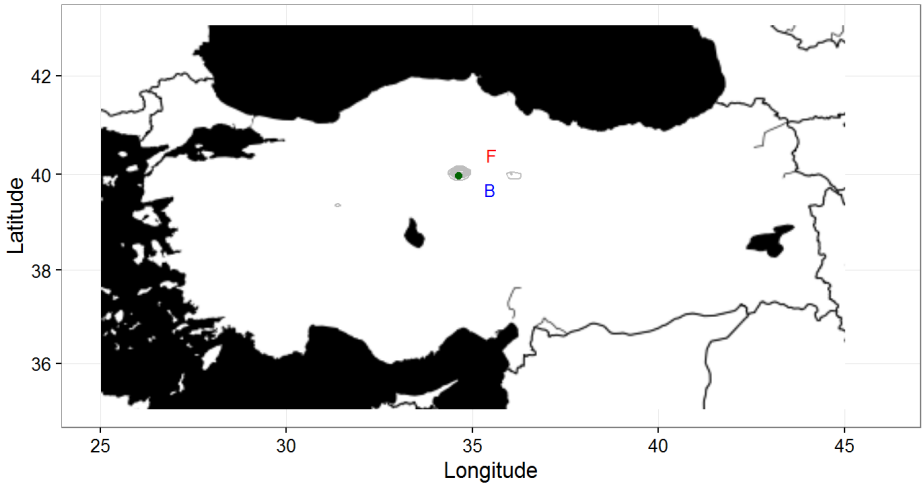

Panel A: Non-directional trade (N_{ij}^{joint})

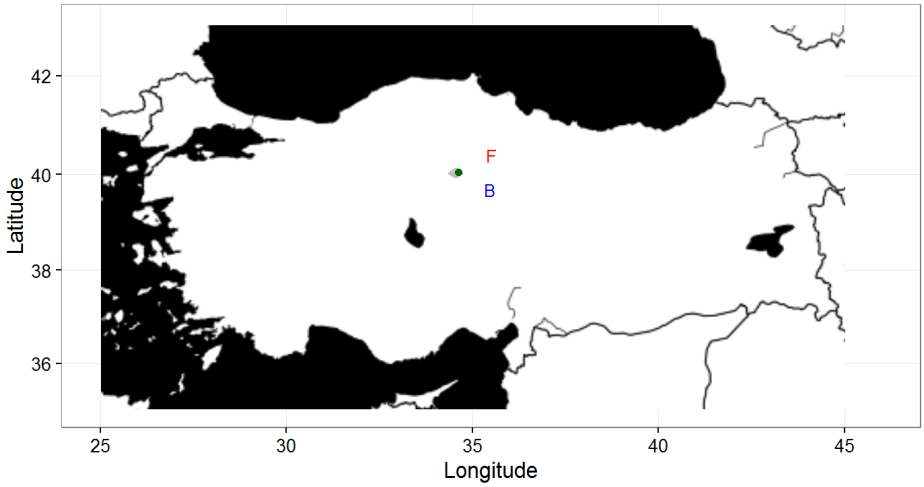

Panel B: Directional trade (N_{ij}^{travel})

Figure B.4: Finding lost cities: *Suppiluliya*.

Notes: The maps show the estimated location for the ancient city of *Suppiluliya*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

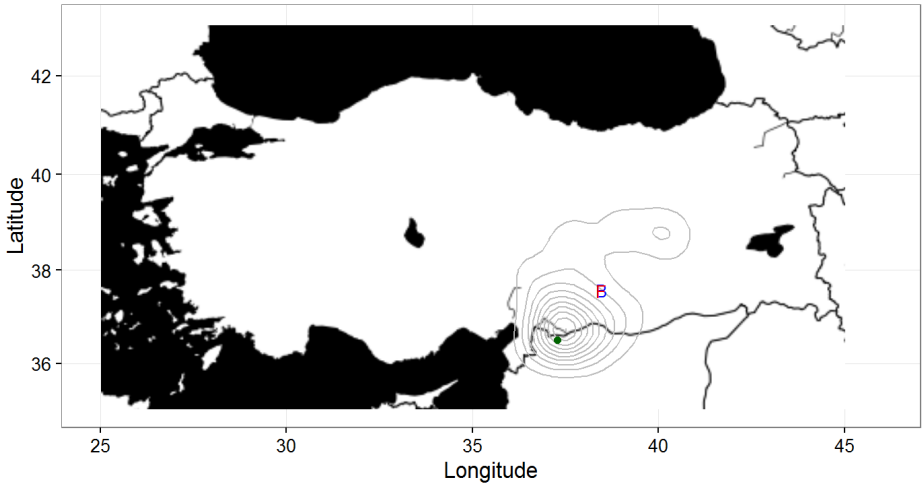

Panel A: Non-directional trade (N_{ij}^{joint})

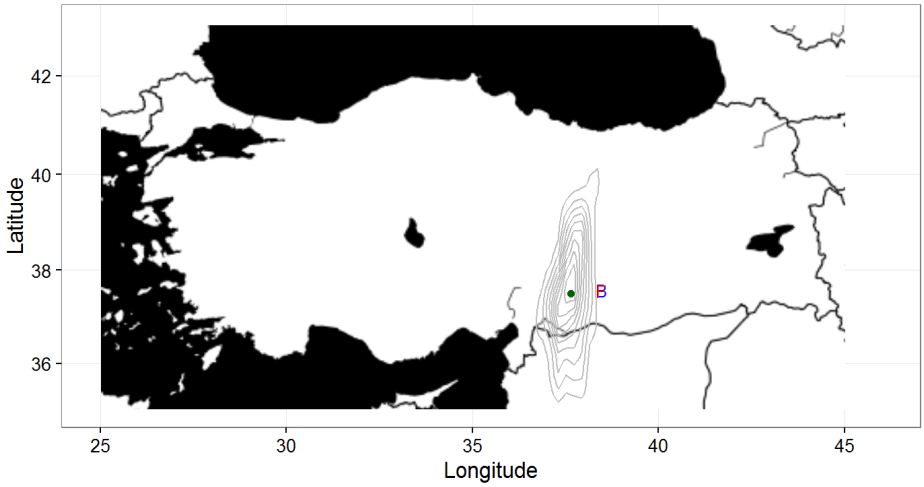

Panel B: Directional trade (N_{ij}^{travel})

Figure B.5: Finding lost cities: *Hahhum*.

Notes: The maps show the estimated location for the ancient city of *Hahhum*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

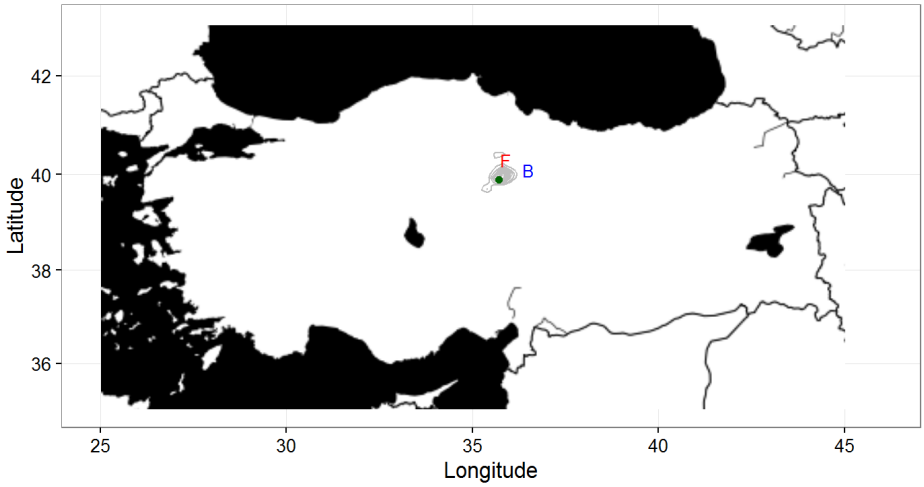

Panel A: Non-directional trade (N_{ij}^{joint})

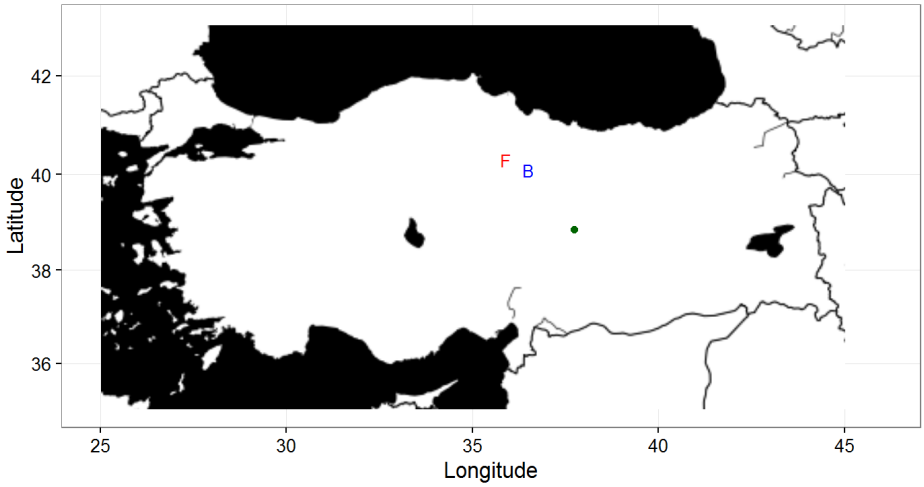

Panel B: Directional trade (N_{ij}^{travel})

Figure B.6: Finding lost cities: *Kuburnat*.

Notes: The maps show the estimated location for the ancient city of *Kuburnat*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

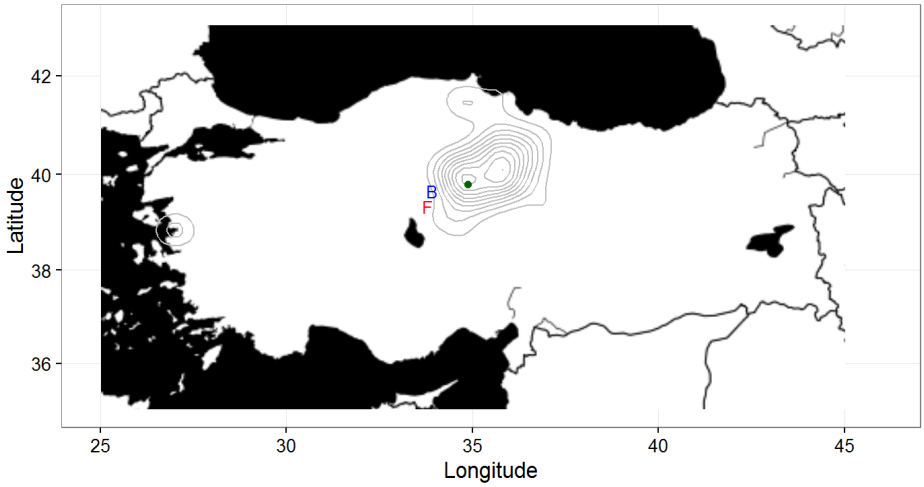

Panel A: Non-directional trade (N_{ij}^{joint})

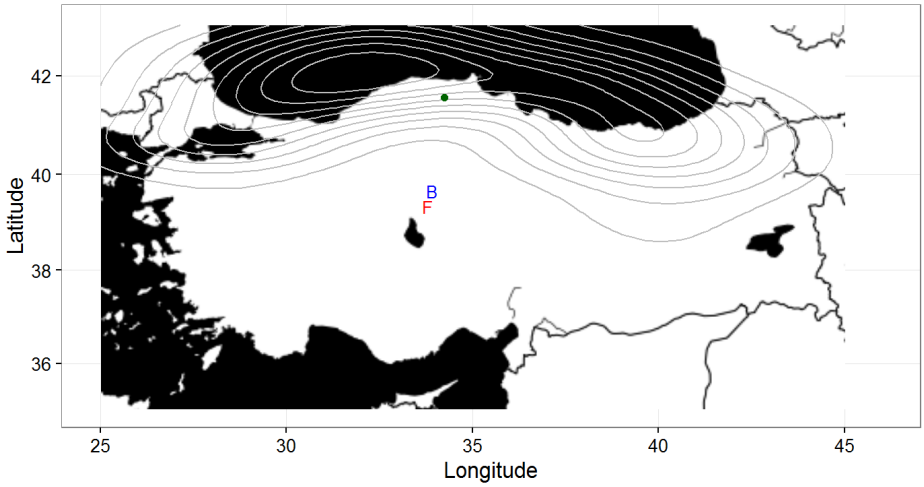

Panel B: Directional trade (N_{ij}^{travel})

Figure B.7: Finding lost cities: *Tuhpiya*.

Notes: The maps show the estimated location for the ancient city of *Tuhpiya*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

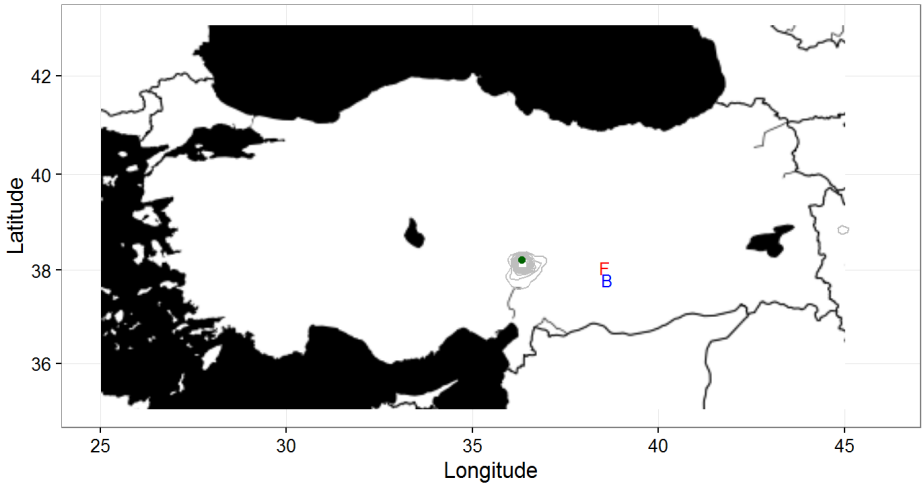

Panel A: Non-directional trade (N_{ij}^{joint})

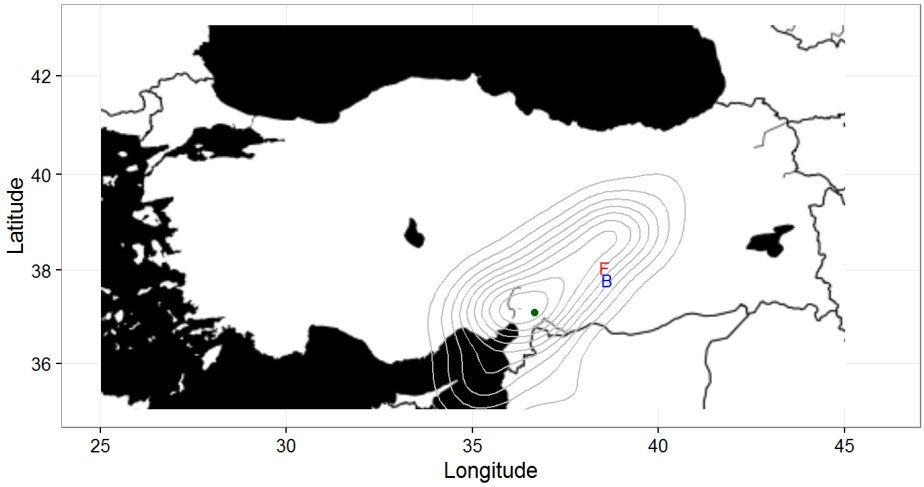

Panel B: Directional trade (N_{ij}^{travel})

Figure B.8: Finding lost cities: *Zalpa*.

Notes: The maps show the estimated location for the ancient city of *Zalpa*. In both maps, the location denoted by "B" corresponds to the site suggested by historian Gojko Barjamovic (Barjamovic, 2011), and the location denoted by "F" to the site suggested by historian Massimo Forlanini (Forlanini, 2008). Both historians base their suggestion on qualitative information collected from historical records, in particular the topography of the site, and references to landmarks identified in more recent historical texts. The green dot correspond to the estimated location from solving a structural gravity model, the estimation in (8). Panel A presents the estimates using non-directional data from joint attestations of city names $\left(N_{ij}^{joint}\right)$. Panel B presents the estimates using directional data on persons' travels $\left(N_{ij}^{travel}\right)$. The contours around the green dot are a contour plot of the confidence area drawn from 83 out of 100 bootstrapped trials (panel A) and 65 out of 100 bootstrapped trials (panel B). *Sources*: Old Assyrian Text Project.

C Imposing constraints on lost cities using merchants' itineraries

We propose to use the information about detailed itineraries of merchants in our dataset to impose a set of constraints on the location of lost cities. Those correspond to the description of the route followed by merchants as they travel between cities. A typical itinerary, which documents travels between both known and lost cities is found in the following excerpt from tablet Kt 83/k117:

> (···) *we [the envoys envoys and the Port Authorities of Wahšušana] sent two messengers by way of Ulama and two messengers by way of Šalatuwar to Purušhaddum to clear the order.* (···)

This letter describes messengers traveling from *Wahšušana* to *Purušhaddum* via *Ulama*, and *Wahšušana* to *Purušhaddum* via *Šalatuwar*. 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 $A \to B \to X$ where A and B are known, and X is lost. Under the assumption that stopping by B on the way from A to X does not represent too much of a detour, this information imposes some structure on the potential location of lost city X: X must be located in such a way that B is somewhere in-between A and X.

We systematically collect all information about itineraries, either of individual merchants or of caravans, involving both known and lost cities. More precisely, we extract information about itineraries of the form $A \to X \to B$, or $E \to F \to X$ where A, B, E, and F are known while X is lost. For each lost city X, we construct confidence regions as follows,

$$\mathbb{X}_\lambda = \{X : ||AX|| + ||XB|| \leq (1+\lambda)\,||AB|| \cap ||EF|| + ||FX|| \leq (1+\lambda)\,||EX|| \cap \cdots\}\,,$$

where the parameter $\lambda \geq 0$ indexes how much of a detour we allow. $||AB||$ is the distance from city A to city B, where we either use shortest distances as the crow flies, or shortest paths taking into account the topography of the land.[14] In plain words, the constraint $||AX||+||XB|| \leq (1+\lambda)\,||AB||$ says that stopping by X on the way from A to B represents at most a λ% detour compared to a direct trip from A to B.

Figures C.1 and C.2 offer a simple graphical example of the construction of the $\mathbb{X}_\lambda$ sets. Assume we have data on the following three itineraries involving lost city X: $A \to X \to B$, $C \to X \to D$, and $E \to F \to X$, where A, B, C, D, E, and F are known cities, while X is lost.

Figure C.1 shows how to use the information contained in the $A \to X \to B$ (left panel) and $E \to F \to X$ (right panel) itineraries. Under the assumption "going through X on the way from A to B does not represent more than a 5% detour from a direct A to B trip", i.e. the $A \to X \to B$ itinerary, city X must be within an ellipse with focal points A and B. Points X_1 and X_2 on the left panel of figure C.1 are two candidates for a 5% detour. The darker shaded concentric ellipses around A and B represent shorter detours. Under the assumption "going through F on the way from E to X does not represent more than a 5% detour from a direct E to X trip", i.e. the $E \to F \to X$ itinerary, city X must be within sort of distorted circle in the alignment of the EF segment. Points X_1 and X_2 on the right panel of figure C.1 are two candidates for a 5% detour. The darker shaded regions represent shorter detours (from 0.3% to 5%).

Figure C.2 shows how to combine the information on different itineraries involving the same lost city X. With a 5% detour at most, each of the three itineraries $A \to X \to B$, $C \to X \to D$, and $E \to F \to X$ generates a different candidate region for the location of lost city X. The intersection of those three regions on the left panel of figure C.2 correspond to the set of locations where the "5%

[14] In this second case, we use Dijkstra's algorithm to define the shortest path between any two points on the map A and B.

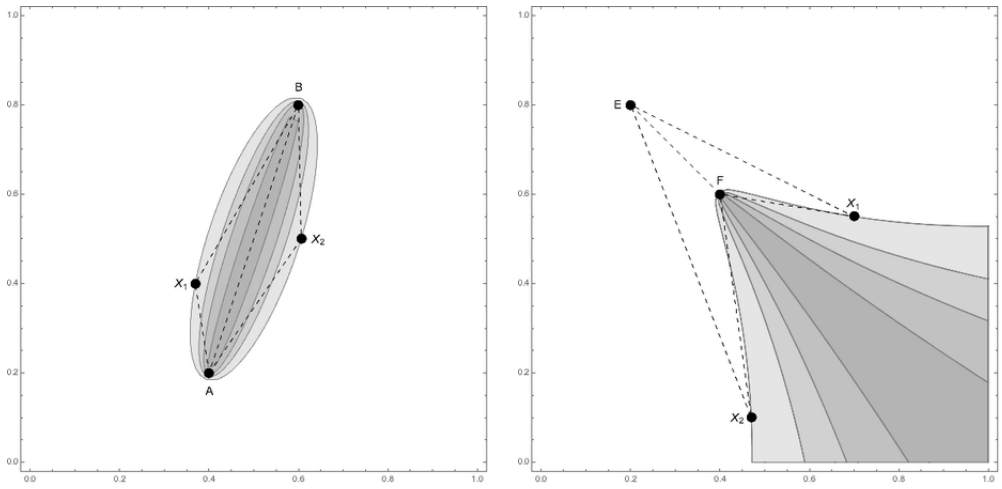

Itinerary $A \to X \to B$ *Itinerary* $E \to F \to X$

Figure C.1: Finding lost cities using merchants' itineraries (I).

Notes: Both figures show examples of how to use data on a single merchants itinerary to constrain the likely location of lost cities. The left panel shows the example of an itinerary of the type $A \to X \to B$, where A and B are know, and X is lost. Points X_1 and X_2 are two possible candidates such that going from A to B via X_1 (or X_2) represents only a 5% detour compared to going straight from A to B. All points inside the shaded area (an ellipse with focal points A and B) correspond to detour of no more than 5%. Darker shades of grey correspond to shorter detours (0.3% to 5%). The right panel shows the example of an itinerary of the type $E \to F \to X$, where E and F are know, and X is lost. Points X_1 and X_2 are two possible candidates such that going from E to X_1 (or X_2) via F represents only a 5% detour compared to going straight from E to X_1 (or X_2). All points inside the shaded area correspond to detour of no more than 5%. Darker shades of grey correspond to shorter detours (0.3% to 5%).

detour at most" constraint is satisfied on all three itineraries. Varying the 5% threshold (from 1.5% to 15%), we draw a series of interlocking regions on the right panel of figure C.2. Darker shaded regions correspond to shorter detours.

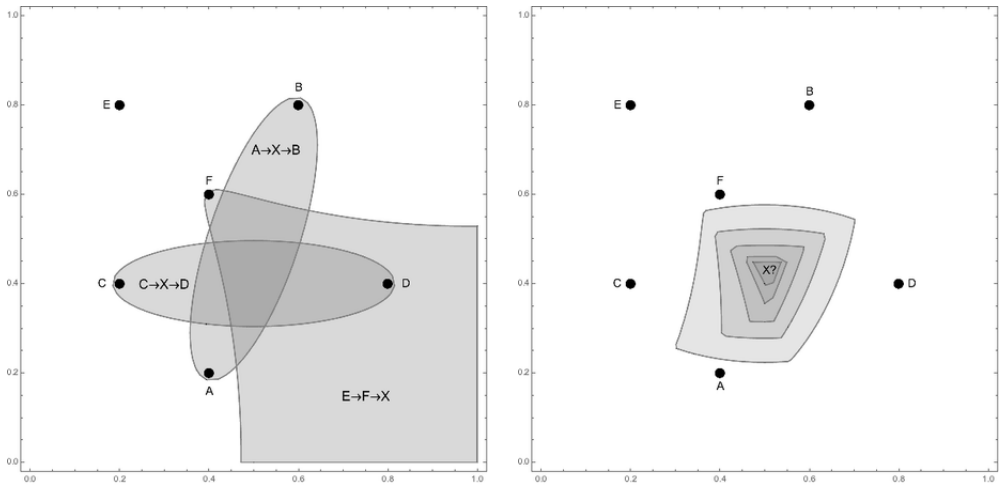

Figure C.2: Finding lost cities using merchants' itineraries (II).

Notes: Both figures show an example of how to combine data on several merchants itineraries to constrain the likely location of lost cities. In this example. we use three itineraries, $A \to X \to B$, $C \to X \to D$, and $E \to F \to X$. The left panel shows the regions corresponding to an at most 5% detour from a straight line separately for each itinerary. The intersection of the three regions corresponds to a set of points where the constraint "no more than a 5% detour" is satisfied for all three itineraries. The right panel shows the intersection of those three constraints for different lengths of detours. Darker shades of grey correspond to shorter detours (1.5% to 15% detours from the straight line).