Riding the Monsoon: Geography and Iron Age Trade in the Indian Ocean

Conrad Copeland*

Working Paper – August 2022

Abstract

This paper exploits ancient textual sources to develop a database of ancient trade in the Indian Ocean and model trade in the region during the Iron Age. Wind-speed data is used to construct a gravity model of trade and is combined with detailed textual data from the Periplus of the Erythraean Sea to analyse historical development trends in the Indian Ocean. While distance was an important factor in maritime trade, the speed of travel was the defining feature with a substantially non-linear effect. Trade is shown to be lower between locations that are close together with similar export baskets, reinforcing an endowment-based, Ricardian specification of the gravity model for this period. Additionally, there is significant evidence for both export-led growth and an ancient version of the 'resource curse'. Cities that export a greater variety of goods exhibited a significant increase in density during the period. Similarly, areas around cities with exports that relied more heavily on manufactures and artisanal goods grew at a faster rate than areas around cities that focussed more on cash crops.

Keywords: International Trade, Economic History, Ancient Trade Networks, Economic Geography.

JEL codes: F14, F15, F63, N75, N77.

---

*Department for International Trade and UCL; conrad.copeland@ucl.ac.uk
I am grateful to Helen Simpson and Yanos Zylberberg for their advice and guidance; and Ferdinand Rauch and Toman Barsbai for helpful discussions and comments. I also thank seminar participants at the University of Bristol, LSE, the WEIA conference, and ASSA/AEA conference for helpful comments. All errors are my own, and identifying them is left as an exercise for the reader.

1. Introduction

Trade in the ancient world was often long, difficult, and dangerous. Small wooden ships would ply the oceans with rudimentary navigational tools in the hope of landing at their destination with their produce intact. Trade costs were high, reducing the attractiveness of long-distance, trans-oceanic trade to all but the most intrepid of merchants. The Indian Ocean in the late Iron Age was an exception to this. The unique climatic elements of the seasonal monsoon made trade from the Horn of Africa to the coasts of India both feasible and attractive. A merchant setting out from one side of the Ocean could, provided they left at the correct time, travel to the other shore with relative ease.

This paper exploits ancient textual sources to develop a database of ancient trade in the Indian Ocean and model trade in the region during the late Iron Age. Using the trade information logged in the Periplus of the Erythraean Sea (Schoff, 1912), I build a comprehensive dataset of ancient trade cataloguing products, origins, and destinations. I supplement this with wind data to calculate sailing times between the various cities along the shores of the Indian Ocean to fit a gravity model to the network. In order to perform a deeper analysis on the cities in the Periplus, ancient development levels are captured by archaeological data on structures and dig sites across the region.

I find that trade between cities along the shores of the Indian Ocean is well-modelled by a gravity framework, implying that long-distance trade between the continents bordering the Indian Ocean was made feasible by the reduced travel time created by sailing conditions and monsoon winds. Further, I find that there is a significant resource endowment effect influencing trade flows throughout the region and creating a non-linear relationship between distance and trade. Cities that are very close together and have similar resource endowments trade less than cities with different endowments. This supplements the traditional technological comparative advantage of a Ricardian model with a resource-derived element of comparative advantage. I also find evidence for a relationship between the composition of exports and the level of development of cities. Cities that export more products and more diverse products tend to have higher development levels. This effect is largely driven by manufacturing products with less developed cities being more dependent on cash-crops, such as spices and incense, implying a possible ancient version of the 'resource curse'.

This paper makes two primary contributions. The first is to model ancient trade in an under-studied region with a gravity equation. The second is to provide insight into potentially trade-driven determinants of development in the region during the late Iron Age. In doing so, it extends the analysis of ancient trade flows into a new geographic region and adds to the discussion of driving factors of ancient trade. The majority of work on modelling trade with gravity models deals with post-eighteenth century time periods.[1] The work that has been done on ancient trade principally deals with the Mediterranean and the Middle East, rarely stretching further east or south than Mesopotamia and Egypt. This paper extends this type of trade analysis to a new region. In contrast to Bakker et al. (2018) which find a significant role for local and regional connectedness in the Mediterranean, the peculiar conditions of the Indian Ocean allowed countries with different resource endowments to engage in sophisticated, long-distance trade with each other; leading to a non-linear relationship between distance and trade. This result is similar to that of Barjamovic et al. (2019) by finding that ancient trade in the Indian Ocean is well modelled by a gravity model but adds a non-linear element based on resource endowments. Further, the evidence contributed in this work on the link between trade and development compliments the work of Bakker et al. (2018). Whereas Bakker et al. (2018) look at the effect of connectedness on development levels, this paper links the composition of exports and trade flows with development levels to illustrate a more direct relationship between export goods and economic activity around cities.

This work relates to the literature that documents the relationship between trade costs, communication costs, and trade flows. The unique climatic conditions of the Indian Ocean created transportation efficiencies that effectively increased the market access of cities along its shores. In this sense, the mechanisms and results of this paper inherit much from the literature on market access. Feyrer (2009) is a good example of this with his study on changes in trade flows due to the closure of the Suez Canal, but similar work can be found in Maurer and Rauch (2019), Juhasz (2018), and Liu and Meissner (2015) among others.

The results in this paper build on the literature around the positive effect of trade on development, including work that discusses the development effects of trade in endowment economies. The wider idea that trade links can lead to growth is a common theme throughout the development literature (see Redding and Venables, 2004; Wacziarg and Welch, 2008; Feyrer, 2019, for examples). I build on this tradition with the analysis of exports and development in the Indian Ocean. Another branch of literature in this area is the analysis of the 'resource curse' as outlined in Auty (1993) and Auty (2001).[2] The evidence in this paper on the relationship between export sectors and development extends the discussion of resource endowments and growth

[1] Work addressing the early-modern period includes Pascali (2017) who analyses trade during the transition from sail to steam shipping.

[2] Ross (1988) provides a survey of the trends in this research.

backwards into the ancient world.

This paper exploits historical data about ancient economies, which often require a different approach than more traditional datasets. Trade in the ancient world is frequently dealt with indirectly (as in Fluckiger et al., 2019; Broodbank, 2006) because of a lack of reliable data. Despite this, textual analysis and other techniques have proved useful in developing reliable data for research. Temin (2013) uses multiple techniques to develop useful price and trade information about Roman grain markets. Similarly, Temin (2006) and Algaze (2008) both use a combination of archaeological evidence and textual data in reconstructing trade in ancient periods. The techniques used in this paper add to this trend by exploiting written sources for trade data.

This paper is divided into six sections. Section 2 describes the data sources used for the analysis. A Ricardian version of a gravity model is outlined in Section 3 and Section 4 contains the results of the gravity estimation. Section 5 discusses the link between trade in the region and ancient levels of development. Section 6 concludes.

2. Data

The data in this paper combine archaeological, textual, and geographic information to develop a comprehensive picture of trade in the ancient Indian Ocean. The following subsections outline the different data sources and their uses in terms of the information they provide for the overall analysis.

2.1 The Periplus of the Erythraean Sea

Economic data for ancient eras are notoriously difficult to come by. Because of this, textual data have become more prominent in the literature (see Barjamovic et al., 2019; Temin, 2013; Algaze, 2008; Temin, 2006, for examples). Building on this tradition, the principal data for the gravity estimation come from a first-century document titled the Periplus of the Erythraean Sea[3] (written by an anonymous author and frequently identified by its Latin title: Periplus Maris Erythraei, PME).[4] The term 'periplus' comes from a Greek term meaning "sailing around". These types of documents were not an uncommon form of writing in ancient periods and were generally written by merchants, sailors, and other travellers to describe their

[3] The name Erythraean Sea was a Greek maritime designation for the Red Sea and Gulf of Aden that also frequently included what is now known as the Persian Gulf and the wider Indian Ocean.

[4] The translation used for this paper is Schoff (1912).

journeys for the benefit of future voyagers (and also to laud their achievements).[5] This periplus was likely written by a Greek merchant in Roman Egypt for the benefit of other merchants looking to ply their trade in the Indian Ocean. It follows a circular route around the coast, mentioning various cities and ports and the goods each trades. Figure 1 presents a map of the cities mentioned in the Periplus.

The information contained in the Periplus is remarkably complete and detailed. The source mentions trade goods, origins, destinations, product qualities, and shipment quantities. Ideally, all this information would be used to construct a comprehensive dataset of trade of a quality similar to the trade data used in modern trade analysis. Unfortunately, since its purpose was one of general description, not all of the information provided in the text is usable in this way. The information on origins, destinations, and products is the most comprehensive and usable for a detailed trade analysis, despite some gaps. The information on quality and quantity is less informative due to it being mostly subjective and relative, preventing the calculation of any value for the trade flows.

Figure 1. World of the Periplus

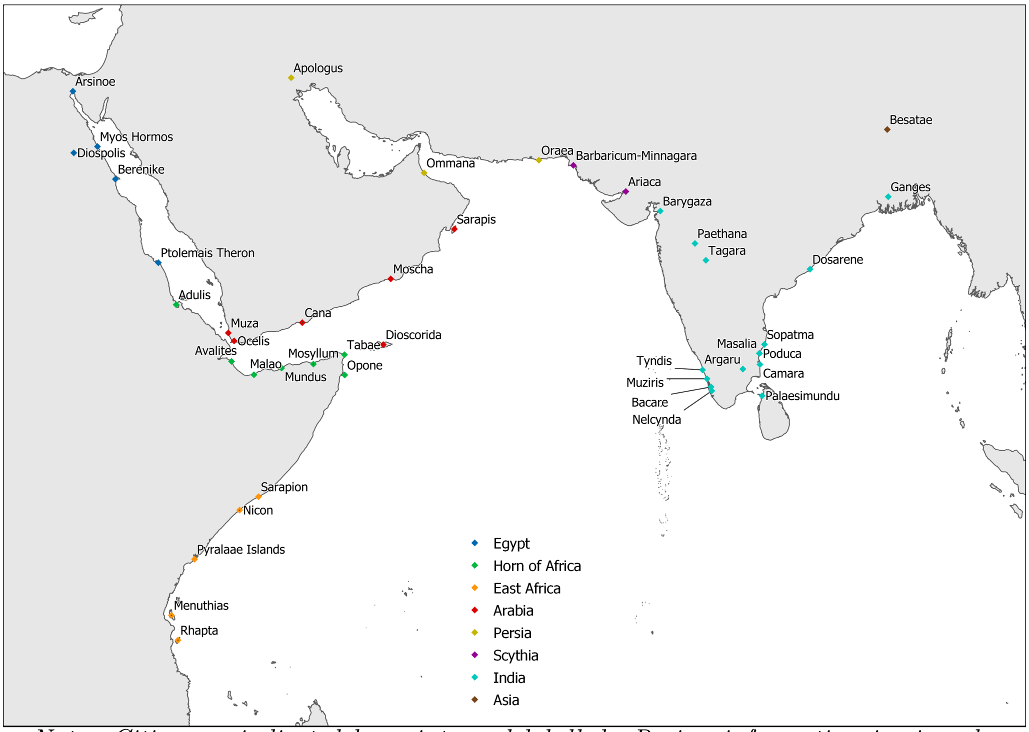

Note: Cities are indicated by points and labelled. Region information is given by colour. (Source: Author's own calculations based on Schoff, 1912).

[5]Notable examples include the works of Hanno the Navigator, Pseudo-Skylax, and Arrian.

Data Description The detail of the document provides a very comprehensive catalogue of goods exported and imported by each city mentioned and the origins and destinations of those same goods. Much of the text records a specific port of destination (origin) for each exported (imported) good:

> *Sailing through the mouth of the Gulf, after six-days' course there is another market town of Persia called Ommana. To both of these market-towns [Apologus and Ommana] large vessels are regularly sent from Barygaza, loaded with copper and sandalwood and timbers of teakwood and logs of blackwood and ebony. (PME, paragraph 36)*

This paragraph outlines specific shipments of copper and various types of timber from Barygaza to the cities of Apologus and Ommana. With this level of detail, identifying origins, destinations, and goods is a straightforward task.

The share of each broad product category in the exports and imports of the different cities is shown in Figure 2. The graphs show a regional distribution of some exports, with different products dominating the export baskets of specific regions. This similarity in export baskets implies a regional structure to production, emphasising the potential role of resource endowments in the products offered for exports (a feature that is explored further in Section 4). This regional correlation of products is less apparent in the imports of cities. The larger variation in import goods emphasises the features of trade in the area: traders were able to travel long distances across the Indian Ocean with relative ease, making a larger variety of goods available for import throughout the network. The monsoon system contributed to the development of this more sophisticated long-distance trading network that allowed for significant expansion beyond local or regional mercantile activity.

The cities mentioned in the Periplus are mostly locatable due to consistency of names across time and comprehensive archaeological work in the region. Where there is no precise location known, there is often a scholarly consensus among archaeologists about city locations. This analysis makes use of these agreed locations of the cities in the data. Figure 1 shows the locations and names of each city mentioned in the Periplus.

Issues Many of the passages of the Periplus follow the format above, where destination, origin, product, and other details are clearly indicated and explicitly mentioned. Despite this, there are cases in which the information is incomplete. Frequently in these cases the destination or origin of the product being traded is given as a region and not a specific city:

Figure 2. Goods Categories as Share of Exports and Imports

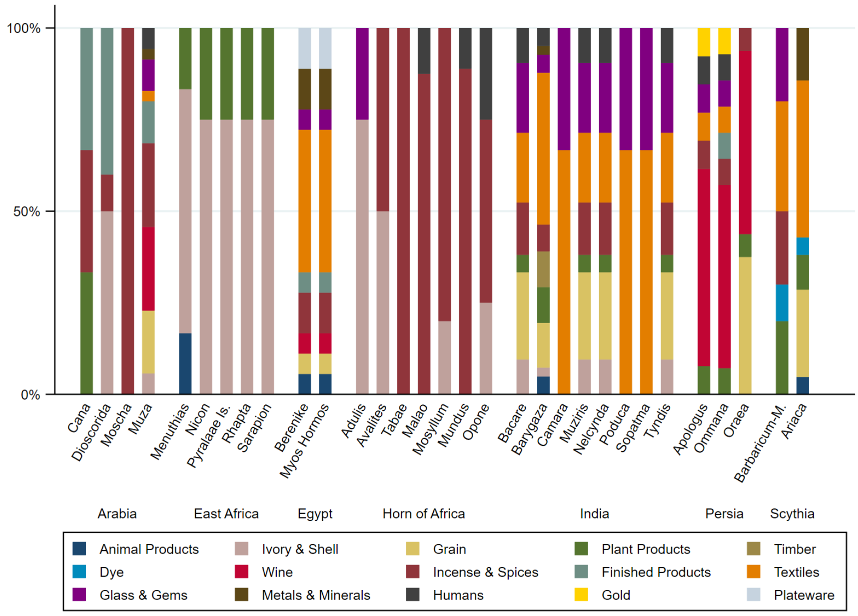

(a) Export Shares

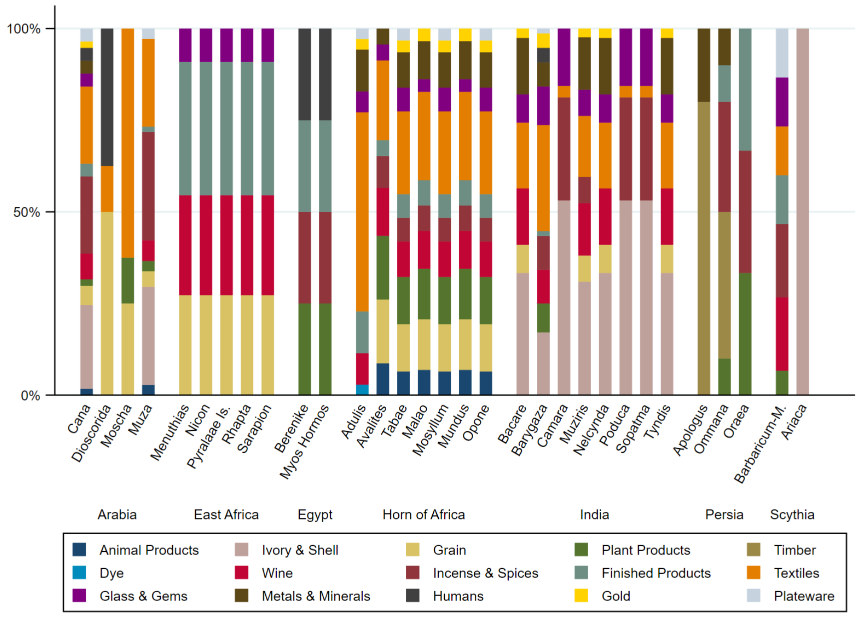

(b) Import Shares

Note: Bar colours display the share of good categories each city exports and imports. Only cities with positive values for both exports and imports are included (Source: Author's own calculations based on Schoff, 1912).

> *After Avalites there is another market-town, better than this, called Malao, distant a sail of about eight hundred stadia...There are exported from these places myrrh, a little frankincense, (that known as far-side), the harder cinnamon, duaca, Indian copal and macir, which are imported into Arabia; and slaves, but rarely. (PME, paragraph 8)*

From this paragraph, there are identified shipments of myrrh, frankincense, and other products to Arabia, a region with no cities explicitly mentioned in this context. In cases such as this, I interpolated the relevant cities using information on the goods traded by various cities in the named region. In this way, if a city located in Arabia is identified in a different paragraph as importing myrrh or frankincense it is designated as a partner city for Avalites and Malao with those goods. Both the non-interpolated data and the interpolated data are used in the analysis.

Due to the nature of the source, the terms used to describe the attributes of the trade goods are inconsistent in many places. While the name of each good is consistent across cities, the quality of the goods traded is generally rather vague and imprecise:

> *There are imported into this market-town, wine, Italian preferred, also Laodicean and Arabian; copper, tin, and lead; coral and topaz; thin clothing and inferior sorts of all kinds; bright-coloured girdles a cubit wide; storax, sweet clover, flint glass, realgar, antimony, gold and silver coin, on which there is a profit when exchanged for the money of the country; and ointment, but not very costly and not much. (PME, paragraph 49)*

As this paragraph illustrates, terms such as 'inferior' and 'not very costly' are used along with named varieties (such as Italian and Laodicean wine) or vague references to exchange rates of precious metals ('on which there is a profit'). While these passages contain significant detail about the quality or value of goods, they tend to be imprecise or relative descriptions. This prevents reliable calculations of value for the products identified in the document.

Quantities are often discussed in equally vague terms:

> *Beyond Mundus, sailing toward the east, after another two days' sail, or three, you reach Mosyllum, on a beach with a bad anchorage. There are imported here the same things already mentioned, also silver plate, a very little iron, and glass. There are shipped from the place a great quantity of cinnamon, (so that this market-town requires ships of a larger size), and fragrant gums, spices, a little tortoise shell, and mocrotu, (poorer than*

that of Mundus), frankincense, (the far-side), ivory and myrrh in small quantities. (PME, paragraph 10)

As with the quality of items mentioned, phrases such as 'very little', 'small quantities', and a 'great quantity' that requires larger ships make it impossible to identify the true quantities of traded products.

These difficulties with both the quality and quantity mean that it is not possible to reliably calculate volumes or prices for goods traded by various cities, and so it is impractical to measure trade flows in value-terms. For this reason, and for the purpose of theoretical consistency that will be discussed in Section 3, trade volume is measured as the number of different goods traded. A list of different products, product categories, and the number of cities that import and export them is given in Table 1.

Trade Variable Construction In order to construct the variable for trade used in this analysis, the number of different products a city imports or exports is summed across all product data. This is a straightforward process for bilateral trade between cities:

$$X_{ij} = \sum_{g} x_{ijg}$$

$$\text{with } x_{ijg} = \begin{cases} 1 & \text{if good } g \text{ is traded between cities } i \text{ and } j; \\ 0 & \text{otherwise.} \end{cases}$$

Where x_{ijg} is an indicator function for whether a given product g is traded between cities i and j.

The variables for export and import diversity for individual cities are constructed as modified Herfindahl indices across the export (import) product categories (as listed in Table 1) of a city:

$$h_i = 1 - \sum_{n} \left(\frac{\lambda_{in}}{L_i} \right)^2$$

Where λ_{in} is the number of products within the category n exported (imported) by city i and L_i is the total number of products exported (imported) by city i. This measure is constructed so that higher values indicate a higher level of product diversity.

Table 2 lists the cities from the Periplus, their coordinates, the number of goods they trade, and the number of other cities they trade with. The trade network mapped by the Periplus is quite heterogeneous; the cities in various regions do not

Table 1. Products of the Periplus

Product	Product Category	X	M	Product	Product Category	X	M
Clarified Butter	Animal Products	6	2	Arabian Clothing	Textiles	1	2
Horses	Animal Products	1	2	Blankets	Textiles	1	2
Mountain Tortoise	Ivory & Shell	1	1	Cloaks	Textiles	6	3
Tortoise Shell	Ivory & Shell	12	9	Coloured Sashes	Textiles	1	2
Ivory	Ivory & Shell	9	6	Girdles	Textiles	7	2
Elephant Ivory	Ivory & Shell	1	1	Local Clothing	Textiles	2	3
Rhino Ivory	Ivory & Shell	6	2	Robes	Textiles	1	1
Rice	Grain	7	7	Skin Coats	Textiles	1	1
Wheat	Grain	14	9	Thin Clothing	Textiles	1	2
Indian Copal	Plant Products	2	1	Tunics	Textiles	5	1
Aloes	Plant Products	1	12	Indian Cloth	Textiles	1	6
Macir	Plant Products	2	1	Cloth	Textiles	2	5
Palm Oil	Plant Products	5	1	Cotton Cloth	Textiles	7	2
Sesame Oil	Plant Products	7	6	Mallow Cloth	Textiles	2	2
Dates	Plant Products	2	2	Monache Cloth	Textiles	7	2
Sour Grape Juice	Plant Products	6	1	Muslins	Textiles	3	4
Sugarcane Juice	Plant Products	6	2	Purple Cloth	Textiles	1	2
Blackwood	Timber	2	1	Raw Silk	Textiles	1	5
Ebony	Timber	2	1	Sagmatogene Cloth	Textiles	7	2
Sandalwood	Timber	2	1	Silk Cloth	Textiles	4	5
Teakwood	Timber	2	1	Silk Yarn	Textiles	1	5
Coloured Lac	Dye	1	1	Carnelian	Gems & Glass	1	1
Purple	Dye	2	2	Coral	Gems & Glass	1	2
Wine	Wine	8	6	Diamonds	Gems & Glass	3	4
Duaca	Incense & Spices	2	1	Pearls	Gems & Glass	5	6
Fragrant Ointments	Incense & Spices	1	2	Sapphires	Gems & Glass	3	4
Frankincense	Incense & Spices	4	13	Transparent Stones	Gems & Glass	3	4
Myrrh	Incense & Spices	3	3	Flint Glass	Gems & Glass	1	1
Spikenard	Incense & Spices	3	4	Glass	Gems & Glass	5	1
Storax	Incense & Spices	1	2	Murrhine Glass	Gems & Glass	1	1
Cinnamon	Incense & Spices	2	3	Copper	Metals & Minerals	3	3
Malabathrum	Incense & Spices	4	8	Iron	Metals & Minerals	1	1
Pepper	Incense & Spices	3	4	Steel	Metals & Minerals	1	1
Spices	Incense & Spices	1	2	Tin	Metals & Minerals	1	2
Awls	Finished Products	5	1	Enslaved Persons	Humans	6	11
Daggers	Finished Products	5	1	Coin/Gold	Bullion	2	2
Hatchets	Finished Products	5	1	Gold Plate	Plateware	1	2
Images	Finished Products	1	2	Silver Plate	Plateware	1	2
Lances	Finished Products	5	1				
Sewn Boats	Finished Products	1	1				
Native Produce	Finished Products	1	12				

Note: This table includes all individual products listed in the text of Schoff (1912) with origins and destinations identified. X indicates the number of cities exporting the product; M indicates the number of cities importing it.

all trade with the same number of partners nor the same number of products. Many cities have only one or few trade partners. The majority of cities trade in multiple goods with only nine cities having fewer than ten products, and only three having a single product. Cities in the Horn of Africa, Arabia, Persia, and India are at the centre of the network with the most trade partners and trading in more products with other regions being largely peripheral to the system.

As indicated in Table 2, not every city in the Periplus trades with every other city mentioned. While there are no zero-trade flows explicitly recorded, these missing connections between cities could be inferred to be zeros in the trade data. This feature of the data is used in Section 4 to include zero-trade flows in the analysis.

2.2 Ancient Sea Routes

The distances between cities in the data are calculated using sailing times with respect to wind speed. These travel times are calculated based on wind speed information and calculations of sailing speeds relative to the wind using the results of sea-trials of a modern reconstruction of a historical trading vessel.

The data on wind speed and direction come from ERS-1 data on surface wind conditions from CERSAT/Laboratoire D'Océanographie Spatiale. These data provide ocean wind speed and direction information for 1 degree-by-1 degree squares of the Earth's surface. Monthly data are available over the period 1991 to 1996. The data used in this paper are a monthly average of wind speeds over the years 1992 to 1995 for the summer monsoon (June to July) and the winter monsoon (January to February). These months are used as they tend to be peak monsoon periods and correspond with the time of year voyages would depart various ports to sail with the winds (Cobb, 2018; Gurukkal, 2016). The vector data are visualised for the Indian Ocean in Figure 3 and nicely illustrate the unique feature of the two monsoon seasons, with consistent south-westerly winds in the summer season compared with steady north-easterly winds in the winter season—facilitating sailing across the ocean between India and Africa/Arabia on a seasonal basis.

The speed of a sailing ship relative to the wind is based on a number of important factors, but primary among these is the angle of the ship's direction to the wind. A method of calculating this relationship is necessary in order to measure the hypothetical speed of a ship travelling with the monsoon winds shown in Figure 3. There are no extant data of this detail on ancient ocean-going vessels, but there do exist modern reconstructions of similar ships with accompanying tests and data. Nomoto et al. (2000) documents the sea-trials of a reconstructed Japanese *higaki kaisen*-style merchant ship, a type of ship from the sixteenth/seventeenth century

Table 2. Cities of the Periplus

City	Region	Coordinates	Goods	Partners
Arsinoe	Egypt	29.97, 32.55	3	7
Berenike	Egypt	23.91, 35.48	22	3
Diospolis	Egypt	25.72, 32.61	3	7
Myos Hormos	Egypt	26.16, 34.24	22	3
Adulis	Horn of Africa	15.26, 39.66	32	3
Avalites	Horn of Africa	11.35, 43.47	22	6
Malao	Horn of Africa	10.44, 45.02	30	9
Mosyllum	Horn of Africa	11.17, 49.11	34	5
Mundus	Horn of Africa	10.9, 46.92	31	9
Tabae	Horn of Africa	11.82, 51.25	32	5
Opone	Horn of Africa	10.42, 51.27	30	7
Menuthias	East Africa	-6.13, 39.33	13	1
Nicon	East Africa	-1.22, 41.84	11	1
Pyralaae Islands	East Africa	-2.27, 40.9	11	1
Rhapta	East Africa	-7.85, 39.78	11	1
Sarapion	East Africa	2.03, 45.33	11	1
Cana	Arabia	14.03, 48.34	36	16
Dioscorida	Arabia	12.51, 53.92	11	10
Moscha	Arabia	17.04, 54.43	5	8
Muza	Arabia	13.32, 43.25	47	11
Ocelis	Arabia	12.77, 43.65	14	3
Sarapis	Arabia	20.47, 58.82	1	1
Apologus	Persia	30.89, 47.58	12	4
Ommana	Persia	24.34, 56.73	16	6
Oraea	Persia	25.21, 64.63	8	1
Ariaca	Scythia	23.06, 70.62	19	7
Barbaricum-Minnagara	Scythia	24.86, 67.01	22	1
Bacare	India	9.59, 76.49	34	6
Barygaza	India	21.71, 72.99	64	15
Camara	India	11.14, 79.86	13	5
Muziris	India	10.16, 76.21	34	7
Nelcynda	India	9.32, 76.54	34	6
Paethana	India	19.48, 75.38	1	1
Poduca	India	11.9, 79.82	13	5
Sopatma	India	12.53, 80.16	13	5
Tagara	India	18.32, 76.13	4	1
Tyndis	India	10.77, 75.9	34	6
Besatae	East Asia	27.33, 88.62	1	5

Note: This table only includes cities that have both goods and partners listed in the text of Schoff (1912). Coordinates are in (Lat, Lon) format; Goods indicates the number of unique products the city trades (both imports and exports); Partners indicates the number of other cities a given city trades with.

Figure 3. Monsoon Winds

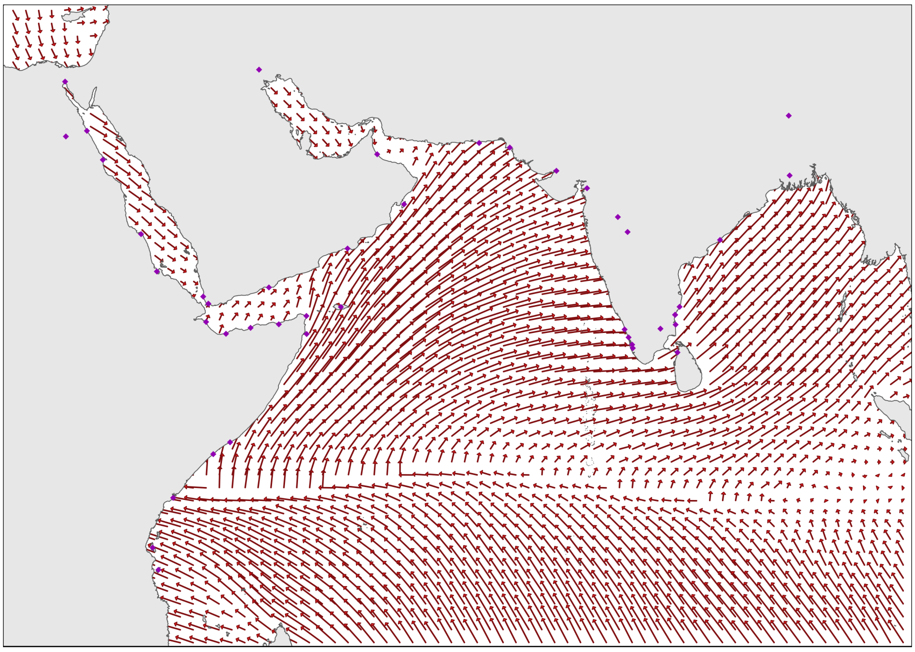

(a) Summer Monsoon

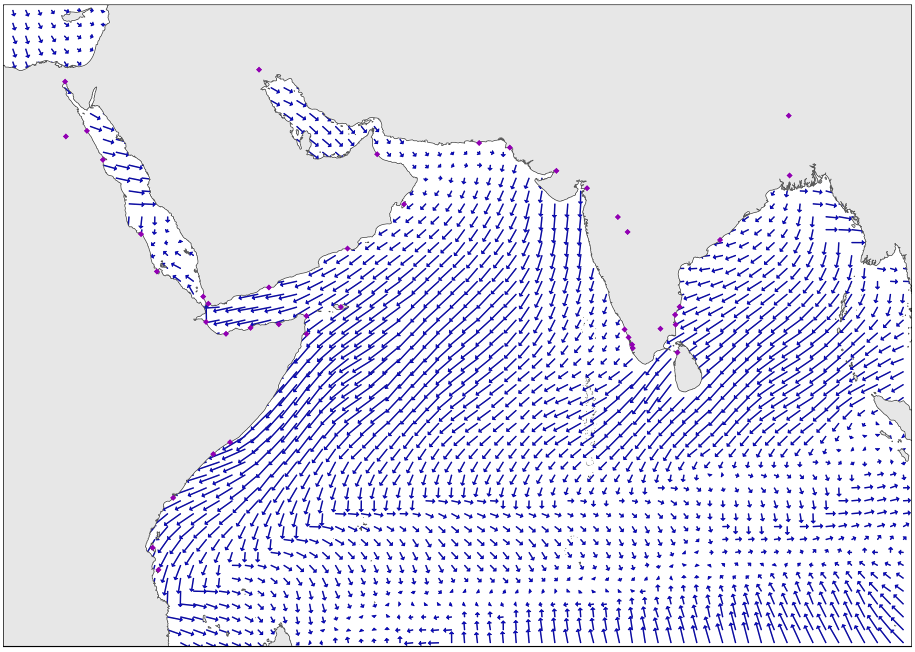

(b) Winter Monsoon

Note: Cities are indicated by points. Arrow direction indicates wind direction; arrow length corresponds to wind speed. (Source: Author's own calculations based on data from CERSAT/Laboratoire D'Océanographie Spatiale).

and very similar in design to ancient oceanic trading vessels. The data from these sea-trials record vessel speed for given wind speeds and directions, then plot them on a polar diagram of the ship. From this information, a functional form for the speed of the ship as a proportion of wind speed for an arbitrary wind direction can be approximated with:

$$s = 0.2 - 0.129 \cos \omega$$

Where s is the coefficient that gives the speed of the ship as a proportion of wind speed and ω is the angle of the wind relative to the direction of travel. The first term of this formula is the average transformation of wind speed to ship speed and the second term is the amplitude around the angle of the wind. This methodology means that a ship will sail four times faster with a tailwind as opposed to a headwind. This formula also implies that a ship will travel at its fastest with a direct tailwind; a feature that, while not fully consistent with the mechanics of sailing, simplifies the process of calculating ship speeds on a grid.

Using this formula to calculate the speed of a ship for any angle and any wind speed, a network of sailing pathways was overlaid on the wind vectors for the Indian Ocean. This created a web of pathways for a ship to travel through the 1 degree-by-1 degree grid squares of the wind vector map. The fastest routes were then calculated using GIS network analysis software to produce a series of travel times and distances along optimal sailing routes (according to wind speed and season) for all the city pairs in the Periplus. These calculated pathways are displayed in Figure 4. Both summer and winter routes were calculated in both directions. The value used for sailing distance between cities in the gravity analysis in this paper is the fastest time (measured in days) to make a return trip between the two cities in the correct sailing seasons.

Table 3. Travel Times Between Cities in the Periplus

Origin	Destination	Travel Time	
		Attested	Calculated
Berenike	Malao	1 Month	36.03 Days
Myos Hormos	Muza	1 Month	30.84 Days
Myos Hormos	Barygaza	>2 Months	64.70 Days
Tabae	Rhapta	<1 Month	17.99 Days

Note: Travel times are calculated according to monsoon wind data in the ideal sailing season for the direction of travel. Attested travel times come from Cobb (2018).

The resulting travel times are in line with attested travel times from the period and region. Table 3 presents selected routes and the calculated unidirectional travel times between them. The similarity of the calculated sailing times to those recorded by ancient sources indicates that the measure of sailing distance developed here is a strong approximation of actual sailing distances between cities in the region.

Figure 4. Sailing Routes Between Cities

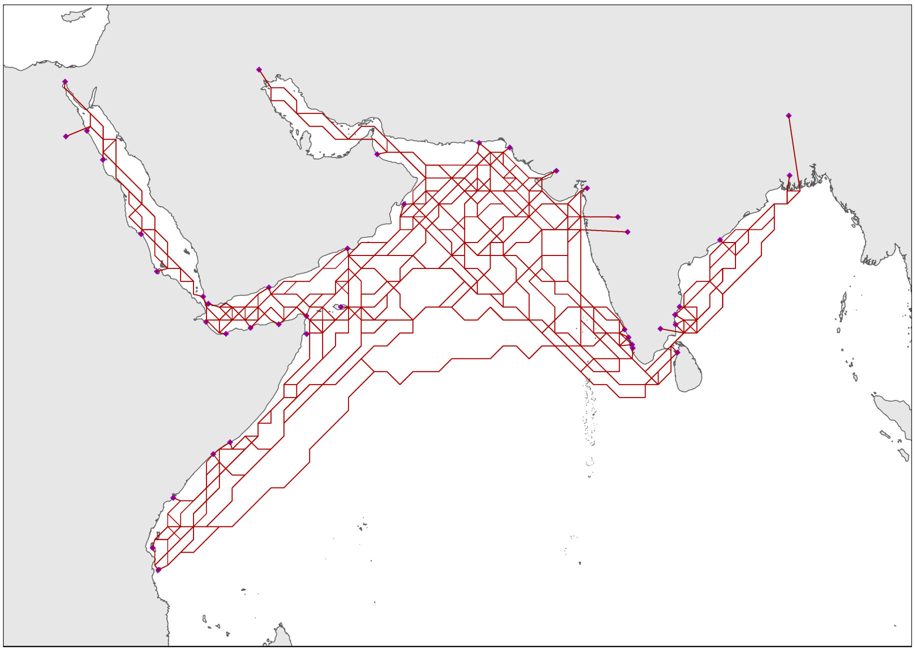

(a) Summer Routes

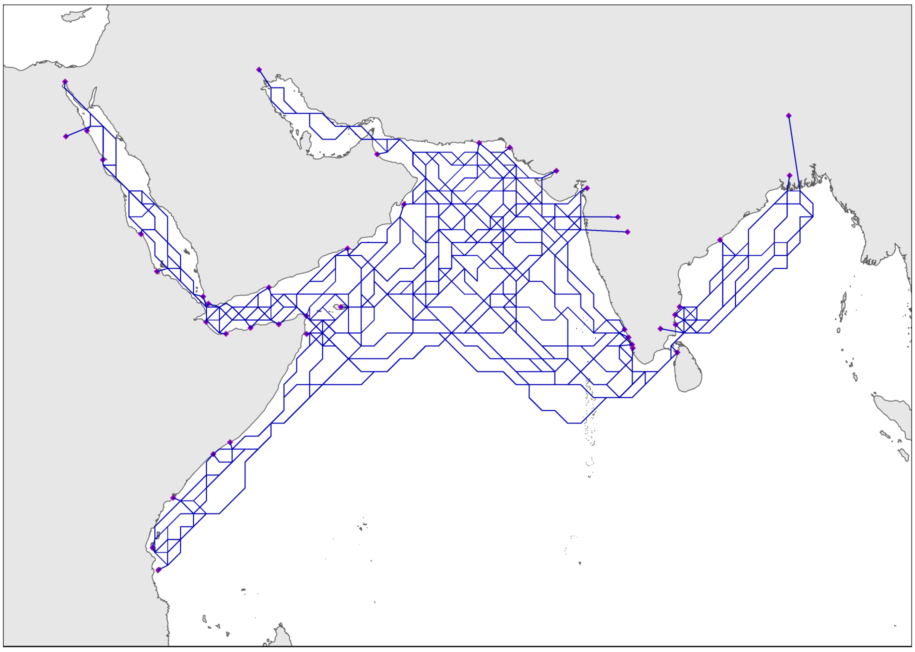

(b) Winter Routes

Note: Cities are indicated by points. Lines represent sailing routes between cities based on fastest sailing time. (Source: Author's own calculations).

2.3 Historical Economic Activity

In order to examine the effect of Indian Ocean trade on the development of the areas around the cities named in the Periplus, I overlay a map of the Indian Ocean with a grid of cells that measure 0.5 degrees longitude by 0.5 degrees latitude (equivalent to approximately 55 km by 55 km). Each grid cell is then determined to be either land, sea, or coast. Land cells are those consisting entirely of land while sea cells are those entirely of sea. Coast cells are determined to be any cell containing a coastline (including islands). The sample is limited to land and coast cells within 500 km of a city (Figure 5).[6] This can be considered to correspond to an effective limit to the 'hinterland' surrounding a city. The area around any given city in the ancient world is largely dependent on the prosperity of that city so that each city-hinterland area creates a core-periphery dynamic. This dynamic weakens as distance to the city increases reaching an upper-limit where the effective influence of the city on communities is minimal.[7]

Figure 5. Sample Grid Space

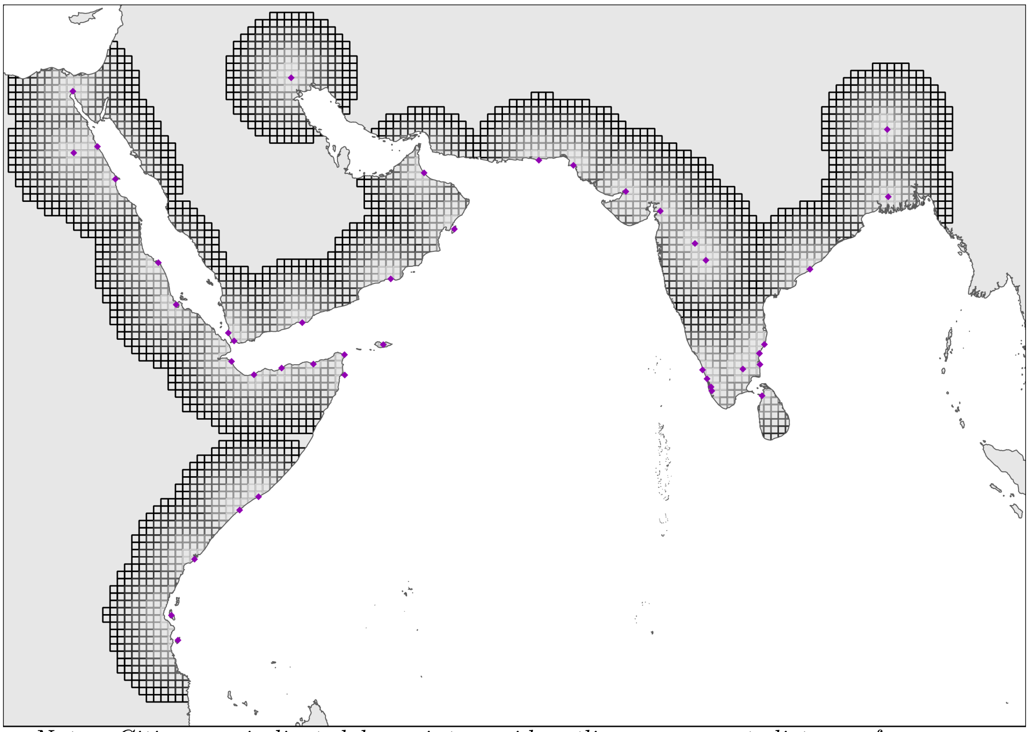

Note: Cities are indicated by points; grid outlines represent distance from nearest city with colour darkening with distance up to 500km. (Source: Author's own calculations).

[6] Effectively encompassing an area bounded by 500 km north of Arsinoe, 500 km south of Rhapta, 500 km west of Arsinoe, and 500 km east of Besatae.

[7] Different upper-limits of this hinterland are explored in Appendix A.

Ancient Growth The data for ancient economic activity come from archaeological sources. To proxy for ancient activity, a measure of the number of built-up structures in a given location is used. This is similar to a commonly used method for determining development in modern times: the use of night-time light levels (see Elvidge et al., 2012; Bluhm and Krause, 2019, for notable examples). The concentration of human-made structures gives an indication of the amount of activity at a location at a given time.

To construct this measure, data on archaeological sites are used. The Pleiades dataset (Bagnall et al., 2016) catalogues various 'places' throughout ancient history. These places can be anything from provinces, temples, roads, and beyond. Each item is geolocated and given a date range for activity. This dataset originally included primarily the Mediterranean area, but its expansion to include other areas of the world (particularly the Middle East and India) makes it well suited for the purposes of this paper. Figure 6 presents the location of the sites contained in the data for the region of the Periplus.

Figure 6. Ancient Archaeological Sites

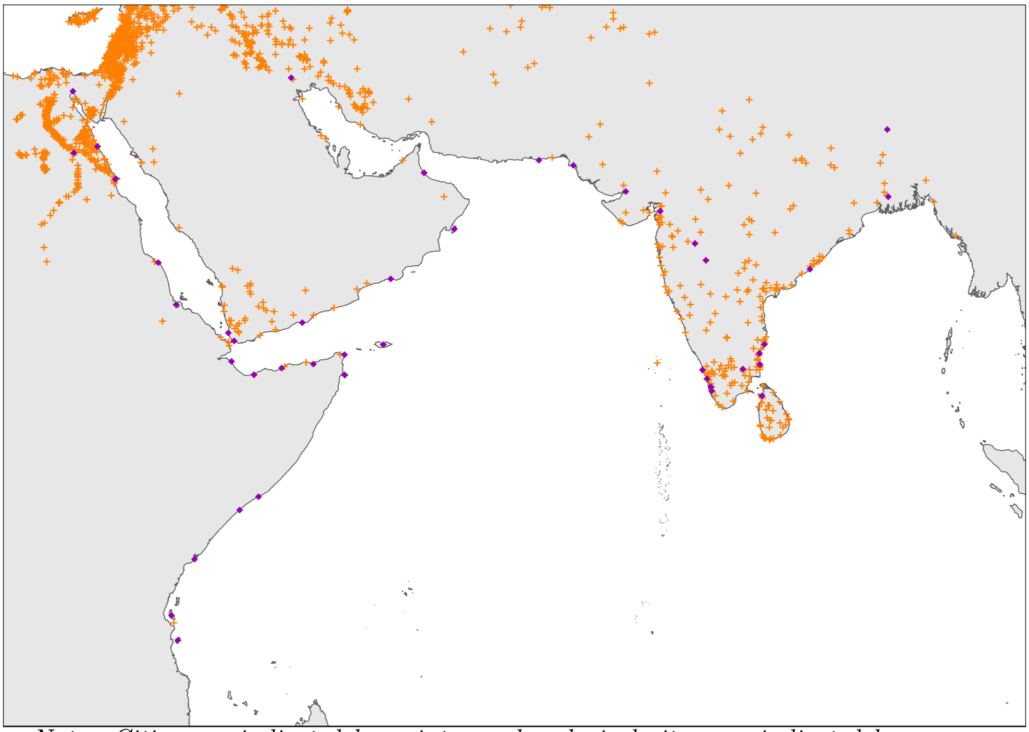

Note: Cities are indicated by points; archaeological sites are indicated by crosses. (Source: Author's own calculations based on Bagnall et al., 2016).

The data used to construct a variable to measure the growth in a given grid cell include all human-made 'places' within the grid cell. These sites are divided up into those that existed prior to 30 BCE (approximately 1632 sites) and those that emerged in the archaeological record after this date until 640 CE (approximately 1852 sites). The date of 30 BCE was chosen as it coincides with the Roman conquest of Egypt and the beginning of Roman access to, and expansion of links with, the Indian Ocean trade network (Cobb, 2018). The cut-off date of 640 CE was chosen as it coincides with the Arab invasion and conquest of Egypt and marks the point at which Roman access to the Indian Ocean was interrupted, signifying a potentially significant disruption to the trading system of the region (McLaughlin, 2018), despite the continued functioning of the system long after (Chaudhuri, 2005). The growth variable y is constructed from the difference between sites (σ^{post}) active after 30 BCE including those founded prior to this date and the number of sites (σ^{pre}) active in 30 BCE in each grid cell c:

$$y_c = \sigma_c^{post} - \sigma_c^{pre}$$

This produces a rough measure of the growth in activity in each grid cell relative to the period prior to the emergence of the trade system documented in the Periplus.

Table 4 shows the distribution of these sites across regions. There are two issues that emerge from the data on archaeological sites: one related to the summing of different sites and the other related to bias in the archaeological record. For the first issue, summing across different sites captures different levels of information. The variety of different 'places' in the data means that summing across all sites would combine cities, towns, structures, geographical features, and political units like provinces. To address this, the type of sites used to construct this variable are limited to human-built features. While this definition includes cities, towns, and structures together, it is the most direct solution to the difficulty introduced by the second issue: a lack of data for the non-Mediterranean world.

The second issue is related to selection bias within the data. Historically, archaeologists have been more interested in, and therefore more focused on, excavating and researching sites related to the Roman Empire or Egypt. This creates a bias within the data since the vast majority of archaeological sites are from around the Mediterranean or within the former Roman province of Egypt, with relatively few outside of these regions. While this analysis is not focused on the area around the Mediterranean, Egyptian cities are included in the data as major ports involved in trade in the Indian Ocean and the sample of grid cells for several cities extends into the Mediterranean littoral. This creates a problem where the number of sites in these areas is artificially large, as can be seen in Table 4 and Figure 6.

In addition to this bias in the data, the fact that Egypt was a constituent part of the Roman Empire at the time introduces a significant confounding variable for

Table 4. Distribution of Archaeological Sites

Region	Sites	Standard Deviation
Egypt	290.200	378.342
Horn of Africa	1.571	1.988
East Africa	0.200	0.447
Arabia	8.500	11.149
Persia	94.333	151.296
Scythia	7.500	3.536
India	13.667	9.076
All Regions	46.907	152.495

Note: Sites is the average number of sites within 500 km of cities in the given region.

the number of sites present: it is reasonable to assume that the number of sites in these grid cells could also be influenced by the presence of a somewhat centralised Roman administration. The other regions, though influenced by trade and economic exchange with the Roman Empire (McLaughlin, 2018), are more distant and separated from direct Roman influence meaning the number of sites is more related to local factors. In an effort to address the issue of bias in site density and the direct influence of Roman administration, I include a control for grid cells being within (either partially or wholly) the territorial limits of the Roman Empire in the analysis of ancient development in Section 5.

3. Gravity in a Ricardian World

Trade costs in the ancient world were notoriously high; with long-distance, trans-oceanic networks exceedingly rare. In this period, distance is frequently thought of as being such a significant barrier to trade that its impact is exponential—approaching prohibitive levels of cost very rapidly. This raises the question of whether a gravity model is suitable, since if distance has a cost-prohibitive upper-bound it cannot be modelled smoothly with a standard gravity equation. While this may be a potential obstacle in other locations, the unique features of the Indian Ocean, notably the seasonal monsoon,[8] mitigate this barrier and make long-distance trade feasible (Chaudhuri, 2005). The regularity of the monsoon and the orientation of landmasses

[8] There could be debate as to the declining impact of this theoretical upper-bound as technology for ocean-going ships developed and changed in the Roman Empire (McLaughlin, 2018). Regardless, the diffusion of this technology was slow and would not have impacted non-Roman cities at the same rate, leaving the climatic factors of the Indian Ocean system a predominant force.

around the Indian Ocean means that regular voyages can leave any given port at a certain time of year and arrive with favourable winds at a destination on an opposite landmass. This seasonal scheduling of voyages is attested to in the Periplus and other extant sources (see Schoff, 1912, for examples) and lends evidence to the theory that predictable monsoon winds allowed for long distance trade in the region independently of advanced nautical technology.

Despite these unique features of trade in the Indian Ocean, modern gravity models that are based on monopolistic competition and demand for variety are not always well specified to deal with pre-modern trade—particularly ancient trade. In this period, production and trade of goods was largely based on technology.[9] In this context, a Ricardian world of comparative advantage predominates where specialisation is promoted through declining transport costs which facilitate trade.

In addition to a Ricardian model being more theoretically sound for ancient trade flows, the mechanics of the model also fit more precisely with the data. Due to the nature of the Periplus data described in Section 2.1 the trade variable of observation is the number of good types shipped between cities. This extensive-margin measure makes models based on the value or quantity of goods traded unworkable. With a Ricardian gravity model, the driving force of comparative advantage means that more competitive[10] cities will trade more types of products (not just more goods) and this fact can be used to fit the data to the model.

I follow Eaton and Kortum (2002) in constructing a simplified version of their gravity model of Ricardian trade.[11] In this model cities have differential access to technology. City j's technological efficiency of producing good g is given by the parameter $z_j(g)$ which is the realisation of random variable Z_j, drawn independently for each g, from a city-specific probability distribution $F_j(z) = Pr[Z_j \leq z]$. City j's efficiency distribution is assumed to be Fréchet (Type II extreme value distribution):

$$F_j(z) = e^{-T_j z^{-\theta}} \tag{3.1}$$

where $T_j > 0$ is a city specific parameter and $\theta > 1$ is a global parameter capturing the variation within the distribution. A larger T_j means a higher-efficiency draw for any given good is more likely and a larger θ indicates less variability.

[9] Though endowments can also play a part, the ability of Eurasian and East African communities to exploit those endowments was determined by the spread of technology such as iron working. Section 4 explores the inclusion of endowment effects to supplement the model.

[10] In this context, competitiveness is related to both the size of the city (labour) and the distance from the destination market (transport cost).

[11] See Eaton and Kortum (2002) for a detailed derivation of the model. A similar model is also derived by Barjamovic et al. (2019).

From this, the cost (and under perfect competition, the price in the destination) of delivering a unit of good g produced in city j to city i is

$$p_{ij}(g) = \left(\frac{c_j}{z_j(g)}\right) d_{ij} \tag{3.2}$$

where c_j is the cost of all inputs, including labour, and d_{ij} is an iceberg representation of transport costs.[12] Consumers in destination i are not limited to goods from origin j and will choose goods from destinations that minimise the prices they face. Therefore, for any given good g its price in destination i will be the lowest across all N sources:

$$p_i(g) = \min\{p_{ik}(g); k = 1, ..., N\} \tag{3.3}$$

Substituting (3.2) into the efficiency distribution (3.1) gives a distribution of prices presented to destination i by origin j:

$$G_{ij}(p) = Pr[P_{ij} \leq p] = 1 - F_j\left(\frac{c_j d_{ij}}{p}\right) = 1 - e^{-[T_j(c_j d_{ij})^{-\theta}]p^{\theta}} \tag{3.4}$$

Because the lowest price for a good in city i will be less than p unless each source j's price is greater than p, the price distribution for what city i actually buys is:

$$G_i(p) = Pr[P_i \leq p] = 1 - \prod_{j=1}^{N}[1 - G_{ij}(p)]$$

inserting the price distribution (3.4) collapses the function into something resembling $G_{ij}(p)$:

$$G_i(p) = 1 - e^{-\Phi_i p^{\theta}} \tag{3.5}$$

where

$$\Phi_i = \sum_{j=1}^{N} T_j(c_j d_{ij})^{-\theta}$$

It can then be shown that the probability that city j provides a good at the lowest price in destination i is

$$\pi_{ij} = \frac{T_j(c_j d_{ij})^{-\theta}}{\Phi_i} \tag{3.6}$$

which is equivalent to the contribution of city j to city i's price parameter. Helpfully, conditioning on the source has no impact on a good's price since a source with a

[12] In order for 1 unit of good g to arrive in destination i, d_{ij} units must be shipped from origin j. $1 - \frac{1}{d_{ij}}$ 'melts' along the way.

competitive advantage exploits this by selling a wider range of goods up to the point that what it sells in a given destination i is the same as the destination's overall price distribution.[13] This means that city i's average expenditure on each good also does not vary by source, and so the proportion of goods that city i buys from city j is equivalent to the proportion of its expenditure on goods from city j:

$$\pi_{ij} = \frac{X_{ij}}{X_i} = \frac{T_j(c_j d_{ij})^{-\theta}}{\Phi_i} = \frac{T_j(c_j d_{ij})^{-\theta}}{\sum_{j=1}^{N} T_j(c_j d_{ij})^{-\theta}} \tag{3.7}$$

Departing from Eaton and Kortum (2002), this expression can then be rearranged into a equation resembling a basic gravity function in Anderson and van Wincoop (2003):

$$X_{ij} = \frac{T_j(c_j d_{ij})^{-\theta} X_i}{\sum_{j=1}^{N} T_j(c_j d_{ij})^{-\theta}} \tag{3.8}$$

using market clearing conditions, this can be turned into a standard gravity equation in the normal way:

$$X_{ij} = \frac{X_i X_j}{X_w} \left(\frac{d_{ij}}{\Phi_j \Phi_i} \right) \tag{3.9}$$

where $X_w = \sum_k X_k$ is total output and Φ_j and Φ_i serve as multi-lateral resistance terms.

Equation (3.9) is therefore a Ricardian structural gravity function and can be estimated in a similar fashion to a traditional gravity function. Following Anderson and van Wincoop (2003) and Silva and Tenreyro (2006), this model can be estimated with a pseudo-Poisson Maximum Likelihood estimator and fixed effects. With this technique, the gravity function is estimated in exponential form, allowing the number of good types to be measured in levels. The estimation equation is given by:

$$E[X_{ij}] = exp\{\beta_0 + \beta_1 ln(d_{ij}) + \mu_i + \mu_j + \epsilon_{ij}\} \tag{3.10}$$

where μ_i and μ_j represent city-specific fixed effects and standard errors are clustered at the city-pair level.

The Ricardian foundation of this model means that at its core is the concept of comparative advantage. This feature, incorporated into the price distributions, means that a city j with a higher comparative advantage (either through technology, input prices, or a shorter transport distance) will sell a wider range of goods (more product types) to city i. As the model illustrates, the proportion of goods sold to city i by city j is equivalent to the proportion of expenditure of city i on goods from

[13] See Eaton and Kortum (2002) for a detailed derivation of these properties.

city j; this is equally proportional to the range of products sold by city j to city i. This unique feature of the Ricardian model is what can be exploited in order to measure the trade between cities in terms of the number of products traded between them, rather than the value of those products as in traditional gravity models. This allows the X_{ij} term of Equation 3.10 to be given as the number of products traded between cities i and j, linking it to the analogous term from Section 2.1.

4. Ancient Gravity Estimation

4.1 Sailing Distance and Trade

The results of the gravity estimations using Equation 3.10 are presented in Table 5. The first column of the table presents the result of the estimation using only data identified in the Periplus and does not include any zero flows. The coefficient on SailDist is on the lower end of expected measures of the distance elasticity, but this is not uncommon under PPML estimation (Silva and Tenreyro, 2006) and could be related to the fact that the elasticity concerns the number of product types instead of explicit trade volume as in traditional gravity estimations. It is negative and significant, indicating that distance has a suppressing effect on trade—in a Ricardian world, distance reduces a city's comparative advantage in markets. In this estimation, a 1% increase in sailing time between two cities reduces trade between them by approximately 0.27%.

Table 5. Gravity Estimation: Distance Measures

	(1)	(2)	(3)
ln(SailDist)	-0.273**	0.334	-8.534**
	(0.114)	(0.241)	(3.156)
sqrt(ln(SailDist))			30.029**
			(10.563)
City FE	Yes	Yes	Yes
N	107	556	556
R^2	0.856	0.225	0.257

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is the total number of good varieties traded; ln(SailDist) is the log of the total round-trip sailing time for the fastest calculated route accounting for wind speed.

The second and third columns of Table 5 present results including all city pairs not explicitly mentioned in the Periplus as zero-trade observations. The coefficient of the standard linear gravity model in column two is of the wrong sign and insignificant, indicating that when zero-trade observations are inserted into the data the linear relationship between distance and trade flows disappears.

Despite this, there remains a non-linear relationship between distance and trade as presented in column three. In this result both coefficients are significant with the linear term being strongly negative and the quadratic term strongly positive creating a relationship that takes the shape of an inverted-U (as illustrated in Figure 7). These results imply that trade is increasing with sailing time over short distances until a critical point is reached and the relationship between distance and trade revert to the expected negative relationship.

Together these results indicate that sailing time seems to be an important factor in determining trade flows. Importantly, the non-linear structure of the relationship with the inclusion of zero-trade flows suggests that the importance of sailing time increases with distance between cities. Trade between cities that are close together is less influenced by sailing time than more long distance relationships, highlighting the importance of the monsoon weather system in facilitating this longer-distance trade. The seasonal and strong winds of the monsoon created sailing efficiencies that played an important part in long-distance trade between cities. The potential drivers of non-linearity in the distance measure are explored further in the following sub-section.

4.2 Endowments and Non-linearity

Including city-pairs not explicitly mentioned in the Periplus as observations with zero trade removes the linearity of the relationship between distance and trade flows due to the number of cities that are close together but do not trade. This leaves a non-linear relationship between distance and trade as illustrated in Figure 7. Importantly, this effect is skewed towards cities that are close together, leading to two possible explanations for this effect. The first is that there could be a bias in the data, with trade between cities that are close together being systematically excluded from the Periplus because of its focus on longer-distance trade. The second is that cities with similar endowments are likely to trade less, and cities that are close together geographically are more likely to have similar endowments.

While the first explanation is plausible, it requires the vast majority of trade to be ignored regardless of the scale of such trade. It is more likely that trade would have been ignored between cities that are close together because trade between such cities was already quite low in value.

This seems to imply that a pure Ricardian model based on technology-derived

Figure 7. Non-linear Trade-Distance Estimates

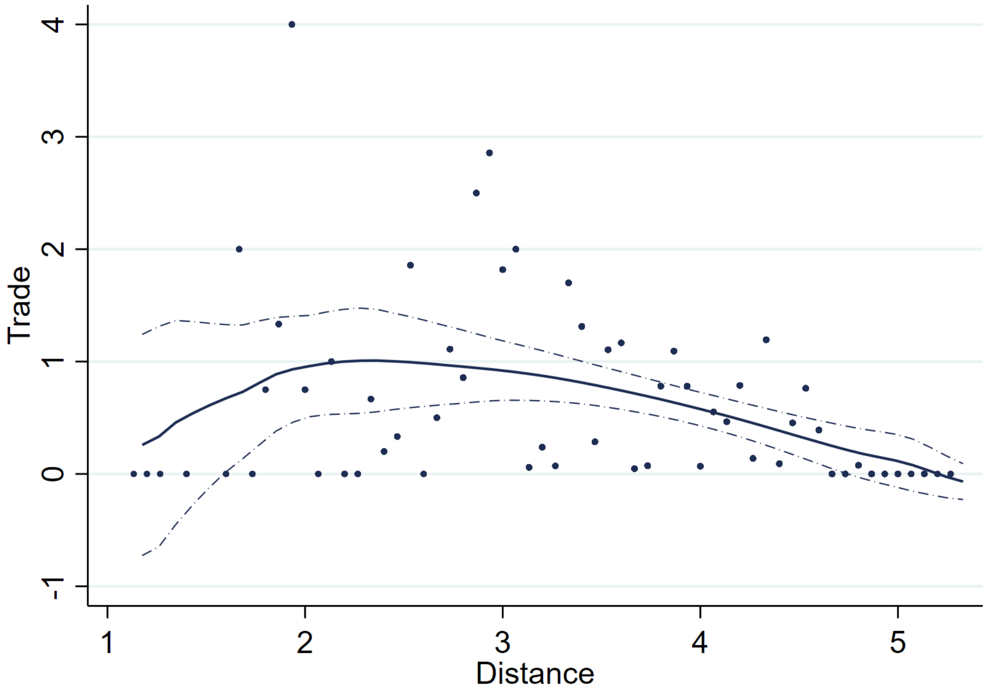

Note: Trade is given by the total number of good varieties traded; distance is the log total round-trip sailing time for the fastest calculated route accounting for wind speed. Missing city pairs are included as zero-trade observations.

comparative advantage is insufficient to fully explain the trade relationships described by the data. I therefore relax the technology-derived comparative advantage assumption of the model discussed in Section 3 and incorporate factor endowments, in line with a Heckscher-Ohlin model. This allows for comparative advantage to also be derived from factor endowments.[14] Such a model is more likely to fit the data and more accurately describe trade flows.

To test this, I create two measures of proximity between cities. The first is a dummy variable indicating whether the two cities are within 14-days' sail of each other (a short return trip that could be covered easily more than once within one sailing season):

$$Neighbour_{ij} = \begin{cases} 1 & \text{if } d_{ij} \leq 14 \text{ days;} \\ 0 & \text{otherwise.} \end{cases}$$

Where d_{ij} is the sailing distance in days between city i and city j.

The second measure is an endowment index controlling for similarity in export goods offered by the cities. It is constructed as the sum across all goods categories

[14] Ruffin (1988) and Ruffin (1990) provide more discussion on supplementing Ricardian models with factor endowments, creating a synthesis of Ricardian and Heckscher-Ohlin trade theory.

of the difference in export baskets:

$$EI_{ij} = 1 - \sum_{n} \frac{|\lambda_{in} - \lambda_{jn}|}{L_i + L_j}$$

Where λ_{in} is the number of products within the product category n (as given in Table 1) exported by city i and L_i is the total number of products exported by city i. Higher values of the index indicate higher levels of endowment similarity.

Table 6. Gravity Estimation: Non-linear Distance

	(1)	(2)
ln(SailDist)	-1.151**	-1.098**
	(0.465)	(0.473)
Neighbour	-2.858**	-1.912**
	(0.734)	(0.797)
Endowment Index		-10.316**
		(1.562)
City FE	Yes	Yes
N	556	556
R^2	0.293	0.430

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is the total number of good varieties traded; ln(SailDist) is the log of the total round-trip sailing time for the fastest calculated route accounting for wind speed; Neighbour is a dummy variable for cities in close proximity; Endowment Index is an index of similarity of endowments.

The results for the estimations that include these two measures are presented in Table 6. The first column includes just the first proximity measure with a negative and significant coefficient, indicating that there is a significant negative effect on trade when cities are close together. The second column supplements this with the endowment index. The coefficient on the index is strongly negative and significant illustrating the highly negative impact of similar endowments on trade between cities, a 1% increase in the similarity of endowments reduces the number of good varieties traded by just under 0.11%. Crucially, the coefficient on the distance measure in both estimations is much more in line with expected distance elasticities in gravity models. This seems to indicate that the lower magnitude of this coefficient in the linear analysis may have been due to this endowment effect.

These results supplement the evidence for a Ricardian gravity system from the linear specification. Where the Ricardian gravity model derived for this analysis focuses on technology and distance as sources of comparative advantage, this inclusion of endowment effects potentially adds another dimension to this: resource endowments. Increasing distance increases trade costs, therefore reducing trade, but cities are also sensitive to endowment effects in that cities with similar product offer baskets are less likely to trade with each other.

5. Trade and Ancient Development

The results of the previous section provide evidence for the Indian Ocean as a highly functioning trade network and the data used as a good representation of that activity. The impact of this trade network on the development of cities occupying the Indian Ocean littoral is the subject of this section.

5.1 Export-led Growth in the Iron Age

I use the data on exports and imports for individual cities (not city-pairs) and the data on archaeological sites from the Pleiades dataset as a proxy for ancient economic activity to test whether the trade balance of a city is related to the level of development in the surrounding area. A frequent theme among ancient writers when discussing trade between Rome and the Indian Ocean was the threat of import dependence and the adverse economic effects of it (Cobb, 2018; McLaughlin, 2018). Generally these discussions focused on the effect of dependence on imports from the Indian Ocean on Rome and its wealth,[15] but implicit in this is the corollary that the cities of the Indian Ocean were growing and benefiting from their exports. Throughout this analysis, controls are included for areas being located within the borders of the Roman Empire for the reasons covered in Section 2.3.

For this subsection, the change in development of a grid cell (y_c) is constructed as the standardised value of the variable discussed in Section 2.3. The relationship is modelled through two equations estimated with OLS. The first estimates the impact of proximity alone:

$$y_c = \beta_1 ln(d_c) + \nu_c + \epsilon_c \tag{5.11}$$

where d_c is the Cartesian distance between the centroid of grid cell c and the nearest identified port in the Periplus. Fixed effects for cell latitude and longitude are given by ν_c and ϵ_c is a normally distributed error term.

[15] Pliny complained at length about the large outflow of money to the cities and countries of the Indian Ocean and how they were becoming rich at the expense of Rome.

The second equation estimates a relationship between the growth of sites in a grid cell and distance weighted measures of the export and import diversity variables outlined in Section 2.1:

$$y_c = \beta_1 \omega_c h_c^x + \beta_2 \omega_c h_c^m + \nu_c + \epsilon_c \tag{5.12}$$

where h_c^x and h_c^m are the export and import diversity measures for the nearest identified port in the Periplus to grid cell c and ω_c is the distance weight given by the inverse of the distance from the cell centroid to the nearest identified port in the Periplus ($\frac{1}{d_c}$). Again, ν_c are cell latitude and longitude fixed effects and errors ϵ_c are normally distributed.

The results for these estimations are presented in the first two columns of Table 7. In the first column, the results indicate a negative and significant relationship between distance and the change in development in a grid cell demonstrating that the further a grid cell is from a port identified in the Periplus, the smaller the change in sites within that cell. While this is somewhat intuitive since ports were often associated with larger urban areas, this gives justification for the weighting of the effects of export and import diversity by distance in order to account for this effect.

The second column displays the results for the export and import diversity variables. The results for both exports and imports are strong and significant, linking a 1% increase in the range of export goods with an increase in development of over 0.3 standard deviations. Conversely, a similar 1% increase in the variety of import goods reduces the development of a grid cell by 0.15 standard deviations.

These results provide evidence supporting the idea that export diversity is associated with higher levels of site growth within grid cells around ancient ports. The distance weighting of the variables means that areas closer to ports will experience these effects much more strongly than those further away, which is consistent with the distance-related effects identified in the first result. Trade in the early first century CE, specifically the diversity of export goods and the dependence on a variety of import goods, seems to have influenced the change in the number of sites present in areas around ports during the following 500 years. This appears to indicate an effect similar to that of modern export-led growth where cities fuel their development and urbanisation through positive trade balances. Despite this, it is possible that this could be linked to a more Ricardian production effect in that more development (a larger population) could be driving increased diversification of products due to an increased comparative advantage. While the results for import diversity temper this possibility somewhat, it cannot be wholly ruled out without more detailed development data.

Table 7. Trade and Ancient Development

	(1)	(2)	(3)
ln(Dist)	-0.094*		
	(0.054)		
Export Diversity		31.273**	
		(11.199)	
Import Diversity		-15.278**	
		(6.691)	
Manufactures			3.115**
			(1.462)
Spices & Aromatics			-22.187*
			(13.262)
Bullion			638.741**
			(315.511)
Rome	0.388**	0.467**	0.465**
	(0.192)	(0.187)	(0.187)
Cell FE	Yes	Yes	Yes
N	3172	3172	3172
R^2	0.065	0.069	0.067

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is the standardised change in ancient place records within a grid cell; ln(Dist) is the distance from the nearest Periplus port; Export and Import Diversity are distance-weighted Herfindahl-style indices of export/import product concentration; Manufactures and Spices & Aromatics are distance-weighted measures of export shares for the given sector; Bullion is the distance-weighted measure of import share of bullion products; Rome is a dummy for the grid cell being a part of the Roman Empire.

5.2 Ancient Resource Curse

The results for export-led effects on development in the previous subsection can be expanded to investigate the specific role of different sectors on the development of areas around cities in the Periplus. In order to accommodate this, a set of sector-specific export and import variables are constructed for selected sectors: manufactures, spices and aromatics, and bullion. These sectors were selected as they encompass three distinct elements of city trade diversity. The first is the reliance on exported manufactured goods which would likely have a large value-added component and would require the input of potentially diverse skilled labour; this could fuel wealth creation, population agglomeration around the port, and consequently growth. The second is the role of high-value, low-skill cash crops in the export offer of a city; spices and aromatics were generally grown on plantations that only required minimal supervision during cultivation and low-skilled labour for harvesting, generating high returns without the economic spillovers implied by manufactures and skilled crafts. Third is the claim by Pliny that the export of gold and silver by Rome, and consequent import of bullion by cities around the Indian Ocean, was making those cities rich.

This set of sector-specific export and import variables are calculated as simple sectoral shares:

$$S(r)_i = \frac{\lambda_{ir}}{L_i}$$

The share $S(r)_i$ of sector r for city i is the number of products λ_{ir} in sector r exported or imported by city i divided by the total number of products exported or imported by city i, L_i. Table 8 gives the composition of each of the sectors used in these variables, including the individual products and product categories from the list in Table 1.

These variables feed into the following estimation equation for the level of growth in sites for a given grid cell (y_c):

$$y_c = \beta_1 \omega_c S(mfg)^x_c + \beta_2 \omega_c S(spc)^x_c + \beta_3 \omega_c S(bul)^m_c + \nu_c + \epsilon_c \qquad (5.13)$$

Where $S(mfg)^x_c$, $S(spc)^x_c$, and $S(bul)^m_c$ are the shares of exports devoted to manufactures, spices and aromatics, and the share of imports composed of bullion respectively; ϵ_i is an error term. The same distance weights ω_c from Equation 5.12 are used to account for similar distance-related effects. Grid cell latitude and longitude fixed effects, ν_c, are included to control for endowment effects of grid cells.

The results of this estimation are presented in the third column of Table 7. The coefficients on the sectoral share variables are all significant but with varying levels of magnitude: the manufacturing share is moderately positively related to development, spices and aromatics are negatively associated with growth, and bullion has a large positive relationship with the change in the number of sites in a grid cell. Specifically, while a 1% increase in the share of exports devoted to manufactures is related to a 0.03 standard deviation increase in development, a similar increase in the share of spices and aromatics decreases development by over 0.2 standard

Table 8. Sectoral Classification of Products

Manufactures:			
Product	Product Category	Product	Product Category
Palm Oil	Plant Products	Robes	Textiles
Sesame Oil	Plant Products	Skin Coats	Textiles
Coloured Lac	Dye	Thin Clothing	Textiles
Purple	Dye	Tunics	Textiles
Wine	Wine	Indian Cloth	Textiles
Awls	Finished Products	Cloth	Textiles
Daggers	Finished Products	Cotton Cloth	Textiles
Hatchets	Finished Products	Mallow Cloth	Textiles
Images	Finished Products	Monache Cloth	Textiles
Lances	Finished Products	Muslins	Textiles
Sewn Boats	Finished Products	Purple Cloth	Textiles
Native Produce	Finished Products	Raw Silk	Textiles
Arabian Clothing	Textiles	Sagmatogene Cloth	Textiles
Blankets	Textiles	Silk Cloth	Textiles
Cloaks	Textiles	Silk Yarn	Textiles
Coloured Sashes	Textiles	Flint Glass	Gems & Glass
Girdles	Textiles	Glass	Gems & Glass
Local Clothing	Textiles	Murrhine Glass	Gems & Glass
Spices & Aromatics:			
Product	Product Category	Product	Product Category
Indian Copal	Plant Products	Spikenard	Incense & Spices
Aloes	Plant Products	Storax	Incense & Spices
Macir	Plant Products	Cinnamon	Incense & Spices
Duaca	Incense & Spices	Malabathrum	Incense & Spices
Fragrant Ointments	Incense & Spices	Pepper	Incense & Spices
Frankincense	Incense & Spices	Spices	Incense & Spices
Myrrh	Incense & Spices		
Bullion:			
Product	Product Category		
Coin/Gold	Bullion		

Note: This table groups individual products from Table 1 into the sectoral categories that are included in the analysis in Section 5.

deviations. A 1% increase in bullion imports, by contrast, increases growth by over six standard deviations, reflecting an extreme sensitivity to specie acquisition. All of these results are robust to adjustments in both the cut-off year for inclusion of sites in the data and the extent of the sample area around the cities.[16]

This seems to imply a divergence in the results for manufacturing (positive with development) and more plantation/extractive production (negative with development). As shown in Table 8, the spices and aromatics sector contain more cash-crop, high yield products and these seem to relate to under-developed areas compared to those that rely more on manufacturing. This effect is analogous to the modern effect of a 'resource curse' where an over-reliance on exploiting natural resources (in this case spices, incense, and other plantation-based aromatics and resins) leads to a cycle of under-development.

The outlier in terms of magnitude is the impact of bullion. This result seems to confirm Pliny's claim that the cities of the Indian Ocean were gaining significantly from the tendency of Roman merchants to pay for products with coin and precious metals. The mechanism at play in this is difficult to determine with the data available, but is likely one of two possibilities (or a combination of both). First, it is possible that the tendency to pay for luxury and high-value goods with specie concentrated wealth in the hands of city elites.[17] These elites could then invest in infrastructure, public works, and other projects that expanded the potential output of their cities and facilitated population growth. The large elasticity associated with the bullion share would indicate a potentially implausibly large propensity to engage in this type of behaviour.

Second, it is possible that the economies of the Indian Ocean were subject to significant liquidity constraints and the influx of Roman specie served to provide monetary stimulus to the economies importing them, supporting larger scale economic activity and growth. This idea is supported by the frequency with which Roman coins and blanks are found re-stamped as local currencies (Cobb, 2018). The significant difference in precious metal content between Roman currency (which was of a higher purity) and local currency also lends support to this theory (Cobb, 2018). While the analysis presented here cannot provide a definitive answer, it does fit with this view of the effect of Roman specie in the economies Indian Ocean.

These results seem to indicate that export sector shares in trade circa 50 CE to 100 CE had significant influence on the change in density of archaeological sites during the subsequent 500 years to 640 CE. While causal statements in this regard

[16] Variations in these values and the related tables are discussed in Appendix A.

[17] High-value products generally had a royal monopoly or were concentrated in the hands of a few powerful aristocrats and these were often paid for with specie (McLaughlin, 2018).

are difficult due to the peculiarities of the data, this evidence is suggestive of features of the Indian Ocean trade network that are analogous to an export-led development effect and an effect similar to the modern 'resource curse'.

6. Conclusion

This paper uses extant textual sources to model and estimate ancient trade flows in the Indian Ocean. The detail of the data allows for the construction of a comprehensive picture of economic interaction, fuelled by climatic conditions, between communities around the Indian Ocean and the way those communities were potentially shaped by that interaction. Two important results are derived from the analysis of these connections.

The first result is that the unique climatic conditions of the Indian Ocean (namely the presence of the seasonal monsoon) facilitated long-distance trade between cities along its shores. This trade is well fitted by a gravity model using sailing times between destinations based on seasonal winds. The relationship between distance and trade flows is decreasing but also non-linear, implying endowment effects consistent with comparative advantage being derived from both technology and resource endowments.

Second are the results for the relationship between trade and development in the region during the late Iron Age. The data and analysis presented in this paper are suggestive of an export-driven development process with more evidence of human activity linked with increased access to export markets and specifically the influence of particular product sectors. This leads to the potential emergence of an ancient version of the 'resource curse' with cities that were more reliant on exporting the ancient equivalent of cash-crops (such as spices, incense, and resins) being more under-developed than cities that were more reliant on manufactures.

Despite these insights, the evidence (particularly for the results related to development levels) remains suggestive. This opens up new avenues for inquiry into the potential for trade in the ancient world to influence development. Particularly inviting further investigation of the potential for and impact of an ancient 'resource curse'.

Ultimately, the insight provided by this paper into the structure of Indian Ocean trade during the Iron Age helps to illuminate the driving forces behind economic connections in the ancient world. The results presented here are consistent with modern trade models and demonstrate the importance of long distance trade networks in the Indian Ocean during the late Classical period. Merchants and economies faced similar challenges (albeit on a different scale) to those faced by current firms and countries as the mechanisms driving the economic relationships then are analogous to those we currently face.

7. References

Algaze, Guillermo, *Ancient Mesopotamia at the Dawn of Civilisation*, Chicago: Chicago University Press, 2008.

Anderson, James E. and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *The American Economic Review*, March 2003, *93* (1), 170–192.

Auty, Richard M., *Sustaining Development in Mineral Economies*, London: Routledge, 1993.

_ , *Resource Abundance and Economic Development*, Oxford: Oxford University Press, 2001.

Bagnall, Roger, Richard Talbert, Tom Elliot, Lindsay Holman, Jeffrey Becker, Sarah Bond, Sean Gillies, Ryan Horne, Gabe Moss, Adam Rabinowitz, Elizabeth Robinson, and Brian Turner, "Pleiades: A Gazetteer of Past Places," http://pleiades.stoa.org 2016.

Bakker, Jan David, Stephan Maurer, Jorn-Steffen Pischke, and Ferdinand Rauch, "Of Mice and Merchants: Trade and Growth in the Iron Age," Working Paper 24825, National Bureau of Economic Research July 2018.

Barjamovic, Gojko, Thomas Chaney, Kerem Cosar, and Ali Hortacsu, "Trade, Merchants, and the Lost Cities of the Bronze Age," *Quarterly Journal of Economics*, Aug 2019, *134* (3), 1455–1503.

Bluhm, Richard and Melanie Krause, "Top Lights - Bright Cities and Their Contribution to Economic Development," Working Paper 7411, Social Science Research Network February 2019.

Broodbank, Cyprian, "The Origins and Early Development of Mediterranean Maritime Activity," *Journal of Mediterranean Archaeology*, 2006, *19* (2), 199–230.

Chaudhuri, K N, *Trade and Civilisation in the Indian Ocean: An Economic History from the Rise of Islam to 1750*, Cambridge: Cambridge University Press, 2005.

Cobb, Matthew Adam, *Rome and the Indian Ocean Trade from Augustus to the Early Third Century CE*, Leiden: Brill, 2018.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, Sep 2002, *70* (5), 1741–1779.

Elvidge, C D, K E Baugh, S J Anderson, P C Sutton, and T Ghosh, "The Night Light Development Index (NLDI): a spatially explicit measure of human development from satellite data," *Social Geography*, 2012, *7* (1), 23–35.

Feyrer, James, "Distance, Trade, and Income - The 1967 to 1975 Closing of the Suez Canal as a Natural Experiement," Working Paper 15557, NBER 2009.

_ , "Trade and Income - Exploiting Time Series in Geography," *American Economic Journal: Applied Economics*, Oct 2019, *11* (4), 1–35.

Fluckiger, Matthias, Erik Hornung, Mario Larch, Markus Lugwig, and Allard Mees, "Roman Transport Network Connectivity and Economic Integration," Working Paper 7740, Centre for Economic Studies July 2019.

Gurukkal, Rajan, *Rethinking Classical Indo-Roman Trade: Political Economy of Eastern Mediterranean Exchange Relations*, Oxford: Oxford University Press, 2016.

Juhasz, Reka, "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review*, 2018, *108* (11), 3339–3376.

Liu, Dan and Christopher M Meissner, "Market potential and the rise of US productivity leadership," *Journal of International Economics*, 2015, *96* (1), 72–87.

Maurer, Stephan E and Ferdinand Rauch, "Economic Geography Aspects of the Panama Canal," Working Paper 1633, Centre for Economic Performance July 2019.

McLaughlin, Raoul, *The Roman Empire and the Indian Ocean*, Yorkshire: Pen and Sword Books Ltd, 2018.

Nomoto, Kensaku, Yutaka Masuyama, and Akira Sakurai, "Sailing Performance of 'Naniwa-maru', Reconstruction of a Japanese Sailing Trader in Early 1800s," *Journal of the Kansai Society of Naval Architects, Japan*, 2000, (234), 115–124.

Pascali, Luigi, "The Wind of Change: Maritime Technology, Trade, and Economic Development," *American Economic Review*, 2017, *107* (9), 2821–2854.

Redding, Stephen and Anthony J Venables, "Economic geography and international inequality," *Journal of International Economics*, 2004, *62*, 53–82.

Ross, Michael L., "What Have We Learned about the Resource Curse?," *Annual Review of Political Science*, 1988, *18*, 239–259.

Ruffin, Roy J., "The Missing Link: The Ricardian Approach to the Factor Endowment Theory of Trade," *American Economic Review*, 1988, *78* (4), 759–772.

_ , "The Ricardian Factor Endowment Theory of International Trade," *International Economic Journal*, 1990, *4* (4), 1–19.

Schoff, Wilfred H., *The Periplus of the Erythraean Sea*, New York: Longmans, Green, and Co., 1912.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Temin, Peter, "Mediterranean Trade in Biblical Times," in Ronald Findlay, Rolf Henriksson, Hakan Lindgren, and Mats Lundahl, eds., *Eli Heckscher, International Trade, and Economic History*, Cambridge: MIT Press, 2006, pp. 141–156.

_ , *The Roman Market Economy*, Princeton: Princeton University Press, 2013.

Wacziarg, Romain and Karen Horn Welch, "Trade Liberalization and Growth: New Evidence," *World Bank Economic Review*, 2008, *22* (2), 187–231.

A. Appendix: Robustness Checks

This appendix contains a number of robustness checks for the results presented in the main analysis. The first subsection includes results for measuring bilateral trade with product shares instead of counts in Section 4. The second subsection includes results for changing the date range of sites used for the development estimations in Section 5. The third subsection covers variations in the size of the sample grid for cities used throughout Section 5.

A.1 Product Shares

In the gravity analysis in Section 4, the dependent variable measuring trade flows was constructed as the sum of exported and imported products between two cities. This captured the Ricardian nature of trade dynamics in the ancient world. Despite this, there are alternative ways to measure trade between two cities while maintaining this Ricardian emphasis on the number of products exported or imported. One of these is through trade shares, where bilateral trade is given as a measure of exports and imports between two cities as a proportion of total exports and imports across the two cities:

$$X_{ij} = \frac{\sum_g x_{ijg}}{Z_i + Z_j}$$

$$\text{with } x_{ijg} = \begin{cases} 1 & \text{if good } g \text{ is traded between cities } i \text{ and } j; \\ 0 & \text{otherwise.} \end{cases}$$

where x_{ijg} is an indicator function for whether a given product g is traded between cities i and j, and Z_i is the total number of goods exported and imported by city i.

Table 9 presents estimations of both the standard linear and non-linear gravity models from Tables 5 and 6 using this share-based measure of bilateral trade. the results are very much in line with those of the main analysis, with only marginal changes to the magnitudes of the coefficients. This strengthens the evidence for both a Ricardian gravity system of trade in the Indian Ocean during the late iron age and the crucial role of the monsoon patterns in facilitating this system.

A.2 Third Century Decline

In the principal analysis, a cut-off of 640 CE was chosen for the presence of archaeological sites used in the data for development levels. This date was chosen as it is the date of the Arab conquest of Egypt. Despite the interruption of Roman access

Table 9. Gravity Model: Product Shares

	(1)	(2)
ln(SailDist)	-0.260**	-1.407**
	(0.104)	(0.533)
Neighbour		-2.371**
		(0.799)
Endowment Index		-12.690**
		(1.897)
City FE	Yes	Yes
N	107	446
R^2	0.850	0.980

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is the share of total good varieties traded; ln(SailDist) is the log of the total round-trip sailing time for the fastest calculated route accounting for wind speed; Neighbour is a dummy variable for cities in close proximity; Endowment Index is an index of similarity of endowments.

being a defining moment in the operation of the Indian Ocean trade network, there is evidence that Roman involvement had been in decline for several centuries prior to this date, beginning in the early third century CE.[18]

In order to address this possibility, the data used for this estimation have a cut-off date of 300 CE—the final date by which the purported decline of the third century would have occurred. Table 10 reproduces the results of Section 5 with the new date. The implications of the results are not different from those presented in the main analysis. Specifically, the effects for export and import diversity are nearly identical; and the results for manufactures, spices and aromatics, and bullion are very similar with increased development linked to the export of manufactures in particular and decreased development linked to the export of cash-crops. Dependence on bullion imports remains linked to significant development. This strengthens the results presented in the main analysis and provides more evidence for the possibility of both an export effect on development and a potential 'resource curse'.

[18] Cobb (2018) and McLaughlin (2018) contain a more comprehensive discussion around the dating of this decline.

Table 10. Third Century: Ancient Development

	(1)	(2)	(3)
ln(Dist)	-0.091*		
	(0.054)		
Export Diversity		31.640**	
		(11.193)	
Import Diversity		-15.517**	
		(6.688)	
Manufactures			3.082**
			(1.461)
Spices & Aromatics			-22.260*
			(13.255)
Bullion			641.890**
			(315.347)
Rome	0.381**	0.458**	0.456**
	(0.192)	(0.187)	(0.187)
Cell FE	Yes	Yes	Yes
N	3172	3172	3172
R^2	0.065	0.068	0.067

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is the standardised change in ancient place records within a grid cell with a date limit of 300 CE imposed; ln(Dist) is the distance from the nearest Periplus port; Export and Import Diversity are distance-weighted Herfindahl-style indices of export/import product concentration; Manufactures and Spices & Aromatics are distance-weighted measures of export shares for the given sector; Bullion is the distance-weighted measure of import share of bullion products; Rome is a dummy for the grid cell being a part of the Roman Empire.

A.3 Variations in Hinterland Distance

There is potential for the development effects analysed in Section 5 to be sensitive to the size of the 'hinterland' radius around a city included in the sample. The main analysis set the limit of the sample at 500 km, the results in this subsection include two different values for this limit: 350 km and 150 km.

Table 11 reproduces the results of Section 5 with the two alternative sample sizes. The results of Table 11 are broadly in line with those of Table 7 for both alternative sample distances. The coefficients have different magnitudes, particularly the ones for the 150 km sample, but the signs remain the same and all are similarly significant, supporting the trends outlined in the main analysis: increased exports are correlated with higher levels of development and the potential for an ancient 'resource curse'.

Table 11. Alternative Sample Distance: Ancient Development

	350 km			150 km		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Dist)	-0.158**			-0.306**		
	(0.074)			(0.144)		
Export Diversity		42.234**			55.579**	
		(12.678)			(17.047)	
Import Diversity		-21.947**			-30.188**	
		(7.574)			(10.151)	
Manufactures			3.024*			2.876*
			(1.577)			(1.572)
Spices & Aromatics			-32.943**			-45.012**
			(15.394)			(21.056)
Bullion			724.274**			1108.955**
			(367.894)			(539.840)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
N	2042	2042	2042	601	601	601
R^2	0.109	0.114	0.111	0.397	0.410	0.404

Standard errors in parentheses

* Significant at 0.10; ** Significant at 0.05

Note: Dependent variable is standardised change in ancient place records within a grid cell; ln(Dist) is the distance from the nearest Periplus port; Export and Import Diversity are distance-weighted Herfindahl-style indices of export/import product concentration; Manufactures and Spices & Aromatics are distance-weighted measures of export shares for the given sector; Bullion is the distance-weighted measure of import share of bullion products; Controls for the grid cell being a part of the Roman Empire are included.