UTRECHT UNIVERSITY

Faculty of Geoscience
Institute of Marine and Atmospheric Research

Climate Physics Master Thesis

Climate change impacts in the Late Roman Empire: a quantitative analysis

First examiner:
Brian Dermody

Candidate:
Sara Filippini

Second examiner:
Claudia Wieners

October 15, 2023

Abstract

The study of past societies' relation with climate is a precious resource to contextualize our own struggles with the current climate crisis, but the field has been left largely unexplored by both historians and climatologists. In this project, we aim at providing a contribution to fill this gap. Our focus is on the late Roman Empire, from the onset of the third century crisis (235 AD) to the fall of the Western Roman Empire (476 AD). This period coincided with climatic change in Europe, as it represented the transition between the Roman Warm Period and the Dark Ages Cold Period. We are interested in exploring the relation between climate change and the Empire's crisis, through the use of a computational simulation of the Empire's agricultural productivity and food staples economy. The project was divided in two phases. First we constructed a simulation that reproduced the known historical trade patterns around 200 AD to a reasonable level of accuracy. In a second phase, we tested this simulation through a different reconstructed climate forcing and evaluated the anomaly between them. We found that the climate transition caused a significant rise in grain import costs in several cities of the Roman East, which could justify some of the Empire's socio-economic issues starting around the 3rd century AD.

Contents

1. Introduction

Climate change has been controversially alleged as a trigger for the collapse of the Western Roman Empire [1, 2]. In this project, we attempt to quantify the impact of climate change on food security, in order to evaluate its importance among other factors contributing to the Roman downfall. Our focus is on the period from the Third Century Crisis (235-286 AD) until the fall of the Western Roman Empire in 476 AD.

In section 1.1 I discuss the relevancy of the topic to the present day, especially in relation to the current climate crisis. In section 1.2 I provide the basic background necessary to understand the project's premises, a more detailed literature review is available in the next chapter. In section 1.3 I formalize our research questions.

1.1 Relevancy of the topic

Throughout the last century, anthropogenic climate change has already meaningfully altered the environment in which our society develops. The expected future changes are even more drastic [3].

Consequently, more and more research has been focused on understanding and simulating climate-society interactions.

A variety of models was produced in recent years, whose primary purpose is to predict the impact of climate change on human society [4].

However, these models are required to operate under deep uncertainty, since the theorethical framework is still being developed and validation sources are scarce. This leads to a multitude of simplifying assumption, not always justified, which erode trust in model's results [4, 5].

Although the threats posed by anthropogenic climate change are likely unprecedented in recorded history, human societies have always dealt with

long term natural oscillations in the climate system. If accurately studied, the resilience of past civilizations to such oscillations can help to contextualize anthropogenic climate change, to ground models into real-world evidence and further our understanding of the relation between human society and its environment [6, 7].

Despite this important incentives, research into past societies' relation with climate has been limited. This is partly due to a lack of interdisciplinary communication between historians and climate scientists, which tend to have different understandings of the impact of climate on historical development [6, 5, 8].

We hope to contribute to fill this knowledge gap through our project.

In particular, as explained further in the next chapter, the Late Roman Empire constitutes a case study of a socio-economic system whose ecological resources were very strained, due to excessive urbanization and reliance on trade [9, 10, 11]. This may have led the Roman Empire to be more sensitive to mild or moderate climate change which would not have caused widespread damage in a different society [2].

1.2 Background

The analysis of various paleoclimate proxies has shown that the Northern Hemisphere is subject to a temperature oscillation of roughly 1°C amplitude, on the millennial scale known as Bond cycle (figure 1.1)

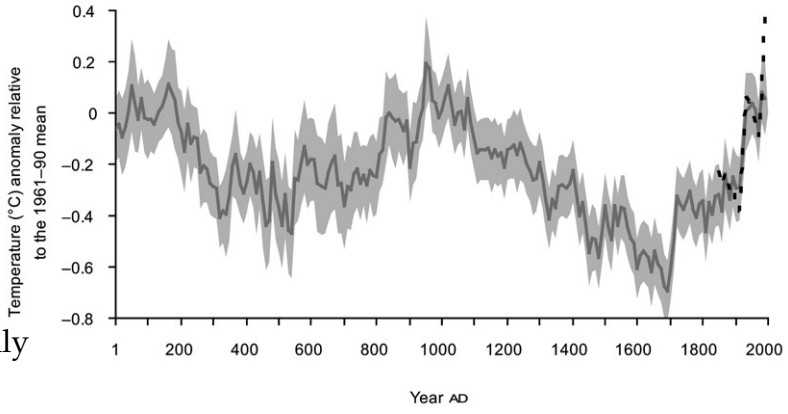

Figure 1.1: Multiproxy reconstruction on Northern Hemisphere extratropical (90° to 60° N) temperature anomaly [12]

[13, 14, 15].

The climatic history of Europe in the last 2000 years can thus be broadly divided into cold and warm phases (also known as Bond events and Bond intervals) [12]. These phases are not uniquely defined in the literature: I provide here the temporal limits which I will be using from now on, but they may appear differently defined elsewhere.

The peak in political power and territorial expansion of the Roman Empire, the so-called *pax romana*, coincides with a warm interval of the Bond cycle, known as Roman Climate Optimum (RCO), or Roman Warm Period (RWP), spanning approximately from 100 BC to 200 AD. During the RWP, most of Europe enjoyed a stabler and warmer than usual climate, as testified by various historical evidence as well paleoclimate proxies [1, 16].

Around 200 AD, the natural proxies begin show an increased precipitation instability, with a higher frequency of extremes, superimposed to a general cooling trend [1, 17]. The period from 200 AD to 500 AD constitutes a transitional phase between the RWP and the Dark Ages Cold Period (DACP), which will last until about 800 AD.

Throughout these three centuries, the historical evidence for droughts, floods, famines and pestilence becomes more frequent [16, 1]. Through intermittent phases, the standard of living declines, as the political and social landscape becomes increasingly unstable [1, 18]. Climate change may have affected directly agricultural productivity, leading to food security issues, other systemic elements of the crisis may be seen as indirect consequences. Hence why we focused here primarily on investigating the impact on food security.

In the fifth century AD, as the Western Roman Empire collapses, we enter the cold phase of the cycle.

The parallel analysis of the Empire's climate and social history leads us to pose the central question of this study: how deeply connected was natural climate change to the Empire's political rise and fall?

1.3 Research questions

Throughout the project, we used a computational model to explore the impact that climate change might have had on the food security of different regions of the Empire. To do so, we will firstly need to construct a simulation that replicates the trade relations in food staples around 200 AD to a reasonable level of accuracy. In a second stage, we will test the simulation through different climate patterns and analyse the anomalies between them.

I divide our research questions into two sections, the first relating to the structure of the trade system around 200 AD, and the second to its resilience to climate change.

1.3.1 Trade network structure/Model validation

1) Is a simple surplus redistribution model, forced with the biophysical and demographic data available to us, sufficient to reproduce the historical trade patterns of the Roman Empire?

We simulate the Roman food staples economy based only on deficit, surplus, and cost distance between regions. By comparing our simulation results with historical trade data, we may determine the extent of our model's accuracy and speculate on the origin behind the inaccuracies. We may also gain insights into the factors influencing trade routes, and determine the significance of cost distance compared to other socio-economic and political factors.

2) What were the key vulnerabilities of the Roman food trade network to climate change?

As described in the next chapter, various potential elements of vulnerability have been identified in the late Roman food staples economy. We want to verify whether any of these elements also emerges in our simulation. Since our model only account for supply, demand and cost distance, the emergence of such vulnerabilities would help identifying them as an inevitable consequence of the population and productivity patterns of the time.

1.3.2 Reaction of the network to the climate forcing

3) What was the impact of the climate transition from RWP to DACP on agricultural productivity, trade routes and importation costs?

We use two climate reconstructions of the Roman Warm Period and the Dark Ages Cold Period as forcing for the model. We then analyse the anomaly in the model output in terms of agricultural productivity and trade patterns. I compare the results with the historical record to gain insight into which known historical developments might have been caused by climate change.

Secondly, we analyse the anomaly in import costs and node degree between the two simulations. I consider the consequences of such anomaly and discuss whether, based on our results, climate change can be blamed for some of the socio-economic issues that began in the 3rd century AD.

2. Literature review

In this chapter, I review the literature on the environmental and historical events of the late Roman period, focusing on their potential connection. In section 2.1 I discuss the historical context on the late Roman empire crises, focusing on systemic issues. In section 2.2 I review the paleoclimate evidence around this period.

2.1 Historical context

2.1.1 Brief overview of the political crises of the late Roman Empire

In the year 235 AD, the emperor Severus Alexander is murdered by his own troops while on campaign in Germany, following an unpopular military decision. The commander Maximinus Thrax is proclaimed new emperor by the legions. This year is traditionally designated as the beginning of the "Third Century Crisis", a period of political, social, economical and environmental instability that will leave profound changes into the Empire's landscape [19, 20].

Maximinus Thrax was the first of a series of "barracks emperors": rulers chosen by the army with no official claim to the succession, who fundamentally acted as warlords over different provinces. The Senate refused to recognise these figures and elected different emperors, who however lacked practical power over most regions. The year 238 AD became known as "Year of the Six Emperors", and represented the peak in the number of contemporary contenders, all of whom were assassinated in the coming years [19, 20].

The political situation returned to a relative stability from the late 3rd/early 4th century, with the Empire now divided into Eastern and Western halves [21, 1].

A second succession crisis in the Western Roman Empire erupted in the 420s AD, in contemporaneity with the mass migration of different German and Eastern European population inside the borders of the Empire, pushed in turn by the Hunnic invasions [1, 21, 20].

The Western Roman Empire eventually collapsed in 476 AD, with the deposition of the last emperor Romulus Augustus [21].

We distinguish then two apparently separate periods of crisis of the late Roman Empire, the third and the fifth century crises. These two periods roughly correspond to two phases of the climate transition, as better described in section 2.2.

2.1.2 Other troubles of the late Roman Empire and their alleged relation to climate change

Of course, it's hard to claim a direct connection between the political and the climatic instability.

However, while the succession crises are the most well documented element of the late roman period, they were only the tip of the iceberg. Most historians agree that various interconnected, systematic issues contributed to the overall lowering of the standard of living from the third century, and most of them have been connected to the climate change in some form [2, 22].

Two important **epidemics** spread across the Empire in the midst and closely before the 3rd century crisis.
The Antonine plague (probably a smallpox ancestor) diffused in various waves between 165 and 190 AD. While there's no agreement on the overall mortality, it is generally recognised that Egypt was the most affected province. Estimates of the population loss in the region range from 10 to 30% [23, 24, 25]. The first outbreak of the plague in Egypt follows closely a two years of unproductive Nile floods (163 and 164 AD). A second outbreak in 170 AD is also preceded by flood failures between 168 and 170 AD [24, 1, 2]. It is possible that the spread of the plague was aided by the food insecu-

rity derived from years of scarce production [24]. A second plague, known as Cyprian plague from the name of one of the witnesses, spread from about 259 to 262 AD, in the midst of the third century crisis, and diffused primarily in those regions experiencing drought conditions [2, 26, 25]. A third large scale epidemics, the Justinian plague, diffused in the Eastern Roman Empire from 541 to 549 AD, during the coldest phase of the DACP. The Justinian plague is usually considered the first recorded example of Bubonic plague [27, 28].

Moreover, from around 230 AD, the northern and eastern *limes* became increasingly vulnerable to **invasion and migration** by different Germanic tribes (Goths, Alemanni and Franks in particular) [29, 30, 22]. Certainly, this is partly a consequence of the political instability and the weakened borders. However, by comparing the recorded events along the border with tree ring growth data, it has been noted that the more than 75% of the events happened in years of relatively low growth rate [31]. It is thus possible that the Germanic tribes migration were led by drought conditions and years of unstable agricultural production. The second peak in invasions was also an attempt to escape the Hunnic invasions, which were in turn driven into Europe by years of extreme drought in Central Asia [1, 32].

Possibly in connection to the German invasion, in the latter part of the third century we observe an exacerbation of **peasant revolts** against the local Roman authorities, especially in Gaul and Spain [33, 34, 35]. While these regions had often experienced similar insurrections in the past, the two peaks in the third and the fifth century AD are evident, and they correlate to climate change, specifically with dry phases in the western provinces [35]. It is possible that food insecurity was part of the trigger for the peasant revolts, while not being the primary motivation.

Beginning from around 160 AD, partly prompted by the insurgence of the Antonine plague, the Empire also begins to experience **inflation** [22]. Prices of wheat, the product over which we have the most reliable information, increase steadily between 160 and 190 AD, and then seem to remain stable until about 250 AD [36, 37]. In the second half of the third century

prices increase more rapidly, and especially from 293 AD to 301 AD we have unsustainable inflation, up to 500% price increase in less than a decade [37]. In 301 AD, emperor Diocletian passes the "edict of maximum prices", which sets the maximum reselling price for the most fundamental commodities, firstly wheat [29]. Data availability reduces in the following centuries, but it is likely that inflation continued at varying rate until the Empire's fall. While the inflation phenomenon, especially the very rapid phase from 293 AD, was partly due to Rome's unwise financial policy [22, 37], it is likely that the more gradual increase throughout the third century was related to real supply limitations. Climate change, and consequent loss in agricultural production, once again appear as a possible candidate.

It has been noted, that also during the period of relative stability of wheat prices, from 190 to 250 AD, other proxies for inflation continued to increase. It's possible that the price of wheat was artificially fixed due to heavy state intervention in this market [36].

In conclusion, most of the troubles of the Roman Empire from the third century onwards can be more or less reliably associated to direct or indirect effects of climate change. This doesn't necessarily imply they were caused by climate change. All of these events can appear in a society without any underlying environmental pressure. However, the correlation is such to motivate a quantitative analysis of potential effect of the climate shift on the agricultural output and food security of the Empire. Any speculation on potential indirect effects should begin with an evaluation of the scale of the direct impact.

2.1.3 The roman food staples economy and its vulnerabilities

By the second century AD, the Roman trade network was an incredibly complex system, in which long distance trade was relatively common, and state intervention was an established practice [11, 38].

Around this period, the Roman Empire was at the peak of its expanse, spanning 33 degrees latitude and 34 degrees longitude. A size rarely ac-

complished by a unitary states in history. The North to South extension is of particular notice if we consider that ecology generally favours East to West expansion [39].

Further, the climate of the Mediterranean region is characterized by a particularly large spatial variability, with many subregional features [40].

The Romans' ability to connect these different subregions, utilizing the spatial climate variability to alleviate the impact of harvest shocks, played a vital role in the Empire's success (and demise). [39, 41]. The Mediterranean Sea, which allowed for relatively cheap long distance transport, was the central element of the resulting trade network [39].

However, many have pointed out that some of the characteristics of the Roman food staple economy, while preserving the system against harvest shocks, also might have hindered its resilience in face of large-scale, long-term climate change [10, 9, 39].

Throughout the first and second century AD long distance trade became increasingly common among the roman provinces [11, 42]. This is testified, among others, in the dating of excavated shipwrecks (figure 2.1) [11]. The rapid globalisation process led many regions to become largely dependent on importation for their survival [43]. In particular, as noted even by contemporaries, Rome was dangerously dependent on Egypt for its grain supply [44, 45]. Since the long distance exporters were relatively few, and quite specialized, the result was a system in which reduced production in a single region could have a strong effect on several importing regions [45, 46].

This situation certainly applied to the grain trades, over which we have the largest wealth of information, but it was also becoming increasingly true for oil trades [47, 48]. We know, for instance, that around 80% of the oil imported towards Rome came from the province of Baetica, in Southern Spain [49, 47].

Further, in regions reliant on importation for primary food staples, hence shifting away from an agricultural economy, urbanisation and population growth were encouraged [43, 9, 11]. It is estimated that, by the third cen-

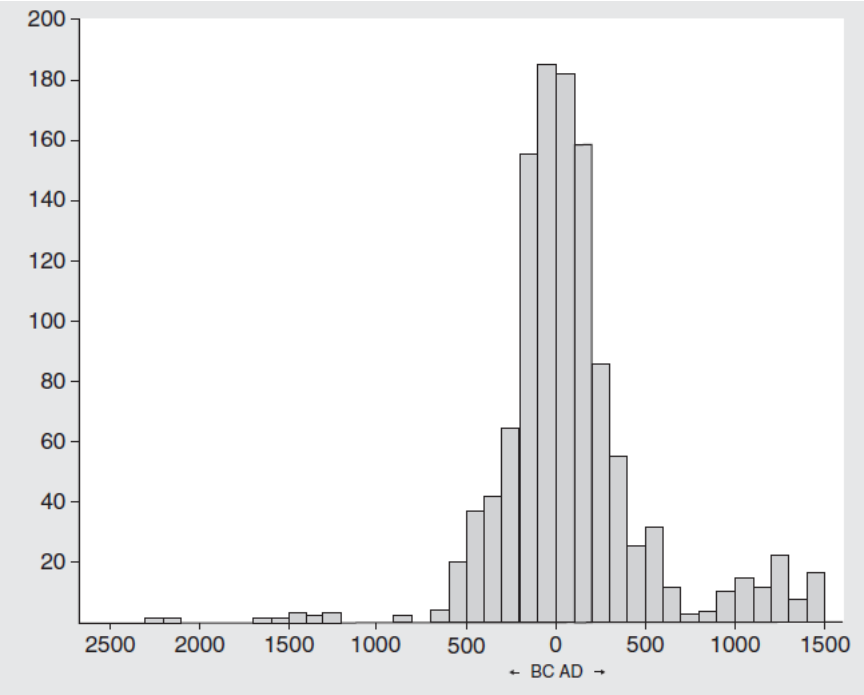

Figure 2.1: Distribution of the dated shipwrecks in the Mediterranean by century, figure by [11]

tury, almost 40% of the Italian population lived in cities, and was not involved in agriculture as their main occupation [11]. In the Augustan period, Rome itself surpassed one million inhabitants, a size that won't be seen again in Europe until the 19th century [45]. This situation was exceptional for a pre-industrial society, and some have speculated that the Empire was approaching the limits of its ecological resources, leaving it vulnerable to climate mean shifts as well as climate instability [10, 9].

Wine, which rose to the status of staple product from the 1st century AD, gradually supplanted cereal production in these urbanised, import-reliant regions [43, 50]. It was mostly produced in Italy, France and Northeastern Spain [51, 49].

Finally, specifically for the grain trades, there's indication for increasing state intervention between the 1st and 2nd century AD [44]. Most of the evidence concerns the supply to Rome, which was the primary concern of authorities [41, 44].

Before the mid-1st century AD, the state intervention was limited to stipulating short term (single year) contracts with *navicularii* who would trans-

port the grain to Rome from the producing provinces, and buying grain for the fixed price distributions to the city's poorest (*frumentationes*) [46]. Later, there was a shift towards ensuring more continuous supply, providing benefits to those who transported grain to Rome for several years in a row. Around the 2nd century AD, we have construction of public granaries (*horrea*) [46, 52].

There is some historical evidence for state intervention in the oil trades towards Rome in the third century AD, but it was likely limited to a brief period [49]. Wine trades appears to have been almost entirely liberalised [50]. The only area of regular state intervention for these products would been the supply to the legions, whenever the staples couldn't be sourced locally [41].

Increasing state intervention in the grain trades was probably a necessity due to the growing urban population. However, it might have reduced the elasticity of the system, by promoting the use of a reduced number of stable importers, and exacerbating the impact of a reduced production by such importers due to, for instance, climate change.

2.2 Climate change in the third to sixth century AD

The period from the third to the sixth century AD encapsulates the transition between the two phases of the Bond cycle [15, 12]. During this time, on average, the European continent experienced a shift towards colder and dryer conditions [14, 1].

The transition developed in three phases. A moderate cooling is reported in the early 3rd century AD, interrupted by a phase of temperature stability from the mid 4th century AD. The temperatures rise slightly throughout the fifth century, to fall abruptly in the early 6th century, entering the coldest phase of the DACP, known as Little Antiquity Little Ice Age (LALIA), spanning from 536 to 660 CE [1, 53, 54].

As for precipitation, a first dry phase began in the 3rd century AD, fol-

lowed by a wet interval in the late 4th and early 5th century. A second harsher dry phase began around 430 AD, leading to drought condition in some provinces [1, 14, 55]. While the cooling trend developed homogeneously throughout the continent, dry conditions were only recorded in some regions in Central Europe and in the East and Southeast Mediterranean [56, 1].

The temporal pattern, in particular that of precipitation changes, overlaps with the phases of the Roman Empire crisis. From the third century crisis, the division of the Empire and relative stability of the 4th century, to the invasions and complete breakdown of the Roman authority in the 5th century.

I discuss here some of the paleoclimate proxies that substantiate this description.

2.2.1 Climate forcing

The leading climate forcings around this period would have been solar irradiance and vulcanism.

The exact mechanism behind the Bond cycle in not entirely understood, however, it has been linked to oscillations in the solar activity, mediated by changes in the North Atlantic circulation [57, 15].

The reconstructed solar irradiance timeseries (figure 2.2) shows a downward trend for the first 500 years of the common era, contemporaneous with an increase in the amplitude of oscillations [58, 1, 17].

The period from the 3rd to the 5th century AD also shows an increase in volcanic activity, especially around the tropics [59, 14, 1]. Volcanic events can have a cooling effect the Earth's climate through stratospheric sulfates injection. The strongest impact of volcanic forcing is clustered after the onset of the Dark Ages Cold Period, in the two decades from 530 to 550 AD. The eruptions of 536 AD and 540 AD correspond to an estimated radiative forcing of $-11.3Wm^2$ and $-19.1Wm^2$ [59, 14]. A smaller but still impactful eruption is placed around 547 AD [59, 14]. The radiative forcing due to

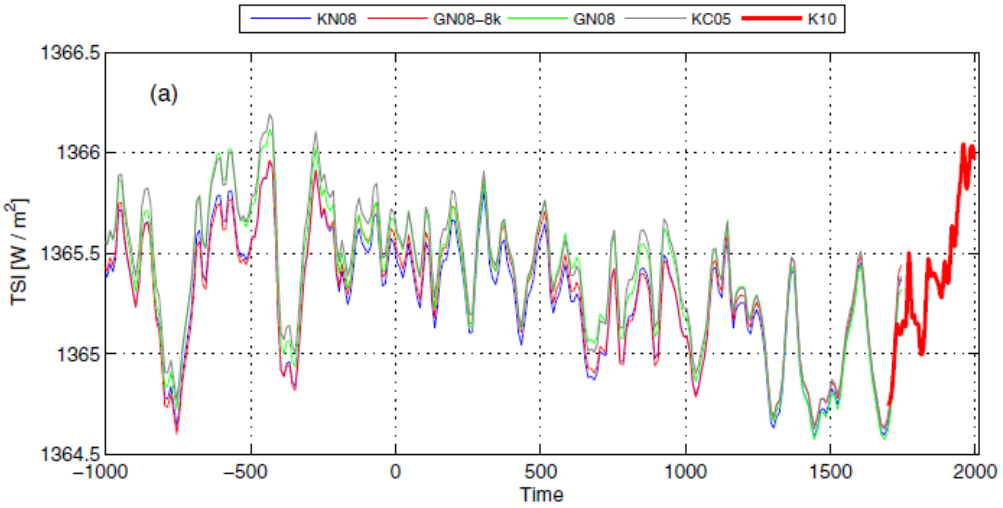

Figure 2.2: Different reconstructions of the total solar irradiance based on ^{14}C isotope measures. Figure from [58].

volcanic eruptions in shown in figure 2.3.

The most recent theory is that the climate transition from RWP to DACP (which was perceived on a global scale) was initiated by solar forcing oscillations, while the volcanic activity was the main trigger behind the coldest LALIA period (which was a more local European phenomenon) [54, 14, 15].

2.2.2 Temperature cooling

A variety of natural proxies as well as climate simulations paints the picture of an overall temperature cooling in the Northern Hemisphere beginning in the 3rd century AD [12, 54].

In the case of Europe, cooling is most reliably testified by glacier expansion [53, 1, 60]. The great Aletsch glacier, the largest glacier in the Alps, experienced a growth phase from about the mid 3rd century until the 8th century, but a brief period of retreat is recorded around the 5th century AD (figure 2.4) [53, 1]. Glaciers respond to climate alterations with varying delays depending on their size, so the growth of very large glaciers is a reliable proxy for sustained temperature cooling [53].

Buntgen et al. also reconstructed summer temperatures based on tree ring width datasets across Central Europe [14], as shown in figure 2.5. This

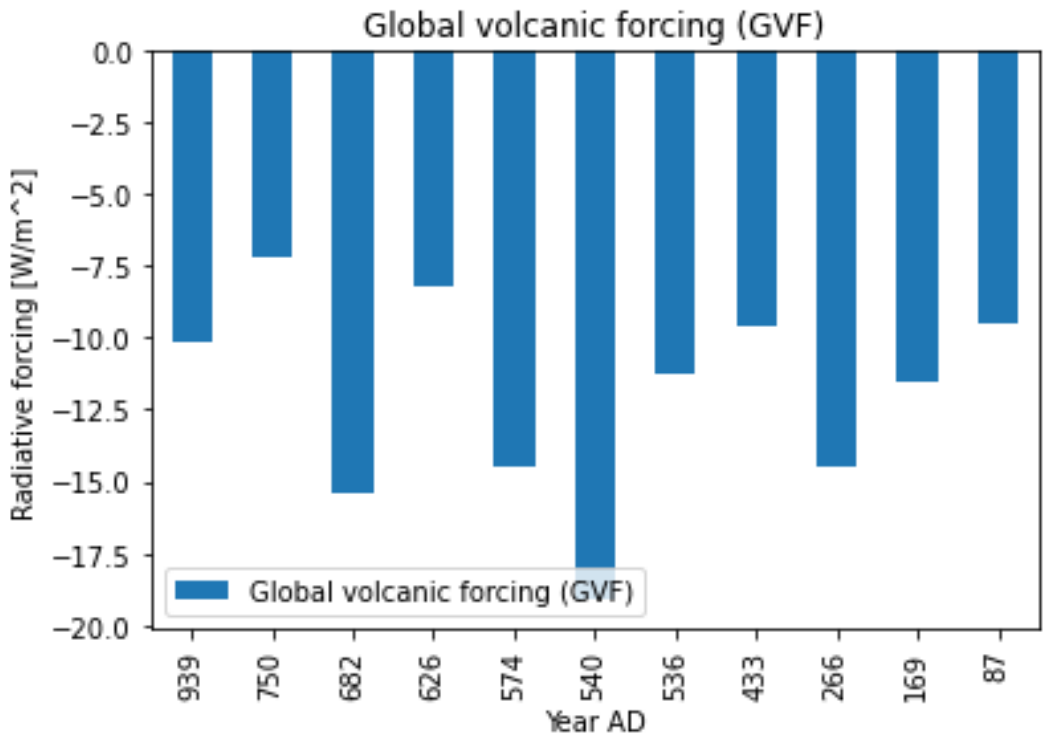

Figure 2.3: Vulcanic forcing in the tropics and Northern Hemisphere, for the first 1000 years of the Common Era, reconstructed from ice core sulfate measurements. Figure created with the data from [59].

record reflects the same described pattern over time, with a gradual decrease from the 3rd century, interrupted by a brief phase of stability in the 4th century.

Overall, the decrease in temperature appears to have been spatially homogeneous throughout Europe [1].

2.2.3 Precipitation pattern

Between the 3rd and the 6th century AD we also observe larger precipitation variability compared to the RWP [1, 14]. Compared to the relatively homogeneous cooling trend, the spatial patterns of precipitation anomaly were more complex, and some regions even experienced an increase in precipitation in the transition period [56, 1].

Buntgen et al. have reconstructed spring precipitation around this period through various tree ring records in central and western Europe (figure 2.6) [14]. We observe two separate dry phases, in the 3rd and the 5th century,

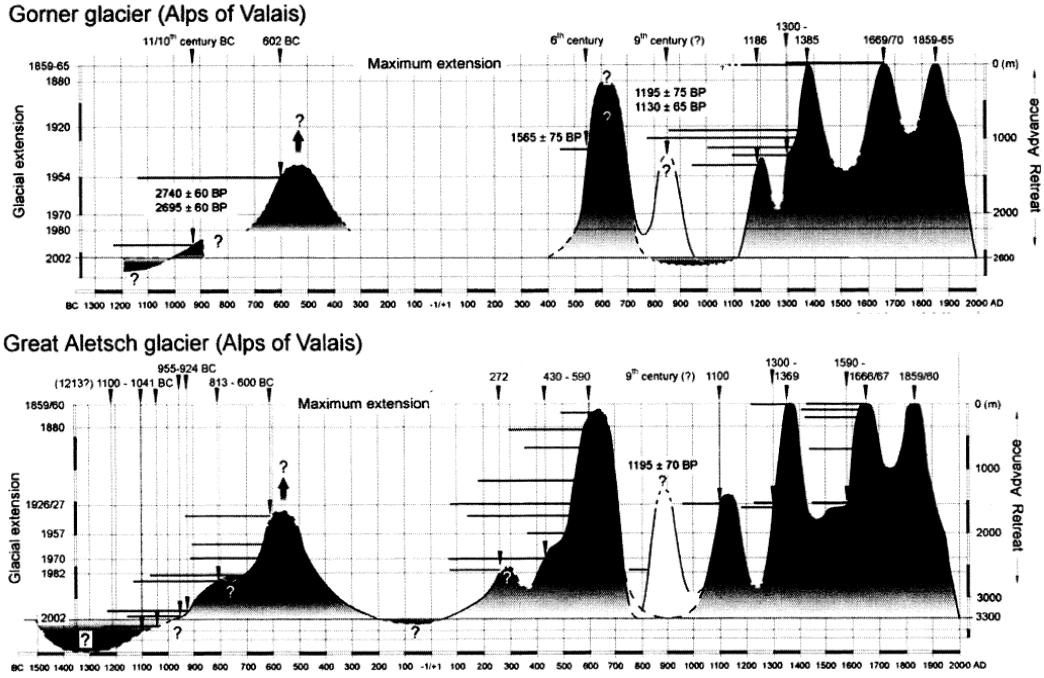

Figure 2.4: Growth period for the Great Aletsch and Gorner glacier in the Swiss Alps, for the last 3000 years. Figure retrieved from [53].

with a wet interval between the 4th and the 5th century.

A similar pattern has been observed in the east and southeast of the Empire, as verified for instance by the reconstructed levels of the Dead Sea [1, 61].

Overall, the spatial analysis of different proxy records seems to reflect a seesaw precipitation pattern [56, 1]. Palestine, Egypt, the Iberian peninsula, and central Europe experienced drier conditions throughout the third and again and in the late fourth and fifth century, with the brief interval of high precipitation in the late third and early fourth century. The central Mediterranean, Turkey and Greece retain a wet climate and even shift towards wetter condition in the fourth century [56].

These precipitation changes reflect an overlap of the negative North Atlantic Oscillation (NAO-) and of the East Atlantic/Western Russia (EA/WR) pressure pattern [56]. The instability of the NAO index in the 3rd to 5th century AD is attested in some natural records [62, 16].

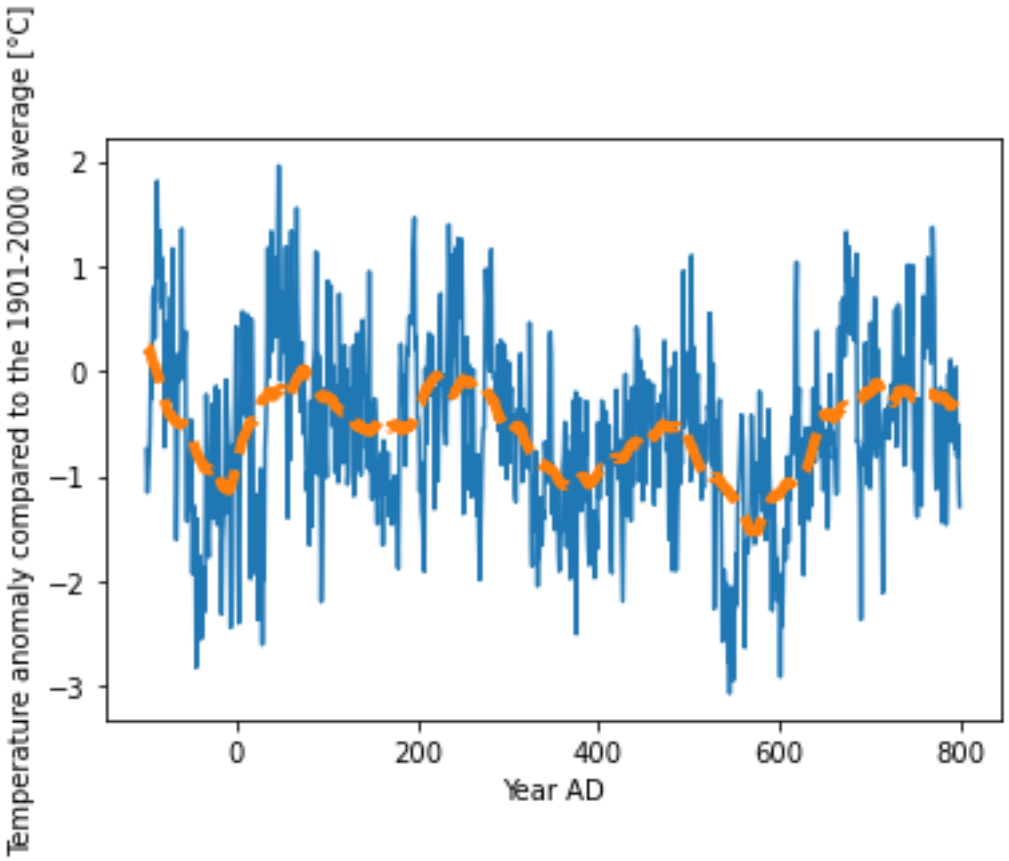

Figure 2.5: Temperature reconstruction based on tree ring width, for the period from 100 BC to 800 AD, for France and Central Europe. The dashed line is a 40 year running average. Figure reconstructed from the data used in [14], available at `https://www.ncei.noaa.gov/access/paleo-search/`.

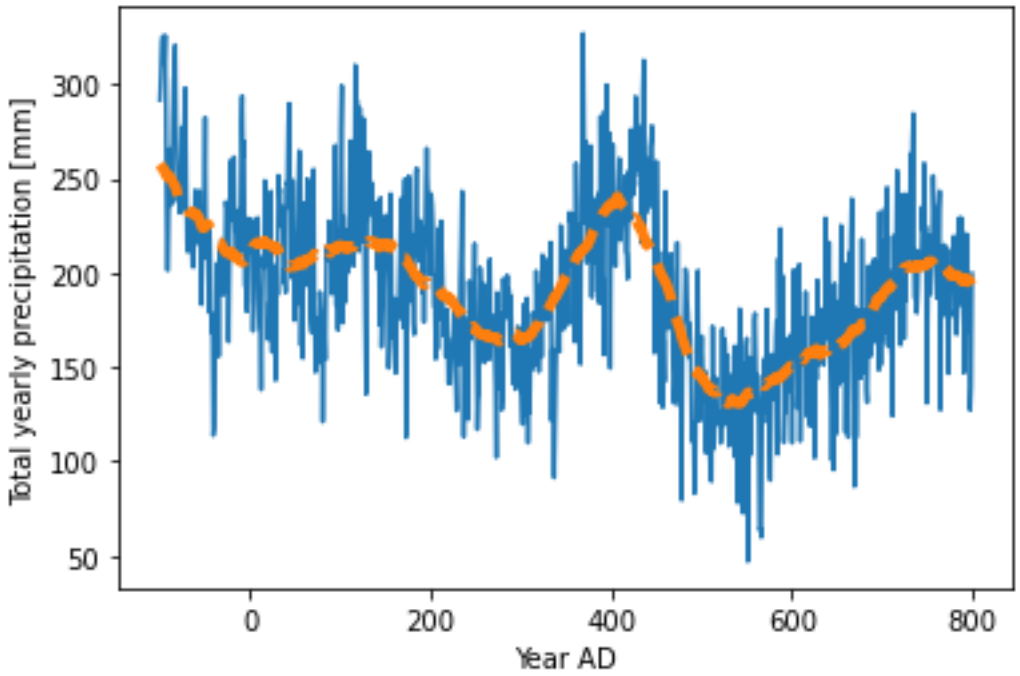

Figure 2.6: Reconstruction of April-May-June precipitation based on tree ring width, for the period from 100 BC to 800 AD, for France and Central Europe. The dashed line is a 40 year running average. Figure reconstructed from the data used in [14], available at `https://www.ncei.noaa.gov/access/paleo-search/`.

3. Methods and data

We employed a computational simulation of the Roman trade network in food staples in order to answer our research questions. We focused on the products that were traded over long distances, and represented staples in the diet of most of the Empire: grain, olive oil and wine. The production was quantified in terms of their raw sources (wheat, olives, grapes), while the trade component of the model uses the finite products. The model has been programmed in Python language.

Our reference date for the model's set up is 200 AD, since we will focus mostly on the switch from the RWP to the more unstable climate of the third century. This date also roughly corresponds to the peak in the Empire's spatial integration: the third and fourth century AD will see a gradual reduction in long distance trade connections in favour of more local production [42, 39, 48]

It's important to stress that we do not expect an entirely accurate reproduction of the Roman trade routes. Our simulation is nothing but a coarse approximation of the Roman economy, that accounts only for the most quantifiable data (population, agricultural production, cost distance). However, the patterns of (in)consistency of our simulation can still provide information on the drivers behind the trade network, or on the importance of certain factors over others.

In section 3.1 I will describe the model's set up. In section 3.2 I will summarise the model's simplifying assumptions. Finally, in section 3.3, I will describe the climate forcing of the three simulations we ran.

3.1 Model description

The model we used is made out of two main blocks. The first computes the agricultural output in each region based on the climate and the population data we provided. The second determines the expected trade flows within regions based on the cost-distance between them. The two blocks are described in further details in section 3.1.2 and 3.1.3 respectively. In section 3.1.1 I describe the required model input. Figure 3.1 summarises the model structure.

3.1.1 Input data

The model relies on three main external drivers: climate, population and cost distances.

Climate forcing is provided as daily precipitation and temperature data. Initially we used reanalysis data of the last a hundred years, later we applied a reconstruction of the two phases of the Holocene climate as described in section 3.3.

Population was divided into urban and rural population. Urban population was taken from the estimates by Hanson et al.[63], while the rural population was reconstructed from archaeological data [64]. The resulting population map is plotted in figure 3.2.

Cost distances, needed for the trade simulation, were derived from the Stanford Geospatial Network Model of the Roman World, known as ORBIS. ORBIS estimates the cost of transporting a kilogram of wheat between any two cities of the Roman Empire, based on a variety of historical and archaeological information [65].

The agricultural land was allocated according to the History database of the Global Environment (Hyde) [66, 67].

Finally, the land which was not allocated to agriculture was assigned as natural vegetation according to the Global Land Cover Characterisation database (GLCC) [68].

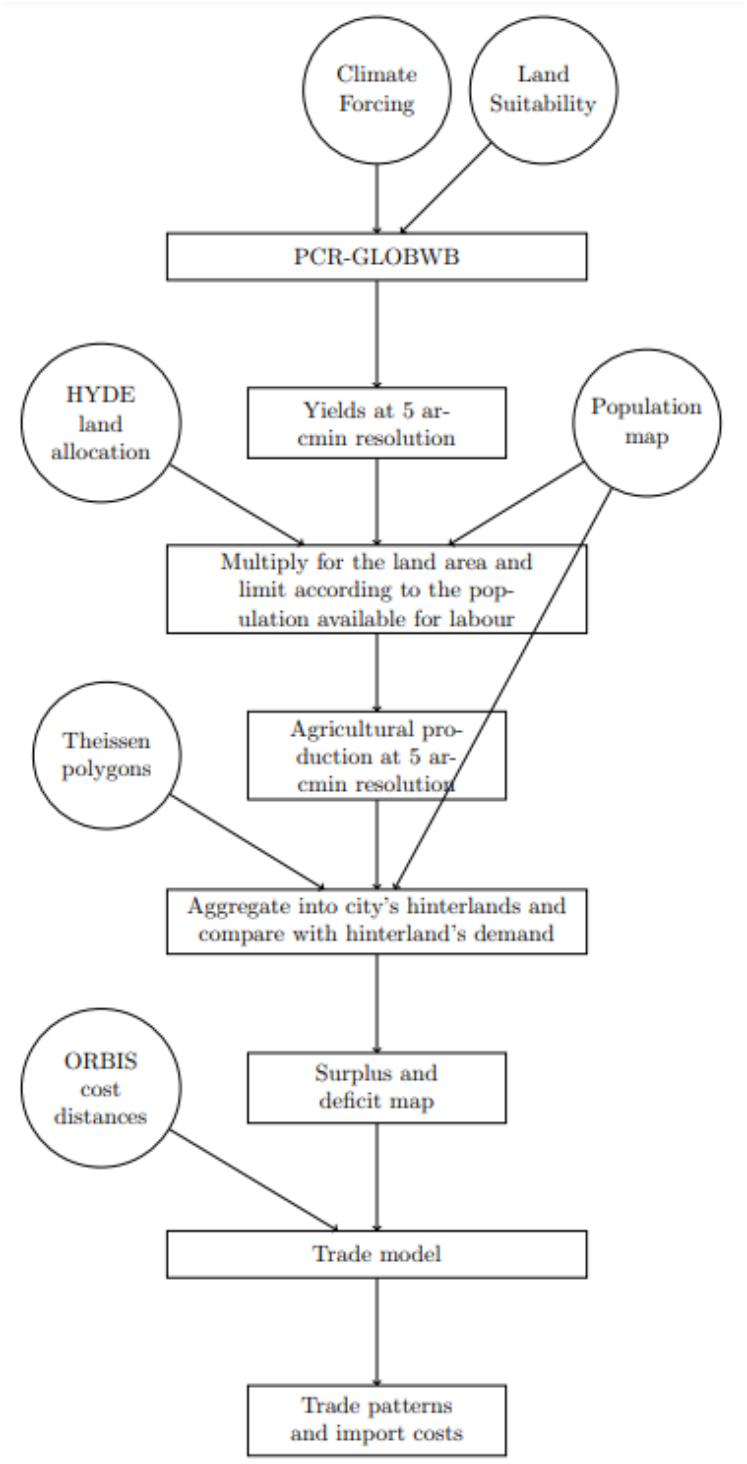

Figure 3.1: Explanatory flow chart of the model

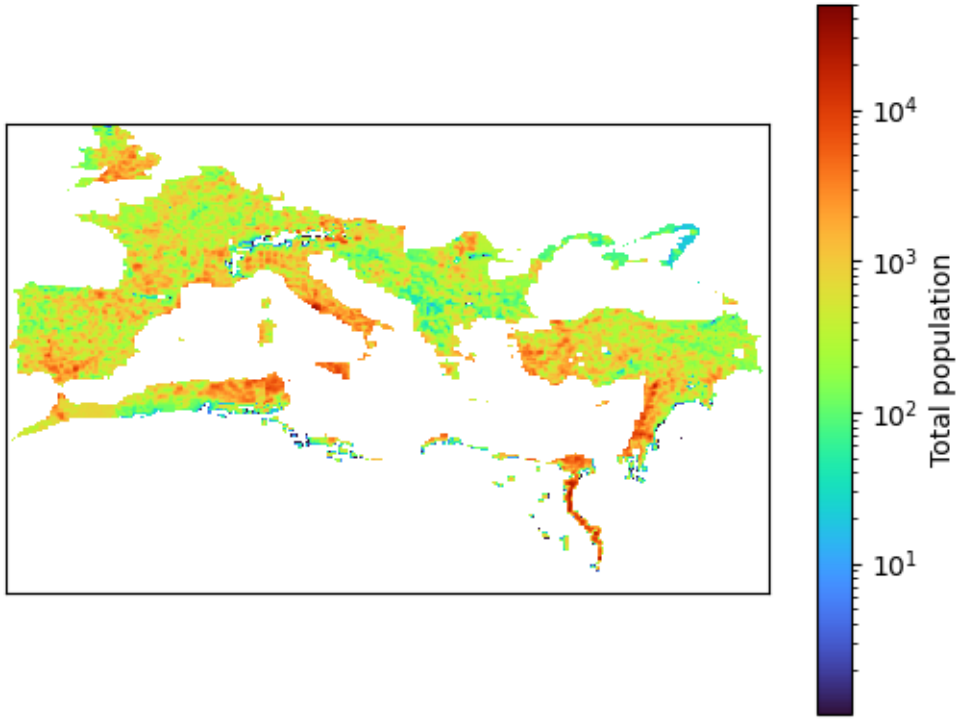

Figure 3.2: Reconstructed population map that was used as input for the model. The urban estimates are derived by Hanson et al. [63], while the rural estimates are reconstructed by Dermody et al. from archaeological evidence [64].

3.1.2 Agricultural production model

The agricultural output is based on the hydrological model PCR-GLOBWB. PCR-GLOBWB takes as input daily climate data on precipitation and temperature, as well as the data on natural and agricultural land cover, and produces simulated evapotranspiration at a resolution of 5 arcmin (roughly 9 km).

PCR-GLOBWB is a biophysically forced model, meaning that the hydrological cycle is here represented explicitely. Where yields can be improved by irrigation, it was included in the model, depending on local river water availability [69].

The data on evapotranspiration were then used to derive the average yields, as kilogram per hectar, based on growth factors dependent on the crop.

We multiplied the average yields for the land allocated to each crop, to calculate the total production of each crop in each 5 arcmin cell.

Finally, the production was limited by the population available for harvest. For this purpose, we assumed 55% of the rural population was available for labour, and each could harvest a total of 6 hectars of wheat, 1.5 hectars of grapes and 2400 kg of olives over one harvest season. This is a reasonable approximation based on historical data.

3.1.3 Simulated trade network

The trade component of the model constitutes a redistribution network, where agricultural surplus flows between different regions based on cost-distance. As mentioned in section 3.1.1, cost-distance refers in this case to the cost of transporting a kilogram of wheat across any two nodes through the ORBIS transport network (figure 3.3). From here on, when I mention nodes being more or less distant from each other, I am referring to their cost-distance rather than their geographical distance.

The redistribution network is applied separately to each product. The trades happen on a yearly resolution, meaning that we compare the yearly

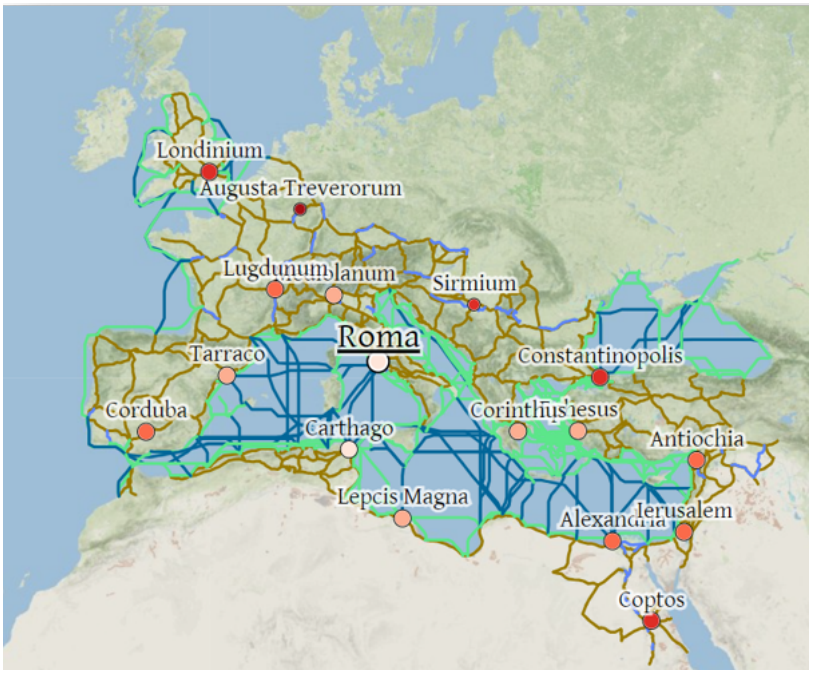

Figure 3.3: The ORBIS transport model of the Roman World, used to determine the cost-distances between different nodes in the model, available at `https://orbis.stanford.edu/`.

demand to the yearly production to determine surplus and deficits. This assumption is reasonable if we consider that most trade took place in the summer when the sea routes were open; however, it also gave rise to some issues which I will discuss later.

As a first step, we aggregate the total yearly agricultural output (produced through PCR-GLOBWB at 5 arcmin resolution) into larger regions which can be considered as independent entities trading with each other. These regions are determined by performing a Thiessen polygon operation. Thiessen polygons are a method used to divide a plane into contiguous regions, each centered around a "seed", in such a way that each point within a region is closer to its regions's seed than to any other's. In our case, the seeds represent 648 cities chosen among the database by Hanson et al. [63], and the Thiessen polygons can be seen as the cities' hinterlands. The resulting territorial division is plotted in figure 3.4.

The yearly demand for each product in each hinterland was determined based on the population data, assuming:

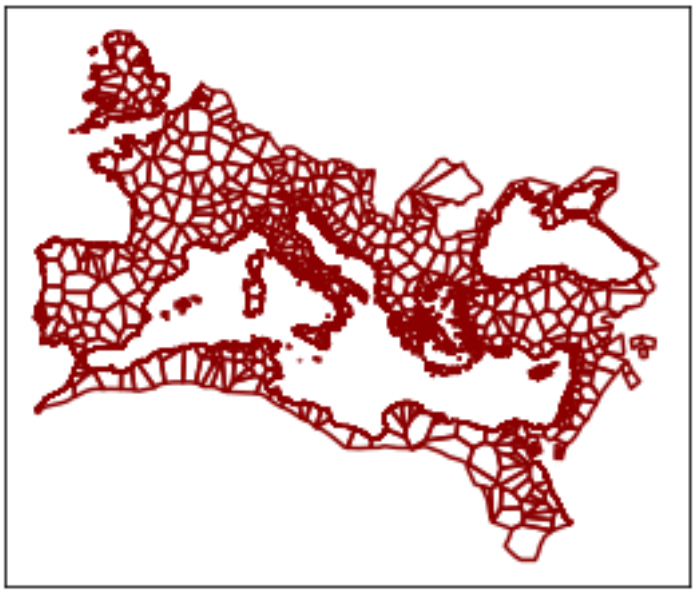

Figure 3.4: Hinterlands of selected cities. The agricultural output is aggregated in these regions, which are considered by the trade models independent entities trading with each other.

- Grain yearly consumption per capita: 220 kg, corresponding roughly to the same quantity in wheat [70].
- Oil yearly consumption per capita: 6 litres, corresponding roughly to 30 kg of olives [70].
- Wine yearly consumption per capita: 185 litres, corresponding roughly to 264 kg of grapes [71]

By subtracting the yearly consumption to the yearly production we determine the surplus or deficit in a certain product for each node.

The redistribution algorithm works as follows.

All the nodes with a deficit simultaneously request the resources necessary to meet their demand to the closest node with a surplus, which responds according to its surplus availability. In a second stage, they all request from their second closest node, whose surplus may now have been limited by exports towards their closest deficit region. And so on until all demand is satistfied or all surplus is exhausted.

This operation is repeated in a series of stages, the number of which can

be set externally as a parameter. Each stage operates only with a fraction of the total surplus, that is, the surplus of each surplus producing region is divided by a value equal to the number of stages, what is left from the redistribution from one stage is passed onto the next one.

The reason behind this fractionation operation lies in approximating the higher complexity of real world systems, by forcing a degree of diversification in importations. In fact, when we simulate the whole network in a single step, the result we obtain is extremely localized, in which each deficit region imports only from one or two close neighbours. This is not historically realistic, since we have a lot of evidence for long distance trade within the Empire's borders. In real world systems, importing regions are not likely to rely on a single exporter, both because it stifles competition and because surplus would not be available at the same time everywhere. The number of stages in the redistribution operation determines how deep in the network an importing region should reach in order to fulfill its demand. A higher number of stages implies a system less constrained by cost-distance, and more importing nodes competing for the same resources.

We use 1000 redistribution stages for our model results. In appendix A I present a sensitivity analysis around this parameter, that justifies this choice of the parameter.

The trade network model operates separately for the three products.

3.1.3.1 Hysteresis

By hysteresis we mean here the tendency to preserve old trade connections, even if the optimal solution would be to create new ones.

This element is implemented in our trade network through year-by-year lowering or increasing of the cost-distance between two regions based on the intensity of their trade connection, in relation to the demand of the importing node. The cost distance modification depends on the fraction of the demand fulfilled through the specific connection, as in:

$$costdistance_{n+1}^{AB} = max(costdistance_{n}^{AB} * (1 - h * (trade_{n}^{AB} / demand_{n}^{A})), costdistance_{n}^{AB} * 0.8)$$

Where n represents the time step (year) of the simulation and h is an adjustable parameter. We place a limit of 80% of the original cost distance, to ensure the system eventually reaches a stable state.

In appendix B we analysed the impact of hysteresis on the trade network and evaluated whether it improved historical accuracy. Since it caused some unstable behaviour I concluded not to apply it to out model results. However, if refined, it may be a useful addition to the model.

3.2 Model assumptions and limitations

Our model stands on a few simplifying assumptions, which are relevant for the analysis of the result. I summarise them and briefly discuss them here.

- Every consumer within the Empire's borders has the same annual consumption for each product. While this is likely sufficiently accurate for grain, it may not be so for oil and wine. And in fact in the next chapter we will observe an inconsistency in the wine trades which is likely the result of this assumption.
- There is no consideration for geographical origin and product quality (eg Italian wine isn't any different from French wine). This also implies that each region with a surplus in a certain product only trades with demand regions (of the same product), and never with other surplus regions.
- The land allocation is determined by modern land suitability maps and local and remote demand within the network.
- Only cost distance (and, when considered, hysteresis) determine the trade flows between surplus and demand regions.
- The system is demand driven. Surplus regions respond to requests by the demand nodes independently of their own convenience.

3.3 Climate forcing and setup of the simulations

The results that I will present in the next chapters are based on three different simulations with three different climate forcing.

- Simulation 1. We used the climate reanalysis ERA6, for the period 1901-2000, without any modification. This climate forcing is referred to as "Historical climate". We use Simulation 1 to test the behaviour of the trade model, and attempt to answer the research questions concerning the trade network structure at 200 AD. The difference in the climate forcing between the last 100 years, and the Roman Warm Period, was not such to alter the fundamental features of the network which we observe in this simulation.
- Simulation 2. The climate forcing for simulation 2 was derived from the "Historical climate" through a selection operation in order to approximate the climate of the Roman Warm Period. Firstly, months were selected which reproduced the RWP precipitation pattern. That is, months with a total precipitation higher than the monthly average in the Iberian peninsula and in the Southeastern Mediterranean, and lower than monthly average in the central Mediterranean. Then we chose days within those months which respected the same pattern in comparison to the daily average. Finally, for the months of December, January and February, days with an average temperature lower than the daily average over the whole territory were selected. The remaining days were put put together in real time order when possible, and looped in order to achieve the same number of years as the Historical climate forcing.
- Simulation 3. The climate forcing for Simulation 3 is meant to reproduce the Dark Ages Cold Period, and it derived through a similar operation as the Roman Warm Period climate, just looking for the inverse trend. Days were selected which showed relatively dry conditions in the Iberian peninsula and the Southeastern Mediterranean, relatively wet conditions in the central Mediterranean, and relatively cold con-

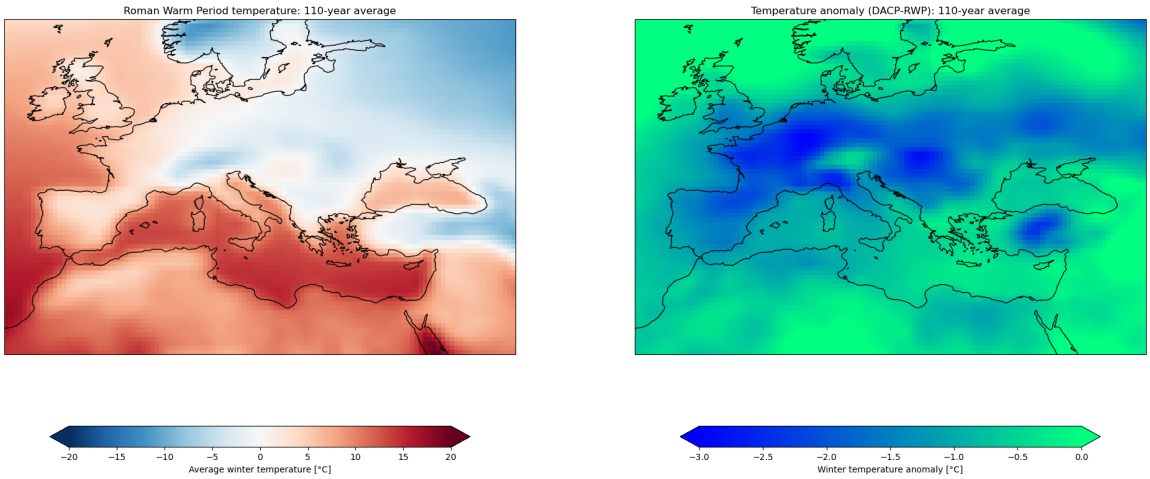

Figure 3.5: Left: 110-year average winter (December-January-February) temperature reconstruction for the Roman Warm Period, obtained through a selection of days from the last century's real climate, in order to approximate the patterns observed in the paleoclimate proxies [1, 12]. Right: anomaly in 110-year winter temperature averages between the DACP reconstruction and the RWP reconstruction (positive values imply a temperature increase, negative values imply a temperature decrease).

ditions over the whole territory.

Figure 3.5 and figure 3.6 show maps of the average temperature and precipitation of the reconstructed Roman Warm Period and the anomaly introduced by the Dark Ages Cold Period reconstruction.

The results from Simulation 1 are analyse in chapter 4, while Simulation 2 and 3 are compared in chapter 5.

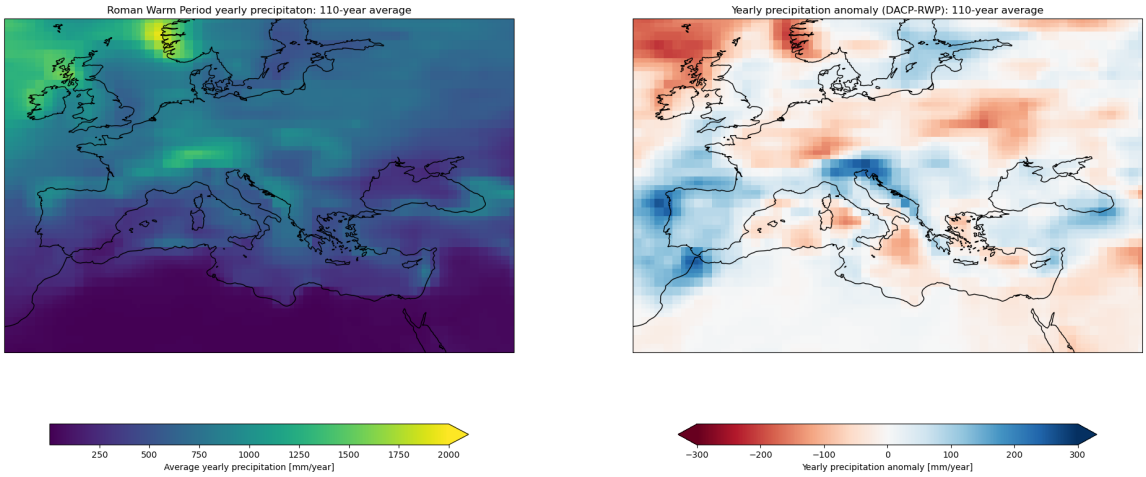

Figure 3.6: Left: 110-year average yearly precipitation reconstruction for the Roman Warm Period, obtained through a selection of days from the last century's real climate, in order to approximate the patterns observed in the paleoclimate proxies [56]. Right: anomaly in 110-year yearly precipitation averages between the DACP reconstruction and the RWP reconstruction (positive values (blue) imply a precipitation increase, negative values (red) imply a precipitation decrease).

4. Model validation

Before testing the sensitivity of the model to different climate forcings, we want to verify that its results are sensible and a reasonable approximation of the Roman trade network around 200 AD. In this chapter we analyse the model results when forced with ERA6 climate reanalysis data (Simulation 1 explained in section 3.3) and we compare them with historical/archaeological sources, discussing the possible source behind model inaccuracies. Note that the difference between the climate reanalysis data and the reconstructed climate forcing of the Roman Warm Period is not large enough to alter the most important aspects of the model results.

In section 4.1 I describe the model results for Simulation 1, while in section 4.2 I evaluate their accuracy. Conclusions on the model validity were drawn in section 4.3.

The comparison between the model results and the historical sources also allows us to produce some observations on the state of the Roman trade network itself around 200 AD, which are presented in section 4.4.

In the Appendix, we tested the sensitivity of the trade model to the hysteresis component and the number of redistribution iterations. We concluded to use 1000 iterations and to exclude the hysteresis component for these results. We found that hysteresis may help solve some of the model's inconsistencies if applied to a limited number of large cities, but it causes unstable behaviour when applied to every node.

4.1 Model results

4.1.1 Agricultural production

Figure 4.1 and table 4.1 reflect the agricultural production output of Simulation 1, which represents the starting point for the trade model.

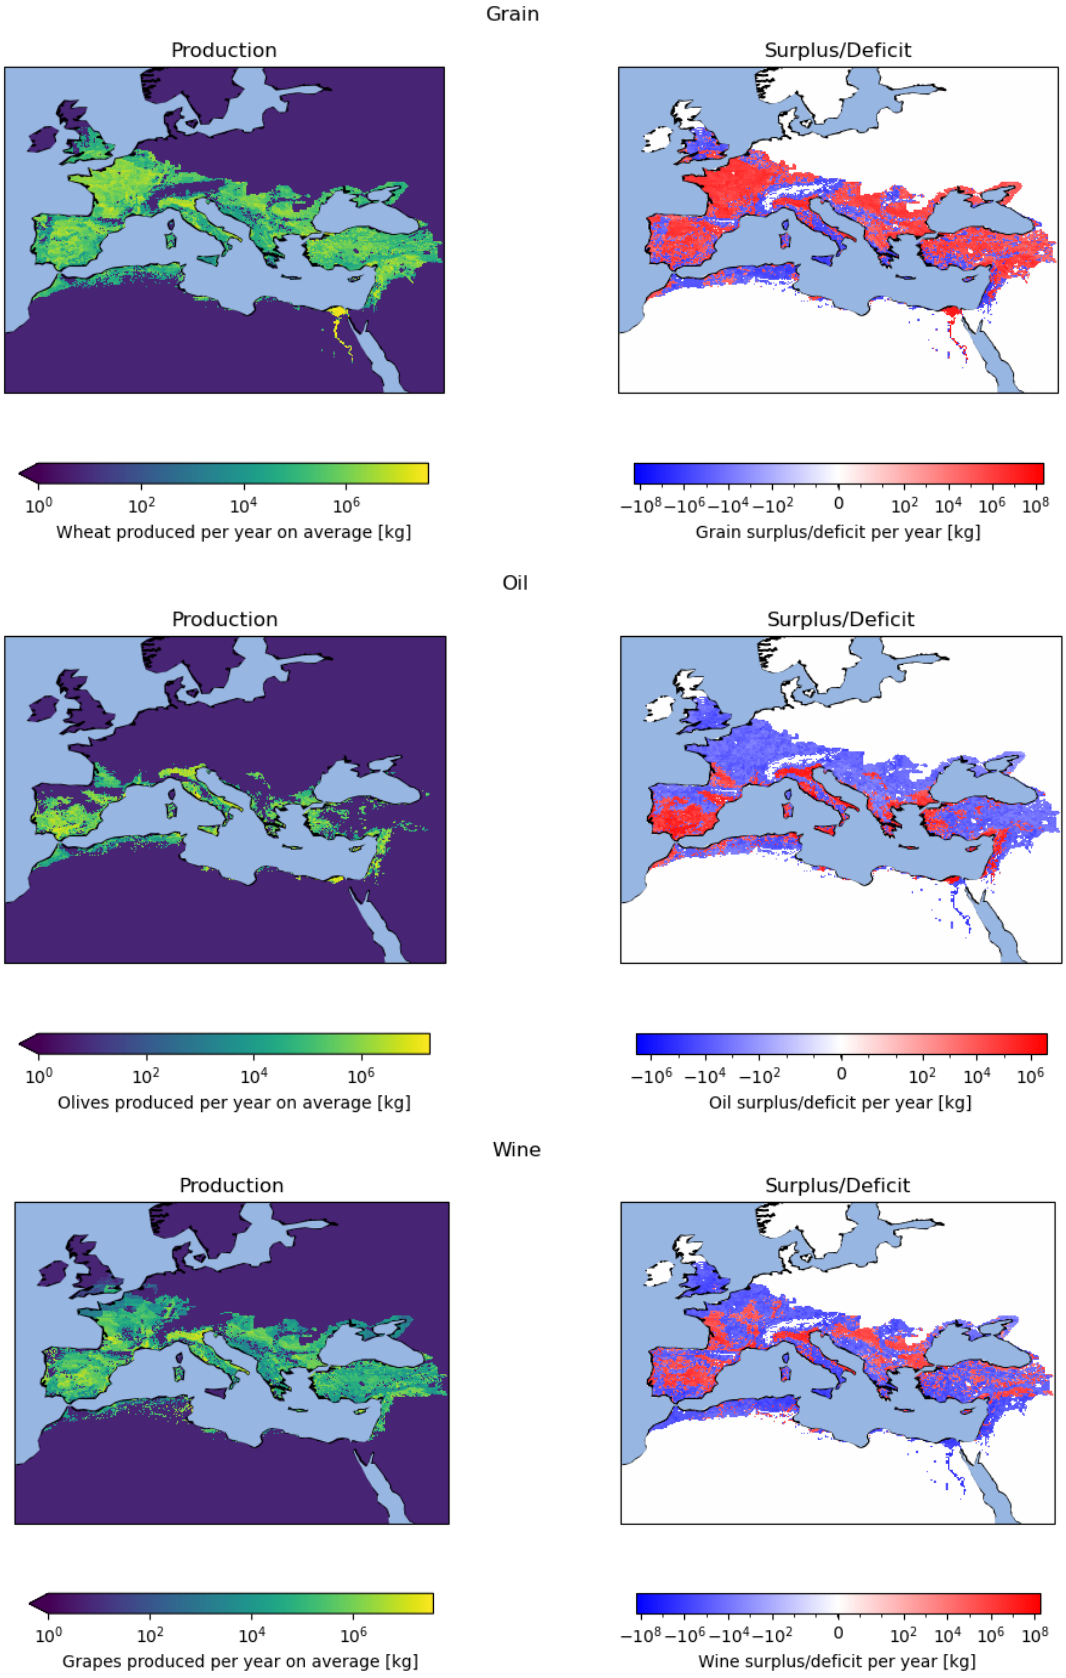

Figure 4.1: Yearly crop production (left) and production-demand (right) at 5 arcmin resolution for all products, 110-year averages,resulting from Simulation 1. On the right hand images, red represents a surplus and blue a deficit. The land is allocated according to the HYDE database. The values were limited by the available local labour.

Simulation 1: Historical climate			
	Grain [kg]	Oil [litres]	Wine[litres]
Yearly production across the whole network	$618 * 10^8$	$32 * 10^8$	$162 * 10^8$
Yearly demand across the whole network	$147 * 10^8$	$4 * 10^8$	$124 * 10^8$
Yearly surplus (production-demand) across the whole network	$470 * 10^8$	$28 * 10^8$	$37 * 10^8$
Yearly surplus as multiple of the demand (no unit)	3.2	6.9	0.3
Yearly production per region	$95.4 * 10^6$	$4.9 * 10^6$	$25 * 10^6$
Spatial standard deviation of production, time average	$192.9 * 10^6$	$10.2 * 10^6$	$54.2 * 10^6$
Temporal standard deviation of production, spatial average	$16.2 * 10^6$	$0.4 * 10^6$	$3.7 * 10^6$
Most productive region	*Zagazig, Egypt*	*Cordova, Spain*	*Toledo, Spain*
Surplus of the most productive region	$2654 * 10^6$	$75.2 * 10^6$	$586.7 * 10^6$
Region with the largest demand	*Rome, Italy*	*Rome, Italy*	*Rome, Italy*
Demand of the largest demand region	$306 * 10^6$	$8.3 * 10^6$	$257.3 * 10^6$
Region with the largest deficit	*Al − Karib, Tunisia*	*Ermopoli, Egypt*	*Ermopoli, Egypt*
Deficit of the largest deficit region	$-125.2 * 10^6$	$-4.5 * 10^6$	$-139.8 * 10^6$

Table 4.1: Table of relevant values of the relevant values of production and demand for the historical climate run. The values are 110-year time averages

Firstly, we observe that the system is demand limited for all products, meaning that the total production covers and exceeds the total demand, especially for oil. However, several regions experience a net deficit in resources and have to rely on imports; the spatial standard deviation of production is large and exceeds the temporal standard deviation in the same node.

The node with the largest demand in the system is predictably Rome for all products. The nodes with the largest deficit are regions which, while having a large demand, are also relatively isolated.

4.1.2 Trade routes

Figure 4.2 and Table 4.2 give an overview of the trade patterns arising from the model. We observe some general features.

Egypt is the largest grain exporter and most of its exports are directed towards Italy, the highly populated regions of North Africa and the Peloponnese. France, Northern Spain and the Po plane are also important grain exporters but they trade on a more local scale, with fewer long distance connections.

The oil trade patterns are dominated by exports from Southern Spain (the region known as Baetica), and in particular we observe the consistent use of the Atlantic route to supply the Britain. In general, oil trade routes move from the coasts towards inland regions.

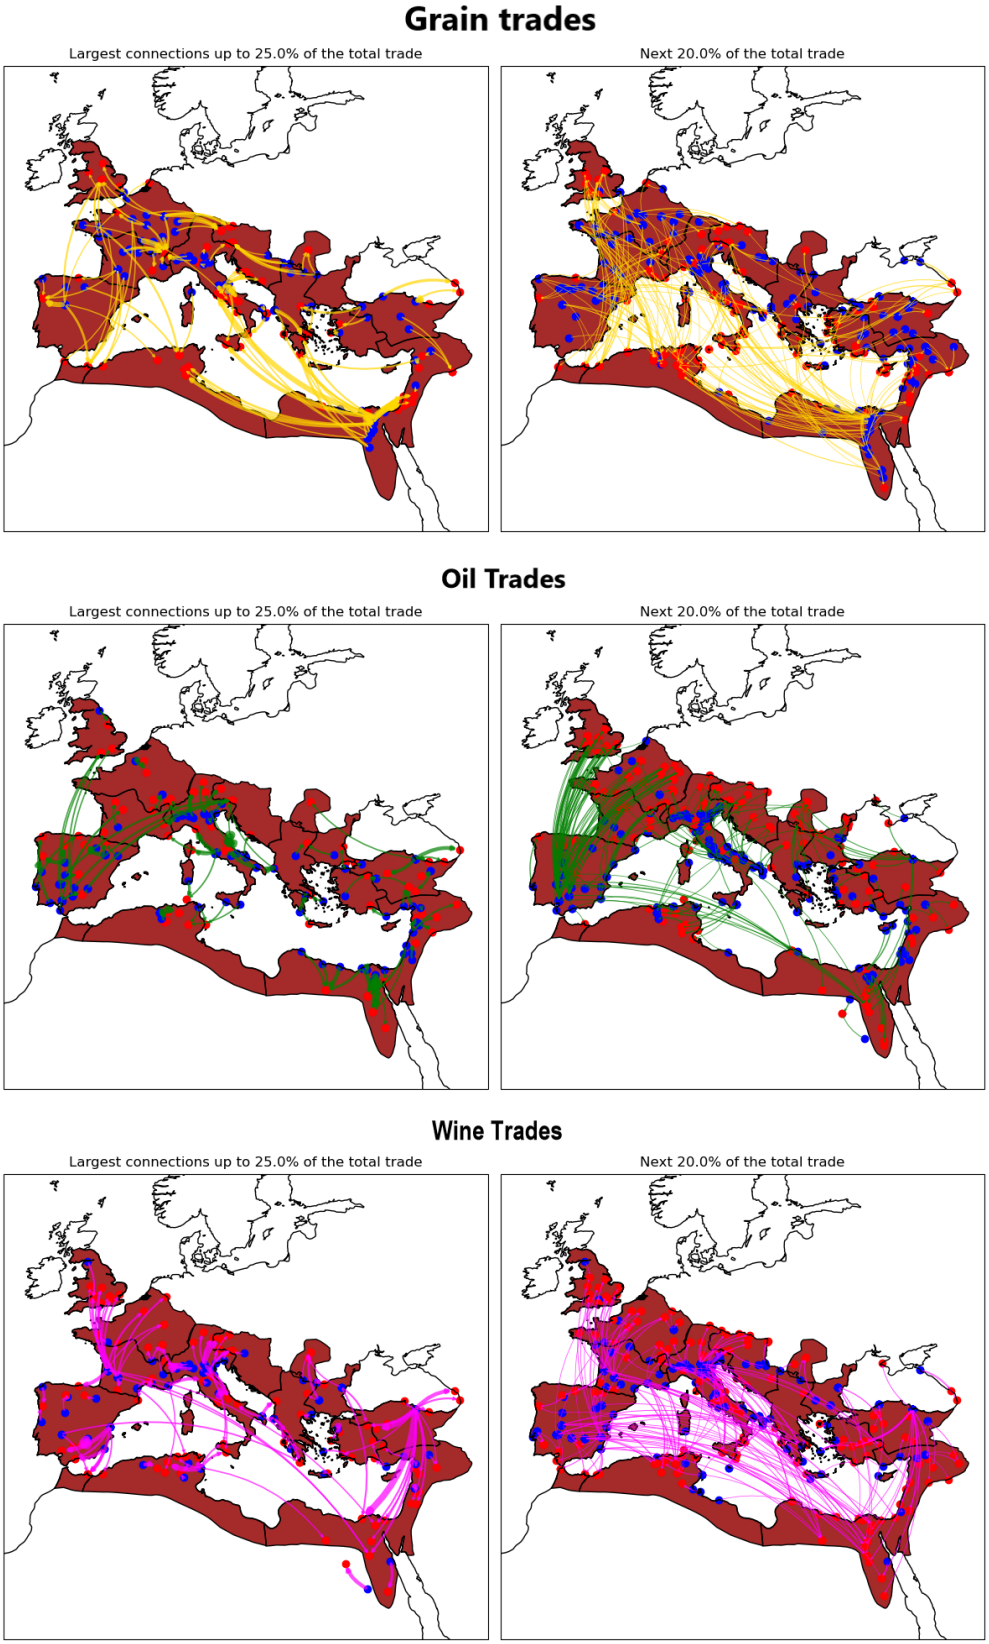

Figure 4.2: Arrow plots representing the trade matrix for Simulation 1, averaged over 100 years. The blue dots represent exporting nodes, while the red dots importing nodes. The size and opacity of the arrows represent the amount traded. For readability, the plots were limited to the largest trade connections, in such a way that the left graphs show 25% of the total trade, while the right graphs show the next 20% other than the connections plotted on the left side.

Simulation 1: Historical climate			
Trade output. 1000 iterations, no hysteresis.			
	Grain	Oil	Wine
Average import/export size	$5.07 * 10^6$ kg	$0.214 * 10^6$ litres	$10.109 * 10^6$ litres
Largest importer	*Al − Karib, Tunisia*	*Ermopoli, Egypt*	*Ermopoli, Egypt*
Import size of the largest importer	$125.202 * 10^6$ kg	$4.533 * 10^6$ litres	$139.778 * 10^6$ litres
Largest exporter	*Zagazig, Egypt*	*Cordoba, Spain*	*Toledo, Spain*
Export size of the largest exporter	$175.196 * 10^6$ kg	$3.544 * 10^6$ litres	$373.249 * 10^6$ litres
Average size of a trade connection	$175.976 * 10^3$ kg	$9.1 * 10^3$ litres	$303.3 * 10^3$ litres
Largest trade connection	*Zagazig, Egypt* *Jerusalem, Israel*	*Itay El − Barud, Egypt* *El − Idwa, Egypt*	*Puebla del Principe, Spain* *Cordoba, Spain*
Size of the largest trade connection	$20.289 * 10^6$ kg	$1.132 * 10^6$ litres	$50.748 * 10^6$ litres
Average node degree	8.9	9.75	12.1
Average import costs per unit product	3.2 sest/kg	5.1 sest/litres	3.8 sest/litres
Region with the highest import costs per unit	*Samandag, Turkey*	*Paris, France*	*Byblos, Lebanon*
Highest import cost	16.8 sest/kg	19.1 sest/kg	15.2 sest/kg

Table 4.2: Table of relevant values of the relevant values of production and demand for the historical climate run. The values are 110-year time averages

The largest wine exporting regions are southern France, northwestern Spain, the Po plane and the Black Sea Turkish coasts. The most well established trade routes are those supplying Britain from Southern France, and Egypt from the Turkish Black Sea. The Po Plane exports mostly travel short distance right across the Alps. We also notice surprisingly large long distance exports from the western Mediterranean towards Egypt.

4.1.3 Trade over cost distance

As we can see in figure 4.2, long distance trade is relatively common in our model results, even with a large overall grain surplus. We are interested in evaluating how the trade system deals with distance. As described in the last chapter, each of the importing nodes is made to contact their neighbours in order of cost-distance, until its demand is satisfied. In the following, the term distance refers to cost distance (cost of transport per kg of product) rather than geographical distance.

Figure 4.3 shows on the left the time-average amount of product traded between any two nodes in relation to the cost-distance between them. The traded amount falls linearly with distance for grain trades, with some outliers related to large importers. For oil and wine trades, the traded amount falls very rapidly over the short distances (0 to 3 sest/kg), and afterwards semi-linerarly.

On the right, the figure shows the cost distance distribution plot of all

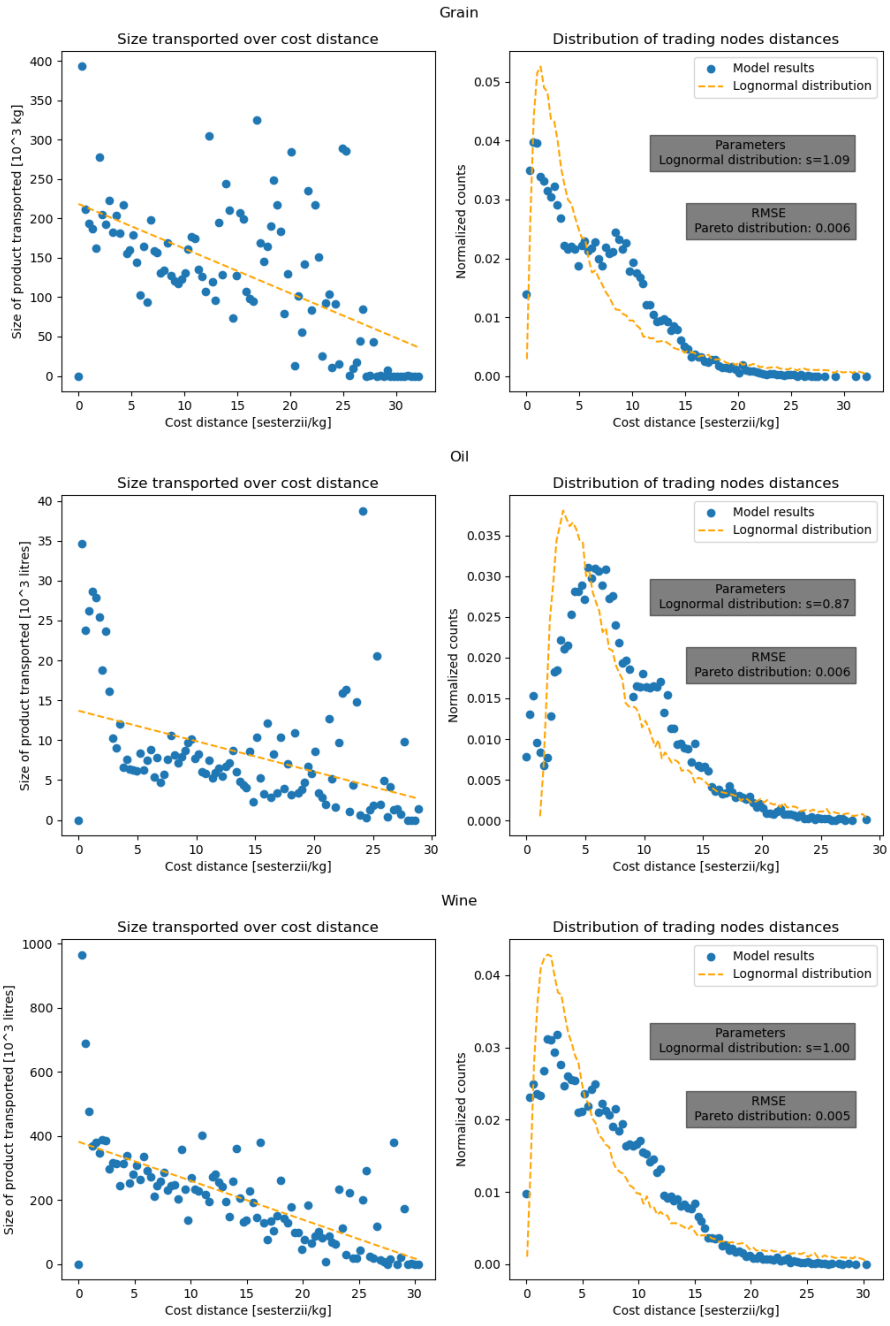

Figure 4.3: On the left, time average amount of traded product over cost distance. On the right, the blue dots represent the cost distance distribution of all the active trade connections as resulting from the model output, while the orange dashed lines are randomly drawn data from a fitted log-normal distribution.

active connections, alongside data drawn from a fitted lognormal distribution. Although not a perfect fit, a lognormal distribution accurately reflects the general structure of the model output for all products.

Most active trade connections develop on the short to mid distance. The distribution peaks around 2 sest/kg for grain, around 3 sest/kg for wine and around 5 sest/kg for oil (for reference the distance between Rome and Naples is estimated in the model at about 3 sest/kg).

At the same time, we observe an heavy tail of long distance connections, for every product but particularly for grain trades. Long distance trade is a necessity in the model to ensure supply to some highly urbanised regions.

4.1.4 Import costs and node degree

Figure 4.4 shows maps of the time average import cost per unit product and maps of the node degree of each hinterland. Import cost per unit product is defined as the total cost of importation divided by the import size of each node, or in other words the average cost-distance travelled by the product. For the rest of the chapter, I will use "import costs" and "import costs per unit product" interchangeably. Node degree refers to the number of active connections for each node (only importing nodes are shown).

We observe how the regions with the highest import costs per unit tend to the situated inland, and isolated from the producing regions. Other more densely populated regions with higher import size (for instance Rome and coastal Tunisia) manage to maintain their import costs relatively low thanks to their geographical position, by increasing their node degree. In fact node degree correlates more closely to import size.

Figure 4.5 shows the relation between import cost per unit, node degree and import size. We observe, expectedly, a positive correlation between all three variables. However, import costs also show a significant vertical spread at relatively low import size values, so that the nodes with the highest import costs are not necessarily the largest importers. Instead, the largest importers in the network are able maintain their import costs relatively low by diversifying their providers, thanks to their geographical position (geo-

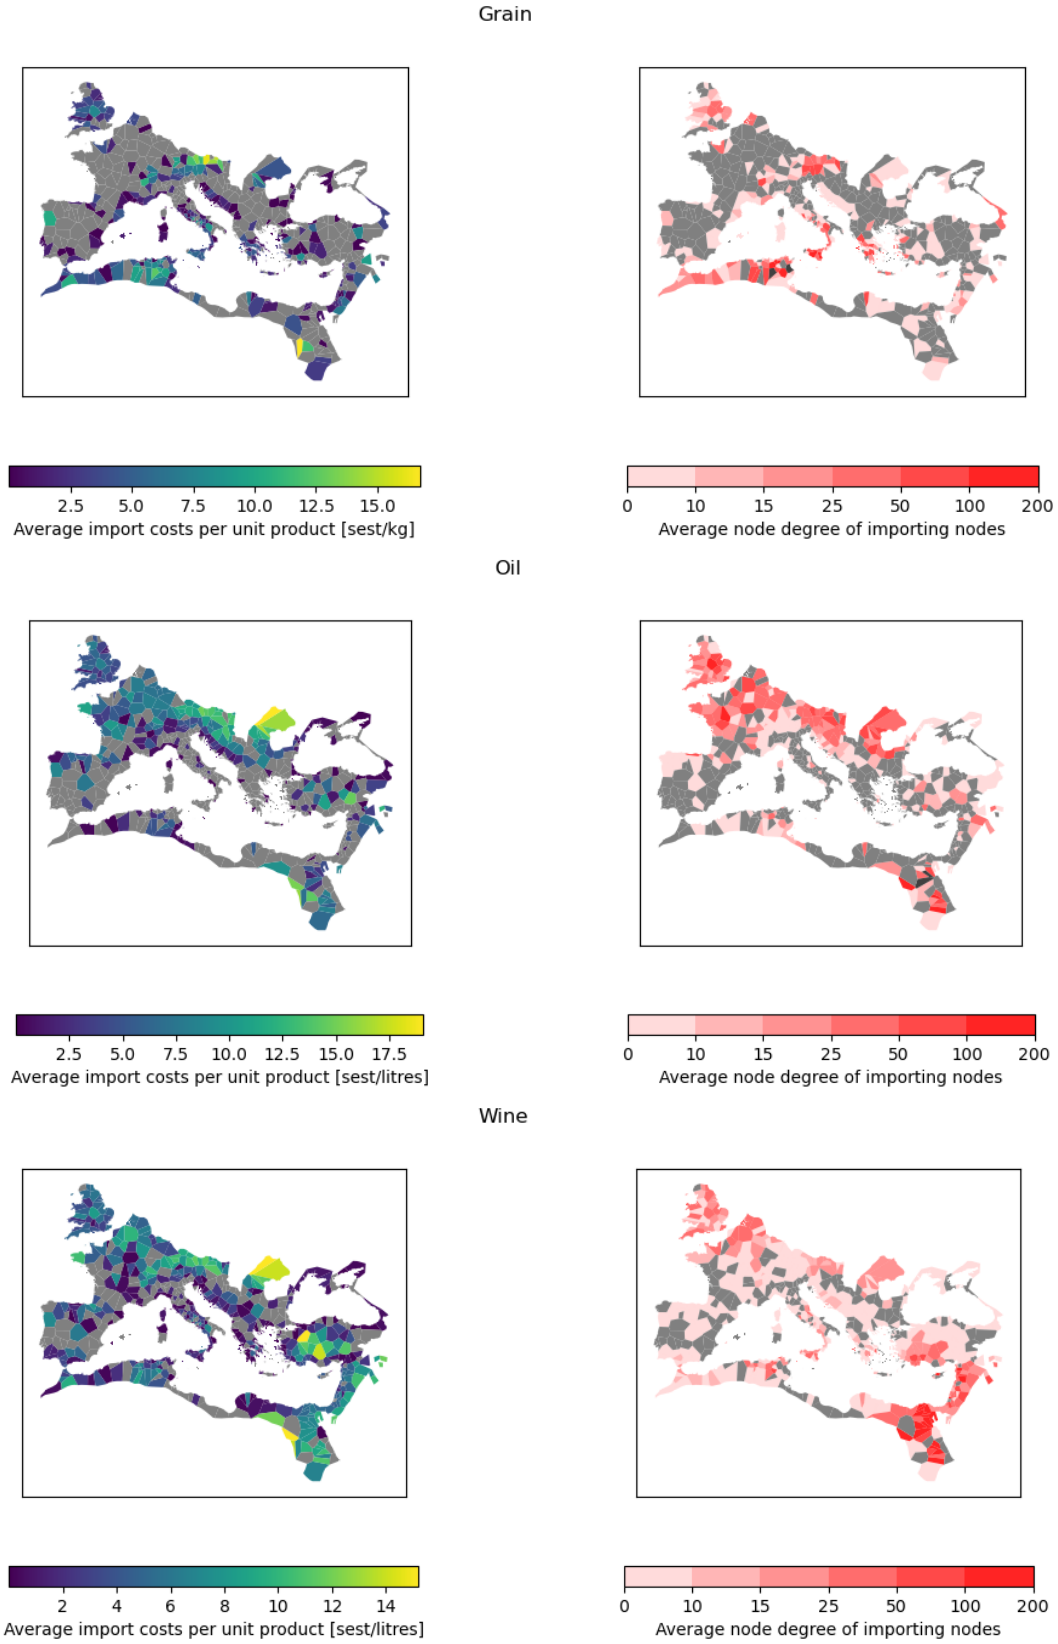

Figure 4.4: Hinterland maps of the time-average import cost per unit product (left) and node degree (right), for Simulation 1. The grey regions are net exporters. We observe that the regions with the largest import costs are isolated or distant from the sea, not necessarily with a very large import size

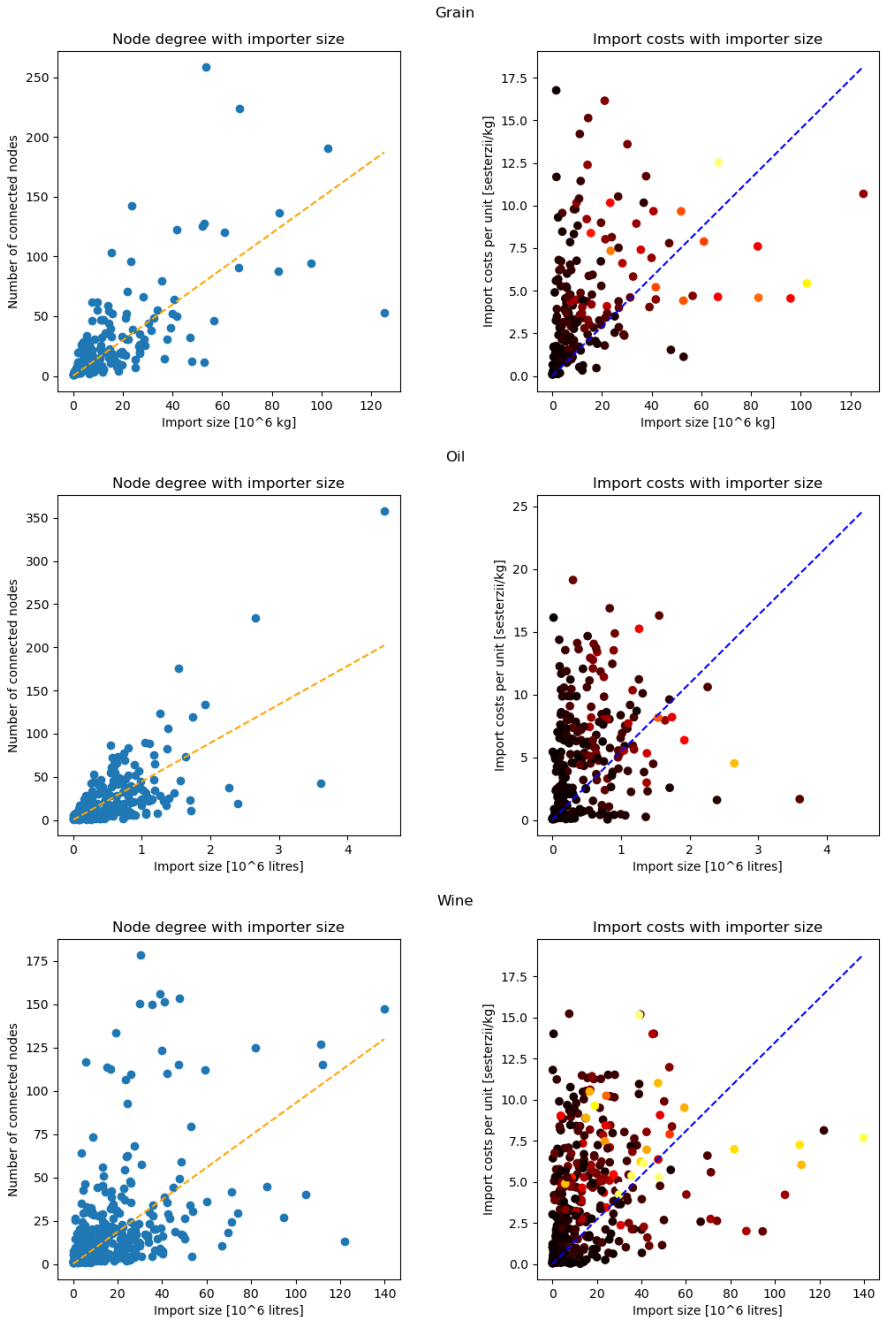

Figure 4.5: Left: Scatter plot of the time average node degree of each node over its time average import size, along with a linear fit. Right: Scatter plot of the time average import costs per unit product of each importing node over its time average import size, along with a linear fit. The lighter colors on the right plot represent nodes with a larger average node degree.

graphical center of the Empire, close to the sea).

4.1.5 Export size distribution

In order to gain information about the market structure, I evaluated the distribution of export sizes across the network. Export sizes refer here to the total product exported by each node, on average for the whole 110 year period.

Figure 4.6 shows the distribution of export sizes resulting from the model output along with random data drawn from a fitted Pareto distribution. I found that in all cases the distribution replicates well the general structure of the model output. The probability density function of a Pareto law distribution is defined as:

$$f(x, b) = b/x^{b+1}$$

Pareto distributions, which often naturally emerge in free market economic systems [72, 47], describe very unequal market structures, in which a small number of very large actors take up the majority of the market. As a general rule, the Pareto principle claims that the largest 20% of actors will tend to control 80% of the market [72]. In our model, this principle is roughly respected, as the largest 20% of the exporters makes up about 70% of the total export.

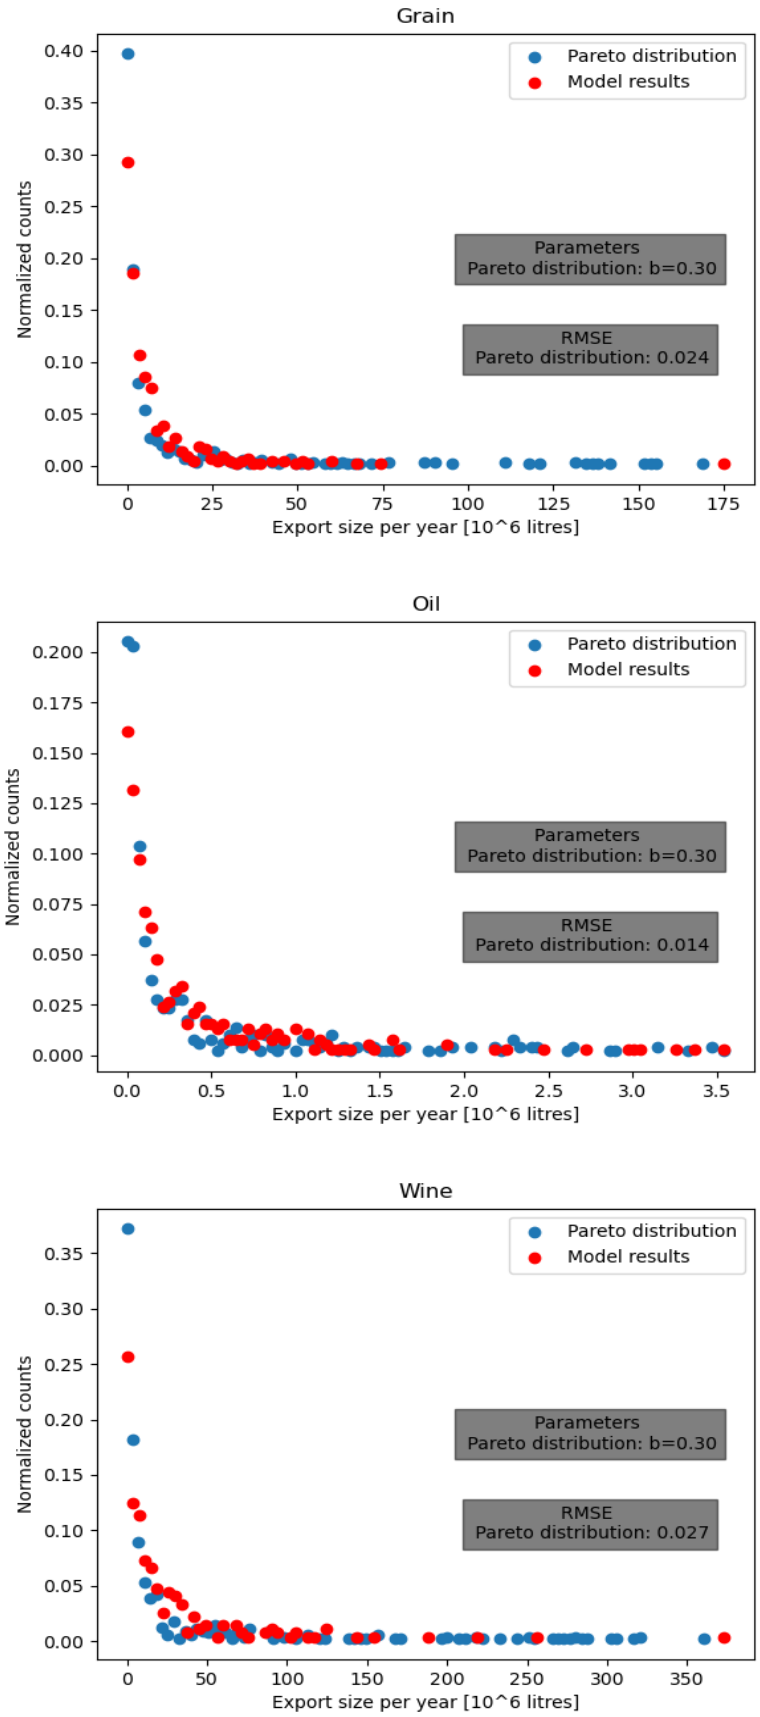

Figure 4.6: Distribution of the time-average export sizes of the model output fitted with a Pareto power law. Red dots: distribution of the time-average export sizes of each node. Blue dots: data randomly drawn from a fitted Pareto distribution.

4.2 Comparison with historical and archaeological sources

We now compare the model results I described with the available historical and archaeological record. Let us note that such validation sources are sparse and their nature is often qualitative, consequently we can only validate some of the model's results. In other cases, we can only assess whether the results are reasonable in relation to present day trade networks data.

4.2.1 Trade routes

4.2.1.1 Grain trades

As is the case for most pre-industrial societies, grain was produced and consumed locally whenever possible [46, 45]. In the model results, the distribution of grain trades over distance peaks at a lower value than the same distributions for oil and wine (see figure 4.3). So that this condition is overall verified.

However, due the rising urbanisation rates of some regions of the Empire, in particular Italy and the Central Mediterranean, long distance grain shipments were a necessity [43, 29, 11].

Archaeological evidence as well as primary written sources show the existence of a few large exporters, that supplied those regions where demand was higher than production. Most of these large shipments developed on sea routes rather than land routes, a feature which we accurately reproduce in the model [46].

The large exporters supplying Italy and the central Mediterranean are usually identified by contemporaries as Sicily, Sardinia, Northern Lybia and Egypt [45, 46, 44].

According to our model, Egypt is indeed the largest grain exporter, although it is rivaled by France. Egypt, however, is better connected than France to Italy, due to the comparative easiness of sea travels, and it is the main exporter of grain to Rome. Conversely, trade exports from France

develop on a more local scale, which may not leave trace in the historical record.

Northern Lybia and Sardinia both appear as grain exporters in our model, but their connection to Italy is not as significant as we would expect.

Sicily appears in the model as a net grain importer rather than exporter. This is definitely an inconsistency with the historical record, and it is likely related to the choice of land allocation. Before the occupation of Egypt, Sicily essentially functioned as an agricultural colony, with urbanisation rates significantly lower than the rest of Italy [41, 73]. Thus, the historical land allocation would have likely been different from the HYDE allocation, which is derived from the modern day distribution of agricultural land.

Finally, we observe a possible inconsistency in the overestimation of agricultural production in the Rome city hinterland. Rome, as the city the largest population, has by far the region with the largest grain demand in the network, twice the second largest city, Ermopoli, on the Nile River. However, its deficit, when accounting for production within the city hinterland, is only the third in the network, less than the smaller (though still sizable) cities of El-Karib and Agrigento.

4.2.1.2 Oil trades

Oil trades didn't leave as much of a trace in the written historical record as grain trades, which were necessary for several regions' survival. However, useful information can be derived from the *amphorae* excavated across the former Empire. These containers were used primarily for oil, and it is generally possible to determine where a certain kind was produced, leading back to the exportation site [49].

The most common origin site among the excavated amphorae is the province of Baetica, in southern Spain [49, 48]. According to the archaeological record, Baetica was the largest oil exporter in the Empire, rivaled by some north African provinces, in modern day Lybia and Tunisia [49, 48]. The latter appear to have grown in importance around the third century AD [49].

Our model accurately reproduces this, with Southern Spain appearing as by far the largest exporter.

Further, there is archaeological evidence for a well established logistical network that supplied the northern borders, Germany and Britain, with Baetican oil [49, 48]. Recent research also suggests that this trade employed mostly the Atlantic route rather than the Rhine river route, as it was previously thought [48]. This is verified quite clearly in our model, with the Atlantic route being among the largest connections in the network (figure 4.2).

However, the importance of North African oil seems to be underestimated in the model results. This is likely to be the results of our land allocation choices as in the case of Sicilian grain. It is also possible that exports of oil from Tunisia rather than southern Spain were guided by other factors such as the comparative stability of the regions especially in the 3rd century AD [48, 49]. This possibility requires the use of an updated land allocation model to be verified.

Aside from the spatial pattern, the surplus in overall oil production appears to be exaggerated in the model results. A degree of overproduction would have been necessary to offset transport or storage losses, but the factor 7 shown in table 4.1 seems excessive. Here, too, a land allocation based on the reconstructed demand could help solve the issue.

4.2.1.3 Wine trades

The largest wine exporters according to historical and archaeological evidence would have been southern France, Italy, Northeastern Spain, Tunisia, and, to a smaller extent, the Black Sea Coast in the Eastern Mediterranean [74, 51, 42, 75].

Southern France and Italy appear to have been the primary competitors, and southern France covered the specialized role of supplying the northern borders through the Rhine river route [43]. The trade of Turkish wine from the Black Sea coasts was probably limited to the Eastern Mediterranean [75].

All of these regions are identified as large exporters in the model, with

the exception of Tunisia, where there seems to be issues with the land allocation.

The historically well known route that provided Britain and the Northern borders with wine from southern France is observed in the model results. France is the largest actor and Italian wine is mostly exported over short distances.

This underestimation could be due to the fact that the exportation of wine from Italy to the provinces was somewhat artificial, either due to the perceived quality of the product, or because of state sponsoring [51].

The expansion of wine production in Italy could also be a consequence of the economic pressure of "globalisation" [43]. With the Empire's territorial expansion, Italy received a strong influx of slave labour (who accounted for up to 30% of the population in the early first century AD [76]), which was usually employed in wine production more than in cereal production. At the same time, the regular and abundant imports of Egyptian grain drew down the prices, making cereals not as profitable both for landowners and for workers, who couldn't compete with slave labour [43, 50, 74].

We observe as expected the exportation of wine from the Black Sea Turkish coast towards the rest of the Eastern Mediterranean.

Finally, in the model we observe very long distance exports towards Egypt, from the western Mediterranean. While Egypt was not a wine producer, this size of trade over such long distances appears unlikely. I could not find any evidence of this trade route. This anomaly in our results could represent a failure of the equal consumption assumption. For a long period before the Roman expansion, wine was not a staple of the majority's diet, but a rather elite good [50]. The downward social expansion in wine consumption was concurrent to the increased production and lowering of prices in the western provinces [50]. If this did not happen in Egypt, due to the regions' unsuitable conditions, wine consumption might have been limited to the richer classes. Beer consumption might have replaced wine consumption in several regions, so that the equal consumption assumption for wine likely needs to be revised.

4.2.2 Trade over distance

In figure 4.3 on the right we observed that in our model the importer-exporter couples tended to form as a lognormal distribution over distance. That is, most of the connections take place at a relatively small distance, but with a heavy tail of long distance trade. This structure of trades is often verified in modern day trade network [77, 78]. In the case of grain trades the distribution peaks at a lower distance, which appears reasonable due to the nature of the product.

While we do not have any explicit validation source on the distribution of trades in the late Roman Period, it is commonly acknowledged that long distance trade of basic commodities was a necessity for densely populated regions [41, 46]. So this "modern" distribution of trades appears reasonable also for this time period.

This result provides an important confirmation of what was mentioned in the literature review chapter: the urbanisation landscape of the Empire was such to require at least some trade on the very long distances.

In the same figure, we plotted the average size traded over each cost distance interval (independently from the number of pairs). We observe that the size traded falls with distance, which is a commonly known feature of most present and past trade networks [77, 78].

4.2.3 Import costs and node degree

In our model results, we observed a dichotomy between large, well connected cities, who maintain relatively low import costs per unit despite their large import size, and isolated cities, with mid to low import size but very large import costs per unit.

According to some theories, the gradual decrease in long distance trade starting from the second century AD was also a consequence of peripherization and increased local production [42]. This result sounds reasonable within this context: if some regions showed very high import costs per unit due to their distance from large producers, this would have encouraged lo-

cal production where possible.

However, the very high node degree of some larger cities (such as Rome, with over 200 different importers), appears a little exaggerated. Our model is likely overestimating the flexibility in trade towards there centers. Likely, each city would have gravitated towards a limited set of importers, creating institutional lock-in over time. This process could be emulated through an hysteresis component (see appendix B), though our attempt in this direction did not achieve the result.

4.2.4 Export size distribution

In section 4.1.5, I evaluated the distribution in the size of exporting regions, as it emerges from our model.

It is possible to utilize archaeological evidence to verify the emergence of this distribution.

As described in the last section, oil was primarily transported in ceramic containers called amphorae. Each container includes a stamp which was used for taxation purposes. Although it is unclear whether the stamps served to identify the producer or the transporter, it is likely that repeated codes referred to the same trade route and the same wider exporting region [49, 47]. Thus, it is possible to use the number of repetitions of the same code as a proxy for the trading power of a region.

Campillo et al. [47], reconstructed the export size distribution according to this method for the sole city of Rome, and reached the conclusion that a Pareto distribution was the best approximation of the archaeological data. I repeated the same test using a wider database, also provided by the same authors in a different paper [48], and compared the archaeological evidence with our model results, as seen in figure 4.7. I find the data to be in good agreement.

This agreement constitutes a validation of our model results, at least in the general structure of the market for oil trades. In our model, we obtained similar results for the export size distribution of grain and wine trades,

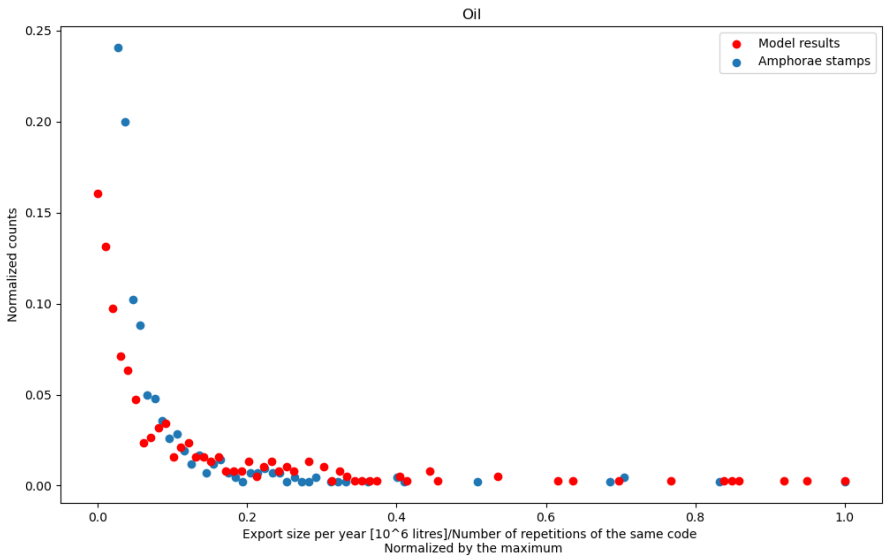

Figure 4.7: Comparison between the export size distribution emerging from our model (blue dots), and the number of repetitions of the same codes on amphorae excavated across the form Empire's territory (red dots). The dataset was made available by Campillo et al. in their published paper [47]. We removed from the dataset the code that repeated themselves less than 10 times, as they may relate to archival errors or excavation biases.

which do not have similarly accessible proxies. However, the Pareto distribution is a common structure for most free market systems [72], so that our results appear reasonable if not accurate. It is possible that the case of grain trades was different because of more intense state intervention in the market [44, 36].

Further, our results supports Campillo et al.'s thesis that the oil trade market was largely unregulated, with little state intervention [47].

4.3 Conclusions on model validity

Overall, we find the model results to constitute a reasonable approximation of the Roman trade network around 200 AD, and may be used to measure its sensitivity to the climate state. However, we have observed some inconsistencies as a consequence of some of the model assumptions.

- Inconsistencies in the agricultural production are likely a consequence of the HYDE land allocation database. Grain allocation appears def-

initely underestimated in Sicily, and oil allocation appears underestimated in North Africa. Grain allocation may be overestimated in the hinterland of Rome.

- The equal consumption assumption is likely not fully accurate for wine. In our model, this limitation causes heavy long distance trade from the western Mediterranean towards Egypt, which appear unlikely and are not verified in the historical record. Several regions would have substituted wine with beer, or wine may have been still a luxury good.
- The trade in the model are more flexible than a real trade network, which results in some large cities having a very large number of trade partners.

4.4 Observations on the state of the Roman trade network around 200 AD

As described in the literature review chapter, the Roman trade network may have been affected by some structural elements of vulnerability which increased its sensibility to climate change. The fact that some of these vulnerabilities emerge also in our results, casts them as natural consequences of the urbanisation and productivity landscape of the period, since these are the only factor we accounted for in the model. The vulnerabilities which our model results show can be summarised in:

- The intense urbanisation rates make long distance trade a necessity to ensure food supply (lognormal distribution of cost-distances). This condition, while relatively common in modern systems, would have been anomalous in a pre-industrial society, and rendered large cities vulnerable to disruptions of the trade system. In fact, most of the food shortages recorded in Rome during the imperial period, are blamed by contemporaries on such disruptions (war, piracy or sea conditions, for instance) [1].
- The market structure appears very unequal, with few very large exporters occupying most of the market. This obviously constitutes a

vulnerability as it exposed a large number of importers to changes in the productivity of a single exporter.

- The import costs vary a lot across the Empire, with some isolated regions having very large import costs per unit, while not being so small to be insignificant. This peripherization of human resources has been cited as one of the systemic issues of the late Empire period [42, 30].

5. Results: Roman Warm Period and Dark Ages Cold Period climates

In this chapter I describe the results of the simulation when forced with the reconstructed climates of the Roman Warm Period and Dark Ages Cold Period. These results refer to Simulation 2 and 3, of the list in 3.3.

Our intention is to compare the effect of a different climate forcing on the agricultural production as well as on the trade patterns, in order to evaluate the overall stress on the system.

5.1 Agricultural production

Figures 5.2 and 5.1 compare the climate forcing used for the two simulations, reconstructed as described in section 3.3

Figure 5.3 shows the maps of the absolute and relative anomaly in crop production introduced by the Dark Ages Cold Period (Simulation 3) on the Roman Warm Period (Simulation 2) (DACP-RWP). We do not provide maps of the Roman Warm Period productivity as they look very similar to those of the Historical Period simulation shown in the last chapter. Table 5.1 gives an overview of the yearly agricultural output for the two simulations.

The anomaly in time-average total productivity summed over the whole territory of the Empire (first line of in table 5.1) is relatively small for all three products. We notice a slight increase in olives productivity in the Dark Ages Cold Period compared to the Roman Warm Period, while grapes productivity declines slightly. Wheat productivity experiences the largest loss, but still quite limited.

Local changes in average production are significant in absolute value (figure 5.3), up to the order to 10^6 kg yearly.

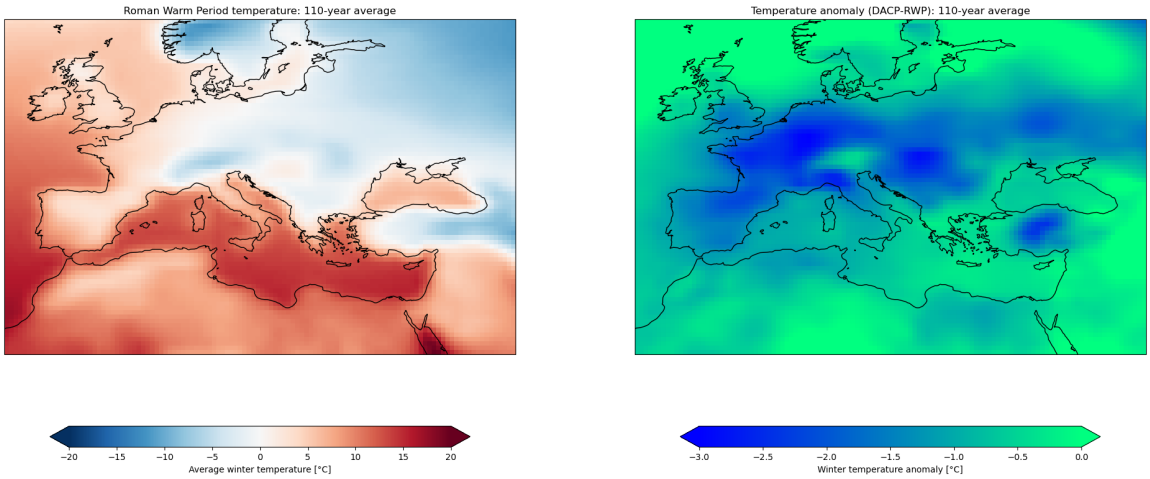

Figure 5.1: Left: 110-year average temperature reconstruction for the Roman Warm Period, obtained through a selection of days from the last century's real climate, in order to approximate the patterns observed in the paleoclimate proxies [1, 12]. Right: anomaly in 110-year temperature averages between the DACP reconstruction and the RWP reconstruction (positive values imply a temperature increase, negative values imply a temperature decrease).

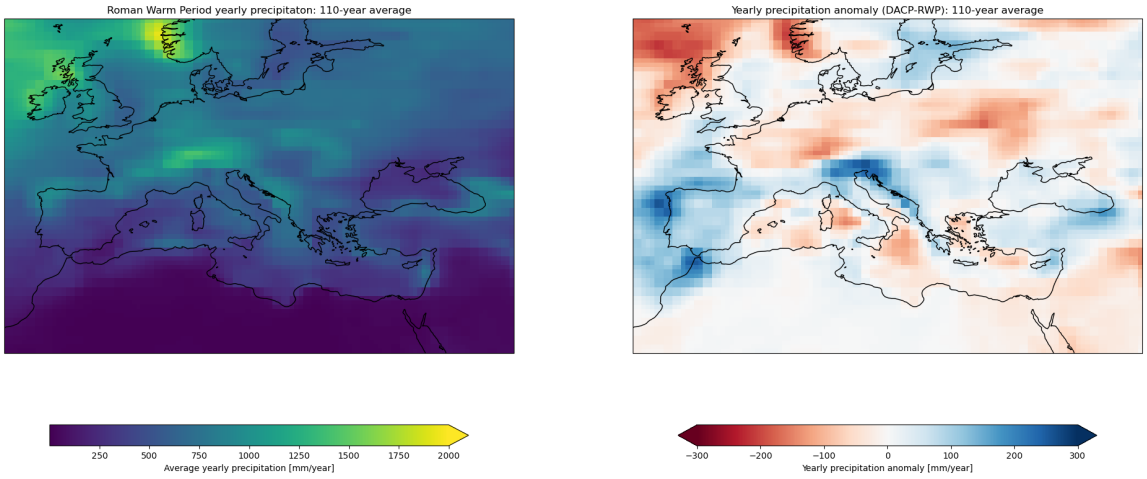

Figure 5.2: Left: 110-year average daily precipitation reconstruction for the Roman Warm Period, obtained through a selection of days from the last century's real climate, in order to approximate the patterns observed in the paleoclimate proxies [56]. Right: anomaly in 110-year daily precipitation averages between the DACP reconstruction and the RWP reconstruction (positive values imply a precipitation increase, negative values imply a precipitation decrease).

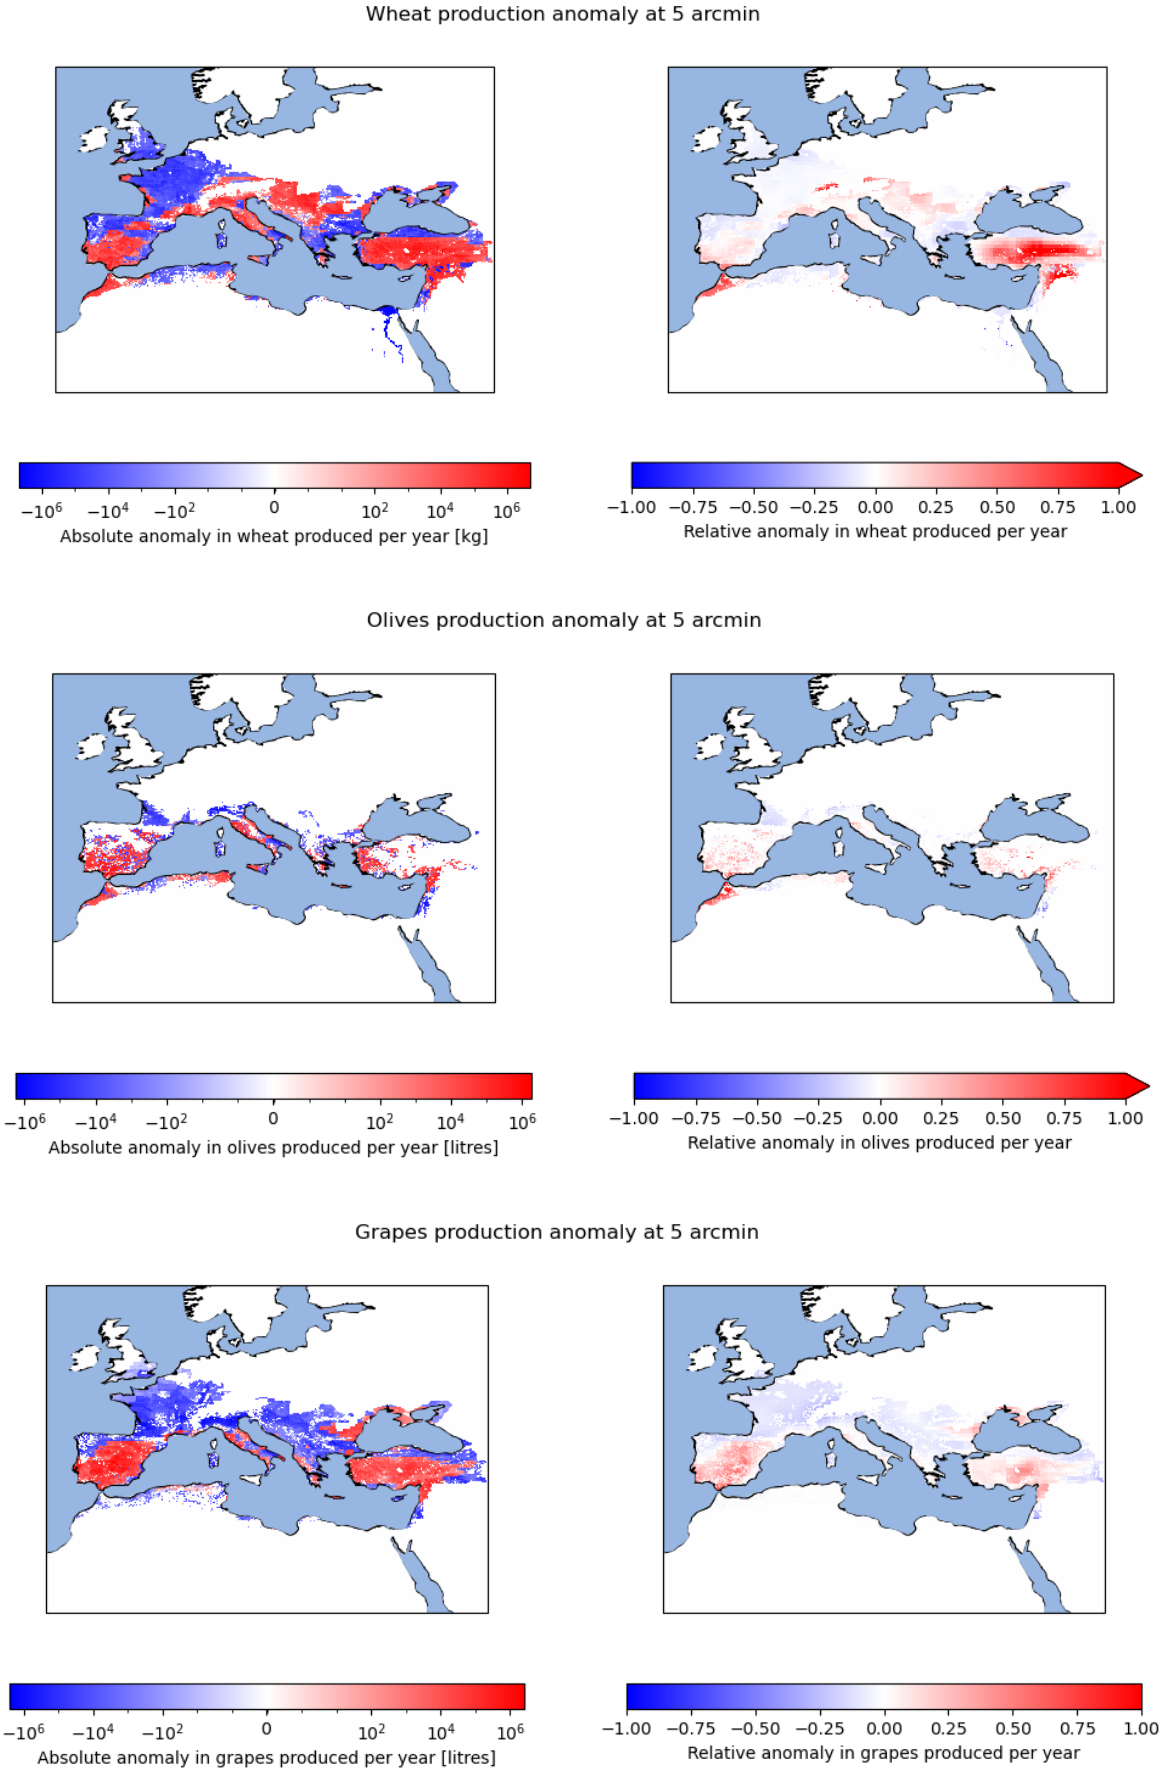

Figure 5.3: Anomaly in time average yearly crop production between the Dark Ages Cold Period and the Roman Warm Period. Absolute anomaly (DACP-RWP) on the left, and relative anomaly ((DACP-RWP)/RWP) on the right. Red stands for an increase in productivity, blue for a decrease.

Simulation 2: Roman Warm Period climate			
	Grain [kg]	Oil [litres]	Wine[litres]
Yearly production across the whole network	$614 * 10^8$	$32 * 10^8$	$162 * 10^8$
Yearly demand across the whole network	$147 * 10^8$	$4 * 10^8$	$124 * 10^8$
Yearly surplus (production-demand) across the whole network	$466 * 10^8$	$28 * 10^8$	$37 * 10^8$
Yearly surplus as multiple of the demand (no unit)	3.2	7	0.3
Yearly production per region	$94.8 * 10^6$	$4.9 * 10^6$	$25 * 10^6$
Spatial standard deviation of production, time average	$197.1 * 10^6$	$10.3 * 10^6$	$52.7 * 10^6$
Temporal standard deviation of production, spatial average	$12.9 * 10^6$	$0.3 * 10^6$	$2.5 * 10^6$
Most productive region	*Zagazig, Egypt*	*Cordoba, Spain*	*Toledo, Spain*
Surplus of the most productive region	$2797.6 * 10^6$	$80 * 10^6$	$483 * 10^6$
Region with the largest demand	*Rome.Italy*	*Rome, Italy*	*Rome, Italy*
Demand of the largest demand region	$306 * 10^6$	$8.3 * 10^6$	$257.3 * 10^6$
Region with the largest deficit	*Al − Karib, Tunisia*	*Ermopoli, Egypt*	*Ermopoli, Egypt*
Deficit of the largest deficit region	$-126.7 * 10^6$	$-4.5 * 10^6$	$-139.8 * 10^6$

Simulation 3: Dark Ages Cold Period climate			
	Grain [kg]	Oil [litres]	Wine[litres]
Yearly production across the whole network	$612 * 10^8$	$32 * 10^8$	$161 * 10^8$
Yearly demand across the whole network	$147 * 10^8$	$4 * 10^8$	$124 * 10^8$
Yearly surplus (production-demand) across the whole network	$463 * 10^8$	$28 * 10^8$	$36 * 10^8$
Yearly surplus as multiple of the demand (no unit)	3.1	7	0.3
Yearly production per region	$94.4 * 10^6$	$5 * 10^6$	$24.9 * 10^6$
Spatial standard deviation of production, time average	$185.6 * 10^6$	$10.7 * 10^6$	$53.6 * 10^6$
Temporal standard deviation of production, spatial average	$13.4 * 10^6$	$0.4 * 10^6$	$3 * 10^6$
Most productive region	*Zagazig, Egypt*	*Cordoba, Spain*	*Toledo, Spain*
Surplus of the most productive region	$2503.3 * 10^6$	$85.9 * 10^6$	$600.5 * 10^6$
Region with the largest demand	*Rome.Italy*	*Rome, Italy*	*Rome, Italy*
Demand of the largest demand region	$306 * 10^6$	$8.3 * 10^6$	$257.3 * 10^6$
Region with the largest deficit	*Al − Karib, Tunisia*	*Ermopoli, Egypt*	*Ermopoli, Egypt*
Deficit of the largest deficit region	$-127.8 * 10^6$	$-4.5 * 10^6$	$-139.8 * 10^6$

Table 5.1: Table of relevant values of production and demand for the Rowan Warm Period and Dark Ages Cold Period climate runs. Unless otherwise specified the values are 110-year time averages

We observe that all three products show similar spatial patterns of yield anomaly. The maps in figure 5.3 show an increase in productivity in Southern Spain, Central Italy, and most of Anatolia except for the Black Sea coasts. Conversely, Northern Spain, most of France, Germany and the Nile Valley experience loss in average yields.

The production anomalies for grain largely correlate with the patterns of precipitation anomaly between the two periods (figure 5.2), with production lowering in regions shifting towards a drier climate.

This is mostly true of grapes and olives as well, except for the Po Plane and the Balkans, where grapes and olives production decreases while grain production increases. Since these regions shifted towards a wetter and colder climate in the Dark Ages Cold Period, this difference may indicate that olives and grapes, as opposed to grain, were temperature-limited and not precipitation-limited here.

Despite the overall similar yield anomaly patterns, the three products are affected differently due to their different land allocation.

Wheat shows the largest absolute decrease in total production, since it was intensely cultivated in the Nile Valley and France.

In the case of olives, the increase in total production is due to the increased productivity of Southern Spain, Northern Tunisia and the Aegean Turkish coast.

Grapes total production decreases slightly due to the fall in yields in France. The favorable yields in Spain cause the increase in the largest producer's size shown in table 5.1.

These quite large absolute anomalies translate into very limited relative change in the largest producers, whose yearly production is in the order of 10^8 kg. The only regions to experience large relative change are small producers, as well as Central Anatolia and Syria, whose grain productivity is almost doubled.

5.2 Anomaly in the time average import and export patterns

We first evaluate the anomaly in the time average output of the trade model for the two periods. The parameters used in the trade model here are the same used in the last chapter, so 1000 iterations and no hysteresis component.

Figure 5.4 shows the change introduced by the Dark Ages Cold Period in comparison to the Roman Warm Period. The export size anomalies are coherent with the production anomalies.

We notice a significant decrease in grain exports from Egypt and France. The loss is countered by increased exports from central Turkey and Southern Spain. However, the decrease in export size from the historical producers (Egypt and France), it's always at least one order of magnitude lower than their export sizes during the Roman Warm Period. This is explained also through figure 5.5, in which I plotted the change in export size of each hinterland versus their export size in Simulation 2 (RWP). We see here that the loss in surplus production increases linearly with the producer's size, in the case of grain, so that the relative loss is always in the order of 10% for these larger producers. We do not observe a significant change in the long distance grain shipments from Egypt (Egypt and France remain the largest grain exporters even if their export size decreases). We note that grain production in the Roman hinterland increases so that its import size declines in the DACP simulation.

Oil exports from Southern Spain increase while those from Italy declines. In figure 5.5 we notice the opposite trend for oil than for grain, although with a much larger spread. The larger producers show larger growth in their export size, still in the order of 10% or less

Wine exports from Spain grow as opposed to southeastern France. In figure 5.5 we see that the changes in wine exports are the largest in relation to the export size during the Roman Warm Period. At the same time,

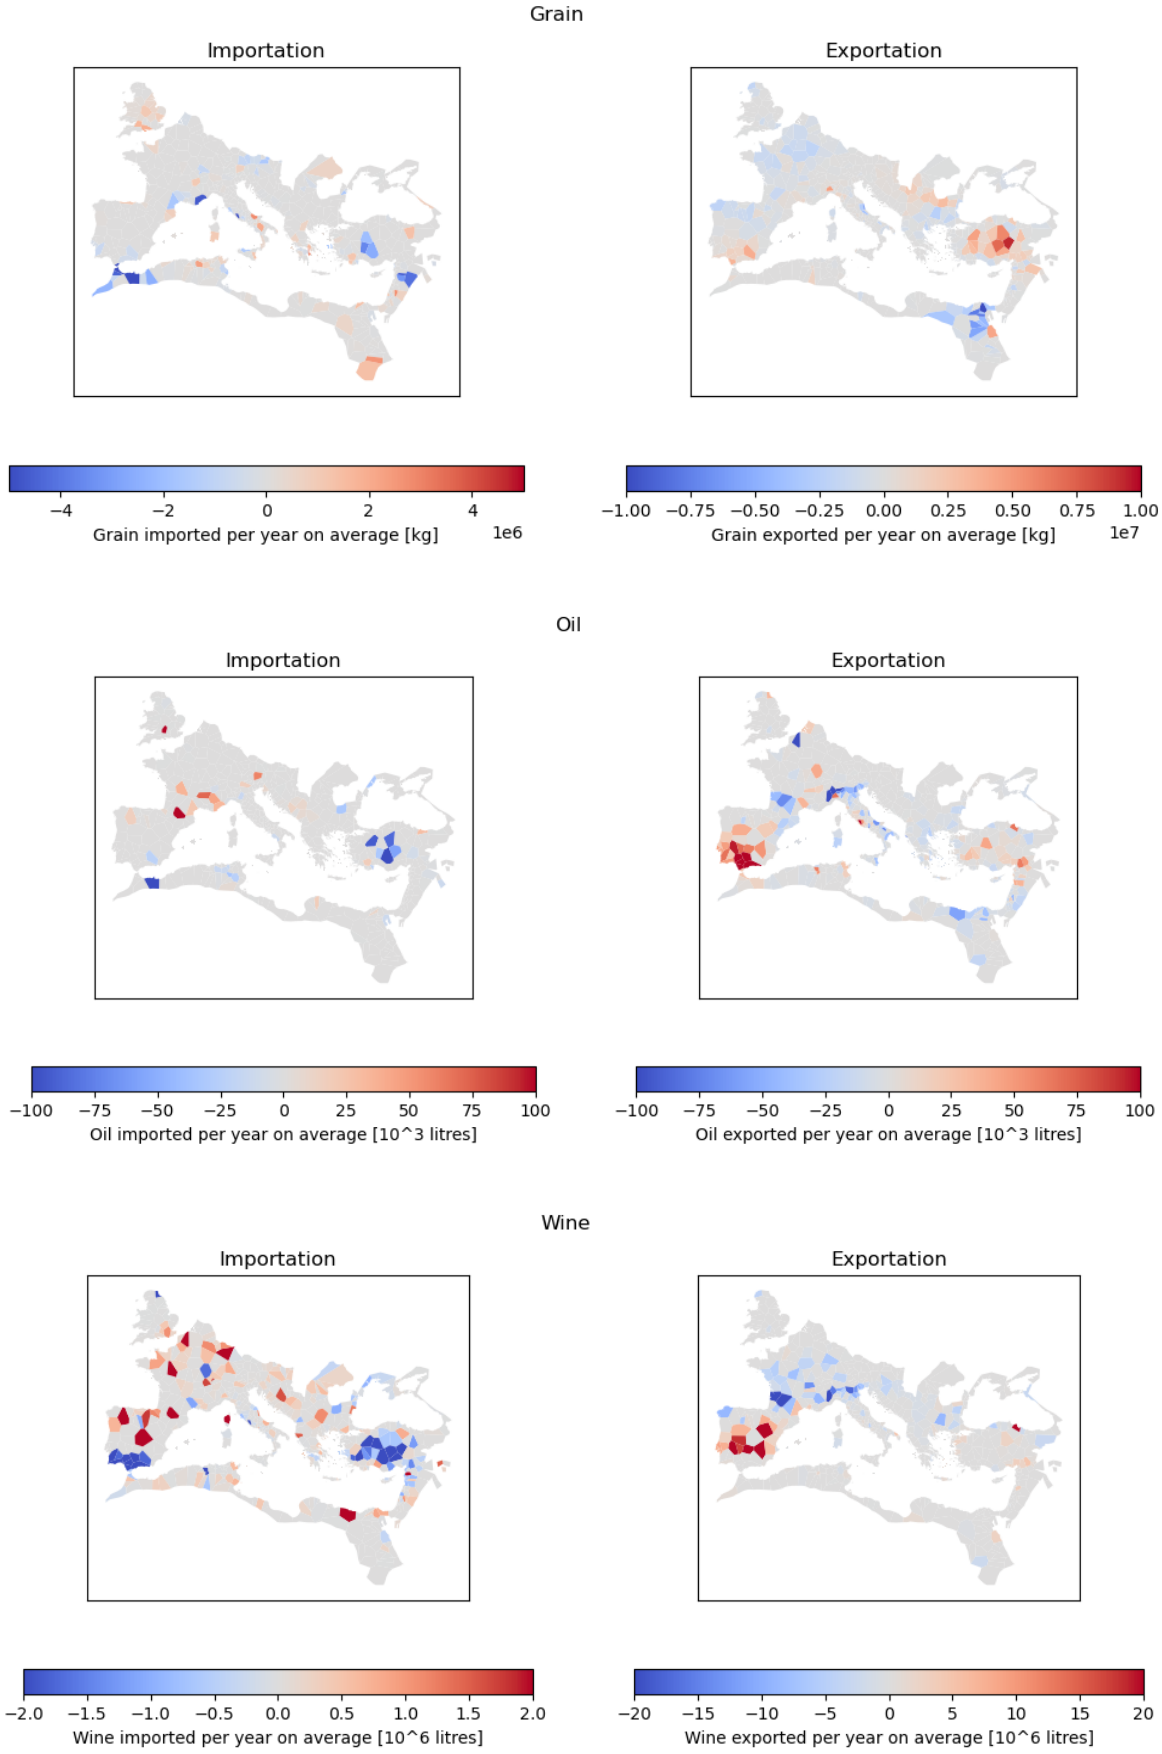

Figure 5.4: Map of the anomaly in imported and exported product per hinterland, between the Dark Ages Cold Period (Simulation 3) and the Roman Warm Period (Simulation 2) (DACP-RWP). Red stands for an increase in exports/imports, blue for a decrease.

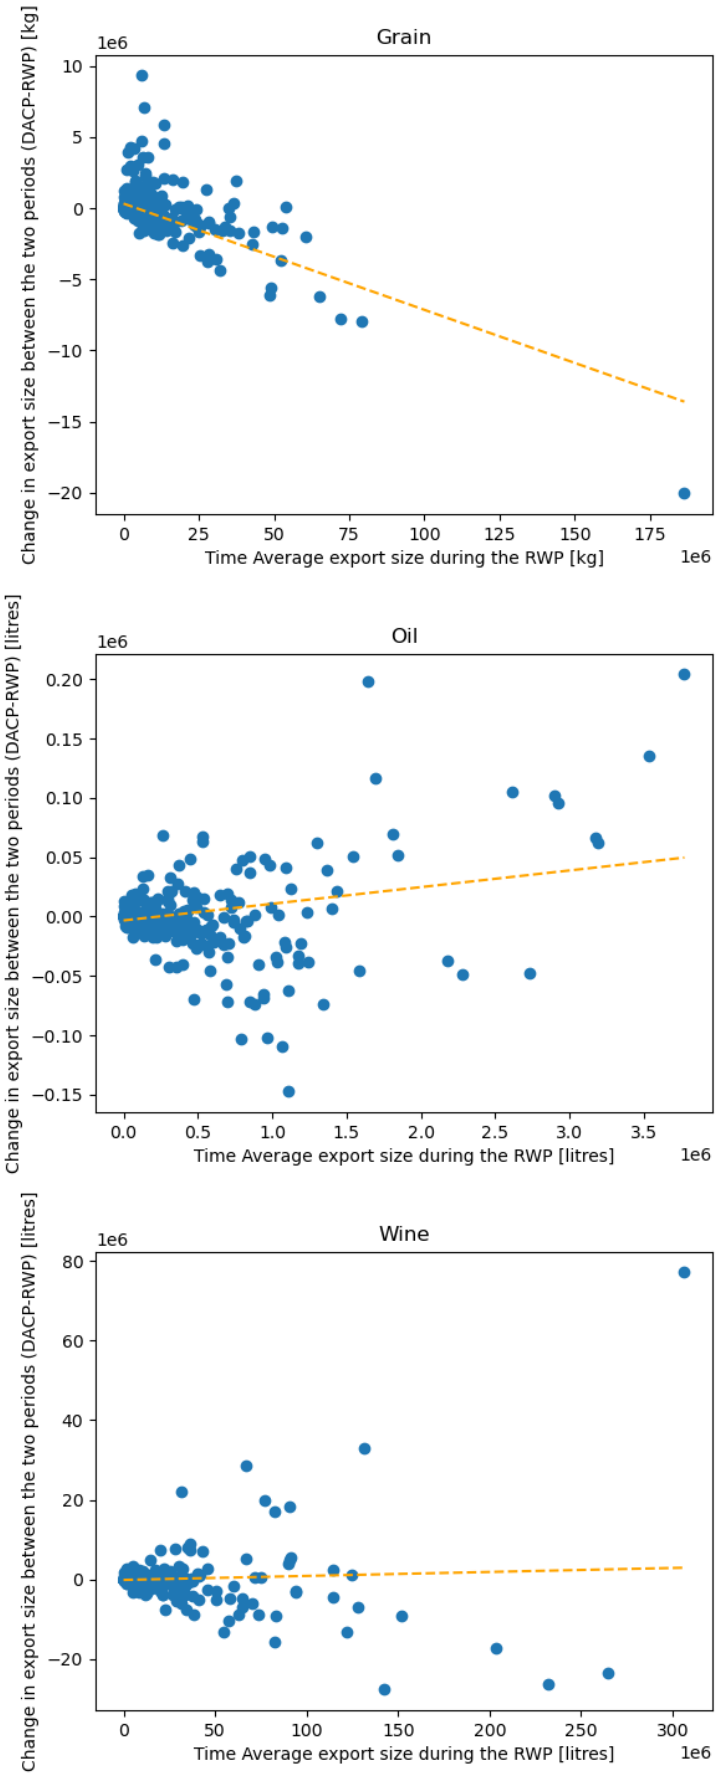

Figure 5.5: Scatter plots showing, for each region, the change in time-average export size between the two simulations (DACP-RWP), over their time-average export size during Simulation 2 (RWP).

Simulation 2 and 3: Roman Warm Period and Dark Ages Cold Period climates. Output of the trade model, grain			
	Roman Warm Period	Dark Ages Cold Period	Anomaly ((DACP-RWP)/RWP)
Average import/export size	$5.09 * 10^6$ kg	$5.03 * 10^6$ kg	−1.16%
Largest importer	*Al − Karib, Tunisia*	*Al − Karib, Tunisia*	–
Import size of the largest importer	$126.735 * 10^6$ kg	$127.788 * 10^6$ kg	+0.83%
Largest exporter	*Zagazig, Egypt*	*Zagazig, Egypt*	–
Export size of the largest exporter	$186.387 * 10^6$ kg	$166.335 * 10^6$ kg	−10.8%
Average size traded per trade connection	$211.7 * 10^3$ kg	$200.4 * 10^3$ kg	−5.3%
Largest trade connection	*Zagazig, Egypt* *Jerusalem, Israel*	*Zagazig, Egypt* *Jerusalem, Israel*	– –
Size of the largest trade connection	$20862.9 * 10^3$ kg	$19900.8 * 10^3$ kg	−4.6%
Average node degree	8.9	9.02	+1.3%
Average import costs per unit	3.43 sest/kg	3.49 sest/kg	+1.7%
Region with the highest import costs per unit	*Astypalea, Greece*	*Cankaya, Turkey*	–
Highest import costs per unit	16.5 sest/kg	16.6 sest/kg	+0.1%

Simulation 2 and 3: Roman Warm Period and Dark Ages Cold Period climates. Output of the trade model, oil			
	Roman Warm Period	Dark Ages Cold Period	Anomaly ((DACP-RWP)/RWP)
Average import/export size	$0.215 * 10^6$ litres	$0.215 * 10^6$ litres	–
Largest importer	*Ermopoli, Egypt*	*Ermopoli, Egypt*	–
Import size of the largest importer	$4.533 * 10^6$ litres	$4.533 * 10^6$ litres	–
Largest exporter	*Cordoba, Spain*	*Cordoba, Spain*	–
Export size of the largest exporter	$3.773 * 10^6$ litres	$3.977 * 10^6$ litres	+5.4%
Average size traded per trade connection	$9.7 * 10^3$ litres	$9.6 * 10^3$ litres	−1.6%
Largest trade connection	*ItayEl − Barud, Egypt* *El − Idwa, Egypt*	*ItayEl − Barud, Egypt* *El − Idwa, Egypt*	– –
Size of the largest trade connection	$1141.7 * 10^3$ litres	$1134.3 * 10^3$ litres	−0.7%
Average node degree	9.7	9.7	–
Average import costs per unit	5.15 sest/litres	5.08 sest/litres	−1.3%
Region with the highest import cost per unit	*Soissons, France*	*ArelauneenSeine, France*	–
Highest import cost per unit	19.1 sest/litres	19.1 sest/litres	–

Simulation 2 and 3: Roman Warm Period and Dark Ages Cold Period climates. Output of the trade model, wine			
	Roman Warm Period	Dark Ages Cold Period	Anomaly ((DACP-RWP)/RWP)
Average import/export size	$10.133 * 10^6$ litres	$10.087 * 10^6$ litres	−0.5%
Largest importer	*Ermopoli, Egypt*	*Ermopoli, Egypt*	–
Import size of the largest importer	$139.778 * 10^6$ litres	$139.778 * 10^6$ litres	–
Largest exporter	*Toledo, Spain*	*Toledo, Spain*	–
Export size of the largest exporter	$306.449 * 10^6$ litres	$383.623 * 10^6$ litres	+25.1%
Average size traded per trade connection	$344.5 * 10^3$	$330.6 * 10^3$	−4%
Largest trade connection	*PuebladelPrincipe, Spain* *Cordoba, Spain*	*PuebladelPrincipe, Spain* *Cordoba, Spain*	–
Size of the largest trade connection	$51111 * 10^3$ litres	$46648.3 * 10^3$ litres	−8.7%
Average node degree	12.8	12.3	−4.1%
Average import costs per unit	2.68	2.7	+0.7%
Region with the highest import cost per unit	*Eskisehir, Turkey*	*Eskisehir, Turkey*	–
Largest import costs	15.52 sest/litres	15.49 sest/litres	−0.26%

Table 5.2: Comparison of the output of the trade simulation for the Roman Warm Period (Simulation 2) and the Dark Ages Cold Period (Simulation 3). The green and red values are show a percentage change of more than 1%.

they involve relatively few regions, with many having almost unchanged productivity.

For all three products, Anatolia becomes increasingly more self sufficient and some cities in this region go from net importers to net exporters.

Table 5.2, gives an overview of the large scale network properties and their anomaly between the two periods.

5.3 Anomaly in the import costs

Changes in the surplus of smaller producers can still have a strong impact on isolated regions which depend on a low number of partners. To evaluate this possibility, I plotted the relative and absolute anomaly in import costs between the two periods (figure 5.6). Note that everywhere in this chapter with the term "import costs" I refer to the 110-year averages in cost of importation per unit product.

We observe that most regions experience a limited absolute variation in import costs (less than 3 sest/kg). Import costs for grain rise the most in the Eastern half of the Empire, particularly in Egypt, Lybia and Dacia. Wine import costs increase in Britain and in northwestern Europe. Oil import costs show very limited change. Central Turkey experiences a decrease in import costs for all three products.

Figure 5.7, shows the import costs anomaly (both positive or negative) plotted against the node's population and connectivity (node degree). We see that the regions with the largest anomalies in import costs for all products are characterised by a low connectivity (generally less than 10 partners for grain trades) and an average population size. That is, they are regions with a relatively large need for importation, but a limited connectivity to other nodes.

While these changes are quite small in absolute value, they represent large relative variation for some regions.

In figure 5.8, on the right side, I locate the regions experiencing a significant relative increase in the import costs on the import size/import costs

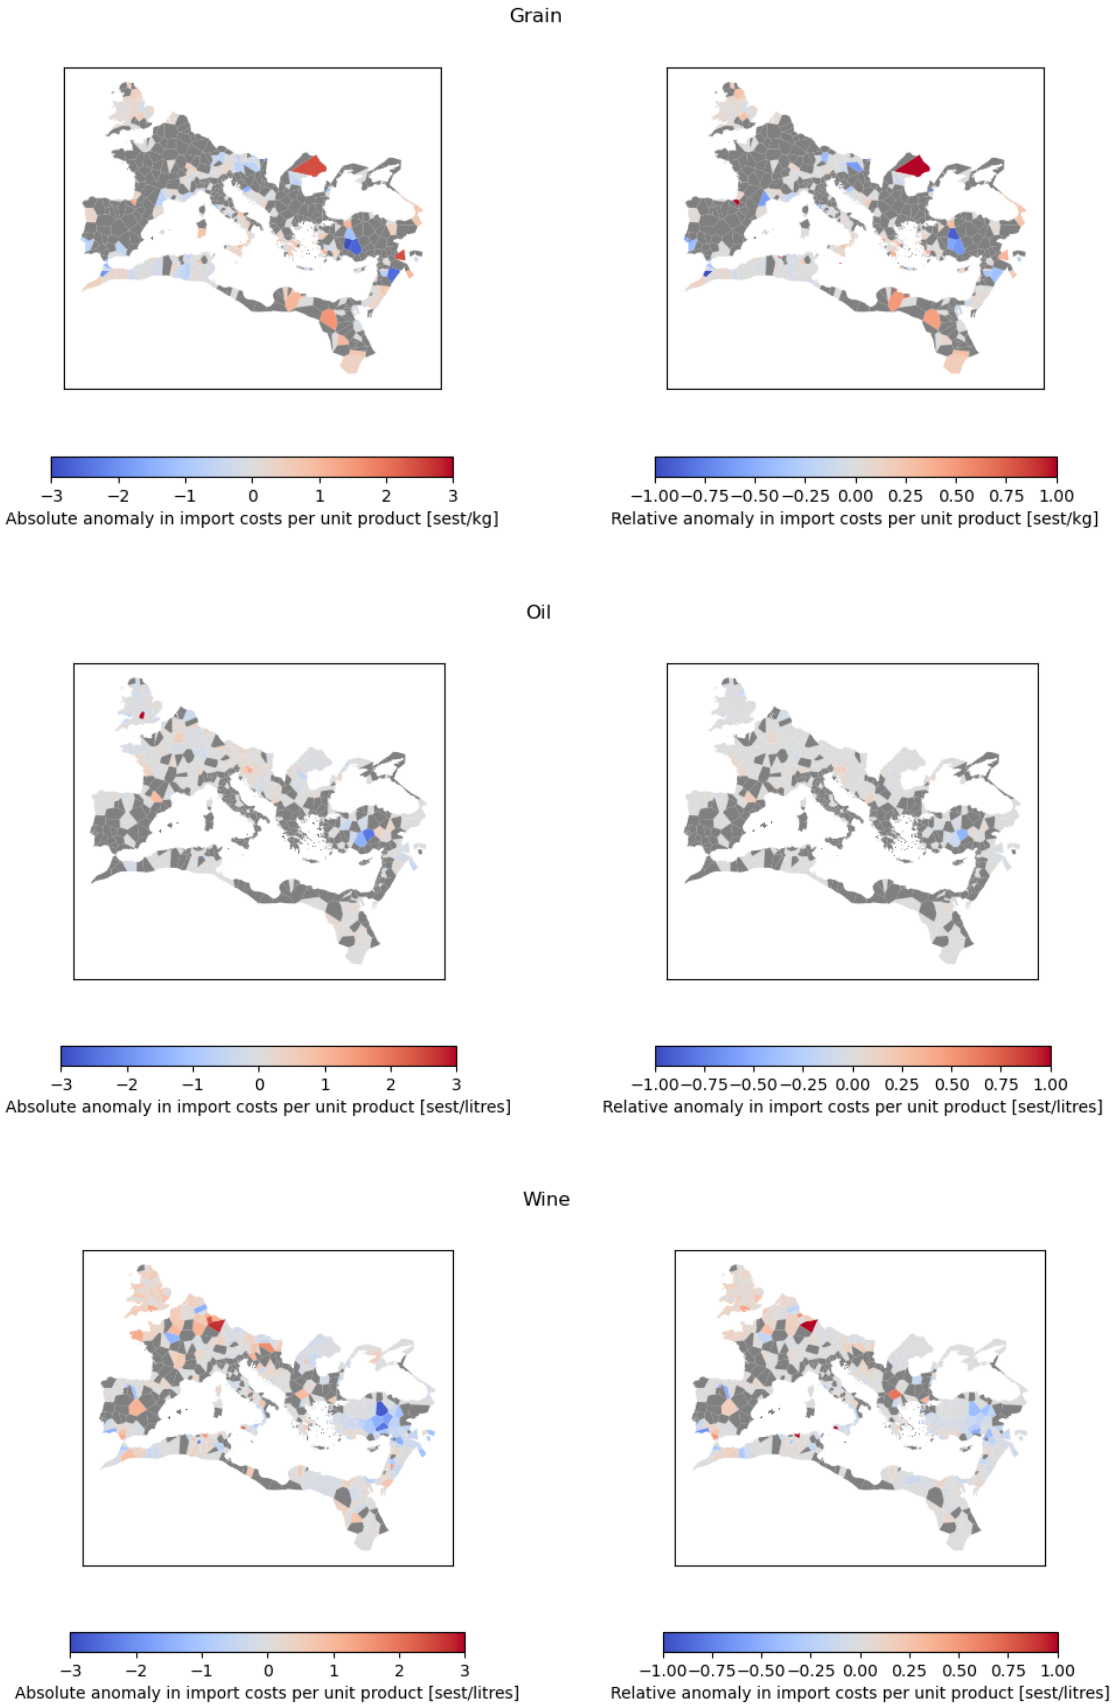

Figure 5.6: Hinterland map of the absolute (left) and (relative) change in import costs following the transition between the transition (DACP time average import costs per unit-RWP time average import costs per unit). Red stands for an increase in costs, blue for a decrease. The export nodes have been set to grey for visibility.

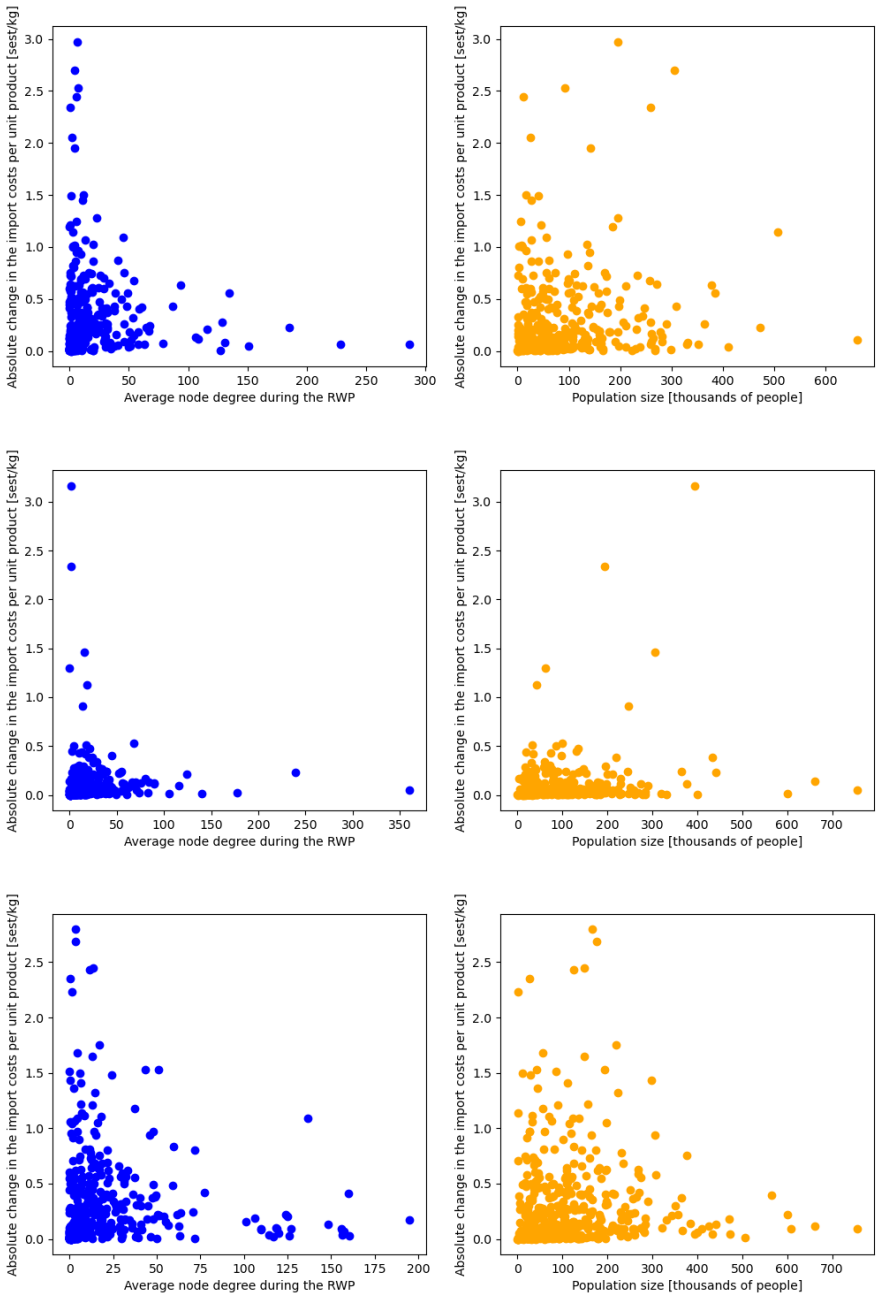

Figure 5.7: Scatter plot showing on the y-axis the anomaly in time-average import costs for each region between the two simulations, in absolute values. The x-axis represents the average node degree in the RWP simulation (on the left), and the population size of the node (on the right). I removed Rome from the right hand plots because its population size (large outlier) rendered the plot less readable. Rome experienced very limited increase in import costs (less than 0.1 sest/kg).

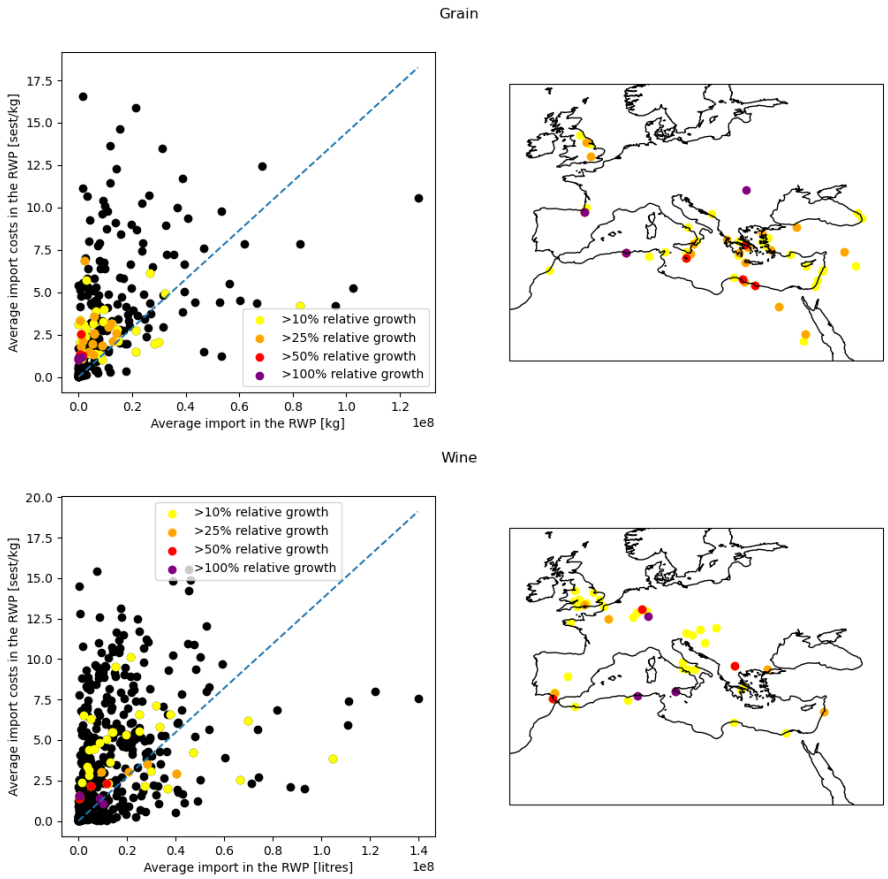

Figure 5.8: On the left. Plot of the time-average import cost per unit product, over the time average import size (both from the RWP simulation). The highlighted nodes experience a significant relative growth in their import costs between Simulation 1 (RWP) and Simulation 2 (DACP), as of the legend. The dashed blue line is a linear fit. We observe how most of the highlighted nodes lie above the linear fit. On the right. Map of the positions of the nodes highlighted on the left plot.

plot for the RWP. We observe that the most impacted regions, especially for grain trades, generally lie above the linear fit in the plot, meaning that the their import costs, while smaller than average, are relatively large compared to their import size. Thus, this are regions which, due to their geographical position, struggle more than others to connect with exporters. The plots are only shown for wine and and grain since oil imports showed very limited relative increase in import costs.

The left side of figure 5.8, shows the geographical location of the points highlighted on the right hand plots. Most of the regions showing a significant increase in import costs for grain are located in the eastern and south-

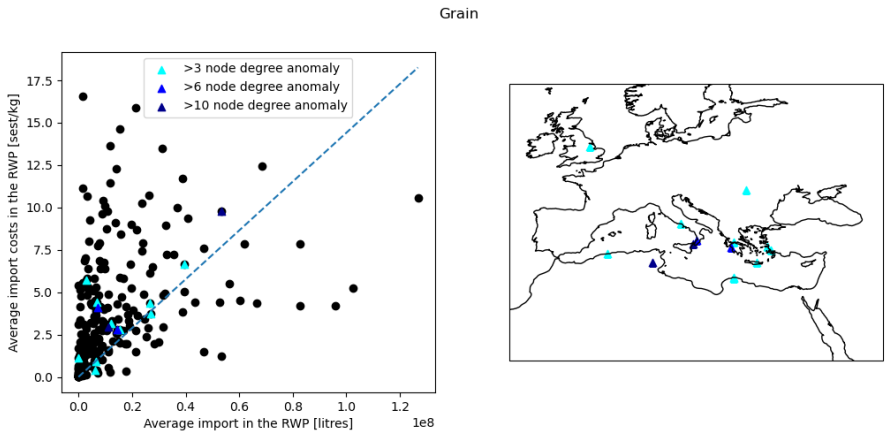

Figure 5.9: On the left. Plot of the time-average import cost per unit product, over the time average import size (both from the RWP simulation). The highlighted nodes experience a significant relative growth in their time-average node degree between Simulation 1 (RWP) and Simulation 2 (DACP), as of the legend. The dashed blue line is a linear fit. On the right. Map of the positions of the nodes highlighted on the left plot.

eastern Mediterranean. For wine trades, the geographical trend in less clear, but the northern borders and western Europe are the most impacted.

Figure 5.9 shows similar plots for node degree (number of contacted exporters) anomaly, only for grain trades. The regions experiencing the largest increase in node degree are located mostly in Tunisia, South-Central Italy and Greece, and they do not correspond to the regions with the highest import costs anomaly. Compared to the latter, these importers have on average a larger import size and they lie closer to the import cost/import size fit, meaning they are more easily connected to the network.

To further substantiate this point, figure 5.10 shows the relation between the import costs anomaly and the node degree anomaly for grain trades. As we can see here, the nodes experiencing the largest nodes in node degree show very limited import costs per unit increase.

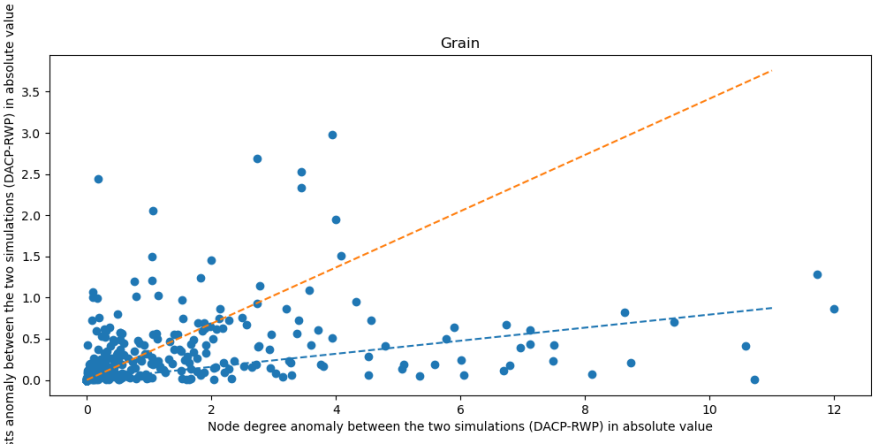

Figure 5.10: Scatter plot of the import costs (per unit product) anomaly between the two simulations, versus the node degree anomaly, both in absolute value. The two linear fit refer to nodes with import costs anomaly higher (orange line) and lower (blue line) than 1 sest/kg.

6. Discussion: Roman Warm Period and Dark Ages Cold Period climates

In this chapter, I first discuss the productivity and trade route anomaly between the two periods, and consider consistencies and inconsistencies with the historical record. Secondly, I evaluate which regions would have been more impacted by the changes in wheat production, and whether these changes justify any of the developments taking place in the third to fifth century AD.

Overall, the impact of the simulated climate shift on the agricultural production of most regions appears quite limited.

However, an important limitation of our analysis lies in the fact that we only considered 110-year averages. Table 5.1 shows that the Dark Ages Cold Period is characterized by a larger temporal variability. It's possible that, by selecting the most unproductive periods in the two simulations, we would observe larger anomalies.

6.1 Change in the productivity patterns

6.1.1 Grain

In the simulation, the largest negative anomalies in grain yields are observed in France and the Nile Valley.

Around this period we do have some indication of declining productivity in the historical record of the Nile river floods. From about 100 BC until about 300 AD the efficacy of the floods was recorded almost yearly [1, 79]. Modern historians collected these data and classified the events in six categories based on their productivity [1]. The percentage of recorded events in the two highest categories (high to very high floods) goes from 17.1% in the

period from 100 BC to 200 AD, to 5.4% in the period from 200 AD to 400 AD [1].

In our model results the Nile valley is still the largest grain exporter. The productivity loss, while quite large in absolute values, is not higher than 10% for the large producers in region. Consequently, the grain shipments to Rome do not appear significantly impacted. This is quite coherent with the historical record which shows that Egypt remained the largest exporter to Rome until its fall.

However, smaller regions at the edge of the river Valley, or coastal cities distant from the Nile Delta, experience significant loss, and some turn from net exporters to importers. In the archaeological record, the third century AD marks the beginning of the gradual abandonment of the Fayum agricultural villages, which lie among the regions experiencing significant yield loss in the model [26].

The relative loss in France appears everywhere very small.

The historical record for yields in France is not as detailed as that of the Nile valley. The rise in peasant revolts in this region in the third century AD [35] could be connected to increased wheat prices. However, our model results do not appear in support of this explanation.

The regions with the largest positive change in grain production are central Anatolia and northern Syria. The relative increase in these regions is also quite high.

Pollen core analysis has revealed an increase in cereal cultivation in Central Anatolia in the late Roman period (4th to 5th century AD) [55]. Archaeological evidence around this same period shows the expansion in cultivation towards marginal lands, which would reflect population growth [55]. Overall, the region appears to become more self sufficient.

Our model results support the theory that climate change was the reason for economic and demographic growth of this region.

6.1.2 Oil

Oil production appears overall positive impacted by the climate transition, in particular in Southern Spain.

Around the 3rd and 4th century AD we have archaeological evidence for a rise in North African oil exports, overtaking Spanish oil exports [49]. This is not justified by the climate transition in our model. This event could be the result of other socio-political factors, like the invasions of Spain by the Franks and other Germanic tribes [49, 47]. However, it's also possible that our model is not picking up this transition because of the previously noted underestimation of North African oil production.

6.1.3 Wine

In the case of wine, in our model we observe a decrease in wine exports from the Po plane and France in favour of Spain.

According to the archaeological record, the 2nd and 3rd century AD witnessed a decline in wine exports from Italy towards the provinces [11, 74, 80]. French wine exports, however, appear to be on the rise in this same period.

It's possible that the climate transition accounts for part of the dwindling Italian wine exports, but other explanations appear more fitting. For example, the demographic growth in Italian cities might have rendered the local markets more appealing than the provincial ones [74, 11]. This is especially true if we consider that wine exports from Italy started declining already in the second century AD, so before the beginning of the cliamte transition [74].

6.2 Growth in grain import costs

In this section we focus exclusively on grain. Grain was the product whose shortages were most likely lead to wider socio-economic issues. It also appears to be the product most impacted by the climate transition, and the one

whose productivity anomalies reflect best the historical record.

In figure 5.8 we observed how some isolated regions in Egypt, Lybia and Palestine experience a large relative growth in import costs per unit. The same applies to some cities on the Turkish coasts, or in the Syrian inland, despite the overall growth in productivity in these regions. In regions reliant on importation (as most of these cities), this would have implied a growth in wheat prices.

In our model, grain is redistributed until all demand is satisfied. So as long as the overall production is higher than the overall demand (which is true for both periods), we do not observe any deficit. In a real life network, a significant rise in average import costs might have also led to periodic shortages.

In the simulation, the cities which experience the highest import cost anomalies are not metropolis, but they are still sizable cities for the time, generally ranging between 10 thousand and 100 thousand people, with some larger outliers. For instance, the import costs in the city of Ezine, on Turkish Aegean coast, which hosted about 300 thousand people, rise by more 40%. Similar increase is observed in Bawiti, Egypt, which hosted about 40 thousand people.

In the city of Aswan, at the mouth of the Nile river, import costs rise by more than 25% on average. Aswan's population is estimated at about 20 thousand during this period, and the city was fundamental for its role on the Nile trade route.

The common denominator among these cities seem to be a low node degree, and import costs larger than expected for their size (figure 5.8). So, they are regions reliant on importation, but not as easily connected to the network as cities in the central Mediterranean.

Conversely, the regions experiencing the largest increase in node degree, as in the number of contacted exporters, are larger in size (usually between 100 and 500 thousand people), and mostly lie in the Central Mediterranean (figure 5.9).

Some of the cities with the largest growth in node degree are: Agrigento, Italy (population about 500 thousand), Vibo Valentia, Italy (population about 100 thousand) and Sbeitla, Tunisia (population about 300 thousand).

The relative rise import costs in these regions is small compared to the first ones (10% or less), although a limited number of cities does experience a large growth in both import costs and node degree (figure 5.6).

The model paints a picture in which large, well connected, central Mediterranean cities are able to cope with the falling productivity by diversifying their sources. More isolated cities in the Roman East struggle to do the same and experience significant growth in wheat prices.

Rome, which was already connected to more than 200 exporters both in the historical and in the Roman Warm Period simulation, does not experience a significant rise in either import costs or node degree.

In order to compare these results with the historical account, we consider the geodatabase constructed by McCormick et al. [1]. It contains the date and location of all recorded food shortages and price spikes over the Empire's territory, for the period from 100 BC to 500 AD.

By analysing these data, I first note that the number of recorded food shortages increases more than three-fold in the period from 200 to 500 AD compared to the period from 100 BC to 200 AD (49 recorded events versus 16). Figure 6.1 shows the location of the food shortages for the second period (200 to 500 AD). All of the recorded events took place in the central and eastern Mediterranean.

The smaller size cities in figure 6.1 (marked in blue) do not exactly correspond to the ones in figure 5.8, but they are often in close vicinity. Our model also predicts shortages on the Greek and Turkish Aegean coasts, which are not verified at least in this dataset. This could be a consequence of the model's underestimation of North African productivity (see chapter on Model Validation).

Compared to our model results, the historical record also shows several

food shortages in major cities, such as Rome, Cartage, Alexandria, Constantinople and Jerusalem (marked in green).

As for Rome, it's important to note that 13 out of the 16 food shortages for the period between 100 BC and 200 AD were located there. For the period between 200 and 500 AD, the city experiences 17 shortages. So the increase is not as large as other regions, and the number of recorded events is partly a consequence of the higher frequency of reporting. Secondly, most of the food shortages in Rome are blamed by contemporaries on exogenous factors such as unfavourable winds, pests contaminating the granaries or piracy, rather than on supply issues.

As for other large cities, it is possible that our model is overestimating their flexibility in imports.

The model results simply state that it was technically possible, under the climate forcing we simulated, for these regions to preserve their supply without a large growth in import costs. But it is possible that the forced diversification of imports would have led to a price spike anyway.

We also note that Egypt is the region with the largest wealth of evidence for inflation during the 3rd century AD [36, 2].

To conclude, our model supports a climate-induced rise in wheat prices in the several cities of the Roman East and Southeast. This price growth might have led to food insecurity and aided the spread of epidemics the last two centuries of the Roman era. However, the developments concerning the Western Empire, such peasant revolts in France, Germanic invasions and price inflation in the western provinces, do not appear to have a reasonable climate-related explanation in our model results.

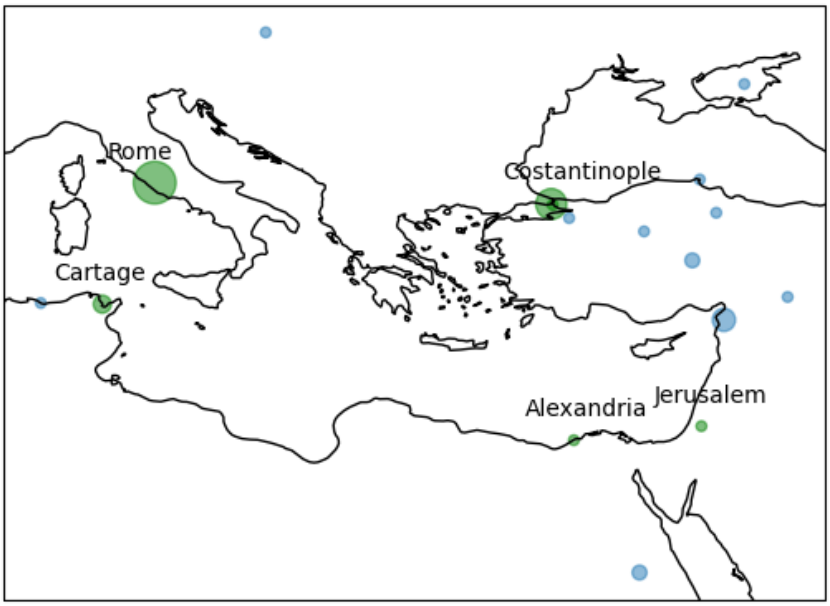

Figure 6.1: Map showing the location of recorded food shortages between 200 and 500 AD, according to the database used by McCormick et al. in their published paper [1], and made available at https://darmc.harvard.edu/data-availability. The size of the point represents the number of recorded events. We highlighted in green the largest cities.

7. Conclusion and outlook

The reasons and mechanisms behind the fall of the Western Roman Empire are among the most commonly debated topics for historians of this period. But most of them today agree that it is misleading to look for a single answer.

Since the very onset of the Empire period, Rome was plagued by systematic issues that rendered it vulnerable to triggers of many kinds, including but not limited to climate change. Some of these vulnerabilities we described in the literature review chapter. Others include, for instance, the peripherization of political, military and economic resources, which left Italy depending on the provinces, eroding gradually its control over them [39].

It is then important to phrase climate change as one of the possible triggers of the Roman downfall, and not the overarching motivation.

Our model results tell us that it is reasonable to imagine that some regions in the East and Southeast of the empire would have struggled with food security and soaring wheat prices due to climate change. Possibly, larger cities in the central Mediterranean might have also had troubles ensuring supply, depending on how elastic they were in their import choices.

Inflation in wheat prices in well attested for throughout this period in Egypt [36, 2]. It could have easily spread to other products, or to other distant regions which depended on Egyptian grain.

Food insecurity in the southeast might have aided the spread of epidemics. The Cyprian plague (249-270 AD) caused in Egypt some of the highest death toll [29]. The Antonine plague which preceded it was probably another trigger for inflation, and might exacerbated the climate-induced issues.

The political instability of the 3rd century AD likely caused food short-

ages on its own rights.

Our model, at least based on the climate forcing we used, can't justify a climate-induced rise in wheat prices in the western Empire. While inflation is less attested in the west than in Egypt, it is regarded as one of the possible underlying issues leading to the military uprising of the 3rd century AD, with military wages unable to keep up with the rise in prices [81, 19].

So, to conclude climate appears as a possible factors contributing to some of the troubles of the Roman Empire around this period, but more work needs to be done to verify whether it was more impactful than other triggers.

The model we employed could be used to simulate other events and their socio-economic impacts. In particular, simulating the impact of the Antonine plague on wheat prices in Egypt, and comparing it with the impact of climate change, could provide significant insights on the relevancy of our own results.

The hysteresis component of the model, when applied more wisely only to large population centres, can be helpful in gauging the capability of large, central Mediterranean cities to diversify their imports when faced with climate change.

Finally, looking at climate extremes, not only time averages, can contribute a lot to evaluate the stress caused by climate instability, which could be significantly larger than that caused by the climate mean state.

Appendices

A. Sensitivity to the number of iterations of the redistribution algorithm

As described in chapter 3, the trade model operates the redistribution of resources over several iterations. The number of such iterations affects the trade network structure.

For a single iteration, every node with a resource deficit tends to import from its one or two closest surplus-producing neighbours. While this produces the most efficient output in terms of cost distance, it does not reflect the condition of real world systems.

By increasing the number of iterations, we encourage competition between different deficit nodes for the resources of the same surplus producing region. In practice, this translates into a larger number of active connections, with a lower average amount traded along each.

However, for our results to be valid, the simulation should not be overly sensitive to small changes of this parameter around the value chosen.

Figure A.1 shows the time-average number of active trade connections and the average size of them as function of the number iterations of the trade model. It is easy to verify that, from around 500 iterations onwards, the simulation reaches a stable point. We chose 1000 iterations as a trade off between stability of the simulation and the computational length.

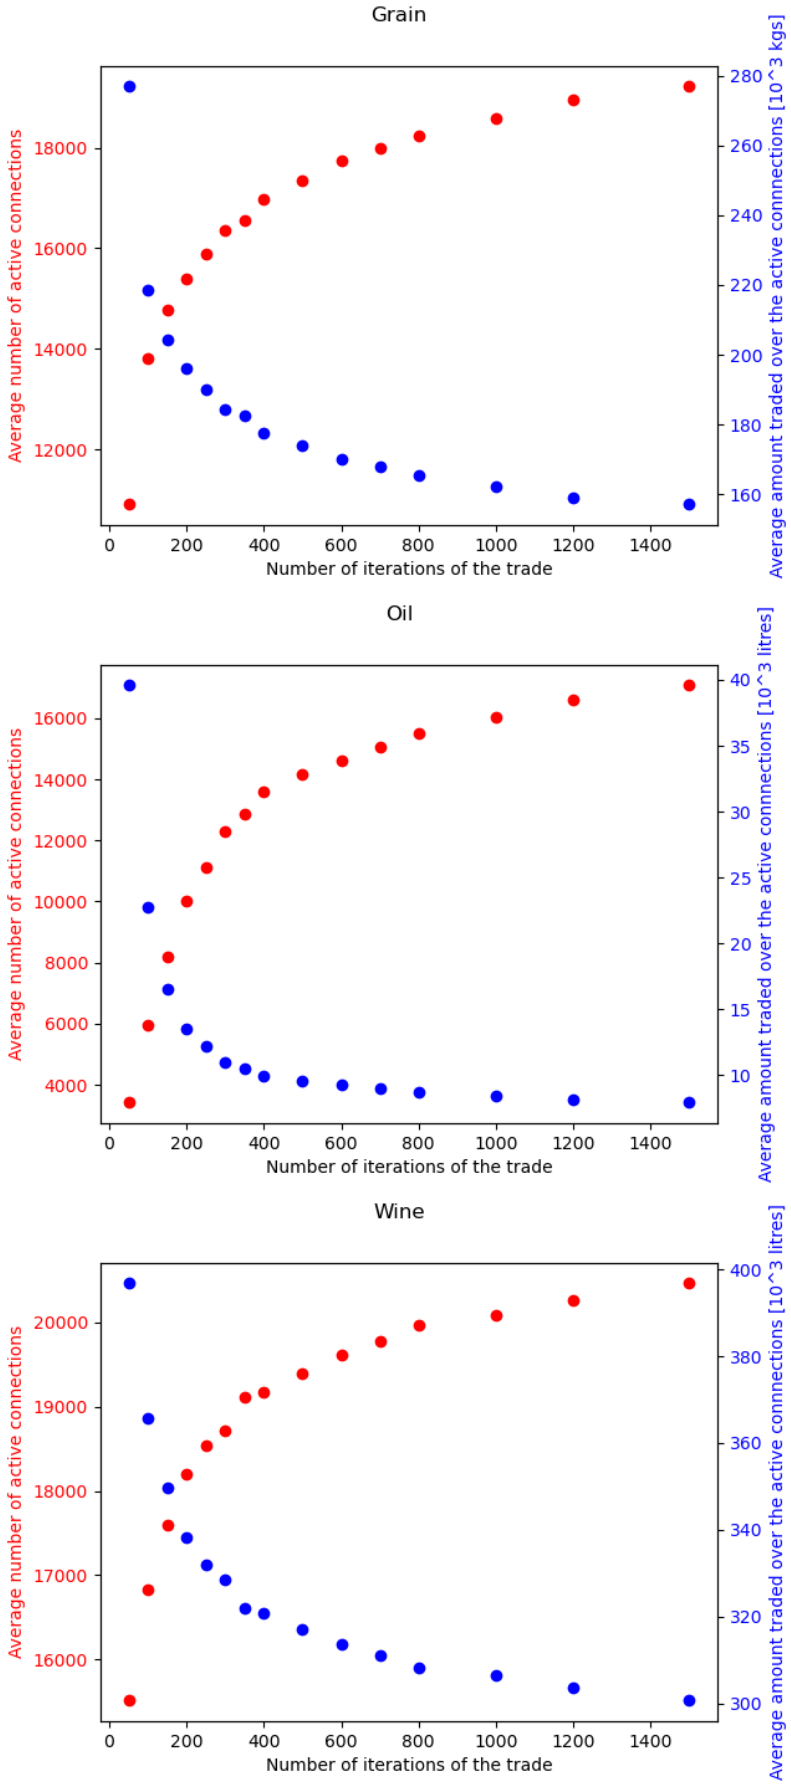

Figure A.1: Sensitivity of the trade model to the number of iterations. Red dots: Number of active connections across the whole network, at any point of the 110-year period. Blue dots: average size of product traded along a single connection, 110-year averages. The model reaches stability around 500 iterations.

B. Hysteresis impact on trades

B.1 Sensitivity of the simulation to hysteresis

The trade network whose results I described operates as memoryless. That is, each year it redistributes the available resources independently from the year before. In order to simulate the development of some established routes, and the preference for regions with a stable yield, we may insert an element of hysteresis, or memory, as described in section 3.1.3.1. This component is controlled by a parameter which determines the maximum percentage of the cost distance dropped for each trade connection every year of the simulation. We apply this new component to the network and evaluate the changes. The possible improvement in historical accuracy is discussed in the following section

Through the application of this additional model element we expect to observe:

- A reduction in the node degree of the importing regions, which will prefer fewer larger exporters over a larger number of smaller ones.
- An increase in the amount traded through each connection.
- A reduction in long distance trade, especially for small amounts. Long distance trades should be reduced to larger, specialized trade connections.
- A change in the geographical trade patterns which favours regions with more stable yields.
- Importantly, we expect trades to be stabilized, so that an importing region is not likely to change its partners frequently throughout the simulation.

We first evaluate the development of some general characteristics of the

trade network over time. Figure B.1 shows the temporal development of the total number of connections and the average amount traded, for different values of the hysteresis parameter.

The hysteresis component of the model causes a rapid decrease in the number of active trade connections over the first 50 years of the simulation, while the ones remaining predictably increase their average size. Thus the average node degree decreases, so that each importing node tends to rely on a lower number of larger partners.

After about 50 years of the simulation, these values become stable for all tested values of the parameter. While the stability point depend on the hysteresis parameter, the sensitivity to it decreases significantly past a value of 0.5%. For the following tests, we will then limit the parameter to a range between 1% and 2%, and we will remove from the results the first 50 years of the simulation in order to exclude the stabilization period.

So, the first two points are verified, we observe an increased specialization as the system sheds small sized trade connections over time.

Figure B.2 shows the cost distance distribution of trades for different values of the hysteresis parameter (same plot as in figure 4.3). While the anomaly is not particularly large, the mean of the lognormal distribution shifts to slightly lower values for all products. The largest negative anomaly little after the distribution peak. The end tail of the distribution is not as strongly affected, so that we still have some very long distance trade connections, but fewer mid distance ones.

Also the third point is verified, we see an increase in the probability of short distance trade in favour of long distance trade.

Next we evaluate whether and the how hysteresis modifies the spatial trade patterns. Figures B.3, B.4 and B.5 show the arrow plots for the three products for different values of hysteresis.

In the case of grain trades (figure B.3), we observe how the region of modern-day Tunisia (at the time just known as Africa), increasingly becomes the system's center of trades, rather than Italy. This region was only

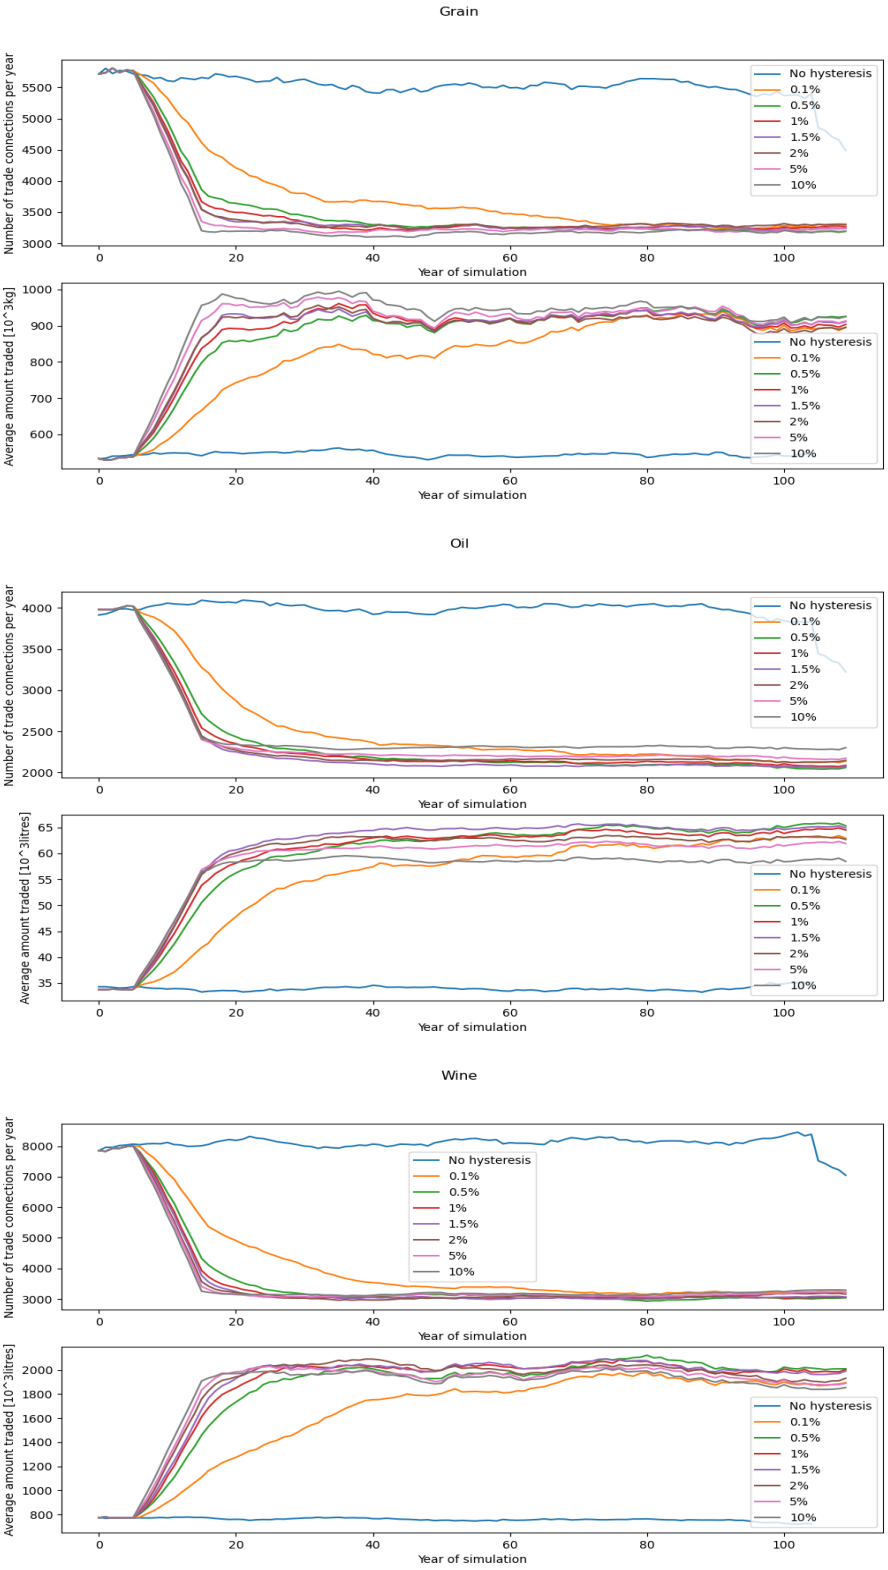

Figure B.1: The plots show the development over time of the number of the active trade connections per year (upper plot) and the average size of a trade connection (lower plots), for different value of the hysteresis parameter. We observe that the simulation reaches a stable state after about 50 years. The simulation shows a small sensitivity to the parameter past a value of 0.5%.

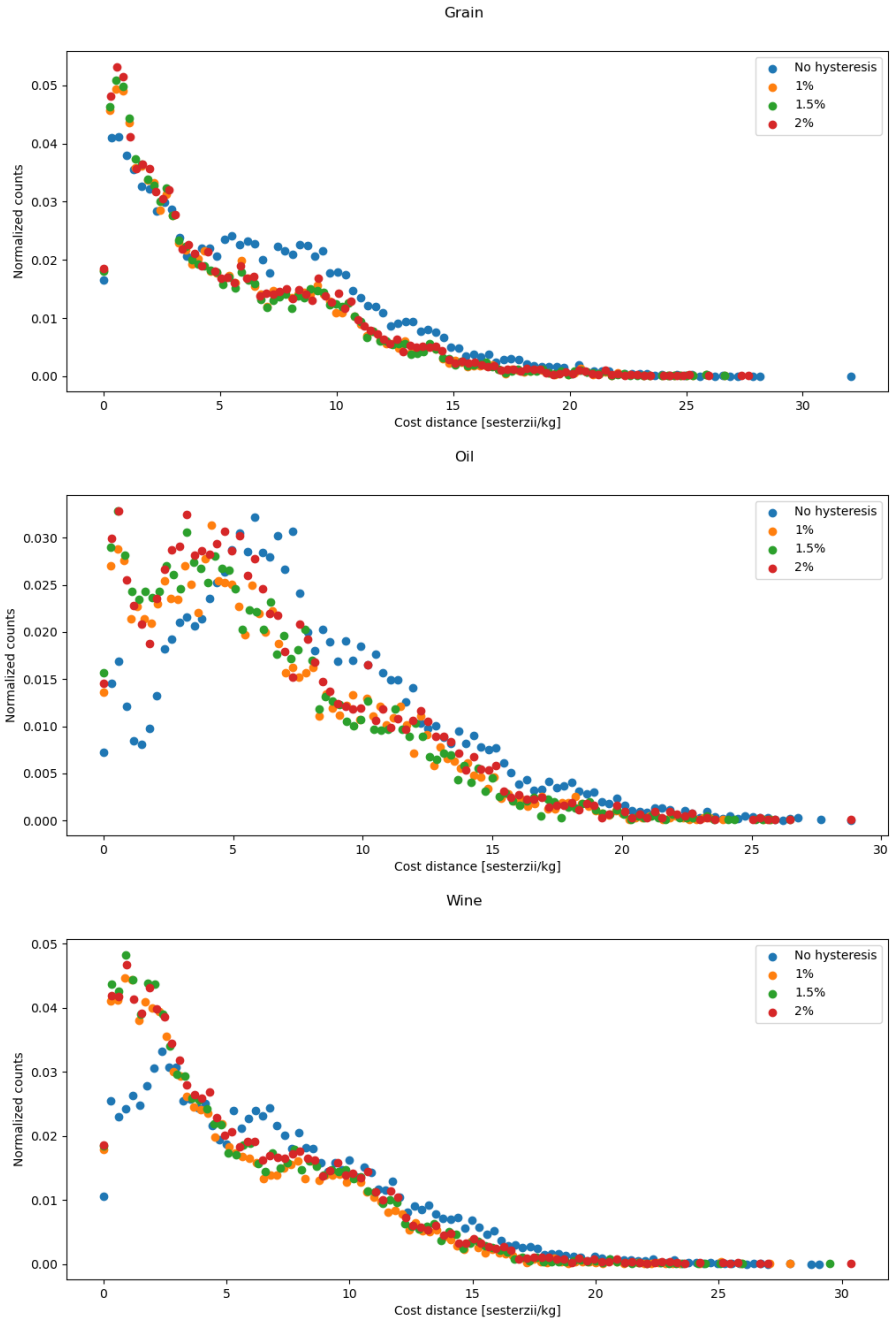

Figure B.2: Cost distance distribution plots for different values of the hysteresis parameter. We observe a decrease in the frequency of mid distance trades, while the tail of the distribution is not as affected.

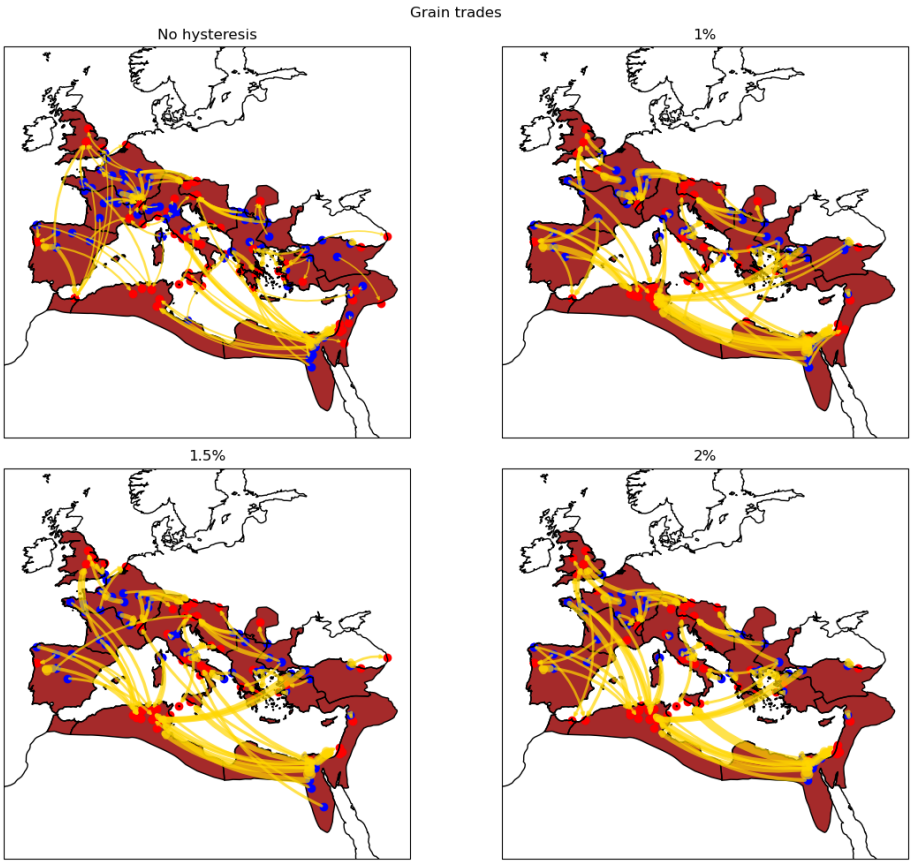

Figure B.3: Arrow plots representing the 110-year averages grain trades for different values of the hysteresis parameter. For readability, all plots are limited to the largest 25% of all trade connections. We notice the increasing importance of Egypt as a long distance exporter, and the loss in importance of northeastern France.

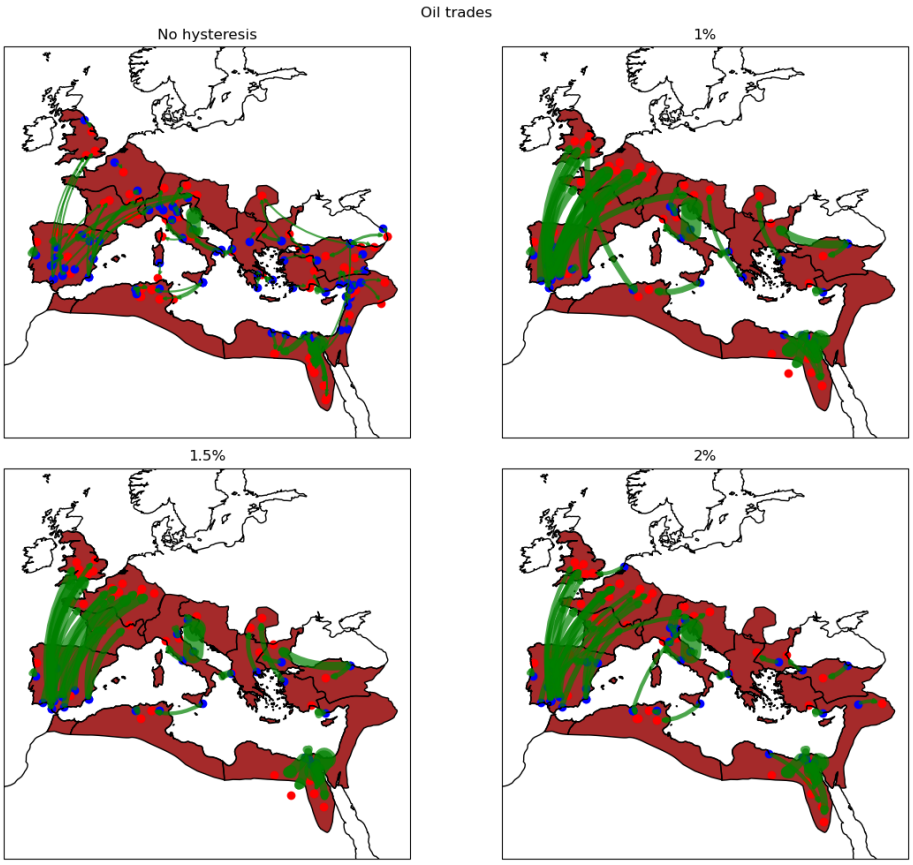

Figure B.4: Arrow plots representing the 110-year averages oil trades for different values of the hysteresis parameter. For readability, all plots are limited to the largest 25% of all trade connections. We notice the growth in the Atlantic trade route.

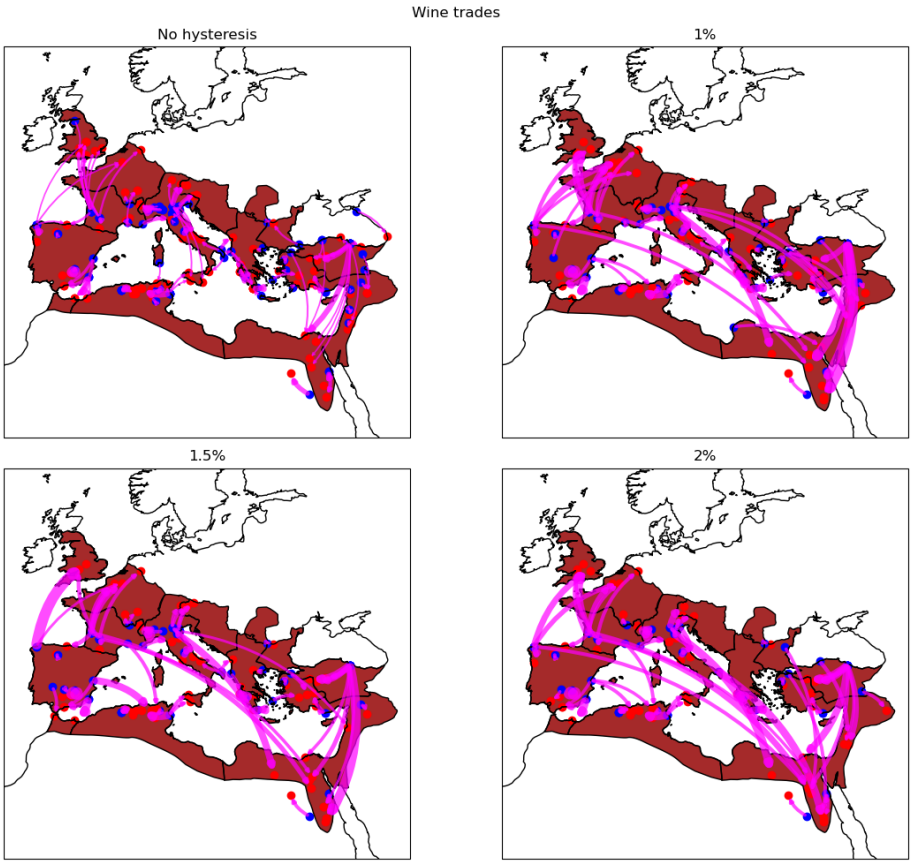

Figure B.5: Arrow plots representing the 110-year averages wine trades for different values of the hysteresis parameter. For readability, all plots are limited to the largest 25% of all trade connections. We notice the growth in long distance trades across the Mediterranean towards Egypt.

moderately populated but was quite poor in resources. The regions of central France and the Po Valley trade significantly less in the hysteresis simulation, while Egypt, Northwestern France and northwestern Spain become more important.

As for the oil trades (figure B.4), the most obvious feature is the increase in the Atlantic oil trade from southern Spain to the British Isles. Other than that, hysteresis just seems to favour exports from larger hinterlands over smaller ones within the same geographical area, so that the overall picture of the trade network changes little.

Finally, for the wine trades (figure B.5) the most relevant feature is the increase in the size of long distance trades across the Mediterranean, from southern France towards Egypt.

To verify whether these changes truly favours regions with more stable yields, we focus on some large importers.

I selected for each product three among the largest importers and I evaluated the temporal evolution of trades towards these nodes with and without hysteresis.

Figure B.6, B.7 and B.8 show these plots respectively for grain, oil, and wine trades. The plots on top show the temporal evolution of the stability of the connected producers; that is, the mean to standard deviation ratio in the agricultural production of the trade partners.

We observe that, in the case of grain trades, an hysteresis component encourages importation from more stable producers towards Rome and Agrigento (similar results are observed in other italian cities), while in Albertville, France it causes very limited change. In the case of oil and wine however, the stability of the importers either declines or remains within the same range.

The bottom plots show the year by year variation in the choice of partners; that is, how many trade partners are switched each year for new ones.We notice that, for almost all the selected cities and products, hysteresis does not produce stability in trade partners. This is ultimately why hysteresis

Rome, Italy

Agrigento, Italy

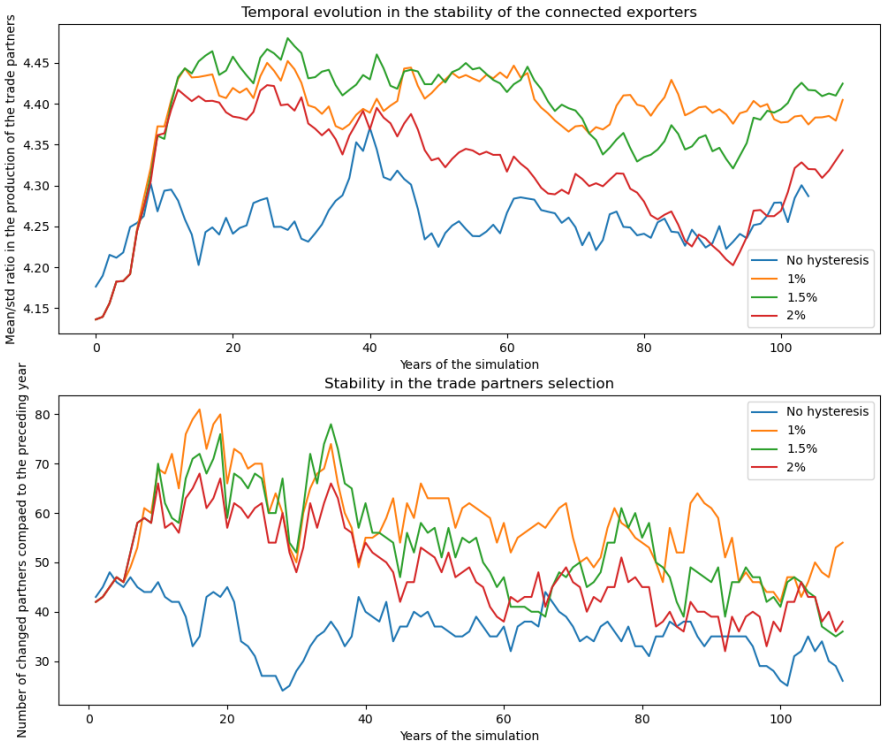

Figure B.6: Plots showing the temporal evolution of trades towards the three largest grain importers (Rome, Agrigento, and Albertville) with and without hysteresis. The upper plots show the stability in the agricultural production of the connected exporters, in terms of mean to standard deviation ratio. The lower plots show the stability in the choice of partners by the importation region, on the y axis it is plotted the number of connections replaced compared to the preceding year. All values are 5 years running averages.

was not used for the results described above. While it provides limited improvement to the results in terms of historical agreement, it also increases the instability in most large importation regions, which would go against its purpose. As described in the next chapter, this may be a consequence of our choice to apply hysteresis uniformly across the network.

B.2 Hysteresis and historical accuracy

The motivation for inserting a memory component in the trade model was to simulate the adaptation of the trade system to the agricultural production pattern. We expect the emergence of some specialized trade routes. As a matter of fact, in the period we are interested in (2nd to 3rd century AD), large cities would be exploiting long term agreements with the provinces for grain provisions, and possibly for other products [44, 45, 46]. The exporting regions involved in this established routes were likely stable producers which could be relied on for several years in a row.

By observing the anomaly introduced by hysteresis in the trade patterns we can gauge whether memory and institutional lock-in would have been an important factor in favouring some trade routes over others. Further, by analysing the temporal development in the imports towards specific large cities, we can determine whether memory, as reproduced in the model, brings forward a stabilization of the trades towards large importers.

Considering the grain trades, as shown in figure B.3, we see that Egypt becomes more prevalent in large-size, long-distance routes, while central and eastern France lose importance. This is an improvement on historical accuracy compared to a simulation without hysteresis, which goes in support of the popular theory that grain exportation out of Egypt were largely derived from taxation of imperial estates, so highly institutionalized [41, 44]. Egypt's yields, compared to France or Spain, are more reliable over time because of the prevalence of irrigation. We also observe, however, that the exports of grain from Egypt towards the densely populated regions of coastal Tunisia are more encouraged by hysteresis than the trades towards Rome and Italy. This could be a result of the overestimate in the agricultural

Ermopoli, Egypt

Eracleopoli, Egypt

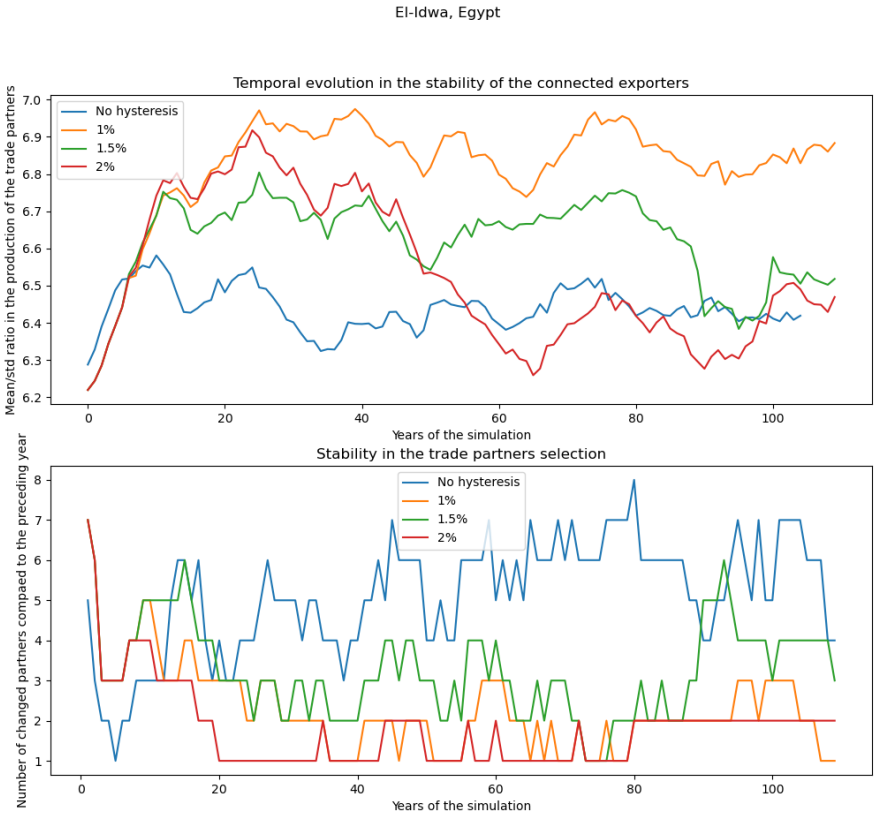

Figure B.7: Plots showing the temporal evolution of trades towards the three largest oil importers (the Egyptian cities of Ermopoli, Eracleopoli, and ìEl-Idwa) with and without hysteresis. The upper plots show the stability in the agricultural production of the connected exporters, in terms of mean to standard deviation ratio. The lower plots show the stability in the choice of partners by the importation region, on the y axis it is plotted the number of connections replaced compared to the preceding year. All values are 5 years running averages.

Alexandria, Egypt

Albertville, France

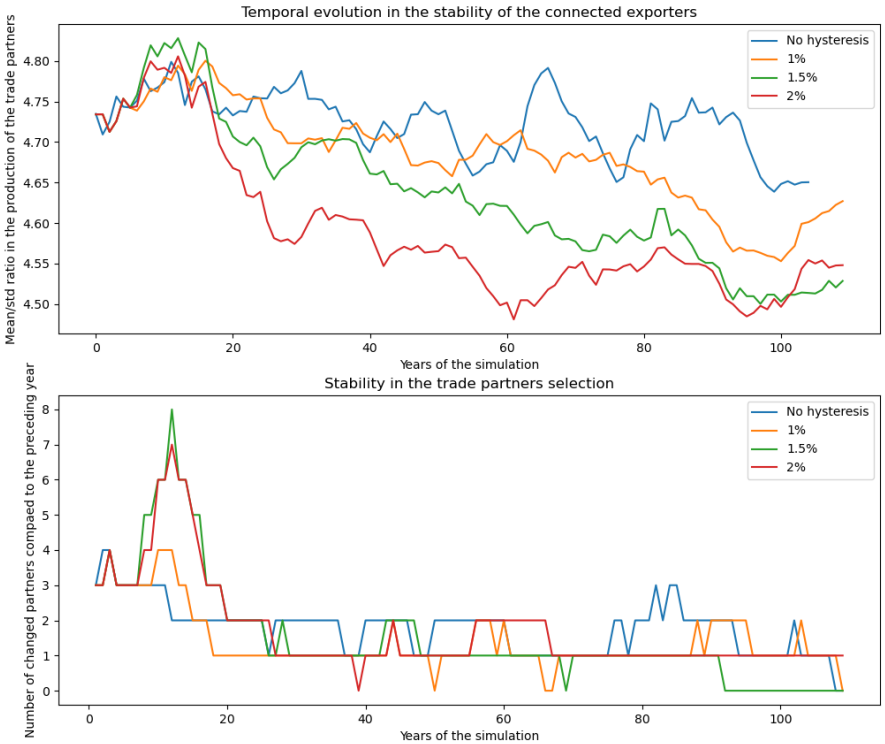

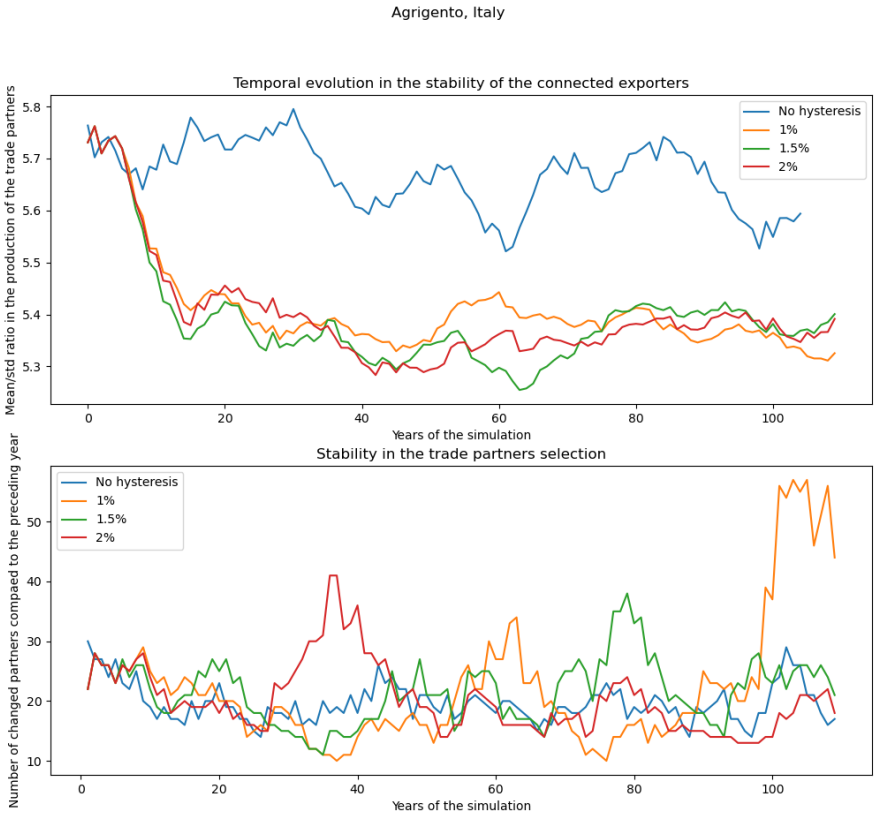

Figure B.8: Plots showing the temporal evolution of trades towards the three largest wine importers (Alexandria, Albertville and Agrigento) with and without hysteresis. The upper plots show the stability in the agricultural production of the connected exporters, in terms of mean to standard deviation ratio. The lower plots show the stability in the choice of partners by the importation region, on the y axis it is plotted the number of connections replaced compared to the preceding year. All values are 5 years running averages.

production within Rome's hinterland, as mentioned.

As for oil (figure B.4), we see that the Atlantic trade route is significantly aided by the presence of a memory component, so that southern Spain tends to export more towards the Empire's borders than towards the central Mediterranean. This is also an improvement in the historical accuracy of the model results. The Atlantic Baetican oil route might have had a large degree of institutionalization since it was used to supply the legions: our result supports this hypothesis [48, 49, 42]. North African oil exports remain underestimated, which leads me to believe that it could be the result of our choice of land allocation.

Finally wine trades (figure B.5), do not overall improve in accuracy through the introduction of a memory component. The only exception is the increase in size of Rhine river route supplying the norther borders, which is well established by archaeological evidence [42]. Possibly this route was also part of the state organized supply to the legions [42]. However, the most obvious feature emerging from hysteresis is the increase in the anomalous long distance trades towards Egypt. Overall, the geographical patterns of wine trades become less plausible when we account for memory. Wine is believed to have been the most liberalised of these three products, largely left to the free market, and our results support this theory.

I then looked into the trades towards some large cities to verify whether hysteresis had the effect of :

- inserting a preference for stable exporters with a higher mean to standard deviation ratio in agricultural production.
- stabilizing the choice of partners over time.

We see that the first condition is largely verified for grain, but not for other products. The second condition however, is mostly unverified. Rather, in many cities hysteresis seems to increase the trades' instability. This is an unexpected result.

I believe that an explanation for this anomalous result lies in the choice of applying the hysteresis component equally to every node in the system,

independently of the population size. This means that each node attempts to stabilize its trades with the same priority as the others. This overall leads to further instability as all importing regions influence each other. This effect only becomes visible when we consider the behaviour of the system over time. When we are looking at 110-year averages, we only observe the overall preference of all nodes for large and stable exporters.

A possible way to solve this issue is to limit hysteresis to large population centers, or establish an hierarchy between importers. This would also be more historically accurate, since institutionalization of trades likely only happened in large cities to a meaningful degree.

Bibliography

[1] Michael McCormick et al. "Climate change during and after the Roman Empire: reconstructing the past from scientific and historical evidence". In: *Journal of Interdisciplinary History* 43.2 (2012), pp. 169–220.

[2] Kyle Harper. "People, plagues, and prices in the Roman world: the evidence from Egypt". In: *The Journal of Economic History* 76.3 (2016).

[3] H.O. Pörtner et al. *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Camridge University Press, 2022.

[4] Carlo Giupponi. "Integrated modelling of social-ecological systems for climate change adaptation." In: *Socio-Environmental Systems Modelling* 3 (2022).

[5] Mike Hulme. "Reducing the future to climate: a story of climate determinism and reductionism". In: *Osiris* 26.1 (2011), pp. 245–266.

[6] George C.D. Adamson. "Re-thinking the present: the role of a historical focus in climate change adaptation research." In: *Global Environmental Change* 48 (2018), pp. 195–205.

[7] Emmanuel Le Roy Ladurie and Barbara Bay. *Times of feast, times of famine: a history of climate since the year 1000*. The Noonday Press, 1971.

[8] Christian Pfister. "The vulnerability of past societies to climatic variation: a new focus for historical climatology in the twenty-first century". In: *Climatic Change* 100.1 (2010), pp. 25–31.

[9] Bruce W. Frier. "More is worse: some observations on the population of the Roman Empire". In: *Debating Roman Demography*. Ed. by Brill. 2001.

[10] B. J. Dermody and R. P. H. van Beek et al. "A virtual water network of the Roman world." In: *Hydrology and Earth System Sciences* 18.2 (2014), pp. 5025–5040.

[11] Alan Bowman and Andrew Wilson. *Quantifying the Roman economy: methods and problems*. Oxford University Press, 2009.

[12] Fredrik Charpentier Ljungqvist. "A new reconstruction of temperature variability in the extra-tropical northern hemisphere during the last two millennia". In: *Geografiska Annaler: Series A, Physical Geography* 92.3 (2010), pp. 339–351.

[13] Stéphanie Desprat et al. "Revealing climatic variability of the last three millennia in northwestern Iberia using pollen influx data." In: *Earth and Planetary Science Letters* 213.1-2 (2003), pp. 63–78.

[14] Ulf et al Büntgen. “2500 years of European climate variability and human susceptibility”. In: *Science* 331.6017 (2011), pp. 578–582.
[15] Gerard Bond et al. “Persistent solar influence on North Atlantic climate during the Holocene”. In: *Science* 294.5594 (2001), pp. 2130–2136.
[16] B. Lee Drake. “Changes in North Atlantic oscillation drove population migrations and the collapse of the Western Roman Empire”. In: *Scientific Reports* 7.1 (2017).
[17] Walter Scheidel. *The science of Roman history: biology, climate, and the future of the past*. Princeton University Press, 2018.
[18] Kyle Harper. *The fate of Rome: Climate, disease, and the end of an empire*. Princeton University Press, 2018.
[19] Lukas De Blois. *The Impact of the Roman Army (200 BC–AD 476): Economic, Social, Political, Religious and Cultural Aspects*. Brill, 2007.
[20] Patricia Southern. *The Roman Empire from Severus to Constantine*. Routledge, 2015.
[21] Stephen Mitchell and Geoffrey Greatrex. *A History of the Later Roman Empire*. John Wiley Sons, 2023.
[22] Arnold Jones. “The third century crisis in the Roman Empire”. In: *Bulletin of the John Rylands Library* 58.2 (1953).
[23] Christer Bruun. “The Antonine plague and the ‘third-century crisis’”. In: *Crises and the Roman Empire* (2007).
[24] Brandon T. McDonald. “The Antonine Crisis: Climate Change as a Trigger for Epidemiological and Economic Turmoil”. In: *Climate change and Ancient Societies in Europe and the Near East*. Ed. by Palgrave Studies in Ancient Economies. 2021.
[25] John Haldon. “Plagues, climate change, and the end of an empire. A response to Kyle Harper’s The Fate of Rome (2): P lagues and a crisis of empire”. In: *History Compass* 16.12 (2018).
[26] Sabine R. Huebner. “Climate Change in the Breadbasket of the Roman Empire—Explaining the Decline of the Fayum Villages in the Third Century CE”. In: *Studies in Late Antiquity* 4.4 (2020).
[27] kaur Harmanjot et al. “A review: Epidemics and pandemics in human history”. In: *International Journal of Pharma Research and Health Sciences* 8.2 (2020).
[28] David Wagner et al. “Yersinia pestis and the Plague of Justinian 541–543 AD: a genomic analysis”. In: *The Lancet infectious diseases* 14.4 (2014).
[29] Walter Scheidel. “Demographic and economic development in the ancient Mediterranean world”. In: *Journal of Institutional and Theoretical Economics (JITE)* 160.4 (2004), pp. 743–757.
[30] Géza Alföldy. “The Crisis of the Third Century as seen by Contemporaries”. In: *Greek, Roman, and Byzantine Studies* 15.1 (1974), pp. 89–111.

[31] David Harms Holt. “Germania and Climate Variability in 3rd and 4th Centuries AD: A Methodological Approach to Dendroclimatology and Human Migration”. In: *Physical Geography* 32.3 (2011).
[32] Peter Heather. “The Huns and the end of the Roman Empire in western Europe”. In: *he English Historical Review* 110.435 (1995).
[33] Grant J. Couper. “Gallic Insurgencies? Annihilating the Bagaudae”. In: *Brill's companion to insurgency and terrorism in the ancient Mediterranean* (1989).
[34] Doi Masaoki. “Peasants and bagaudae in Roman Gaul”. In: *Classical views* 28.3 (1989).
[35] Doi Masaoki. “Bagaudes Movement and German Invasion”. In: *Klio* 71.17 (1989).
[36] Paul Kelly. “Third-Century Inflation Reassessed”. In: *Theoretical Roman Archaeology Journal* 4.1 (2021).
[37] Alfred Wassink. “Inflation and financial policy under the Roman Empire to the price edict of 301 AD”. In: *Historia: Zeitschrift fur Alte Geschichte* (1991), pp. 465–493.
[38] Taco Terpstra. “Communication and roman long-distance trade”. In: *Communication in the Ancient World* (1991), pp. 45–61.
[39] Walter Scheidel. “The shape of the Roman world: modelling imperial connectivity.” In: *Journal of Roman Archaeology* 27 (2014), pp. 7–32.
[40] Piero Lionello, Paola Malanotte-Rizzoli, and Roberta Boscolo. *Mediterranean climate variability*. Elsevier, 2006.
[41] Paul Erdkamp. *The grain market in the Roman Empire: a social, political and economic study*. Cambridge University Press, 2005.
[42] Tyler Franconi et al. “From Empire-wide integration to regional localization: A synthetic and quantitative study of heterogeneous amphora data in Roman Germania reveals centuries-long change in regional patterns of production and consumption”. In: *Plos one* 18.1 (2022).
[43] Ryan M. Geraghty. “The Impact of Globalization in the Roman Empire, 200 bc—ad 100.” In: *The Journal of Economic History* 67.4 (2007), pp. 1036–1061.
[44] Giovanni Geraci. “Feeding Rome: the grain supply.” In: *A Companion to the City of Rome 101*. 2018, pp. 219–245.
[45] Peter Garnsey and Richard Saller. *The Roman Empire: economy, society and culture*. University of California Press, 2015.
[46] G. E. Rickman. “The grain trade under the Roman Empire”. In: *Memoirs of the American Academy in Rome* 36 (1980), pp. 261–275.
[47] Xavier Rubio-Campillo et al. “Bayesian analysis and free market trade within the Roman Empire”. In: *Antiquity* 91.359 (2017), pp. 1241–1252.
[48] Xavier Rubio-Campillo et al. “The ecology of Roman trade. Reconstructing provincial connectivity with similarity measures”. In: *Journal of Archaeological Science* 92 (2018), pp. 37–47.

[49] José Maria Blazquez. "The latest work on the export of Baetican olive oil to Rome and the army". In: *Greece Rome* 39.2 (1992), pp. 173–188.
[50] Alexander Conison. "The Organization of Rome's Wine Trade". Doctoral dissertation. University of Michigan, 2012.
[51] Nicolas Bernigaud et al. "Understanding the development of viticulture in Roman Gaul during and after the Roman climate optimum: The contribution of spatial analysis and agro-ecosystem modeling". In: *Journal of Archaeological Science: Reports* 38.103099 (2021).
[52] Caroline Cheung. "Managing food storage in the Roman Empire." In: *Quaternary International* 597 (2021), pp. 63–75.
[53] Hanspeter Holzhauser et al. "Glacier and lake-level variations in west-central Europe over the last 3500 years". In: *The Holocene* 15.6 (2005).
[54] Feng Shi et al. "Roman Warm Period and Late Antique Little Ice Age in an earth system model large ensemble". In: *Journal of Geophysical Research: Atmospheres* 127.16 (2022).
[55] E. Xoplaki et al. "Hydrological Changes in Late Antiquity: Spatio-Temporal Characteristics and Socio-Economic Impacts in the Eastern Mediterranean". In: *Climate change and Ancient Societies in Europe and the Near East*. Ed. by Palgrave Studies in Ancient Economies. 2021.
[56] B.J.Dermody et al. "A seesaw in Mediterranean precipitation during the Roman Period linked to millennial-scale changes in the North Atlantic." In: *Climate of the Past* 8.2 (2012), pp. 637–651.
[57] Gerard Bond et al. "The North Atlantic's 1-2 kyr climate rhythm: relation to Heinrich events, Dansgaard/Oeschger cycles and the Little Ice Age". In: *Geophysical Monograph-American Geophysical Union* 112 (1999), pp. 35–58.
[58] Feng Shi et al. "Evolution of the solar irradiance during the Holocene". In: *Astronomy Astrophysics* 531 (2011).
[59] M. Sling et al. "Timing and climate forcing of volcanic eruptions for the past 2,500 years". In: *Nature* 523 (2015).
[60] Heinz Waneer et al. "Mid-to Late Holocene climate change: an overview". In: *Quaternary Science Reviews* 27.19-20 (2008).
[61] Heinz Waneer et al. "Holocene climate variability and cultural evolution in the Near East from the Dead Sea sedimentary record". In: *Quaternary Research* 66.3 (2008).
[62] Jesper Olsen. "Variability of the North Atlantic Oscillation over the past 5,200 years". In: *Nature Geoscience* 5.11 (2012).
[63] J. W. Hanson and S. G. Ortman. "A systematic method for estimating the populations of Greek and Roman settlements". In: *Journal of Roman Archaeology* 30 (2017), pp. 301–324.
[64] Brian J. Dermody et al. "A Model of Grain Production and Trade for the Roman World". In: *Simulating Roman Economies: Theories, Methods, and Computational Models*. Ed. by Oxford University Press. 2022.

[65] Walter Scheidel and Elijah Meeks. *ORBIS: the Stanford Geospatial Network Model of the Roman World*. Available at https://orbis.stanford.edu/. URL: https://orbis.stanford.edu/.

[66] W3Techs. *PBL Netherlands Environmental Assessment Agency*. Available at https://www.pbl.nl/en/image/links/hyde. URL: https://www.pbl.nl/en/image/links/hyde.

[67] Kees Klein Goldewijk. "Estimating global land use change over the past 300 years: the HYDE database". In: *Global biogeochemical cycles* 15.2 (2001).

[68] *Global Land Cover Characterization (GLCC)*. URL: https://www.usgs.gov/centers/eros/science/usgs-eros-archive-land-cover-products-global-land-cover-characterization-glcc#overview.

[69] L. P. H. van Beek et al. "Global monthly water stress: Water balance and water availability". In: *Water Resources Research* 47.7 (July 2011).

[70] Helen Goodchild and Robert Witcher. "Modelling the agricultural landscapes of Republican Italy". In: *Agricoltura E Scambi Nell'italia Tardo Repubblicana*. 2010.

[71] Tim Unwin. *Wine and the vine: an historical geography of viticulture and the wine trade*. Routledge, 2005.

[72] Berry Arnold. *Pareto distributions*. International Co-Operative Publishing House, Fairland, 1983.

[73] Jonathan Prag. "Sicily and Sardinia-Corsica: the first provinces". In: *A Companion to Roman Imperialism*. Ed. by Brill. 2013.

[74] Paul Arthur and David Williams. "Campanian wine, Roman Britain and the third century AD". In: *Journal of Roman Archaeology* 5 (1992), pp. 250–260.

[75] Alan Bowman and Andrew Wilson. *The Roman agricultural economy: organization, investment, and production*. Oxford University Press, 2013.

[76] P.A. Brunt. *Italian Manpower, 225 B.C.–A.D. 14*. Clarendon Press, 1971.

[77] Matias Berthelon and Caroline Freund. "On the conservation of distance in international trade". In: *Journal of International Economics* 75.2 (2008).

[78] Paul Kelly. *The international trade network in space and time*. http://dx.doi.org/10.2139/s 2012.

[79] Danielle Bonneau. *Le fisc et le Nil: incidences des irrégularités de la crue du Nil sur la fiscalité foncière dans l'Égypte grecque et romaine.* Cujas, 1971.

[80] Nicholas Purcell. "Wine and wealth in ancient Italy". In: *The Journal of Roman Studies* 75 (1985), pp. 1–19.

[81] Arnold Jones. "Inflation under the Roman Empire". In: *The Economic History Review* 5.3 (1953).