Adopting a new technology in the Bronze Age - Examining the transmission and adoption of ceramic innovations in Western Anatolia and the Aegean using a cultural evolutionary framework

Christina Alam

A thesis submitted for the degree of Doctor of Philosophy

Institute of Archaeology

UCL

February 2024

I, Christina Alam, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

During the Bronze Age, novel cultural traits were transmitted between potters in the Aegean, including pottery-making techniques that ethnographic research suggests were most often taught from parents (or other close relatives of the parental generation) to offspring. Following a phase of slow adoption, they were rapidly adopted in the context of palace and elite-consumed pottery. I investigate how geography, population size, and community-level factors such as elite demand affected innovation adoption using cultural evolutionary theory and simulation models.

I simulate trait transmission through a network and explore both ‘toy’ and ‘realistic’ models. Using toy models, I describe the dynamics emerging from the assumptions of the model, which incorporates the effects of travel costs, the local rate of innovation adoption (LRA, which could represent multiple factors affecting the local rate, e.g., a measure of elite control over local adoption or the presence of social boundaries affecting inter-group adoption), population size, and random trait loss. Their outputs suggest that relative population clustering affects the shape of the adoption curve and that different combinations of travel costs, LRA, and population density result in spatial diffusion propagating along different types of wavefronts.

‘Realistic’ models are run on networks reflecting the Aegean’s interesting geography, consisting of two land masses connected by an archipelago, with natural coastal and inland topographies that create relatively isolated sub-populations; they additionally incorporate the effects of different landscape types and the location of the innovation’s origin. Their outcomes are quantitatively compared to empirical adoption curves for the potter’s wheel innovation and suggest that the

location of origin can determine the shape of the curve, especially its position on the East – West axis defined by the landmasses of Greece and Anatolia.

The study showcases the potential of simulation modeling for addressing questions about cultural transmission using archaeological data. It shows that abstract models with empirically supported assumptions can be informative, even when the empirical data are recorded at varying levels of granularity and originate from a scholarly tradition that might otherwise inhibit large-scale quantitative approaches.

Impact Statement

Within academia, this thesis will provide useful insights to all archaeologists interested in the diffusion of innovations in space and time, including pottery specialists who are interested in computational models addressing large scale and geographic factors affecting the transmission of techniques between potters. The study will also be of interest to archaeologists working in the Aegean who want to use simulation modeling to enhance their understanding of long-term change, and particularly those who want to use legacy data or data from secondary sources such as published studies. The thesis also produces insights regarding intercultural exchange and its relationship with technological innovation that could be of interest to them. Outside the discipline, the insights from this study will benefit researchers working in disciplines such as cultural evolution, who are interested in the transmission of information between networks of communities when the strength of the communication between them is affected by their distance in space.

Outside academia, the thesis can inform the design of educational tools in museum settings that enable audiences to visualise and experiment with models of communication networks between past communities and explore how such networks affected the evolution of technology and material culture in the past.

To Vasso, Maria, Shamsul, and Nasos

Contents

Acknowledgements

This study would not have been possible without the support of my supervisors Stephen Shennan, Mark Lake, and Andy Bevan. I would like to thank them for believing in the potential of the project, even when the challenges that research, life, and the pandemic brought my way filled me with doubt. They were the best mentors I have ever had, and they deeply shaped the way I think about archaeology, giving me the tools to pursue my research questions in the face of archaeological uncertainty. I will always treasure our intellectually stimulating conversations which made me identify and improve weaknesses in my thinking and develop confidence in my own voice. Working with them was truly a privilege and any mistakes included in this thesis are entirely my responsibility. I would also like to thank James Conolly and Borja Legarra Herrero for being excellent examiners and for suggesting new ways to think about my research.

The Institute provided a very nurturing environment, and I felt like I found my intellectual home there. I'm very grateful for the opportunity to have valuable conversations on both the computational and the pottery-related aspects of my research with Mark Altaweel, Kevin MacDonald, Patrick Quinn, Bill Sillar, and Ulrike Sommer. I would also like to thank Lisa Daniel for providing administrative support.

I would especially like to thank Bill Sillar for giving me the opportunity to help him teach the 'Sites and Artefacts' and 'Interpreting Pottery' courses which were incredible learning experiences – Bill's insights regarding the technical and social aspects of pottery production influenced my modeling approach, and I found his energetic delivery of lectures and his generosity towards his students to be deeply inspiring. Being a TA for 'Sites and Artefacts' also gave me the opportunity to teach alongside Charlotte Frearson, Ceri Shipton, and Paul Wordsworth and learn from their approach to teaching.

The project also benefitted from the input of archaeologists specializing in Aegean pottery and computational archaeology. I would like to thank Evangelia Kiriatzi for an incredibly valuable conversation which shaped the course of the project, where she suggested examining the transmission of multiple technological traits (iron-reducing decoration and kilns) in addition to wheel-made pottery. She was extremely generous with her suggestions and time. I would also like to thank Enrico Crema for generously providing feedback during the first stages of this project's conception.

Andy Bevan with Brenna Hassett, Olga Vassi (Greek Archaeological Service) and Despoina Tsardaka (Chios Ephorate) gave me the opportunity to participate in the Emborio Hinterland project, where I got to see some of the pottery from Chios and the vicinity of Emborio in person during the surface survey. I learnt a lot from the project, and I was also lucky enough to have conversations on prehistoric pottery production with Jannis Kozatsas who generously shared his expertise — I would also like to thank Jannis for teaching me basic pottery forming techniques (he is not at all to blame for my being so bad at it!). This project was also a refreshingly social experience which was much needed in the middle of an isolating pandemic PhD. It meant doing research with a team of highly capable and kind researchers and making new friends. Anastasia, Flora, Giorgos, Juan, Katerina, Matt, and Vicky were wonderful team members and I really admire their enthusiasm and dedication for computational and field-based research.

I would like to thank my wonderful friends Dafni, Dimitra, Elena, Fil, Giorgos, Katerina, Marianna, Marisia, Monal, and Morgan for supporting me throughout the PhD with fun conversations, long phone calls and warm hugs and for being patient with my PhD-related disappearing acts and frequent complaining.

I owe the biggest thanks to my parents and to my sister for supporting me emotionally and financially, for always being my shoulder to lean on, and for taking care of my dog Tsoureki when I was away from home. They have always supported my dreams and made the biggest sacrifices for my well-being and happiness — I don’t know how I could ever repay them. I also owe a big thanks to Kalliopi Notias, without the support of whom I would not have been able to start this PhD.

Finally, I want to thank Athanasios Giannenas for being such an amazing partner. He supported me through all the hard times both practically and emotionally and he has been a constant source of inspiration for me with his enthusiasm for simulating and understanding complex systems. He is also a fantastic teacher who helped me understand basic concepts in computer programming and mathematical modelling. Thanks for everything and I’m looking forward to our future adventures!

1. Introduction

1.1. Problem statement

This study explores the transmission of ceramic innovations in the Aegean during the Bronze Age and attempts to trace adoption trajectories over time for three novel traits that appeared over this period: the use of the potter's wheel for forming, the use of kilns for firing, and the use of iron-oxide based slips for decorating pottery. The study focuses on data reuse in that it leverages previously published data in order to construct empirical estimates of adoption over time. Its main contribution is in the development of simulation models that attempt to generate the phenomenon of interest in an 'in-silico' setting and explain the adoption and transmission process by exploring the simulations under different conditions. The code for all simulations and analyses is available at: https://github.com/christinaalam/Innovation_transmission_Aegean and is written in R (R Core Team 2021), an open-source environment that allows other researchers to explore and re-evaluate the model.

As will be discussed at length in chapters 4 and 5, the diffusion of the potter's wheel is the best documented case of innovation adoption in the Aegean literature, and the observation period and spatial extent of the study area was selected primarily with a view to tracing this process. Specifically, the observation period starts from the Early Bronze Age II period, when the first traces of the wheel appear on Kastri/Lefkandi I pottery (Choleva 2012) and lasts until the Late Bronze Age III phase. The study area includes mainland Greece, the Aegean islands, Crete, and western Anatolia. This spatiotemporal context offers the opportunity to study the adoption of a technological innovation in a geographic region characterized by boom-and-bust population cycles

(Dalfes *et al.* 2013), migrations and, especially after the establishment of palaces, varying degrees of urban organisation (Branigan 2001) which may or may not have affected diffusion rates on multiple scales; on the micro scale, using the "site" as a unit of analysis and the macro scale, which treats the entire Aegean study area as a system.

Social context appears to have been a major factor affecting the cultural transmission of traits over the Eastern Mediterranean: in some regions, like the Northern Levant and Crete, the wheel was progressively adopted, with quantities, shapes, sizes and functional categories of wheel-made pottery increasing gradually over time to almost completely replace handmade techniques, whereas in other regions, like the Southern Levant and Cyprus, the wheel remained a minority forming technique (Roux & Jeffra 2015), as shown in Figure 1. Roux attributes this resistance to innovation to the existence of clear social boundaries that did not allow forming techniques to be transmitted between different ethnic groups (Roux & Jeffra 2015: 178).

A similar pattern of inhomogeneous adoption arises when the same process is examined on a smaller scale. Knappett (2016) distinguished between two processes of diffusion for the potter's wheel in the Aegean: the first, through contacts with Anatolia in a time of westward migrations and associated with the Kastri/Lefkandi I feasting wares, and a second wave through contacts with Crete, where the new technology was already well established. The former seems to be in line with the pattern observed for Cyprus and the Southern Levant, and the wheel is in some cases rejected after an initial period of adoption. The latter is considered a more "successful" example of adoption, and the wheel becomes "embedded" (Knappett 2016: 107) in local production.

Roux has suggested that the initial uptake of innovations depends on socioeconomic elites, who can bear the cost of supporting standardised production (Roux 2013). Elite transmission bias

(Henrich 2001) can also account for the fact that "elite" wheel-made feasting wares made up only a small percentage of local pottery assemblages in the initial stages of adoption.

It is also generally acknowledged that there was a spatial component to the adoption process for the potter's wheel innovation, since its origin is frequently placed in Western Anatolia (Choleva 2020: 64) and in the literature its transmission is sometimes seen as a result of the movement of craftspeople between different communities in the Aegean. The Bronze Age was a time when technological improvements in sailing technology (Broodbank 2002) increased the flow of communication between communities separated by sea barriers. Therefore, it is very likely that geographic boundaries and the costs of travelling associated with different sailing technologies affected the way and speed in which new ideas and technologies were taught between craftspeople living in different communities. It is also possible that the number of communities and craftspeople affected the direction and speed of the transmission of new traits between communities of potters – the importance of large population sizes for the emergence and spread of technological innovations has been shown through both empirical studies and theoretical modelling (this is discussed in more detail in chapter 2).

Finally, it is possible that the diffusion of novel traits in pottery depended on the interaction between different stages of the operational chain of pottery production. For example, the fact that finely levigated clay (clay where naturally occurring inclusions have been removed) is often used to produce wheel-made ceramics, could suggest that the potter's wheel could not be successfully transferred to communities of practice that did not possess the 'know-how' to produce finely levigated fabrics. A dependence between traits in transmission processes has already been studied at an abstract level through theoretical and empirical models of cultural linkage (Yeh et al. 2019; Morin & Milton 2018) and thoroughly discussed in ethnographically grounded archaeological discourse (Sillar & Tite 2000). As will be shown in Chapters 4 and 5, the adoption processes for kilns and iron-oxide based slips have received less attention and are less clear to us as a result. I have attempted to bring together information from the literature to suggest possible adoption trajectories in space and time.

This study adopts a cultural evolutionary framework which is discussed at length in Chapter 2, and the specific approach as to how cultural evolutionary theory informed the central assumptions of this research is described in Chapter 6. Adopting such a theoretical perspective is a clear deviation from the stance typically taken in Aegean archaeology, where there is a tendency to associate evolutionary frameworks with a disregard for "human ontogeny" (Knappett & Van der Leeuw 2014: 67; Wimsatt and Griesemer 2007: 227). As shown in Chapter 2, this framework may focus on system-level outcomes, but it is equally concerned with developing microlevel foundations for the behaviour of individuals that gives rise to said system-level outcomes. Insights from cultural evolution and ethnoarchaeology (in particular, studies using the chaîne opératoire framework) will be used to inform assumptions about the mechanisms of cultural transmission and the "regularities" (Roux 2019: 260) that permeate the decisions of potters and consumers.

In terms of the empirical evidence available for the Aegean, an increasing number of pottery specialists (Berg 2007; Berg 2011; Berg 2015; Caloi 2011; Choleva 2012; Choleva 2015; Choleva 2018; Jeffra 2013; Gorogianni et al. 2016; Spencer 2007) have been using Roux and Courty's identification framework which allows identifying variations in the forming technique and its use of Rotative Kinetic Energy from a rotative device (Roux & Courty 1998). Identification is either achieved through macroscopic visual inspection or a combination of macroscopic and microscopic methods with the aim to quantify the prevalence of each forming "method" for different chronological phases per archaeological site. X-radiography has also been an effective method for identifying forming techniques in a relatively cost-effective way, with the added advantage that it is possible to identify primary forming techniques with greater precision than through simple macroscopic analysis, where it is difficult to discern between diagnostic features of surface treatment and secondary forming from primary forming technique (Berg 2011: 2). Due to their focus on quantification, such studies have formed the backbone of the empirical dataset that was used to compare the simulation models to reality.

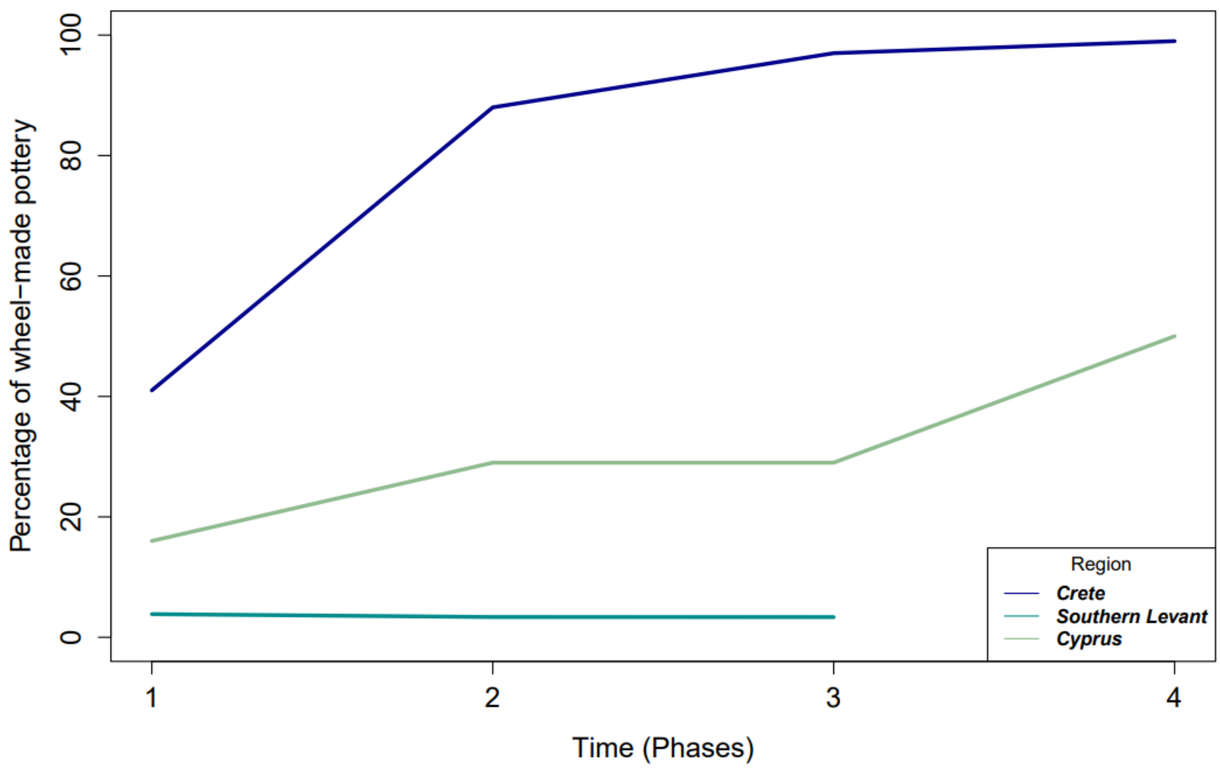

Figure 1: Percentages of wheel-made pottery per phase in Crete, Cyprus, and Southern Levant. Phases in Crete: Middle Minoan IB-Late Minoan IA (a process lasting for approximately 300 years from 1900 to 1600 BC), Cyprus: Late Cypriot IA-Late Cypriot IIB (a process lasting for approximately 400 years from 1600 to 1200 BC), Southern Levant: EBA IIIA-IIIC (a process lasting for approximately 500 years from 2700 to 2200 BC). Data from Roux & Jeffra (2015:170,172,175).

1.2. Research questions

Bringing together the insights above, it is possible to formulate the following research questions which focus on macroscopic system-level factors but also consider the role of community-level decisions.

1. What was the role of geography, inter-community distance, and population size in the adoption of ceramic innovations in the Bronze Age Aegean?
2. What was the relative contribution of geography, population size, and local response to innovation that caused the adoption patterns for ceramic innovations in the BA Aegean?".
3. Was the trajectory of innovation adoption affected by the decisions of local elites?

1.3. Thesis structure

Chapters 2 and 3 describe the theoretical framework and epistemological approach of the study. Chapter 2 introduces the theoretical framework that informed the study's assumptions and general inferential strategy. It focuses on two strands of theory: cultural evolution and insights from ethnoarchaeological studies. Chapter 3 then takes an epistemological approach and shows that simulation models can provide different types of explanations for the research questions that were posed here.

Chapters 4 and 5 examine the empirical data in greater depth. Chapter 4 attempts to describe the adoption process for the three innovations through the lens of the literature on the Aegean Bronze Age with a qualitative approach. It also includes a review of some of the influential studies on the topic. Chapter 5 focuses on the potter's wheel and kiln data and integrates legacy data to produce three datasets that are then analysed with Monte Carlo methods (used for the potter's wheel data) and aoristic analysis (used for kiln data) in order to produce empirical trend lines for the adoption trajectory of each technological trait over time.

Chapters 6, 7, and 8 describe the design and development of a simulation model and attempt to throw light on the innovation adoption process by running the model under different conditions. Chapter 6 describes the design of the model focusing on aspects of the model's ontology and the main assumptions that inform the behaviour of the entities in the world of the model and how they relate to the theoretical framework of the study. Chapter 7 explores 'toy' versions of the model which are relatively simple in terms of the complexity of the environment of the simulation. Insights drawn from such abstract modelling are meant to sharpen our intuitions about the spatial transmission processes in various environments. Finally, Chapter 8 involves running the model under more 'realistic' conditions in an environment that incorporates aspects of Aegean geography such as natural barriers and landscape zones associated with different probabilities of habitation. The chapter includes various experiments that explore the role of geography, communication costs, and community level factors affecting trait adoption and their outputs are quantitatively compared to empirical trend lines for the potter's wheel data.

Chapter 9 is the concluding chapter of the thesis, and presents the insights gained from the modelling process in light of the results presented in chapters 7 and 8 and a critical discussion on the nature of the insights drawn and how they address gaps in the literature. Other points of discussion include the compatibility of legacy data from the Aegean literature with simulation modelling and ways that the study could be extended in other directions to explore the cultural transmission of other types of innovations.

2. Cultural transmission as an evolutionary process

2.1. Culture evolves

Culture is a complex phenomenon which varies greatly with respect to space and time. While empirical observation of the material manifestations of cultural variation is relatively straightforward, identifying the source(s) of this variation is more challenging, as is identifying mechanisms which ensure the survival of variations over a timespan exceeding one generation.

In order to study cultural change, it is first necessary to select a working definition of culture that adequately summarises the characteristics of culture that can be observed empirically. As pointed out by Mesoudi (2016), culture can be described as "socially transmitted information" or in the words of Richerson and Boyd (1984: 430) "information acquired by imitating or learning from other individuals and able to affect an individual's phenotype, usually behaviour". Seen under this light, variations in cultural traits can be brought about by a process of "descent with modification" from the parental to the next generation, in a similar fashion to Darwin's evolutionary process described in the "Origins of Species". The similarities and differences between biological and cultural evolution were examined in a systematic way by studies in the 'Dual Inheritance Theory' tradition, which sought to describe the interaction of these two strands of "inheritance" with mathematical models (Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981).

Cultural evolution places emphasis on different modes of social learning as the transmission mechanism (Shennan 2012: 15) that introduces variation in cultural traits within a given population and affects their selection. In an archaeological context, this involves identifying cognitive biases that result in an alteration of the original piece of information through ethnoarchaeology, experiment and ethnography (Shennan 2012: 16). Under the Dual Inheritance model of biological and cultural evolution (Richerson and Boyd 1985:285), processes of cultural evolution can generate cultural variation and the frequency of cultural variance can be affected by natural selection. Biological and cultural evolution can be 'asymmetric' in that they can have different transmission patterns — in general, in biological evolution the genetic parents have an equal proportion of variation that they inherit to their offspring, but in cultural evolution other members of the parental generation can also transmit, even to a higher extent than the genetic parents, and it is also possible for one parent to contribute more than the other parent (Richerson and Boyd 1985:286). It is possible for cultural traits to be naturally selected or culturally selected. In Richerson and Boyd's words "the most interesting effect of asymmetric transmission is the way natural selection can act on asymmetrically transmitted variants. If models occupying social roles other than that of the genetic parent are important in cultural transmission, and if cultural variation affects which people attain these roles, then selection can increase the frequency of cultural variants that reduce ordinary genetic fitness (Richerson and Boyd 1985:286).

The framework generally places a strong emphasis on transmission as a replication process, which has attracted criticism from scholars advocating for an "Epidemiology of Representations" approach (Sperber 1985: 74-75), which calls for more systematic study of cognitive micro-processes and cognitive causal chains which give rise to macro-scale phenomena like cultural traditions, with Sperber's epidemiology acting as an equivalent to diffusion, a notion frequently encountered in the cultural evolution literature. This strand of theory is also sometimes referred to as Cultural Attractor Theory (Buskell 2017). While both approaches hold relative advantages and disadvantages as to their explanatory power, the focus here is on Dual Inheritance Theory and cultural evolution in general, because of the greater emphasis placed on empirical observations of material culture and population level patterns, which makes it more compatible with population and “time-averaged” (Premo 2014: 105) archaeological data.

This definition of culture encompasses all forms of socially transmitted information, including animal cultures. In this respect, research in anthropology has focused on identifying what makes human culture advantageous in terms of survival and reproduction. Some believe the ‘secret’ lies in the cumulative nature of human culture (Henrich et al 2008; Henrich, 2016), where new traits, which frequently have reproductive advantages, are copied or creatively recombined to produce new traits, as is the case with composite technologies like the bow and arrow (Lewis and Laland 2012), and accumulated over time, creating new selection pressures for populations. New traits are copied between peers through horizontal transmission and intergenerationally through vertical and oblique transmission (vertical transmission occurs when teachers and learners are related). This presents the relative advantage of not having to assess the costs and benefits of new traits, e.g. survival strategies, manufacturing techniques, through individual learning processes like trial-and-error experimentation, which is time consuming (Henrich, 2016) and often dangerous — a good example provided by Henrich is a review of the adaptational advantages of food taboos, which protect populations from consuming dangerous substances without them having to resort to dangerous experimentation. This process of relatively random experimentation, followed by the elimination of maladaptive behaviours or “pieces of knowledge”, is consistent with concepts from Campbell’s (1974) “evolutionary epistemology”, where blind variation is followed by “selective retention” (Campbell, 1960: 380-381) of units of knowledge.

The cultural evolutionary literature acknowledges that, in reality, variation is not likely to be completely random and that the blind variation model is only relevant in certain contexts (Mesoudi 2021: 16) but ignores directionality which is an important factor affecting cultural variation in other contexts. The reliance on blind variation is also a key point of contention between the cultural evolutionists following Boyd and Richerson's (1985) framework and the theorists following Cultural Attractor Theory (CAT), who place more emphasis on the cognitive aspects affecting variation-producing processes - Morin (2016) summarises the debate from the perspective of CAT in a discussion where he suggests multiple points of contention with the theorists following Boyd and Richerson's (1985) framework. According to Morin (2016:451), "cultural attraction theorists believe that cultural evolution is driven by a myriad of widely shared tastes and preferences that bring people in many different societies to favour similar traditions". Morin seems concerned about the emphasis placed by the Richerson and Boyd (2005) framework on the "dichotomy" between context and content-dependent biases, which he believes are not the major factors affecting the copying process that underpins cultural evolution (Morin 2016: 452). the distinction between the two is not as clear-cut as the Richerson and Boyd framework suggests (Morin 2016: 453).

It is important to make a distinction about whether evolution, in terms of processes of drift and selection, acts on cultural traits or on the human actors who produce them. In this case, as argued by Shennan, cultural evolution operates on both levels, and it is important to examine lineages of traits, such as technological lineages, and the human populations who produced them independently, because different evolutionary processes may be identified for each one. From the trait or the "technical lineage's" point of view, transmission is always vertical (Shennan, 2013: 145). Especially regarding archaeology, it may be argued that identifying technological lineages may be easier, due to the nature of the data that usually survives taphonomic processes than reconstructing genetic lineages, although this could change given the advance of archaeogenetic research.

One mechanism for introducing variation in cultural evolution, that acts in a similar way to mutation in biological evolution, is innovation (O'brien & Laland, 2012: 437). Again, there are many definitions of what constitutes an innovation, and often it is necessary to distinguish it from invention. Generally, invention is seen as the appearance of a novel trait, and innovation is seen as the diffusion of a variation of an existing trait through a population, an evolutionary process which renders this variation subject to selection pressures (Walsh *et al*. 2019: 51).

Another way to define invention is as an accumulation of small variations in the micro-scale that gives rise to new forms in the macro scale, a view that links the introduction of new traits to the tempo and scale of evolutionary processes. This, in turn, has given rise to the view that changes in the tempo of cultural evolution may result from shifts in population size, with larger populations acting as "distributed minds", in the sense that they consist of specialists whose combined efforts result in an exponential increase of new traits (Bentley & O'Brien 2012: 3-4).

Henrich (2001), attempted to model the diffusion of innovations using a mathematical model. Much like the results of a previous model (Bass 1969), the resulting cumulative adoption curve was S-shaped (logistic) in the case of biased transmission, that is, social transmission where a trait is preferred over others because of its prevalence in a population(conformist-bias) or because the person who transmits the traits is preferred over another teacher/person with social prestige (prestige-bias), and r-shaped when individual learning takes place. It must be noted that subsequent research (Steele 2009) demonstrated that other factors, apart from biased social learning can result in S-shaped curves, like economic inequality or a spatial component to diffusion, leading the author to deduce that using a "pattern" to infer a process is misleading, precisely because of this equifinality.

Drift

If invention and innovation are evolutionary forces that introduce variation in cultural systems, an opposing force known as 'drift' "destroys variation" (Neiman 1995: 10). In order to comprehend the notion of drift, it is important to understand the stochastic nature of cultural transmission. Chance plays a big role in this regard — it is not necessarily the case that every teacher in a generation will transmit their trait to a student. It is entirely possible that some teachers will have no students, while others will teach a large number of students. Therefore, it is to be expected that the traits of teachers with no students will disappear in the following generation. This mechanism results in the loss of some traits, and it does not cause the emergence of new traits, so on the whole it tends to reduce variation in cultural systems.

Neiman (1995) employed mathematical simulations to explore the effects of drift on cultural transmission and showed that small populations are more vulnerable to the effects of drift — they lose traits at a faster rate compared to large populations (Neiman 1995:10) — which results in diminished cultural diversity.

The demographic hypothesis

Henrich (2004: 209) and Shennan (2013: 151) have also stressed the importance of an "effective" population size for a trait to be successful in terms of evolution, linking larger population sizes with a higher probability to invent a new trait and a higher probability for a trait to be transmitted, due to the size of the communication network being larger, resulting in more nodes and edges, although in this case network structure is an important thing to consider. Some have argued that the concept of an "effective" population size has varied in the past (Premo, 2016: 609), with an influential definition being that of Neiman (1995), who conceptualized it as the number of individuals possessing a trait that is being copied by the population, pointing to the importance of the number of "teachers" in a population relative to the census size of the population. Premo's (2016) study highlights the importance of a working definition of "population" and indicates three different concepts of population from genetics: census population size, meaning the total number of living individuals at a given time, breeding population size, meaning the individuals who pass on genetic material to subsequent generations (Hedrick 2011:205), and "effective" population size. The definition Premo uses for the effective population size is directly taken from the population genetics literature and "is a measure of the effect of random genetic drift on a population genetic feature of interest in a real population", thus incorporating the effect of drift (Premo 2016: 606).

'Fragile networks', while initially acting as breeding grounds for innovation due to craft specialisation, are highly unstable and prone to collapse and loss of technological knowledge, when central nodes in a network cease to exist. Roux (2008: 98) has made reference to such fragile systems in the context of the transmission of the potter's wheel focusing on the size of the learning network, which if "too small, then the technological system is fragile and will not resist important historical events such as those that transform socioeconomic structures". In other words, historical contingency, is seen as a driving force of technological evolution as it impacts directly on the learning networks that are necessary for transmission to occur. According to Roux, one way to infer the size of learning networks is through using quantitative proxies for producer skill and rates of production (Roux, 2008: 99). The effects of population size on cultural trait accumulation have also been demonstrated through empirically validated (with data from the Late Pleistocene) simulations (Powell et al. 2009: 1301).

Another simulation-based study of innovation by Crema and Lake (2015) focused on transmission and network properties which act as “cultural incubators” for beneficial, neutral, and ‘suboptimal’ innovations and provided further evidence that robust networks, both in terms of structure and size, are crucial in the initial stages of diffusion. The authors re-examined basic transmission processes from Dual Inheritance Theory, such as model-based bias, in terms of their differential evaluation of trait fitness. Under the lens of fitness evaluation, certain transmission processes, such as prestige-biased transmission, allow a very indirect assessment of trait fitness, as agents can base their assessment on the “success” of the model instead of the trait they copy. A “false” assessment could be due to the fact that the model’s success is due to other factors. An increase in payoff uncertainty results in model-based copying being less successful (Crema and Lake 2015: 28)

Furthermore, the authors argue that on the basis of the simulation results, population size alone is not a certain predictor of increased rates of cumulative evolution, as in cases where the payoff of a trait is uncertain, the rate of copying increases even with small population sizes, on the condition that the sample size also shrinks (Crema and Lake 2015: 34). Therefore, population structure is also important, especially concerning the diffusion of beneficial innovations, and smaller groups that act as “cultural incubators” (Crema and Lake 2015: 16) within a larger population are the optimal way for beneficial traits to spread. The results also indicate that decisions taken by those who first adopt new traits are crucial, because opting for “suboptimal” traits instead of beneficial traits can lead to such traits prevailing over beneficial ones (Crema and Lake 2015: 31). In this context, “suboptimal” traits could be traits that are more costly to adopt (e.g., take a long time to learn) relative to the benefits they present (Crema and Lake 2015: 5).

A simulation-based study by Derex et al. (2018) also explored the effects of population structure on cultural evolution, focusing on its effects on innovation rather than cultural transmission. Specifically, the study was concerned with the effects of population fragmentation on cumulative culture, which is the result of the accumulation of cultural traits and their recombination into increasingly complex ones. The authors set up an agent-based model which allowed experimenting with different levels of population fragmentation, migration rates, the probability of innovation, learning error, and the extent to which innovation depends on cultural diversity (Derex et al 2018: 4).

They found that when innovation depends on cultural diversity, "intermediate" levels of population fragmentation can help sustain the emergence of more complex traits or as the authors define it, "cultural complexity" (Derex et al. 2018: 6). While in theory large non-fragmented populations can sustain complex traits (since they are more resilient in the face of random loss of cultural traits), they are less likely to produce them under this scenario compared to moderately fragmented populations. However, fragmentation has the opposite effect when innovation does not depend on cultural diversity (Derex et al. 2018: 5).

Trait linkage and cultural hitch-hiking

An important consideration regarding trait transmission is whether and how cultural traits are linked and how this might affect their evolutionary trajectories. Sometimes referred to as trait "hitch-hiking" in biology (Smith & Haigh 1974) or "wholesale copying" (Morin & Milton 2018: 392), the process refers to traits that are transmitted as part of a package, rather than independently.

In cultural transmission, this means that one trait is being intentionally copied, and other traits linked to it are copied unintentionally. Researchers have recently attempted to identify transmission mechanisms for linked traits empirically and through simulation, as well as define what the implications would be for studies using population-level frequency data to infer transmission mechanisms.

The results of Yeh et al's (2019) simulation-based study used a simple model, where trait linkage occurred randomly, suggest that different frequencies of trait linkage can create very different population-level frequencies for cultural traits. Low trait linkage in a population that can transmit "packages" of varying size can lead to low cultural diversity, while high trait linkage means that most trait are part of a package, which could be considered as a single trait and, therefore, resulting in frequencies similar to unbiased transmission. In other words, this could lead to a situation where a researcher falsely interprets population-level frequencies as an indication of conformism, when in fact, the pattern occurred due to the presence of links. In terms of pay-off, non-functional and even "detrimental" traits can spread because they bear multiple links to functional ones, although this depends on how large the pay-off difference is. To complicate things further, package size also affects diffusion rates (Yeh *et al*. 2019: 6-8).

Niche construction and cultural ecologies

While Dual Inheritance Theory has produced robust models on the evolutionary mechanisms that pertain to the interaction of organisms, genes and cultural traits, less attention has been paid to how they modify their environments. Niche construction theory (Odling-Smee *et al*. 2003: 1-3) was developed to draw attention to the mechanisms through which organisms modify their own habitats and in so doing, create new selection pressures for themselves, their offspring, and others occupying the same habitats, generating adaptational feedback loops. Odling-Smee theorised niche construction using ideas from Lewontin's study (1983) critiquing contemporary evolutionary theory for overemphasizing the passive adaptation of organisms to their environments (Laland et al. 2016: 191).

Special reference will be made here to cultural niche construction, a subfield of niche construction focusing on how such cultural processes, whether intentional or unintentional, create cultural ecologies which affect natural and cultural selection. Interestingly, certain cultural processes seen under this light could be considered to slow down the rate of natural selection (Odling-Smee *et al*. 2003: 337). Specifically in the case of human cumulative culture, niche construction can generate two different feedback routes: route 1, which can result in modifications of cultural selection, and in turn, lead to a new cultural response (or no response), and route 2, which can result in modifications in natural selection, altering the gene pool, which can lead, again, to a cultural response (or no response). Ultimately, the main difference between routes 1 and 2 is that cultural niche construction either has an impact on the human genome (route 2) or not (route 1) (Odling-Smee *et al*. 2003: 338-339).

Route 1 can be used to trace the cultural feedbacks generated by cultural niche construction, in terms of favouring the spread of new technologies at the expense of others (Odling-Smee *et al*. 2003: 339) and is possibly a useful framework for archaeological studies of past technologies.

Scholarly interest in cultural niche construction has varied depending on the type of niche construction and feedback. Shennan's (2011) study on wealth as cultural niche construction shed light on cultural evolutionary processes linked to social institutions, property rights and inter-generational wealth accumulation, a phenomenon linked to the increasing sedentism and agriculture-based wealth brought forth by Neolithization. Wealth inequality is seen as "ecological inheritance" which results in increased reproductive success in a Darwininan sense (Shennan 2011: 918). The focus is on inter-generational wealth, as the niche construction framework is best applied on socioenvironmental modifications that affect multiple generations. Within this framework, the nature of the resources that constitute wealth needs to be such that it is possible to ration them and claim some sort of private ownership over them (Dyson-Hudson & Smith 1978). The role of social institutions, culturally constructed niches in themselves, in this case, is to define who has access and ownership over resources. North (1990), using a behavioural ecology approach, saw property rights as emerging from an assessment of the relative costs and benefits of enforcing them, and constant renegotiation of property generates feedback that triggers renegotiation of social institutions (Shennan 2011: 919). Simulation models for technological niches have also demonstrated that a cost-benefit assessment is a vital component of intentional niche constructing behaviours (Mohlenhoff & Codding, 2017: 224).

It must be noted that while a number of researchers, the aforementioned included, have effectively theorized and modelled the effect of wealth niche construction on the evolutionary and reproductive fitness of humans, ranging from agricultural practices (Netting 1981), social monogamy (Fortunato & Archetti 2010), and aristocratic endogamy (Summers 2005), less attention has been given to how wealth and social inequality niches affect cultural selection processes from the perspective of traits, which could be a fruitful avenue for future research, especially on the subject of technological lineages. Diffusion theory (Rogers 2003) has indicated that early adopters of innovations are frequently members of socio-economic elites, and comparative archaeological research on Bronze Age Danish and Korean metalworking (Kim 2001: 473) has indicated that their decisions at early stages of adoption determine the diffusion rates, which points to an important connection between socio-economic niche construction and the cultural evolutionary trajectory of technology.

It appears that there are methodological issues that need to be addressed for cultural niche construction to be formally incorporated in cultural evolutionary analysis. As Yeh and colleagues argue (2019: 2), niches involve multiple traits, making further research into cultural linkage necessary.

Archaeological case studies of cultural transmission

Within archaeology, there is a growing number of studies of cultural transmission that aim to infer the social learning processes that must have taken place to produce observed frequencies of cultural traits whenever those frequencies deviate from the outputs of a null model, that is, frequencies that could have resulted from a purely random process, e.g. (Shennan & Wilkinson 2001; Bentley & Shennan 2003). These studies, both using empirical data in the form of frequencies of decorative motifs of Linearbandkeramik pottery from Merzbach valley, Germany, deduced different mechanisms of cultural transmission took place. Shennan and Wilkinson’s (2001) study produced results that suggested that anti-conformism, a preference for less common traits to be copied, was the main mode of cultural transmission, while Bentley and Shennan’s (2003) subsequent study pointed to unbiased transmission as being more likely, with a brief stage of anti-conformism in earlier phases.

Recent research combining archaeological simulation and Bayesian inference (Crema *et al*. 2016) examined the Merzbach assemblage with more rigorous methods, whilst shedding light on some problems inherent in archaeological datasets, like the effects time-averaging that occurs when archaeological data are split into discrete chronological phases which often last for centuries, resulting in older traits being considered contemporary to earlier ones. The study also addressed the fact that most studies of this kind assumed that cultural systems under analysis had reached an equilibrium, expressed as the result of system parameters remaining stable for a long time (Crema *et al*. 2016: 1-2). The outputs of the simulation pointed at different modes of transmission occurring at different times, (Crema *et al*. 2016: 7) with anti-conformism as a copying strategy in the earlier phases, and conformism in the later phases.

These studies demonstrate what is one of the great advantages of using cultural transmission to identify the mechanisms driving cultural change. As argued by Eerkens & Lipo (2007: 240), it "provides an explicit quantitative framework for modeling the evolutionary process, allowing researchers to generate specific and testable hypotheses about material culture change".

2.2. Insights from ethnoarchaeology

Chaînes opératoires

A largely francophone tradition of archaeological technological analysis which focuses on reconstructing operational chains or chaînes opératoires of pottery assemblages (often interpreting such operational chains as representative of group or ethnic identity) has produced another strand of theory that is relevant when it comes to questions of technological adoption. Proponents of this school often reject evolutionary approaches on the grounds that they place too much emphasis on the mode of transmission rather than the content of what is transmitted (Gosselain, 2008: 151). Indeed, even from an evolutionary point of view, one important issue concerning cultural transmission is whether the mode of cultural transmission is content dependent. Researchers in this francophone tradition have used ethnoarchaeological, archaeological, and experimental methodology to produce important insights, which can be greatly beneficial to evolutionary research, especially for the purpose of reconstructing technological lineages.

From the perspective of niche construction, such studies could shed light on cultural niches built by generations of craftspeople, which affect the transmission of traits linked to one stage of the operational chain, for example shaping for pottery, to a preceding stage, such as clay preparation and tempering. This is especially relevant to wheel-thrown pottery, which is usually accompanied by fine clays and smaller inclusions.

Ethnoarchaeology

An assumption frequently encountered in the literature is that certain operations, like fashioning techniques, which depend on the practitioner developing specific motor habits- the set of repeated gestures that become internalised though daily practice and result in high fidelity copying-through a long apprenticeship- are less likely to change (Gosselain 2002: 79; Rice, 1987: 462) and remain relatively stable over time compared to other traits. Data from ethnographic studies suggest that techniques which leave visible traces on the vessels, such as decoration, firing, processing techniques that result in specific fabric colours and some preforming techniques are more easily transmitted between different producers than fashioning techniques (Gosselain, 2000: 191-2).

Gosselain argues this is due to the visibility of these techniques, which makes their resulting products easier to copy visually, without needing to invest time in apprenticeship with a tutor. In this case, feedback from consumers could have a direct impact on the visual appearance of pottery in this respect, as producers could opt to copy a popular decorative motif introduced by another potter to secure more clients. Furthermore, this aspect of visible techniques makes them easily transmissible through space, thus offering an insight on more “superficial, situational, and temporary facets of identity” (Gosselain, 2000: 191). However, it could also provide insights on the social context of the demand on pottery and on consumers (Gosselain, 2011: 222) who, in turn, become visible through their stylistic preferences. This type of horizontal transmission through peers is considered to result in more abrupt change in cultural traits, at least from the perspective of cultural evolution (Cavalli-Sforza and Feldman 1981; Reyes-García *et al*.2016).

Fashioning techniques are generally seen as “invisible”, at least in the eyes of contemporary consumers. This means that the probability that a consumer offers feedback on the fashioning process itself is more limited. It also means that any potters who wish to master a new fashioning technique must be taught by an expert who will teach them through often non-verbal ways. A student will have to imitate the hand positions of a tutor and will be corrected when their movements fail to produce the desired results. In consequence, the probability of transmission error, which could produce a higher degree of variation in the shape of the vessels, is lower.

While Gosselain's stance on the cultural evolutionary notion of cultural transmission is generally critical, his ethnographic research in Cameroon pointed at the significance of vertical transmission in the case of pottery techniques. The author traced the spatial distribution of displaced female potters – due to "ethnic endogamy" and marriage networks - and identified a clear correlation between the distribution of potters, ceramic styles and ethnic languages (Gosselain, 2002: 29). This, in turn, led to a conclusion that language equals ethnic association, and that mobility must have occurred within the confines of these ethno-linguistic barriers. The association with ceramic styles must be attributed to vertical transmission (Gosselain, 2002: 27), in other words, from parent-to-offspring. Offspring who migrated because of marrying into another community must have also taught their children. Gosselain supports his conclusions with quantitative data procured through interviews with 77 artisans in South Cameroon, and the vast majority reported having been taught pottery by someone in the nuclear and, less often, the extended family. Only 0.06% of the interviewees had been taught pottery by someone they were not in any way related to (Gosselain, 2002: 23). It should be noted that there are other documented cases in the literature, like the Luo in Kenya, where women who marry into a new community must be "relocalised" by their mother-in-law and learn how to make pottery in the "local way" (Herbich, 1987: 200-201) even if their mothers had taught them a different way. This is an interesting case of vertical transmission which then paves the way for conformist transmission to take place.

In a more recent paper, Gosselain provides another hypothesis on the tendency of potters to not change forming techniques other than maternal affiliation and motor habits: it could be interpreted as a sign of "deliberate conservatism" (Gosselain, 2011: 221), an attempt by the potters to affiliate themselves with the tutor who provided instruction, or with a group "identity". In this case, seen from an evolutionary perspective, this mechanism would be akin to conformist transmission.

However, this has interesting implications for cultural transmission research in communities with mobile potters. In a hypothetical study of a pottery assemblage partially produced by migrant potters settling in a new community but retaining maternal/paternal forming techniques, this conformist tendency would result in cultural frequencies that would indicate an anti-conformist bias. This process could then result in a material culture 'enclave' within the new community.

Archaeology

A study by Gomart and Burnez-Lanotte (2012) draws methodological inspiration from Roux's framework for identifying forming techniques (Roux 1994) as well as from ethnographic studies such as the ones mentioned above (linking specific motor habits to affiliated artisan groups), for LBK (Linearbandkeramik) and non-LBK pottery from Rosmeer, a site in Limburg, Belgium, and is presented here as a useful case study of research attempting to identify operational chains purely on the basis of archaeological observations.

The authors performed a two-fold technological analysis of the assemblage, grouping it into two distinct Chaînes opératoires, ROS1 and ROS2 that they considered to represent two different but coexisting groups of producers residing in the village (Gomart, 2012: 231). Initially, grouping was based on morphometric straits which indicated forming and finishing techniques, while the secondary analysis included traits relating to tempers, vessel shape, and decorative motifs. LBK vessels were made using the ROS1 Chaîne opératoire, while non-LBK pottery can be attested in ROS1 as well as ROS2. The main proxy used to discriminate between the two traditions is the way coils were formed and overlapped to form the vessel and, most significantly, the variation in coil heights which, in the case of ROS1 points to a greater number of artisans, while in ROS2 to a smaller number of producers (Gomart, 2012: 236).

Moreover, macroscopic and microscopic analysis of the temper used in both Chaînes opératoires identified three types of temper (powdery, hematite, and bone) and after a correlation analysis of fashioning techniques vs temper, the authors distinguished between sub-groups in each Chaîne opératoire: ROS1.1, with only powdery tempers, ROS1.2 with all the aforementioned tempers, ROS 2.1, again with only powdery tempers, and ROS2.2 with mostly bone inclusions and occasionally hematite and powdery ones (Gomart 2012: 237-8). ROS1.1 outnumbers ROS1.2, while ROS2.2 outnumbers ROS2.1.

In terms of fabric colour, most non-LBK vessels in both Chaînes opératoires are characterized by orange surfaces indicating an oxidising phase whilst being fired, but there is a small number of pots with slip applied post-firing (Gomart 2012: 241).

A consequent stylistic analysis of decorative traits then resulted in different subgroups for ROS1 and 2, since the discrimination based on temper was not correlated with different groups of stylistic traits. ROS1A, ROS1B, ROS2A, ROS2B. The authors claim that there was no significant discrimination in two clear groups corresponding to the two Chaînes opératoires in terms of style and thematic, but there was some correlation between decorative technique and Chaîne opératoire (Gomart 2012: 243).

The ROS2 tradition is, according to the authors, a discrete, smaller group of specialised producers making vessels of a specific function who had a different apprenticeship trajectory than the ROS1 potters (Gomart 2012: 247). The spatial distribution of ROS2 fragments suggests that production took part in a specific part of the settlement. Finally, Gomart believes there was some loose imitation of ROS2 morphostylistic traits by ROS1 potters which suggests that non-LBK pots had a specific "cultural role". However, no further attempt is made to elaborate what is meant by "cultural role".

Overall, Gomart's study of Rosmeer pottery has numerous strengths. First, analysis takes place on multiple levels, initially focusing on the distinction of chaînes opératoires based on forming and finishing technique and is then expanded to include criteria based on tempers and decorative traits, both in terms of style as well as technique. It serves to illustrate that researchers should be careful when grouping vessels into rigid producer groups, as discriminating based on different criteria could result in different groups. This was clearly illustrated by the fact that using tempers vs stylistic traits as discriminants resulted in different subcategories, ROS1.1 AND ROS1A. Nevertheless, combining proxies leads to the following deductions: first, that stylistic traits are more easily imitated between different groups of producers, a pattern that has also been identified by ethnographic studies, as discussed above and second, that temper preferences are also easily transmitted between groups, as bone inclusions, a typical characteristic of the ROS2 Chaîne opératoire, can also be attested in the ROS1 tradition.

One weakness to the authors' approach, however, is the assumption that the direction of transmission was from ROS2 to ROS1 on the basis that bone tempers are more abundant in the former, using examples from non-LBK pottery in other sites (Gomart 2012: 247). However, this argument is made relatively weaker in light of the fact that the study lacks an analysis of clay composition which could have ruled out imports as part of the assemblage, a fact that is, however, acknowledged in the study (Gomart 2012: 235). It should not be automatically deduced that the source of transmission was the group where the trait is more abundant (the simulations in this study suggest that it is possible for the origin community to have relatively low abundance of the trait). It is equally possible, especially in contemporaneous traditions like ROS1 and ROS2, that osseous tempers were simply more desirable by producers (or consumers) of one tradition.

Overall, the most important contribution of the study is the insights on craft specialisation: the authors conclude that ROS1 pottery, while seemingly produced by non-specialised potters who mostly produced LBK vessels for various domestic functions, also includes pottery that could be characterised as specialised on the basis of morphostylistic traits (Gomart 2012: 248), but possibly imitated elements of another, smaller, more specialised tradition of non-LBK pottery. ROS1, seen as a community of practice, is a good example of non-specialised artisans who adopt traits from a different, specialised community. This suggests that craft specialization exists a spectrum rather than a clear-cut distinction between specialized and non-specialised potters. It is also possible that the same potters participated in multiple operational chains.

Experimental approaches

A seminal study by Roux and Corbetta (1989) combined ethnoarchaeology and experimental approaches to assess the cost of learning how to throw pottery on a wheel as an investment in time, with subjects ranging from complete novices to potters of varying degrees of expertise. The results, which have been generally accepted by researchers, indicated that mastering wheel-throwing, the process of forming a vessel using only Rotative Kinetic Energy (RKE) from a wheel, is time consuming and takes 10-15 years of apprenticeship under a tutor, and that certain morphometric traits in pottery, such as height, are good proxies for inferring expertise in wheel-throwing.

Recent research has made use of this ethnoarchaeological/experimental to produce new insights on the nature and mode of "motor skill adaptation" (Gandon 2019). Gandon's experiments with expert potters in India were aimed at measuring the "cost" of adopting a new trait, in order to, eventually, identify the optimal conditions for the adoption and diffusion of new traits. Gandon's test subjects were expert potters belonging to different groups and specialised in different shapes and wheels (Gandon 2019: 37). They participated in three experiments, producing familiar and unfamiliar shapes with familiar and unfamiliar wheels.

The study produced important insights on the effect that pre-existing motor skills have when adopting new traits, as potters tend to use hand positions that they are already familiar with when approaching the production of a new shape. It provides empirical evidence on the notion of "embodied" learning that is often mentioned in the literature on pottery apprenticeship, as well as a robust framework for defining expertise in terms of motor skills: experts who produced more "diverse" pottery in terms of shape variation, possessed a larger "repertoire" of hand positions, which allowed them to adapt more quickly to new shapes (Gandon 2019, 237), and that potters who produce only small open shapes have hand position repertoires that are too limited to allow them to tackle making new shapes. In terms of what the implications for archaeology are, it follows that greater variation in vessel shape in an assemblage signifies greater artisan expertise and is also positively correlated with a higher probability to adopt new shapes. On the other hand, according to the author, domestic production is probably negatively correlated with a high probability to adopt new shapes because the cost of adopting is greater in such contexts, where potters consistently produce a small number of shapes.

Another important issue addressed by this study is whether different kinds of wheel, in this case the kick-wheel and the stick-wheel, would present a significant cost of adaptation for the potters. The results indicated there was no significant cost of adaptation (Gandon 2019: 237). This could provide an answer to the question of whether the evolution of the wheel, from the tournette to a high-speed wheel, had a significant effect on the diffusion of wheel-throwing as a technique.

In terms of modes of transmission, Gandon interprets the expert potter’s reproduction of new shapes as an individual-learning transmission, where the maker adapts their own personal repertoire of motor skills to produce the desired result without observing others, and claims that social learning would be necessary for less expert potters to produce the same result (Gandon 2019: 238) and that different transmission modes may present different costs of adapting new motor skills (Gandon 2019: 230).

2.3. Conclusion

If one's goal is to study the change of cultural patterns over time, there is much to be gained from adopting a Darwinian evolutionary approach, which treats culture as a replication system and provides a solid theoretical foundation based on empirical observation, mathematical modelling, and computer simulation work. While this strand of theory is not without gaps — notable ones being the difficulty of applying the models on time-averaged archaeological data and the lack of consensus on how traits are transmitted as part of a package — current research openly acknowledges their existence and is actively focused on addressing them (e.g., Miller-Atkins & Premo 2018, Perreault 2018, Premo 2014, Yeh et al. 2019).

Ethnoarchaeological and experimental research are promising avenues in terms of complementing the cultural evolutionary models, which tend to be more abstract in their focus, with additional information on the technical and social aspects that affect the process of cultural transmission. They provide empirical evidence for the prevalence of different transmission mechanisms and have helped formulate a set of expectations in terms of the signatures that certain types of social learning leave on material culture. Much of this work has focused on the social learning of pottery-making, which is also the most abundant archaeological material.

Both frameworks offer multiple avenues for addressing empirical questions about innovation diffusion. How one applies them depends on the type of explanation they are interested in - this is explored at length in the following chapter.

3. Simulation as a method for understanding cultural transmission

Introduction

As discussed in the previous chapter, there is a long-established theory of cultural transmission which draws elements from Darwinian evolution and has generated fruitful research with various archaeological applications. The theory is useful in terms of describing processes and mechanisms at the micro scale that give rise to the macro-scale patterns of cultural variation observed in the empirical record, and it can also generate predictions that we might expect to hold true in real-life processes of cultural transmission. However, while the processes the theory describes are dynamic, the empirical record is notoriously static, time-averaged (e.g., see Perreault 2018 on the implications of time-averaging in the observed rate of cultural evolution and Bevan and Crema 2021 on the effects of archaeological periodisation on inferences about the past), and is often a reflection of processes other than the ones we are interested in. For example, a survivorship bias arising from the differential preservation of materials might have nothing to do with processes of cultural evolution in the past, but it might lead one to falsely infer that at a specific point in time, trait diversity in pottery was higher than trait diversity in basketry because the former survives better than the latter. Finally, understanding any process involves having some knowledge of its causes. This means that we need a method to establish causal links between the theory and real-world data.

Simulation has emerged as one the most productive theory-informed methods of understanding cultural transmission and its application has become standard practice in evolutionary archaeology (Lake 2015: 27). This is because it is possible to use simulation to establish causality as well as observe the outcomes of the micro-process that we think happened in the past over time through iteration. This provides the desired link between the theory and the empirical record and allows us to compare the outputs of a simulation and real-world data to answer specific questions of the empirical record (e.g., "Did a certain type of cultural transmission bias impede the spread of the potter's wheel in the Bronze Age Aegean?").

Simulation can also be useful in terms of expanding the theory by generating new theoretical expectations (e.g., "what types of spatial patterns should we expect when cultural transmission takes place in insular environments?"). Lake (2015: 23) argues that while a distinction between empirical and theory-building models exists in the literature, it may not be particularly useful as in many cases, even if a simulation model is designed with the aim to address an empirical question, it can still be used to inform theory building. Therefore, one can use simulation to not only identify specific factors affecting particular episodes of cultural transmission, but also to gain a deeper understanding of the cultural transmission process itself.

The discussion in this chapter has an epistemological character in that it is concerned with the types of understanding that we can acquire by using simulation models (Section 3.1). It also presents general rules of thumb that can lead to the design of informative simulations (Section 3.2) paying particular attention to the notion of "level of abstraction", which characterizes our research questions, models, and the explanations that they inform (Section 3.3). Finally, I broadly describe the main elements of the modelling approach adopted for this study (Section 3.4). Details on how this general approach relates to the implementation of the model can be found in Chapter 6.

3.1. Using simulation to go from how-possibly to how-actually understanding

The problem of linking theory to incomplete and biased empirical data is shared by most historical sciences, including history, geology, biology, and, more importantly, archaeology and it affects the nature of our explanations and the types of understanding that they afford us. When we use models to explain the empirical record, our explanations typically fall under the umbrella of 'how-possibly' explanations, whose aim is to propose plausible causes, as opposed to 'how-actually' explanations, whose aim is to say exactly what happened that led to the observed phenomenon.

The concept of 'how-possibly' understanding was first introduced by Dray (1957) to describe the nature of certain types of explanations in the field of history. According to Dray (1968: 407) it is the incomplete nature of the empirical evidence which necessitates the use of how-possibly explanations, since "the account the historian gives at any one time is always built up piecemeal out of elements, for many of which he has independent evidence. What he claims to know at any stage of the constructive process-which is, of course, a never-ending one-can always raise the problem of how one event could have occurred in the light of some other occurrences for which there is also good evidence". This certainly applies to the study of dynamic processes of cultural change over time; in most cases we want to understand the causes of a specific episode of cultural change, but the process of cultural evolution is itself 'never-ending' and the evidence at hand is always fragmentary. To make matters even more frustrating, there is also the problem of equifinality which, according to Barrett (2019: 130) is "the possibility that in a dynamic, open system, the final observed state can be arrived at via different mechanisms and from different initial states". For the intents and purposes of computer simulation, this means that there can be multiple models (and parameter value combinations) that result in the same outcome.

In other words, most of the time we cannot say exactly what happened that caused the observed phenomenon, but we can suggest plausible scenarios that could have resulted in the picture that emerges from the empirical data. Brainard (2020:2) stresses the fact that explaining 'how-possibly' does not only involve demonstrating that something is possible, but that it requires showing the ways in which a certain outcome is possible.

Indeed, it is the range of potential explanations that can be produced by empirically informed simulation models that lends simulation its theory-building potential. Simulations can suggest multiple "ways" in which an outcome is possible, and this should not scare away those interested in how-actually explanations. I would argue that this enables us to pursue fruitful explanatory avenues even in the face of equifinality which, in turn, enables us to "learn" through "experimentation" and explore alternative scenarios that allow showing, according to Lake "what conditions must have been met for that phenomenon to have occurred if the model is correct" (Lake 2015: 24 citing Premo 2008: 50). Chapman and Wylie (2018: 108) also see merit to this kind of approach, seeing it as a way to "enrich" archaeological "inquiry" by suggesting multiple working hypotheses.

Another aspect of explaining how-possibly according to Dray, is that explanations that are "given in the face of a certain sort of puzzlement" (Dray 1957: 165-166), in other words, they are 'unexpected'. This is particularly important in the case of simulation models, where we start out with certain assumptions and a theory and seek to explain a phenomenon of interest without building it into the model itself (this will be discussed in greater depth in Section 3.2). If the simulation generates the phenomenon of interest, the modelling exercise is possibly successful in helping us identify some its main causes (the ones included in the formulation of the model) provided that there is no alternative mechanism that could have given rise to the exact same output. However, it can be argued that the opposite case, where a simulation model that is empirically well-supported in its assumptions fails to generate the phenomenon of interest can be even more informative in that it can suggest new avenues of inquiry for addressing the research question. This means that the how-possibly explanations arising from simulation models can provide powerful insights into the processes we are interested in, even if they are not an exact match to the empirical record.

Simulation is also useful because in the case of dynamic processes, we do not fully know the implications of our models until they are iterated over time and in a particular environment. Very often, simulation outputs suggest a counter-intuitive, 'unexpected' explanation for a process or event (e.g., Romanowska *et al*. 2022: 43), one that could not have been foreseen before running the simulation. According to Forrester (1971:111 and 139), this is because social systems exhibit feedback loops whose outcomes cannot be easily explained by human intuition without using computer simulation. Furthermore, seen in the light of complexity theory, social systems exhibit emergence, a property that characterizes systems that are made up of multiple separate parts whose interaction at a microscopic level produces global-scale outcomes that are irreducible at the level of the individual entity — this is often summarized with the phrase 'the whole is more than the sum of its parts' (e.g., by Ratter (2012: 96)). Simulation studies are often successful at capturing such emergent phenomena and at inferring the micro-level causal mechanisms behind them — a famous example is Schelling's (1969) 'segregation' model, which established that segregation at the level/scale of the urban neighbourhood can emerge out of individual level neighbour-preferences. Note how the segregation model does not rule out other mechanisms that could have produced neighourhood segregation. It simply offers a ‘how-possibly’ explanation that is formulated at the level of the individual.

Some researchers (e.g., Brandon 1990, Bokulich 2014) argue that starting from a ‘how-possibly’ explanation does not rule out eventually acquiring a ‘how-actually’ explanation. Brandon (1990) was concerned with the validity of how-possibly explanations in biology, especially explanations regarding the evolution of biological traits. According to Brandon, a good example of a how-possibly explanation is Darwin’s explanation for the evolution of the eye — Darwin posited that the eye arose out of “an optic nerve merely coated with pigment” (Brandon 1990: 177) followed by successive modifications. In Brandon’s (1990: 179) view, how-actually explanations are “potential explanations, none of whose explanatory premises contradict or conflict with "known facts"” and are “based on generalizations or laws we have good reason to believe are true, but whose initial conditions are speculative”. In Darwin’s example, he based his explanation on a well-supported generalization (that traits can evolve in successive steps via a process of descent with modification) but had limited information on the initial conditions that set in motion the process of the evolution of a specific trait, in this case, the eye.

This means that we cannot arbitrarily accept any explanation as a valid how-possibly explanation for a phenomenon — it has to be consistent with a theory and our current knowledge. Additionally, since such explanations “can be rigorously formulated so that they do have testable consequences” (Brandon 1990:180), it is possible to evaluate them using empirical data. Ultimately, for Brandon (1990: 184) “a how-possibly explanation is one where one or more of the explanatory conditions are speculatively postulated. But if we gather more and more evidence for the postulated conditions, we can move the how-possibly explanation along the continuum until finally we count it as a how-actually explanation”.

This nicely captures what makes the insights gained from archaeological simulation valuable; in most cases, when one is studying cultural change using archaeological data, one may know a lot about one or more aspects of the process and very little about others - Brandon's "postulated conditions" (1990: 184) — and also be in a position to make reasonable 'law-like' assumptions on certain fundamental aspects of human behaviour (for example, humans reproduce sexually, they do not live under the sea, they tend to copy their parents, etc.) or use an established theory (in our case, cultural evolution). In such cases, it is possible to build a simulation using this information to discover alternative scenarios that could have given rise to the phenomenon of interest, compare the simulation to empirical data, and select the most plausible 'how-possibly' explanation which is most consistent with one's domain knowledge and the data. This does not stop someone coming in with new data, perhaps regarding some aspect of the model that was previously completely unknown and turn the previously 'how-possibly' explanation into a 'how-actually' explanation.

Considering how novel techniques in data collection and analysis such as radiocarbon dating have traditionally led to the re-evaluation of archaeological models (for an overview see Manning 2014: 130) it is very likely that if an archaeological simulation produces a how-possibly explanation with testable implications, future researchers will eventually come along to evaluate it with new data, turning it into a 'how-actually' explanation or discarding it from the pool of reasonable potential explanations (this kind of hypothesis reduction approach is discussed at length by Kandler and Powell 2018). This, in turn, means that the epistemic value of archaeological simulations depends to some extent on their results being reproducible and their conceptualization, assumptions, and implementation strategy being openly available to other researchers.

3.2. Optimising the explanatory potential of simulations of cultural transmission processes

Regardless of whether we want to achieve how-actually or how-possibly understanding, it is important to design a simulation model according to certain principles in order to maximise its explanatory potential. Archaeological simulations, like most archaeological methods, borrow many of their methodological and technical aspects from other disciplines such as complexity science (Brughmans et al. 2019), statistical physics (e.g., the model by Knappett, Evans, & Rivers 2008) and evolutionary biology (e.g., Mesoudi & O' Brien 2008). This means that design principles typically emerge on an ad-hoc basis, according to the specific needs of individual researchers. However, there is a literature of methodological and review papers (Lake 2015, Barcelo & Castillo 2016) that address the issue from a more general perspective. Here I focus on the suggestions made by Lake (2015), who suggests rules of thumb that can inform the design of simulations with good explanatory potential. Lake has an extended discussion on each of those main points and how they arise from current discourse in epistemology, philosophy of science, and the archaeological simulation literature. Here I will briefly present his suggestions and discuss their implications from the perspective of how-possibly and how-actually explanations.

Lake's suggestions can be broadly summarised in three points:

1. A focus on experimentation instead of mere "emulation" (Lake 2015: 23).
2. Ensuring that the simulation generates the phenomenon of interest (it arises in simulation outputs as a result of formal "micro specifications") rather than pre-programming it into the simulation (Lake 2015: 24).

3. Simple models and ‘realistic’ models offer different advantages — parsimony (Lake 2015: 26) and generality (Lake 2015: 27) in the case of the former, “good predictions” in the case of the latter (Lake 2015: 27 citing Crutchfield 2008: 274) — which should be gauged before deciding on an appropriate level of abstraction.

The first point is an argument in favour of a simulation design that allows discovering multiple how-possibly explanations for a phenomenon and against a design whose aim is to simply reproduce the phenomenon. Lake agrees with Premo (2005, 2007, 2008) in that mere emulation does not equate to explanation, mostly because it does not deal with the problem of equifinality. I would add that even if the problem equifinality was somehow eliminated, an experimental design can still be more informative than one which is only geared towards the emulation of a phenomenon, and I will use an example to demonstrate it.

Say that we were interested in learning what caused the collapse of the Mycenean palatial system, and we speculated that it was due to some degree of internal infighting, economic instability, and external hostilities (which we would include as parameters in the model). If we had enough information and good empirical data, we could in theory identify the parameter value combinations that result in an outcome that fits the observed phenomenon exactly and say, for example, that the collapse of the Mycenaean palatial system was due to high economic instability, high internal hostilities, and to a lesser degree, external hostilities. In an ideal scenario, where there is no other combination of parameter values that produces the empirical outcome (no equifinality), we could come to the conclusion that this is indeed what caused the fall of that particular palatial system in this particular scenario. However, apart from knowing that this was the exact combination of factors that led to the observed outcome, this exercise does not help us understand the process much more beyond that. For example, we cannot answer what would have happened if there were no external hostilities. Would the Mycenaean palatial system have been resilient in the face of internal instabilities had it not been for the factor of external hostilities? We cannot know that, because we have not explored the model in conditions beyond those that reproduce the observed outcome exactly. In other words, we know what the process was, but there are more interesting insights that could be gained from exploring the implications of alternative scenarios.

The second point is that simulations produce informative explanations if the outcome that we want to explain is ‘generated’ by running the simulation and is not pre-programmed into the simulation beforehand. This is, again, a distinction that is easier to grasp through an example. Say that we were interested in learning what gave rise to the emergence of the palatial system itself. If we build a simulation that includes a parameter that represents a global probability of a palatial structure (with all associated functions of the palace, such as complex administration) randomly emerging, it is inevitable that given enough time (running the simulation for a sufficient number of timesteps) a palatial structure will emerge in the simulation. This explains nothing about how palatial structures emerged from entities in the micro-level (e.g., villages), and it is akin to circular reasoning, where we are trying to explain something starting with itself as an initial premise/assumption. In this example, it is the equivalent of saying that if there is a probability for a palace to emerge, it will emerge, and this explains the emergence of palaces.

Finally, the third point is that there are different benefits associated with models of different levels of complexity/abstraction, but Lake leans toward the view that simpler models are more likely to fit the criteria discussed above (potential for experimentation and “being generative with respect to the problem at hand”) (Lake 2015: 27). Simple models are more likely to describe a wider range of possible scenarios, and therefore allow us to experiment with multiple potential explanations. Simple models are also less likely to contain pre-programmed inputs that are the same as the intended outputs, which would result in the model not being generative with respect to the process of interest. Avoiding circularity is also relevant here; the mere fact that the number of inputs for more complex models is larger than the number of inputs for simpler models makes it more likely that a modeler will unintentionally include information that results in a simulation with outputs that lead to undesirable circular explanations.

Finally, the principle of parsimony (also referred to as 'Occam's razor') suggests that if we are interested in getting a specific explanation for a particular phenomenon, we should begin with the simplest explanation and add complexity only when necessary. Whether a researcher chooses to employ Occam's razor is a matter of preference, but scientists have been reducing problems down to simple models successfully for a very long time (Sober 2015: 2) and it makes sense from a practical standpoint - the more complex the model the more data we need to validate it and the more assumptions we have to make. As argued by Scerri et al. (2018:588), "more complexity means more parameters and more ways for a model to differ from reality. This means that unless informed a priori by secure information, or fitted to substantial quantities of conditioning data, more complex models can be more wrong, not more realistic".

3.3. The implications of selecting the appropriate level of abstraction for simulations on our understanding of cultural transmission processes

The distinction between simplicity and complexity frequently comes up as a quality that characterizes our models: models are described as either simple or complex, but research in epistemology suggests that simplicity and complexity also characterize our explanations and research questions, and that this has implications for the types of understanding that we acquire by modelling a phenomenon. In that literature, simple models and explanations are the ones that represent higher levels of abstraction compared to more complex ones.

Bokulich (2014) frames the problem in the context of 'how actually' and 'how possibly' explanations, as discussed above, whilst partially adopting Brandon's (1990) stance that 'how possibly' explanations differ from 'how actually' explanations in their degree of empirical support and that it is possible to turn how possibly explanations into how actually ones by acquiring more empirical evidence. However, she also believes that Brandon's view lacks an important element that determines the potential of any model to produce a 'how actually' over a 'how possibly' explanation: levels of abstraction.

More specifically, Bokulich suggests that "there are different levels of abstraction at which the explanandum phenomenon can be framed and which correspond to different explanatory contexts" (Bokulich 2014:333). In other words, there are different levels of abstraction at which we define our research question and that we require of our explanations and, more crucially, they play a major role in determining where our explanations fall in the how-possibly/how-actually spectrum. Bokulich explains this distinction with reference to the "tiger-bush", a long-studied phenomenon of vegetation patterning in semi-arid conditions (Bokulich 2014: 325-326) but I will attempt to do the same with an archaeological example.

In our Mycenean palace example, if we ask:" Was the collapse of the Mycenaean palatial system due to internal economic instability or due to generalized conflict in the area?" we are effectively framing the question in a way that asks us to evaluate two alternative potential explanations defined at a high level of abstraction. It should be straightforward to build a simple a highly abstract model that tests for both scenarios and selects the one that is a better fit to the data, ending up with an explanation that satisfies our research question enough to count it as a how-actually explanation and say, for example, that the collapse of the Mycenaean palatial system was due to internal economic instability. However, if our research question is defined at a finer level of granularity, we may at best get a how possibly explanation. For example, if we instead ask" which were the factors driving the internal economic instability that led to the collapse of the Mycenaean palatial system", we are framing the question in a way that requires evaluating multiple alternative factors and modelling their interactions using a more complex fine-grained model. Such a model could suggest a mechanism driving economic instability such as an increasing emphasis on the production of sheep for wool leading to a collapse in the available pasture. Given that we have the same empirical data to evaluate both simulations, e.g., the geographic coordinates of the main palatial centres and their dating, it is more likely that the simple abstract model will provide what appears to be a definitive 'how-actually' explanation compared to the more complex finer grained model, which will yield a how-possibly explanation (but note that following Brandon and Bokulich, it would still be possible to turn it into a how-actually explanation if we gathered more data).

3.4. A modelling approach to study innovation adoption in the Bronze Age Aegean

The insights from the previous sections helped establish rules of thumb that informed the modelling approach I have adopted for this study. In my opinion, the distinction between how-actually and how-possibly explanations is a useful one in terms of establishing expectations on what the answers to our research questions will be. Section 3.3, in particular, highlighted the effect that the framing of the research question has on the nature of our explanations. The archaeological simulation literature suggests that simple abstract models that avoid pre-programming outputs and adopt an experimental design are useful for addressing research questions in a field with biased and incomplete data.

The research questions for the present study were described at length in Chapter 1, but it is useful to re-iterate them here to show how they relate to the level of abstraction adopted for the simulation model and address the types of understanding that can be gained from running it. The first research question was: “What was the role of geography, inter-community distance, and population in the adoption of ceramic innovations in the Bronze Age Aegean?”. The framing of this question is open-ended, meaning that it allows room for exploring multiple facets of geography and population and their effect on novel trait adoption. This makes it compatible with an exploratory research design as suggested by Lake. The benefit of such an approach is that it can offer interesting scenarios for Aegean archaeologists, who have long speculated on the role of geography in innovation adoption in the context of communication networks between potters (e.g., Choleva 2020) and produced extensive verbal models describing the movement of pottery along existing long-distance networks (e.g., Sherratt 1999). What these studies lacked was a tool allowing them to iterate their short-term models in time and experiment with aspects of geography itself. Hopefully, the insights gained from such an experimental simulation can inform future data-collection efforts and perhaps even highlight new avenues for empirical research.

The second research question was “What was the relative contribution of geography, population size, and local response to innovation that caused the adoption patterns for ceramic innovations in the BA Aegean?”. The framing of this question suggests a need for a how-actually explanation, or a well-supported how-possibly explanation that establishes what caused a specific episode of cultural transmission. This means that our explanation should not contradict known facts, and this is where the empirical archaeological data comes in. Some of these known facts will be used as assumptions informing model design (they are discussed in detail in the Chapter 6) and since my simulation models is designed to have testable outcomes, I will also use some of this data to evaluate how the model fits the real world. To avoid circularity, I will use different lines of evidence to support the assumptions of the model and to compare the model’s output to the empirical record — as discussed above, the goal is to ‘generate’ the phenomenon without preprogramming its outputs into the simulation.

The success of this exercise depends on the nature and quality of the empirical data available which — given that they are published by other researchers with their own individual research goals and theoretical approaches — should be expected to demonstrate considerable variability. Chapters 4 and 5 deal with the empirical data in greater depth. Chapter 4 addresses the issue qualitatively, looking at the discourse in Aegean archaeology and reviewing the most relevant studies, and Chapter 5 is a quantitative analysis of the dataset that was produced by combining the datasets from the studies discussed in Chapter 4. Chapter 5 will also evaluate the compatibility of the dataset associated with each innovation (wheel-made pottery, kiln, iron-oxide based decoration) with simulation modelling.

4. Pottery traditions and technological change in the Bronze Age Aegean

4.1. Introduction

In previous chapters, it was shown that pottery traditions evolve over time by retaining some aspects inherited from previous generations, losing some aspects due to chance, and incorporating new elements that are introduced either by chance or by intentional efforts to solve a specific problem or to address a specific need. The pottery traditions that emerged in the Aegean during the Bronze Age evolved in a similar manner; there are observed instances of technical conservatism - cooking pots frequently exhibited little change over time (see examples by Lis (2010: 235) and Menelaou & Kouka (2021: 9)) — as well as other instances where new elements were introduced and retained for shorter or longer periods of time in local traditions.

The potter's wheel (an example of wheel-made pottery is shown in Figure 2), kilns (an example shown in Figure 3), and iron-oxide based lustrous decoration (an example shown in Figure 4) were new elements that were introduced in Aegean pottery traditions at some point during the Bronze Age and were more widely adopted by the end of the period. The rate at which they were adopted by local communities over time varies. The potter's wheel was first introduced in the region during the EBA II period, and its use increased very slowly (and in some cases even dropped) over time until the MBA, at which point it started increasing more quickly due to its use for Minoan and Minoanising pottery (see discussion below). After the first introduction of kilns (possibly in the EBA (Jones 2021), kiln numbers increased during the MBA and LBA. Lustrous decoration appeared early in the MBA Kamares ware in Crete and was widely adopted in the LBA through its incorporation in a Mycenaean set of practices for producing pottery - or, using terminology from cultural evolution, a trait ‘package’.

This chapter will attempt to trace the adoption trajectory of each innovation through the lens of Aegean archaeology and the narratives woven by researchers working in the field. This largely qualitative discussion is meant to be complementary to the quantitative analysis of the empirical data emerging from some of the studies mentioned here (among others) that is included in the following chapter. Section 1 examines the potential mechanisms that enabled the transmission of new technical knowledge between communities separated by smaller or greater distances in space. I put emphasis on the concept of ‘itinerant production’ and on whether there is evidence that itinerant potters enabled the transmission of the novel traits under discussion. Sections 2, 3, and 4 are brief reviews of the literature on the adoption of wheel-made pottery, kilns, and lustrous decoration in the Aegean. Finally, section 5 gathers the different lines of evidence from sections 2-4 to identify similarities and differences in the transmission trajectories of the three innovations.

Figure 2: A wheel-made conical cup. Source: The Metropolitan Museum of Art, New York. Bequest of Richard B. Seager, 1926. URL: www.metmuseum.org.

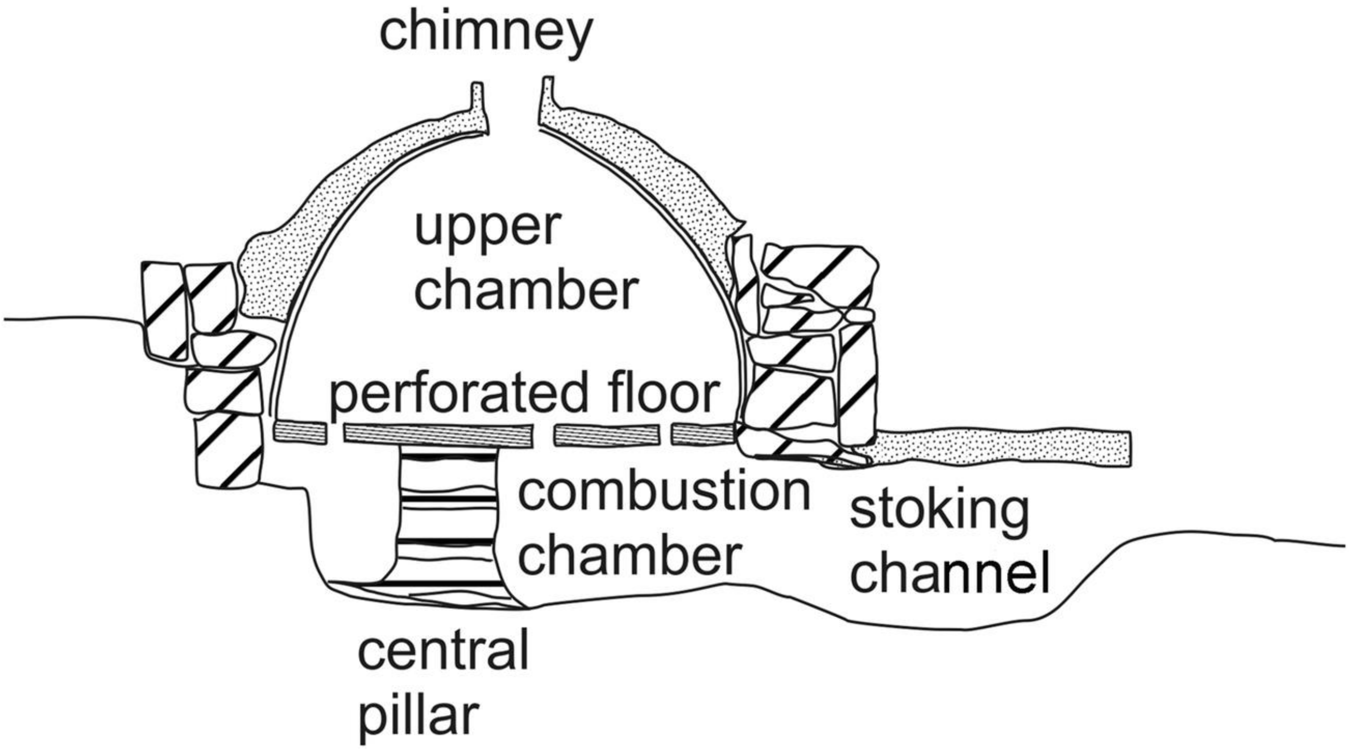

Figure 3: A sketch of a kiln by Karkanas et al. (2019: 5764).

Figure 4: An example of iron-reduction used to produce lustrous decoration on a Mycenaean stirrup jar. Source: The Metropolitan Museum of Art, New York. Fletcher Fund, 1925. URL: www.metmuseum.org.

4.1.1. Mechanisms enabling the transmission of novel traits in space

There are two main points of interest regarding the mechanism enabling the transmission of new traits by communities: who transmitted the information and who received it, and how they interacted in space. The first point of interest has been investigated through a chaîne opératoire framework, which has been, in some cases, used as a proxy for identifying specific groups of potters. Special attention has been given to identifying groups of specialised producers. This discourse is generally placed in the context of 'emerging social complexity', and it attempts trace the socioeconomic factors that resulted in the more 'centralised' organisation of the Cretan and mainland Greece palatial systems of the Middle Bronze Age.

Even though the notion of complexity is used frequently in this literature, definitions vary, with scholars using it interchangeably with notions of wealth inequality, architectural characteristics like town planning, craft specialization (Cultraro 2007: 64) and "sophisticated" cultural "phenomena" such as weighting systems and administrative practices (Rahmstorf 2016: 226). The adoption of new traits in pottery is sometimes linked to specialised production potentially operating under palatial control or satisfying demand from palatial centres. For example, Betancourt et al (2021: 114) see the cooperation of workshops with highly specialised potters as a necessity for the production and wide dispersal of wheel-made lustrous decorated pottery in Crete.

The second point of interest regarding the mechanism driving the transmission of novel traits between communities located in different locations is sometimes referred to as the "technological mobility" patterns (Knappett & Kiriatzi 2016: 15) of Bronze Age craftspeople. If traits were taught from one potter to the other, we can assume one of three alternative mechanisms that enabled it; either there existed groups of itinerant potters who were not associated with a specific location but who taught the new techniques to local potters, or information was transmitted via word of mouth between potters, or potters associated with a specific location (empirically, a ‘site’) would relocate for extended periods of time to other communities. As discussed in Chapter 2, we typically expect techniques that require long apprenticeships (like forming techniques) to require extended stays at a community different from the potter’s original community, but other traits such as decorative motifs could be acquired by simply looking at an imported vessel or through word-of-mouth.

Long apprenticeships in other communities could have taken place in the context of marriage networks, where female potters trained in the tradition of their community of origin moved to the community of their male spouse — ‘patrilocality’, an example observed ethnographically among the Luo in Kenya (Herbich: 1987) — or male potters trained in their original communities moved to their female spouse’s community (‘matrilocality’). The question is, is it possible to assess whether craftspeople mobility in the Bronze Age was linked to itinerant production or potter re-settlement using empirical data?

A volume edited by Kiriatzi and Knappett (2016) includes multiple studies describing potential scenarios that could result in potter ‘mobility’ from the perspective of Aegean archaeology and using empirical data. According to Knappett & Kiriatzi (2016: 15), elite involvement is consistently identified as a possible driver enabling the movement of potters around the Aegean, since they probably provided both the “means” (Knappett & Kiriatzi 2016: 15) and “ends” (elite-driven demand for specific types of pottery) for travel. However, it is also noted that non-elite potters travelled as well (Knappett & Kiriatzi 2016: 16). For instance, Kiriatzi and Andreou argue that expert potters using the wheel probably moved between the Mycenaean core in southern Greece to central Macedonia at some point during the LBA, since some of the pottery made in that area was made using local materials but with a "foreign technique" (using the wheel, relatively fine calcareous clays, and iron-oxide based decoration) (Kiriatzi & Andreou 2016: 141-142). According to Kiriatzi and Andreou (2016: 144), this suggests that the pottery was not just an imitation of Mycenaean pottery, and that the potters producing such pottery operated locally but had been trained in the Mycenaean tradition, implying that either potters from Macedonia moved to the south to complete lengthy apprenticeships, or that potters from the south moved to central Macedonia.

In the context of the adoption of Minoan ways of making pottery (e.g., wheel-fashioning) in areas outside of Crete, Nikolakopoulou and Knappett (2016: 103) argue that in the literature there has been a sharp distinction between a colonisation model, where Minoan elites with their attached potters settled in colonised areas (and produce Minoan-looking pottery), and an acculturation model, where communities outside of Crete imitated Minoan material culture to gain prestige. They see both models as extremes with potentially useful elements — mobility in the case of colonisation models and a focus on adoption for acculturation models — (Nikolakopoulou and Knappett 2016: 104) and advocate for an approach that considers both the movement of potters in space and the mechanics of adopting a new technique from the more micro-level perspective of individual potters and their local communities of practice.

Overall, these studies identify specific instances where potters must have moved from one community to the other, and they convincingly use empirical data to support their claims. Between two parties (as in Kiriatzi and Andreou's (2016) example) it is not often clear who moved to whose community, and what the social context behind the interaction was. Nevertheless, evidence for potter relocation is valuable in terms of developing preliminary models of potter interaction in space.

Regarding itinerant production, it is difficult to assess whether itinerant potters played any part in the transmission of the new traits, because itinerant production itself is difficult to identify using empirical data: what types of evidence would we expect to see in the pottery that would suggest that it was made by travelling craftspeople rather than potters associated with a specific geographic location?

From the perspective of ethnoarchaeology, there are cross-cultural "regularities" (Roux 2019: 260) that characterise itinerant production, and which we can use to define diagnostic criteria to identify pottery made in the context of itinerant production. Van der Leeuw (1977) produced a typology of different modes of pottery production focusing on economic organisation using information arising from his ethnographic research among modern potters in Tibet and Turkey (van der Leeuw 2020: 246). Itinerant potters fall under van der Leeuw's "individual industry" category, which consists of itinerant craftspeople (individual potters) who work with local materials in a wider region, sell/exchange their pottery in the context of a regional market, and use temporary firing structures (van der Leeuw 2020: 245-246). Following this framework, itinerant potters differ from stationary potters in the larger spatial extent in which we might identify their pottery (compared to pottery made by potters with a permanent base), the fabrics that they use — a wider range of fabrics made from local materials for itinerant potters as opposed to a single fabric — and the firing structures that they use (temporary contraptions rather than kilns).

In the Aegean, there are few cases where researchers have securely identified itinerant potters using pottery-based proxies. When researchers do attempt to identify itinerant production using proxies on the pottery, they do not make reference to a specific theoretical framework that would help establish consistent criteria for the identification of itinerant producers (such as van der Leeuw's typology). This means that the term 'itinerant potter' is used on an ad-hoc basis and that, even if there were greater numbers of studies using the term, it would be hard to say whether they meant the same thing.

For example, Christakis argues that two pithoi from Ayia Triada, Crete, were made using local manufacturing techniques and fabrics, but the shape of the pithoi and certain decorative aspects were more reminiscent of pithoi from Knossos, suggesting the presence of itinerant potters (Christakis, 2004: 302). However, it is also acknowledged that it is also likely that these were local "imitations" (Christakis, 2004: 302).

Recently, researchers have used fairly sophisticated methods to throw light on possible instances of itinerant production in mainland Greece during the EH period. Katsonopoulou *et al*. (2016) focused on large storage jars (pithoi) found in Early Helladic II-III site Helike, in order to assess whether they were made by the same group of potters who produced the rest of the pottery, mostly tableware, found on site. In this multimethod study on n=63 pithoi samples, macroscopic analysis was complemented by petrographic analysis, SEM (Scanning Electron Microscopy) and XRD (X-Ray Diffraction) analysis to identify standardisation patterns in the forming technique, clay procurement and processing strategies, and firing conditions (Katsonopoulou *et al.* 2016: 13-14).

Macroscopic analysis indicated that the pithoi in Helike were built using a rough clay base and overlapping coiled walls and fired in oxidising conditions and the low variation present in morphometric traits visible in wall thickness and rim diameter suggests their production was standardised, with some room for individual preference in decoration (Katsonopoulou *et al* 2016: 13-15). Additionally, petrographic analysis pointed to the use of locally procured clay and temper and Scanning Electron Microscopy carried out on a smaller number of samples (n=14) indicated low firing temperatures, which the authors attribute to using open fires instead of a kiln. Finally, two fabric groups, A and B, were identified as part of a larger study including tableware, both featuring local clay. Pithoi were classified in fabric group B (Katsonopoulou *et al* 2016: 14-15).

The authors argue that the potters who produced the Helike storage pithoi were probably specialised craftspeople, who produced only large storage vessels with small variations in shape made from local clays and fired them in situ, which the authors see as evidence for their being produced by itinerant artisans rather than local potters (Katsonopoulou *et al* 2016: 15-16). The tableware found next to the pithoi suggests that the storage jars were located inside areas intended for food consumption. The spatial distribution of the pithoi is interesting: they were found inside buildings intended for habitation, aligned with the inner walls, rather than in designated storage rooms. This led the authors to suggest that pithoi were placed in areas designated from the planning stage of the buildings, and that their manufacturers were involved in this planning process to some degree (Katsonopoulou *et al* 2016: 15-16). However, the authors do not provide architectural evidence which would render their argument more robust.

Concerning the “itinerant artisan” narrative, the authors acknowledge that this hypothesis has been based heavily on ethnography, in particular on “pitharades”, mobile potters making jars in multiple sites, who however, fired their products on central kiln sites (Katsonopoulou *et al* 2016: 15) and not on open fires like the Helike artisans, achieving high temperatures. The authors also note that ethnographic parallels have also been used by Christakis (1996) to support the idea of itinerant specialists in Minoan Crete. This implies that the authors see using ethnography to inform such archaeological interpretations as a weakness. In my opinion, if there is a slight weakness in Katsonopoulou et al’s (2016) and Christakis’ (1996, 2004) approach, it is not in the use of ethnographic data per se, but in the use of isolated examples rather than systematic frameworks based on cross-cultural regularities (like the one by van der Leeuw (1977)). Overall, the Helike study offers interesting insights on the nature of craft specialisation in Early Bronze Age Peloponnese, especially on the fact that the focus of said specialisation was to create products whose function was linked to the built environment. As aforementioned, the connection of these specialised products with food storage and “innovative building techniques” is seen as indicative of its connection with increasing surplus accumulation by elites and, therefore, wealth inequality and social complexity. (Katsonopoulou *et al*. 2016: 13). It also makes a well-supported case for the existence of itinerant potters in the EBA.

Overall, archaeological research confirms that there were many ways in which potters could communicate at a distance. Travelling was a possibility, and a well-supported one on the basis of empirical evidence. For instance, an analysis of strontium isotope ratios of eleven skeletons in Grave Circle A at Mycenae confirmed that only two skeletons that were buried in the grave were locals, and all others had come from different regions (Nafplioti 2009: 279). Whether travelling was in the context of marriage, relocation to an affiliated settlement (e.g., in Mycenaean networks of interaction), or itinerant production is less clear, and it is unlikely that empirical evidence from pottery alone will lead us very far in distinguishing between such alternative scenarios. Combining pottery data with different lines of evidence such as ancient DNA and strontium isotope analysis of biomolecular data will probably be a more fruitful avenue for future research in this area, and recent studies have shown the potential of such methods for throwing light on the movement and marital patterns of individuals in the Aegean. A good example is a recent study by Skourtanioti et al. (2023) using ancient DNA methods on a sample from 102 individuals from mainland Greece, Crete and the Aegean islands has produced important insights regarding the population movements and reproductive patterns of people in the Bronze Age. The authors confirmed that the first farmers in Crete were of Anatolian origin (Skourtanioti et al. 2023: 297) and showed that there was an influx of genes from the east in other Aegean islands. They also identified cases where the evidence suggests that communities "regularly practiced first-cousin intermarriages" (Skourtanioti et al. 2023:296) and found that endogamy occurred from the Neolithic until the end of the Bronze Age throughout the Aegean, which the authors note is inconsistent with prehistoric marital patterns in other parts of the world. Additionally, they found that during the 'Myceneanisation' period in Crete, Cretan individuals living in ports reproduced with individuals from outside the island (Skourtanioti et al. 2023:297).

4.2. Wheel-made pottery

As discussed in Chapter 2, the introduction of the potter's wheel has been identified as an interesting case of cultural transmission in many different contexts due to the fact that its mode of transmission (vertical transmission) is often seen as a mechanism that ensures the survival of the trait over long periods of time and possibly impedes the introduction of novel traits in culture. If parents teach their children how to use a specific forming technique, and those children then teach their own children the same technique, there is little room for variation to be introduced. If, however, potters travelled between communities, it is possible that the new forming technique would be introduced in the new community by a migrant potter and that the technique would still be used in the community after their death by their offspring.

This is what the literature suggests happened during the first introduction of the potter's wheel in the Aegean. According to Choleva (2020), wheel-made pottery was introduced by migrants from Western Anatolia, who used the new technique to produce a specific set of eating and drinking vessels of Anatolian origin that comprise the well-known "Lefkandi-Kastri" assemblage. In most sites, the technique enjoyed little popularity compared to hand-making techniques, suggesting a conservative tendency to copy one's parental generation and not teach the technique beyond the immediate group. This changed after the first palace period, when wheel-made pottery became widely used in the context of making Minoan and Minoanising pottery — and later, Mycenaean and Mycenaeanising pottery. This suggests that the technique was then associated with a different mode of production and a different transmission mode (horizontal and oblique transmission, where the technique is taught more widely and beyond the family group).

4.2.1. Identification of wheel-made pottery

There is a growing number of studies by Aegean pottery specialists that use Roux and Courty's (1998) observations to trace distinct ways of combining coil-fashioning and the (increasing) use of Rotative Kinetic Energy to produce wheel-fashioned pottery, which is seen as the 'precursor' of wheel-made pottery in most sites.

Spencer's PhD thesis (2007) examined pottery from the Early to Middle Helladic phases of two sites in mainland Greece: Lefkandi and Asine in terms of technical aspects which point to specialised production in the case of Lefkandi, and household production in Asine. Spencer used macroscopic and petrographic analysis to ensure that the sample consisted of locally produced ceramics. Among the data published are relative percentages of wheel made (using Roux's method) and handmade pottery for consecutive stages lasting 100-200 years (Spencer 2007: 284-287, 303) as well as diachronic observations on other proxies of pottery standardisation, like indications of firing control, levigated fabrics and polished surfaces (Spencer, 2007: 303).

More recently, Choleva published multiple macroscopic analyses of pottery from large sites like Lerna (Choleva 2012), Palamari (Choleva 2015), and Lefkandi (Choleva 2018: 58) tracing each chaîne opératoire and reporting the relative percentages of each one of Roux's methods within each assemblage, whilst also reporting the relative percentages of handmade pottery vs pottery generally considered to be wheel-made (including all of Roux's methods as a single category), as in Pefkakia and (Choleva 2018: 62). The analyses from Palamari and Lerna are reviewed below as case studies, illustrating Choleva's methodology and attempt at inferring the social processes that resulted in the identified operational chains. One of Choleva's more recent publications (Choleva

2020:71) reports percentages of wheel-made pottery for multiple sites around the Aegean (Troy, Emporio, Heraion, Palamari, Lefkandi, Aya Irini, Pefkakia, Mitrou, Thorikos) and has been one of the most valuable sources of information regarding the spatial transmission of wheel-made pottery in the Aegean during the EBA. In the following sub-sections, I will also review case studies by Berg (2013) on pottery from Phylakopi.

4.2.1.1. Palamari II-III

Macroscopic analysis was carried out by Choleva (2015) on a small sample of handmade and wheelmade (wheel-coiled) pottery consisting of shallow bowls (distinguishing between "Type 1" and "Type 2" shallow bowls) and plates that belong to the stratigraphic horizon Palamari II and III, with the intention of tracing chaînes opératoires linked to three stages of production: clay preparation, vessel formation and firing (Choleva 2015: 146). These shapes are the most commonly occurring wheel-made shapes in Palamari (Choleva 2015: 164), where the first evidence for the use of the potter's is dated to the Early Bronze Age II (Palamari II-III).

Analysis of wheel-coiled pottery (n=65) indicated the presence of two fabric groups and three forming groups. In terms of fabric and tempers, two "clays" were identified:

Clay A. Fine with small/medium inclusions, brown to red/brown, mostly used for plates and shallow bowls. This is the most commonly occurring clay (Choleva 2015: 165).

Clay B. fine with fewer inclusions, red, sporadically used for plates and shallow bowls. As acknowledged by the author, both fabric groups were identified solely through macroscopic inspection and the results should not be considered conclusive, as petrographic analysis is needed to confirm them. (Choleva 2015, 155). The groups were:

Group 1 (n=48): the vessel was mostly formed through coiling, and a rotational movement was used at the final stages of forming the rim, smoothing the interior and top part of the pot, followed by surface treatment to make forming operations less apparent.

Group 2 (n=14): “stretched surfaces” and “undulations” caused by applying “pressure to the coil joints” are diagnostic features of RKE, which was used to thin out the walls, similar to Roux’s method 2. Subsequent surface treatment probably erased wheel marks on the external surface.

Group 3 (n=3): RKE was used more than in Groups 1 and 2 and employed from an earlier stage of joining the coils, followed by thinning the walls using RKE. These operations resulted in diagnostic features like “undulations” and striations” visible on the surface. As in Group 2, surface treatment possibly erased wheel-marks. (Choleva 2015, 165).

Based on these observations and the correlation between forming groups and “clays”, the following chaînes opératoires were identified. The author argued that artisan groups correspond to each operational chain:

Chaine opératoire 1 consists of pottery made using Clay B by potters using techniques linked to forming Group 3, on plates and Type 1 shallow bowls. Choleva argues that the employment of RKE from early stages of forming, and the rarity of this chain suggests the presence of specialised craftspeople. Products of the chaîne opératoire are treated as imports on the claims that Clay B bears macroscopic similarity to fabrics from Lefkanti, Ayia Irini, and Pefkakia. Lefkandi, especially, is considered by the author a “production centre” for wheel-made plates (Choleva 2015: 159).

Chaîne opératoire 2 pottery is made from the relatively common Clay A and is associated with forming Group 2, mostly to produce Type 1 shallow bowls and a small number of plates. Potters linked to this chain are considered locals, and less specialized in using the potter’s wheel than the Group 3 potters, because RKE is employed at a later stage of forming, possibly indicating attempts by local specialists to incorporate the wheel in local practices and adapting the wheel-coiling technique to pre-existing coil-based hand-based techniques (Choleva 2015: 166). It is difficult to discern whether there was an intentional attempt to copy the work of the more highly specialised potters of Group 3, although this is possible considering the fact that both groups are contemporary.

Finally, chaîne opératoire 3, linked to the abundant Clay A and featuring vessels of all known shapes of shallow bowl, is seen as indicative of local mass production, and its largely handmade character is seen as indicative of non-specialisation. Its most interesting feature in terms of the effect the potter’s wheel had on local production, is that there is evidence that locals attempted to imitate surface features of wheel-coiled pottery through surface treatment imitating the “horizontal striations” usually associated with the potter’s wheel (Choleva 2015: 166).

The sample of handmade shallow bowls and plates was relatively larger n=124, of which 6 were plates, 103 were Type 1 shallow bowls and 15 were Type 2 shallow bowls. According to the author, they were mostly made using a coiling technique (Choleva 2015:159). In terms of the relative percentages of wheel-made and handmade pottery from the sample (n=189), 65 (34%) were wheel made and 124 handmade (66%).

Viewed through a cultural evolutionary lens, this narrative of local potters imitating highly specialised potters, although not recognised as such by the author, possibly points to horizontal and oblique transmission, with local tendencies towards anti-conformism in the form of a prestige-based bias, with Group 3 acting as a model in the social learning process.

4.2.1.2. Lerna IV

Lerna is located in the Peloponnese, in mainland Greece. The wheel was first introduced during the EBA III period, on occupation level Lerna IV, which has been divided into three chronological phases (Choleva 2012: 350, Rutter 1995: 3-10). Choleva re-examined all pottery sherds that had been characterised as "wheel-made" by Rutter (n=48), except 11 rim fragments which were deemed inappropriate for analysis (Choleva 2012: 358, 360) using Roux's framework in order to identify different stages of wheel-coiling. Two groups were identified on the basis of forming technique using RKE:

Group 1 features the equivalent of Roux's Method 1 (found in 22 sherds), where RKE is employed minimally at the late stage of the forming of the vessel "rough-out". It also includes 8 sherds produced using Roux's Method 2, which employs RKE from the previous stage of forming, that of thinning the walls (Choleva 2012: 364-365). Group 2 fragments are relatively scarce (n=4) and indicate increasing employment of RKE. 2 sherds were produced using Method 3: the use of RKE to join the coils, thin out walls, and shape the rough-out. 2 sherds bear traces of Method 4: using the wheel from the initial stage of forming the coils. The author also identified three (out of the initial sample of n=48) vessels that were considered coil-built but mostly handmade, with minimal use of RKE for surface treatment (Choleva 2012: 367).

A significant contribution of this study is the fact that the author reports the relative abundance of each fashioning method per shape and fabric group. Overall, the picture that arises from the study is that the rate of adoption was slow, leading the author to assume local potters were “resistant to change” (Choleva 2012: 376). An operational chain linked to the production of the “Fine Gray Burnished wheel-fashioned pottery” is identified and considered the predecessor of “Minyan” pottery. Its most characteristic traits are grey fabrics linked to reducing firing conditions, burnished surfaces, and specific shapes (Choleva 2012: 373-374). This could be interpreted as an attempt to imitate the appearance of metallic vessels. Furthermore, Choleva sees this operational chain as evidence for a group seeking to differentiate themselves from other potters, producing products of high aesthetic value (Choleva 2012: 376).

Overall, the study has relative strengths compared to the analysis of pottery from Palamari and offers promising avenues for future research. Percentages of wheel-made pottery over handmade pottery are provided for each sub-phase (1, 2, and 3) of Lerna IV, as well as percentages of the four different forming methods per sub-phase, and quantitative data on the prevalence of specific shapes (bass bowl, tankard, kantharos, rim-handled cup, and one-handled cup) per phase (Choleva, 2012: 371-373). While it must be noted that the sample is relatively small, the analysis includes quantitative data that could be integrated with those from other studies in a secondary dataset.

Reference should also be made to studies utilizing Roux and Corbetta's (1989) observations to infer the expertise of potters using the wheel, mainly based on observations on vessel height, whilst carrying out more 'traditional' analysis in terms of identifying wheel-made from handmade pottery without examining whether there are intermediate stages of coil-building. Berg's (2007) study on pottery from MBA Phylakopi is a good example of this kind of approach and is reviewed below.

4.2.1.3. Phylakopi

It is generally considered that the wheel was introduced in Phylakopi, an island in the vicinity of Crete, in the Middle Bronze Age, through contacts with Minoan Crete, although it is argued that complete adoption occurred very gradually (Berg 2007). In order to trace the trajectory and local adaptations of wheel-based pottery techniques, Berg carried out macroscopic analysis on pottery (n= 74,557 sherds "and 479 complete vessels") from the 1974-77 excavations of Phylakopi, dated from the Middle Cycladic to the Late Cycladic II phases, focusing on decorative motifs, forming techniques, fabric, and shape. The author makes some important remarks on certain aspects of the Phylakopi excavation which impede contextual understanding of the pottery: the small size of the trenches meant that no complete architectural units were excavated (which made securely identifying workshops impossible), and the pottery, although well preserved and retained in its entirety, contained a large number of small fragments, which proved problematic for identifying shapes, decorative motifs, and hybrid forming techniques. Another drawback is the fact that the Middle Cycladic is not as well represented in the dataset as the Late Cycladic phases (Berg 2007: 242, 248).

In terms of fabric, the author identified two groups via macroscopic analysis: “Local Fabric Group” and “Cycladic White Fabric Group”, which had also been identified through petrographic analysis by other researchers. The first refers to a local, orange-brown, coarse fabric, which was identified on pottery of all shapes and surface treatments and includes a sub-group (“conical cup”) made of softer clay, which is associated with small undecorated Minoan shapes (wheel-made). The second refers to a whiter fabric of “finely levigated clay”, which was used for local wares following the Melian tradition (Berg 2007: 243).

The study produced good quantitative data on the percentages of wheel-made pottery per phase, including the sample size provided per phase, as well as relative percentages of each fabric group per phase, including sherd counts (albeit provided only for wheel-made shapes). Finally, quantitative data were provided identifying the most common forming technique per phase (handmade or wheel-made). Emphasis was also placed on identifying Minoanizing shapes (Berg 2007: 244-245).

Overall, Berg identified two local pottery traditions, one linked to producing Minoan imitations (with Minoan decorations and the use of slips to imitate the appearance of Minoan fabrics), and the other producing handmade Melian pottery (with decorative motifs stemming from local traditions) and deduced that the wheel was mostly used to produce small shapes of Minoan origin, even though there are also large handmade minoanizing shapes made using the wheel. There is also a correlation between forming technique and fabric: wheel-made shapes were made using a softer variant of a local reddish fabric, Cycladic white was used to produce Melian handmade decorated pottery. The author suggests that local potters used the wheel only when they wanted to imitate Minoan pottery, which probably became a status symbol because of its appearance and function in feasts, but did not incorporate it in the local tradition (Berg 2007: 247-9). Interestingly, in other Cycladic settlements like Ayia Irini in Kea, the wheel was used to produce both local and Minoanising shapes (Berg 2007: 249), which suggests that it was better incorporated into local traditions.

Finally, it is interesting that the wheel was initially used to produce small open vessels, and gradually used to produce larger and more closed forms (height measurements are provided). The author observes that this pattern in Phylakopi is strikingly similar to that of other settlements not only in the vicinity of Melos, like in Ayia Irini, Akrotiri, and Knossos, but also as far away as Toumba in Macedonia, and that the pattern might be better explained in terms of physical and motor-skill constraints as the ones identified experimentally by Roux and Corbetta (1989), who observed that small open vessels present fewer difficulties in wheel-form than larger and more closed shapes (Berg 2007: 243). This also allows an indirect approximation of the skill level of Phylakopi potters per phase which suggests that local artisans were probably not "experts": Berg found evidence of non-standardized production in Phylakopi, on the basis of morphometric similarity (Berg 2004: 82).

Considering the findings of the recently published aDNA study by Skourtanioti et al. (2023) regarding patterns of endogamy in the Aegean Bronze Age, it would be interesting to examine whether the difference between the levels of integration of the wheel in local production between Melos (Phylakopi) and Kea (Ayia Irini) can be explained in terms of differences in marriage patterns. In this case, for example, the lack of integration of the wheel in local production that Berg reports for Phylakopi could be due to a tendency of potters trained in Minoan traditions to only marry within their group, and therefore restricting the transmission of the wheel beyond the group. In the final subsection dedicated to data on wheel-made pottery, I will briefly mention studies that combine Roux and Courty's (1998) framework with an experimental approach.

4.2.2. Experimental approaches

In order to examine the transmission of wheel-made pottery on Crete, Jeffra (2013: 33) adopted an experimental approach where she first built experimental vessels based on Roux and Courty's (1998) framework of wheel-fashioning methods and which had the diagnostic criteria "macroscopic traces" associated with each forming method. This was followed by macroscopic inspection of pottery from Knossos, Palaikastro, and Myrtos-Pyrgos comparing each sherd to the experimental vessels to infer what their forming method must have been on the basis of their similarity to the experimental vessels (Jeffra 2013: 33). Evidence from the pottery suggested that the generalized use of the wheel was linked to elites in the context of feasting, as the most widely produced wheel-made shape was the conical cup,

Jeffra's goal was to test the hypothesis that the adoption of the potter's wheel was linked to a "small, closed group of specialists" (Jeffra 2013: 43) with the assumption being that if that was the case then empirically one would observe high adoption of a single method of RKE (a method from Roux and Courty's framework) near the origin (Jeffra 2013: 32) and that a later stage of "diversification of methods as other potting groups learned through trial and error".

On the basis of her analysis, Jeffra concluded that instead adoption in Crete happened gradually over time and was characterised by a progression from the use of diverse methods of wheel-coiling to a generalized use of Roux's and Courty's (Roux and Courty 1998: 751) method 3 (Jeffra 2013: 43), which involves a heavier reliance on Rotative Kinetic Energy produced by a rotative device in order to shape the vessel compared to other methods.

4.3. Kilns

As mentioned previously, in the Aegean during the Bronze Age, new technologies were also developed in the context of pottery firing by introducing more stable contraptions to achieve closed firings. We tend to associate such contraptions with built structures such as kilns, but there is evidence that less formally structured contraptions such as the ones used during clamp firings (Ballut et al. 2017, Jones 2021: 82) were effective ways of controlling the firing atmosphere — oxidising (oxygen-rich) versus reducing (oxygen-poor) — during the Neolithic.

It is generally acknowledged that there are no physical remains of kilns that are definitively dated to the Neolithic, and that the evidence for the use of such contraptions can mostly be found on the finished products themselves (Jones 2021: 82). This is where experiments have been useful in attempting to reconstruct the contraptions that were used for firing the pottery as well as the firing process itself (e.g., by Vitelli 1994). Nevertheless, it is quite certain that kilns were used in the EBA (Jones 2021). Kilns became more prevalent during the MBA and continued to rise in the early LBA, dwindling towards the end of the period (see Figure 6 in the following chapter) – which is, perhaps not coincidentally, the period when the palatial systems in Crete and Mainland Greece collapsed (Eder & Lemos 2019:135).

In this section I will attempt to trace the process of kiln adoption during the Bronze Age through the narrative that has been built by scholars specialising in firing technologies. This qualitative picture will be enhanced by the quantitative analysis of the empirical data emerging from some of those studies, which is included in the following chapter.

4.3.1. Incentives driving kiln adoption

The incentives behind each individual decision to adopt kilns may remain unknown to us, but it is reasonable to assume that generally the incentive relates to the material properties and aesthetic traits of the finished products. There is a large literature on whether and how firing temperature, the heat distribution inside the firing area, and the rate of change of the firing temperature differs for pottery fired in a kiln versus pottery fired in other contraptions such as pit kilns, and I will briefly touch on the main facets of that literature in the section below. However, it is important to note that, at least as far as ethnographic evidence from modern potters suggests, potters may adopt kilns for reasons other than the immediate technical aspects relating to the firing conditions inside the kiln (Roux et al. 2018). For example, in Roux and Lara's study (2016) some potters rejected kilns, among other reasons, because they were not interested in changing the type of fuel they had been using for open firings and because firing using a kiln would require them to sit next to the kiln and not complete other chores during the firing stage as they had done with open firings (Roux et al. 2018: 1047). Additionally, interviews with expert potters indicated that, when evaluating the technical advantages of kilns over open firings, they focused less on the technical advantages of kilns and more on the economic and organisational advantages of firing using kilns (Roux et al. 2018: 1048).

This suggests that, if one is interested in studying kiln adoption at the level of the community, the performance characteristics of the kilns are only one side of the coin. The other side relates to the organisation of production, cost-benefit decisions made on the basis of resource availability, and factors associated with social learning aspects of adopting a new firing technology – for example, Roux et al. (2018) found that inexperienced potters were more likely to adopt kilns when most of the other potters had done so (a conformist bias) than expert potters, who adopted on the basis of technical-economic-organisational factors.

Regarding the performance characteristics of kilns, controlling the firing atmosphere in terms of oxygen content and controlling the rate at which the temperature rises in the kiln emerge as the most significant advantages of using kilns. Experiments have shown (Ther 2014) that Bronze Age kilns did not generally achieve higher temperatures compared to open firings, but the assumption that kilns equate to high temperatures is still prevalent in the literature (e.g., Rice 1987: 153). Kilns allow reliably alternating between an oxidising and a reducing atmosphere by allowing air to enter the kiln via openings on the kiln chamber in the case of oxidised firing, or by restricting the flow of air inside the kiln by shutting off the openings in the case of reduced firing. Oxidised firing results in reddish and brown colours in the kiln -the redness of a fired pot also depends on factors like how much iron is in the clay (Orton et al. 1993: 69), as iron turns red when it reacts with oxygen. Reduced firing, on the other hand, results in more greyish and black colours on the surface of the ceramic (Maniatis 2009: 8). This allows achieving polychromic effects on the surface of the pot by using slips with a slightly different consistency to the clay used for the body of the vessel (see section 4.3 on the iron-reduction technique (IRT) which involves using an iron-rich slip to produce lustrous dark-coloured decoration).

Another significant advantage of using kilns is that they allow controlling how quickly the firing temperature rises during firing (Sillar & Tite 2000: 13, Gosselain 1992). Extreme rises in firing temperatures can lead to thermal shock, which results in fractures on the pot (Tite et al. 2001: 313). This is especially important when the fabric is fine, as experimental studies have confirmed that the addition of temper improves thermal shock resistance (Tite et al. 2001: 321). Interestingly, most wheel-made pottery in the Bronze Age was on the finer end of the scale in terms of fabric coarseness; for example, Choleva (2020: 85) states that the fabric of wheel-made pots was very fine in Emporio and ranged from fine to medium coarse in Heraion in EBA II. Therefore, it is very likely that the adoption of kilns had some kind of amplifying effect on the adoption of wheel-made pottery, at least at the level of individual communities - Sillar and Tite (2000: 5) argue that this effect has been observed in multiple contexts but is not “universal”. Unfortunately, it is hard to say whether this effect was pronounced enough so as to affect the transmission process at the macroscopic scale of the Aegean; as I will briefly discuss below (and in greater detail in the following chapter), the temporal resolution of dated kilns is usually coarser than the resolution of dated wheel-made pottery, making it difficult to assess whether the rate of kiln adoption and the rate of wheel-made pottery adoption covary.

4.3.2. How is kiln use identified in the archaeological record?

Kiln use is difficult to quantify, especially in terms of quantifying relative kiln use as a percentage of the population that uses a kiln over the percentage that does not. When kilns are identified, they indicate presence on a presence absence scale; in such cases, we definitely know that a kiln was being used at that specific point in time. The problem is that when a kiln is absent, we cannot necessarily infer that a kiln was not being used. As the saying goes, “absence of evidence is not evidence of absence” (Sagan 1977: 7). Researchers have sometimes been able to infer the presence of a kiln-like contraption by looking at the pottery — a good example is the study of Late Neolithic and Early Bronze Age pottery from Phaistos, Crete, by Mentesana et al. (2019). Mentesana et al's results suggest that a kiln was used based on the fact that the pottery shows signs of increased control of the firing conditions, but the authors are careful not to claim that the results suggest the use of a kiln per se, but rather that: "firing procedures did change in terms of potters' control over the firing conditions towards the beginning of the EM IA in terms of temperature rate" (Mentesana et al. 2019: 891) and that it this was due to new type of "equipment".

Regarding the spatial location of confirmed kilns, the main source of information is Hasaki's (2002) 'Atlas of Kilns', a database that includes published kilns dated to the Bronze Age. Hasaki (2002) provides information on the location of each kiln and other morphometric measurements, and assigns each one to her one typological schema based on the morphological characteristics of the kilns. Most kilns in the database are assigned to one of three-time bins; Early Bronze Age, Middle Bronze Age, and Late Bronze Age, while others are dated within a narrower time slice. This makes any attempt to quantify the relative abundance of kiln use over time challenging and necessitates the use of statistical treatment in order to account for the in-phase uncertainty of time assignment for each kiln (discussed in detail in the following chapter). Jones also suggests that the difficulty of identifying kilns also varies per time period, stating that "the identification of kilns during the EBA remains problematic" (Jones 2021: 87-88). This means that there is greater uncertainty in earlier periods, and that any attempt to plot the number of kilns over time will indicate deceptively low numbers for the Early Bronze Age.

Linking each kiln to a specific potting community is also challenging. While there are sites with a single kiln, there are also central 'kiln sites' with multiple kilns — e.g., Kato Gouves, (Chatzi-

Vallianou 1997) — which were probably sites where potters from multiple communities gathered to fire their pottery. To make matters more complicated, the presence of a kiln in one site does not mean that every potter in that community was using a kiln — it is just as likely that some of the potters in the community were using other contraptions or open firings, which leave little archaeologically recoverable evidence. Therefore, unlike the published data on wheel-made pottery, there is no way of knowing what the relative use of a kiln was at the level of the community, making it difficult to compare the datasets for the two innovations — even though, as discussed above, it is very likely that kiln adoption had some effect on the adoption of the potter's wheel.

4.4. The iron-reducing technique for decoration (IRT)

Novel techniques were also developed in pottery decoration. One of the most interesting cases is the development and production of iron-based clay slips which were combined with certain firing techniques to produce lustrous decoration — as opposed to the matt effect produced by clay slips containing manganese (Jones 2021). Out of the three innovations discussed here, the iron-reducing technique for decoration was possibly the one whose adoption was most affected by the choices of consumers: while the technique requires significant technical expertise, which suggests that specialised potters made the final decision the adopt, it also results in observable aesthetic traits (lustre and colour), which means that consumers could visually identify it and partially select it through their preferences, even as an indirect influence.

Jones (2021:67), defines the new technique as the "iron-reducing technique" which involves the use of a clay rich in iron, or an ochre with iron-oxides. According to Jones (2021:67) the process during firing works like this: "during the reducing phase of the firing, iron oxides of dark colour are formed. The final colour depends on the nature of the paint material's composition, the thickness of the paint layer, and the atmosphere during the final stage of firing and during cooling". The new technique took advantage of the changing levels of oxygen in the kiln to produce brown, red, and black colours, with more reddish colours being the result of "an overall oxidising atmosphere" (Jones 2021:69) and it was possible to achieve variations in the colour of the slip by varying the thickness of the slip layer: "a thin layer coloured black during the reduction phase may be subsequently re-oxidised either partly (to brown) or wholly (to black)" (Jones 2021:69). This was due to the fact that "a thicker layer remains black because the more effective sintering that takes place during the reduction phase protects it from later reoxidation" Jones (2021:69). For ease of reference, I will adopt Jones' terminology and refer to this trait as IRT for the rest of this chapter.

4.4.1. The spatial and temporal trajectory of the transmission of the IRT

According to Jones, an interesting aspect of the IRT in the Bronze Age is the fact that it "was one of the few techniques to make progressive technological advance" (Jones 2021:67). Indeed, in his long-term account on the development of Greek pottery firing and decorative techniques, he sees the neolithic as a period of experimentation, which is followed by the development of "high status high quality" pottery in Crete during the Middle Bronze Age (Kamares ware) (Jones 2021: 121) and during the end of the period, the production of such "high quality" pottery in a standardised manner. The view that the standardised lustrous decorated pottery spread to other parts of the Mediterranean because of its quality is shared by Betancourt et al. (2016: 64) who state that: "It is not surprising that this fine technique of producing pottery would have spread to the mainland of Greece by LH I and LH IIA" (Betancourt et al. 2016: 64). For Betancourt et al., the IRT trait is part of a technological package involving the use of the potter's wheel, a diluted iron-based slip, and three-phase firing (oxidation-reduction-re-oxidation) whose success depends on its desirable technical traits (or using a term from the cultural evolution literature, a 'content-based bias'). However, elsewhere in the same discussion, they suggest a spatial aspect to the transmission of that package, as they argue that the Mycenaean lustrous decorated pottery package, which was descended from the Cretan one, 'replaced' local technological packages : "the vessels produced in this way were more desirable for fine dining and drinking than earlier vessels, and even if the copies did not achieve the outstandingly lustrous effects of the original MM Ill to LM lA southeast Cretan products, they were still more desirable than many of the materials they replaced"

(Betancourt et al. 2016: 64). It is not unreasonable to describe this process as the spatial transmission of a set of practices in pottery-making which involves local potters making a mutually exclusive decision to accept or reject the innovation where accepting it means that they stop making pottery using the set of practices they had been using previously.

In other words, the narrative by Betancourt et al. implies that the package originated from southeast Crete, spread to the rest of the island without significant changes (but note the relative lack of lustre as one moves further away from the origin) and then to the mainland, replacing other ways of making pottery. The process where a package of traits originates from a specific location and is transmitted in space over time with most of the traits linked together but with some traits missing or modified - in this case, modifications could consist of things such as the lustrous trait becoming less pronounced — as one moves away from the origin brings to mind the notion of "isolation-by-distance" from evolutionary biology, which has also been used in archaeological applications of cultural evolutionary models (Shennan, Crema, and Kerig 2015).

The notion is described in greater detail in chapter 6 (it is one of the main starting assumptions of the simulation model explored in chapters 7 and 8), but Shennan, Crema, and Kerig's 2015 (2015) description of the implications of this model in the spatial distribution of cultural traits is relevant here: "a general rule of thumb is to start by assuming that the frequency of cultural transmission is mostly characterised by a distance decay, where the greater the inter-distance between the donor and the recipient, the less likely is the occurrence of a transmission event. This implies that, other things being equal, the similarity of cultural traits should also decline over distance" (Shennan, Crema, and Kerig 2015: 104). This fits Betancourt et al's description of the transmission of the package from southeast Crete to the Greek mainland and its development into the Mycenaean lustrous decorated pottery which retained most of the same traits (wheel-made, kiln-fired, three-phase firing) except for the high degree of lustre.

It must be noted, however, that in a different discussion on this technological package in the site of Bramiana in south-east Crete, which Betancourt et al. see as one of the core areas/origin of the package, Betancourt et al. (2021) attributes the relative lack of lustre for pots made using the same technology in other parts of the Aegean to the lack of calcareous and illite-rich clays that were, according to the authors, instrumental in achieving a highly lustrous appearance without the surface bloating (Betancourt et al. 2021: 114). In other words, Betancourt et al. partially associate the 'lustrous' effect with local geology, which is another factor generating archaeological signatures that resemble the results of an isolation-by-distance process. It is best to quote Betancourt et al. directly here: "By LM I, lustrous dark-on-light pottery was being produced across the Aegean. The imitators could duplicate and improve on the style of decoration and the firing, but they never understood all the complexities of the process, and they did not have access to the highly calcareous and special clays of Southeast Crete with their high illite content. As a result, the LM I and Late Helladic I potters of other places never quite matched the finest surfaces of the MM III potters who invented the new workshop processes in the first place." (Betancourt 2021: 114).

This implies that potters in other areas lacked some of the information that would have helped them copy the package faithfully, and it suggests that the strength of communication between potters in Crete and the mainland was weaker compared to that between potters within Crete. Of course, it is unlikely that Cretan potters were as aware of the chemical processes occurring during the vitrification of the slip as modern archaeologists are, let alone linking them to local geology in a reliable manner, but it is possible that they could recognise the relative merits of different types of clay and communicate them to other potters. Even if local geology partially determined how successful cultural copying was, a lack of direct communication between Cretan and mainland potters — and by direct, I mean communication in the context of an apprenticeship or an extended period of staying in another community so as to be able to recognise the ideal clay/tempering materials visually — could be the reason why we see less similarity as spatial proximity decreases. It becomes clear that, as far as IRT decoration is concerned, its transmission was highly dependent on the development of firing techniques and specific firing structures. Therefore, it is difficult to disentangle the IRT decoration trait from traits related to firing and investigate the transmission of the IRT trait in isolation from the transmission of firing techniques and structures. We can, however, attempt to look at how it spread in space through time in a qualitative manner. It is fortunate that there is a number of comprehensive reviews of the available data such as Kiriatzi's (2010), Betancourt et al's. (2016), and Jones' (2021) that have identified the phenomenon in specific places and times during the course of the Neolithic and the Bronze Age and brought together the available published evidence from various sources (note, e.g., Kiriatzi 2010: 685-690 and Betancourt et al. 2016: 64 where the authors provide a comprehensive list of site locations and the publications associated with each site).

The use of lustrous iron-based slips produced using a reducing atmosphere is very old, and some scholars place its origin in Mesopotamia (Jones 2021: 69 citing Noll and Heimann 2016:225–27). However, the scope of the present study is the Aegean and Western Anatolia, so I have only attempted to trace the trait in this region chronologically by reviewing the information included in the aforementioned studies starting from the Neolithic all the way to the LBA. It should be noted that out of the three reviews mentioned, only Jones (2021) identifies and traces the IRT trait separately from other technological aspects of the pottery — Betancourt et al. (2016:6) discuss IRT as a trait characterizing Minoan dark-on-light ware and therefore trace the entire package/ware in space and time, while Kiriatzi (2010) discusses IRT as a trait characterising the 'Minoanising' pottery that borrowed elements from the Minoan dark-on-light and spread to the coasts of the Peloponnese possibly through Kythera — Kiriatzi considers 'Lustrous Decorated' ware and 'Red Silver Micaceous' ware as instances of this technological package in the Peloponnese (Kiriatzi 2010: 685). Thus, Kiriatzi's discussion also traces the 'Minoanising' package as a whole — but note that she uses IRT as a diagnostic criterion distinguishing Lustrous Decorated from Matt-Painted ware e.g., Kiriatzi 2010: 686, note no. 17).

Jones (2021) argues that the use of such lustrous decoration, which he considers a sintered "silicate paint", began from the "Urfirmis of the Neolithic and Early Bronze Age" (Jones 2021: 69) and was still being used into the Bronze and the Iron Age. He notes that Urfirmis ware can be found in multiple areas of the Peloponnese ("Asea, Franchthi Cave, and Lerna, and probably extending to Kouphovouno") with significant stylistic uniformity unlike other types of pottery made in this period (Jones 2021:82) and sees this as evidence that the potters who made this type of pottery and were based in different sites interacted with each other. Some also argue that the Urfirmis ware in each site was made by same group potters who were moving between these different locations (Jones 2021: 82, Vitelli 1984:126) possibly in the context of marriage (Cullen 1985: 96) in a patrilocal system (female potters relocating to the community of their spouse).

During the Bronze Age, one early instance of this trait was in Crete during the beginning of the Middle Minoan period. More specifically, the trait occurs on the Kamares ware, a group consisting of pots probably intended for elite consumption (Day et al. 2006, Jones 2021: 84), which was also imitated in a Cycladic site with heavy Minoan influence (Akrotiri), but with less 'glossy' results (Jones 2021:87).

Kiriatzi's (2010) study provides valuable information regarding the locations and relative abundance of the Minoanising package that involves the IRT trait, and she also highlights some of the factors that make identifying it challenging. As the author herself states: "It is not easy to trace the spatial distribution of this pottery since its identification or description has not been consistent through time and across the different sites, while in many cases there are only rough assessments of quantity" (Kiriatzi 2010: 686). However, Kiriatzi argues that it is possible to qualitatively discern a coastal pattern to the spread of this new type of decoration/pottery.

Fortunately, Kiriatzi's (2010) narrative can be used to inform a qualitative reconstruction of the temporal trend of the transmission of this 'Minoanising' pottery (which includes the IRT trait): "Concerning chronology, in the few sites where relevant data are available, the mainland 'Minoanising' pottery first appeared in MH I, possibly together with vessels imported from Crete (Agios Stephanos, Asine, Aspis, Kolonna, Lerna, Nichoria) and became more common in MH II (Asine, Kolonna, Lerna, Nichoria), while differentiated patterns appeared in MH III-LH I, when, for example, it becomes more popular in Agios Stephanos but drops in Asine, Aspis and Nichoria" (Kiriatzi 2010: 688-689). The picture drawn here is, like Betancourt et al.'s (2016) narrative on the spread of the lustrous decorated package in Crete, consistent with a process of 'isolation-by-distance', only this time the entire island of Crete functions as the origin point. The package arrives to the Peloponnese, the part of the mainland closest to Crete, enjoys increased popularity in the next few hundred years (MH II) and then dwindles in the areas in the Peloponnese that are located the furthest away from Crete (Asine, Aspis, and Nichoria) but increases in areas that are located near Crete (Ayios Stephanos). The 'loss' of the package in faraway areas compared to areas closest to the 'origin' can be explained in multiple ways, but the notion of cultural drift is probably relevant here.

As discussed in Chapter 2, cultural drift is a process through which the frequency of traits in a population rises and falls purely due to chance. In this example, it is not unreasonable to assume that potters in Asine, Aspis, Nichoria, and Ayios Stephanos would be equally likely to 'lose' the trait by chance: a prominent teacher dying before teaching apprentices, pottery with the trait falling 'out of fashion', etc. However, if those communities/sites enjoyed varying levels of communication with potters from Crete based on their distance in space (e.g., if being near Crete meant more frequent visits from Cretan potters or greater numbers of Cretan potters visiting one's community), we should expect communities located at a greater distance to Crete compared to others to be more likely to not re-gain the trait after losing it due to chance.

Overall, it looks as though the IRT trait first appeared in the context of the Neolithic Urfirmis pottery in the Greek Mainland. Afterwards, it re-appeared in the context of the 'high-status' Kamares ware in Crete and spread to the rest of the island and then to the Peloponnese, possibly through Kythera (as suggested by Kiriatzi (2010)). Finally, it became part of a Mycenaean technological package, which sped up its transmission in the rest of the Aegean.

4.5. Synthesis

While the picture presented here is slightly complicated, it is worth making some preliminary observations on the adoption trajectories of the three innovations before delving into the quantitative analysis of the (sparse) quantitative data emerging from the literature. All innovations are technologically linked to each other to some extent, except perhaps wheel-made pottery and IRT, even though it is possible that were transmitted as part of the same package at some point during the Bronze Age. Out of the three innovations, the IRT trait is the one which could throw more light on the contribution of consumers and their preferences on innovation adoption. However, it is also the least consistently identified by specialists and there is a lack of published quantitative data on its prevalence at the site level in the literature. While it probably emerged during the Neolithic in the mainland, it is unclear whether its transmission trajectory was linked to the Cretan IRT of the Bronze Age, or whether the knowledge on how to produce it was lost and then it was randomly re-invented in the Bronze Age. However, its appearance in the Bronze Age (Kamares ware), was almost definitely the beginning of a transmission process that unfolded through the rest of the Bronze Age, spreading to the Peloponnese in the mainland and then north all the way to Macedonia as part of a Mycenaean technological package including the other two innovations.

Kilns are linked to both wheel-made pottery (due to the fine fabric of wheel-made pots which create the need for better control of the firing temperature's rate of change), and to IRT (kilns provide control over the firing atmosphere and allow reliably alternating between oxidising and reducing phases). Their identification of kilns belonging to early periods is patchy, and there are few ways of inferring kiln absence in the archaeological record. It is fortunate that scholars have published databases with information on the location and technical characteristics of kilns, which enable quantitative investigation (included in the following chapter).

The most securely identified innovation is the potter's wheel, especially because of recent efforts to quantify the relative proportions of wheel-made pottery over hand-made pottery over time. Such efforts are supported by a long-established set of ethnographically informed diagnostic criteria that are visible through macroscopic inspection of pottery sherds, and by decades of ethnographic research that have provided information on the social context of the transmission of pottery forming techniques. The scholars that apply this set of diagnostic criteria on pottery from the Bronze Age Aegean have been very productive in recent years and provide enough quantitative data to perform quantitative analysis on a secondary dataset built by integrating their primary data and even to inform the development of a simulation model.

5. Empirical analysis of adoption trends and population size data

This chapter describes the process of integrating primary data from the literature into secondary datasets to trace the adoption patters for the innovations discussed in chapter 4. Section 5.1 introduces a method for integrating data from different regional chronologies. Section 5.2. is divided into two sections that present and analyse two secondary datasets produced by integrating primary data from the literature. Special attention is given to addressing uncertainty by using Monte Carlo simulation. Section 5.3. presents the dataset used to trace the adoption of kilns over time, mainly using data from Hasaki (2002). Temporal uncertainty in the dating of kilns is addressed by performing aoristic analysis. In section 5.4 I review the available evidence on population levels in Aegean during the Bronze Age and explore a number of approaches for acquiring an empirical estimate of population size over time and linking it to specific sites/locations. Finally, section 5.5 summarises the main findings of the analyses included in sections 5.1-5.4.

5.1. Chronological framework

In order to trace the increase in the prevalence of a trait over the Aegean in time it is important to establish whether dates are given in terms of an absolute or a relative chronology. Aegean archaeology's reliance on relative chronologies (Minoan, Helladic, Cycladic, Anatolian) constructed by establishing regional ceramic typologies is a fact generally acknowledged (Doumas 2009) and pottery is dated using relative chronologies. Several issues arise because of this reliance on pottery typologies, including the uncertain placement of some phases — such as the Lefkandi I- Kastri phase (Angelopoulou 2008: 149) — in relative chronologies and the difficulty of conducting chronological comparisons between assemblages belonging in different regional typologies.

There have been attempts to produce a general chronology linking different regions and to produce chronological bins providing absolute date ranges for each phase, like Manning's (2010) chronology which was used here with the following modifications: the absolute date bins provided for each of the three regional chronologies were aggregated to produce a single phasing system comprising 14 phases, using the following rule:

1. For start and end dates: If two dates were provided: e.g., "2450/00", the mean value, 2425, was used as a start date.
2. After step 1, a mean value was procured from all three start dates, and the same process was used for all three end dates per bin.

Some of the strengths of this approach are that it produces no overlap between different phases, and that it allows comparing the arrival time of each innovation across all the sites included in the study, regardless of its location. This system also allows quantifying the length of each phase and incorporating it into the analysis. On the other hand, this method heavily relies on using estimated date ranges for some regions on the basis of subdivisions for other regions. For example, the system introduces a tripartite subdivision for the Middle Helladic in Mainland Greece using the Cretan subdivision into Middle Minoan I, II, and III to integrate data from the Mainland with Cretan data — in reality, it is unlikely that there would have been a clear association between the subdivisions of that period in the mainland and in Crete. As Spencer notes, this is "a system developed for use within a Cretan framework and not necessarily consistent with developments on the Greek mainland" (Spencer 2007:12). However, the fact that the study does not focus on the dating and evolution of typological traits (whose dating would depend more heavily on such typologically based subdivisions) probably means that this weakness will not affect the results significantly.

Phase	Cyclades_pot	Mainland_pot	Crete_pot	Crete	Mainland	Cyclades	Aggregate_start	Aggregate_end	Aggregate_duration
2650-2475	ECII (Keros-Syros phase	EHII	EMIIA	2650–2450/00	2650–2500	2650–2500	2650	2475	175
2475-2217	Kastri	EHII late/Lefkandi I	EMIIB	2450/00–2200	2500–2200	2500–2250	2475	2217	258
2217-2075	Kastri to Phylakopi I	EHIII	EMIII	2200– 2100/2050	2250– 2100/2050	na	2217	2075	142
2075-1912	MC-Phylakopi I	MH	MMIA	2100/50– 1925/00	2100/2050-	2200-	2075	1912	163
1912-1863	na	MH	MMIB	1925/00– 1875/50	na	na	1912	1863	49
1863-1725	na	MH	MMII	1875/50– 1750/00	na	na	1863	1725	138
1725-1687	na	MH	MMIII(A-B)	1750/00– 1700/1675	na	na	1725	1687	38
1687-1615	LCI	LHI	LMIA	1700/1675– 1625/00	1700/1675– 1635/00	1700/1675– 1625/00	1687	1615	72
1615-1470	LCII	LHIIA	LMIB	1625/00– 1470/60	1635/00– 1480/70	na	1615	1470	145
1470-1415	na	LHIIB	LMII	1470/60– 1420/10	1480/70– 1420/10	1420/1400-	1470	1415	55
1415-1380	LCIII	LHIIIA1	LMIIIA1	1420/10– 1390/70	1420/10– 1390/70	na	1415	1380	35
1380-1323	na	LHIIIA2	LMIIIA2	1390/70– 1330/15	1390/70– 1330/15	na	1380	1323	57
1323-1195	na	LHIIIB	LMIIIB	1330/15– 1200/1190	1330/15– 1200/1190	na	1323	1195	128
1196-1062	na	LHIIIC	LMIIIC	1200/1190– 1075/50	1200/1190– 1075/50	na	1196	1062	133

Table 1: Manning's (2010) absolute chronology and 14 derived phases, including the duration of each phase.

5.2. The potter's wheel dataset

5.2.1. Presence-absence data

It was possible to create a dataset recording the presence and absence of wheel-made pottery at each site, focusing on the best stratified and fully published sites in the dataset, preferably with long habitation sequences. Given that the authors of the original studies dated the first appearance of wheel-made pottery at each site using traditional regional chronological schemas, I initially recorded each presence/absence using the same system. The sites and the regions they belong to can be found in Table 2 below, along with the references to the publication reporting the presence/absence of wheel made pottery for each site.

The map of locations in Figure 5.1 shows that Mainland Greece is slightly overrepresented compared to the other regions (but then this region is significantly larger than the Aegean islands and Crete in terms of inhabitable surface area), and that sites are typically located near the coast –

it is hard to say whether this reflects a genuine habitation preference or whether it is due to some form of recovery or publication bias.

As shown in Figure 5.2, the trend over time of the percentage of inhabited sites in the dataset where the potter's wheel is recorded as present resembles a logistic curve. The trait first appears in the 2475-2217 BC phase, increases slowly until the 1912-1863 BC phase, at which point it increases rapidly (this is roughly the time when palatial systems first made their appearance in Crete). After that it continues increasing less rapidly until it reaches an equilibrium during the 1687-1615 BC phase. Table 3 also shows the number of sites with presence of wheel-made pottery for each phase

Site	Region	References	Information dated based on Table 1 absolute chronological framework
Akrotiri	Aegean Islands (Thera)	Hilditch (2008)	1725-1687
Aphrodisias	Western Anatolia	Joukowsky & Benjamin (1986)	2475-2217
Asine	Mainland Greece	Spencer (2007)	2475-2217
Ayia Irini	Aegean Islands (Kea)	Gorogianni, Abell & Hilditch (2016)	1912-1863
Ayios Stephanos	Mainland Greece	Spencer (2007)	2075-1912
Cesme-Baglaraci	Western Anatolia	Şahoğlu et al. (2018)	2475-2217
Deriziotis Aloni	Mainland Greece	Stocker (2003)	Absence of wheel-made pottery 2217-2075
Emporio	Western Anatolia (Chios)	Hood et al. (1981) Hood et al. (1982)	2475-2217

Eretria	Mainland Greece	Müller Celka et al. (2018)	2475-2217
Eutresis	Mainland Greece	Goldman (1931)	2475-2217
Helike	Mainland Greece	Cline & Forsén (2012)	2217-2075
Heraion	Western Anatolia (Samos)	Menelaou, Kouka, & Day (2016)	2217-2075
Kastri (Kythera)	Aegean Islands	Broodbank & Kiriatzi (2007) Kiriatzi (2010)	1912-1863
Knossos	Crete	Berg (2009) Jeffra (2013)	1912-1863
Kolonna	Aegean Islands (Aigina)	Gauss & Kiriatzi (2011)	1687-1615
Kommos	Crete	Branigan (1991)	1912-1863
Koukounaries	Aegean Islands (Paros)	Katsarou & Schilardi (2004)	Absence of wheel-made pottery on 2650-2475
Lefkandi	Mainland Greece	Spencer (2007)	2475-2217
Lerna	Mainland Greece	Choleva (2012)	2217-2075
Liman Tepe	Western Anatolia	Şahoğlu (2005) Türkteki (2014)	2475-2217
Markiani	Aegean Islands (Amorgos)	Marangou et al. (2008)	Absence of wheel-made pottery on 2650-2475 and 2475-2217
Menelaion	Mainland Greece	Catling et al. (2009)	2075-1912
Mitrou	Mainland Greece	Hale (2016) Kramer-Hajos & O'Neill (2008) Van de Moortel, Zachou & Rutter (2018)	2217-2075
Mochlos	Crete	Doudalis (2018)	1912-1863
Myrtos	Crete	Cadogan (2011)	1912-1863

		Jeffra (2013)	
Myrtos-Pyrgos	Crete	Jeffra (2013)	1912-1863
Nichoria	Mainland Greece	McDonald (1972) Spencer (2007)	2075-1912
Orchomenos	Mainland Greece	Spencer (2007)	2475-2217
Palaikastro	Crete	Jeffra (2013)	1863-1725
Palamari	Aegean Islands	Choleva (2015)	2475-2217
Panormos	Aegean Islands (Naxos)	Angelopoulou (2008)	Absence of wheel-made pottery on 2650-2475 and 2475-2217
Pefkakia	Mainland Greece	Maran (2007)	2475-2217
Phaistos	Crete	Caloi (2020)	1912-1863
Phylakopi	Aegean Islands (Melos)	Berg, I. (2004)	1725-1687
Poliochni	Aegean Islands	Brea (1956) Cultraro (2005)	2075-1912
Proskynas	Mainland Greece	Zachou (2004)	Absence of wheel-made pottery on 2650-2475 2475-2217
Psathi	Crete	Hood (1965) Nodarou (2003)	
Serraglio	Aegean Islands (Kos)	Vitale (2016)	2475-2217
Sitagroi	Mainland Greece	Renfrew, Gimbutas, & Elster (1986) Renfrew et al. (2003)	Absence of wheel-made pottery on 2650-2475 2475-2217 2217-2075
Sklavouna	Aegean Islands (Paros)	Katsarou & Schilardi (2004)	Absence of wheel-made pottery on 2650-2475

Thebes	Mainland Greece	Dakouri-Hild (2001)	2475-2217
Thorikos	Mainland Greece	Choleva (2012)	2475-2217
Toumba-Thessaloniki	Mainland Greece	Andreou & Psaraki (2007)	1323-1195
Trianda	Western Anatolia (Rhodes)	Marketou (2012)	1912-1863
Troy	Western Anatolia	Blegen, Caskey & Rawson (1951) Choleva(2020)	2475-2217
Yortan	Western Anatolia	Kiamil (1980)	2475-2217
Zakros	Crete	Evely (1988)	1863-1725

Table 2: Sources for the empirical data on the presence and absence of wheel-made pottery in each site of the Aegean. This data was then used to calculate the proportion of sites with presence of wheel-made pottery over time (shown in Figure 2).

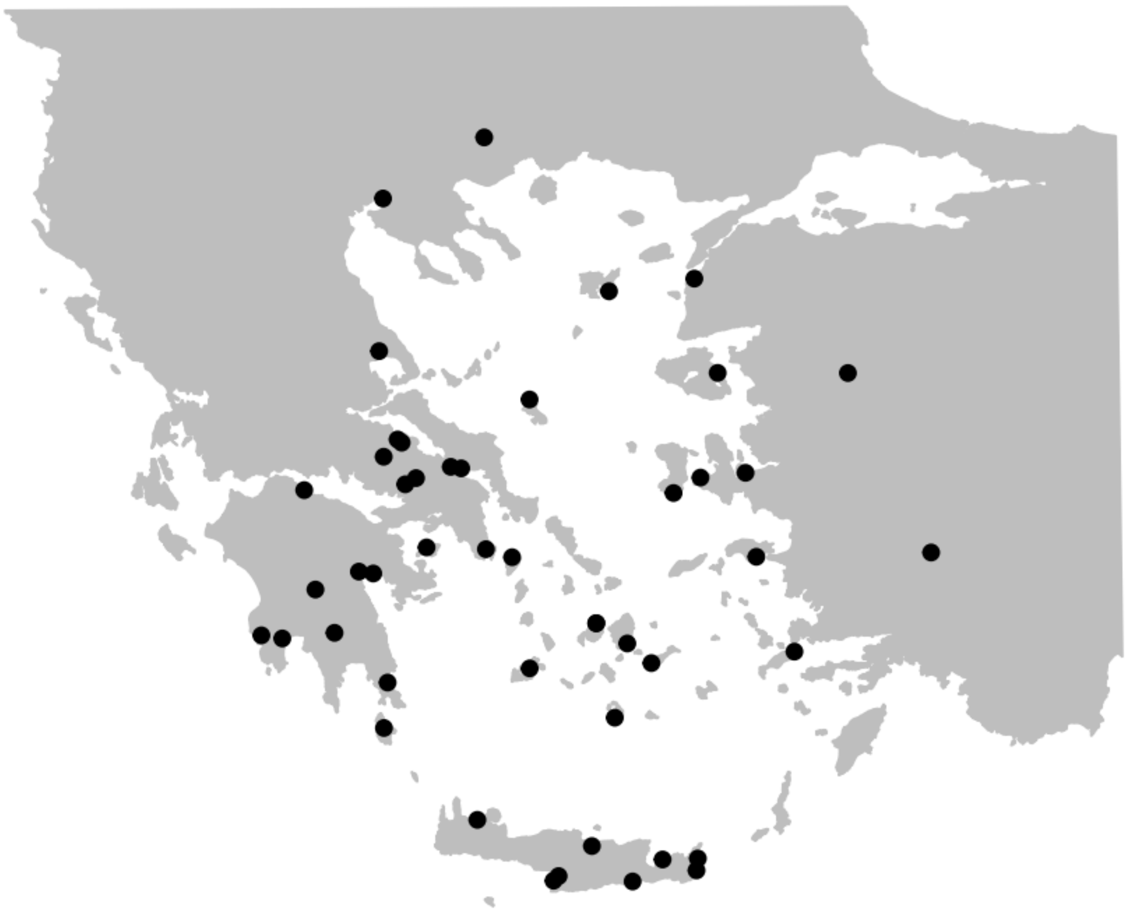

Figure 5: Map showing the locations of the sites with information on the presence/absence of wheel-made pottery. The dataset includes 10 sites from the Aegean islands, 8 from Crete, 20 from Mainland Greece, and 9 sites from Western Anatolia.

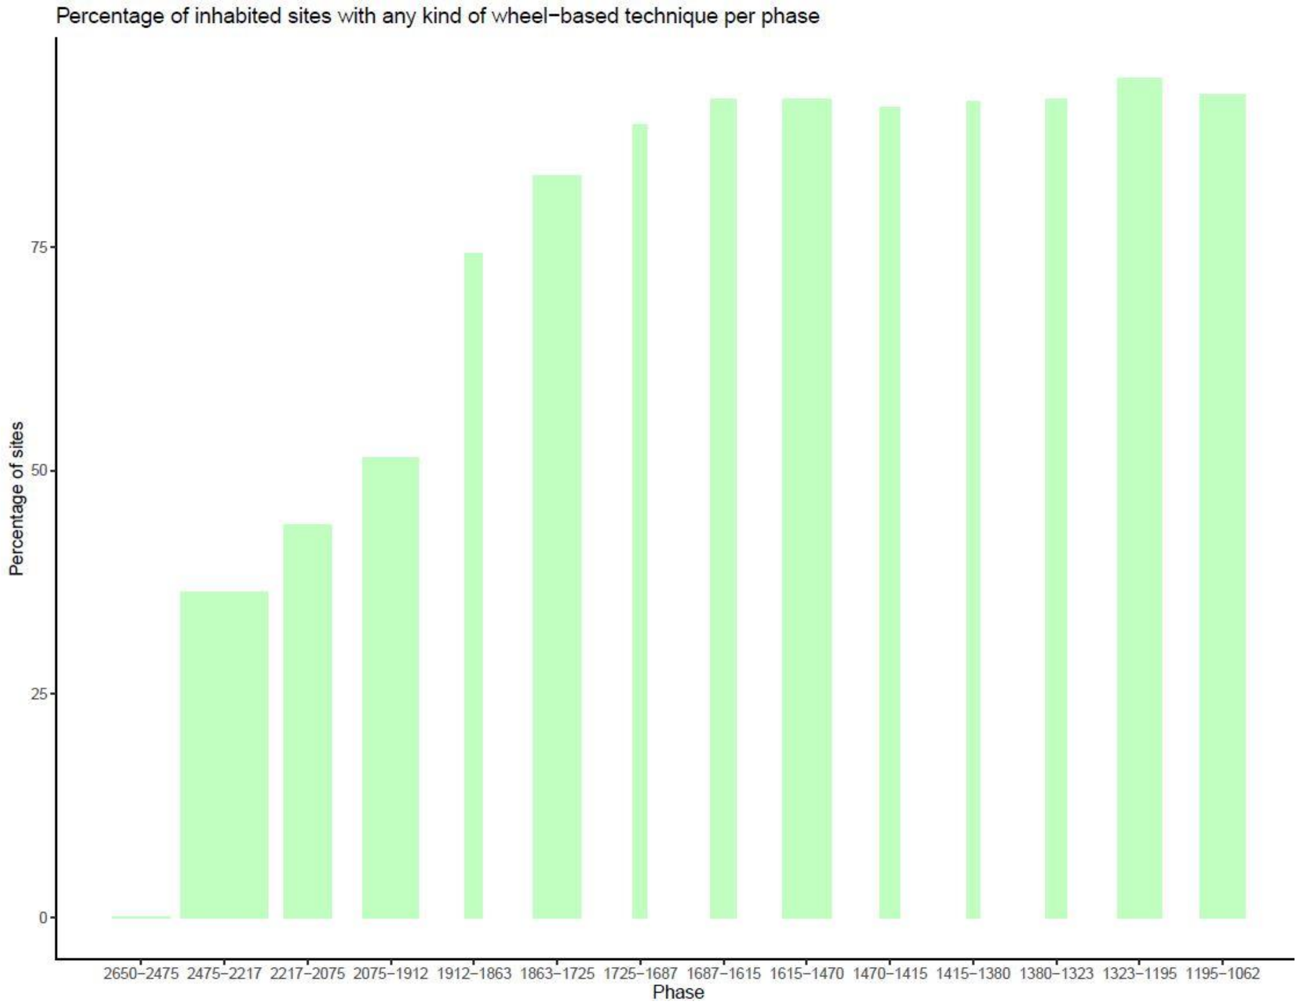

Figure 6: Histogram showing the percentage of inhabited sites in the dataset with presence of wheel-made pottery over time. The width of each bin represents the length of each phase in calendar years.

Phase	Crete	Mainland Greece	Western Anatolia	Aegean Islands
2650-2475	0	0	0	0
2475-2217	0	10	9	1
2217-2075	0	13	11	1
2075-1912	0	13	12	2
1912-1863	5	16	13	2
1863-1725	7	19	13	2
1725-1687	7	21	13	4
1687-1615	7	22	13	5
1615-1470	7	22	13	5
1470-1415	5	21	13	4
1415-1380	6	21	13	5
1380-1323	7	22	13	5
1323-1195	6	22	14	3
1195-1062	5	19	10	4

Table 3:The number of sites with presence of wheel-made pottery per region for each chronological phase

5.2.2. Proportion data

The dataset collated from published sources to trace the mean proportions of wheel-made pottery over time is built from limited data sources and datasets with small sample sizes that were published using diverse reporting standards. Data aggregation — especially with regards to each data point's position in a temporal bin — required making certain decisions that possibly affected the reported trend. These two factors — small sample size and reporting diversity – sometimes act in a synergistic manner: the impact of an aggregating/binning decision for each data point is even more impactful when the sample is small, since summary statistics such as the mean produced from smaller samples can be sensitive to outliers (Agresti 2018: 137).

For example, the decision to assign Phase 7 (which corresponds to one of the MBA sub-phases) with a 52 % proportion of the trait based on Joukowsky's (1986) reported value of 52 % for the entire Middle Bronze Age possibly affects the mean for that phase more than it does the mean for another MBA subphase, Phase 6, which has also been assigned a 52 % value, due to the fact that the former has a smaller sample size than the latter.

The following tables (tables 4 and 5) provide information on each published data point containing information about the proportion of wheel-made pottery at different sites in the Aegean per phase.

Site	Source
Aphrodisias	Percentages reported by Joukowsky (1986). Where the original publication used wide assignments to larger bins, I have repeated the same value for each sub-phase. For example, the reported LBA percentage was 76 %. This value was reported for Phases 8-14 in the dataset.
Asine	Percentages reported by Spencer (2007).
Aspis	Percentages reported by Philippa-Touchais et al. (2020).
Ayia Irini	For Phase 2 (late EB II) Choleva's (2020: 71) estimate is based on percentage of Kastri wares. For Phases 4-7 (MBA), I used the estimate produced by the analysis of inventoried pots as Gorogianni et al. (2016:211) argue this is likely to be more accurate.
Emporio	For Phase 2, the reported value is based on 2 proxies: the first uses my estimate of 21 % based on Hood's (1982) catalogue, the second is 4 % based on Choleva (2020:71). The reported value is the mean value of the two estimates.
Heraion	For Phase 2, 4 out of 44 sherds (9.09090909 % which has been rounded to 10 %) dated to Heraion II-III were "wheel-made or wheel-finished according to macroscopic evidence (HT12/14, HT12/26, HT12/29, HT12/32)." (Menelaou et al. 2016).
Knossos	Percentages reported by Jeffra (2013: 38).
Lefkandi	Phase 2 corresponds to Spencer’s (2007) Phase 1: 8.1 % "wheel-fashioned and 8.7 % "wheel-finished". Phase 3 corresponds to Spencer’s (2007) Phase 2 and 2-3: 20-21 % wheel-fashioned. Phases 4 and 5 correspond to Spencer’s Phases 3 and 4: 46 % wheel-fashioned. Phase 6 corresponds to Spencer’s Phase 5: 50 % + wheel-fashioned. More specifically, for Phase 2, the 16.8 % reported is the sum of the reported 8.1 % "wheel-fashioned and 8.7 % "wheel-finished". The 20.5 reported for Phase 3 is the mean of Spencer's estimate range, which is 20 to 21 %.
Lerna	Percentages reported by Choleva (2012).
Malia	For Phase 5, I used Knappett’s catalogue for the Secteur Pi pottery, which represents pottery dated "MM IB or MM IIA" (Note the chronological uncertainty): 0.44354838709 based on the 120 out of 248 catalogued sherds marked as wheel-made (“wm” in catalogued entries). For Phase 6, the reported percentage was my quantification based on the Quartier Nu Catalogue.

Mitrou	The survey catalogue does not distinguish between different subphases for the MH, so I have reported the same trait proportion value for Phases 4-7. For the LH pottery (Phases 8-14) the catalogue distinguishes between subphases, so I have calculated a percentage for each subphase, disregarding sherds with a wide temporal assignment ("LH").
Myrtos-Pyrgos	Percentages reported by Jeffra (2013: 40).
Palaikastro	Percentages reported by Jeffra (2013: 39).
Palamari	Percentages reported by Choleva (2020: 71).
Pefkakia	Percentages reported by Choleva (2020: 71).
Phylakopi	Percentages reported by Berg (2007:244)
Thorikos	Percentages reported by Choleva (2020: 71).
Troy	Choleva's assessment (2020:71) of "10-20%" wheel-made pottery for "Troy IIb–early III" was assigned to Phase 1, using the mean of 10% and 20% to get a single value summary statistic.

Table 4: Sources for the quantitative data on the proportions of wheel-made pottery for each site in the Aegean that were used to produce the empirical trend curve used for this study.

Site	Phase 1 2650 - 2475 n = 12	Phase 2 2475-2217 n = 14	Phase 3 2217-2075 n = 9	Phase 4 2075-1912 n = 6	Phase 5 1912-1863 n = 9	Phase 6 1863-1725 n = 10	Phase 7 1725-1687 n = 7	Phase 8 1687-1615 n = 6	Phase 9 1615-1470 n = 3	Phase 10 1470-1415 n = 2	Phase 11 1415-1380 n = 2	Phase 12 1380-1323 n = 2	Phase 13 1323-1195 n = 2	Phase 14 1195-1062 n = 2	References
Aphrodisias	32	16	13	44	52	52	52	76	76	76	76	76	76	76	Joukowsky (1986)
Asine	0	4	4.9	6	9	7.9	NA	NA	NA	NA	NA	NA	NA	NA	Spencer (2007)
Aspis	NA	NA	NA	NA	20	8	NA	NA	NA	NA	NA	NA	NA	NA	Philippa-Touchais et al. 2020
Ayia Irini	NA	0.1	NA	NA	58	75	92	NA	NA	NA	NA	NA	NA	NA	Choleva (2020)
Emporio	NA	12.5	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Hood et al. (1982)
Heraion	0	0.1	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	MA	Menelaou, Kouka & Day (2016)
Knossos	0	0	0	NA	95	86.57	93.069	100	NA	NA	NA	NA	NA	NA	Jeffra (2013)
Lefkandi	0	16.8	20.5	46	46	50	50	NA	NA	NA	NA	NA	NA	NA	Spencer (2007)
Lerna	0	0	2.42	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Choleva (2012)
Malia	0	0	0	0	0.443548387	40	NA	NA	NA	NA	NA	NA	NA	NA	Schoep & Knappett (2003)
Mitrou	0	0	20	11.76	11.76	11.76	11.76	94.44	100	100	100	100	93	100	Kramer-Hajos & O'Neill (2008)
Myrtos-Pyrgos	0	0	0	0	100	94.642	NA	100	NA	NA	NA	NA	NA	NA	Jeffra (2013)
Palaikastro	0	NA	NA	NA	NA	100	93.75	100	NA	NA	NA	NA	NA	NA	Jeffra (2013)
Palamari	0	10	34	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Choleva (2020)
Pefkakia	NA	1.2	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Choleva (2020)
Phylakopi	NA	NA	NA	NA	NA	NA	4.9	63.3	69.1	NA	NA	NA	NA	NA	Berg (2007)
Thorikos	NA	0.3	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Choleva (2020)
Troy	15	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	Choleva (2020)

Table 5: percentage of wheel-made pottery per chronological phase.

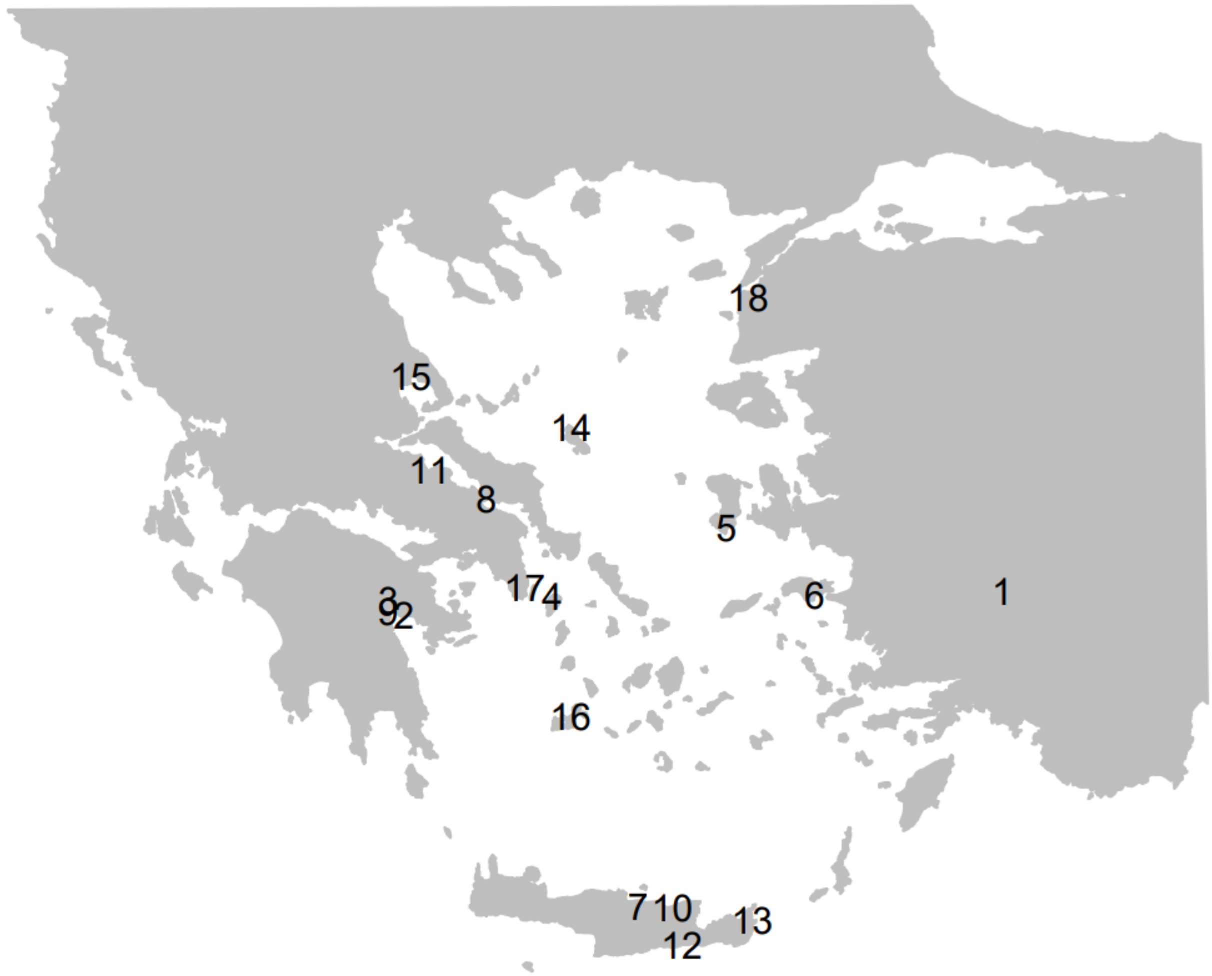

Figure 7: Map of locations where there is published information on the relative use of the wheel to make pottery (sites in Table 4). 1. Aphrodisias, 2. Asine, 3. Aspis, 4. Ayia Irini, 5. Emporio, 6. Heraion, 7. Knossos, 8. Lefkandi, 9. Lerna, 10. Malia, 11. Mitrou, 12. Mytos-Pyrgos, 13. Palaikastro, 14. Palamari, 15. Pefkakia, 16. Phylakopi, 17. Thorikos, 18. Troy.

As can be seen in Tables 3 and 4, the available empirical evidence is patchy -in Table 4, note the large number of cells with NA (not available) values. Therefore, any attempt to produce a trend using a summary statistic such as the mean should be made with a strong awareness of the uncertainty inherent in the data. In an attempt to communicate this uncertainty to the reader, and as an analytical tool that will later be used to evaluate the simulations' fit to the empirical data, I have calculated confidence intervals for the mean percentage of wheel-made pottery over time using the method suggested by Shennan (1997: 83).

More specifically, to calculate 95 % confidence intervals for the population mean using the sample mean for each one of the 14 archaeological phases, I used a t value drawn from Student's t distribution rather than the frequently used Z score drawn from a normal distribution in the following formula, since the sample size n for each archaeological phase was smaller than 40 (Shennan 1997: 82-83):

$$\bar{x} \pm t_{a,d.f.} \frac{s}{\sqrt{n}}$$

Where $\bar{x}$ is the sample mean, $t_{a,d.f.}$ is the appropriate value drawn from Student's t distribution for the degrees of freedom (n-1, where n is the sample size) of the sample and a desired probability – in this case .95, and s is the standard deviation of the sample.

The left panel of Figure 8 shows the empirical trend with a 95 % confidence envelope, where it is evident that the small sample sizes for phases 9-14 result in extreme uncertainty over the value of the 'true' mean. However, given that the literature suggests that adoption of the wheel was very well established after Phase 9, and the fact that by that time the empirical trend appears to oscillate around an equilibrium, it is safe to assume that the empirical trend can be used for some sort of quantitative comparison to the simulated trends. In fact, the reason why so few researchers publish wheel proportion data for these phases is that by that time handmade pottery is found in such small quantities compared to the abundant wheel-made pottery (linked to Mycenaean-style production) that the discipline as whole has tended to assume such quantification does not merit much attention – but note cases where research projects focused on making that distinction for the Late Bronze Age, such as work by Kiriatzi et al. (1997), attempting to throw light on the persistence of handmade production in Northern Greece. This is an interesting case where disciplinary certainty in well-established qualitative assessments has resulted in quantitative trends with regions of high uncertainty.

Another major characteristic of the empirical trend is that it is the result of archaeological periodisation. In other words, trait proportions are not reported in regular, equal intervals, but instead they correspond to pottery phases whose length is unequal. Therefore, while the proportion of wheel-made pottery when plotted as a value on the y axis remains the same, the slope of the trend from each phase to the next changes when one takes into consideration the length of each phase, as shown in the right panel of Figure 8, where Time shown on the x axis consists of points corresponding to the exact calendar year when each phase ended, as indicated by absolute chronological schemes such as Manning (2010).

It should be noted that I am combining data on wheel-made pottery from the mainland and from Crete to produce the empirical trends in Figure 3. However, research suggests that the transmission of the technique in the mainland and in Crete should be considered independent processes (Knappett 2016) given that the evidence suggests that wheel-made pottery was introduced Crete in the Middle Minoan period. Given that the goal of the present study is to study the transmission process at the scale macro region including both the mainland and Crete as a single unit, I will use a single curve combining empirical data from both regions. This does not mean that I consider the transmission process to have propagated from the mainland to Crete (and the simulation models explored later do not impose a directed propagation of the trait from one region to its neighbours except from using Western Anatolia as an early origin). There is not enough empirical information to definitively confirm whether RKE techniques in Crete originated from the mainland or from some other origin in Western Anatolia (or from both regions). I have combined evidence from different regional typological phases to produce temporal bins within a single absolute chronological framework to map the adoption of wheel-made pottery over time at the scale of the macro region that includes mainland Greece, western Anatolia, the Aegean islands, and Crete. However, this does not mean that I consider each measurement/event to have taken place at exactly the same moment within each temporal bin.

As discussed at greater length in chapter 3, simulations are not perfect replications of observed phenomena, but simply abstractions that capture some elements/details of the phenomenon in order to explain aspects of interest relative to the research question. As the research questions in the study are concerned with the transmission of cultural traits in the entire macro-region that includes Crete, I have designed simulation models that do not model the transmission in each region individually, but instead model cultural transmission throughout the entire study area. Using a single empirical curve that is the result of aggregating data from sites in multiple regions allows me to compare the results of a simulation that models transmission within a single macro-region to the empirical record. However, it should not be assumed that the outputs of the simulation perfectly replicate the empirical record - they only do so to the extent that we can explain some patterns of interest at the spatial and temporal scale of interest.

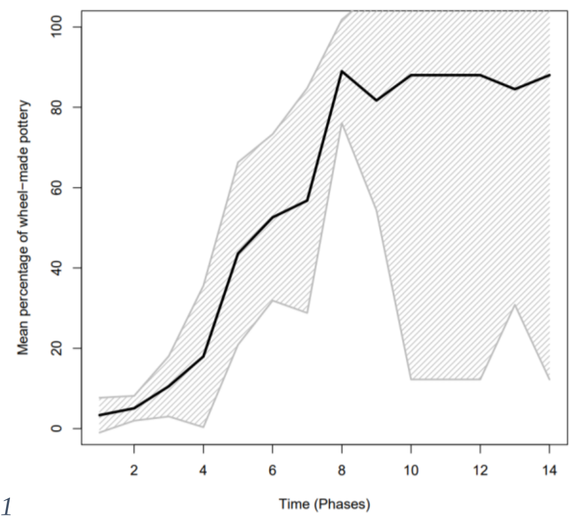

8.1

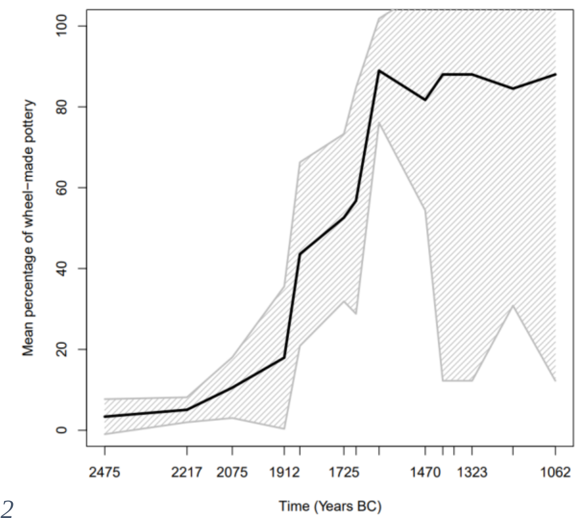

8.2

Figure 8: Different ways of calculating a 95 % confidence interval for the temporal trend produced by taking the mean of each phase using the data on Table 4.

Slope is important in an important aspect of the empirical trend, because we are interested in how quickly trait proportions change over time — the literature suggests that there was a resistance/hesitation to adopting wheel-made pottery at first, and that there was a period of fast uptake that corresponded to the rise of the first palaces. This is reflected in the new trend (Figure 8.2), where the initial phase of slow adoption is relatively longer than the one for the trend that was produced without taking phase length into consideration (Figure 8.1), resulting in a curve whose shape is more consistent with the S-shaped curve typically associated with innovation adoption (see Rogers 2003 for the theory behind the adoption processes that generate S-shaped curves). This procedure also confirms that the region of high uncertainty (1470 to 1060 BC) corresponds to a shorter time period than what was suggested by Figure 8.1, and that allows one to place greater trust in the sample mean over the entire study period.

During the 2475-2217 BC interval, uncertainty is relatively low due to the larger sample size (compared to other periods) produced by including statistics from Choleva's studies on the initial slow adoption of the wheel in different parts of the Aegean (2012, 2020). Therefore, one can securely assume trait adoption to be lower than 10 % in a prolonged period (roughly 200 years) of slow adoption. The sample mean suggests that the next phase saw a sharp increase in trait adoption, but the width of the confidence interval, especially on the lower bound, suggests that this may be due to uncertainty. A period of slower increase follows that, but wheel-made pottery truly takes off after the first palace period (1912 BC) and reaches an equilibrium at around 1600 BC. Even if the "true" mean is somewhere near the lower bound of the confidence interval, it is still safe to assume that trait adoption was fastest between 2217 and 1600 BC.

Having established a confidence envelope around the empirical trend provides an opportunity to perform quantitative comparisons between the simulations and the observed pattern.

A Monte Carlo Simulation approach to temporal and sample size uncertainty

An alternative method to address the presence of uncertainty in multiple aspects of the empirical mean is to use Monte Carlo simulation, a set of random sampling methods that allow one to randomly draw from a known distribution or interval within which one knows contains the “true” value of interest. The strength of the approach lies in numbers — randomly sample one time and the chances of drawing a value near the “true” value are low, but randomly sample thousands of times and there is a high chance that the “true” value was captured in one of the simulations. The term simulation used here should not be confused with the use of the term “simulation” used throughout the text referring to generative simulation modelling. The term Monte Carlo simulation used here refers to the simple act of repeatedly sampling from a known distribution.

As aforementioned, there are two types of uncertainty that affect the empirical mean. First, the mean for each phase was estimated using a very small number of sites, a fact that was addressed by calculating one 95 % confidence interval for each phase. Second, the temporal assignment of the mean for each phase should be treated with some suspicion: while researchers report a trait proportion value for each archaeological phase, it is not clear whether this value is representative of a uniform aggregation of the trait proportions for different sub-phases in each phase, or whether it is more representative of the trait proportions of early/later sub-phases within the phase. The approach taken to address this was to assume that each reported value was drawn from the “end” of each phase, but this is a big assumption to make when the periodization processes that affect empirical measurements in the Bronze Age Aegean are so unknown to us.

Our understanding of sample size and temporal uncertainty can be enhanced by introducing statistical methods such as Monte Carlo simulation to the analysis of the empirical data. Regarding sample size uncertainty, the calculation of a single confidence interval can be supplemented by calculating multiple confidence intervals for each reported value — in this example I have calculated 1000 confidence intervals for each phase — and then randomly drawing a trait proportion value from each confidence interval using a uniform probability distribution to produce 1000 trend lines over time. Having used a 95 % confidence interval means that one should expect 950 out of 1000 intervals to contain the "true" value. Figure 9.1 shows 1000 trend lines produced by randomly drawing a value from a confidence interval for each phase. As in the previous example (Figure 8.2), I used the year each phase ends as the time the measurement refers to in order to observe the effects of using Monte Carlo simulation on one aspect of uncertainty in isolation – this example does not take temporal uncertainty into consideration.

The second problematic factor is temporal uncertainty about the exact moment (year) within each phase that the measurement (in this case trait proportions) refers to. Bevan and Crema (2021: 2) refer to this problem as "within-phase uncertainty" and (among other tactical simulation approaches) demonstrate the application of Monte Carlo simulation to construct "envelopes" that can illustrate the effects of "within-phase uncertainty" per phase. The experiment in Figure 9.3 draws inspiration from Bevan and Crema's (2021) approach: for each reported phase-mean, a date (year) is sampled at random from a uniform probability distribution within the interval that is defined by each phase's start and end date. As in the previous experiment, I have not taken sample size uncertainty into consideration in order to observe the effects of the Monte Carlo approach on the temporal dimension in isolation.

Finally, it is possible to produce Monte Carlo simulations that capture the effects of both within-phase uncertainty and sample size uncertainty. Figure 9.2 shows the envelopes resulting from such an experiment. Overall, it appears as though all types of uncertainty assessment methods establish that the empirical trend is a logistic curve (S-shaped) that reaches an equilibrium value sometime after 1800 BC. Note that the time after 1800 BC is one of higher uncertainty than previous periods when taking sample size uncertainty into account, as illustrated by the transparency of the envelope. It is uncertain how informative the Monte Carlo estimate is for later periods given that it is well-established in the literature that wheel-made pottery is found in extremely high proportions during that time.

The envelope produced by incorporating within-phase uncertainty produces a double S-shaped curve. The combined uncertainty type simulations (Figure 9.2) do not rule out the possibility of a double S-curve either, but it appears just as likely that the “real” trend was a single S-curve.

Overall, the Monte Carlo simulations that incorporate within-phase uncertainty appear to be the ones most straightforwardly compatible with direct comparison to the outputs of the equation-based models presented here. The results of quantitatively comparing the two are presented in Chapter 8.

The temporal-uncertainty Monte Carlo approach can be taken in multiple future directions: here I have presented a case where one is willing to admit complete ignorance about the probability of the year/moment that each event (in this case, a measurement of wheel-made proportions for a pottery assemblage) took place within each phase-interval — the use of a uniform probability distribution reflects this. In other cases, however, it is possible that we have some domain knowledge that allows making the assumption that a trait was more likely to appear at the beginning or the end of a phase. Regarding wheel-made pottery, in particular, the date when it appears in the Aegean is qualitatively referred to as “the end of EBA II” (Choleva 2012: 375). Incorporating this information can be as simple as choosing a different probability distribution to sample from, for example a left-skewed distribution with a peak towards the end of the interval.

9.1.

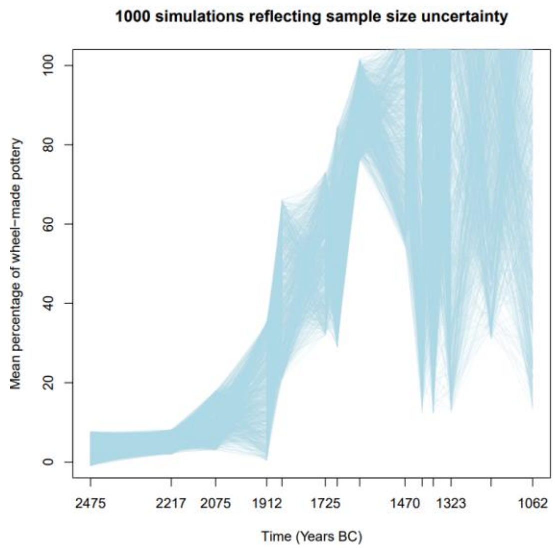

9.2.

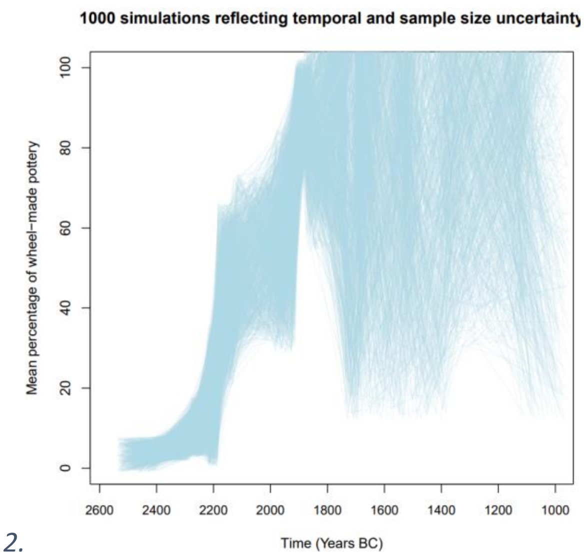

9.3.

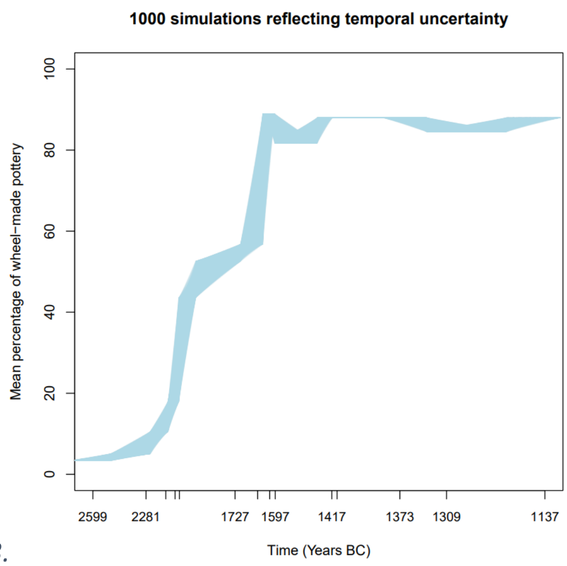

Figure 9: Three approaches to modelling uncertainty in the transmission trends over time for the potter's wheel dataset (proportion data shown in Table 4) using Monte Carlo simulations. The approach shown in 9.1 involves running 1000 simulations where the value on the y axis is randomly drawn value from the confidence interval calculated for each phase using the method discussed above. Since the approach does not consider temporal uncertainty, I use the end of each phase (using the chronological framework in Table 1 for the value on the x axis representing time (the approach considers uncertainty associated with the size of the sample). The approach in 9.2 involves using the same approach for the value on the y axis as used in 9.1, but I also randomly sample a date from the range of dates defining each phase in the chronological framework for the x-axis value representing time (this considers uncertainty associated with the size of the sample and temporal uncertainty). Finally, the approach in 9.3 involves only using randomly drawn values for the x-axis value (as in 9.2), but for each of the 1000 simulations I use the mean percentage of wheel-made pottery calculated for each phase (this approach only considers temporal uncertainty).

5.3. The kiln dataset

Hasaki's database (2002) includes 79 kilns dated to the Bronze Age located in Mainland Greece, the islands in the Aegean and the Ionian Sea, and Crete, which are included in the final dataset as seen in Table 1. As discussed previously, the date ranges for each kiln differ, with some kilns being associated with a very wide time interval, for example the Middle Helladic Period (e.g., the kiln at Marathon), to others dated to narrower intervals (e.g., one kiln in Priniatikos Pirgos which is dated to the approximately 70 year-long LMIA period). I excluded kilns which were not assigned a specific time interval, as well as one contraption in Siopata, Gavdos dated to the Final Neolithic-Early Minoan period, whose definition as a "kiln" is considered uncertain (Kopaka & Theou 2018: 454). Overall, 73 out of Hasaki's 79 Bronze Age kilns were retained in the final dataset. Kilns from Western Anatolia from Massa (2016: 152) were also added in the dataset to ensure more even coverage of the phenomenon over the study area (all three are dated to the Early Bronze Age). Massa (2016: 152) mentions 7 kilns from Kumyer Mevkii (Tırpan and Gider 2011:386-387), one kiln from Karataş (Warner 1994:187), and "several" kilns from Seyitömer Höyük (Bilgen et al.2013:205) — for consistency, wherever multiple kilns are reported without a specific number, three kilns were included in the dataset for the same location and date range.

The final dataset consists of the 84 kilns shown on the map in Figure 10. There is a considerably larger number of kilns in Crete compared to other regions — whether this reflects greater popularity of kilns in Crete or is simply an artifact of a research bias in Minoan archaeology compared to archaeology in mainland Greece is hard to tell, but it is not unreasonable to suggest that there is some connection between the palatial economy and more standardised pottery production requiring the use of a kiln. It should be noted that Figure 5 shows all the kiln locations in space as if they were in existence/use at the same time, but the Bronze Age is a very long period of time, and it is important to keep in mind that the chronological uncertainty inherent in the date of each kiln (kilns are dated using long time intervals) makes it hard to establish what the relative popularity of kilns was over time and across space.

However, in recent years, archaeologists have adopted methods to incorporate certain types of temporal uncertainty into analyses of presence/absence data. Aoristic analysis (Ratcliffe 2000) is a popular approach, which was adopted by Johnson (2004) and, more recently, by Crema (2010, 2012). Crema's (2012) implementation was particularly influential and inspired other archaeological applications (e.g., Yubero-Gómez et al. 2016, Orton et al. 2017). In the following section, I will briefly explain the main premise and how aoristic analysis was used to reconstruct the trend of the relative adoption of kilns over time and in space, using functionality from the archSeries package (Orton, 2017) for the R statistical environment.

Site	Date	Long	Lat
Polychrono	EH	23.53	40.02
Marathon-1	MH	23.96	38.16
Lerna-1A	MH	22.71	37.53
Lerna-1B	MH	22.71	37.53
Lerna-1C	MH	22.71	37.53
Lerna-1D	MH	22.71	37.53
Mycenae-1	MH	22.76	37.73
Sparta-3A	MH	22.45	37.06
Sparta-3B	MH	22.45	37.06
Eretria-3	MH	23.8	38.39
Kirra-1A	MH	22.44	38.43
Kirra-1B	MH	22.44	38.43
Kirra-1C	MH	22.44	38.43
Zarkos	MH	22.12	39.61
Aigeira-1	LH	22.36	38.15
Asine A	LH III	22.99	37.52
Asine B	LH III	22.99	37.52
Berbati-2	LH	22.8	37.15
Tiryns A	LH III	22.8	37.6
Tiryns B	LH III	22.8	37.6
Pylos	LH IIIB	21.69	37.03
Thebes-1	LH III	23.32	38.32
Dimini	LH III	22.89	39.36
Pherae-4	LH IIIC	22.73	39.34
Naxos A	LH	25.38	37.11
Naxos B	LH	25.38	37.11
Cos A	MM	27.29	36.89
Cos B	LM IA	27.29	36.89
Phaistos-2	MM IIB	24.81	35.05
Zakros-1	MM IIIA	26.22	35.11
Vathypetro	MM IIIA-LM IA	25.16	35.21
Zou	MM IIIB	26.11	35.16
Stylos	LM IIIB	24.12	35.43
Kato Gouves A	LMIIB	25.31	35.33
Kato Gouves B	LMIIB	25.31	35.33
Kato Gouves C	LMIIB	25.31	35.33
Kato Gouves D	LMIIB	25.31	35.33
Kato Gouves E	LMIIB	25.31	35.33
Kato Gouves F	LMIIB	25.31	35.33
Kato Gouves G	LMIIB	25.31	35.33

Kato Gouves H	LMIIB	25.31	35.33
Kato Gouves J	LMIIB	25.31	35.33
Kato Gouves I	LMIIB	25.31	35.33
Kato Gouves K	LMIIB	25.31	35.33
Knossos-3	LM II	25.16	35.3
Knossos-4A	LM IA-IB	25.16	35.3
Knossos-4B	LM IA-IB	25.16	35.3
Knossos-4C	LM IA-IB	25.16	35.3
Phaistos-3	LMIB-LMIIIC	24.81	35.05
Agia Triadha		24.79	35.06
Metropoli, Gortyna		24.96	35.06
Kommos	LM I A-IB	24.76	35.01
Achladia		26.05	35.17
Zakros-2		26.22	35.11
Mochlos A	LM IB	25.91	35.18
Mochlos B	LM IB	25.91	35.18
Palekastro		26.25	35.2
Kavousi	LMIIIC	25.86	35.12
Armeni	LM III	24.46	35.3
Zakros-2A	LM IA	26.22	35.11
Zakros-2B	LM IA	26.22	35.11
Priniatikos Pyrgos-1A	EM II - MM I	25.73	35.13
Priniatikos Pyrgos-1B	LM IA	25.73	35.13
Mochlos - Chalinomouri	LM IB	25.94	35.18
Makrygianni plot-1B	MH II	23.73	37.97
Makrygianni plot-1B	MH II	23.73	37.97
Raches Mosholouriou	LH I-II	22.09	39.35
Kakeika-1	MH-LH	21.64	38.37
kakeika-2	MH?	21.64	38.37
Mallia-Quarter Theta	MM IB	25.49	35.29
Ialysos- Margariti plot	LH IIB - LH IIIA2	28.16	36.41
Pigadia, Tsekou-Trempela plot-1A	LH III	27.21	35.51
pigadia, Tsekou-Trempela plot-1B	LH III	27.21	35.51
Khamalevri- Pateras plot	LM IIIC	24.6	35.37
Chalasmenos	LM IIIC	25.83	35.08
Zominthos	LM I	24.89	35.25
Petroto	MH?	21.43	39.32
Sparti-Stavropoulou plot	EH II	22.42	37.07
Siopata	FN-EM	24.1	34.84

Seyitömer Höyük	EBA
Kumyer Mevkii	EBA
Karataş	EBA

Table 6: The kiln dataset.

Figure 10: The locations of the 84 Bronze Age kilns included in the kiln dataset. The relative opacity of each node relates to the number of kilns in that location: more transparent nodes represent locations with a smaller number of nodes compared to more opaque nodes.

Aoristic analysis of kilns to reconstruct the trend of kiln adoption over time

At first glance, Table 5 does not suggest any problems with the date assignment for each kiln: there is a relative date associated with each kiln, and indeed some kilns are dated to a narrow time frame that beats even the finest resolution of some Bronze Age pottery. However, not every kiln is dated to as narrow a time interval; some kilns are dated within 70 year-periods, while other kilns are dated within 200-year periods. Therefore, it is safe to assume that for kilns dated to narrower time intervals there is less uncertainty as to when they were in use, while for kilns dated to very wide time intervals, there is greater uncertainty. Furthermore, the width of the time interval is not necessarily linked to how long we think a kiln was being used for — and, at least as I am aware, there is no ethnographic evidence suggesting how long a kiln is used for in general. One could go into the reasons behind the uncertainty of each date, which will, no doubt, differ per each individual case, but this would mean that they would have to come up with a way to quantify the uncertainty for each individual case, and even if all the information was there (unlikely), this exercise would become more challenging as the sample size got larger.

Instead, it is preferrable to adopt an approach that treats each data point consistently (in this case, a data point is an entry/row in the dataset in Table 5). We will assume that, making the least number of assumptions, it is equally likely that at a kiln was used at any time within the interval that the kiln was dated in. This means that, if the probability of a kiln existing between 2500 and 2000 BC was equal to 1 (we definitely know that a kiln existed in that time period), the shape of the probability density function would resemble a uniform distribution — in other words, an equal probability of existence at each moment within that time period.

If we treat each kiln as an event whose occurrence represents a probability equal to 1 (100 %) and its absence a probability of 0 we can split this probability into bins that represent the probability that the event occurred during the time period represented by the bin. Thus, in the previous example, if we know a kiln was present at a site sometime during the 2500-2000 BC interval and we know that kilns are typically used for 50 years, it would be possible to split the interval into ten 50-year bins ,which would be assigned a 1/10 = 0.1 aoristic weight of kiln existence (since 1 represents the total probability and we are dividing it equally into 10 bins).

If we do the same thing for each kiln site, it is possible to get the sum of those aoristic weights for each time bin and get something like the plots in Figures 11.1-11.4. Each of the plots uses a different assumption regarding the length of use for each kiln — 11.1 is the most unrealistic and assumes that kilns are only in use for a year. The trend over time varies slightly depending on the length of the time bin used, but overall kiln presence tends to increase over time in a non-linear fashion, with punctuated episodes of very high kiln presence. High kiln presence starts occurring at around the time of the first palatial systems in Crete (around 2000-1900 BC).

Using the aoristic weight for each site over time also allows examining the relative intensity of kiln use in space over time. Figures 12.1-12.9 show the system at 9 different snapshots in time and suggest that early during the Bronze Age kiln use was more intense in Western Anatolia (12.2), while later kilns became prevalent in Mainland Greece (12.3) and later in Crete (12.4, 12.5) and in regions associated with the influence of the Minoan and Mycenaean palatial systems (12.6, 12.7, 12.8, 12.9). It should be noted that the relative lack of Anatolian kilns is probably due to a bias in data collection, since there was a far smaller number of published kilns in the English literature in Western Anatolia compared to all other regions in the dataset — publications for Western

Anatolian kilns are likely to be written in Turkish and therefore not included in databases like Hasaki (2002).

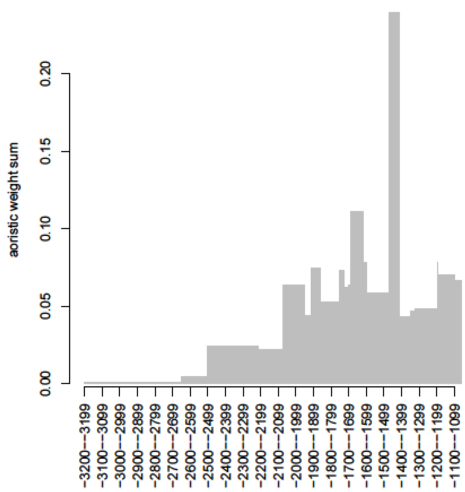

11.1.

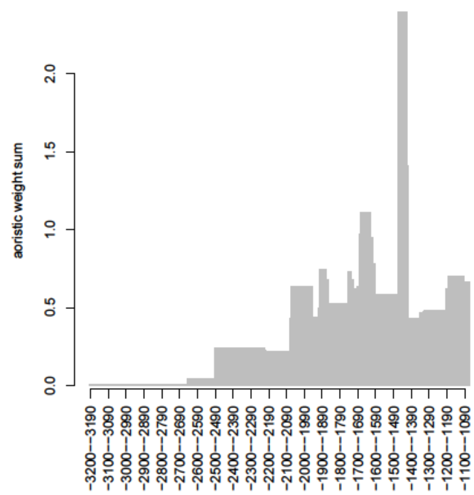

11.2.

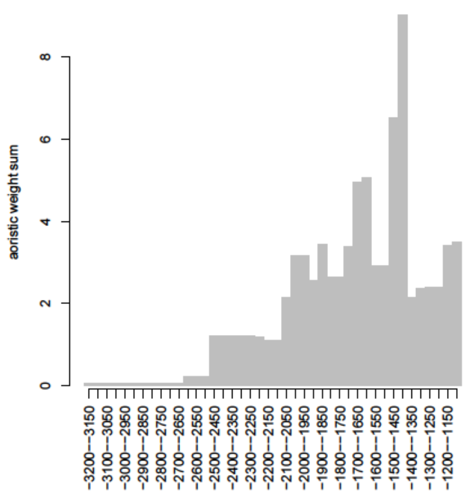

11.3.

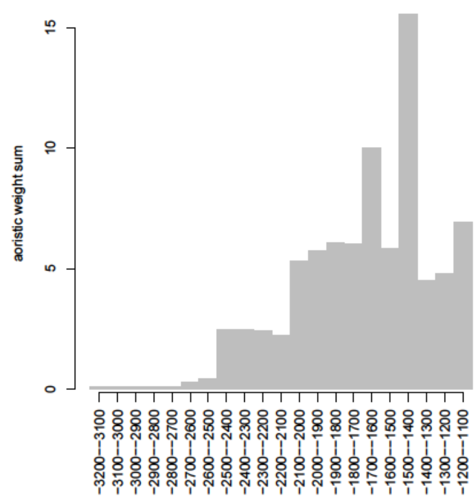

11.4.

Figure 11: The sum of aoristic weights of all kiln locations in the dataset over time calculated using the aorist function from the archSeries package (Orton 2017). Note that the value on the y axis is not to be interpreted as a rate, a probability, or a ratio. It is simply the sum of aoristic weights, which does not have a straightforward interpretation, but can be used as a relative measure of kiln presence probability over time. Each panel shows the sum of aoristic weights calculated using a different bandwidth: 11.1 uses a 1-year interval, 11.2 uses a 10-year interval, 11.3 uses a 50-year interval, and 11.4 uses a 100-year interval.

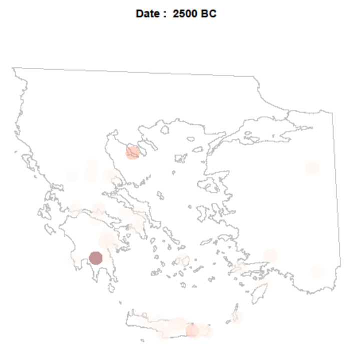

12.1.

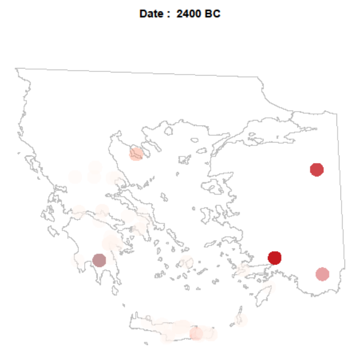

12.2.

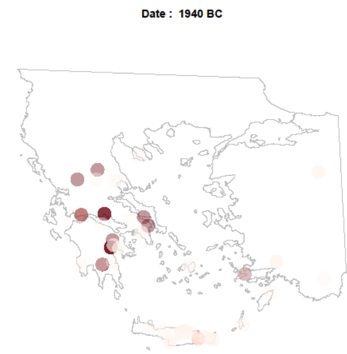

12.3.

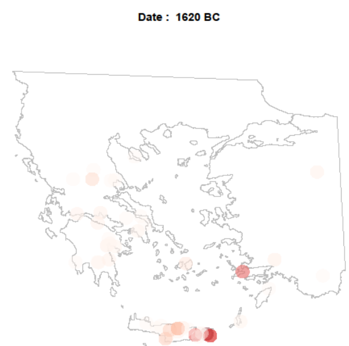

12.4.

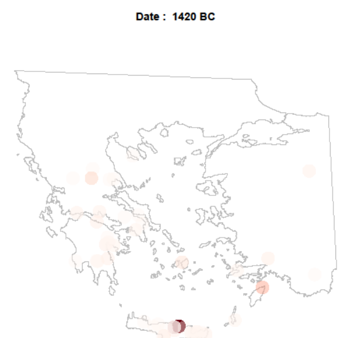

12.5.

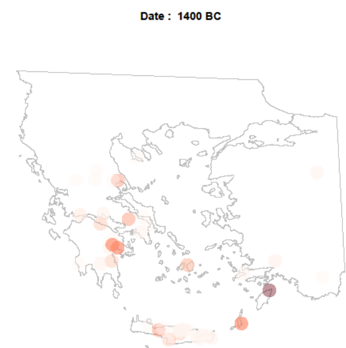

12.6.

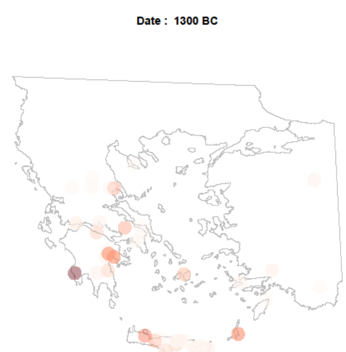

12.7.

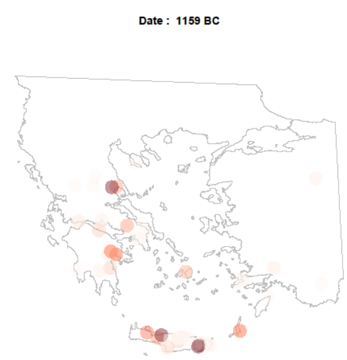

12.8.

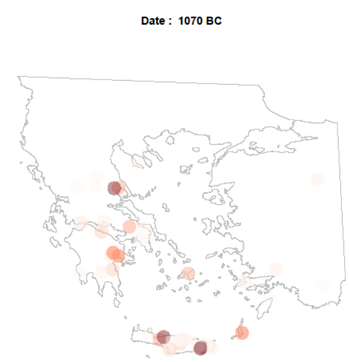

12.9.

Figure 12: Maps showing the aoristic weight of each kiln location at 9 different snapshots during the Bronze Age. The colour of each node represents the aoristic weight on a white to dark red scale – darker nodes have higher aoristic weights.

5.4. Empirical data on population levels for the Aegean Bronze Age

As discussed in chapter 2, theoretical modelling studies point to demographic patterns and population size as factors that greatly affect innovation transmission. Therefore, to assess whether demography affected the transmission of pottery-making technologies in the Aegean Bronze Age, it is important to collect empirical information on demographic patterns for the entire study area during that time.

A useful distinction concerning population size and demographic patterns is the scale at which we measure population. Empirically, population size could be captured by proxies for the number of individuals living in an area, but it could also be reflected by the number of communities in the study area (typically associated with the unit of the 'site' in archaeological investigations). Given the patchiness of the archaeological record and the effects of time-averaging, it is more likely that empirically one would observe population as the number of communities/sites than as the number of individuals. The assumption that a site equals a settlement/community is fairly reasonable, but what can it tell us about the size of the population? As Drennan (2017: 20) notes, the relationship between the number of settlements and population can be affected by "cultural, environmental, and economic" factors and using site numbers as a proxy for regional population size does not work well for "assessing in relative terms how many people lived in which locations or parts of a region."

In theory, a method that would achieve both a broad spatial and temporal coverage and an assessment of the number of communities over time for a given landscape is surface survey - there have been a number of surface survey projects that have produced a good picture of regional demographic patterns over long periods of time for certain parts of the Aegean such as Kythera

(Broodbank and Kiriatzi 2007: 260-261, who mention 24 sites dated to the FN/EB, 37 sites dated to the EB II period, and 20 dated to a ‘First Minoanising’ period), Antikythera (Bevan, Conolly, and Tsaravopoulos 2006, Bevan and Conolly 2013), and in Crete in the area around Knossos (Whitelaw, Vasilakis, and Bredaki 2018) and the Bay of Mirabello (Haggis 1996, Hayden 2003, Hayden 2004, Hayden 2005, Watrous et al. 2012). In practice, however, using this approach to get population size and distribution estimates for the entire Aegean and Western Anatolia region presents significant challenges, because it requires a process of secondary integration of data from multiple surveys. In the following subsections I review a number of studies that attempted to integrate such data into a single dataset or for comparative purposes.

5.4.1. Synthetic and comparative datasets compiled from regional surveys

Systematic extensive and intensive survey has been used consistently in Aegean archaeology since the 1970s (Alcock and Cherry 2004), and, as Stellatou (2007: 53) notes, methodologies have been changing, resulting in an inhomogeneous collection of surveys, often designed to address specific research questions and of small geographic and temporal scales. According to Alcock and Cherry (2004: 7), there is a tendency in Mediterranean archaeology to focus on “increasing spatial resolution” to the detriment of large-scale spatial and temporal coverage. This means that any attempt to identify larger scale patterns must be accompanied by synthesizing multiple datasets.

However, publication strategies can be problematic in terms of data re-use, as most published surveys do not include “raw data,” only the outputs of statistical analysis and subsequent interpretation of the results (Cherry et al. 1991). Finally, as shown in the case studies below, since there is no single standard on how surveys are conducted and published, any attempt to synthesize or compare survey data is a demanding and time-consuming process that requires either extensive coarse graining of data, or simulation to extract meaningful models of settlement patterns.

Synthetic Approach

Spencer and Bevan (2018) compiled a synthetic dataset from three regional intensive surveys in Crete to perform point process modelling and gain insight into first order and second-order locational preferences for settlements in the Cretan Bronze Age including the Gournia survey (Watrous et al. 2012), the Vrokastro area survey (Hayden 2003, 2004, 2005), and the Kavousi area survey (Haggis 1996).

The authors define these surveys, as well as most surveys conducted and published before the 21st century as legacy datasets, as the lack of artefact-scale data and their reliance on non-georeferenced distribution maps create challenges for future synthesis: sampling bias and survey intensity is not easily quantifiable, and the assessments of surveyors on certain site attributes (Spencer and Bevan 2018) cannot be critically assessed due to the coarseness of published data. Additionally, the fact that the assessment of site attributes like size is arbitrary and is not a result of explicitly defined criteria on the basis of sherd density/type per tract results in datasets that can be used in isolation but cannot easily be compared in terms of their methodology, let alone used to produce synthetic datasets.

The approach adopted by Spencer and Bevan (2018: 73) was to identify three contemporary surveys conducted using similar methodologies and sampling strategies, whose coverage areas were in close proximity to each other in the area of Mirabello, north-eastern Crete, and which use the ‘site’ as the unit of analysis. One approach taken to address the patchiness of data was to assign “arbitrary buffer(s)” around single point coordinates when site area had not been estimated/reported by surveyors (Spencer and Bevan 2018: 73).

Overall, this strategy achieves the stated aim of producing site-level data for a broader region, greatly supported by recent high-quality survey datasets. Similar approaches could be used to produce site-level data for broader regions, but that would only be feasible if such survey data existed for the entire Aegean.

Southern Aegean – Comparative Approach

Regarding the use of data from multiple surveys for comparative studies, a doctoral dissertation by Stellatou (2007: 53) has addressed the difficulties in comparing regional survey data, with the additional complication of comparing original regional-scale data (from the Kythera Island Survey Project) to legacy datasets in a study spanning the Final Neolithic and Early Bronze Age in the southern Aegean. It was noted that the lack of digitized data from published surveys led to the author deciding to include a smaller number of surveys in the synthetic dataset than intended, because of the time constraints of doctoral research and the time-consuming nature of digitizing legacy data to process them in a GIS.

Stellatou compared data from the Methana Survey Project (Mee & Forbes 1997), the Argolid Exploration Project (Jameson et al. 1994), and the Kythera Island Project (KIP). Sampling strategies differed between studies (or did not exist in the case of the first season of the extensive survey part of the Argolid Exploration Project), as did the distance between fieldwalkers and the definition of sites: in the Methana survey, scatters above 5 sherds per square meter were given site status, and the edges were defined using a minimum cut-off of 2 sherds per square meter (Stellatou 2007: 74), in the Argolid survey there was no explicit criterion for defining a scatter as a site, but edges were defined using a cut-off of one sherd per 10 square meters (Stellatou 2007: 76). Overall, data quality for the KIP was superior to that for the Methana and Argolid surveys: data was digitized and incorporated in a GIS, the sampling strategy was partly probabilistic and site edges were defined in terms of relative sherd density.

However, there were some key factors that made the surveys comparable: all site scatters were revisited to verify their site status, the relationship between on-site and off-site artefact densities was formally investigated, and 1:5000 maps from the Hellenic Geographic Army Service were used to plot tracts and sites, making their subsequent incorporation into a GIS easier.

Overall, the aim of Stellatou's thesis was to compare settlement patterns for the Final Neolithic and Early Bronze Age on the basis of different landscapes, not to synthesize data to produce an overall site distribution for the Southern Aegean. It was necessary to identify specific criteria that could render these surveys comparable.

Given the time-consuming nature of digitizing and synthesizing legacy data for only three microregions covering an area of roughly 100 sq.km (intensively surveyed), and the fact that artefact density data are rarely published in Aegean surveys, it is safe to assume that using similar methods to extract site distribution-data for the entire Aegean, including Crete and Western Anatolia (the study area selected for this project is nearly 200000 sq. km.), would be unfeasible given the time-constraints of a doctoral project. I explore a different approach for estimating the number and location of communities/sites, as well as the size of each community in the section below.

5.4.2. Legacy data from government databases

As aforementioned, given the scope of the project, it is best to adopt a cost-effective approach to collect the site locations necessary validate a network model. One alternative to integrating data from surface surveys is to turn to archival material from official sources, such as a list of protected monuments in countries where legislation is put in place to record and protect all known archaeological sites, wherever this is possible. The "Ongoing Catalogue of the Listed Archaeological Sites and Monuments of Greece" (http://listedmonuments.culture.gr/info.php, 2020) has been published by the Hellenic Ministry of Culture and Tourism since 1993 and includes all known "immoveable monuments, archaeological sites and historic places that required a specific legal act of designation, demarcation and protection" identified since 1921 in an online dataset, that includes information on the location, date, site type, and Government Gazette (the official journal of listed legislation). This sampling approach is biased in that it does not include a large region of the study area — Western Anatolia falls under the modern state of Turkey and is therefore completely missing from the database.

In the catalogue, sites dated to the Bronze Age are usually placed within one of three time-slices: "Early Bronze Age", "Middle Bronze Age", and "Late Bronze Age", or a wider interval like "Bronze Age". Regarding site location, for monuments listed in the Government Gazette after the 1970s, a list of coordinates demarcating a protected area polygon are provided using the Greek Geodetic Reference System 1987 (GGRS87) or coordinates on specified maps from the Hellenic Geographic Army Service (usually 1:5000 scale).

In terms of sampling strategies, the catalogue includes all archaeological sites "known" to the Greek government, including both chance discoveries and sites confirmed through more formal means like academic excavations and surveys. Therefore, both purposive and random sampling could be considered to have contributed to compiling the dataset. Unfortunately, no information on the circumstances that led to the discovery of archaeological remains is included in most Gazette documents. Nevertheless, they are always signed by the Director of the local Ephorate of Antiquities, which means that government officials and archaeologists have verified the presence and dating of archaeological remains listed there.

The presence and extent of recovery bias must be assessed, at least informally if not explicitly quantified, as the significant economic cost of delaying construction or the use of plots where archaeological remains have been found frequently results in owners and construction companies either failing to report preliminary findings to the Archaeological Service, or refusing to let go of their property in cases where the Archaeological Service deems the remains worth protecting — which in itself is very costly (Catling 1980: 3) and thus not frequently pursued by the state. This problem occurs both in the context of urban construction (Μπρουσκάρη 1993: 456) and in the case of rural sites (Kaliampakos et al. 2011: 159-160). Additionally, urban sites are more likely to be subject to rescue excavation by the Archaeological service, and thus not included in Government Gazettes aimed at "protecting" monuments by prohibiting future construction activities. Therefore, a strong bias against urban areas should be present in the catalogue of listed sites.

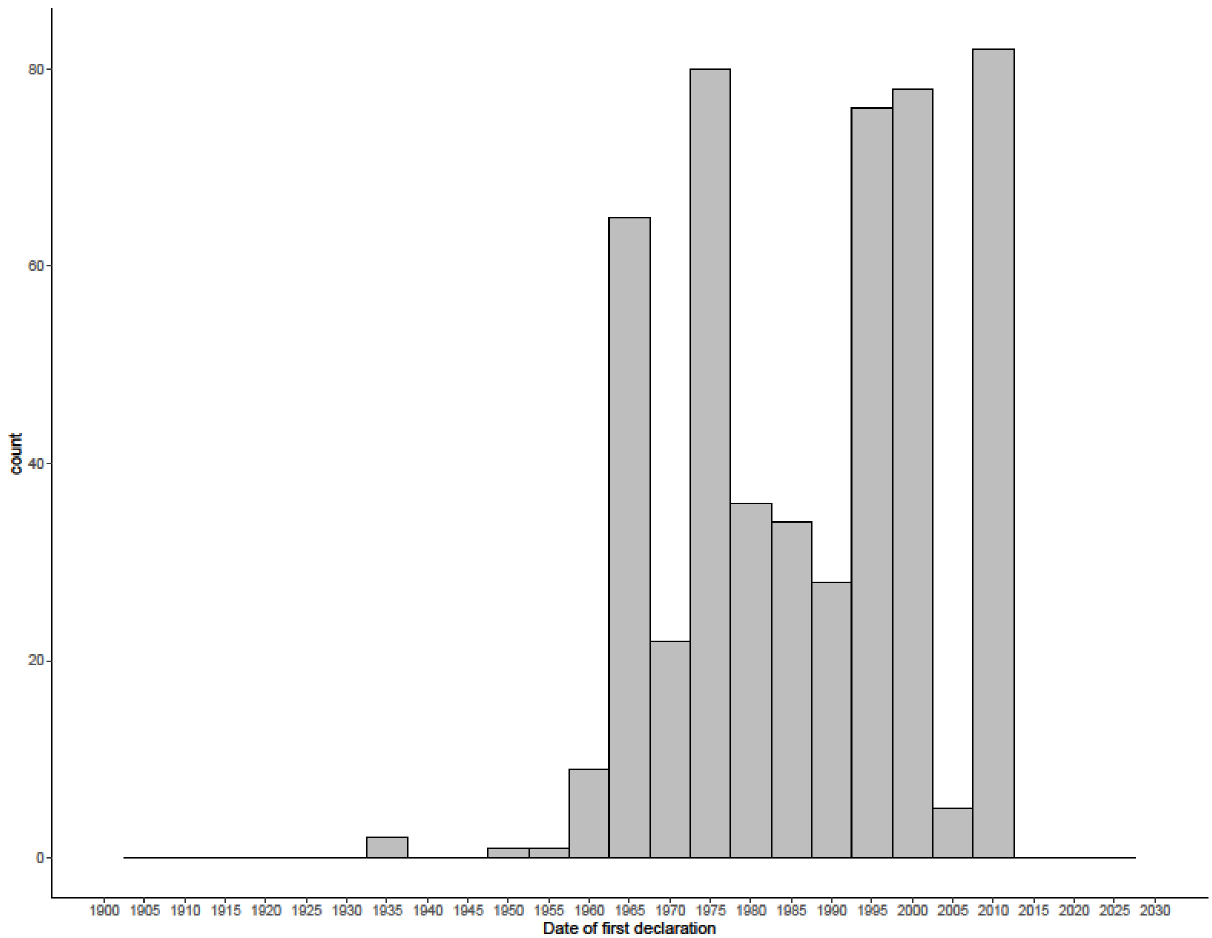

Figure 13: Histogram showing the date for the first declaration of each protected monument in the Catalogue. The peaks around the period before and after the Junta, and during the 90's – periods when there was a boom in development in Greece.

The "Ongoing Catalogue of the Listed Archaeological Sites and Monuments of Greece" includes information on the date when each site was declared a protected monument in the Government Gazette. As shown in Figure 13, there was an increase in the intensity of government affiliated archaeological research in the early '60s and late 70's, with a sharp decline during the Greek Junta, which lasted from 1967 to 1974. A similar trend can be observed after the mid '90s, with a break during the recent economic crisis. The fact that those periods were associated with a boom in development suggests that the reason each monument was declared was because it was already in danger due to ongoing development in the area. This also means that there could be multiple sites that were not declared as protected monuments because there was no development happening in the area — a sampling bias that probably means that regions of Greece that have high densities of modern populations will be overrepresented in the database — so in this sense and in contrast to what was discussed above, the sampling bias was probably for, not against urban areas.

I include a short description of the variables recorded in the catalogue for each site:

1. Coordinates
2. Geodetic Reference System
3. Dating

Coordinates

As shown in Figure 14 out of 521 sites included in the database, 332 entries (63.72 %) included coordinates demarcating the boundaries of the polygon for each protected area. Additionally, as shown in Figure 15 most polygons include a single archaeological site.

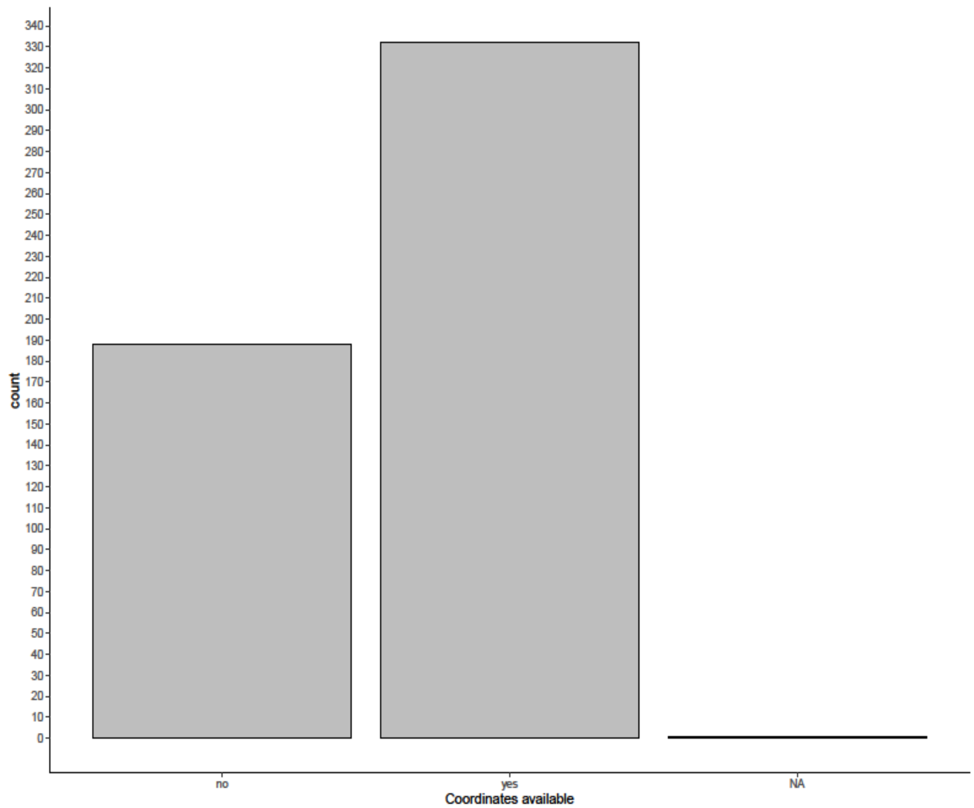

Figure 14

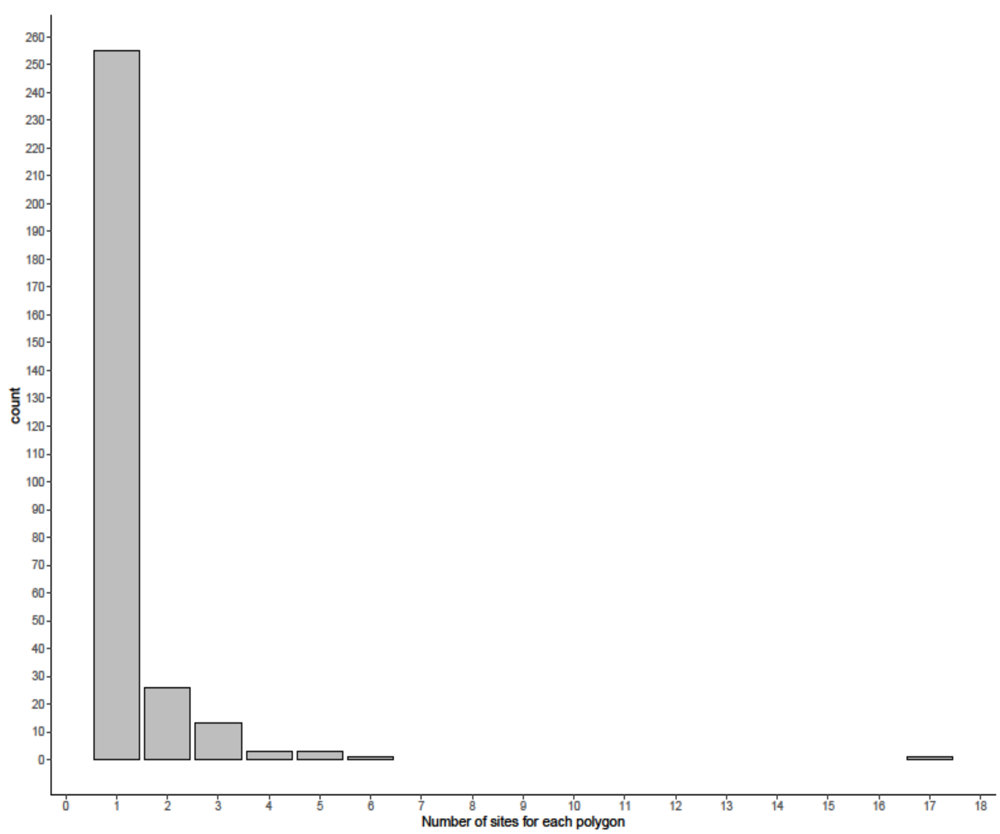

Figure 15

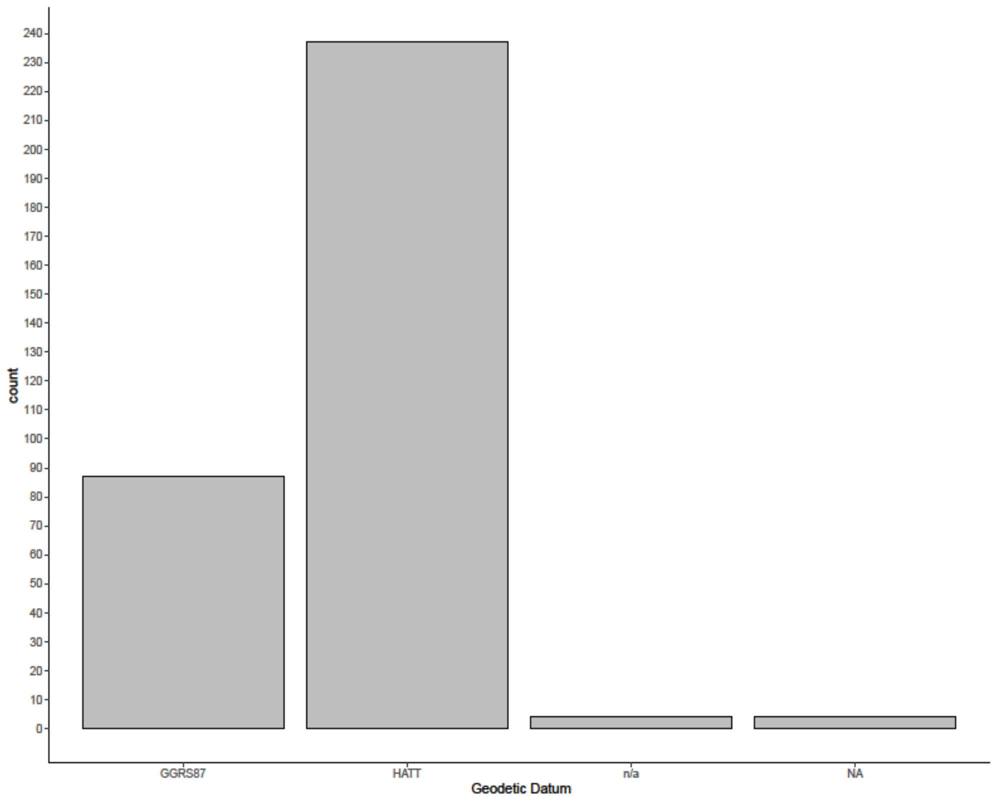

Figure 16

Figure 16 shows that out of the 332 polygons, 237 (roughly 71.39 %) use the older HATT geodetic datum, and 87 (26.2 %) use the Greek Geodetic Reference System 1987 (GGRS87) datum, which replaced the former after 1987.

Geodetic Datum	Ellipsoid	Projection
GGRS87	GRS80	Greek Grid (Transverse Mercator)
GR-Datum (HATT)	Bessel	HATT

Dating

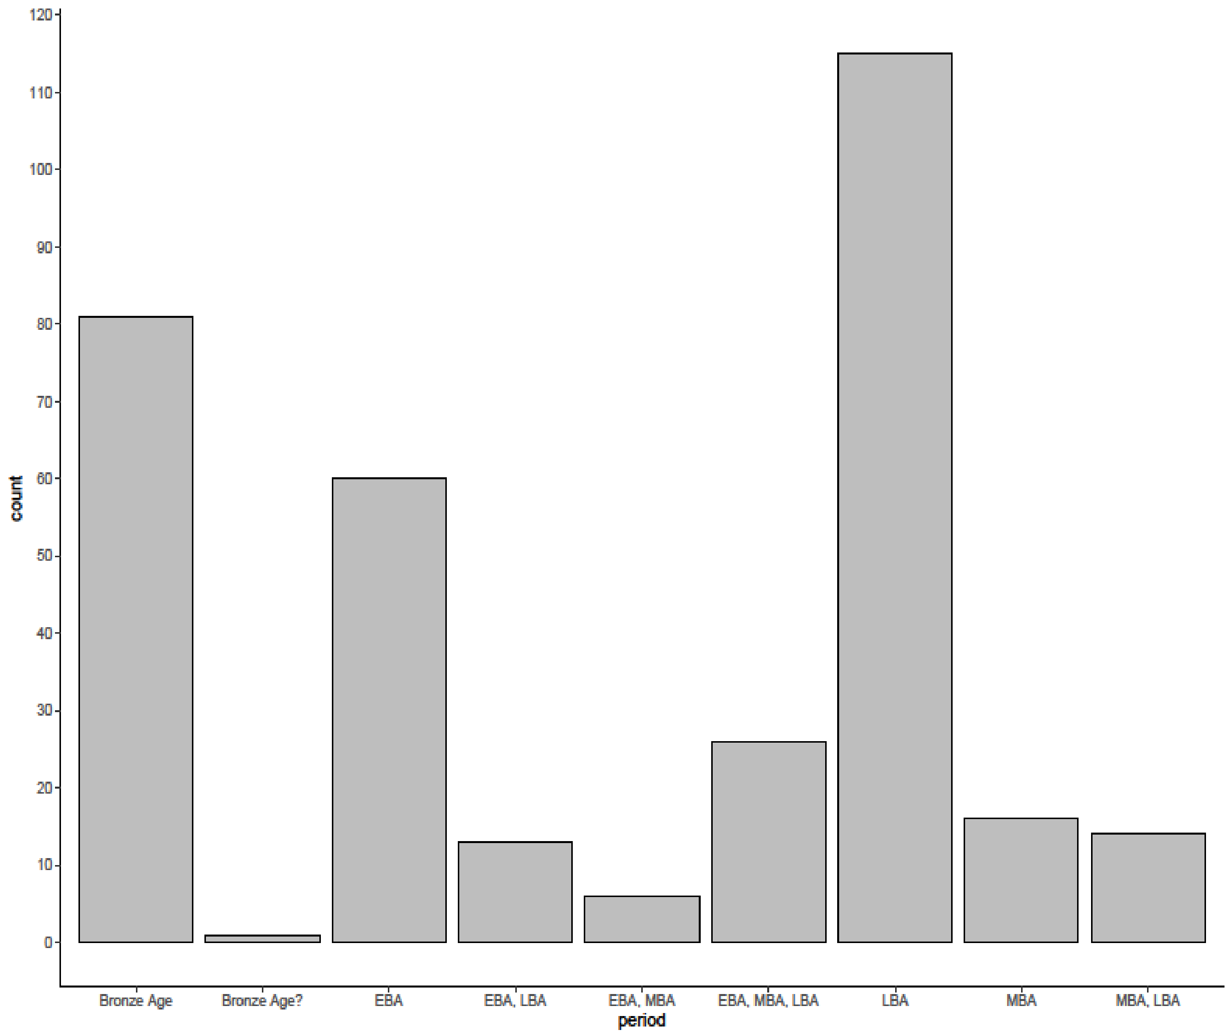

Figure 17

Looking at Figure 17 it is possible to get a rough estimate for the relative number of communities belonging to each of the three Bronze Age subphases. Late Bronze Age (LBA) is the period with the greatest number of sites, followed by Early Bronze Age (EBA), and then Middle Bronze Age (MBA). The categories “EBA, LBA”, “EBA, MBA”, “EBA, MBA, LBA”, and “MBA, LBA”, refer to multi-period sites or multi-period polygons. The “Bronze Age” category includes sites and polygons that were not securely dated to one of the three main sub-divisions. Figure 13 shows the locations of protected areas belonging to the three subphases of the Bronze Age and represented by their centroids.

The maps confirm that based on archaeological data from the Greek government, population levels in the Aegean measured as the number of communities decreased from the Early Bronze Age to the Middle Bronze Age and then boomed during the Late Bronze Age. Because of the coarse temporal resolution, this dataset fails to identify within-phase variation in population levels – for example, it fails to identify a well-established population bust at the end of the Late Bronze Age (see section 5.4.3). In terms of spatial sampling, it appears that the Peloponnese, Attica, Euboea, the Macedonian plain, and Crete are well represented in the database, while other parts of the study area are completely missing — Western Anatolia is obviously missing due to being part of Turkey, but the Aegean islands are also underrepresented (except perhaps for the Early Bronze Age). Overall, site intensity is higher in coastal regions — it is hard to tell whether this reflects a ‘real’ spatial pattern that was occurring in the Bronze Age or whether it is caused by sampling bias.

Figure 18: The locations of protected sites in government records dated to one of the three Bronze Age subphases (Early, Middle, and Late Bronze Age) and represented by their centroids.

5.4.3. Demographic information

As discussed in Chapter 2, changes and fluctuations in population size at the level of the community should affect rates of innovation adoption in the study area, and ideally a suitable proxy capturing the size of the population at the level of the community/site and linking it to specific locations would be necessary both for the formulation of the model, as well as for comparing the outcome of the simulation to the archaeological record. An increase in the number and size of settlements might reflect an increase in population at the macro-scale, though a coarse temporal resolution might produce artificially large numbers of sites for a timestep/phase due to time-averaging.

Furthermore, especially regarding population increase in the Bronze Age Aegean, researchers using information from surface surveys have convincingly argued that population increase might be reflected less in an increase of all sites, most of which would be small villages/farms, and more in the increase and greater size of communities with “urban” characteristics (Whitelaw 2017).

Based on the analysis of data from a Greek government database, there was an increase in the number of sites in the Late Bronze Age, and population levels dropped from the Early Bronze Age to the Middle Bronze Age. However, this is probably as far as this dataset can take us in linking population levels to site locations and site sizes for the entire study area. Prior to data collection and analysis, I expected that the area of the polygon for each protected area might reflect the size of each site. However, soon after completing data collection from the “Ongoing Catalogue of the Listed Archaeological Sites and Monuments of Greece”, it became evident that polygon area could not be considered a reliable proxy for site size.

After a preliminary inspection of each protected polygon using satellite imagery from Google Earth, I observed that polygon area was not always an accurate representation of the extent of Bronze Age settlements, as the edges of polygons sometimes overlap with those of landscape features such as hills/mountains. In other cases, the area of polygons was readjusted after further archaeological investigation confirmed that the area of the Bronze Age settlement was not as large as originally hypothesized. Interestingly, there were also cases where following the reassessment and re-declaration of a protected area, private properties were built in the areas that were recently removed from the declaration. This suggests that there may have been other factors affecting the selection of the area of each protected site other than the archaeology itself (e.g., pressure from developers and landowners).

The lack of explicit criteria — or explicit description of the criteria — for declaring a protected area in the gazette means that it is not possible to identify cases where such sampling biases affected the area of each declared site and remove them from the final dataset or use statistical treatment to reduce their effect. It was decided that due to the lack of good data, polygon area would be an unreliable proxy for settlement area and the relative population density at each site.

An alternative approach would be to qualitatively identify specific cases in the government database where there is more information on the kind of archaeology included and where it could be safely assumed that there was a correlation between settlement and polygon areas. The area for each of these polygons could then be used to estimate a probability distribution for Monte Carlo sampling, in order to simulate site area values for all polygons to be used as a proxy for settlement size, which would then be used as a proxy for the local population size at each community. However, this approach is time consuming, highly subject to the interpretation of biased government records and would only provide data of a very coarse temporal resolution (the traditional tripartite division of the Bronze Age in Early, Middle, and Late Bronze Age).

Whitelaw's study on Minoan urbanisation (2017) provides frequency distributions for site sizes per phase for Mainland Greece and Crete collated from comparable surveys in terms of site size and dating. However, this only covers half of the study area chosen for this study — it is missing the regions of Western Anatolia and the Aegean islands — and it also fails to associate site location with site size, which would be necessary if one was to use those locations and associated local population sizes as an input for the simulation models.

Overall, the analysis of site locations from Greek government sources suggests that there were fluctuations in the size of the population in the Aegean during the Bronze Age, which resemble a boom-and-bust pattern with booms in the Early Bronze Age and Late Bronze Age (the phase with the highest population levels) and a bust phase occurring sometime during the Middle Bronze Age.

This picture is confirmed by another study (Weiberg et al. 2019) that has brought together locational data from the Aegean and produced regional and 'global' estimates, and which benefits from data of higher temporal resolution. As shown in an example of regional estimates for the Peloponnese in Figure 19, Weiberg et al. (2019), used multiple proxies to infer population levels over a long period stretching from the beginning of the Early Neolithic to the end of the Roman period. A boom-and-bust pattern can be observed for the Bronze Age periods, with an initial boom at some stage in the EBA, a bust at the end of the EBA/beginning of MBA, a boom during the palatial period and most of the LBA, and a final bust at the end of the LBA. This pattern can be observed in cases where site numbers and surface area are used as proxies for population size.

Considering the information available, this is our best estimate for the relative population levels in the wider Aegean region over time during the Bronze Age (as aforementioned, there are other good regional estimates, but they focus on smaller regional units). How such macroscale fluctuations relate to patterns in settlement location is less clear due to the challenges of combining multiple lines of evidence that were discussed above.

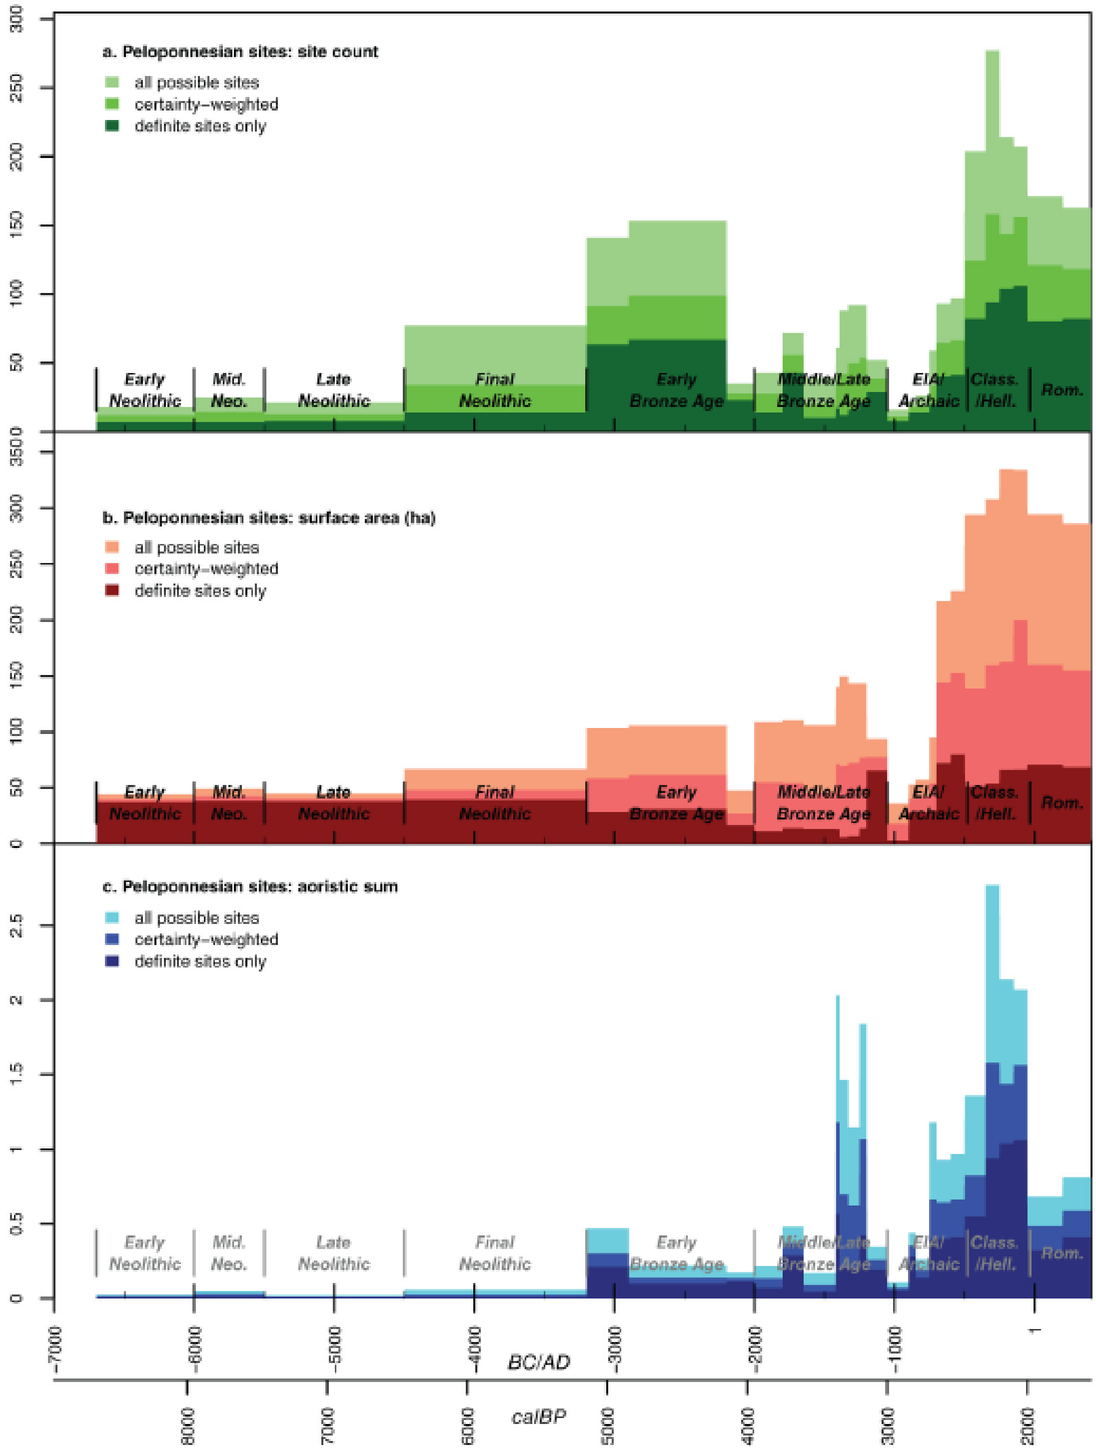

Figure 19: Source: Weiberg et al. (2019:750).

5.5. Compatibility of empirical trends with simulation modelling

Looking back at the insights gained from the analysis in sections 5.1-5.4 it is possible to assess how one can use the empirical estimates for the different innovations and the population patterns to inform the design of a simulation model. Compared to kilns and the IRT trait, the potter's wheel has received significantly more attention regarding the collection of quantitative data to establish adoption patterns over time. Traditional approaches consisted of simply reporting the presence of wheel-made pottery at the site-level over time. Combining such legacy information from 40 sites around the Aegean allowed a rough estimate of the system-level adoption trends over time and indicated that the rate of adoption was slow at first and then increased during the Protopalatial period, which coincided with the emergence of the first palatial systems in Crete. Comparing this empirical trend to simulation models would require that the output of the simulations took the form of presence/absence data at the site level.

More recent attempts at quantification have largely been driven by the adoption of Roux's methods and are limited in number but provide more information at the level of the site — relative abundance rather than simple presence/absence data. When accounting for temporal uncertainty, it is possible to construct empirical trends with Monte Carlo envelopes and compare them to the outputs of simulation models for validation purposes, provided that the outputs of those simulation models were in the form of ratio data/percentages of wheel-made pottery at the site level.

Kiln locations which were mainly drawn from Hasaki (2002) are sufficient information for a coarse-grained look into the intensity of kiln presence over the study area, which is shown to vary over time and in space. The relative importance of Crete over the study area becomes more pronounced after the Protopalatial period, paralleling the patterns from the potter's wheel data.

Aoristic analysis was used to establish a system-wide adoption tend over time through aoristic weight sums, which suggests that kiln use increased over-time in a roughly non-linear fashion. However, aoristic sums are not straightforwardly interpretable, which makes it difficult to compare this proxy to the outputs of a simulation in a way that would enable a direct quantitative comparison.

Regarding demographic data, the discussion in section 5.4 shows that there is a lack of good empirical data on the exact number of communities and their location in the Bronze Age Aegean, and information on the relative size of each community is either unreliable (e.g., the area of protected monuments/sites from government records) or focuses on a small part of the study area (e.g., the information from Whitelaw 2007). However, multiple lines of evidence suggest that population levels rose during the Early Bronze Age, dropped sometime during the Middle Bronze Age, and then rose to their highest during the Late Bronze Age, the final phases of which were marked by another drop in population levels.

Overall, the empirical trend with the Monte Carlo envelope using percentage data for the potter's wheel appears to be the best choice of an empirical trend that can be quantitatively compared to the outputs of simulation models. The choice of this empirical trend has also informed certain aspects of the design of the simulations, especially regarding numerical outputs — the following chapter describes design of the simulations and the assumptions that inform them.

6. Assumptions, ontologies, and formalisms

Introduction

As with any model, a simulation model creates a simplified version of the real world in order to address specific questions. However, the process of simplifying reality is no trivial task, particularly when the ‘reality’ one is interested in existed thousands of years ago. The process of simplification requires making assumptions about what ‘matters’ in the past, and what does not in terms of identifying aspects of past reality that should be included in the model to provide meaningful answers to the research question. In the process of building an in-silico reality, the separate entities comprising the model are also simplifications of structures or elements of the physical world, and it is important to define what their role in the model is, how they relate to each other, and what assumptions were made to define them ontologically (i.e. basic issues of how best to categorise real world things or processes, and how problematic the borders of such classification are). Finally, formalising the model requires coming up with specific rules as to how the separate parts and entities comprising the model are supposed to interact.

In the first section of this chapter, I will explore the main assumption of the model and investigate its validity on the basis of past research. In section 2, I will explore which methods can be used to define entities reliably and transparently in a simulation model. The formalisations and assumptions used to ontologically define said entities are discussed in section 3. Finally, section 4 describes the mathematical formula used for the simulations explored in the following chapters. The behaviour of the model described here is explored in the following chapter by running simulations in a simple environment.

6.1. The assumption that distance matters

As discussed previously, the social learning that occurs between different communities can take many forms, occur at the intra-community or inter-community scale, and result in different degrees of adoption (or non-adoption) with respect to the incorporation of new traits in local pottery assemblages. However, as highlighted in Chapter 3, the aim of this study is to identify the main factors affecting the transmission of pottery traits in a parsimonious manner, which means starting with the simplest explanation and building in complexity only when necessary.

In the simplest model, the process that results in the diffusion of a trait in space can be reduced to a single variable: geographical distance. The main assumption in this case is that individuals and communities that are located close to each other will interact more and as a result are likely to be more similar to one another than individuals and communities that are located further away from each other. This can be due to a number of reasons. For example, the transmission process, which could be facilitated by long-term learning between communities or exchange in a market setting, will require a greater amount of effort and time in terms of travelling between different locations when the distance between said locations is larger.

This idea has been explored in archaeology, with applications ranging from case studies investigating the potential of an identifiable spatial pattern to the "Maya collapse" (Kvamme 1990: 197), to others studying the connectivity of Roman amphora trade through similarity measures (Rubio-Campillo et al. 2018). The theoretical framework that was used in those studies to explain the relationship between geographic proximity and cultural similarity is referred to as "Tobler's First Law of Geography" (Kvamme 1990: 198-199), often while quoting the following axiom: "everything is related to everything else, but near things are more related than distant things"

(Tobler 1970: 236), without explicit reference to the effects of an important element of Tobler's simulation study: reproduction at the level of the population.

For all intents and purposes, in archaeology Tobler's axiom can be applied to cultural similarity arising from two distinct processes that produce a similar signature: spatial autocorrelation that emerges from a process where similarity between different locations is due to their sharing a common value for an environmental variable — this is sometimes referred to as "exogenous" spatial autocorrelation (Bonada et al 2012: 57) — which we might expect in cases where cultural traits such as the choice of material to produce an object are affected by geography (e.g., the use of shell as temper on coastal locations) — and spatial autocorrelation as the effect of an "isolation by distance" spatiotemporal process (Wright 1943), which we might expect in cases where a cultural trait is passed on intergenerationally and by means of horizontal transmission between different communities, and where the strength of transmission depends on the relative cost (material resources and time) of travelling to and from said communities (this can be considered "endogenous" spatial autocorrelation (Bonada et al 2012: 57)). This study is concerned with the latter.

To be more specific, the term 'isolation-by-distance' (Wright 1943) originated in evolutionary biology and population genetics and incorporates the element of reproduction more formally. Its basic premise is that given a certain environment, populations that are located closer to each other will share more genes in common than those located further away from each other.

Archaeological studies with a cultural evolutionary agenda have used the concept of isolation-by-distance to either test whether it produces significant deviations from the expectations of a neutral model (Premo & Scholnick 2011) of cultural transmission, or to construct a null-model against which to evaluate competing models about macroevolutionary change (change in trait diversity over time) (Shennan et al. 2015). The potential of the isolation-by-distance concept for addressing archaeological macroevolutionary questions has been explored both through theoretical simulation (in the case of the Premo & Scholnick study) as well as through the direct comparison of archaeological data to an isolation-by-distance null model.

Premo & Scholnick's (2011:164) main aim was to establish whether varying the spatial scale of cultural transmission produces a different trait diversity "signature" to what would be expected under a neutral model of social learning. The neutral model, in this case, is the "Fisher-Wright" model, which the authors describe as a model where reproduction takes place randomly in a "finite" population whose size does not change and where there is no selection or overlap between generations (Premo & Scholnick 2011: 165). When certain conditions hold — e.g., when the rate at which new traits appear in the population (i.e., innovation) is equal to the rate at which traits disappear randomly due to stochastic effects (i.e., drift) — the situation is one of unbiased transmission (Premo & Scholnick 2011: 165). Such neutral models, which were first introduced in archaeology by Neiman (1995), have been used as a null hypothesis against which to evaluate the possibility of alternative social-learning mechanisms and transmission biases (i.e., non-random teachers/traits copied) in previous studies by Shennan and Wilkinson (2001) — where anti-conformism was the alternative hypothesis- and Kohler et al. (2004) — where conformism was the alternative hypothesis.

What Premo & Scholnick (2011) suggested was that the spatial scale of the analysis is an important factor affecting social learning, and that the assumption of "global" social learning – elsewhere referred to as panmixia (Premo & Scholnick 2011: 165) — does not always hold in the context of prehistoric cultural transmission, where “long distances inhibited the transmission of cultural variants” (Premo & Scholnick 2011:166). With the aim to explore the degree to which altering the scale of social learning on the neutral model can produce a signature similar to biased transmission, they built a “spatially explicit” (Premo & Scholnick:175) agent-based model where it was possible to define neighbourhoods within a distance around each individual (in Euclidean space), and where each individual was equally likely to learn from any individual (including themselves) in the neighbourhood (Premo & Scholnick 2011:166-168). This distance parameter allowed exploring the effect of spatial scale on unbiased transmission, and when this scale was set to “global”, it represented the neutral model described above. The authors also studied the effect of violating other assumptions of the null model, such as varying population size and density (Premo & Scholnick 2011:166). Ultimately, the study established that when the neighbourhood where each individual selects their teacher is much smaller than the size of the entire study area, trait diversity can be higher than expected from the aspatial neutral model described above, and it can falsely lead one to interpret it as an indicator of biased transmission (Premo & Scholnick 2011:173).

Shennan et al’s (2015: 107) study of European Neolithic morphological trait variation used isolation-by-distance as a null model against which to evaluate the possibility that the total inter-assemblage similarity between different sites could be attributed to transmission processes operating between distinct groups in the population and was not simply a function of their distance in space. The authors collated data on pottery and ornament traits at the site level and computed a similarity index to compare how similar pottery and ornament assemblages were between different sites (Shennan et al. 2015: 105-106) and presented convincing evidence of the existence and effect of population structure in inter-assemblage similarity.

The aforementioned study applied the isolation-by-distance model on a scale very much comparable to the one that Wright's original model was formulated on: the site-level used by Shennan et al.'s (2015) study was treated as the equivalent of Wright's sub-population-level, and its focus was on trait frequency over time — an equivalent to "gene frequency" that was the basis of Wright's (1943) model. Finally, all sites were located on the same landmass, which could be considered a "continuous area" as per Wright (1943: 117) if one assumes that natural barriers and terrain ruggedness did not greatly affect population dispersal. Most importantly, the study established the potential of isolation-by-distance as a simple model that can be used as a null hypothesis to alternative models of transmission in cultural macroevolution.

Overall, the literature suggests that geography can greatly affect the process of trait transmission — but it must be noted that the aforementioned studies focused on macroevolutionary patterns of trait diversity rather than the transmission of a single trait in isolation. The assumption that distance in space inhibits social learning between communities is a reasonable one given what is suggested by the outcomes of simulations (e.g., Premo and Scholnick 2011). Therefore, this study will adopt the assumption that distance matters, and attempt to explore the effects of distance under a range of different scenarios. The formulation of the model (which is discussed in more detail below) allows exploring a case where distance does not matter, so it is possible that simulation outputs will question the validity of this assumption if they do not produce a similar signature to the empirical trend for the potter's wheel innovation.

6.2. Model ontology

The archaeological and ethnographic records provide us with plentiful but fragmentary information on the aggregate outcomes of the decisions and actions of craftspeople in the past in the case of the former, and a more dynamic view of the community-wide processes that characterise pottery making at shorter time scales in the case of the latter. As discussed in chapter 2, both viewpoints — archaeological and ethnographic — provide evidence for diversity in pottery-making in terms of processes and sequential stages of production (chaînes opératoires), the types of relationships and social learning occurring between potters, and the different trajectories of novel traits as they are copied (or not) between potters. However, if one is to identify the main macro-scale factors affecting technological adoption in the Bronze Age Aegean, it is worth attempting to 'abstract away' from specifics and identify aspects of the empirical reality that are more relevant in addressing the research questions.

The archaeological record rarely allows identification (Sillar and Tite 2000:9), let alone description, of the decisions and actions of individual potters, and the main unit of analysis in Aegean archaeology is usually the site, which typically describes the preferences of a community comprised of consumers and producers/potters, or site-based workshops — e.g., in an Aegean context, Quartier Mu at Malia during the earlier 2nd millennium BCE. Therefore, rather than using individual potters and consumers as entities in the model, a practical approach is to move upwards to a more macro-scale entity and consider the community as a discrete unit in order to ensure that the outputs of the simulation are comparable with the empirical data, which are collected at the level of the site.

Given that the research questions focus on the effects of geography in trait transmission, it makes sense to consider communities (our chosen entity/unit of analysis) as being located in an environment that incorporates aspects of geography that could affect trait transmission in the world of the model. Another characteristic of the entities is whether they can move within that environment, or whether they are stationary — in this regard, communities of itinerant potters could be characterised as different types of entities to communities of potters who are associated with a single site and who only move in the vicinity of that site to collect different types of clay or fuel.

The choice of entities that represent aspects of reality captured by the empirical record is crucial when it comes to formulating and running simulation models that produce outputs that are compared to the empirical data. To put it simply, one needs to ensure that they are comparing like and like, even if, as suggested by Kohler and van der Leeuw, "a good model is not a universal scientific truth but fits some portion of the real world reasonably well, in certain respects and for some specific purpose" (Kohler and van der Leeuw 2007:3).

There are two concepts that are key in this process and that can be used in tandem to inform the design of models with meaningful outcomes: abstraction and formalisation. Abstraction is necessary in order to filter out the details of the "real" world and produce categories of things that can be considered to describe roughly the same entity/process, and formalisation follows the abstraction process by setting up a set of rules that must apply in order for an object/person etc. in the real world to fit into one of the categories that were created. Abstraction is essentially two-fold: abstraction of the real world to produce entities that can be sampled to produce an empirical dataset — in this case the 'site' was the entity of choice for the archaeologists producing the primary datasets that were explored in Chapters 4 and 5 — and further abstraction to produce a representation of those entities that can be used in the context of a computational model (for instance, using a point in 2-dimensional space to represent a site location).

Finally, scale is important in terms of defining entities. As in the example above, it is possible to reduce a system to its most basic physical components and define entities that represent individual human beings, but it is also possible to define entities at higher scales, such as communities, if the goal is not to explain how such high-level entities came into being (Lake, 2015, 24). According to Lake (2015:25), if the goal is to formulate a model that produces a phenomenon of interest, it is useful to incorporate interactions between agents and between agents and their environments that are "relevant to their scale". This methodological stance can be applied to the definition of all entities at the "appropriate" (for lack of a better word) scale, whether they represent active agents (as in the case of most agent-based models) or features of the environment.

In Section 6.3, I define the entities that were used to formulate the simulations in following chapters from an ontological perspective: what their properties are, how they relate to each other, and how they are represented in the model. I will also describe how they relate to the archaeological and ethnographic record and attempt to justify the main assumptions that were made in the process of defining them. Finally, I will present the empirically unknown aspects of reality that I will experiment with in the simulations included in the following chapters (model parameters) and how they relate to the entities in the model and present a mathematical formula that represents what happens to the entities at each moment in time in the model (section 6.4).

6.3. Ontologies and formalisms

6.3.1. Entities

Communities and the strength of craft transmission between them as nodes and edges in a network

The main entities in the model are communities and the relationships between them and they are represented as a spatial network. Networks provide good abstractions for systems where we know that relationships between different entities exist and "matter" (Brughmans 2013: 625), and where one is interested in the outcomes of those relationships, particularly when there is a spatial component that affects them. Here, communities are represented by nodes, points in space that are connected by edges that represent relationships. Each edge has a weight associated with it, a numerical value that represents the strength of transmission of information between a pair of nodes as a function of the spatial distance between them. If a pair of nodes is near each other, they will affect each other more than they will affect other nodes that are located further away, and the edges between them will hold a higher edge weight relative to their edges to distant nodes. Adopting a network abstraction to represent the entities in the model is useful from a practical perspective — it is easy to identify communities/nodes that are important or central in terms of trait transmission through the use of network centrality measures (on the topic of centrality measures see Collar et al. (2015)).

Communities are defined in the model as stationary entities that represent a community of potters and consumers that is linked to a permanent settlement at a specific location. Each community can produce wheel-made or handmade pottery, which allows it to have a feature representing the relative proportion of wheel-made pottery over handmade pottery. Given that most published empirical studies with reported percentages of wheel-made and handmade pottery do not explicitly measure the relative use of each technique for different shapes and vessel sizes, I do not incorporate such distinctions in the model.

I define the community entity as stationary based on the assumption that the production of wheel-made pottery was associated with pottery production in permanent settlements. There are several points that suggest that this assumption is valid. First, one can take it as a given that permanent settlements were the primary mode of habitation in the Bronze Age Aegean. Additionally, while the introduction of the potter's wheel is associated in the literature with communities migrating from Western Anatolia, and while it is possible that there existed itinerant potters making pithoi during the Bronze Age (Katsonopoulou et al. 2016), there is little evidence to suggest that the making of tableware — the type of pottery the potter's wheel is typically associated with — was associated with itinerant production. Furthermore, wheel-made pottery is consistently associated with specific types of fine fabrics in terms of coarseness — very fine in Emporio on Chios and fine to medium coarse at the Heraion on Samos in EB II according to Choleva (2020: 85) —, which are, as ethnographic studies suggest, frequently (but not always) linked to a stage in the operational chain that is associated with production on settlements near the clay extraction location — for example, some examples of modern African potters travel distances of 3 km to procure raw materials (Gosselain and Livingstone-Smith 2005: 35), but note that there are more 'extreme;' cases where people involved in pottery production travelled as far as 15 km to procure different types of clay (Nicklin 1979: 441). The process of clay extraction and processing to produce fabrics is also frequently associated with specific groups but not individual potters, e.g., in Northern

Cameroon (Livingstone-Smith 2000:33). Even if one assumes that clay was procured and processed by groups in permanent settlements and that the potters using the wheel were travelling between communities and were only involved in the forming and firing stages of the operational chain, there is still the problem of fuel procurement and firing: the picture that emerges from ethnographic studies is one where in the context of specialised production, fuel procurement is a particularly laborious task that involves moving large quantities of heavy fuel, and is therefore typically carried out at a local level because it is more cost-effective (e.g. wood procurement from somewhere in the vicinity of a permanent settlement) as in Arnold's "Ceramic Resource Threshold Model" (Arnold 2006:3).

Finally, the fact that wheel-made pottery in the Aegean is relatively fine suggests that the lack of added temper, or at least the limitation in the hardness, quantity, and size of such inclusions, makes this type of pottery less resistant to thermal shock, which occurs when the firing temperature rises too quickly and results in fractures on the pot (Tite et al. 2001: 313), as experimental studies have confirmed that the addition of temper improves thermal shock resistance (Tite et al. 2001: 321). This means that in the absence of tempering, additional effort should be made from craftspeople during the firing process to ensure that the contraption used for firing is designed in a way as to mitigate that effect, for example by using kilns (Sillar and Tite 2000: 5).This suggests that potters involved in the forming stage of wheel-made pottery production were either the same as or worked alongside potters involved in the firing stages, and possibly lived in the same communities. Given the community investment required to build Bronze Age kilns and considering that it is possible that that they were used for multiple years/decades (at least on the basis of ethnographic data from mainland Greece e.g., Blitzer (1990: 695)), suggests that at least the firing stage of the operational chain was associated with permanently settled potters. Therefore, the assumed lack of portability of the pottery package consisting of kiln firing/wheel-made pottery and lustrous decoration is reflected by making the communities in the model static.

However, the decision to define communities as stationary does not mean that the model does not consider the outcomes of the movement of individuals and exchange of ideas, know-how, and raw materials between them. This is simply taken into consideration by incorporating the concept of inter-community transmission into the model. The micro-level mechanisms facilitating this transmission and relevant entities defined in micro-scale such as the existence of individual potters that travelled between different communities to enable it are not explicitly modelled due to the lack of secure information on such patterns in the empirical record (see chapter 4). Instead, inter-community transmission is represented by adopting a network abstraction, where communities are represented by nodes and the direction and strength of transmission between different nodes is represented by edges and edge weights respectively.

The temporal aspects of the process are introduced in the model by incorporating a time property to the communities/nodes: each node has a trait proportion value for each timestep in the model, and that value is updated in the next timestep. Edges, however, do not have a time property: here I am making the assumption that the strength of transmission remains the same over time since nodes do not change their position — which would require re-calculating the strength of transmission between a pair of nodes when one or both of them moved — and because it is likely that travel technology did not change significantly enough during the Bronze Age to affect processes of cultural transmission between potting communities. This is a simplification that mainly served practical aspects of the modeling and model exploration process, by ensuring that the model could be adequately explored within a reasonable timeframe. Making communication costs a function of time (e.g., by making communication costs smaller over time as an effect of improved sailing technology) would introduce further complexity in the model. However, the transition from rowing to sailing boats during the Middle Bronze Age — which can be inferred from the fact that material originating from Egypt and the Near East was found in Crete in the Early Bronze Age (Wiener 2014: 10) and through the depiction of sails in Minoan art during the Middle Bronze Age (Broodbank 2010) — certainly made travelling faster (and possibly safer), so future work could involve introducing a temporal aspect to the edge weights in this model.

Environment

The environment in the model is an abstraction of the Aegean in that it is defined as a surface that determines where communities can be placed. The area of the surface that represents the sea is not populated by nodes. When the environment is divided into land and sea zones, the location of communities/nodes is purely random on the land zone.

Apart from this simple two-fold division, it is possible to further divide the land into landscape zones that determine the location of communities in a probabilistic manner: the probability that a node is located on a specific zone in the environment is equal to the probability attached to that zone. This allows for the carrying out of experiments that explore the effects on craft trait transmission of different habitation patterns in relation to the communities' preferences to settle in different types of landscapes. For instance, what would the spread of the potter's wheel look like if communities of potters preferred settling on the coast?

More specifically, landscape zones correspond to different types of terrain and distance from the coastlines. In what follows, the terrestrial zone is split into a coastal flat, a coastal steep, an inland flat, and an inland steep zone. Coastal land is defined as land within a certain distance from the coast, typically 0-5 km, and flat land is defined as land where the slope is below 10 degrees with the assumption that flat land would be, in general, more favourable for habitation in the context of permanent settlements than steep land – flat land being, in general, easier to cultivate and build structures on although steep land can offer defensive advantages. The 10 degree cutoff used to define land as flat or steep is, again, based on an assumption made about habitation preferences in the past (e.g., Bevan, Frederick, and Krahtopoulou 2003) and the 1-5 km cut-off is, accordingly, based on an assumption made from ethnographic information on the average distance that people will cover to cover their subsistence needs, at least in the context of hunting and foraging (Kelly 1992: 47). As discussed by Triantafyllou et al. (2008: 3032), while stable isotope analyses frequently point to agriculture and animal husbandry as the main subsistence strategy in the Aegean during the Bronze Age, "experimental research has shown that certain marine resources, especially those of low trophic level such as sardines and anchovies — often consumed in the traditional Mediterranean diet — might fall below the detection threshold of the methods used". When subsistence strategies are not inferred on the basis of stable isotopes but on the basis of the ubiquity of shell and fish remains the picture changes, suggesting that people partially covered their subsistence needs by fishing — such data suggest that there was regular consumption of marine mollusks, crabs, and some types of fish in Crete and Thera (Livarda et al. 2021: 9, Mylona 2020). Therefore, I consider as coastal the regions where reaching the sea would require at maximum a day's walk and allow the inhabitants of each community to frequently cover some of their subsistence needs by fishing/maritime trade.

As in the case of the "edges" entity, the landscape zones do not have a temporal property (they remain stable over time), even though it is possible that in reality habitation preferences changed over time.

Time

Time is a concept that was previously mentioned in passing but deserves a more thorough discussion. As a fundamental aspect of simulation models (Lake 2015:6), the modeller's treatment of it can affect the way we observe a process. In the context of a simulation, time can be treated as discrete or continuous. Here, I am treating time as discrete and represent the passage of time as a sequence of a finite number of timesteps, episodes of equal length where the modelled process takes place.

The duration of each timestep represents the passage of one calendar year. The main advantage of this approach is that it provides a good point of contact with the empirical record. Archaeological periodisation frequently results in observations being reported in phases, time intervals whose duration ranges from years to centuries. A single calendar year is a unit of time whose length allows comparing the outputs of the simulations with coarser-resolution empirical data relatively easily (for example, an archaeological phase that is 50 years long would correspond to 50 timesteps in the simulation), and it also appears to be a reasonably long period of time for the modelled process to take place – one year is long enough to allow travelling and knowledge exchange between neighbouring communities, even if it does not cover the entire 10-year apprenticeship period linked to forming technique transmission by ethnoarchaeological studies (e.g. Roux 2003:15).

6.3.2. Parameters

As discussed in the Chapter 3, one of the benefits of using simulation in archaeology is that it allows exploring unknown aspects of a process that one believes occurred in the past. This is done by incorporating parameters into the model: numerical values that represent aspects of past reality that are not measurable empirically, and that one can vary at each simulation run in order to explore different scenarios – though sometimes it is possible to have a rough idea of the range of values such parameters can take from archaeological or ethnographic observations.

Population size

Population size, N, is a parameter that determines the number of communities that are placed on the environment and should not be confused with a measure of the size of each community, which is not explored in this model due to the lack of empirical information on the population sizes corresponding to published sites (see Section 5.4 for a more extended discussion on this topic). In the model, communities/nodes represent centroids of population districts of similar size so when simulations are run on a 'realistic' environment representing the Aegean, there is a higher population density in the main plains by virtue of higher probability of nodes in those areas. This in a sense adds in some basic population variation without varying site size on a per community/node basis.

The parameter can only take integer values, given that the community is a discrete entity – it does not make sense to have "half" a community or "half" a node in the model. Furthermore, with the exception of the simulations in Section 8.4, in most of the simulations included in this study the parameter is not a function of time, meaning that I am only exploring the model under conditions where the population is stable (in equilibrium).

Penalty for crossing large distances *k*

The penalty for crossing large distances k is a measure of the communication decay associated with large spatial distances between communities (the mathematical formula incorporating this parameter is shown in section 6.4). Values of k larger than 0 create a negative exponential relationship between the Euclidean distance in space between two communities and the strength of transmission between them. Roughly speaking, this means that the greater the distance between two communities is, the transmission of information between them will decrease by an amount multiple times larger. Varying this parameter allows exploring what the effect of different travelling methods and technologies is on the transmission of traits between communities in the simulation. When its value is set to be equal to 0, all edge weights are equal to 1 and, therefore, distance has no effect on trait transmission, allowing one to test the starting assumption that ‘distance matters’.

Rate of local adoption *a*

The rate of local adoption, *a*, is a measure of how quickly the proportion of the trait in each community/node increases between the previous timestep and the current one regardless of the contribution of other communities/nodes i.e., regardless of the contribution of "horizontal" transmission between them (the mathematical formula incorporating this parameter is shown in section 6.4). Varying this parameter allows exploring what trait transmission would look like if each community was very receptive, a little receptive, or not receptive to the new trait - but where "receptivity" is treated as being continuous. For example, high *a* can reflect cases where the community is so invested in the new trait so as to prefer it to the alternative and partially abandon the alternative trait.

The parameter can be used to reflect the effects of a factor limiting the adoption of the trait at the level of the community. This allows modeling the effects of the decisions of elites who are sometimes discussed as partially determining the proportion of potters producing wheel-made pottery due to the innovation only being available to potters attached to them (e.g., Roux 2003).

The simulations in the following chapters examine variations in *a* in two different scenarios: one where *a* is the same for all communities, and one where *a* takes different values for different communities. The latter possibly reflects reality more accurately than the former, but it is only by running the simulations that one can get an idea of whether this makes a big difference in practice -and under which conditions it does so.

Table 7: Entities included in the model, their properties and relation to each other, and the parameters affecting them in the model.

Entity	**Properties**	**Representation**	**Relation to other entities**	**Change over time**	**Parameters acting on entity**
Community	Discrete Proportion of pottery with trait Timestep Location	Node	Environment determines location	Yes	Population size N, *a*, habitation probabilities
Contact between Communities	Edge weight	Edge	Connects Communities	No	*k*
Environment	Habitation probabilities	Vector maps	Determines the location of Communities	No	Coastal distance and slope thresholds defining regions (default at 5 km and 10 degrees respectively)

6.4. Mathematical formalisms

The following formula describes how the trait value for each node/site is calculated for each timestep in the main model:

$$c_{it} = (c_{it-1})a + \sum_{i=j}^{N} c_{it-1}\, w_{i,j}$$

$$w_{i,j} = \begin{cases} e^{-\kappa d_{i,j}}, if\ i \neq j \\ 0, if\ i = j \end{cases}$$

Where c_{it} is the trait proportion value at community/node i at time t, a is a globally defined parameter determining the strength of local adoption at every community/node (meaning that the rate of local adoption is the same for every community), N is the number of communities in the study area, and $w_{i,j}$ is the edge weight between nodes i, j , and $d_{i,j}$ is the Euclidean distance between nodes i, j. Self-weights for nodes are set to be equal to 0, but this does not mean that the model does not incorporate the effect of each node increasing on its own – this effect is captured by setting a value for the parameter a that is higher than 1. If a is equal to 1, the trait proportion value at each node does not change on its own. If a is lower than 1, local adoption is delayed.

Table 8: The main parameters in the model and their value ranges. Note that most experiments explore smaller ranges.

Parameter	Represents	Range
a	Rate of local adoption	0.2 to 1.6
κ	Penalty for crossing large distances	0.2 to 1.6
N	Population size	500 to 2500

7. Developing understanding with 'toy' models

Introduction

In the previous chapter I defined a simulation model describing the effect that spatial interactions of varying strength between communities of practice and consumption have on global adoption patterns for a single cultural trait (wheel-made pottery) over discrete timesteps (e.g., timestep 1, timestep 2). Several assumptions drawn from the archaeological and ethnoarchaeological literature informed the construction of an appropriate abstraction to represent said communities as well as the definition of an ontological schema describing the relationship between different communities in space in the artificial world of the simulation model.

However, the mere act of ontologically defining a model in terms of its micro-scale interactions in a single timestep does not mean that one has a complete understanding of how the process will unfold over time and in the macro scale, and it is impossible to know in advance whether a complex simulation model can reproduce the empirical pattern in a satisfactory manner (no matter how well-supported the assumptions underpinning it are) without actually running the simulations. This is, fundamentally, what makes running simulation models so useful in the first place — they enable a deeper understanding of systems evolving over long time periods (hundreds or even thousands of years) even if one only has information on events and processes unfolding over shorter timescales.

To fully understand the outcomes generated from the model, it is useful to adopt the previously discussed strategy of starting from the simplest possible model and adding in complexity only when it is necessary — i.e., when the simple model does not produce outputs that are considered similar enough to the observed pattern. This requires experimentation on two aspects: first, the formulation of the model — the mathematical formula defining the process occurring at every timestep and the environment in which this process takes place - and second, the optimal range for each parameter in order for the simulation to generate informative outputs (this process is sometimes referred to as "sensitivity analysis" (Popper and Pichler 2015: 93)). Because the interaction of different parameters can create a large number of unique combinations of parameter values, the process of disentangling what the effect of each parameter is can become very complicated very quickly as soon as one increases the number of parameters or the complexity of the environment of the model. Ideally, one would first want to see what type of process the model generates in the simplest environment possible, where it is easier to identify the effect of each individual element/parameter.

Experiments with simple models can also be useful during "model verification" (Popper and Pichler 2015: 91), an important step in the process of designing a simulation model which ensures that the output of the simulation model is what was expected initially, when the modeler started developing the mathematical formula or the algorithm of the model. Of course, it is possible for simulation outputs to be counter-intuitive (for example the some of the outputs of the simulation by Crema 2015: 187), and this can potentially be informative, but it is necessary to rule out the possibility that those unexpected results were due to some internal inconsistency or a coding error. Obviously, the more complex the simulation models become, the more probable it is that such coding errors (sometimes called "bugs") will occur, and various stages of testing were required during the development of the more complex models in subsequent chapters to ensure that such errors were identified and corrected – elements of the model verification process for the simulation models are described in greater detail in Appendix 2. The main advantage conferred by using simplified models in the development and testing stage of a simulation model is that by reducing the model to a small number of elements in a simple environment, it is possible to manually calculate what the simulation outputs should be. Ultimately, verification using such simple models renders the simulation process and the entire modelling exercise less opaque, allowing one to place their trust in the outputs of the simulation and the understanding that they can provide in explaining the phenomenon of interest.

A good way to proceed with such experimentations is to construct ‘toy’ models — very simple models that fit certain aspects of the system one wants to capture but simplify or modify other aspects of reality to study the effects of the aspects of interest in isolation. It should be noted that the term ‘toy model’ can be used to subjectively describe any model that one considers to be abstract, simple, or unrealistic enough. As discussed in previous chapters, all models are abstractions of reality, and this certainly applies to the more ‘realistic’ models that will be explored in the following chapters. Therefore, the term toy model is used here to denote the lower end of a spectrum ranging from model simplicity to complexity.

In this chapter, I will explore three toy models that are simplified either in the features of the environment or in the values of the parameters of the simulation model — for example, running the model using a very small population of communities (captured by the parameter N). Each toy model presented here is a simplified version of the base model described in the previous chapter that has been designed to serve a specific function in the process of model verification and sensitivity analysis. Because of their generality, the models can also be used to inform ‘expectations’ on the spatial patterns generated by certain empirical cultural transmission processes (for example, see section 6.3).

In section 7.1, I will explore the simplest toy model of the three (Toy Model 1) which explores what impacts population size, the rate of local adoption, and the penalty for crossing large distances have on system-wide trait adoption without explicitly taking into consideration local variations in the suitability of the environment for habitation (something which we might expect in real life). This is accomplished by running Toy Model 1 in an environment where communities (nodes) and the connections between them (edges) comprise a fully connected network, where all nodes are placed on a circle in roughly regular intervals and each node is connected to every other node in the network. The rationale for this is discussed at length below, but basically this provides an abstract environment which is very roughly reminiscent of the shape of the coastline around the Aegean albeit one that is nowhere near a 'realistic' landscape. In one experiment, I will only increase the number of nodes (N), keeping the edge weights between them equal, in order to explore the effect of increasing the number of communities in isolation (without taking into consideration variations in other parameters). In another experiment, I will assign edge weights whose values correspond to some degree of geographic distance between different communities.

But what about the locations of communities with respect to that of other communities? In reality we might expect some communities to belong in clusters and others to be relatively more isolated, either due to an irregular distribution of resources in the environment (a first-order effect) or due to a process of attraction/dispersal between communities (a second-order effect). Such patterns can lead to some communities being relatively more 'central' than others, and we might expect this to affect how quickly a trait gets adopted by all the communities in the network, especially when the community of origin is part of a local cluster (associated with faster initial adoption), or when it is relatively isolated from all other communities in the network (associated with slower initial adoption). These dynamics are explored in section 7.2, where I conduct experiments on Toy Model

2. This model is like Toy Model 1 in all aspects apart from the location of each community on the circle, which is random, resulting in relative local clustering in some regions in the network. This also introduces an element of stochasticity in the simulations (unlike Toy Model 2 which was purely deterministic): the location of communities is random. This means that, purely by chance, running the simulation using a different seed for R's random number generator could generate different outputs. It is reasonable to assume that the system-wide outputs of simulations with smaller numbers of nodes will be more sensitive to the effects of stochasticity - meaning that running a simulation with small N with a number of different seeds will produce more variable outputs than running a simulation with large N with different seeds.

In section 7.3. I will explore a toy model that incorporates the aspect of population density more explicitly and assumes that the location of each community is purely random. This is, again, a stochastic model that generates slightly different (but not significantly different system-wide) results when using different seeds for the random number generator. The main difference to Toy Model 2 is that this time the model runs on a less abstract environment — a square surface rather than a circle- and with larger N values.

Finally, section 7.4. explores a model where every community loses a random proportion of the trait at each timestep in order to assess the impact of random drift on the cultural transmission process.

7. 1. Toy model 1: Fully connected network of communities located on a circle at regular intervals, varying the size of the population size N, the penalty for crossing "large" distances *k*, and the rate of local adoption *a*

Introduction

As discussed previously, exploring a simulation model involves experimenting with simulations using different values for the parameters of the model to observe the effect that changes in parameter values have on the output of each simulation — here, the main parameters are the size of the population of communities N, the penalty for crossing "large" distances *k*, and the rate of local adoption *a*.

This process is sometimes called "sensitivity analysis" and it can take place in the context of running simulations for the more 'realistic' model described in Chapter 6. However, the advantage of conducting a preliminary sensitivity analysis using a toy model in advance of a sensitivity analysis on the "realistic" model is that one can build a toy model that eliminates the effects of the environment, allowing one to isolate the effects of parameter changes without geographic 'noise' distorting the picture (in the case of a geographically grounded model).

Therefore, the toy model I will explore in this section is one where instead of using a realistic-looking environment like a map of the Aegean to represent the environment, I will use an abstract space where communities are located on a circle and in roughly regular distances to their immediate neighbours as the environment. In keeping with the ontological definition of the model as discussed in chapter 6 and with the assumption that all communities have some effect on each other's trait proportion value, in this toy model every community is connected to all other communities in a fully connected network. We already know that the number of connections/edges between nodes in a fully connected network can become very large the more one increases the number of nodes (communities in this model) — if the number of nodes is N, the number of unique edges (only one edge per pair of nodes) to be $N(N-1)$ if we don't allow connections between each node and itself, and N^2 if we do. One might also expect that increasing the number of communities in the model will lead the trait proportion value at every node as well as the overall mean trait proportion of all communities to increase faster over time as well. Experiment 1 involves increasing the number of communities (population size N) without taking into consideration the distance between communities – this is accomplished by setting the penalty for crossing "large" distances $k = 0$, which ensures that all edges/connections are treated as equal in importance, with an edge weight equal to 1 (since $e^0 = 1$). Experiment 2 involves increasing the number of communities while considering the effect of their distance in space by setting $k = 0.5$, which results in each edge having a different weight – in this case, edges connecting more distant communities having a lower weight than those connecting communities located near each other.

Finally, in Experiments 3 and 4 I will explore the effects of increasing the penalty for crossing large distances and the rate of local adoption respectively, whilst keeping the number of communities constant at N = 20 and the initial trait proportion value for the "origin" community/node at 0.1.

Experiments on Toy Model 1

7. 1. 1. Experiment 1: k=0, a = 1, origin = 0.000000001, observed over 10 Timesteps, varying N

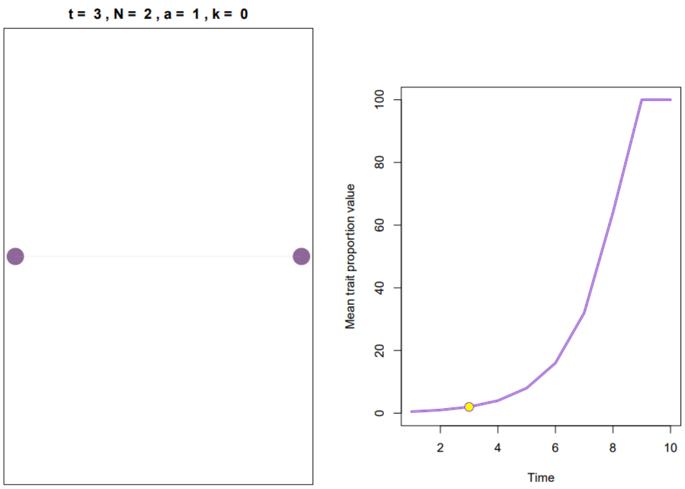

20.1.

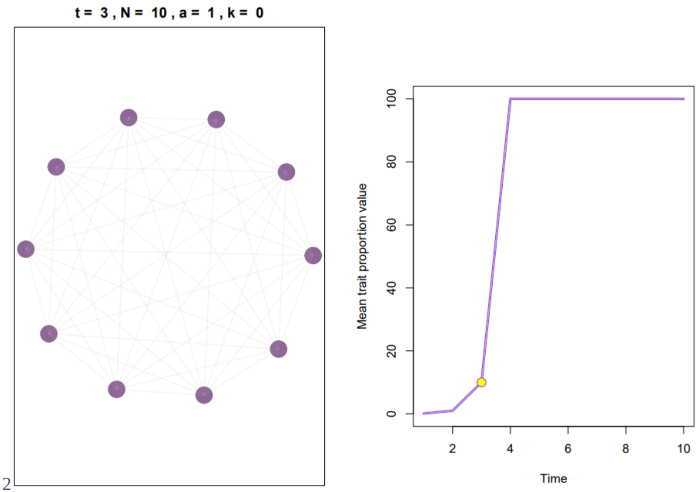

20. 2

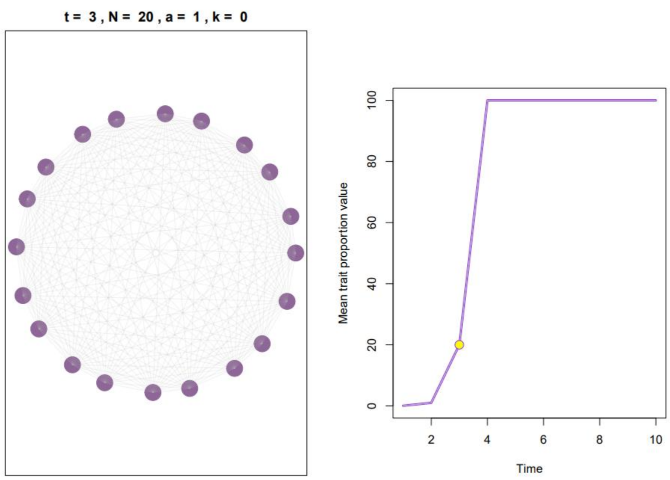

20.3.

Figure 20: Outputs of four simulations, each with a different value for the population of communities/nodes N, and the same value for the rate of local adoption a = 1, the "penalty for crossing large distances" parameter k = 0, and the value at the origin community/node at Timestep 1, which is equal to 0. 000000001 (a very low value that ensures that the mean trait proportion does not reach the 100 % equilibrium too rapidly). In all cases the system is shown at timestep 3 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

Results from Experiment 1

Figure 20 shows the simulation outputs resulting from the first experiment using Toy Model 1. The simulation output for a very small case with a population N of only two communities (Figure 20.1) suggests that, under this model, the mean trait proportion of the system increases non-linearly over time and resembles an S-shaped curve. If one increases the number of communities/ nodes from two to ten (Figure 20.2) this effect becomes even more pronounced, and the system reaches the 100 % equilibrium (the maximum value mean trait proportion can take and given that the model does not include trait loss, the value at which the system remains stable) in Timestep 4 rather than Timestep 8 as in the case with two communities. Finally, running the simulation with a population of twenty communities (Figure 20.3) does not change the timestep when the system reaches equilibrium compared to the case with ten communities. However, the rate at which the mean trait proportion increases between Timesteps 2 and 3 is noticeably higher in the case with N = 20 compared to N = 10. In all cases, however, the trajectory is always non-linear and an increase in the number of communities results in "speeding-up" the trait transmission process, as expected.

7. 1. 2. Experiment 2: *k* = 0.5, *a* = 1, origin = 0.1, observed over 100 Timesteps, varying N

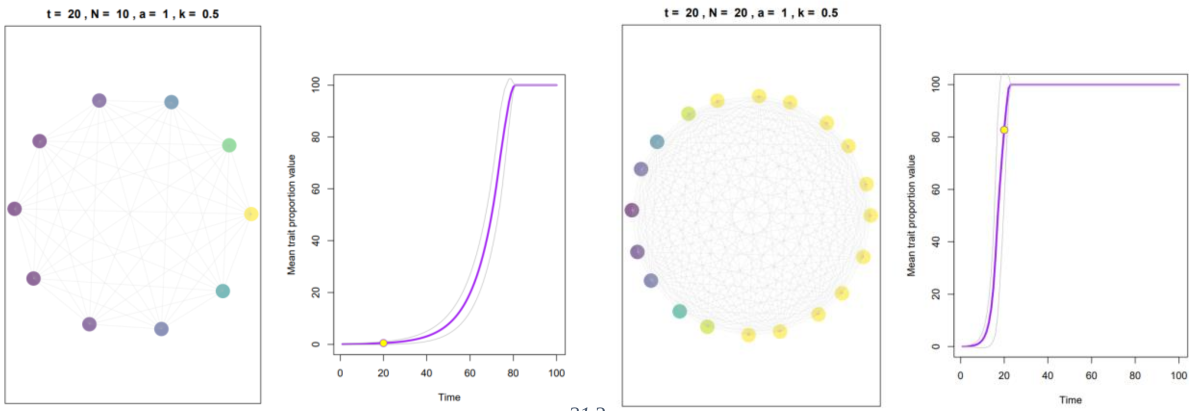

21.1. *21.2.*

Figure 21: Outputs of two simulations, each with a different value for the population of communities/nodes N, and the same value for the rate of local adoption a = 1 and the "penalty for crossing large distances" parameter k = 0.5, and the value at the origin community/node at Timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 20 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

Results from Experiment 2

The simulation outputs shown in Figure 21 demonstrate the effects of increasing the number of communities in the transmission network while factoring in the effects of their distance in space – note that in the previous experiment, the edge weights for all edges were equal to 1, resulting in the model disregarding the effects of distance between communities. As in the previous experiment, all trajectories are non-linear and resemble S-shaped curves.

7. 1. 3. Experiment 3: N=20, origin = 0.1, observed over 100 Timesteps, *a* = 0.1, varying *k*

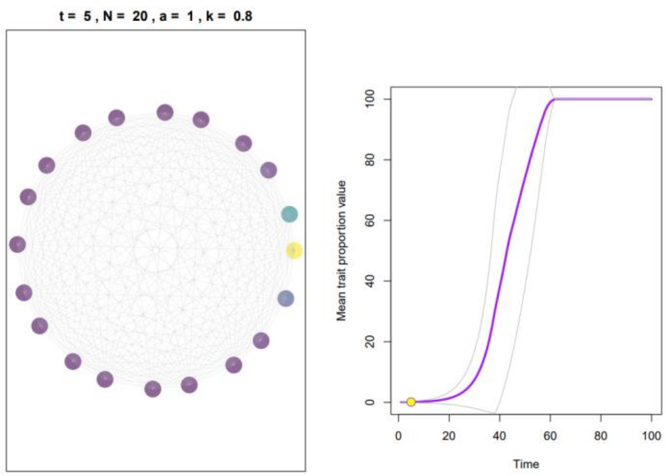

22.1.

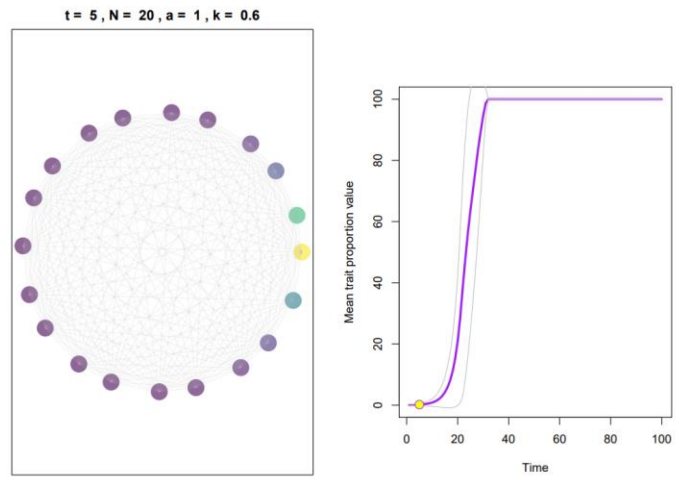

22.2.

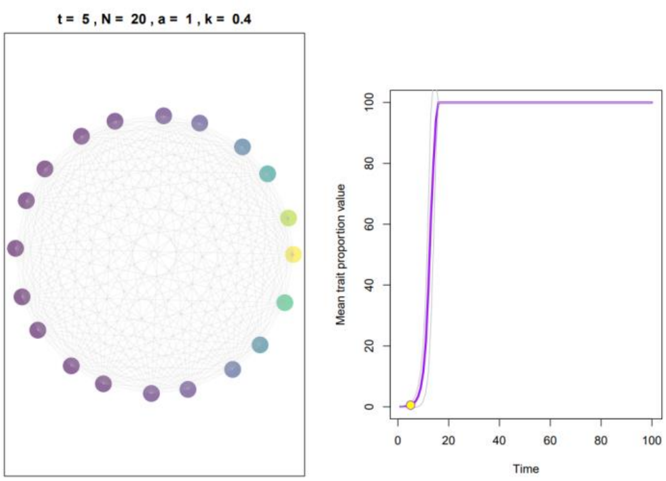

22.3.

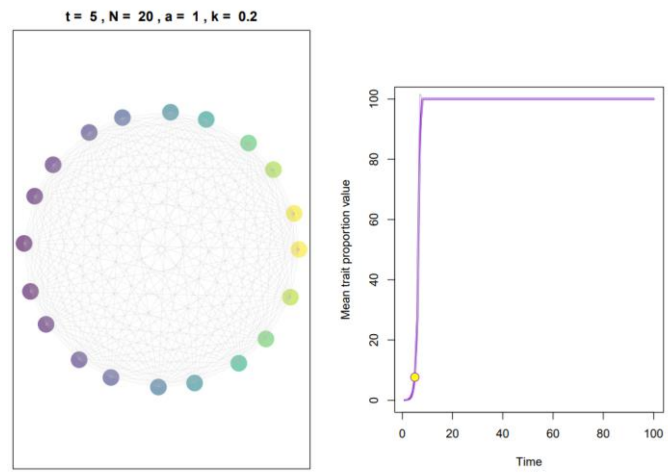

22.4.

Figure 22: Outputs of four simulations, each with a different value for the "penalty for crossing large distances" parameter (k), the same value for the rate of local adoption a = 1, the population of communities/nodes N = 20, and the value at the origin community/node at Timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 5 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

Results from Experiment 3

The outputs of Experiment 3, which involved running four simulations with different values for the parameter *k* — which reflects the magnitude of the non-linear communication decay between different communities/nodes - suggest that increasing *k* results in S-shaped trajectories where the mean trait proportion value reaches the 100 % equilibrium later in time, as shown in Figure 22. The overall shape and steepness of the curve does not change noticeably — the most pronounced effect is that the initial period of slow trait transmission is significantly prolonged with higher values of *k* (note, for example Figure 22.1, with *k* = 0.8).

7. 1. 4. Experiment 4: N=20, origin = 0.1, observed over 100 Timesteps, *k* = 0.5, varying *a*

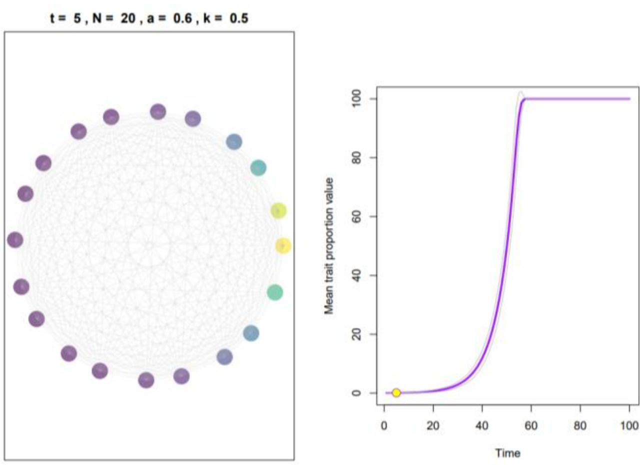

23.1.

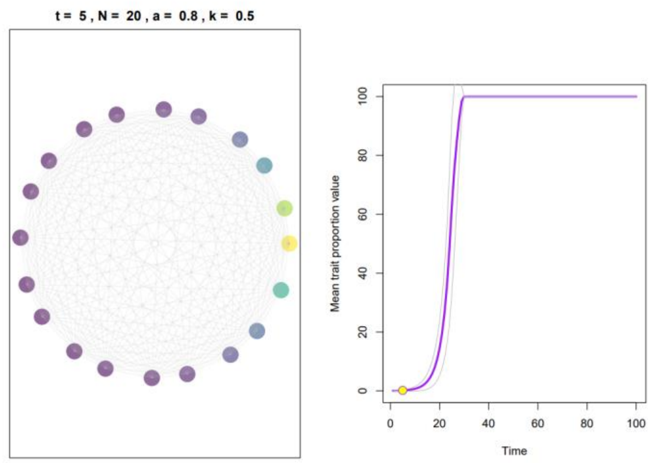

23.2.

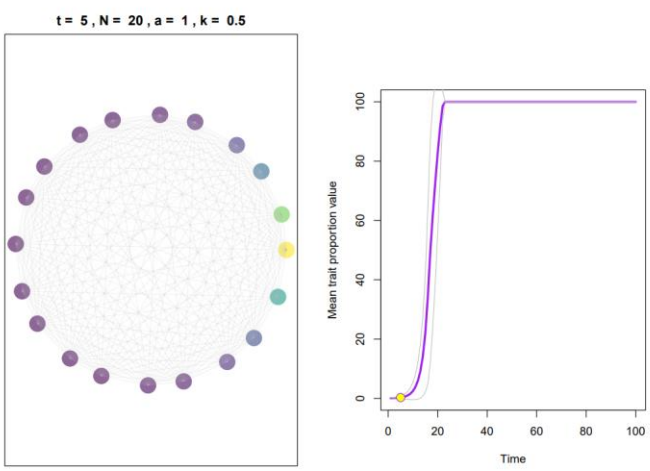

23.3.

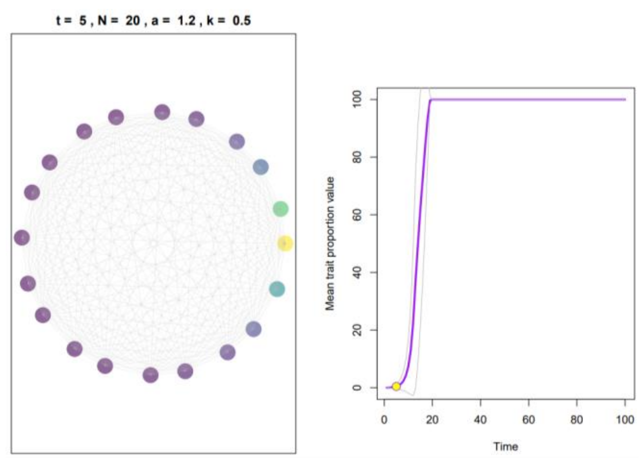

23.4.

Figure 23: Outputs of four simulations, each with a different value for the rate of local adoption a (0.6, 0.8, 1, and 1.2) and the same value for the "penalty for crossing large distances" parameter k = 0.5, the population of communities/nodes N = 20, and the value at the origin community/node at Timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 5 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

Results from Experiment 4

A similar picture to the one from Experiment 3 emerges from Experiment 4, which investigated the effect of increasing the rate of local adoption *a* (Figure 23). The simulated trajectories of the mean trait proportion value over time resemble steep S-shaped curves and increasing *a* results in the initial period of slow adoption being prolonged rather than the slope of the curve becoming less steep. However, the region between $a = 0.6$ (Figure 23.1) and $a = 0.8$ (Figure 23.2) produces a more pronounced effect on the simulated curve than the region between $a = 0.8$ (Figure 23.2) and $a = 1$ (Figure 23.3). This suggests that there is non-linearity in how the values of the local adoption rate parameter (*a*) affect the overall system, and that increasing the value of this parameter by a certain value does not always result in a proportional effect in trait adoption trajectories.

Finally, a key takeaway from all experiments from Toy Model 1 is that *a* and *k* have opposite effects – increasing *a* speeds up the adoption process while increasing *k* slows it down — this is an effect consistent with the formulation of the model, where we might expect higher communication costs associated with *k* to result in slower system-wide adoption.

7. 2. Toy Model 2: Fully connected network of communities located on a circle in random intervals, varying the size of the population size N, the penalty for crossing "large" distances *k*, and the rate of local adoption *a*

Introduction

Having explored the effects of increasing the values of the main model parameters in a very abstract environment (a configuration of a small number of communities/nodes in regular intervals around a circle) it is worth considering whether a more "realistic" environment would lead to different trait adoption trajectories and whether the effects of *N*, *k*, and *a* would be different in a slightly different environment.

But what defines a configuration of communities in space as 'more realistic' in this case? In the absence of environmental factors that could lead to different areas being preferable for habitation, as well as factors determining the relative position of each community with respect to its neighbours, leading to community dispersal (e.g., competition over resources) (Bevan, 2020: 71) or clustering (e.g., cooperation), one might expect the configuration of communities in space to be random. Realistic settlement patterns are characterised by the existence of clusters at local scales, as well as sub-regions with relatively dispersed and isolated communities, and random settlement patterns manage to capture this element of inhomogeneity well, which makes them a good solution when the goal is to simulate a realistic settlement pattern in the absence of empirical information about locational preferences.

In this section I will explore a toy model (Toy Model 2) which is like Toy Model 1 in all aspects apart from the relative position of nodes, which will be random rather than regularly spread over the environment. As shown in Figure 24, this leads to some areas on the circle having clusters of communities, and others to be relatively less densely populated. This has interesting implications regarding the interaction of different parameters. For example, a scenario where the penalty for crossing large distances, k is relatively high, could lead to regions with local clusters increasing their trait proportion values relatively faster compared to regions with a lower population of communities. However, we should expect this effect to become weaker if a high k is combined with a high value for the rate of local adoption a; since each community depends less on the impact of its neighbours to increase its trait proportion value, this should make clusters (that mitigate the effects of high communication costs associated with high k values) less important.

In the following sections I will explore Toy Model 2 by varying the population size N and parameters k and a. Given that this toy model was designed to study the effects of community/node location on the outputs of the simulation, I will not explore the case where $k = 0$, as this would lead to completely disregarding the effects of the relative distance between communities/nodes, rendering the entire exercise pointless.

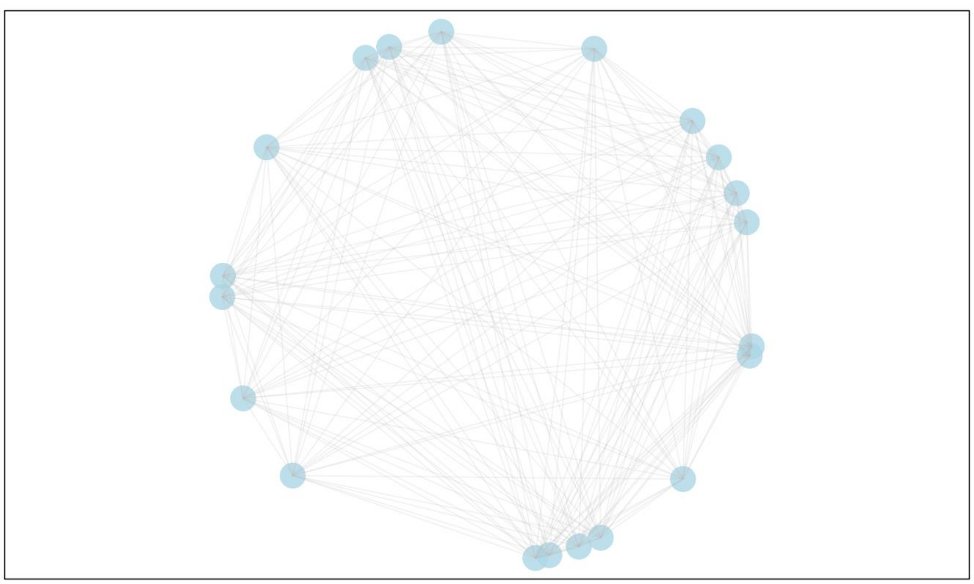

Figure 24. The configuration of communities for Toy Model 2, using a population of twenty communities (N = 20). Note that, even in a purely random pattern, there are regions on the circle with a higher node density compared to others.

Experiments on Toy Model 2

7. 2. 1. Experiment 1: k = 0.5, a = 1, varying N, origin = 0.1, observed over 100 Timesteps

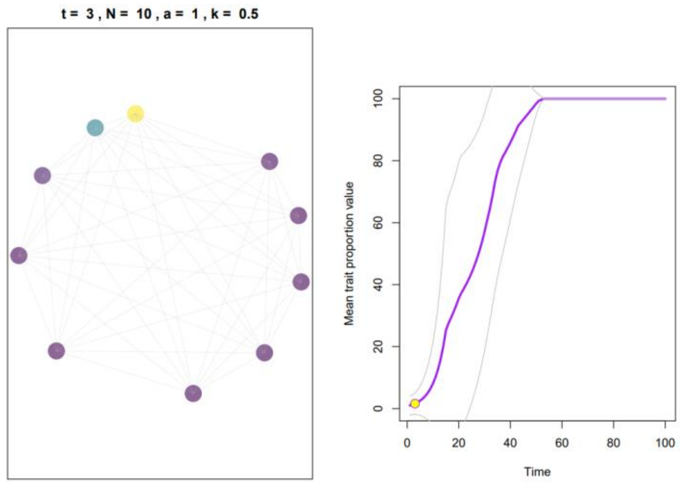

25.1.

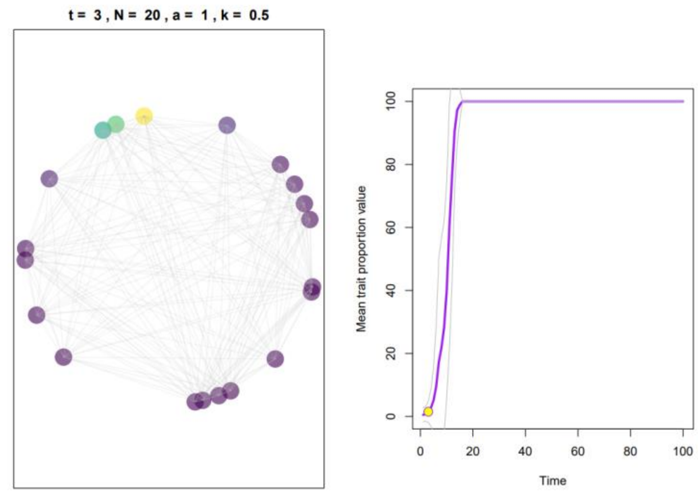

25.2.

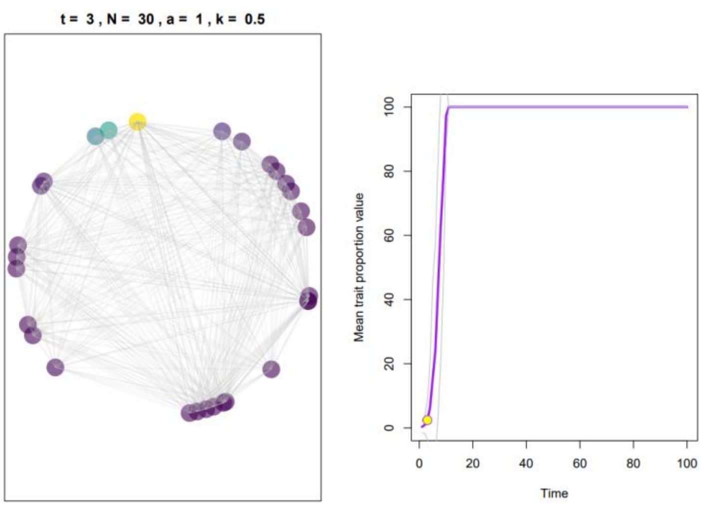

25.3.

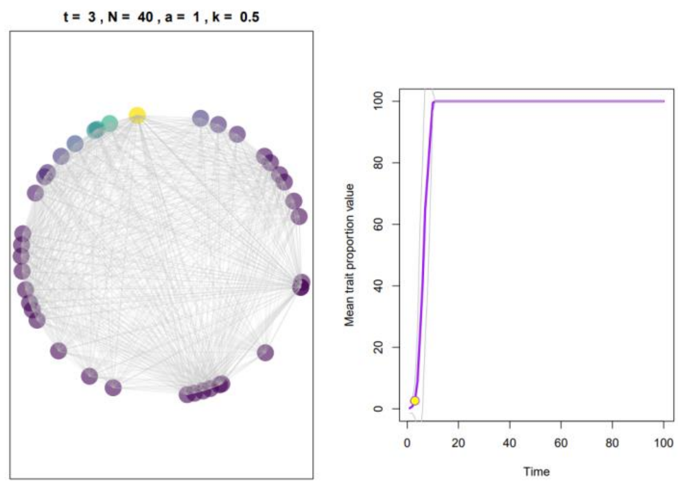

25.4.

Figure 21: Outputs of four simulations of trait transmission in a fully connected network of communities located randomly on a circle. Each simulation uses a different value for the population of communities/nodes N, and the same value for the rate of local adoption a = 1, the "penalty for crossing large distances" parameter k = 0.5, and the value at the origin community/node at Timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 3 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

7. 2. 2. Experiment 2: N=20, *a* = 1, varying *k*, origin = 0.1, observed over 100 Timesteps

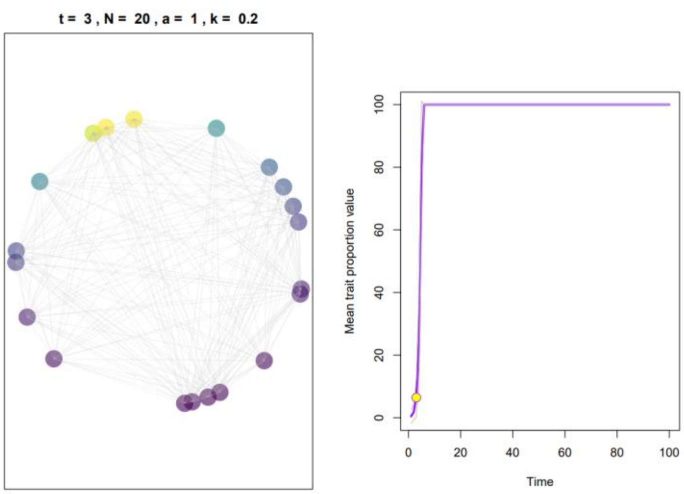

26.1.

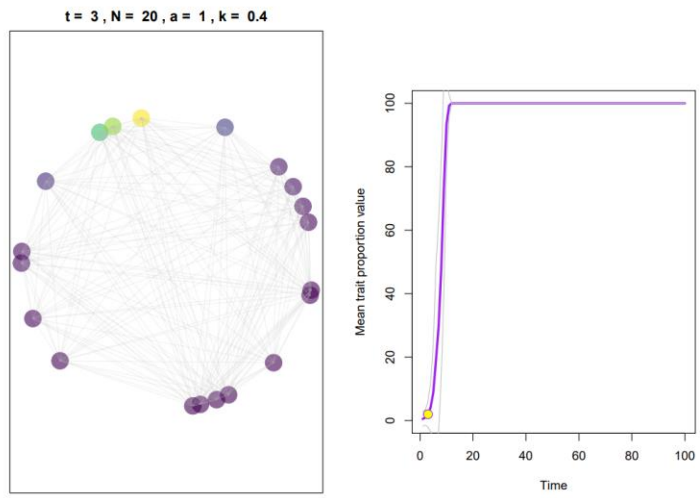

26.2.

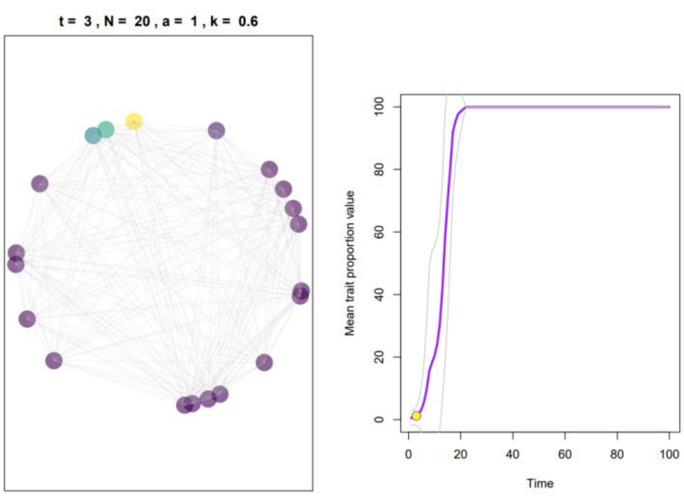

26.3.

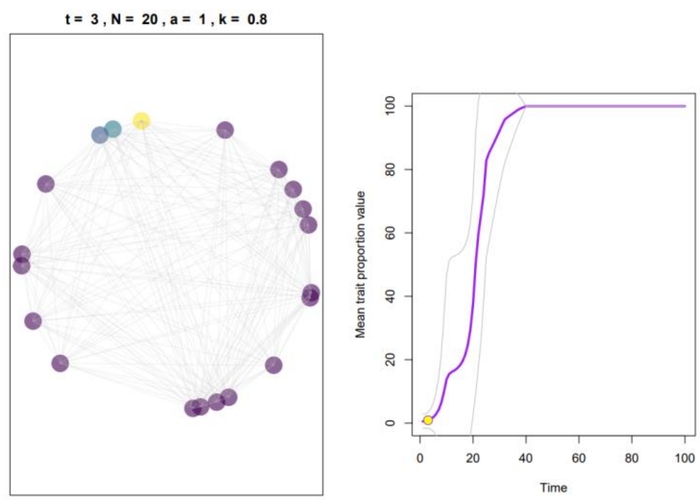

26.4.

Figure 26: Outputs of four simulations of trait transmission in a fully connected network of communities located randomly on a circle. Each simulation uses a different value for the "penalty for crossing large distances" parameter (k), and the same value for the rate of local adoption a = 1, the population of communities/nodes N = 20, and the value at the origin community/node at Timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 3 on the left panel, and the right panel indicates the mean trait proportion value at that timestep, indicated as a red dot plotted on top of the trend over time (shown as a purple line). At timestep 3, the cluster which includes the community that is "origin" of the trait has a higher trait proportion value relative to the values of all other communities in the network. Grey lines define an envelope one standard deviation above and below the mean.

7. 2. 3. Experiment 3: N=20, *k* = 0.5, varying *a*, origin = 0.1, observed over 100 Timesteps

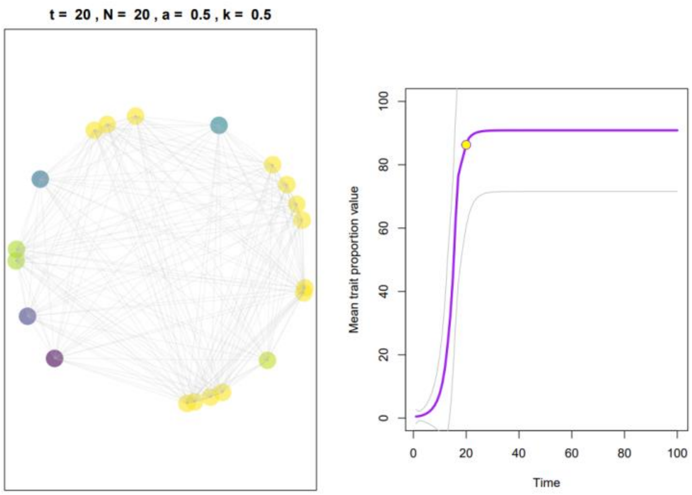

27.1.

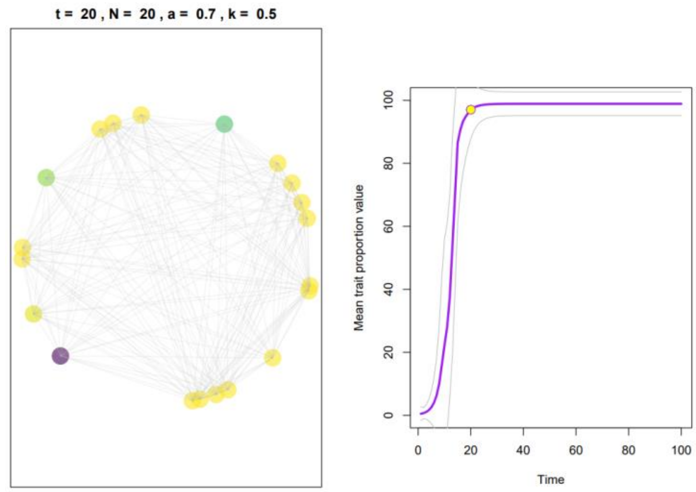

27.2.

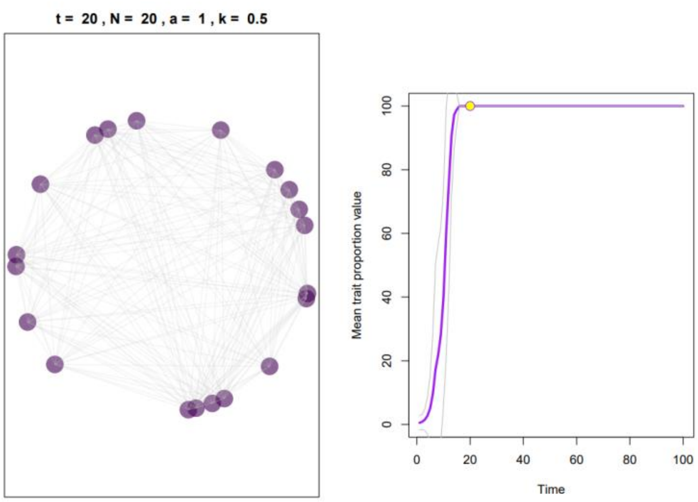

27.3.

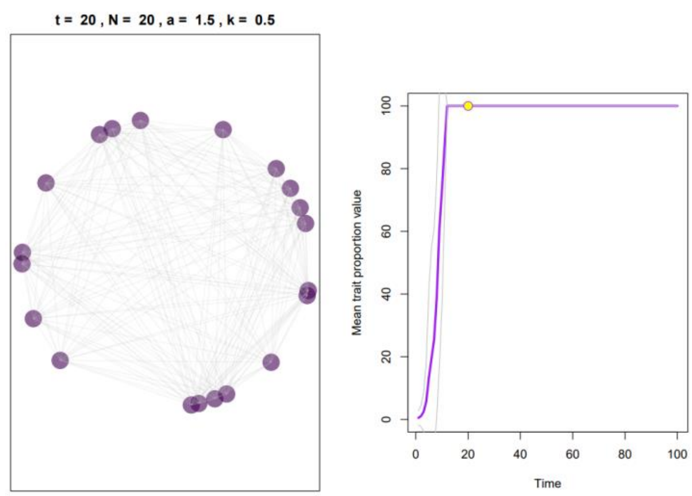

27.4.

Figure 27: Outputs of four simulations of trait transmission in a fully connected network of communities located randomly on a circle. Each simulation uses a different value for the rate of local adoption a (0.5, 0.7, 1, and 1.5) and the same value for the "penalty for crossing large distances" parameter k = 0.5, the population of communities/nodes N, and the value at the origin community/node at timestep 1, which is equal to 0.1. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. In all cases the system is shown at timestep 20 on the left panel, and the right panel shows the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). At timestep 20, the cluster which includes the community that is "origin" of the trait has a higher trait proportion value relative to the values of all other communities in the network in the simulations where a is low. When a is high, all nodes have a trait proportion value equal to 100 %. Grey lines define an envelope one standard deviation above and below the mean.

Results from Experiments 1, 2, and 3

The outputs of experiments 1, 2, and 3 confirm that varying the population size N (Figure 25) and the parameters *k* (Figure 26) and *a* (Figure 27) can affect trait proportion trajectories in that the initial period of slow adoption becomes shorter, as was the case for Toy Model 1

However, the presence of local clusters and regions with isolated nodes can sometimes lead to the model exhibiting qualitatively different behaviours to Toy Model 2 (simulations with the same parameter values). For instance, if one compares the top left panels of Figure 4 and Figure 27, which include simulations for Toy Model 1 and Toy Model 2 with the same parameter values, it becomes clear that the presence of a cluster near the origin node in the case of Toy Model 2 leads to the trait reaching almost 100 % adoption near timestep 25, whereas in the case of Toy Model 1, the trait barely exceeds 10 % at the end of the 100 Timestep observation window. This suggests that the presence of local clusters can lead to faster propagation of the trait across the entire network, even when the rate of local adoption is low (a = 0.5). In this example, the penalty for crossing large distances is relatively neither high nor low (k = 0.5), which effectively means that after the trait has been well-established in the cluster, it is relatively easier for it to "jump" to more isolated nodes in the network compared to a scenario with higher *k*. It should be noted that in some cases where the presence of relative clustering results in fast uptake of the trait early, it is also possible to see relatively slower adoption at the latest stages of the process– shown as a more pronounced shoulder in the adoption curve (e.g., Figure 27, bottom right panel) — when the only communities/nodes left to adopt are the most isolated in the network.

The bottom right panel of Figure 26 is particularly illuminating regarding the importance of *k* in this toy model. Increasing *k* to 0.8 results in an almost double-S shaped trait adoption, which is possibly due to the effect of the trait becoming established in the cluster around the origin relatively quickly (the first “exponential” phase of the trend), which is then followed by a delay when the further propagation of the trait depends on a hop from the cluster to the more isolated communities to the left and to the right of the cluster (the hop will become more expensive if *k* increases). Figure 28 shows the network at four timesteps t = 10, 20, 30, and 40 and indicates that the first ‘exponential’ phase at timestep 10 (Figure 28.1) is indeed associated with an episode of fast transmission in the cluster that the origin node belongs to. This is followed by an episode where the rate at which the mean trait proportion increases slows down (timestep 20, Figure 28.2) which is finally followed by a second exponential phase (timestep 30, Figure 28.3) which levels off at the stage when all nodes apart from the most isolated ones have adopted the trait.

Overall, the experiments on Toy Model 2 suggest that relative variation in the density of communities over the environment can generate slightly different trait transmission trajectories compared to a more regular distribution of communities (Toy Model 1) and have interesting effects when combined with certain values for the parameters of the model. For example, the fact that the origin community/node is part of a local cluster combined with a certain value for the parameter that determines the cost of longer distances over shorter ones, *k*, affects qualitative aspects like the shape of the curve (a double S- curve rather than a single S-curve). The presence of a cluster near the origin community/node also leads to almost complete innovation adoption in cases where the local rate of adoption *a* is low (Toy Model 2), something that would not have happened in the case of a more homogeneous distribution of communities (Toy Model 1).

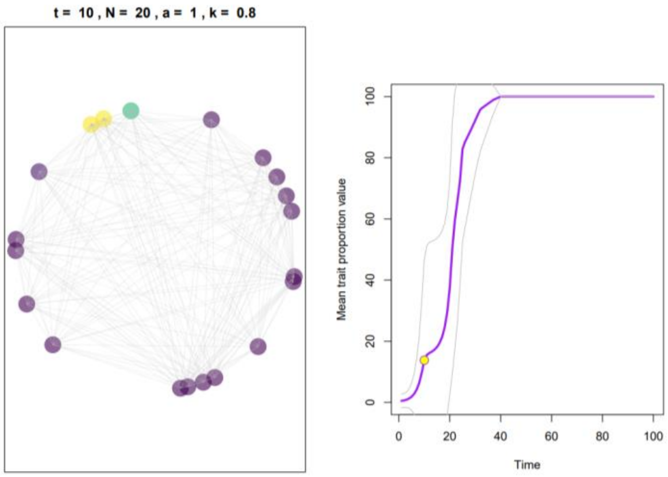

28.1.

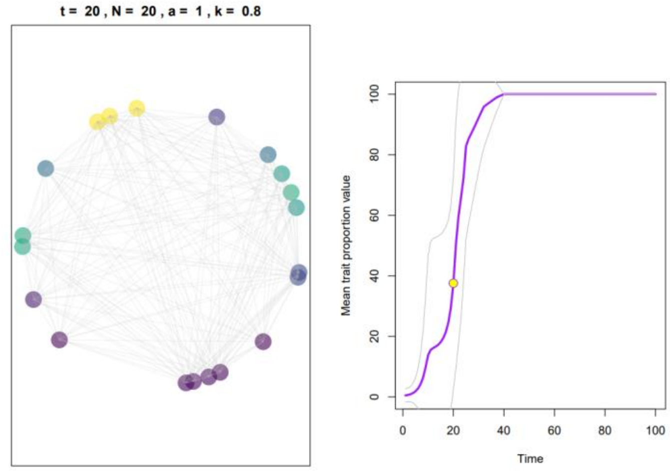

28.2.

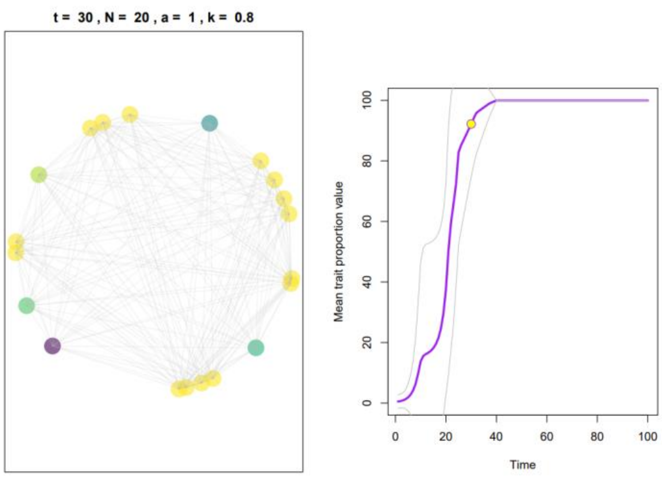

28.3.

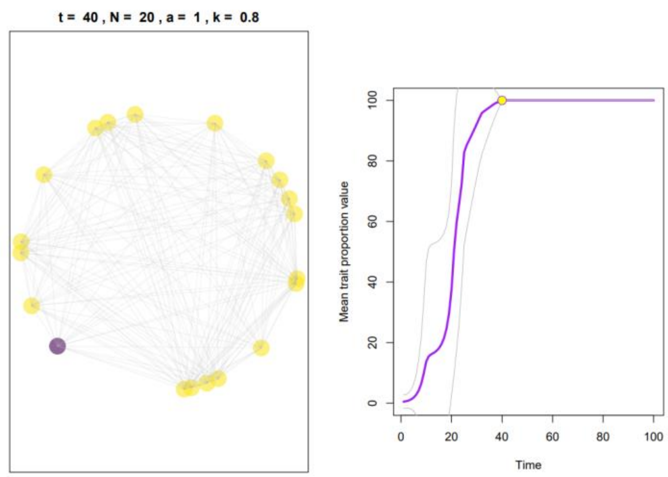

28.4.

Figure 28: The outputs of one simulation with N = 20, k = 0.8, a = 1, and with the origin node's value at timestep 1 = 0.1. The left-hand sides of the four panels show the network at timestep 10, 20, 30, and 40. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. The right-hand side of each panel shows the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean.

7.3. Toy Model 3: Fully connected network of randomly located communities on a square, varying the size of the population N, the penalty for crossing “large” distances *k*, and the rate of local adoption *a*

Introduction

As discussed previously, in the absence of other information about the location of communities in a landscape — like the presence of natural barriers like the sea, the existence of processes of attraction-dispersal between communities, etc. — the most reasonable assumption is to consider the location of communities in space to be random. For the previous model I used a highly abstract environment to represent a landscape consisting of a sea (the Aegean) surrounded by land (Mainland Greece and Western Anatolia). However, a significant part of the trait adoption process takes place on the main landmasses, where one might expect the existence of patches of relatively lower density of communities over others, but not as dramatic as the one imposed by the environment in Toy Model 2.

That is not to say that the outputs of Toy Model 2 were not informative. They suggested that the presence of areas of relative higher node density can, even in a purely random settlement pattern, function as “clusters” that drive faster innovation adoption. The presence of such relative clustering sometimes led to qualitatively different adoption curves (double S-shaped over single S-shapes ones) of the mean trait proportion value over time. This effect was stronger when the parameter associated with the penalty for crossing large distances *k* was higher and when the population size N (number of communities/nodes) was lower. The former makes intuitive sense; making crossing large distances very expensive makes the overall system reliant on areas with a relatively large number of nodes that are separated by short distances. The latter is possibly due to the fact that decreasing the population size N results in decreasing overall node density when the environment is a finite area in space. Again, it makes sense that the presence of relative clustering drives faster trait system-wide adoption in generally less densely populated areas more than it does in very densely populated ones. In the case of the former, a scenario where a trait starts from an origin node located in a cluster unfolds like so: the trait spreads relatively faster in the cluster and then its spread gets delayed when the trait crosses the large distance separating the cluster and the other, more isolated communities in the environment. In the case of the latter, however, the delay caused by the need to cross the distance from the cluster to the isolated communities would be shorter, because the distance between the cluster and its neighbours tends to be smaller — and given that in this model, the relationship between distance and communication decay is non-linear (with its slope dependent on *k*), we should expect even a small difference in the distance between a cluster and its neighbours to have a relatively large effect when *k* is high.

This points to a relationship between the relative density of communities in the environment, and *k*. Specifically, it points to the importance of the relationship between the overall population density and the existence of relative clustering as an effect of increasing communication costs captured by *k*. In order to better explore this relationship, I will conduct an experiment on a new toy model, Toy Model 3, which will incorporate community/node density in a more explicit manner by defining the environment as a square surface with a side equal to 100 units in length and an area equal to 10000 square units. Node density is then defined as the ratio between the number of communities/nodes N and the area of that surface. This way, a population size of N = 200 communities/nodes corresponds to node density equal to 200 communities per 10000 square units of length.

This approach allows producing experimental outputs that can be compared to the outputs of simulations for the more "realistic" models explored in the following chapters in relation to the effects of population density and the penalty for crossing large distances.

The experiment explores the conditions under which the diffusion wave that forms over the landscape propagates as a 'radial' wavefront (a case where trait transmission resembles a wave expanding radially from the origin), as a 'path-driven' wavefront (a main wavefront that propagates in the direction of 'paths' formed by areas of higher node density (such paths emerge purely randomly) as well as the conditions under which the wavefront is 'cluster-driven' in that it advances via clusters located near but not adjacent to the wavefront — a more robust definition of wavefronts will be provided below. Given that this sort of qualitative characteristic is difficult to describe in quantitative terms, I observed how each simulation unfolded over time and assigned the qualitative labels "radial", "path-driven", and "cluster-driven" to each simulation.

This means that there is some degree of subjectivity in my assessments that should be acknowledged, particularly in cases where a simulation output is 'between' two qualitative states. Such cases were rare and whenever encountered, I assigned a qualitative label that best described the propagation of the wavefront for the longest amount of time in the simulation.

7.3.1. Experiment 1: Radial, path-driven, and cluster-driven wavefronts associated with different combinations of *a*, *k*, and N, origin = 0.01, observed over 100 Timesteps

When the model runs on a 2-dimensional surface, it becomes clear just by looking at the simulation outputs that the trait spreads through the study area by means of a propagating wavefront that starts from the community of origin and dies out when the trait has become established in every other node. But what exactly is a wavefront? Fort and Méndez (2002:897) provide a useful definition of wavefronts in the context of trait diffusion: "wavefronts are special solutions to the dynamics of systems with several equilibrium states ... usually defined as solutions such that they travel with constant shape and connect an unstable (initial) and a stable (final) state. In practice, this simply means that before the propagation of the front the system is in an unstable state, and after the front has swept along the system, its state is stable. The relevant variable depends on the system and process one wishes to describe".

Following their definition, in our system of interest stable states are associated with regions where every node has a trait proportion value equal to 100 % and unstable states are associated with region where nodes have yet to reach the 100 % equilibrium. The spatial process unfolds in time in this manner: before the wavefront reaches a node, its trait proportion value is less than 100 % and as soon as the wavefront reaches the node its trait proportion value increases, changing its previous state. Therefore, its previous state was unstable — given that in this model all nodes will eventually increase their trait proportion value even if they do so very slowly. Note that because I do not allow trait loss in this model, once a node reaches the 100 % state it will never change, since trait loss is not 'allowed', and it is impossible to increase a proportion above 100 %. This means that for each node, the 100 % state is a stable state.

The experiment explored here involved running simulations with the same seed (30) for R's random number generator, with the value at the origin node set at 0.01, and varying N from 200 to 700 by 100, and *k* and *a* from 0.4 to 1.2 by 0.2. Regarding the shapes of the wavefronts produced by each simulation, I was able to identify three main shapes/types of propagation wavefronts: radial, path-driven, and cluster-driven.

Results from Experiment 1

Figures 29.1-29.3. show snapshots of the propagation of a radial wavefront in time: At Timestep 5 (Figure 29.1), the node with the highest value is the origin node, and several nodes located in a circle surrounding the origin also have high values compared to every other node in the area. This reflects a process similar to the effects of isolation-by-distance. After 5 Timesteps (Figure 29.2), the trait has propagated along a radial wave advancing in concentric circles around the origin cluster. By Timestep 15 (Figure 29.3), the trait has been widely adopted in nearly half of the study area, again advancing radially. This type of wavefront propagation fits Fort and Méndez's (2002) definition well; the shape of the wavefront remains constant over time. Radial diffusion wavefronts are well-documented in archaeology in the context of the spread of agriculture in Europe, starting from Ammerman and Cavalli-Sforza's 'wave of advance model' (1971) which, according to Ammerman and Cavalli-Sforza (1979: 275) takes the form of a population wave expanding outward at a steady radial rate.

29.1. 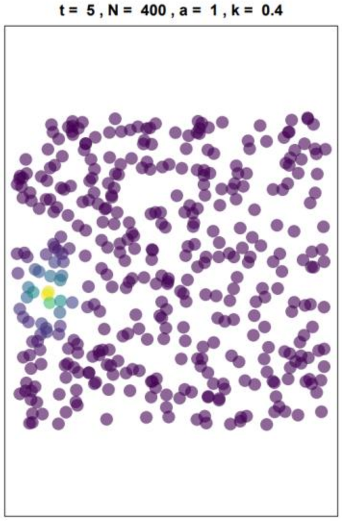

29.2. 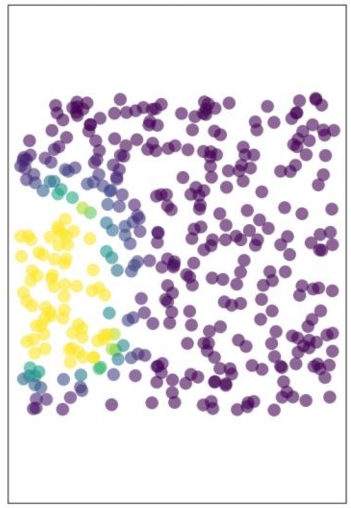

29.3. 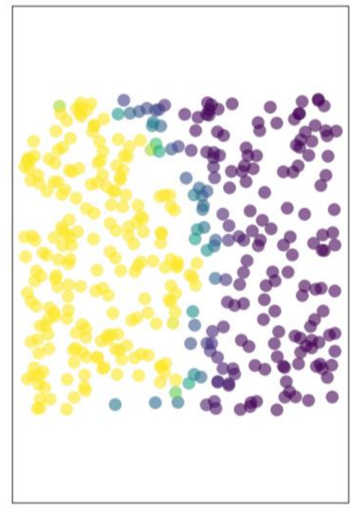

Figure 29: The propagation of a 'radial' trait adoption wavefront in space over time.

In contrast to 'radial' wavefronts, 'path-driven' wavefronts exhibit a different type of behaviour. Figures 30.1-30.3 show the propagation of a 'path-driven' wavefront along the same instantiation of randomly placed nodes as in the previous example (N = 400 and using the same origin community/node). This time, the adoption rate is the same as in the previous case, but the penalty for crossing large distances *k* is higher (1 rather than 0.4). This makes crossing large distances more expensive and leads to the formation of a path-driven wavefront, which propagates along paths where the relative density of nodes is higher compared to other areas in the environment. In the early stages of the process, at Timestep 50 (Figure 30.1), the trait has become established in a cluster near the origin node and is propagating along a high-density path to the south-west of that cluster. At Timestep 70 (Figure 30.2), the path to the south-east has become connected to the main wavefront, which is also propagating along a path to the north and north-east. At Timestep 100 (Figure 30.3), the wavefront is still propagating along paths in the landscape. Note that even though the trait has become established in almost half of the study area, there are still nodes near the origin cluster towards the west side that have relatively low trait proportion values. This is because these nodes are not located on high-density paths and are generally relatively isolated from other nodes.

30.1. 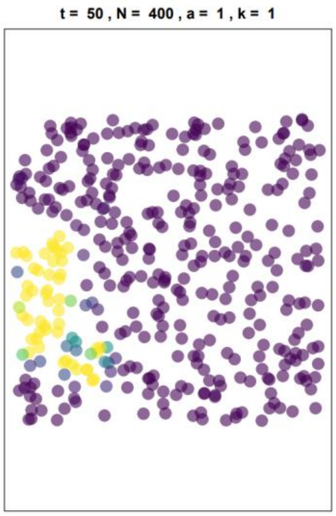

30.2. 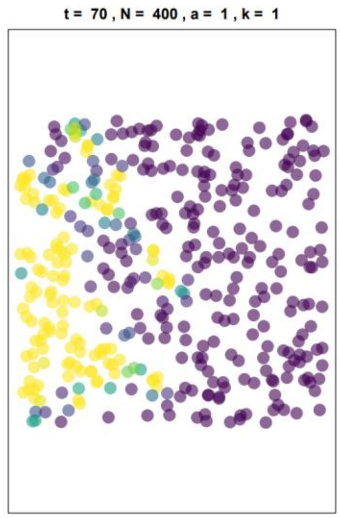

30.3. 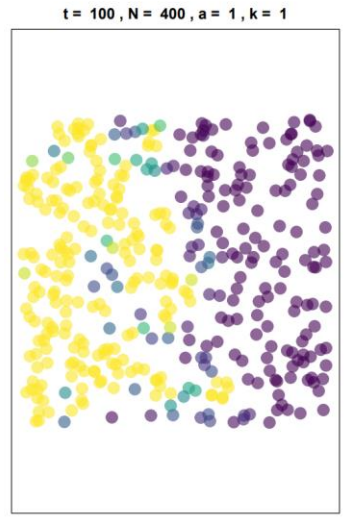

Figure 30: The propagation of a 'path-driven' trait adoption wavefront in space over time.

Finally, increasing *k* and keeping all other parameter values constant leads to the formation of cluster driven wavefronts (Figure 31). Here, a wavefront is formed around the origin node early in the process at Timestep 40 (Figure 31.1). By Timestep 60 (Figure 31.2) this cluster has grown, and a new cluster is forming to its south-east. Finally, by Timestep 80 (Figure 31.3) the two clusters have joined into a single wavefront propagating to the east, and new clusters are forming to the north. This process repeats until the end of the observation window. It is perhaps inaccurate to describe this wavefront as one with constant shape as per Fort and Méndez's (2002) definition — indeed its main characteristic is that it is formed by a single cluster that is then joined with other clusters developing near the wavefront but not adjacent to it. However, the mode of the propagation process *is* constant over time.

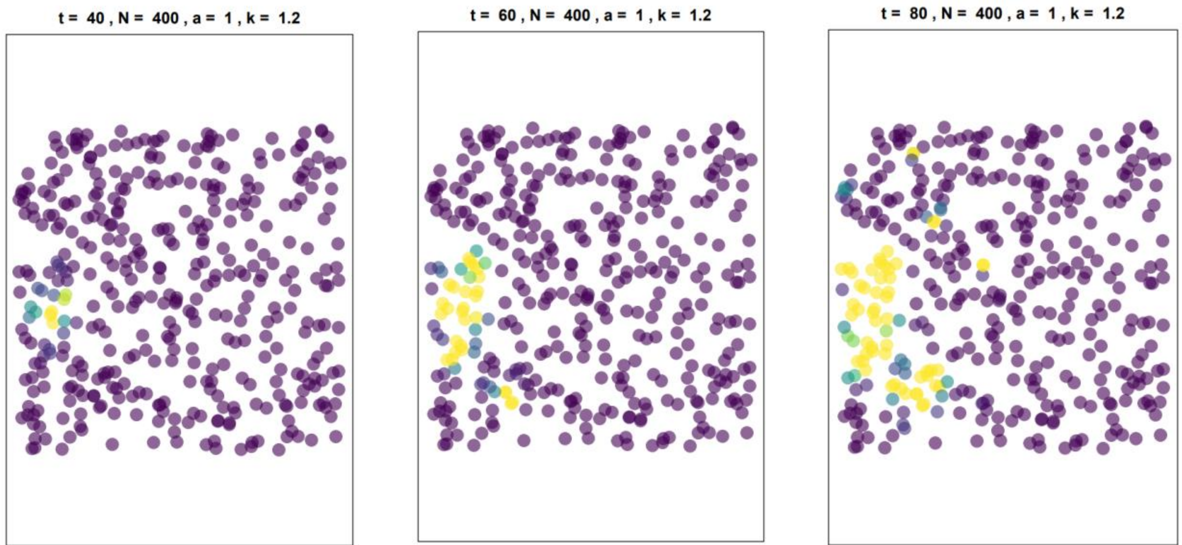

31.1. *31.2.* *31.3.*

Figure 31: The propagation of a 'cluster-driven' trait adoption wavefront in space over time.

An important insight gained from this experiment is that the three wave propagation modes are also associated with trait adoption trajectories occurring at different rates. Radial wavefronts are, generally, the 'fastest' of the three modes, followed by path-driven trajectories, and cluster-driven wavefronts are the 'slowest' of the three. Figures 32.1-32.3 show snapshots of the three types of wavefronts for a simulation with N = 500, the origin node trait proportion value equal to 0.01, and different values for *a* and *k*. It takes the 'path-driven' wavefront 30 timesteps (Figure 32.2) to become established in an area the same size as the one occupied by the radial wavefront at Timestep 10 (Figure 32.1). The cluster-driven wavefront propagates at an even slower pace, occupying the same area at Timestep 60 (Figure 32.3).

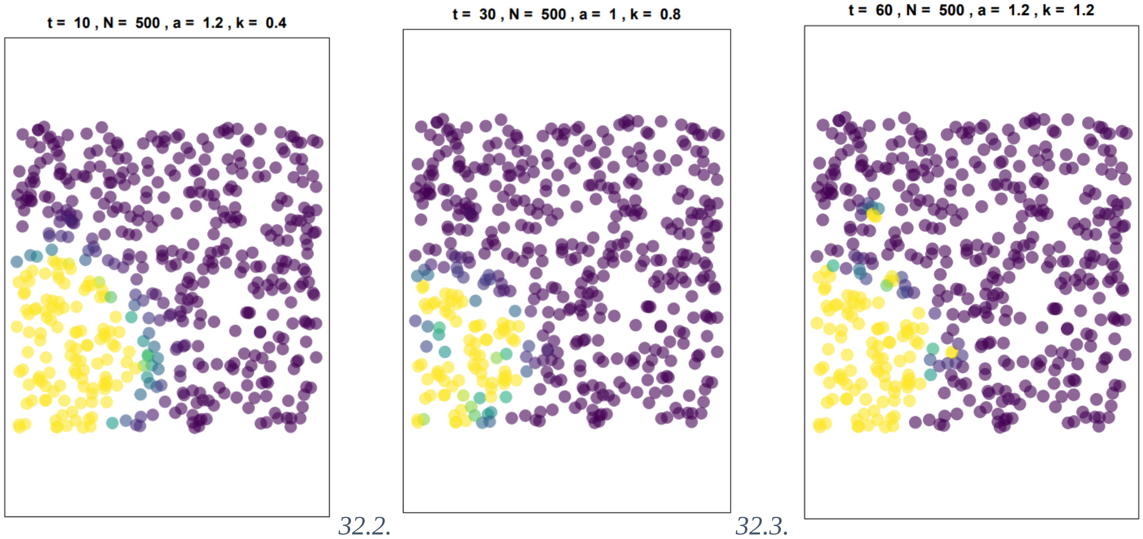

32.1. *32.2.* *32.3.*

Figure 32: Radial (32.1), path-driven (32.2), and cluster-driven (33.3) trait adoption wavefronts.

The parameter space map in Figure 33 clarifies what the effects of the interactions between N, *a*, and *k* are and how they lead to different types of wavefront propagation. It includes different panels for different of values of N (increasing N essentially captures the effect of increasing overall population density in the study area). The x and y axes of each panel correspond to simulations using different values of *a* and *k*. The map confirms that ‘radial’ propagation is associated with faster system-wide trait adoption; pixels in light grey consistently occur at the top tight areas of each panel. These areas are associated with low *k* values – which drive faster system-wide adoption due to low communication costs- and high *a* values – which, again, drive faster system-wide adoption due to high local adoption rates. Additionally, increasing population density (by increasing N and keeping the size of the study area constant) results in a shift from mostly ‘cluster-driven’ wavefronts (e.g. the case where N = 200) to ‘radial’ wavefronts becoming the modal wavefront propagation method (as in the case where N = 700). Path-driven propagation typically occupies a region diagonal to the x and y axes and represents a transition from the ‘radial’ areas to the ‘cluster-driven’ areas of the parameter space map — but note that there is an exception for N

= 300, where a case of path-driven propagation is located inside the ‘cluster-driven’ area of the map. Additionally, path-driven propagation does not occur for all population densities; when the population density is at its lowest (N = 200) the only possible outcomes are cluster-driven, radial propagation and no observable process in space (white pixels). It first occurs at N=300, where it describes 7 out of 25 simulations. As the population density increases, its occurrence decreases (but N = 600 breaks this pattern).

The general picture that emerges is one where all parameters have an effect on the way that wavefronts propagate across the system. The effect of increasing population density is most pronounced in the range from N = 200 to N = 300, and it becomes more subtle beyond that. These results are not surprising; when the population density is high, the relative distances between nodes generally become shorter, making the effects of non-linear communication decay (captured by k) less important. Increasing the local adoption rate effectively speeds up the process even more.

This pattern has interesting implications which should generalise well in the face of empirical cultural diffusion processes in (preferably isolated) regions without internal environmental barriers and boundaries (preferably flat landmasses). While the model is highly abstract and does not incorporate many aspects that characterise the transmission of cultural traits over time — such as trait loss as an effect of drift (see chapter 2 for a definition of drift) — its outputs can be used to inform an ‘expectation’; all else being equal, we should expect empirical cases of cultural diffusion to be associated with radial wavefronts (and speedy adoption) when population densities are high and local adoption rates are equal to 1 (e.g. a case of cultural traditions with high-fidelity copying over time) or higher.

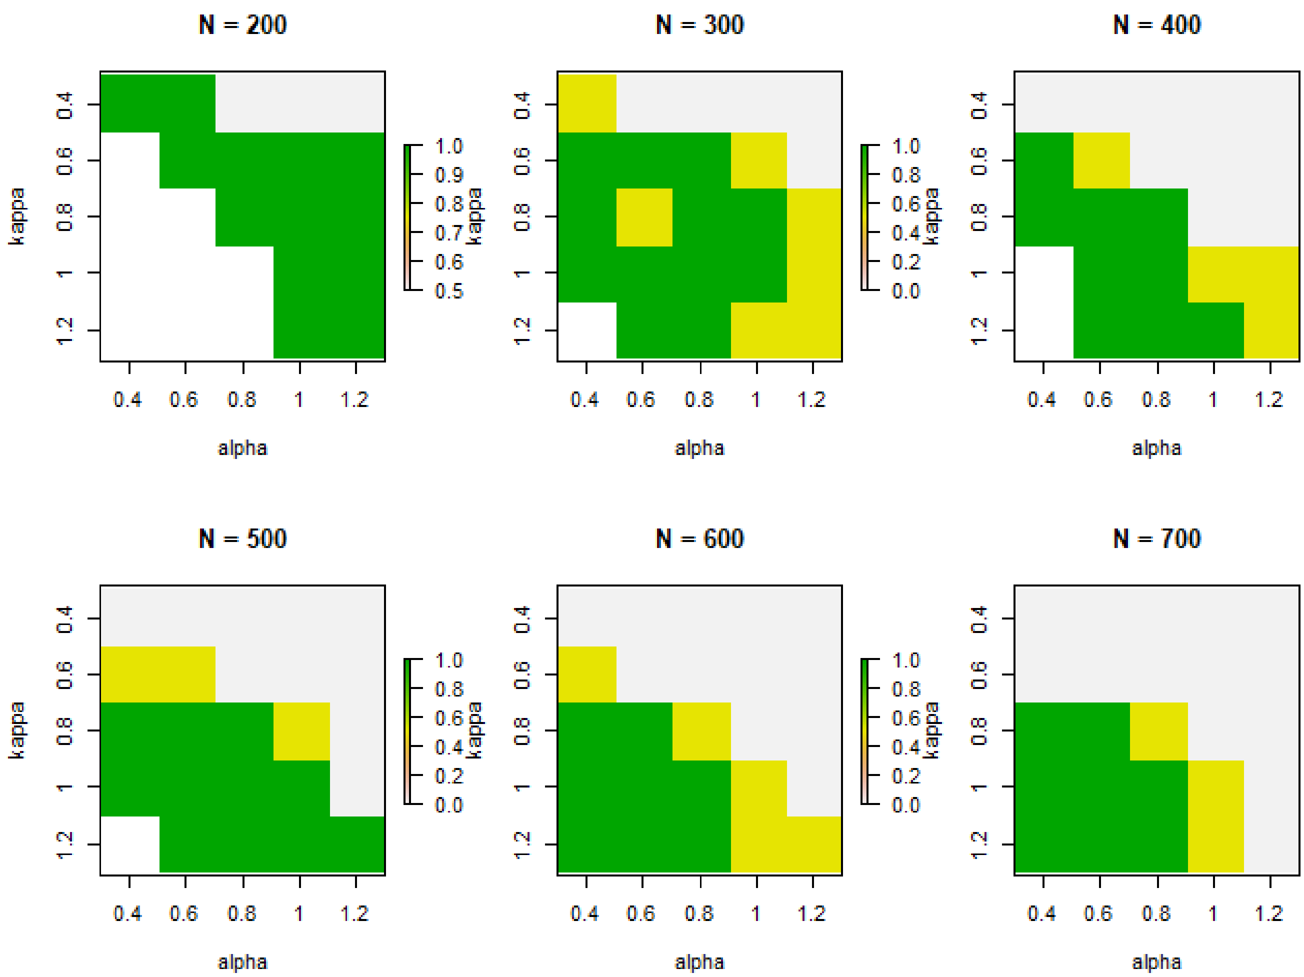

*Figure 33: Parameter space map with values indicating whether each simulation output produces radial (in light grey), path-driven (in yellow), and cluster-driven (in green) trait adoption wavefronts. Each panel is associated with different population densities produced by setting N = 200, 300, 400, 500, 600, and 700 and keeping the size of the study area constant at 100 * 100 units in length. White pixels represent simulations where there is no observable wavefront due to the fact that the trait does not spread beyond the origin point at any point in the observation period (100 timesteps).*

7.4. Toy Model 4: Random trait loss model, varying the environment

Introduction

It is reasonable to assume that in most real-life situations of cultural transmission within and between communities, sometimes some of the population loses the trait due to chance. This could happen, for example, when a prominent teacher dies before their apprentices complete their training. In order to test whether the insights gained from the previous models still hold in a scenario that involves the loss of traits due to chance, I will explore a final toy model, which includes random trait loss at each timestep.

Specifically, at every timestep I allow each community/node to lose a proportion of the trait, which means that some of the potters in that community lose the trait. For each community, the magnitude of its trait loss is a value drawn randomly from a uniform probability distribution between 0 and its current trait proportion value. This value is then subtracted from the current trait proportion value of the community.

I conduct three experiments using the trait loss model. Experiment 1 investigates whether the proximity of the origin node to a cluster of communities protects the trait against system-wide loss. Experiment 2 involves comparing an environment on which communities are located randomly, resulting in relative clustering (the toy model 3 environment) and an environment on which communities are located at regular intervals (the toy model 2 environment) in terms of their providing more resilience against system-wide trait loss. Finally, Experiment 3 investigates the effect of random trait loss on the wavefront types observed in the previous section.

7.4.1. Experiment 1: Simulation on a fully connected network of communities located on a circle at random intervals using the model with random trait loss, N=20, a = 1, k = 0.5, origin = 0.1, observed over 100 Timesteps, using two different locations for the origin

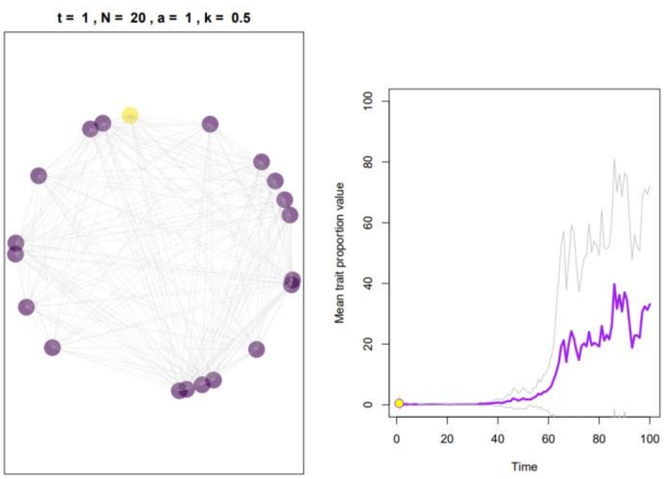

34.1.

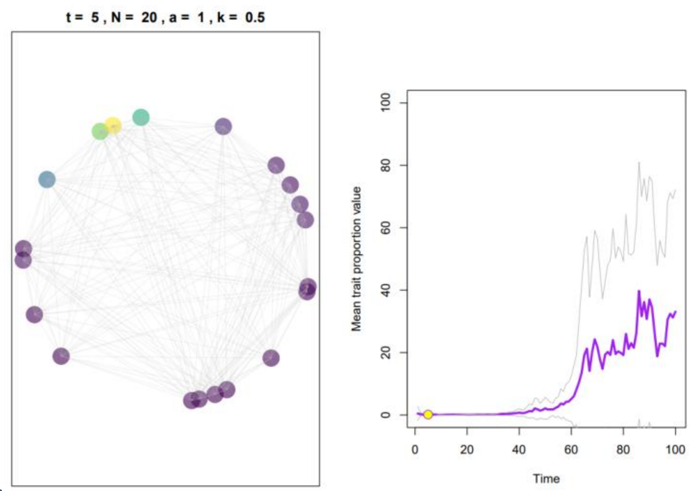

34.2.

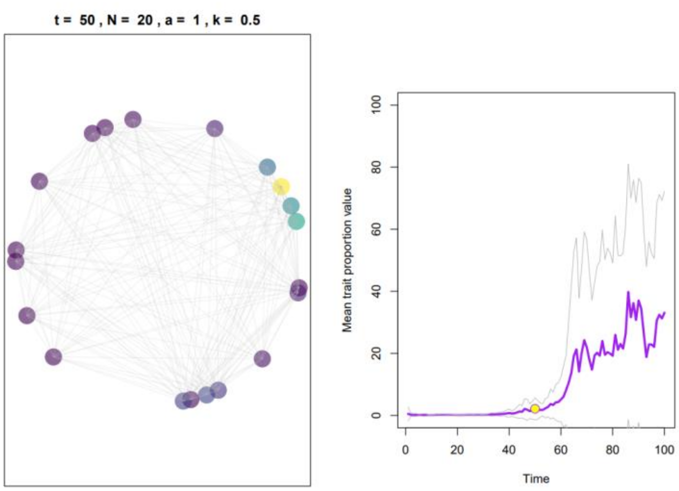

34.3.

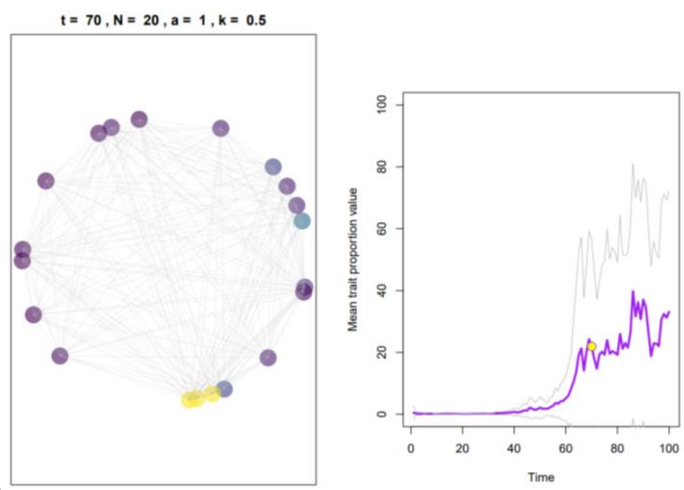

34.4.

Figure 34: The outputs of one simulation with N = 20, k = 0.5, a = 1, and with the origin node's value at timestep 1 = 0.1. The left-hand sides of the four panels show the network at timestep 10, 20, 30, and 40. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. The right-hand side of each panel shows the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

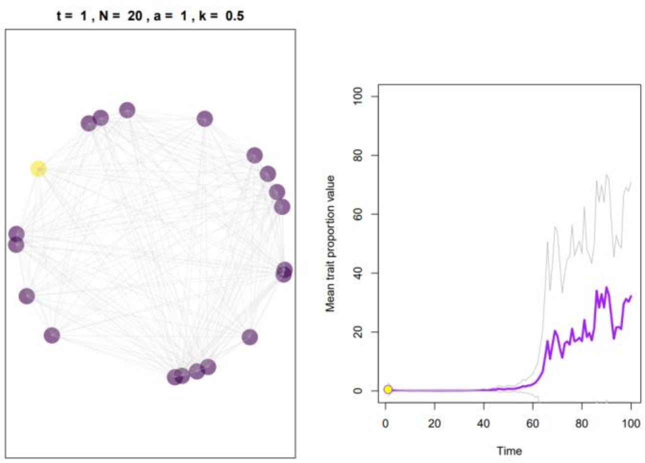

35.1.

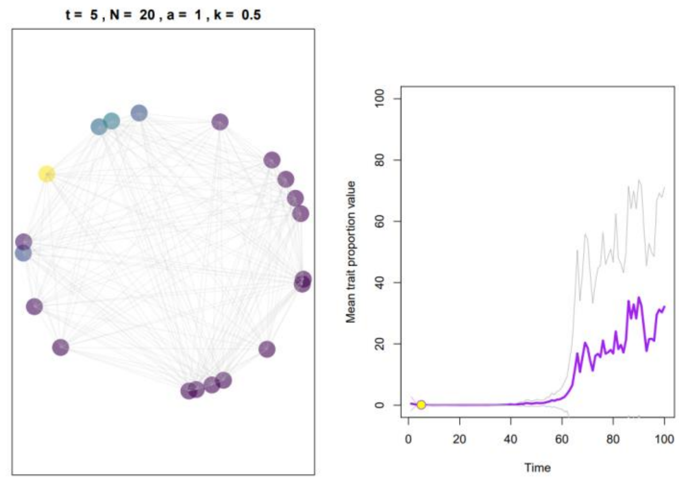

35.2.

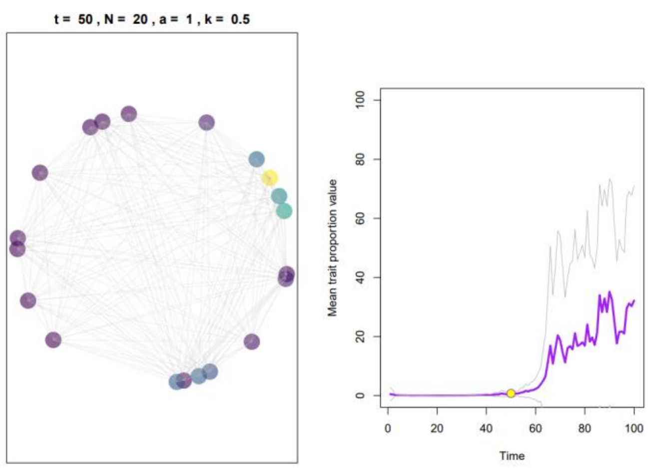

35.3.

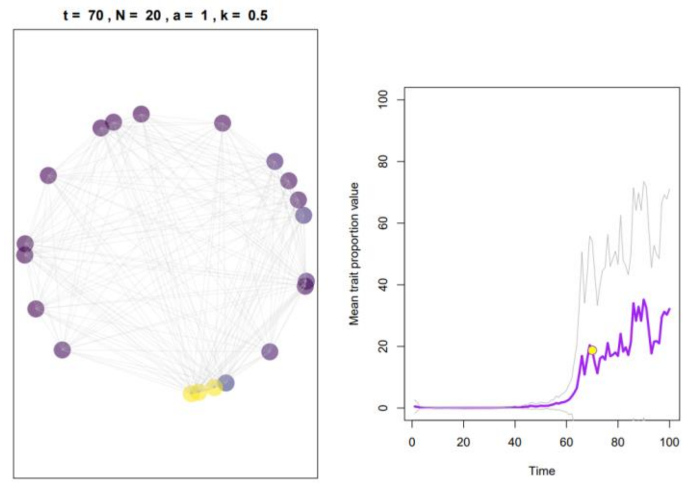

35.4.

Figure 35: The outputs of one simulation with N = 20, k = 0.5, a = 1, and with the origin node's value at timestep 1 = 0.1. The left-hand sides of the four panels show the network at timestep 10, 20, 30, and 40. The colour of each node reflects the relative trait proportion value of each node compared to the values of other nodes in a spectrum ranging from purple to yellow, where purple signifies lower values relative to yellow, which signifies higher trait proportion values. The right-hand side of each panel shows the mean trait proportion value at that timestep, indicated as a yellow dot plotted on top of the trend over time (shown as a purple line). Grey lines define an envelope one standard deviation above and below the mean. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

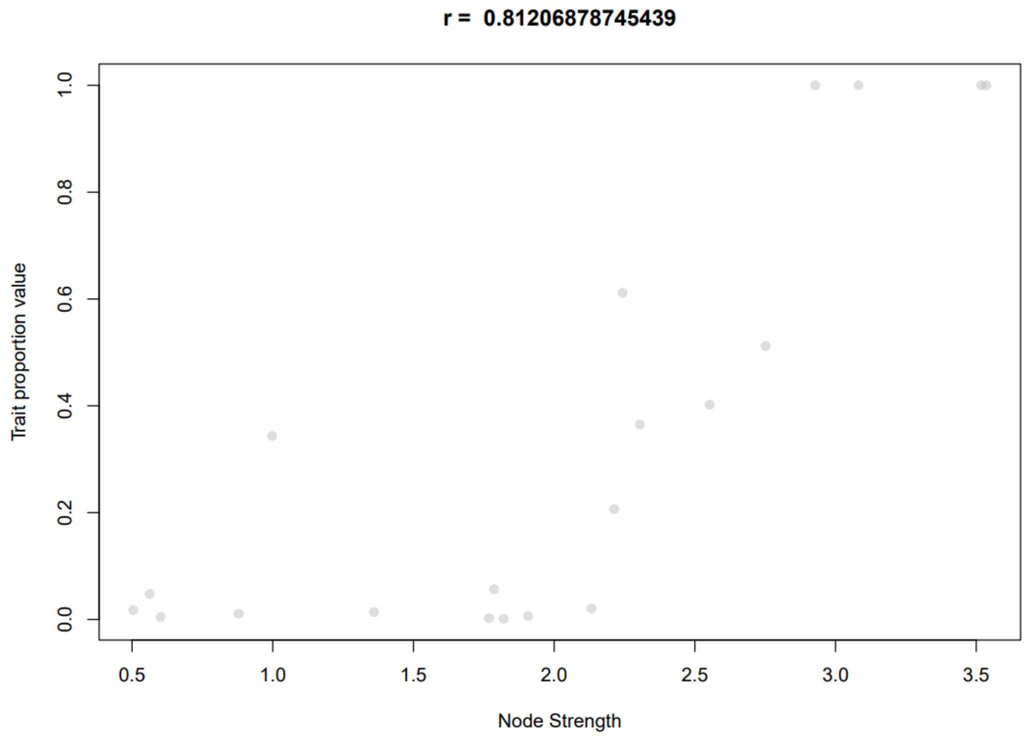

Figure 36: Scatterplot of each community/node's strength (degree centrality including edge weights in the calculation) and its trait proportion value at the end of the simulation in Figure 35. The higher the community's strength, the higher its trait proportion value at the end of the simulation.

Results from Experiment 1

As can be seen in Figures 34 and 35, whether the origin node is located near a cluster of communities or not does not make much difference in terms of the system-wide survival of the trait. What matters more is the fact that a large enough cluster of communities exists somewhere in the environment. Looking at the two simulations, the process that unfolds is similar; the trait becomes established in the cluster nearest to the origin at first and is transmitted through intermediate nodes to the largest and most tight-knit cluster in the environment (the southernmost cluster in this case). The origin node and its cluster lose the trait relatively quickly, but the trait survives in the large cluster even when individual communities in it occasionally lose the trait – they regain the trait easily through the other communities in the cluster. The system-wide mean trait adoption proportion then oscillates around an equilibrium, rising above it when the trait spreads to other communities beyond the large cluster and dropping below it when some communities in it randomly lose the trait.

This means that the presence of relative clustering in the environment matters more than the location of the origin near a cluster, suggesting that initial conditions are not as important as the structural properties of the environment. While it is very possible that this insight only applies to very simple environments with small populations, it is worth exploring how the presence of clusters protects the system against total innovation loss (Experiment 2). Finally, a relevant network statistic that can throw light on why the trait finally survives in the southernmost cluster is node strength, which is essentially node degree but with edge weights taken into account in the calculations. Node strength reflects the sum of all the edge weights, in this case the sum of all "communication strength values" represented by each edge weight. This means that the more information a community gets from all others in the network, the more probable it is that a trait will survive and be more widely adopted (community-level high trait proportion values) in that community.

7.4.2. Experiment 2: Simulation on a fully connected network of communities located on a circle at regular intervals using the model with random trait loss, N=20, a = 1, k = 0.5, origin = 0.1, observed over 100 Timesteps.

37.1. 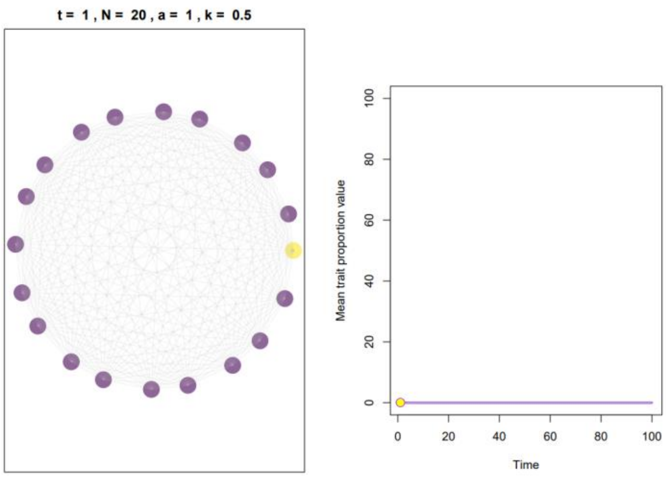

37.2. 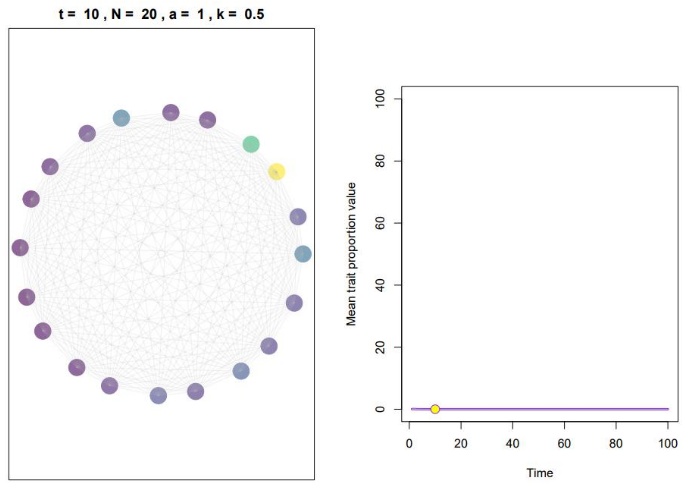

37.3. 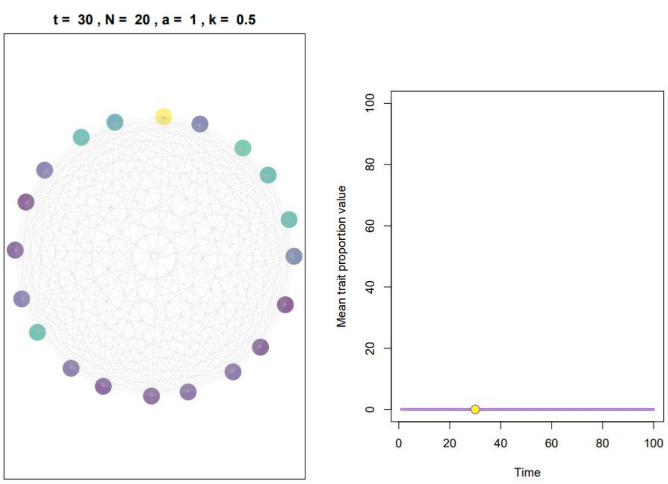

37.4. 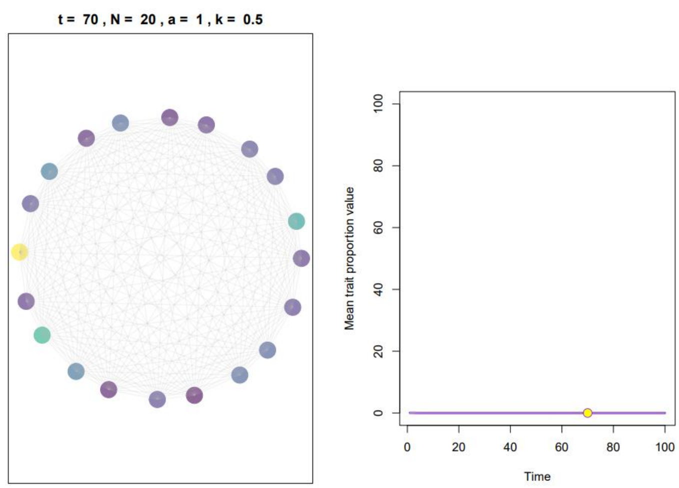

Figure 37: When the communities are regularly placed on the environment – meaning that there is no relative clustering -, the trait is vulnerable to random trait loss at the level of the community and the system-wide mean trait proportion value remains very low. The trait is never entirely lost – at least not in the 100 timestep interval used for this simulation. It spreads from the origin to its adjacent communities on the circle but is quickly lost in every community. The system-wide value remains at an extremely low equilibrium value. Grey lines define an envelope one standard deviation above and below the mean.

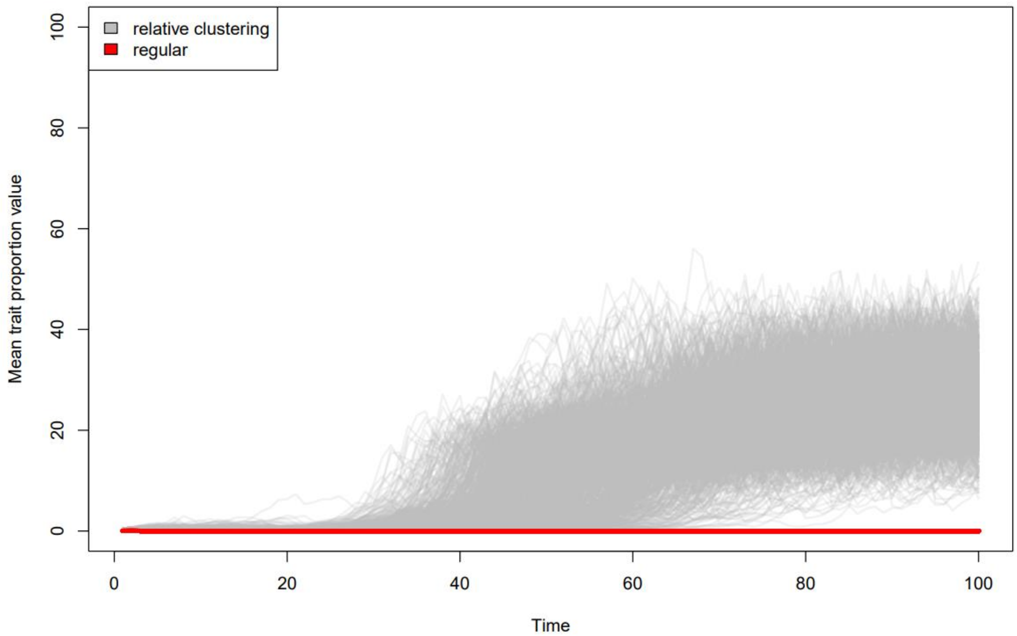

Figure 38: 1000 simulations (in grey) of the model with random trait loss in the environment with relative clustering (the environment in Figures 34 and 35) vs 1000 simulations (in red) with the same parameter values (a = 1, k = 0.5, N = 20) of the model with random trait loss in the environment with regular placement of communities in the environment (the environment in Figure 37). When trait loss occurs randomly, the trait does not spread to the rest of the system in an environment where communities are regularly placed no matter how many times one repeats the simulation. However, when there is relative clustering of communities, the trait spreads to the rest of the system to some extent (even reaching 50 % mean trait adoption for some simulations) depending on the effects of stochasticity. The relative clustering environment consistently outperforms the regular environment in terms of protecting the trait against random loss, possibly due to the process shown in Figures 34 and 35, where the presence of clusters ensures that the trait survives within the cluster even when other communities in the environment lose it.

Results from Experiment 2

The outputs of experiment 2 confirm that the presence of relative community clustering in the environment ensures that the trait becomes established in a larger area than if communities were distributed more evenly across the environment. The latter results in the mean trait proportion value never increasing above its value at Timestep 1 — the transmission process is shown in Figure 37, which shows a simulation using the same values as the ones in Figures 34 and 35. To ensure that this difference is not purely due to chance (note that random trait loss introduces an element of stochasticity in the simulations) I repeated 1000 simulations in both environments using the same parameter values (Figure 38). Even though the simulation outcomes in the environment with relative clustering show greater variation in outcomes compared to the environment with no relative clustering, it always ensures that a greater system-wide mean trait proportion value will be reached by the end of the simulation.

7.4.3. Experiment 3: Radial, path-driven, and cluster-driven wavefronts associated with different combinations of *a*, *k*, and N, origin = 0.01, observed over 100 Timesteps

In order to assess whether trait loss has any effect on the qualitative aspects of the spatial transmission process and the wavefronts described in section 6.3 I will run the trait loss model using the environment and parameter values used for the experiment in section 6.3.1 (the value at the origin node set at 0.01, and varying N from 200 to 700 by 100, and *k* and *a* from 0.4 to 1.2 by 0.2, and using seed (30) for R's random number generator). I will compare the simulations with trait loss to those without and describe how it affects the shape of the transmission wavefront.

Results from Experiment 3

Figures 39.1-2 show that when the transmission process includes random trait loss, the speed of the transmission wavefront is slower, and its shape can be different compared to the model where trait loss does not occur (Figure 39). In the case of parameter values that resulted in radial propagation in the no loss model, including trait loss makes the propagation of the wavefront more 'path driven'. Similarly, in the case of parameter values that led to path driven propagation in the no loss model (e.g., Figure 30), they lead to 'cluster-driven' propagation in the trait loss one (Figure 40).

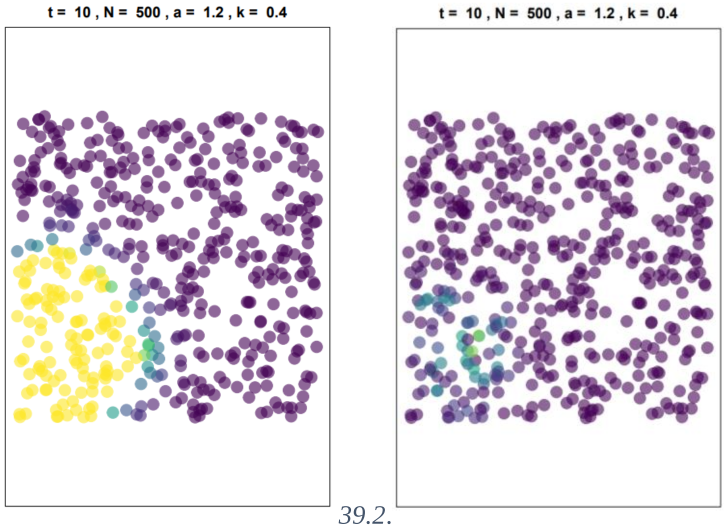

39.1. *39.2.*

Figure 39: Transmission wavefronts at Timestep 10 for simulations with the same parameter values except for random trait loss (random trait loss occurs in the simulation shown in 39.2). Note how in the simulation where trait loss does not occur (39.1) the transmission wavefront propagates radially and has covered a larger area compared to the simulation on the right. The simulation with trait loss propagates in a more path-driven manner compared to the simulation without trait loss.

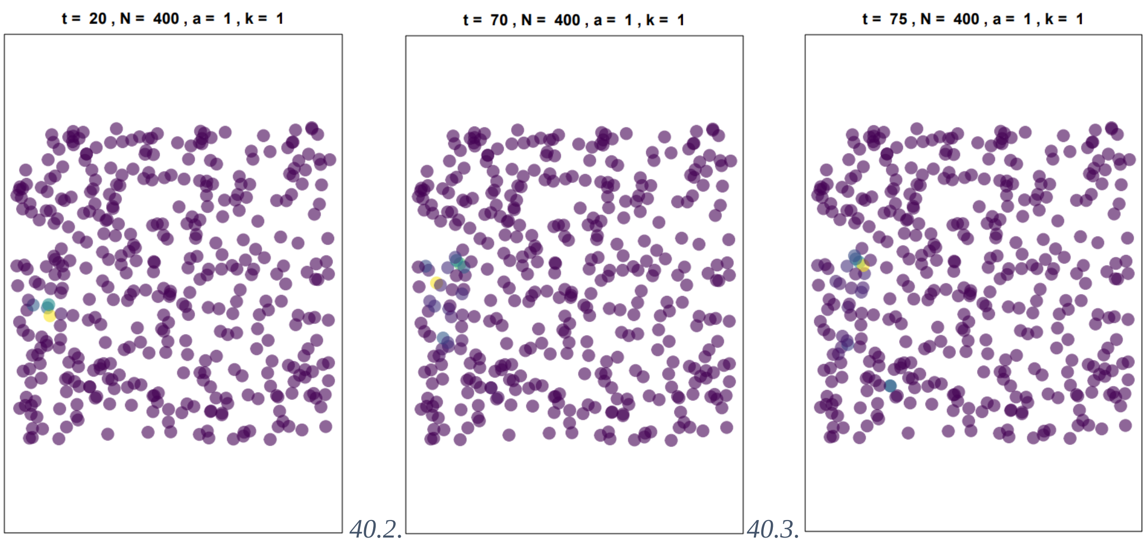

40.1. *40.2.* *40.3.*

Figure 40: Transmission wavefronts at Timesteps 20, 70, and 75 for a simulation with N = 400, a = 1, k = 1 where random trait loss occurs at each community at every timestep. Note how there is no discernible diffusion wavefront – instead, the trait 'jumps' between relatively well-connected clusters located near the origin point.

7.5. Discussion

As shown in this chapter, experimenting with simple models can hone one's intuitions about the simulated process and help identify the elements of the model and the environment that are more relevant to addressing one's research questions. The experiments explored here showed that the relative clustering of communities in the environment is a feature that can ensure higher system-wide trait adoption even when the local rate of adoption in each community is low — provided that the isolation-by-distance assumption holds and that communication costs increase non-linearly with distance. Relative clustering can also help ensure the system-wide 'survival' of the trait against chance events like random trait loss. In such cases, the location of origin of the innovation is less important than the existence of relative clustering somewhere in the environment.

Furthermore, there is a relationship between communication costs, the local rate of adoption, and the population size (number of communities) which determines the speed, shape, and type of the transmission wavefront that propagates in space over time. When the population size is large and the average inter-community distance decreases, transmission is radial (a pure isolation-by-distance process). This also happens when communication costs are low. Low population sizes with high costs and lower local rates of adoption cause the transmission to be more dependent on corridors of tightly knit communities or clusters of well-connected communities in higher density areas of the environment.

Taking on board those insights, it is worth considering what this tells us about empirical cases of trait transmission under similar conditions. In cases where a trait/innovation is transmitted between a small number of communities of practice, it might be more beneficial if some of those communities are located in a cluster than if they are all evenly spread out in the environment, as this can protect the trait from loss due to chance events and it can ensure that the trait keeps spreading even if there is some community level factor limiting the number of craftspeople/adopters that can use the innovation (which is captured by setting the parameter $a < 1$ in this model). In the case of Bronze Age innovations, which are sometimes discussed in the context of production driven by elite demand (e.g., Knappett 2016), this can mean that even if the community-level adoption of a trait was heavily determined by elites, it still had a better chance of spreading to other communities if the area had a relatively high density of communities. Another factor limiting the number of people that can adopt the new trait at the community level is the presence of social boundaries that limit the spread of cultural elements from one group to another — a case well documented ethnographically among communities of potters (Roux 2015:1).

Roux has also discussed 'fragile systems' with low population sizes as being vulnerable to collapse (chance events that cause the loss of an innovation). The results show that when the number of communities is low, the configuration of the network across which transmission takes place can determine whether the trait will survive somewhere in the system. Note that, perhaps counterintuitively given that the isolation-by-distance assumption holds, in this case this 'somewhere' might not be a community near the origin of the innovation. Instead, when we suspect that a trait was being transmitted in space and time and was vulnerable to being lost due to chance events, we should look at large and well-connected clusters in the environment as possible candidates for the trait surviving there after the origin lost it, and as candidates for being new 'origin nodes' that drove system-wide transmission in later stages of the transmission process. This has implications on the way we infer the origin of an innovation using empirical data and the 'isolation-by-distance' principle; we might not always be able to reliably infer the 'true' origin of an innovation (but we can still infer 'an' origin for the part of the transmission process of interest). This might be even more relevant for traits that do not survive well in the archaeological record, e.g., traits that characterize artefacts made of organic materials.

Turning to the more immediate question of trait transmission in the Aegean, the importance of relative clustering in driving trait transmission is relevant here as well. As discussed previously, the Aegean comprises large landmasses as well as natural boundaries (sea, mountains) that affect local population densities and create pockets of community clusters, especially the archipelago connecting the two landmasses. In the case of traits transmitted between a small number of communities, we should expect such traits to survive better in regions where there is a high density of relatively tightly knit communities.

Crete could be such an example, and as discussed in Chapter 4, there are examples of innovations that originated elsewhere, reached the island, were lost or enjoyed little popularity in the area of origin but survived in Crete, which operated as a new origin for more 'successful' system-wide transmission: the potter's wheel appears to be one, as does IRD in the realm of potting technology. Crete is large enough to host a large cluster of communities and is connected to the two landmasses — through Antikythera and the Peloponnese for Mainland Greece and through Kasos/ Rhodes to Western Anatolia — to receive the trait from a distant origin and to spread it to other regions. Given that the environment does not incorporate the geography of the Aegean this is purely speculative in this context and is explored more formally in the following chapter.

Whether a trait is prone to loss due to chance events or not is also relevant here. For traits that are transmitted vertically, such as pottery forming techniques, the reliance on clusters for protection against trait loss might be weaker. For traits that are transmitted more unreliably (e.g., writing systems depending on a small effective population of scribes) this reliance may be greater. Of course, in the models shown here the distinction between trait loss and lack of trait loss was explored as two-fold division rather than as a spectrum, so this is, again, more of a speculative thought rather than an insight clearly emerging from the exploration of the model.

8. Furthering understanding with Aegean-based models

Introduction

Having established some preliminary expectations on the basis of how the model works in abstract environments through simple toy models, a natural next step is to add further realism to the simulations by using a more ‘realistic’ environment that is a closer representation of Aegean geography and by using more ‘realistic’ parameter ranges. Given that the parameters k and a are ‘free’ parameters which are not calibrated by empirical data, it is only possible to use potentially realistic values for the parameter N, which represents the number of pottery-making communities in the Aegean, as well as use a realistic number of timesteps for the simulation.

Simulations will be quantitatively compared to the potter’s wheel trend from chapter 5, since this innovation is the only one out of the three (potter’s wheel, kilns, IRT) for which there is data of sufficient granularity and quantity to produce an empirical time series. However, the simulation model defined here is abstract enough to produce outputs that can be *qualitatively* compared to the transmission processes for the kiln and IRT innovations.

The chapter is divided into four sections: section 8.1 explores simulations where there is no community-level trait loss and section 8.2. develops methods to compare the outputs of the simulations in section 8.1 with the empirical record. Section 8.3 explores simulations of the model with trait loss. Subsections of 8.1. and 8.3 involve experiments that are run on different environments and under different parameter values or using variations of the main model. I also explore a more coarse-grained model without trait loss that introduces fluctuations in the number of communities over time in section 8.4.

8.1. Aegean based simulations without trait loss

Figure 41 shows the effects of varying the number of communities in the environment (parameter N), suggesting that linear increases in N do not result in linear increase in the speed of the transmission process reflected by the trend of the mean trait proportion over time. The shapes of the curves are similar (double-S curves) and increasing N beyond 500 results in similar trends over time, which suggests that the model is not as sensitive to N beyond this value. Therefore N = 2000 is a value that will be held constant in order to explore the effects of the other parameters in the following experiments.

Figures 42.1-42.6 show the output of a simulation with randomly placed nodes N = 2000, $a = 1.2$, $k = 0.6$ running for 1588 timesteps, selecting the easternmost node as the origin point for the new trait, and with the value at the origin being equal to 0.1 at Timestep 1. Timestep 1 is shown in Figure 2.1. Within 100 years, the trait has reached a cluster of nodes around the origin point (Figure 42.2), and the trend is still in its initial phase of slow increase just before a sharp increase phase. By Timestep 200 (Figure 42.3), the trait has reached most of the southeast part of Western Anatolia, propagating along two wavefronts, one to the north and one to the west of the cluster. At this stage, system-level adoption increases very quickly.

At Timestep 500 (Figure 42.4), mean trait proportion values increase more slowly. Judging from the spatial process on the map, this is due to the trait having spread to the entire Western Anatolian mainland, and because both wavefronts (north and west) have stumbled upon barriers: to the north, the wavefront has stopped at the Sea of Marmara, and to the West it has stumbled upon the Aegean, where it has reached the islands of the eastern Aegean that are located within a small proximity of the Anatolian coastline, but not the islands of the Central Aegean. Within another 100 years, the trait has spread northwards towards Thrace through the Dardanelles Strait and established a new west-facing wavefront that is moving towards Macedonia (Figure 42.5). Additionally, the trait has reached high values at Euboea, after “crossing” the Aegean from the original western wavefront – at this stage, the trend enters a new exponential phase. By timestep 800 (Figure 42.6), the trait has spread to Macedonia through the northern wavefront, and the Euboean cluster has developed into another wavefront moving westwards through mainland Greece. At this stage system-wide adoption increases very fast until it reaches the 100 % equilibrium.

The picture that emerges is one where geography plays a major role in terms of increasing or decreasing the speed of transmission. The spatial process is one where the new trait originates from a point inland located on a large landmass (Western Anatolia). It spreads very fast until it has been transmitted to all communities within the landmass, and then there is a delay until it reaches another large landmass (mainland Greece) either via a long-distance maritime route (the Euboea node) or by a small hop through the Dardanelles strait.

It is worth noting that Crete — a smaller landmass that is relatively isolated from the other two and depends on centrally located islands in the Aegean for communication — receives the trait relatively late in the adoption process and is certainly not a driving factor in spreading the trait to other parts of the Southern Aegean. This is in stark contrast to the picture painted by the empirical data and the literature on Crete’s role in establishing ceramic innovations like the potter’s wheel and iron-oxide based decoration in the Aegean during the Bronze Age and could be interpreted in multiple ways. First, it is possible that for N = 2000, the minimum inter-node distance is small enough that the diffusion process moves along wavefronts and space acts as a continuous surface -at least until each wavefront reaches the sea. Intuitively, it seems likely that high population densities would result in the diffusion process being driven to a higher degree by land travel than by maritime travel. Alternatively, it is possible that Crete becomes important only if travelling to and from Crete was relatively easier compared to travelling between other regions, which would be possible if faster means of travelling were only used by communities in Crete, which is not unreasonable given the literature's emphasis on sailing being associated with Minoan communication networks. This would possibly enable the trait to reach Crete faster, become well established in a cluster relatively quickly (because of the high *a* value which means trait proportion increases on its own), and then propagate along a wavefront that travels to the south part of mainland Greece through a connection via Kythera.

Overall, the simulation outputs suggest that geography is an important factor driving the spread of the trait, and that the adoption process resembles a travelling wave whose propagation across large landmasses (like Western Anatolia) corresponds to phases of high rates of increase on the adoption curve. This process is punctuated when the wavefront reaches the coastlines of said landmasses.

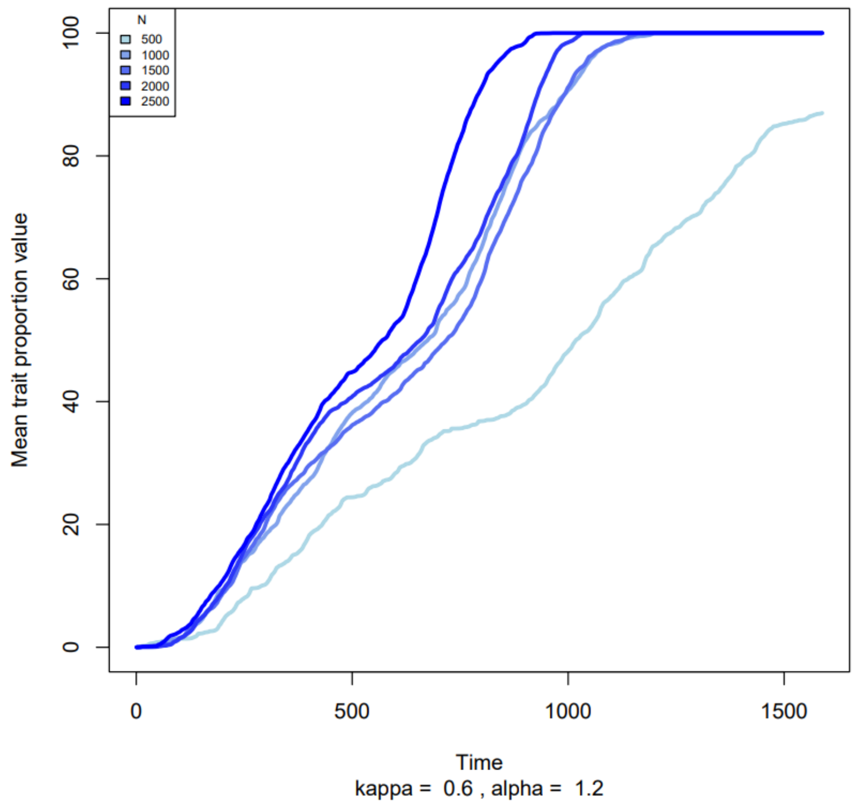

Figure 41: Simulations varying population size. The easternmost community/node is selected as the origin of the trait with a value equal to 0.1 at Timestep 1.

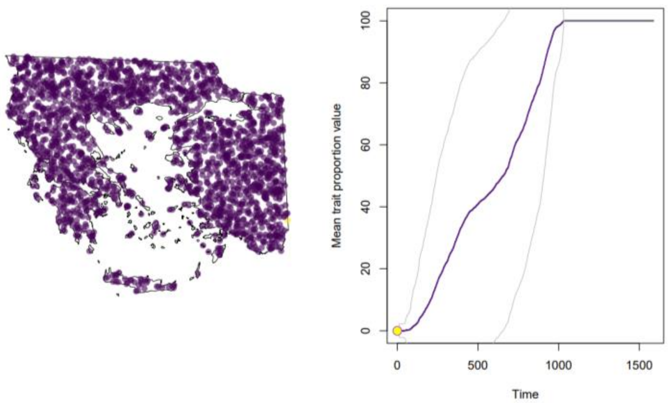

42.1.

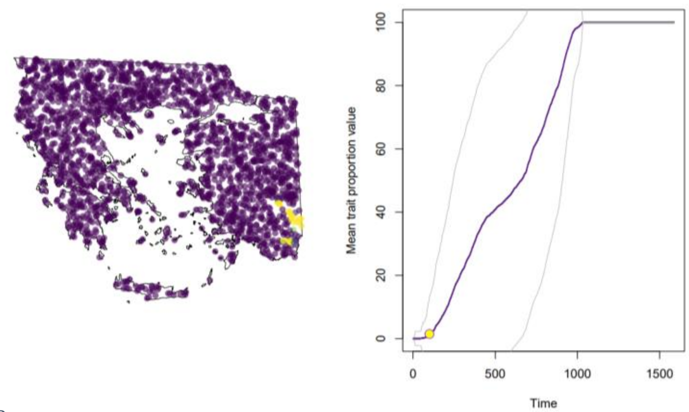

42.2

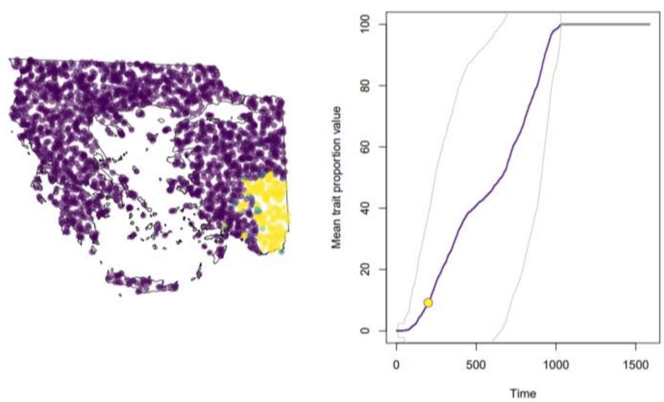

42.3.

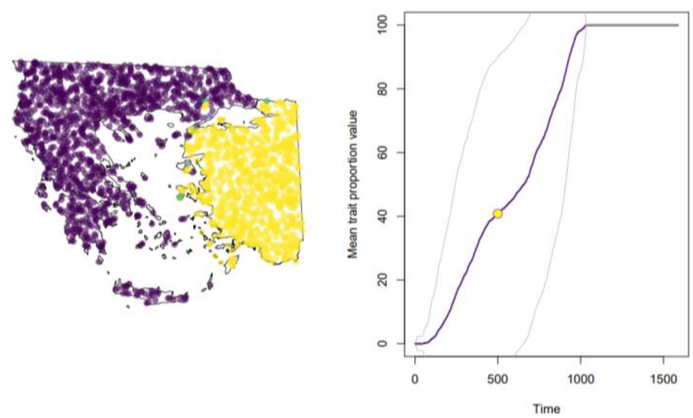

42.4.

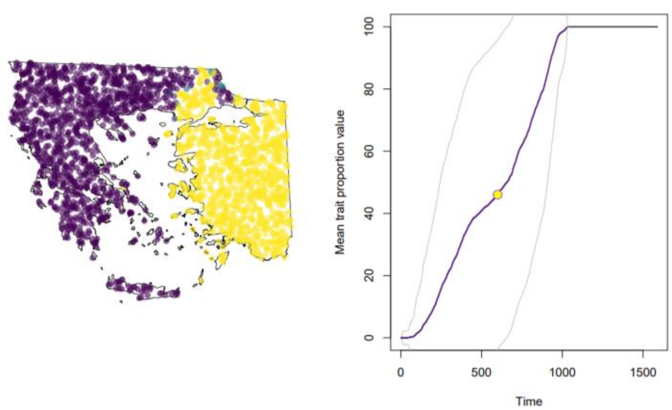

42.5.

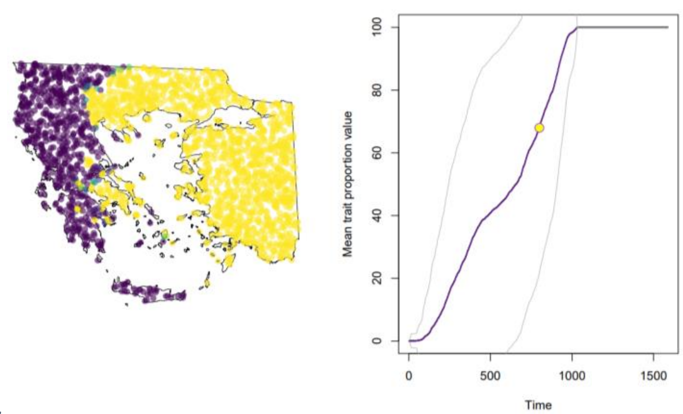

42.6.

Figure 42: Trait adoption at Timestep 1, 100, 200, 500, 600, and 800. Left, map of the Aegean, with colour representing trait proportion value at each node. Right, mean trait adoption value and standard deviation per timestep. The yellow dot indicates the mean trait proportion value for the entire study area at the timestep which is shown on the map.

8.1.1 Experiment 1: Varying the location of the origin point

As aforementioned, the simulation outputs above suggest that trait adoption spreads fastest on large landmasses and is punctuated by slower adoption episodes that start when all nodes within a landmass have adopted and which last until the trait reaches a new landmass, either via a terrestrial or maritime route – the latter happens faster when islands are used as “bridges”. This suggests that, overall, geography is an important factor determining trait diffusion trajectories. A question that remains, however, is whether geography is as important in the beginning of the process as it is in the intermediate or final stages. In this section I will explore the effect of varying the location of the origin node using locations outside Western Anatolia — a choice of location that was based on the hypothesized origin of the potter’s wheel (discussed in detail in Chapters 4 and 5). The aim of this experiment is to investigate the effect of different initial conditions — the location of the origin community — on system-wide adoption trajectories but it can also be applied to study the transmission process of other traits, such as kilns and IRT, whose adoption trajectories were possibly Cretan-Minoan driven.

I define 9 origin points in the northwest, west, southwest, south, southeast, east, northeast, north, and centre of the study area. Figure 43 shows the locations of the origins — note that they have not been symmetrically placed on the study area, as this would result in some points falling on the sea. Each simulation used the same parameter values as in the previous examples: $k = 0.6$, $a = 1.2$, N = 2000, running for 1588 timesteps –the only element that changes in each simulation is the location of the origin point.

The trait adoption trajectories in Figure 44 suggest that when the origin node is located near the eastern or western edge of the study area — as is the case for the NW, W, SE, E, and NE origins -

the resulting trend exhibits the double-S curves that were discussed previously. Note that this also happens when the node is located on an island near a major landmass, as in the case of the Western node in Zakynthos (island to the west of the Peloponnese). In such cases, the propagation of the wave is delayed when it reaches the bottleneck that is crossing the sea to Western Anatolia (when the origin is in mainland Greece/Balkans) and to mainland Greece/Balkans (when the origin node is located in Western Anatolia).

The fastest global adoption process happens when the origin node is located in the centre of the study area, on an island near mainland Greece, and when the origin node is located in the North. Both cases produce S-shaped curves, but in the case of the former (Figure 45) the adoption process is initially slow, because it depends on the strength of communication between relatively isolated nodes on islands (those nodes are sparser compared to the average density on landmasses). However, placing the origin on such a central location means that this initial slow phase eventually ends when the trait reaches the Greek mainland and Western Anatolia landmasses, and from that point onwards the trait propagates along two waves on each landmass, increasing at a steady high rate. In this scenario, Crete remains a laggard, as in previous simulations. However, it is now the central part of Crete that receives the trait first.

Crete is interesting in that, when used as an origin point (Figure 46), both nodes located there (nodes SW and S) produce S-shaped trait adoption trajectories that reach the 100 % equilibrium at around the same time as the double-S curves produced by the NW, W, SE, E, and NE origin nodes. There is a prolonged phase of very slow trait adoption up to timestep 400, and then trait adoption increases exponentially for about 100 timesteps. Crete's relative geographic isolation is the main factor behind the slow initial uptake, but Crete's central position on the East-West axis means that when the trait leaves Crete, it has the potential to propagate along two wavefronts, one to the East and one to the West.

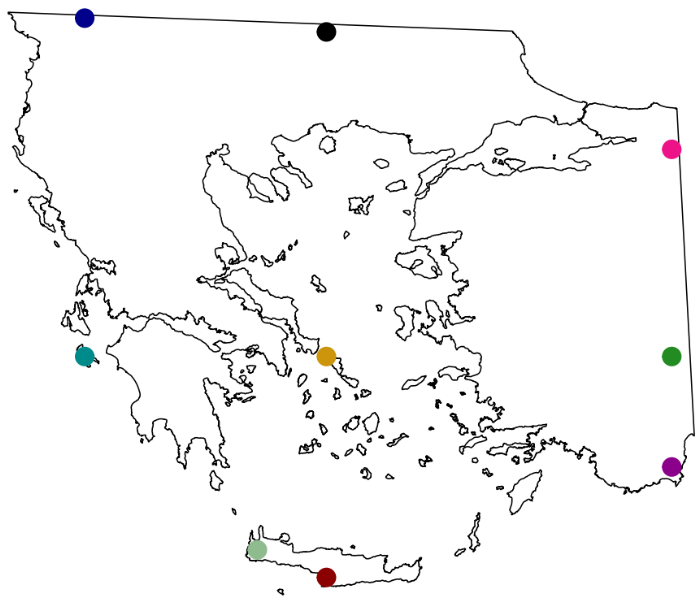

Figure 43: Origin points used for each simulation.

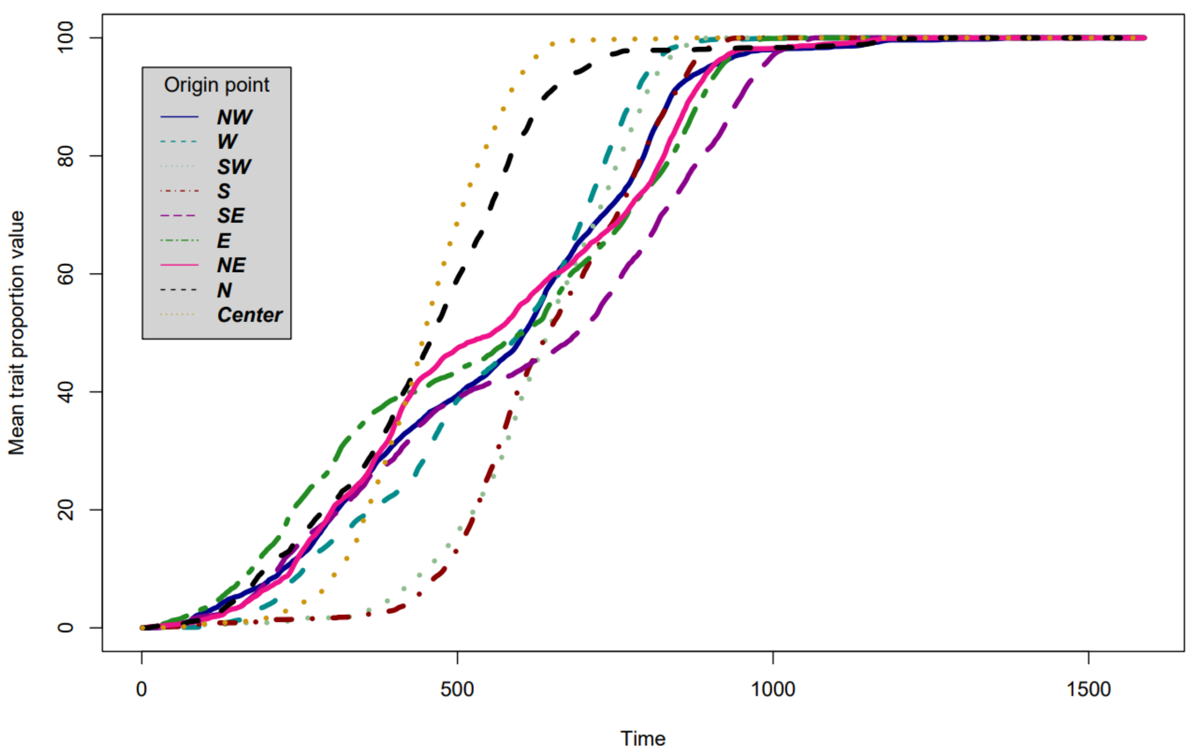

Figure 44: Mean trait adoption per timestep for simulations with different origin points and k = 0.6, a = 1.2, N = 2000.

Origin Point: Center

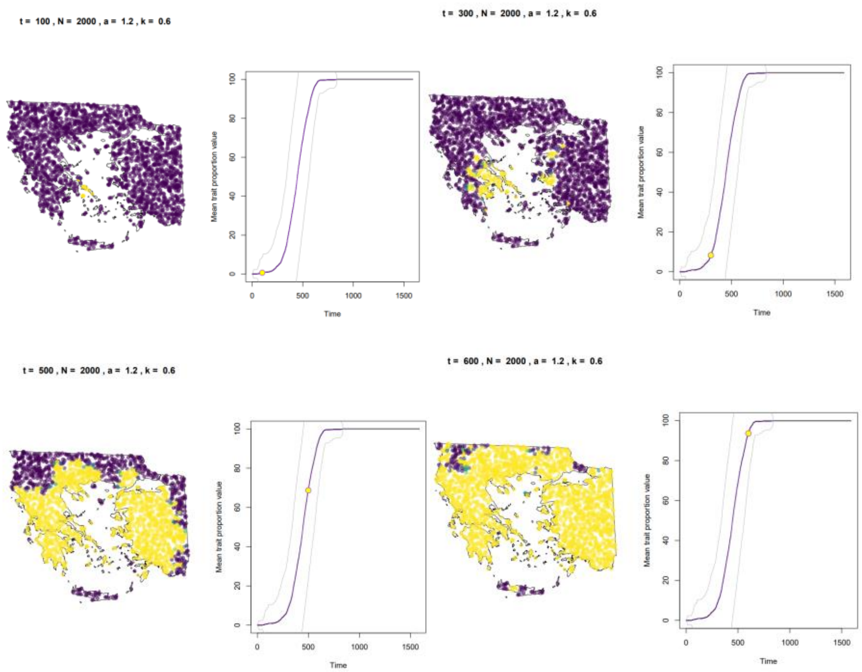

Figure 45: Trait adoption at timestep 100, 300, and 600. Origin point is "Central" (the island of Andros) and k = 0.6, a = 1.2, N = 2000.

Origin Point: Southwest (Crete)

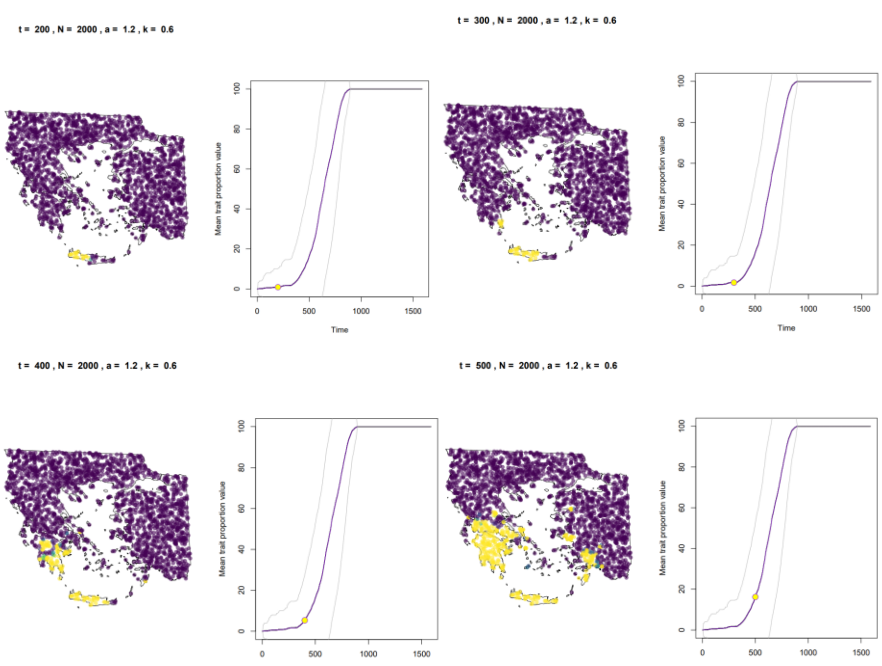

Figure 46: Trait adoption at timestep 200, 300, 400, and 500. Origin point is "Southwest" (Crete) and k = 0.6, a = 1.2, N = 2000.

8.1.2. Experiment 2: Coastal vs inland habitation preference

In the previous chapter it was established that for a relatively simple abstract model incorporating the communication costs associated with non-linearly weighted distances and a high rate of local adoption, geography and population structure can affect trait adoption trajectories quantitatively — changes in the rate of change per timestep —– and qualitatively — diffusion that propagates as a circular wave vs a main diffusion wave propagating along nearby clusters.

The simulation described in this section explores the effect of habitation preferences for different types of landscapes using a simple two-fold division between coastal and inland landscapes. Coastal landscapes are defined as the regions of the study area that are located within 1 km from the coastline, and inland landscapes are defined as the areas that do not fit this criterion. The location of nodes is no longer purely random — instead, I have placed nodes at random on each zone (coastal/inland) using different probabilities for each zone whose sum is equal to 1. For example, coastal preference would correspond to a probability of coastal habitation $p > 0.5$, where inland habitation probability would be $1-p$. Figure 47 shows an extreme case where each node has a 0.95 probability of falling within the coastal zone, and a 0.05 probability of falling within the inland zone. It is reasonable to assume that coastal areas would be more densely populated than regions deep into the landmasses — both landmasses include highly mountainous areas with steep slopes that are not explicitly factored in as geographic boundaries in this model, but possibly impeded continuous occupation over the landscape.

The effect of coastal preference is first explored using this extreme case, because if it proves to be that the model is not sensitive to the extreme case, it is unlikely that it will be more sensitive to a less extreme case of coastal preference, allowing one to safely reject coastal preference as a relevant factor affecting transmission and move on to other, more relevant environmental parameters. As in previous simulations, the parameter values remain the same: $k = 0.6$, $a = 1.2$, $N = 2000$, with each simulation running for 1588 timesteps and using the easternmost point as the origin node for a coastal preference equal to 0.95.

It becomes immediately apparent that the double-S trend produced by simulations where nodes were plotted purely randomly over the study area is not produced by the new simulation shown in Figure 48. The first 300 years exhibit slow growth, and then the rate of spread increases exponentially and then retains a high but steady growth rate until full adoption is reached around timestep 600. As would have been expected, having a higher density of communities around the coast means that as soon as the trait reaches the coast, a new wave that propagates parallel to the coastline is the main driver of trait adoption globally.

Initially, the trait spreads relatively slowly in Western Anatolia because nodes on the inland zone are sparser compared to coastal ones. As the coastal wave advances, it starts moving inwards towards the inland zones. Coastal preference also ensures that islands receive the new trait relatively earlier compared to the simulation in Figure 42, as fast adoption on islands depends on the trait reaching the coast as fast as possible. For example, Crete has already received the trait by timestep 700 via the southern coast of the Peloponnese, whereas in the simulation with completely random node locations, Crete had not yet received it on timestep 800 (Figure 42.5).

Strangely enough, while one would expect ‘laggard’ nodes to be found strictly on inland regions – and this is certainly the case for the cluster of laggard nodes on the northwest of the study area at timestep 800 (Figure 48.4), it is also possible to identify laggards on coastal locations on islands – for example, the single purple node on Kasos, near Crete. This suggests that even though coastal preference can lead to islands receiving the trait earlier, it is sometimes the case that for relatively isolated islands like Kasos, if a node is located far away from the nearest neighbouring island — in this case, Karpathos to the east — it is possible for a node to exhibit delayed adoption compared to other nodes in its vicinity. In fact, it looks as though the fast adoption on islands is mostly due to the coastal wave that develops on large landmasses rather than "coastality" enabling faster communication between islands. All the above suggest that double S-curves can be associated with geographic barriers like landmasses with maritime boundaries and steep mountainous areas that limit human habitation, but this effect becomes less pronounced (and even non-existent) when coastal preference is relatively high.

If one wished to further develop the model in the direction of greater realism, they would have to consider the fact that islands are likely to interact with their nearest neighbours. Even if an island is relatively isolated, it is still reasonable to assume that some of its resources would be dedicated to enable the exchange/movement of people and raw materials with other islands. The more limited the resources of each island, the more likely it would be that those resources would be used to communicate with a smaller number of preferred nearest neighbours. The model explored here does not consider the relative ranking of distances for each node's neighbours, for calculating the inter-node weights. A model that incorporates this effect is explored in section 8.1.5.

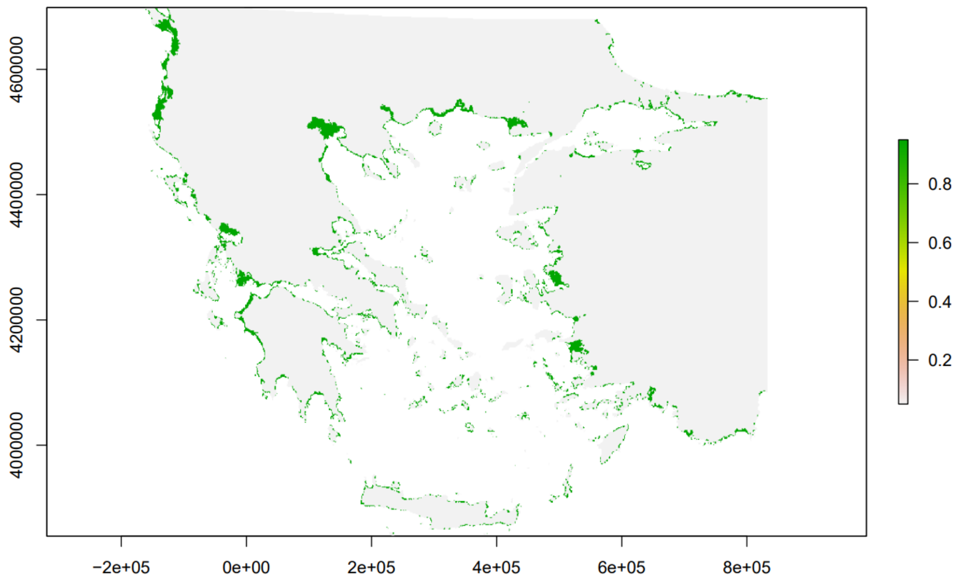

Figure 47: Study area divided into an inland and coastal zone. Nodes have 0.95 probability of falling on the coastal zone, and 0.05 probability of falling on the inland zone.

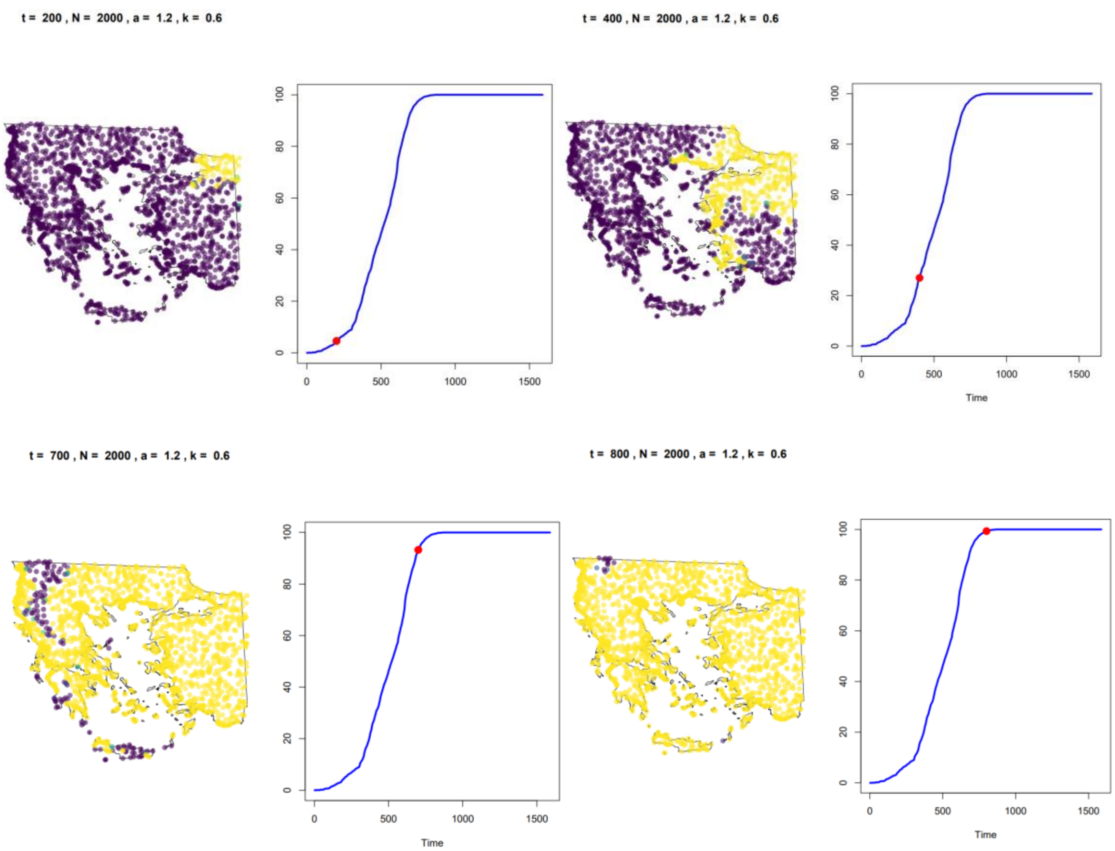

Figure 48: Trait adoption at timestep 200, 400, 700, and 800. k = 0.6, a = 1.2, N = 2000 and coastal habitation probability is equal to 0.95. The easternmost node is the origin of the trait, with a starting value equal to 0.1. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

8.1.3. Experiment 3: Habitation along coastal flat, coastal steep, inland flat, and inland steep zones

The previous experiment explored the effect of varying population density in coastal and inland zones and assuming that in reality population density was probably lower in the steep regions. The landscape in that environment reflected an extreme case where 95 % of the population inhabits a narrow coastal zone within 1 km of the coast, leading to a diffusion that develops as a coastal wave. This section explores a model that includes geographic boundaries within the inland zone and coastal zones — the inland zone is now divided into a flat inland zone, where the slope is below 10 degrees, and a steep zone, where the slope is above 10 degrees. The coastal zone is extended to a region 5 km beyond the coastline, a scenario which is more flexible in terms of allowing coastal populations to "move" from flat areas immediately next to the coastline to mountainous areas near the coast. The coast is again divided into steep and flat zones which present different relative advantages; for coastal steep areas, a high vantage point over the coast provides protection against hostile incomers. Regardless of this advantage, it does not seem likely that coastal steep zones were as densely inhabited as coastal flat areas, since coastal steep areas probably rendered agricultural production difficult or at least less productive, and therefore less successful in sustaining a large population size per community as coastal flat areas — a preference for lowlands is reported for Bronze Age Crete by Argyriou et al. (2017:20) and, as Whitelaw (2000: 152) states, "site locations in prehistory throughout the southern Aegean … clearly favoured low-slope land, which could be cultivated without requiring slope terracing for soil conservation". A preference for coastal land is also attested for 'Minoanising' sites in Kythera (Broodbank and Kiriatzi 2007: 262) and a preference for 'agriculturally favourable land" is noted for Antikythera during the Bronze Age (Bevan and Conolly 2013: 122). There is some recent evidence from an extensive surface survey that suggests that mountainous areas became more densely inhabited during the Neopalatial period in eastern Crete (Kalantzopoulou 2022), but it is mostly based on inorganic architectural evidence (which certainly suggests the presence of human populations, but its absence should not be used to rule out the presence of human habitation, especially as houses could also be made of organic materials). The model explored here does not incorporate temporal fluctuations in landscape zone probabilities, but it assigns a relatively low probability of habitation to zones with high slope values. As in the previous experiment, the probability of each node falling on one of the three zones is parameterised, and the probabilities assigned to all zones sum to 1.

The simulation shown in Figure 49 is performed on a landscape where each community/node has a 0.45 probability on falling on the coastal flat zone, a 0.15 probability on falling on the coastal steep zone, a 0.30 probability of falling on the inland flat zone, and a 0.10 probability of falling on the inland steep zone. Parameter values follow the baseline set in previous sections: k = 0.6, a = 1.2, N = 2000, keeping each simulation running for 1588 timesteps and using the easternmost point as the origin node.

The outputs suggest that for those landscape probabilities - where population density is highest on coastal flat regions, followed by medium densities at inland flat regions, low densities around coastal steep areas and very low densities on inland steep areas — the trend is still double-S shaped, with initial fast adoption delayed when the wavefront reaches the mountainous area in Macedonia. Figure 50 shows that when coastal flat habitation preference is even more extreme at .70 — the simulated trend is very similar to the one produced by the simulation in Figure 48.

Having established the landscape zones of interest, it is possible to investigate the effect of the location of the origin points when population density differs between landscape zones. Using the origin nodes shown in Figure 43, and the landscape probabilities shown in Figure 49, it is possible to discern (Figure 51.2) a similar pattern to the one in experiment 1 (Figure 51.1) — origins in the North and Center produce the fastest, single-S shaped adoption trajectories, nodes located on Crete (S and SE nodes) produce single-S curves with a more prolonged phase of initial slow adoption, and all other origin nodes produce double-S curves that reach the equilibrium at the roughly same time as the Cretan nodes. However, those double-S curves are less pronounced, and their trajectories are not as similar to each other as in the example in Figure 51.1. This is due to the fact that each trend reaches the region of slow adoption that produces the punctuated double curve at different times, possibly when a landscape obstacle like a mountain range is reached.

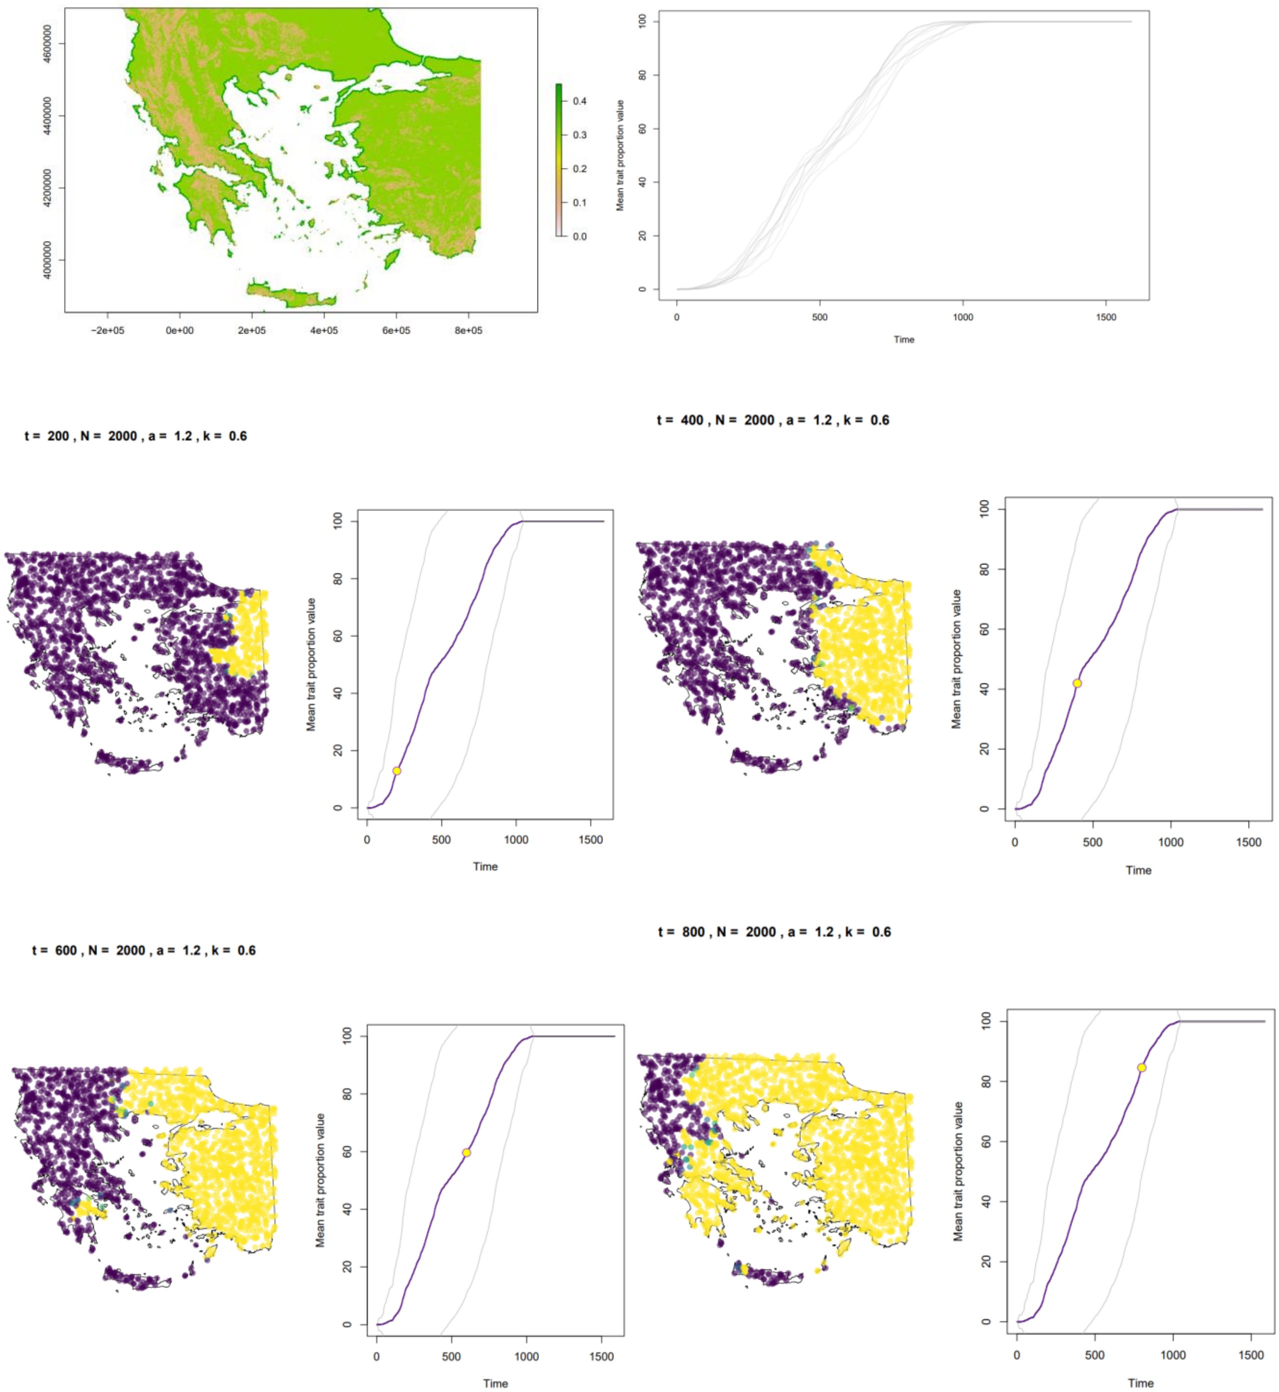

Figure 49: Study area divided into an inland flat (light green), an inland steep (light orange), a coastal flat (dark green), and a coastal steep (orange) zone. Nodes have a 0.45 probability of falling on the coastal flat zone, a 0.15 probability of falling on the coastal steep zone, a 0.30 probability of falling on the inland flat zone, and a 0.10 probability of falling on the inland steep zone. The top right panel shows the outputs of 10 simulations using the same parameter values and suggests that stochasticity does not cause significant changes in the shape of the curve. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

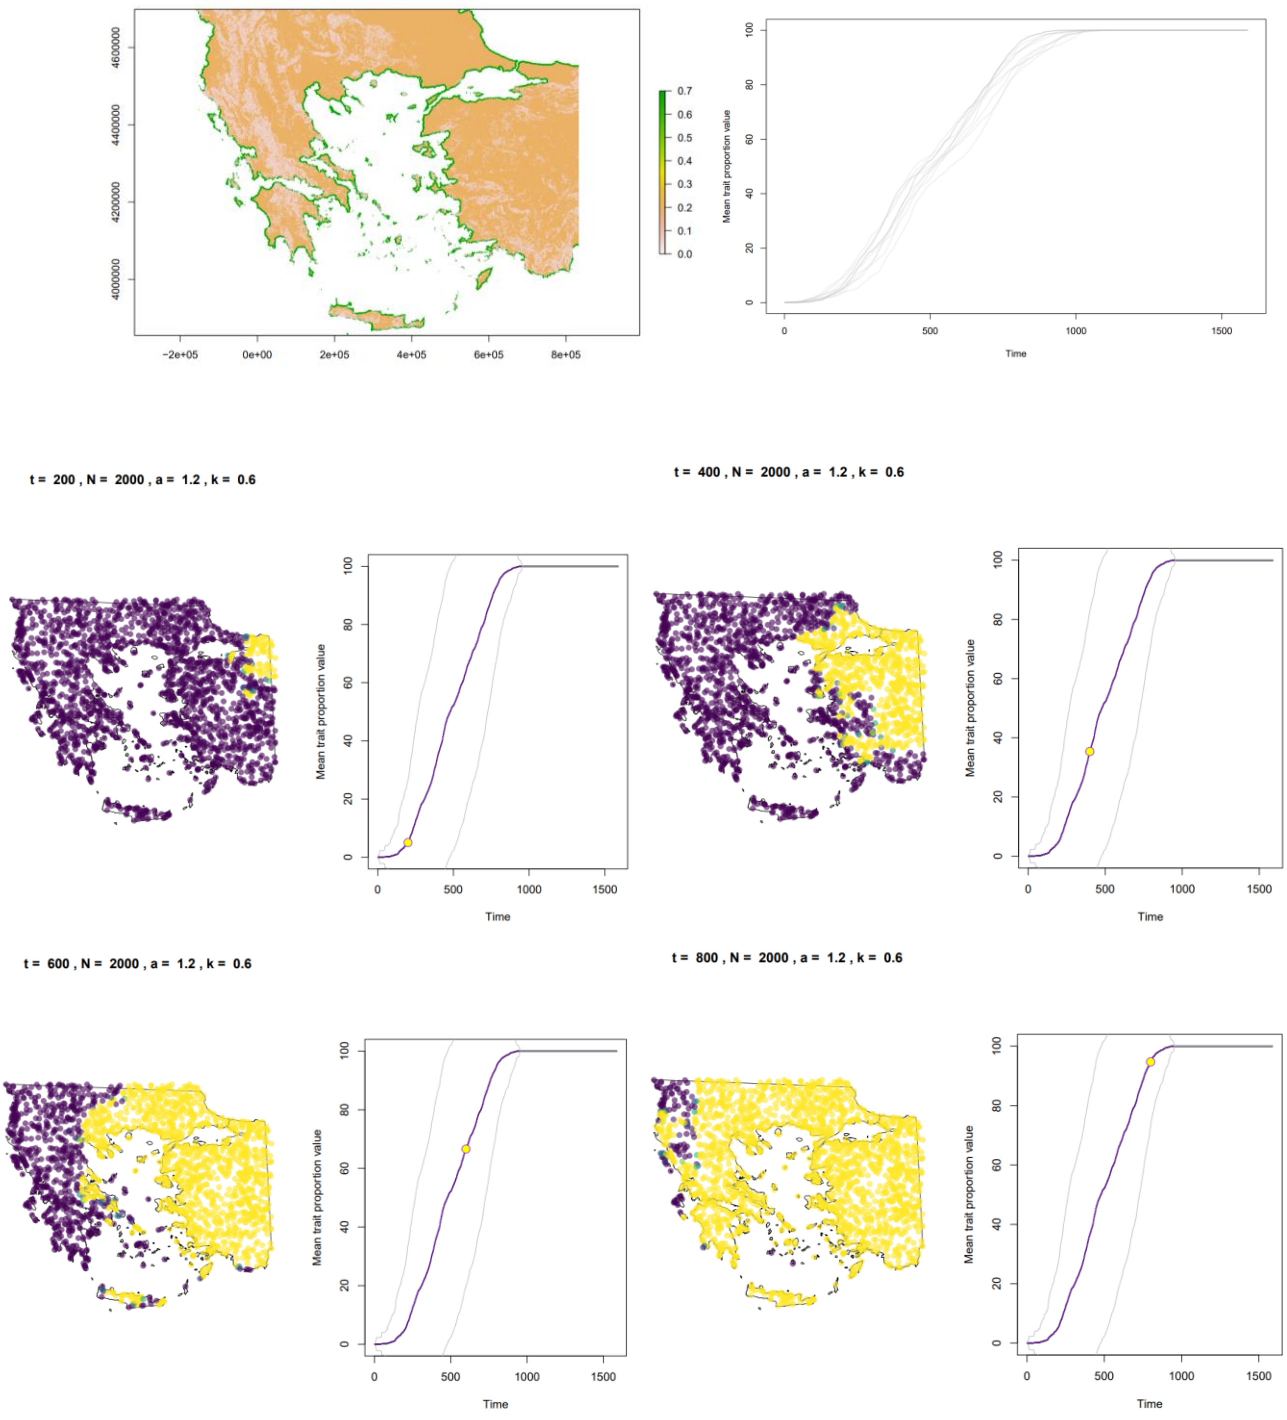

Figure 50: Study area divided into an inland flat (light orange), an inland steep (orange), a coastal flat (dark green), and a coastal steep (orange) zone. Nodes have a 0.70 probability of falling on the coastal flat zone, a 0.05 probability of falling on the coastal steep zone, a 0.20 probability of falling on the inland flat zone, and a 0.05 probability of falling on the inland steep zone (coastal steep and inland steep have the same probability). The top right panel shows the outputs of 10 simulations using the same parameter values and suggests that stochasticity does not cause significant changes in the shape of the curve. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

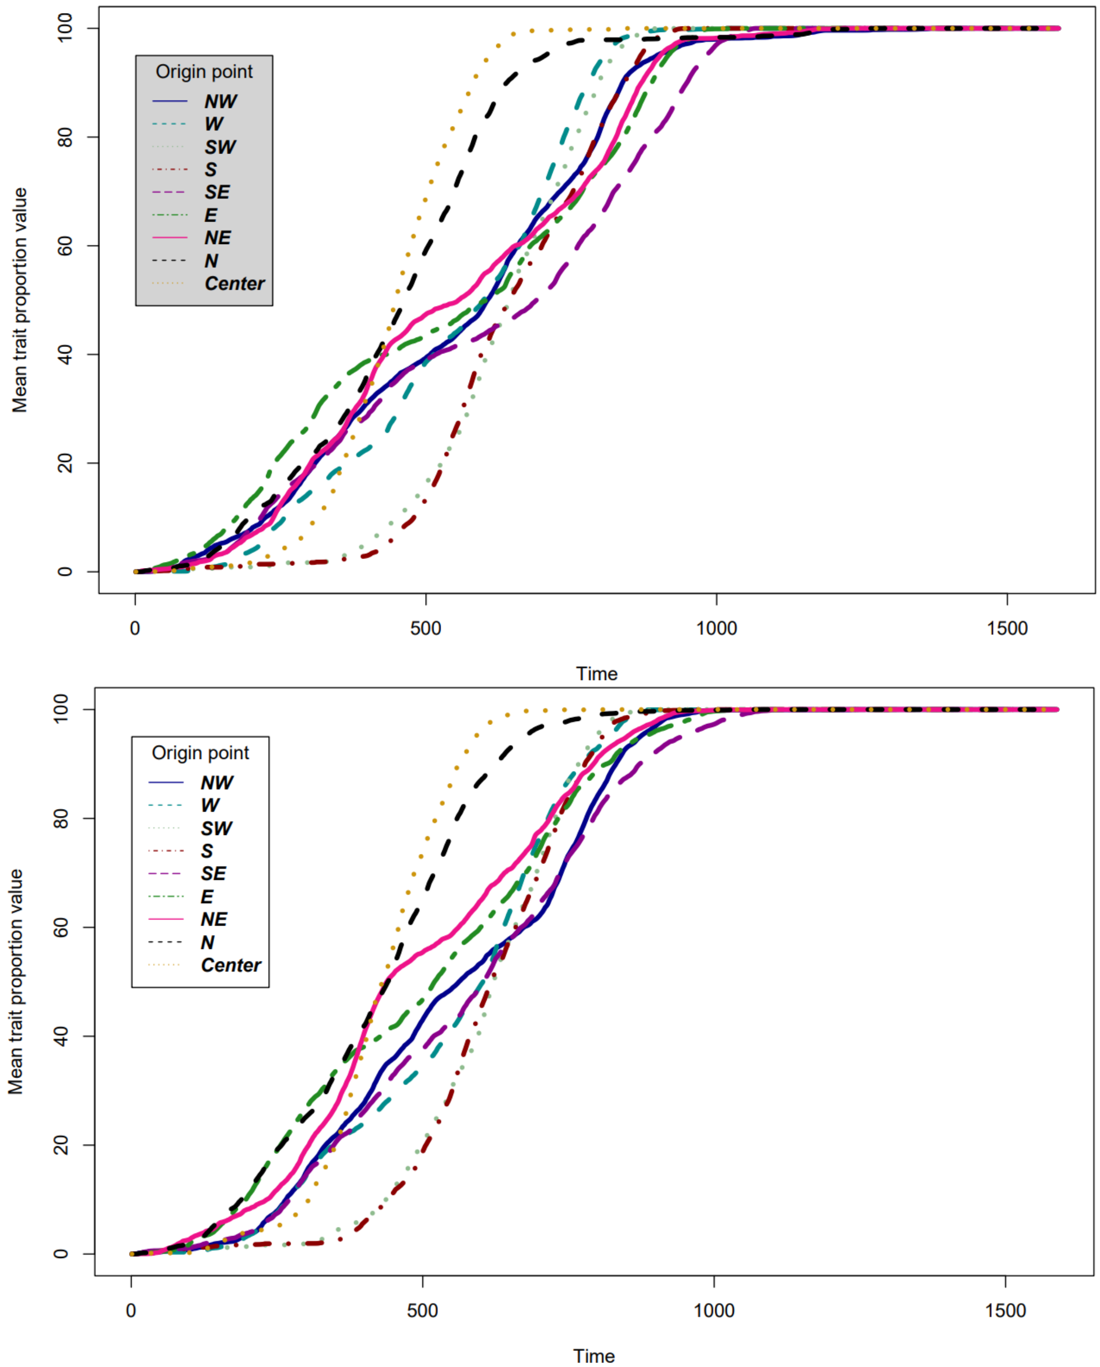

Figure 51: 51.1 shows the mean trait adoption per timestep for simulations with different origin points and k = 0.6, a = 1.2, N = 2000 with equal probabilities of habitation for the four landscape zones. 51.2 shows the mean trait adoption per timestep for simulations with different origin points and k = 0.6, globally defined a = 1.2, N = 2000 with different probabilities of habitation for each zone, with coastal flat habitation being more likely at .45, followed by inland flat habitation at .30, coastal steep at .15 and inland steep at .10. Note how the trend lines for origin points in the Center, North, South, and Southeast remain relatively similar when probabilities become unequal, but trends for all other origins become less similar to a double-S curve.

8.1.4. Experiment 4: Localised rate of local adoption a

It seems unrealistic that every community would have the same propensity towards adopting new traits. In a real-life scenario, one might imagine one community to be more open to adopting new ideas — for example, an urban settlement with a more diverse population of potters who might be more willing to adopt ideas coming from other locations due to a tendency of migrant potters to adopt techniques originating from their original communities — and, conversely, other communities that would be less willing to adopt the new trait and incorporate it into the learning process when training new potters.

Therefore, it would be fitting to re-formulate the model by incorporating a local adoption value a_i for each node and replacing the global parameter a with it. For each node i, a_i is drawn stochastically from a log-normal probability distribution.

The following formula describes how the trait value for each node/site is calculated for each timestep:

$$c_{it} = (c_{it-1} a_i) + \sum_{i=j}^{N} c_{it-1}\, w_{i,j}$$

$$w_{i,j} = \begin{cases} e^{-\kappa d_{i,j}}, if\ i \neq j \\ 0, if\ i = j \end{cases}$$

Where c_{it} is the innovation value at node i at time t, a_i is a node-specific variable determining the strength of local adoption, N is the number of communities in the study area, and $w_{i,j}$ is the weight between nodes i, j. $d_{i,j}$ is the Euclidean distance between nodes i, j. Self-weights for nodes are set to be equal to 0, but this does not mean that the model does not incorporate the effect of each node increasing on its own – this effect is captured by setting a_i > 1. If node i has a_i = 1 then the trait proportion value at that node does not change on its own. If a_i is lower than 1, local adoption is delayed.

To reiterate, for each node i, the case 0 > a_i > 1 describes cases where the new trait value decreases locally in the new timestep compared to the previous one, but this does not affect the influence of external nodes to the overall trait proportion value of the node. In cases where a_i = 1, the node retains the trait proportion value it had in the previous timestep - plus the contribution of other nodes. Finally, a_i > 1 describes cases where the node trait proportion value increases locally regardless of the contribution of the other nodes.

As aforementioned, a_i is drawn randomly from a log-normal probability distribution for the variable A. This prevents local adoption from having negative values, because negative trait proportion values make no sense in this case - no adoption/loss of the new trait is signified by a_i = 0. I expect the most common value to be a_i = 1, because the learning process (vertical transmission) that occurs between members in a community of practice of potters has been traditionally considered to favour faithful copying as a result of acquiring specific motor habits (Gosselain 2000: 192) – but note that in regards to copying shapes, it is sometimes the intention of copying the end result rather than "high-fidelity" copying of gestures that results in faithful copies (Gandon et al. 2020: 12). Cases where a_i > 1 could represent cases where the new trait is favoured due to the members of each node-community abandoning another trait in favour of the new trait – such cases will be treated as relatively rare, as will cases where 0 < a_i > 1.

Log-normal distributions are typically described by two parameters μ and σ, with μ being the location parameter, which would correspond to the mean of a normal distribution if the values in the x axis were logged, and the scale parameter σ representing the standard deviation of the logged values, respectively (Forbes et al. 2010). From an initial exploration of different values for μ and σ, it seems like μ = 0 and σ = 0.25 is a good starting point in terms of getting a probability density function that has no negative values, has a mode = 1, and where a values higher than 2 are rare — intuitively, it seems unlikely that cases where trait adoption doubled within one year (each timestep represents one year in the model) were common.

Figure 52 shows an example of stochastically drawn *a* values for a simulation with N = 2000.

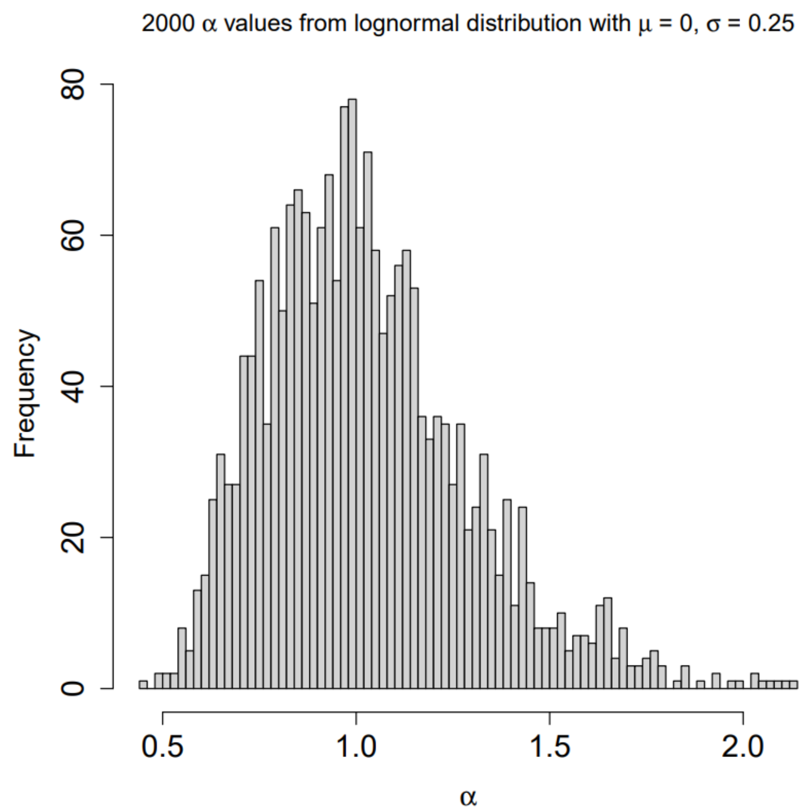

Figure 52: a values for N = 2000 communities drawn from a lognormal probability distribution with μ = 0 and σ = 0.25

Alternatively, one can relax the assumption that *a* cannot be negative and allow *a* to take negative values, indicating nodes where there is trait loss at each timestep. To avoid negative *a* causing meaningless negative trait proportion values, I have included a rule in the code implementation of the model where if a node has a negative trait proportion value, its trait value becomes 0. To allow for negative *a* values, I use a normal distribution with mean = 1, which allows modelling a population where high fidelity learning is still the most common tendency, but where there is a smaller or larger number of cases where there is trait loss or extreme trait adoptiveness — it is possible to control the degree to which each value is near the mean = 1 by changing the value for the standard deviation of the distribution. Unlike the log-normal distribution which is right-skewed, the normal distribution is symmetric around the mean. This means that values for *a* higher than 1 should be more common than when using a lognormal distribution. For an experiment with N = 2000, sampling from such a distribution produces *a* values shown in Figure 53.

In order to compare the two approaches, I conduct an experiment with two simulations using the baseline *k* = 0.6 and N = 2000, with the easternmost node as the origin point (with a value equal to 0.1), using the samples shown in Figures 52 and 53. Figure 54 compares the two simulations at timestep 400, where the simulated trends produce similar-looking double-s curves, but the simulation with the lognormally distributed sample for a reaches the equilibrium at a slightly higher value than the one with the normally distributed one. Strangely enough, the wavefront of the latter has already crossed the Marmara strait into Thrace and has already become better established in Crete than the one with the lognormal sample, even though it contains negative *a* values that can lead to local innovation loss. This is probably due to the fact that the log-normal distribution used for the simulation in Figure 54.2 is right skewed and has a lower proportion of values higher than 1 compared to the symmetric normal distribution used for the simulation in Figure 54.1.

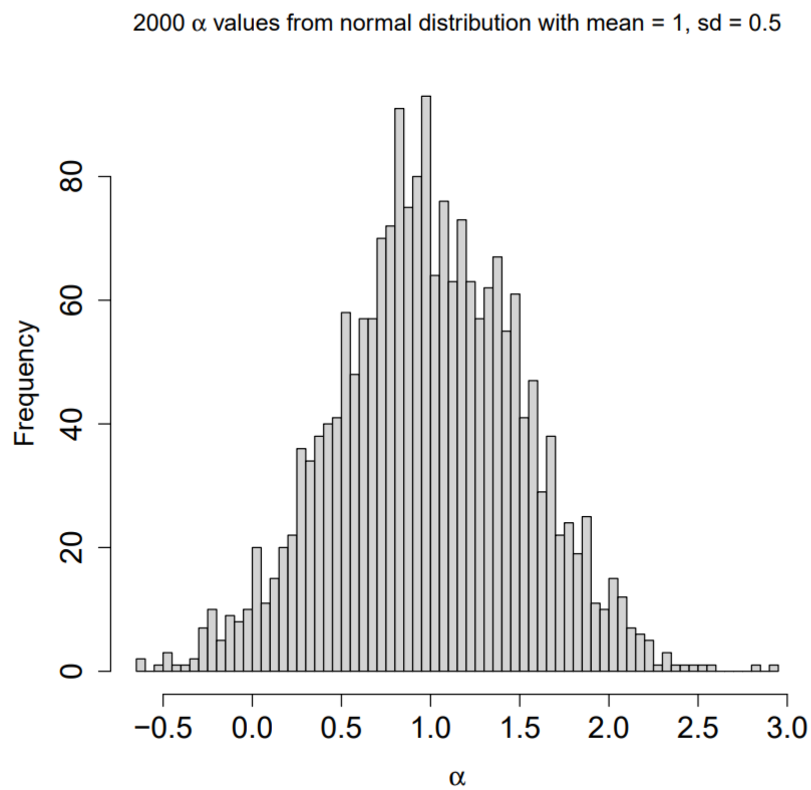

Figure 53: a values for N = 2000 communities drawn from a normal probability distribution with mean = 1 and standard deviation = 0.5.

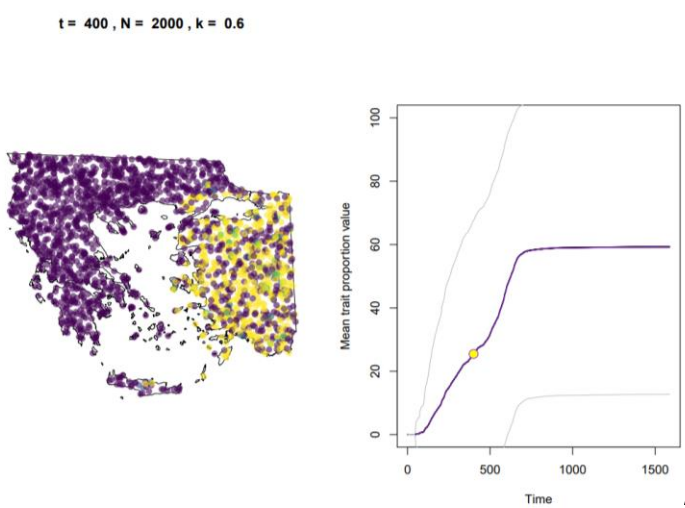

54.1.

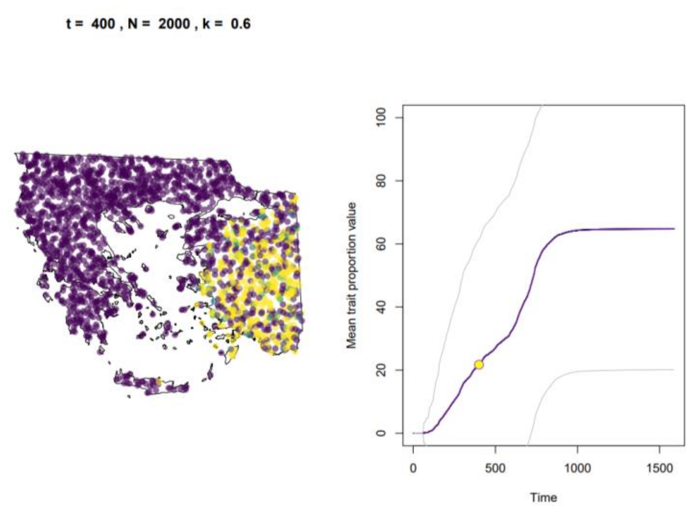

54.2.

Figure 54: 54.1, simulation with a values drawn from a normal distribution. 54.2, simulation with a values drawn from a lognormal distribution. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

8.1.5. Experiment 5: Effects of preferential communication with neighbours

The picture that emerges from the previous model with a globally defined local adoption parameter *a* is one where certain characteristics of the study area and the origin node's geography affect trait adoption trajectories significantly. However, given that for all sections apart from sections 8.1.2 and 8.1.3, inland terrain and boundaries like mountains were not included in the environment — but maritime boundaries were — the simulated process is largely driven by inland waves, and even though theoretically travelling between islands should help introduce the trait to a new landmass (e.g., mainland Greece), travelling between islands appears to be less important than communication between neighbouring communities/nodes on the landmasses. This effect was somewhat mitigated by increasing the density of communities around the coast, but it resulted in inland coastal travel being the main driver — not hops between islands. While some degree of coast-hugging should be expected in the case of Early Bronze Age travel using canoes and longboats (Broodbank 2000), it is safe to assume that after the advent of sailing technologies open sea travel between islands was preferred to extended coast-hugging journeys which would have been more time-consuming and would protect ships from crashing into rocks when weather conditions were adverse.

In reality, it is possible that neighbour preferences affected local decisions on who to interact with to some extent, especially in the case of small islands with only 1-2 communities. The main assumption here is that when resources that enable travel are limited, a cost-benefit assessment of who to interact with is to be expected. A neighbour ranking would be useful in such cases — if one has some knowledge of the local geography, they can visit the neighbours closest to them, preferring to visit the nearest one, and then the second nearest, and so on. The more limited the resources at each node, the smaller the number of preferred neighbours one can visit. However, since I have not included information on the resources of each community in the model, they will not be incorporated in the model presented in this section.

Some Aegean-based network models like Broodbank's (2000) proximal point analysis have used a set arbitrary number of nearest neighbours for every node. This might work well for sub-regions where there are small islands separated by large distances, but the present study area includes subregions with large landmasses, where the assumption of nearest neighbour preference is less likely to hold true. One way to incorporate neighbour preference and keep the isolation-by-distance aspect of the process (with a non-linear effect on distance captured by *k*) is to produce a new set of weights that capture the effect of neighbour ranking and multiply them by the weights that were used in the previous models.

The following formula describes how the trait value for each node/site is calculated for each timestep:

$$c_{it} = (c_{it-1})a + \sum_{i=j}^{N} c_{it-1}\, w_{i,j} r_{i,j}$$

$$w_{i,j} = \begin{cases} e^{-\kappa d_{i,j}}, if\ i \neq j \\ 0, if\ i = j \end{cases}$$

$$r_{i,j} = \begin{cases} \dfrac{1}{n_{i,j}}, if\ i \neq j \\ 0, if\ i = j \end{cases}$$

Where c_{it} is the innovation value at node i at time t, a is a globally defined parameter determining the strength of local adoption at every community/node, N is the number of communities/nodes in the entire study area, and $w_{i,j}$ is the weight between nodes i, j , $d_{i,j}$ is the Euclidean distance between nodes i, j. Self-weights for nodes are set to be equal to 0, but, as in previous sub-models, this does not mean that the model does not incorporate the effect of each node increasing on its own — this effect is captured by setting a value for the parameter a that is higher than 1. If a is equal to 1, the model corresponds to the previously discussed model, and the trait proportion value at each node does not change on its own. If a is lower than 1, local adoption is delayed. The new addition to the model is the new set of weights $r_{i,j}$ which is the reciprocal of the rank $n_{i,j}$ of each neighbour j for each node i. $n_{i,j}$ is produced by ranking each nodes neighbours in descending order on the basis of their Euclidean distance $d_{i,j}$. The following section explores a simulation using the same parameter values used for the simulations included in previous sections: k = 0.6, a = 1.2, N = 2000, using the easternmost node as the origin point with a starting value equal to 0.1.

An experiment using only rank weights is included in Appendix 3 and explores the effect of those weights in isolation of the other parameters in the model.

Outputs

The simulation outputs produce a similar trend based on the mean trait adoption value per timestep to the one observed for the main model with a globally defined local adoption parameter. In terms of its shape, the double-S curves associated with geographic bottlenecks are also encountered here (Figure 55), meaning that the process is still mainly driven by large waves on the landmasses.

The outputs of the new model differ slightly in certain qualitative aspects from the outputs of the previous model. As shown in Figure 56, the trait now reaches Crete faster through its island connections to Western Anatolia compared to when it reaches other nodes, where the trait had already become well established in the Peloponnese and the Ionian by the time it got to Crete. This suggests that the new model succeeds in adding more realism in the way the trait spreads between nodes in an archipelago on a local scale, without changing the overall picture that emerges from the "global" geography of the region — e.g., the importance of two large landmasses separated by the sea.

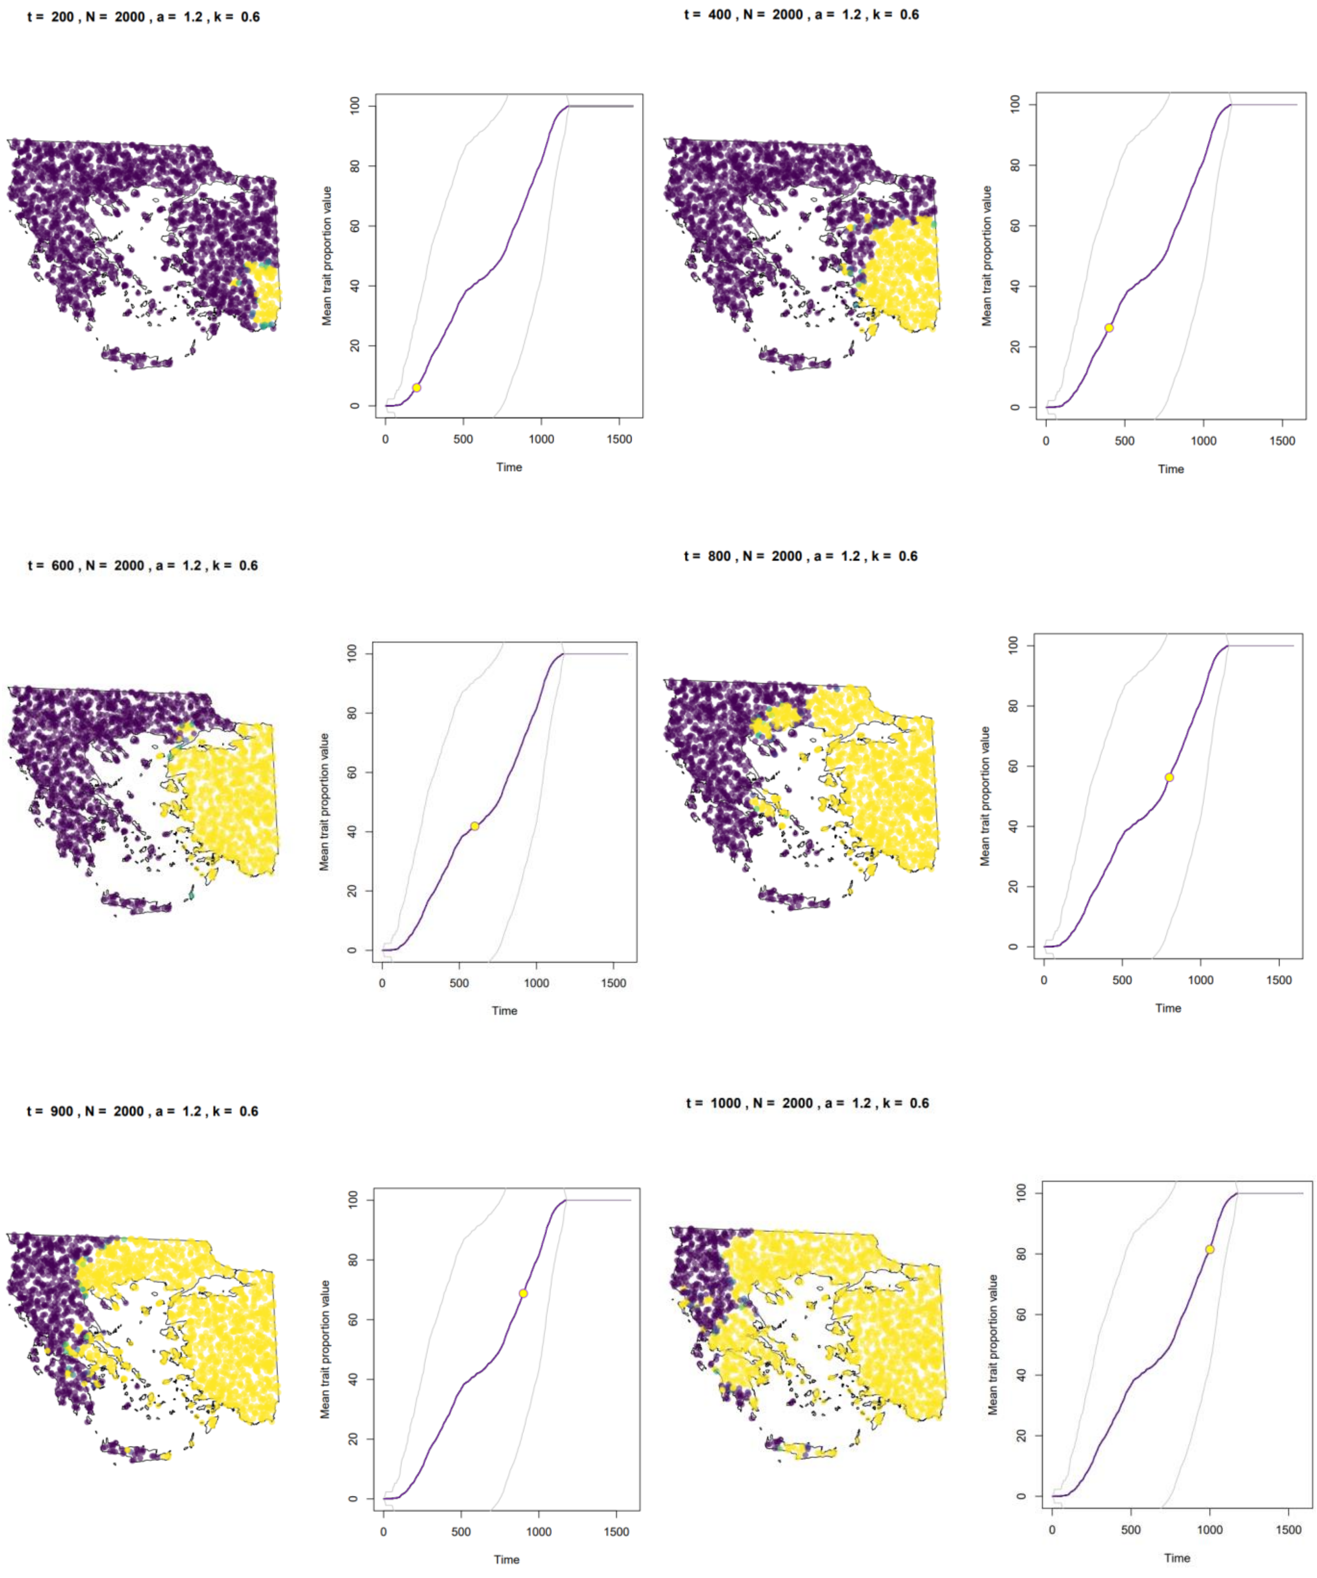

Figure 55: Simulation including a globally defined local adoption rate parameter a and a new set of weights determined by the Euclidean distance ranking of nodes. Trait adoption at timestep 200, 400, 600, and 800, 900, and 1000. k = 0.6, a = 1.2, N = 2000. The easternmost node is the origin of the trait, with a starting value equal to 0.1. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

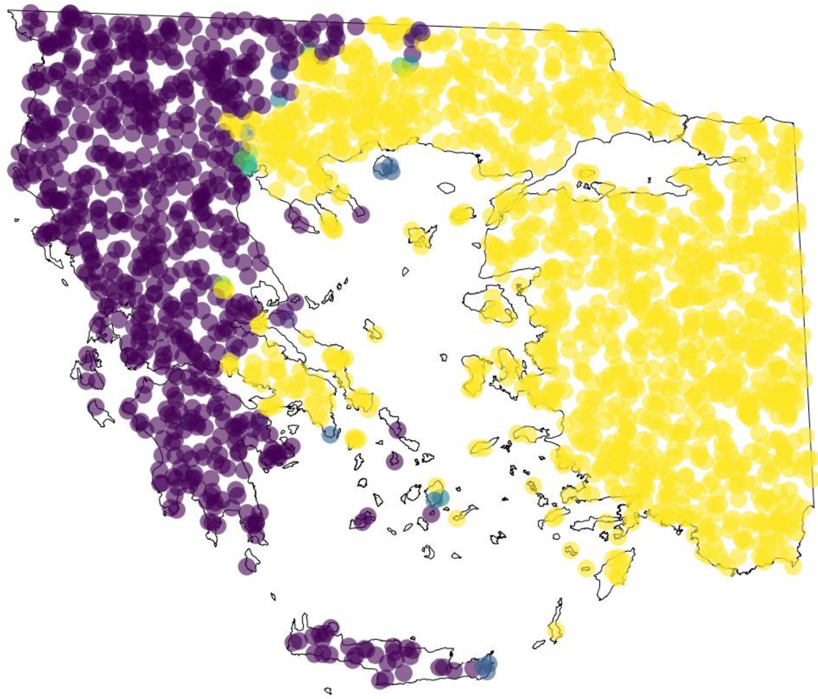

56.1.

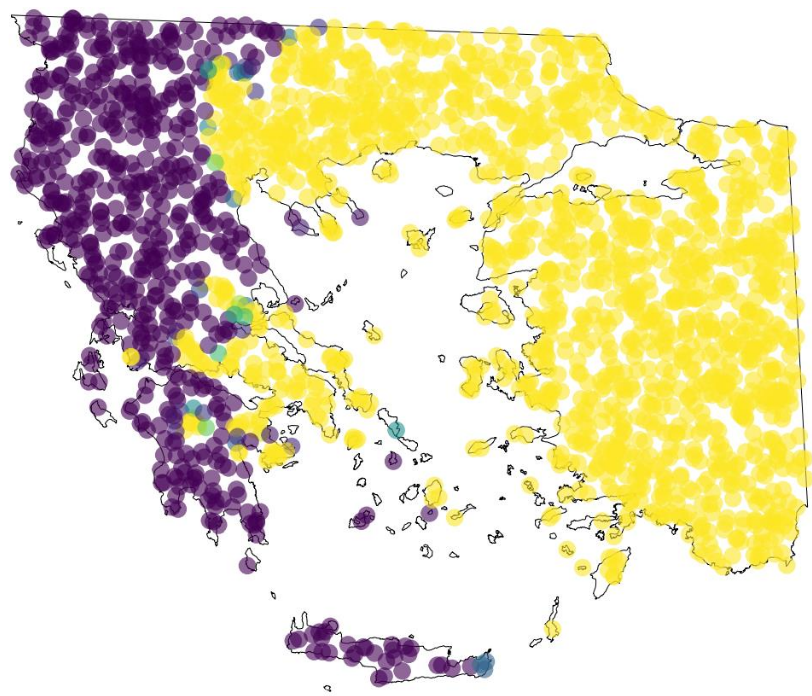

56.2.

Figure 56: Figure 56.1, simulation in Figure 55 at Timestep 860 – considering node ranking- when the trait first reaches Crete. Figure 56.2, Simulation using the same parameter values as in the simulation in Figure 56.1 but without considering node ranking-shown here at Timestep 815 when the trait first reaches Crete. Note that even though the trait reaches Crete slightly later in the first example, it does so before reaching the Peloponnese and the Ionian.

8.1.6. Discussion

The simulations explored in the previous sections throw light on quantitative and qualitative characteristics of the transmission process. In terms of their shape, the curves produced are single and double-s curves, with the latter being the result of the diffusion process' reliance on terrestrial travel/communication that is facilitated by local clusters of communities. This reliance is associated with episodes of slower adoption when the travelling waves reach geographic boundaries like the sea or a narrowing of the stretch of land that is located next to the wavefront. This suggests that, under this model, geography is an important factor affecting the speed of system-wide trait transmission.

Experiment 1 explored how the location of the origin node affects global trajectories, and it appears that geography also plays an important role in this respect. The origin node's position on the East-West axis can produce qualitatively different diffusion trajectories even if it does not affect when the global trend reaches the 100 % equilibrium — in other words, when the trait starts spreading from Crete it can reach the point of being adopted by everyone in the population at the same time as in the case where it starts from Western Anatolia, but it will appear as a more sudden uptake from the perspective of global adoption. Double-s curves were observed for origin nodes located on the far west and far east of the study area. It seems like the "sweet spot" for an origin node to lead to the fastest adoption is if it is located halfway between the east and the west — note the steep s-curves for the North and Central nodes. It also helps if such central nodes — central with respect to the east-west orientation — are located on a landmass (e.g., North) or on a well-connected island near the mainland (e.g., the central node). When such nodes are located on isolated islands like Crete, there is a prolonged phase of slow global adoption until the trait reaches the coast of mainland Greece and Western Anatolia. The importance of the East-West axis is possibly due to the fact that this is the axis on which the main landmasses are defined. There seems to be no great difference in transmission trajectories when origin nodes are located on opposite sides of the North-South axis — for example, the trends produced by the Southern Anatolian (SE) and Northern Anatolian (NE) origin node are remarkably similar.

Experiment 2 explored the effect of locational preferences for different landscape types on global adoption using an extreme coastal preference as a case study. The simulation output suggested that the reason behind the double-s curves in the previous examples is indeed the reliance on inland travel: when the population density is higher around the coast, the trait adoption process develops into a coast-hugging wave that avoids the bottleneck associated with the trait having exhausted the area of the landmass of the origin node. A similar process was observed for the simulations in experiment 3, which involved incorporating a distinction between flat and steep zones (above 10 degrees of slope). When this distinction was made with a medium-high coastal preference, the trend produced was a double-s curve associated with a bottleneck when the wavefront reached a steep zone. However, this effect was not observed when coastal preference was high. Experiment 4 showed that when there is variation in the values of *a* (a different value for each community) this lowers the equilibrium value of the system, but it does not change the shape of the curve. Experiment 6 showed that including a weight that rewards communication between nearest neighbours has no strong quantitative or qualitative effects on the transmission process.

Overall, the experiments point to the adoption process' strong reliance on geography, even when the local adoption rate *a* (globally defined to be the same for every node) is high.

8.2. Quantitative comparison of simulation to empirical data (no trait loss model)

Previous sections explored simulated outputs through a visualization-based approach, where the simulated curves were described in terms of their qualitative characteristics. For example, curves were described as S-shaped or double S-shaped curves and little attention was paid to the exact timing of arrivals at different regions, or how the simulated outputs compare to the empirical trend on a quantitative basis. This section moves in the direction of a more formal quantitative description of the fit between simulation and reality. The question is, what defines a good fit?

One idea of a good fit is one where, between two curves (one representing the empirical trend and one representing the simulated trend), there is an exact match between the value on the x-axis and the value on the y-axis for each point the curve is passing through. The problem is, however, that given how uncertain the value on the y-axis is — note the width of the confidence intervals for the empirical mean in Chapter 5 — it is very likely that the 'true' empirical value is not exactly captured by the empirical trend, especially when sample sizes are so small, and that the simulated fitted curve describes the small sample that is available and not the overall process that one is interested in. Were new data points to be added to the dataset, such a model might stop being a good fit (this is sometimes called 'overfitting').

Another approach to getting a good fit is one that is based on a timestep-by-timestep comparison between the empirical and simulated trends — this has advantages and disadvantages and is explored in sub-section 8.2.1.

It is also possible to compare the shapes of the empirical and simulated trends: while in previous sections curve shape was treated as a qualitative characteristic, it is possible to summarise the shape of a curve by fitting a known function to it. This function is defined by certain parameters, whose numerical values can be compared (subtracted) to get a quantitative estimate of the difference between the shape of the empirical and the simulated trends. This method is explored in sub-section 8.2.1 combined with a Monte Carlo approach to incorporating temporal uncertainty (as explored in Chapter 5).

Finally, apart from the system-wide characteristic of the diffusion process that is captured by the mean trait proportion value per timestep, there are micro-level details that characterize the diffusion process. For example, while trend lines for the mean might be similar between two simulations, the time at which the trait arrives at different regions may differ, as shown by the simulations in the sections above, where adding a nearest-neighbour favoring weight changes arrival times but not the overall trend. Since there is empirical information on when the wheel-made pottery trait arrived at different parts of the Aegean, one can compare the empirical arrival times with the arrival times from simulations for different parameter values, identifying combinations of parameters that produce the closest fit to reality. Section 8.2.3 develops such an approach, again in conjunction with a Monte Carlo method to address within-phase uncertainty for the empirical data.

8.2.1. Phase-by-phase measures of fit between simulations and the empirical trend

When trying to compare between two trend lines composed of discrete measurements at different phases on a phase-by-phase basis, as in this case, there are two aspects that can be taken into consideration. First, the value that both trend lines have per phase and two, how quickly each trend value changes from one phase to the next, which describes the rate of change of the trend.

Therefore, a first measure of fit between the simulation and the empirical trends is to compare the y-axis value of each one per phase, and then sum the values of each ‘comparison’ to get an overall picture of how similar the trends are for the entire observation period. This can be done by subtracting the y-axis value of the simulated trend from the y-axis value of the empirical trend, taking the absolute value of that, repeating for every timestep, and then taking the sum of those absolute values. The lower the value of the sum, the better the fit of the simulated trend to the empirical one.

A second measure of fit is to measure how quickly values change from one phase to the next for each trend (the per phase slope of the simulated and empirical trends), perform a comparison between the two values on a per phase basis, and then sum the values of each “comparison” to see how similar the trends are over time. This can become slightly more complicated than the calculations done for the first measure.

When each phase has a different duration, the duration of each phase will partially determine the slope of the trend from each phase to the next. Therefore, for each phase, I have used the common formula for the slope:

$$\frac{dy}{dx}$$

Where for each phase t, dy is the difference between the value at time t and the value at time *t-1*, and dx is the length of the time interval from *t-1* to *t*, and the value of the fraction represents the slope. For each segment, the absolute difference of the slope for the simulated and the empirical trend was calculated, and the sum of those values is a measure of how well the simulated trend resembles the empirical one in terms of the rate of change over time and the shape of the curve. Again, the lower the value of the sum, the better the fit of the simulation to the data.

Finally, it also possible to produce a third measure of fit that combines the raw-value based and the slope-based measures into a composite measure: a sum of the two measures, which can help identify simulations that produce similar outputs to the empirical trend in both aspects: raw values and slopes. It should be noted that in order to compare like and like, both sets of values — sums of absolute differences and sums of absolute differences in slope — were normalised to a 0-1 range using minmax normalisation.

Figure 57 shows the results of parameter sweeps for different combinations of values for *k* and *a* in a population with N = 2000, with randomly placed nodes, and where the origin node is the easternmost node in the population. As evident from the parameter space maps for the raw value (Figure 57.1) and the slope measure (Figure 57.2), the simulation that fits the empirical trend best (*a* = 1.6, *k* = 1) in terms of raw values is not necessarily the one that fits it best in terms of the least difference in slopes — this is the simulation where a = 1.4, k = 1.4. In the composite measure (Figure 57.3), the best fitting combination is a = 1.4, k = 0.6, which is also being indicated as a good fit by the slope measure.

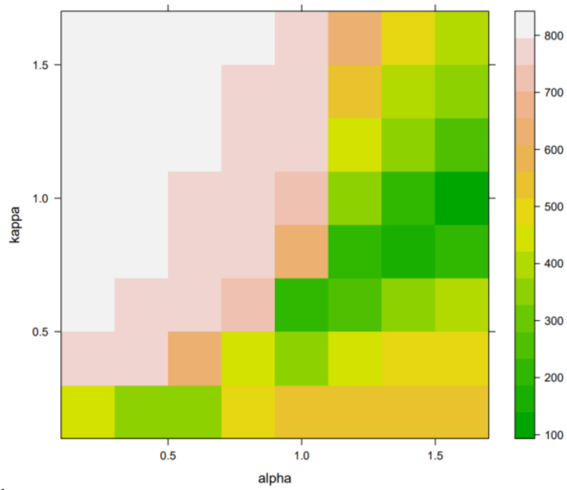

57.1.

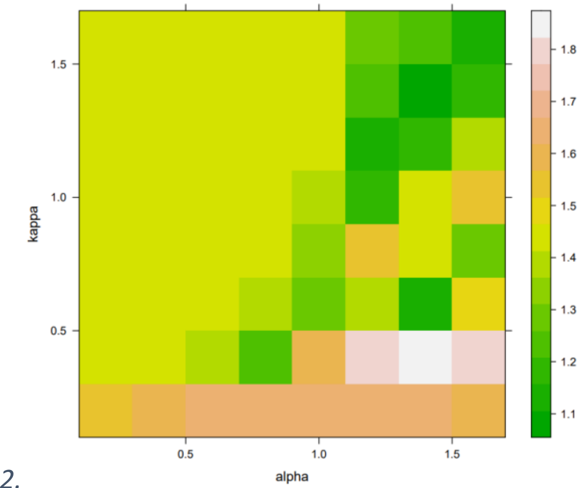

57.2.

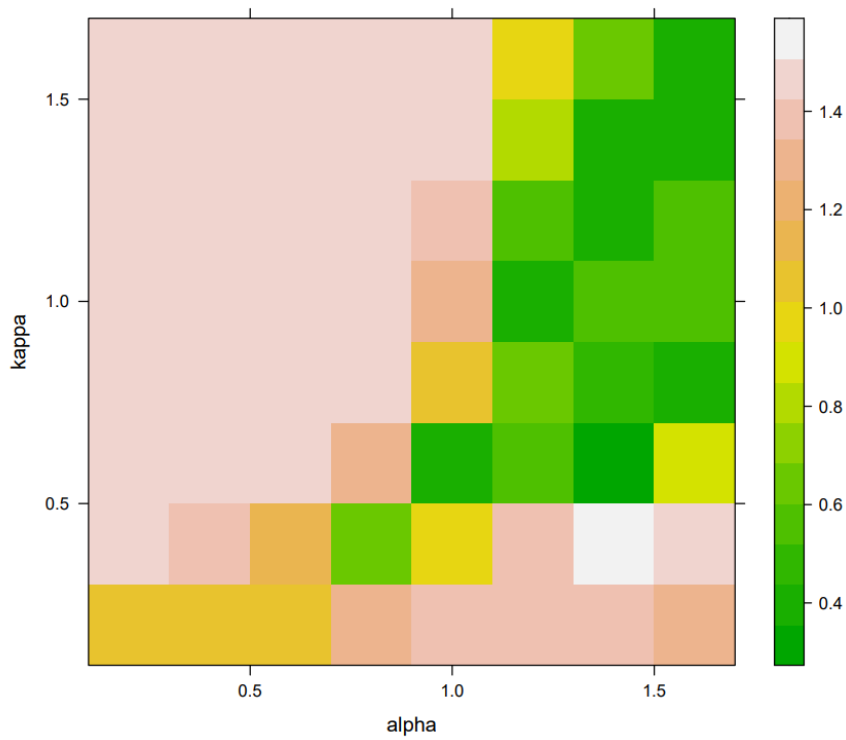

57.3.

Figure 57: Parameter sweeps for a and k, range for a: 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, and 1.6 and for k: for a: 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, and 1.6 – 64 combinations in total. 57.1 shows values for the sum of absolute differences between the simulation produced for each combination of parameter values and the empirical pattern shown in Figure 18. 57.2 shows values of the sum of absolute differences in slope between each simulation and the empirical pattern. 57.3 has values for a composite measure of fit, where each value is the sum of the normalised (minmax normalisation) sums of absolute differences and the normalised sums of absolute differences in slope. Lowest values in dark green represent the best fit for each measure.

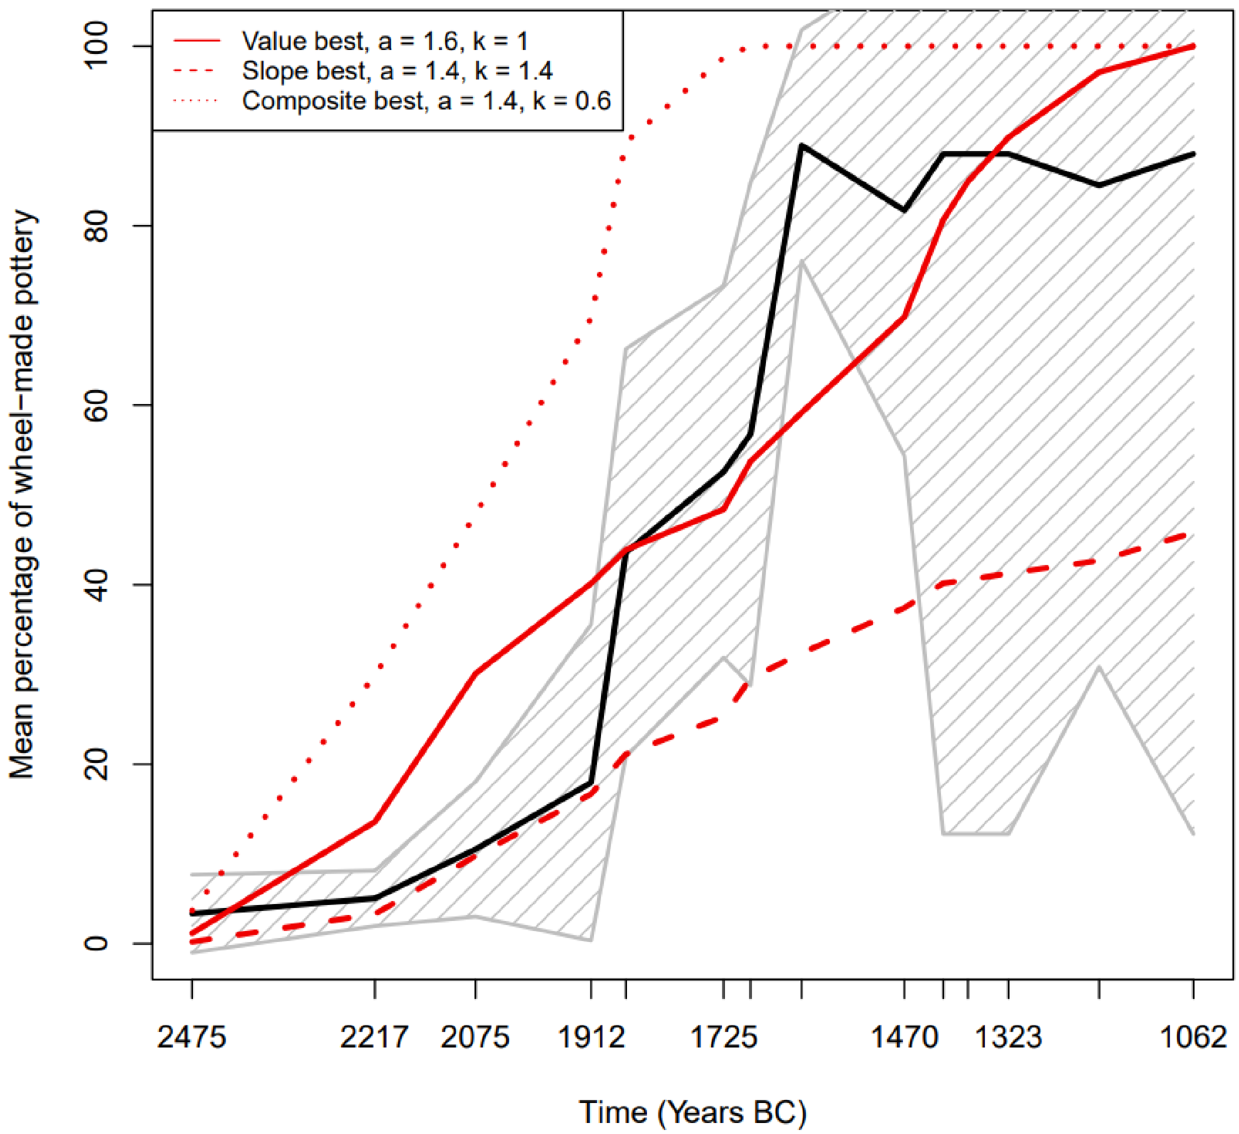

Figure 58

When comparing the three combinations of parameters indicated by each measure as producing the best fit (Figure 58), it is evident that, at least when using simple visual inspection, the trend produced by the raw value measure performs the best over the entire study period. However, this trend does not produce the initial slow adoption curve indicated by the empirical trend, and for most time intervals, the value on the y axis falls outside the confidence interval. On the other hand, the simulation indicated by the slope measure as the best fit follows the empirical trend very closely and within the confidence interval for the first three phases, but then deviates from the steep curves at the time intervals 1912-1863 and 1687-1615, the two intervals that represent the episodes of fastest adoption. Finally, the simulation indicated by the composite measure arguably performs the worst compared to the other two, suggesting a need for a different type of composite measure.

8.2.2. A model-fitting and Monte Carlo simulation approach to comparing empirical to simulated data

The methods developed above have few advantages and significant disadvantages. One of the advantages is that comparing the outputs of simulations to the empirical record can be done on a qualitative basis, through simple visual inspection: in most cases it is possible to overlay the simulation indicated as the best fit by different measures over the confidence envelope to see where each simulation deviates from the interval that one can be 95 % confident the "real" mean is located in. However, the measures either failed to identify cases where the simulated trend is qualitatively similar to the empirical one (the "raw" value measure) or they were so sensitive to noise-errors that they only managed to find good fits for simulations that were qualitatively similar to the empirical trend in the first phases of adoption, but not beyond (the "slope" measure).

More importantly, as discussed in Chapter 5, this type of confidence interval for the empirical trend only takes into consideration the uncertainty arising from using small sample sizes to estimate the mean per phase — it says nothing about the temporal uncertainty inherent in the empirical trend. We already know that each length has a different phase, but we do not know what types of periodization processes resulted in the aggregated measurements we have for each phase — and taking phase length into consideration when constructing the envelope did not fully address the problem: it was possible to identify phases where the high level of uncertainty around the mean was countered by the small length of the phase (meaning that a "bad" estimate of the mean for that phase would not greatly affect our estimate over the entire study period), but it was impossible to compare the empirical trend to the simulated one without resorting to a method that aggregates the results of the high temporal resolution simulations to produce coarser-grained curves, "masking" the concavity of the curve during first stages of adoption in the process. This effect was particularly pronounced because of the fact the first few empirical phases are some of the lengthiest in the study period. This suggests a need to move away from phase-to-phase comparisons between the simulated and the empirical trends, towards a method that compares the values of the parameters that define the shapes of the curves themselves.

It has been established that the outputs of the simulations frequently resemble logistic curves for most combinations of a and k — and when they don't, they should probably be considered a bad fit given the similarity of the empirical trend to a logistic (S-shaped) growth curve — as shown in Chapter 5. Logistic growth is a process that starts with a period of slow growth, followed by an exponential growth phase where the growth rate is at its highest, and then a tapering off when the process approaches an equilibrium value. In the context of trait adoption and innovation diffusion in general, the equilibrium can characterize the limit of the population, for example, cases where everyone has adopted and there is no one left to adopt in a community — in this case the equilibrium value would be 100 % — or cases where there is some other reason why the trait does not exceed an equilibrium value. These general characteristics hold true for most of the simulated S-shaped trends, but sometimes trends differ in how steep the S-curve is, how long the tail is, and what the equilibrium value of the system is.

Therefore, a practical way to describe their shape in a robust manner that allows quantitative comparisons to the empirical trend is to fit logistic curves to both the simulated trends and the empirical one. It is then possible to extract from the fitted curves the parameter values that define their shape: the asymptote representing the equilibrium and the scaling parameter, which roughly describes how quickly the process reaches the equilibrium using the same method that was used in chapter 5 to describe the empirical trend and the Monte Carlo simulations derived from it. This allows comparing the scaling parameter of each simulation output (shown in Figure 59.1) to the scaling parameters of the 1000 Monte Carlo simulations of the empirical trend taking into account temporal uncertainty (section 5.2.2). The parameter space map in Figure 59.2 allows comparing how good a fit each simulation is to the observed data: each cell represents the mean of the differences between each simulation arising from a unique combination of values for *a* and *k* and the 1000 Monte Carlo simulations (reflecting temporal uncertainty) associated with the empirical values. There is a significant number of cells that have no values assigned to them (shown in white) — reflecting cases where it was impossible to fit a logistic curve to the simulation outputs and should be treated as cases that represent a bad fit to the empirical data.

Overall, the region where both the local rate of adoption and communication costs are low seems to produce the worst fits to the empirical data, although there are cases where very low communication costs can counteract the low rates of local adoption (k= 0.2, a = 0.2, 0.4) and produce a relatively good fit to the data, at least compared to other combinations of parameter values. The region that produces the best fits is the one where *k* is low (0.4 and 0.6), and where the rate of local adoption *a* is medium to high (0.8 to 1.6). In fact, the best-fitting simulation using this method of comparison is one where trait proportions do not increase or decrease on their own (a = 1, k = 0.4) and where communication cost is low.

The special case where a = 1 has been discussed previously as a case that represents a lack of a local component affecting the rate of adoption at each community. In such cases, each community transmits a trait without increasing the number of learners that are taught at any given time period (which would be the case for a < 1) and the transmission strength is not affected by a local factor which sets a limit to the number of potters that produce pottery with the trait of interest, such as local elites (the case where a > 1). A similar tendency is well-established in the ethnographic literature on the transmission of forming techniques at the level of the community, which suggests that forming techniques are typically transmitted by a process of vertical transmission between teachers and learners, without a community level factor affecting this process (as discussed in chapter 2). Therefore, the fact that the simulations indicate this type of transmission as the most likely mode of social learning that produced the empirically observed picture agrees with the literature and could serve as a preliminary confirmation that the formulation, ontological definition, and execution of the simulations is valid and can generate useful interpretative suggestions that can complement current archaeological discussion on the adoption of the potter's wheel in the Aegean.

For example, specialists suggest that in many cases the local reaction to the new trait was to reject it, or at least not to immediately replace the current mode of production (hand-building techniques) with the wheel, and that the "success" of the trait in later periods should not be associated with a local enthusiasm to adopt it (e.g., Knappett 2016). The results presented here partially agree with this verbal assessment: it is rare to see a case where low local receptiveness to innovation produces a picture similar to reality, but if communication costs are low enough, it is possible. The main takeaway from this experiment is that it is not necessary for local receptiveness to be high for the potter's wheel to spread in the Aegean in the way that it did and over a long period of time. All that is required is that local communities of potters receive the new trait, retain it for some time (roughly one year) and remain open to external influence (low communication costs afforded by low k values in this model). Low communication costs can result from having efficient means of travelling, but they can also emerge when settlements are located close to each other — as k penalizes long distances over shorter ones. In the context of elite-imposed limitations in the number of potters producing wheel-made pottery or in the sheer quantity of wheel-made pottery (which can be captured by setting $a < 1$), this means that, if the model accurately reflects the real-world process, the observed empirical pattern of the adoption of the wheel in the Bronze Age Aegean was not caused by the decisions of elites.

Interestingly, there are two regions of the parameter space map that produce good fits to the observed data, albeit the former generally produces better fits than the latter: a region where communication cost is low and local receptiveness is medium to high, and another region where communication cost is high and local receptiveness is medium to high. To complicate things further, there is a region diagonal to the axes that produces simulations with a "bad" fit, where k = 0.6 and a = 0.8 and 1, k = 0.8 and a = 1, 1.2, and 1.4, and where k = 1 and a = 1.2 and 1.4.

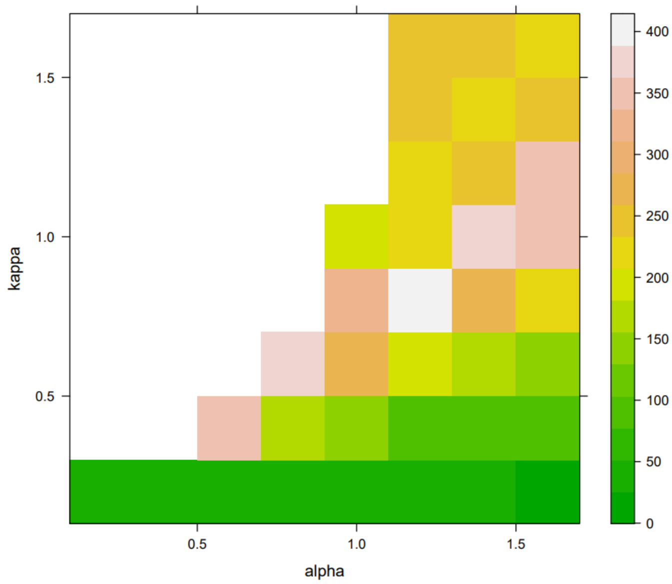

59.1.

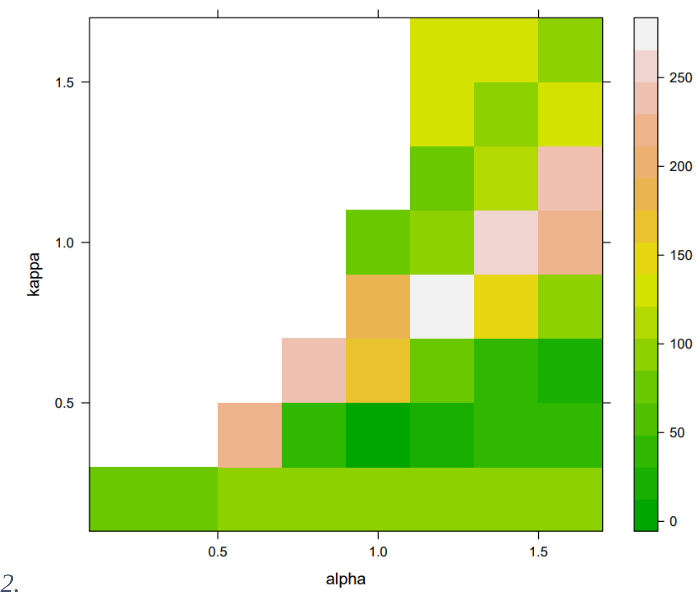

59.2.

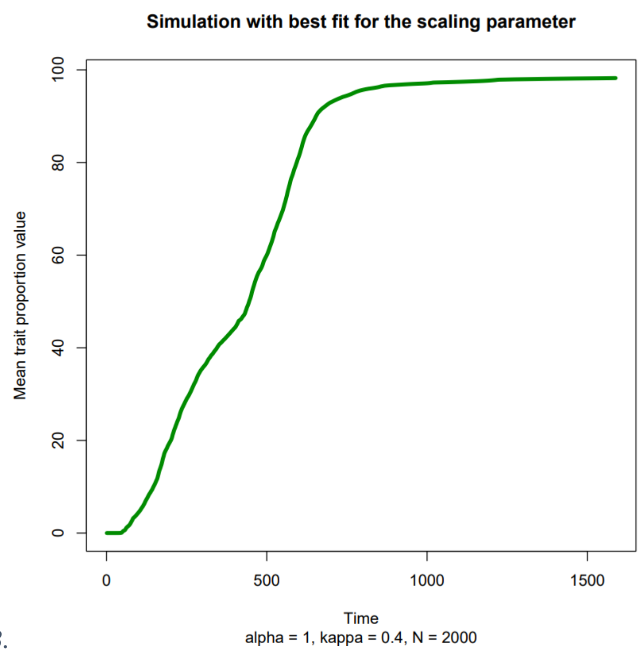

59.3.

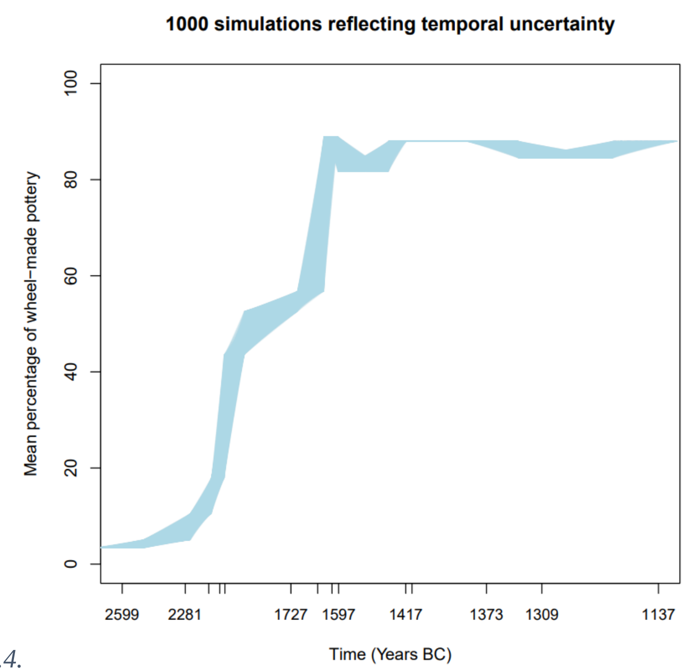

59.4.

Figure 59: 59.1, Parameter space map where each cell holds the value of the scaling parameter for a logistic curve fitted to a simulation output (64 different combinations of a and k, N = 2000). 59.2, Parameter space map where the value of each cell is the difference between the scaling parameters shown in the top left panel and the mean of the scaling parameters of the logistic functions fitted to each one of the 1000 Monte Carlo simulations based on the empirical data (reflecting temporal uncertainty, 19.4). 59.3, the simulation output indicated as the best fit to the observed data, where a = 1, k = 0.4, N= 2000.

8.2.3. Comparing simulated arrival times to empirical arrival times per region

As discussed previously, a simulation can fit the empirical record in terms of system-wide dynamics like the mean but fail to capture micro-level characteristics of reality such as the time at which a trait arrives at known locations. Capturing such micro-level details is particularly important in the case of models that are formulated with the intention to capture the geographic aspects of innovation diffusion, such as the one explored here.

The arrival times for Crete, Euboea, and Mainland Greece for the 64 combinations of *a* and *k* in the simulations explored above are shown on the left panels of Figures 60, 61, and 62 respectively. A trait is considered to have arrived at a region when there is at least one node with a value higher than or equal to 0.1. For all regions, there is a large region on the parameter space maps that results in the trait arriving very late (low *a*, and medium *a* with high *k*). Those simulations are sometimes cases where it was not possible to fit a logistic curve to the trend. In most cases, the lower the *k* and the higher the *a*, the earlier the arrival of the trait at each region, which is to be expected given that these cases represent extremely fast trait diffusion. Overall, Crete appears to receive the trait relatively later compared to Euboea and Mainland Greece, and Euboea gets it slightly later than Mainland Greece.

When comparing the simulated arrival times to 1000 Monte Carlo simulations for the empirical arrival times - see Figures 60.2, 61.2, and 62.2, where each cell represents the mean difference between each simulated arrival time and the 1000 Monte Carlo simulations for the empirical arrival time-, the region with low *k* (around 0.5) and *a* around 1 produces the best fits to the observed arrival times. Crete is an interesting case in that there are three combinations of *a* and *k* that produce an equally good fit. However, the fact that the same *a* and *k* combinations that produce a good fit in terms of arrival times also produce a good fit in terms of the macro-scale mean, as discussed in the previous sub-section, suggests that the simulations succeed at capturing the characteristics of the observed process at multiple scales and that it is worth placing greater confidence in the model.

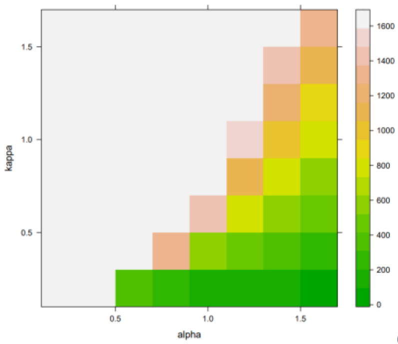

60.1.

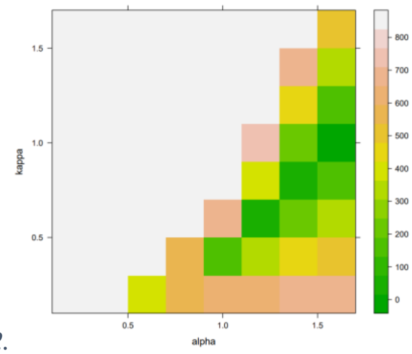

60.2.

Fig 60: 60.1, arrival times for wheel-made pottery for 64 simulations with different combinations of a and k for Crete. Figure 60.2, parameter space map where the value of each cell is the difference between the arrival times in Figure 60.1 and the mean from 1000 Monte Carlo simulations of arrival times based on empirical data from Crete (reflecting temporal uncertainty).

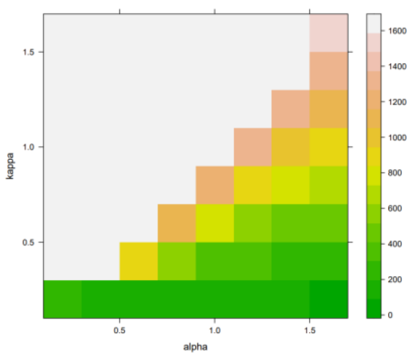

61.1.

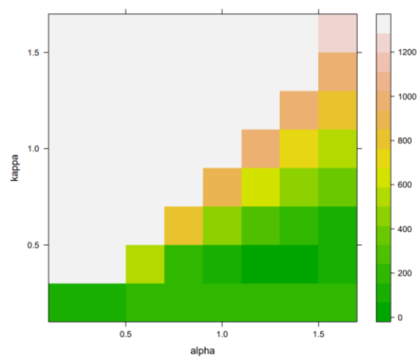

61.2.

Fig 61: Figure 61.1, arrival times for wheel-made pottery for 64 simulations with different combinations of a and k for Euboea. Figure 61.2, parameter space map where the value of each cell is the difference between the arrival times in Figure 61.1 and the mean from 1000 Monte Carlo simulations of arrival times based on empirical data from Euboea (reflecting temporal uncertainty).

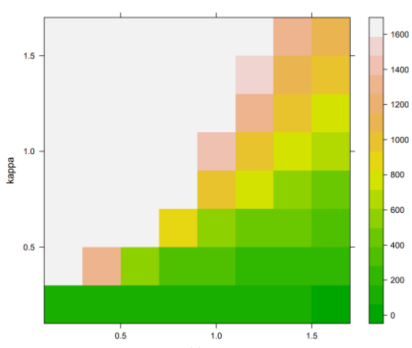

62.1.

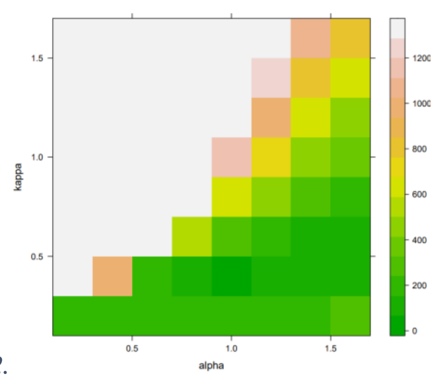

62.2.

Fig 62: Figure 62.1, arrival times for wheel-made pottery for 64 simulations with different combinations of a and k for Mainland Greece. Right panel: Figure 62.2, parameter space map where the value of each cell is the difference between the arrival times in Figure 62.1 and the mean from 1000 Monte Carlo simulations of arrival times based on empirical data from Mainland Greece (reflecting temporal uncertainty).

8.3 Simulations with trait loss

Having observed the behaviour of the model in a fairly realistic environment and compared it to the pattern that emerges from the empirical archaeological data, it is worth exploring a model with potentially even more realistic assumptions about whether communities retain the trait after they have received it. That is, that it is more realistic to assume that each community would lose some of the trait (some of the potters in the community would switch to another way of making pottery or that the potters with the trait would leave the community/pass away). I explored such a model in a toy environment in chapter 7, but in this section, I perform simulations on the realistic environment that was defined in section 8.1.1. To compare the outputs of the model with loss with the model with no loss, I repeat the experiments in the previous section using the same parameter values where possible. Adding trait loss introduces further stochasticity in the simulations – note that previously the only stochastic element was the location of each community (using landscape zones made the selection of location probabilistic rather than stochastic) as well as the local adoption rate for the model in section 8.1.5. Here the calculation of the magnitude of trait loss for each community at every timestep is random - at every timestep I subtract from each community's trait proportion value one random value drawn from the range: 0 (no loss) to the community's current trait proportion value (total loss) - and I explore the effects of this stochastic element by repeating certain experiments multiple times to assess the effect that stochasticity has on the outputs of each experiment.

Section 8.3.1 explores the model in an environment where all communities have an equal probability of being placed somewhere on the landscape (except the sea). Section 8.3.2 explores the outputs of the trait loss model in the two landscapes that were defined in section 8.1.3. Finally, section 8.3.3. explores the parameter space more formally.

The literature on innovation diffusion in the Aegean suggests Crete was the origin of many innovations or at least functioned as a location from which pre-existing traits spread to multiple regions with greater speed (like the potter's wheel). Previous sections explored the role of Crete as the origin of the innovation and suggested that transmission processes starting from Crete produce a qualitatively different signature to transmission processes starting from other locations in the Aegean. However, another element that could make Crete important is whether it protects the trait against random loss at the system level (and in so doing functions like the self-sustaining clusters in the toy model experiments). I formally explore this in section 8.3.2, which includes simulations for the trait loss model using various origins, including Crete.

8.3.1. Experiment 1

Repeating the experiment in section 8.1.1 with the same values for all parameters results in no observable spatial transmission process. This suggests that adding random trait loss overall decreases the rate at which the transmission process occurs or even stops it completely, which is not surprising. However, it does present a complication in that it means it is more difficult to compare the outputs of the model with no trait loss compared to the model with trait loss.

In order to better observe the outputs of the trait loss model I used $k = 0.2$ rather than $k = 0.6$ and kept the same values for the other parameters (N = 2000, $a = 1.2$). Figure 63 shows the outputs of the simulation at Timesteps 1, 250, 350, 450, 550, and 750. The trend of the mean over time is still a double-S shaped curve, and punctuation in the trend is due to transmission of the trait beyond the landmass of Anatolia, as observed in the model with no loss. This time, however, the trait never reaches the 100 % equilibrium – it oscillates around an equilibrium at 22 %. It also never reaches certain regions in the study area: Crete and most islands never receive the trait in the 1588 timestep observation window. Spatial transmission depends on clusters in the centre and coasts of Western Anatolia. From there the trait becomes established in clusters in Thrace, Macedonia, and Thessaly between timesteps 250 and 350. Then the trait spreads to clusters in Delphi, Kea, and the Argolid in the Peloponnese between timesteps 350 and 450. The final large clusters to receive and retain it are located in central Euboea (near Lefkandi) and the southernmost tip of the Peloponnese.

Figure 64 shows the mean trend emerging from repeating the same simulations 10 times and suggests that stochastic effects may introduce slight variations in the equilibrium value of the system and the timing and volatility of the punctuation. However, the shape of the curve is always a double-S shape, so there are no major qualitative differences in the simulation outputs.

Figure 65, which involves running the model with the same values for *a* and *k* but varying the value for the number of communities in the environment (N), suggests that the shape of the curve can change as a result of increasing the number of communities above a certain value — note how the double-s curve becomes more step-like when N is equal to 2500. Increasing the number of communities also increases the value at which the system reaches an equilibrium.

Figure 63 shows that there is significant variation around the mean trend over time, which persists even when the system has reached the equilibrium – this can also be seen qualitatively by the fact that there are communities with very low trait proportion values on the map (shown in purple) while there are others with very high values (shown in yellow) and the latter retain the trait over long periods of time even when they occasionally lose some proportion of it. This can be seen in a histogram of the trait proportion values at the end of the simulation which reflects a bimodal distribution (Figure 66), where most communities have very low values (or even 0), and a smaller number of communities have extremely high values (and even complete adoption in some cases). Strikingly, there is a lack of communities with values around the mean of the system. The distribution of values suggests that the system mean might not be the best measure of central tendency.

What is it about the communities with very high trait proportion values that makes them retain the trait at high levels compared to other communities? Given that this simulation uses the same value of a for all communities, this suggests that it is something about their position in space that makes them unique. Figure 67 suggests that there is a correlation between each community/node's strength (the sum of edge weights connecting each node to all other nodes in the network). This means that what is key in terms of retaining the trait is having strong enough connections to other nodes, since this is a way for a community to get the trait back soon after it has lost it. Seen from the scope of the fragile and non-fragile systems of transmission that the literature suggests determined the success of innovation diffusions in the Bronze Age (Roux 2008), this suggests that it is not only the size the learning network that protects an innovation from loss, but its spatial structure also matters when the strength of transmission is determined by the distance separating communities of potters.

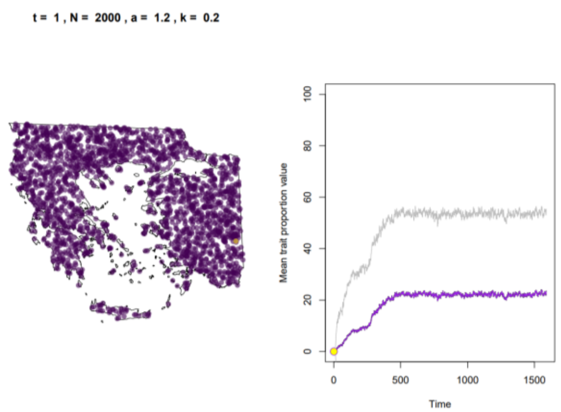

63.1.

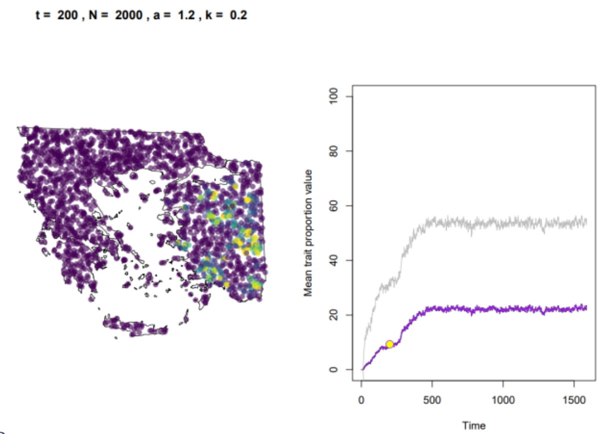

63.2.

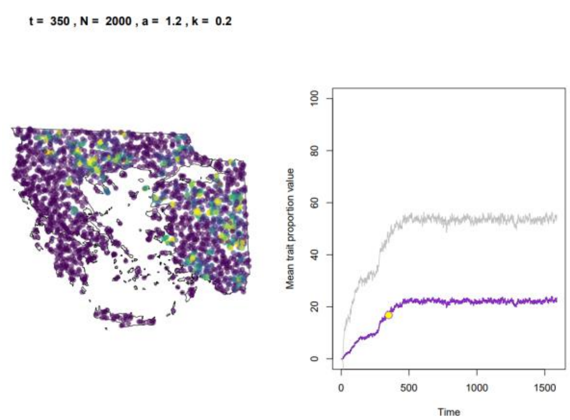

63.3.

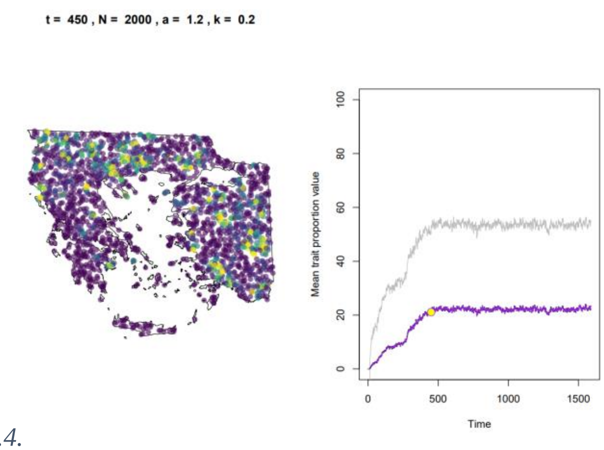

63.4.

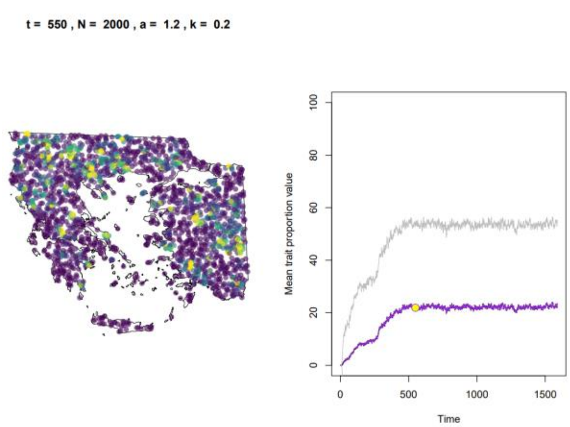

63.5.

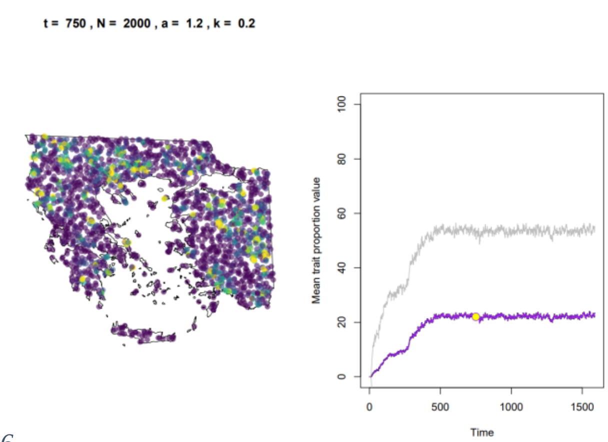

63.6.

Figure 63: The spatial transmission process emerging from a simulation with N = 2000, a = 1.2, k = 0.2 and the value at the origin node set at 0.1 at timesteps 1, 200, 350, 450, 550, and 750. Communities with higher trait proportion values are yellow, and communities with lower values are purple. The right-hand side of each panel shows the mean trait proportion value over time and an envelope of one standard deviation – note how there is a large variation in trait proportion values even after the system starts oscillating around the 22 % equilibrium value. Grey lines define an envelope one standard deviation above and below the mean. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

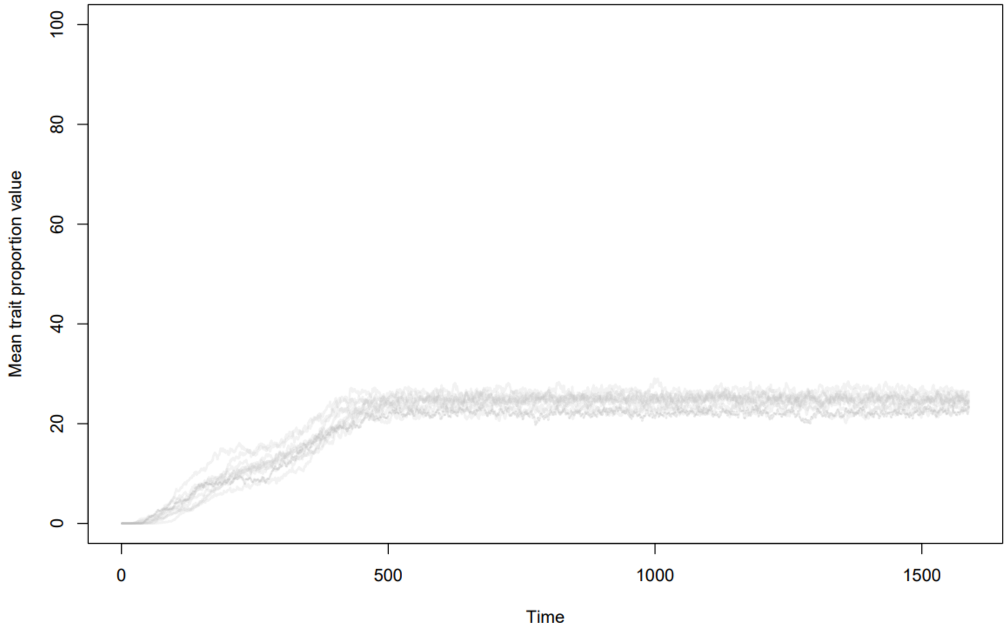

Figure 64: The mean trait proportion value trends over time for 10 simulations using the same instantiation of nodes and parameter values as the simulation in Figure 63. Variation is due to the stochasticity inherent in the simulations.

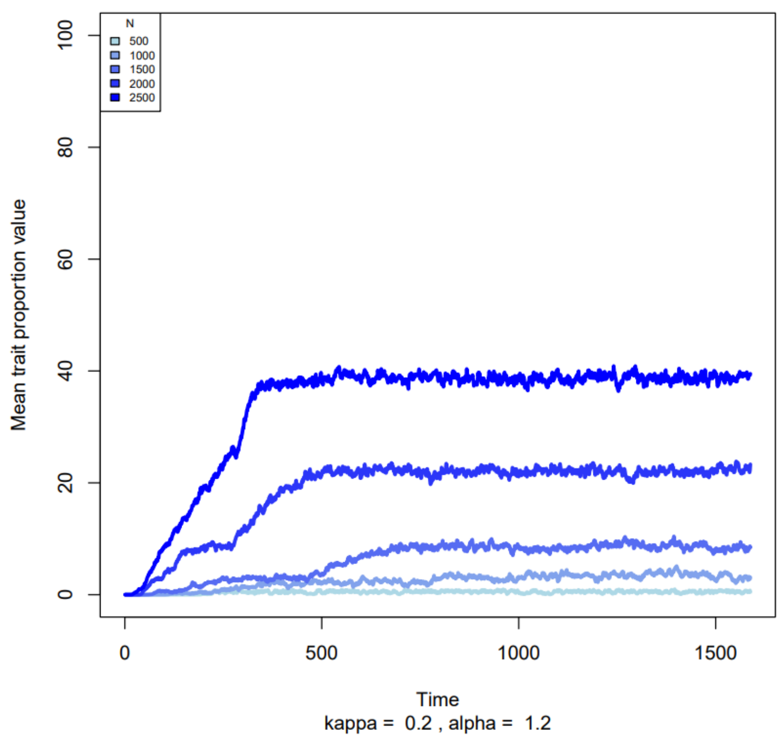

Figure 65: Mean trait proportion values over time for simulations run with varying values for the parameter N (the number of communities in the environment) and k = 0.2, a = 1.2.

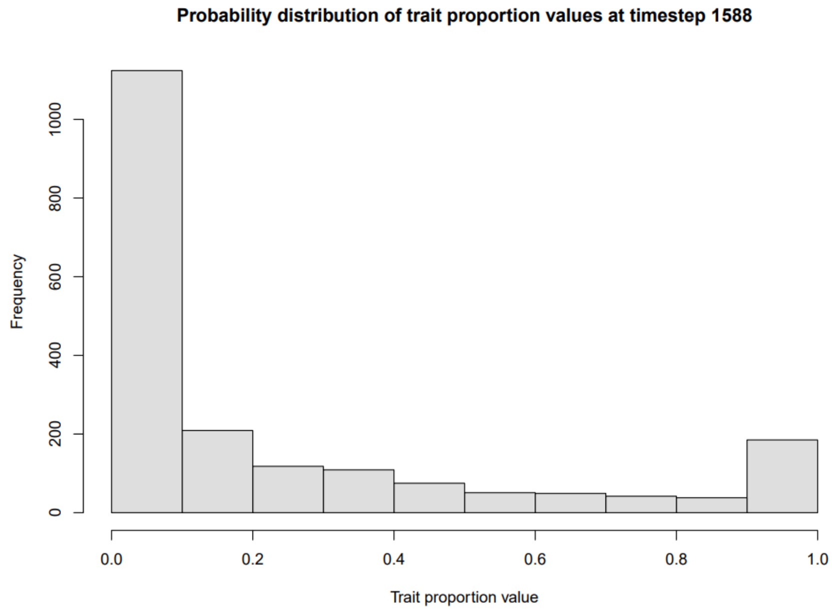

Figure 66: Histogram of the trait proportion values of all communities/nodes at the end of the simulation in Figure 23, when the system is at equilibrium. Note that the distribution is bimodal, and the main mode is around very low trait proportion values near 0 %, while the secondary mode is around extremely highly values (almost complete adoption at 100 %).

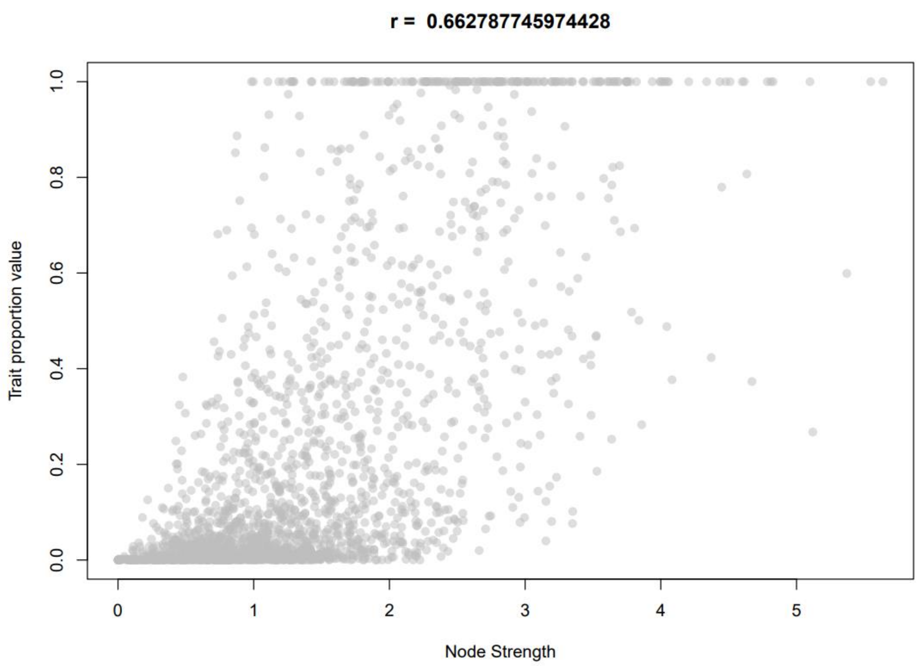

Figure 67: Scatterplot showing the relationship between each community/node's strength (the sum of edge weights connecting each community with all others in the network) and its trait proportion value at the end of the simulation in Figure 23. There is a moderate positive correlation between strength and trait proportion value (calculated using Pearson's test).

8.3.2. Experiment 2: Different origins

Experiment 2 involves running a simulation with N = 2000, a = 1.2, and k = 0.2 using different origin points (their location is shown in Figure 3). The shapes of the curves produced by each origin point in Figure 68 are roughly the same as in Figure 44, which shows the outputs of the model with no loss: origins in the landmasses of Anatolia and Greece result in double-S shaped curves (with the punctuation happening when the trait reaches the other landmass), while origin nodes in Crete, the central Aegean and central Macedonia lead to single-s shaped outputs.

The most striking aspect of the transmission process that emerged from the experiment is that the clusters that retain the trait when the system reaches the equilibrium are the same for all origin points (Figure 69). In Western Anatolia, such areas are the area south of the Urla peninsula to the south of Cesme, the area to the south of Troy, the area near Karatas. In the Greek mainland such clusters are located in central Macedonia, Corfu, Thessaly, central Euboea, near Delphi, Kea, the Argolid, the gulf of Messenia and central Peloponnese. Interestingly, even when Crete is the origin point, it loses the trait fairly early on in the process as soon as the trait reaches the Peloponnese – this is possibly due to the fact that the population of communities in Crete is not large or well-connected enough to form a self-sustaining resilient cluster that protects the trait against loss at the level of the cluster — it is possible that different parameter values for the rate of local adoption or smaller communication costs associated with lower values for k would result in the trait surviving in Crete. This example suggests that in this system it is not guaranteed that the origin of the trait will have it at the final stages of the process when the system reaches the equilibrium. The locations of the self-sustaining clusters are not unique in that they are located near the innovation's origin. They are unique in that they are structurally placed in the network in such a way that allows the communities that comprise them to regain the trait from their neighbours and to give it back to them when they lose it in return.

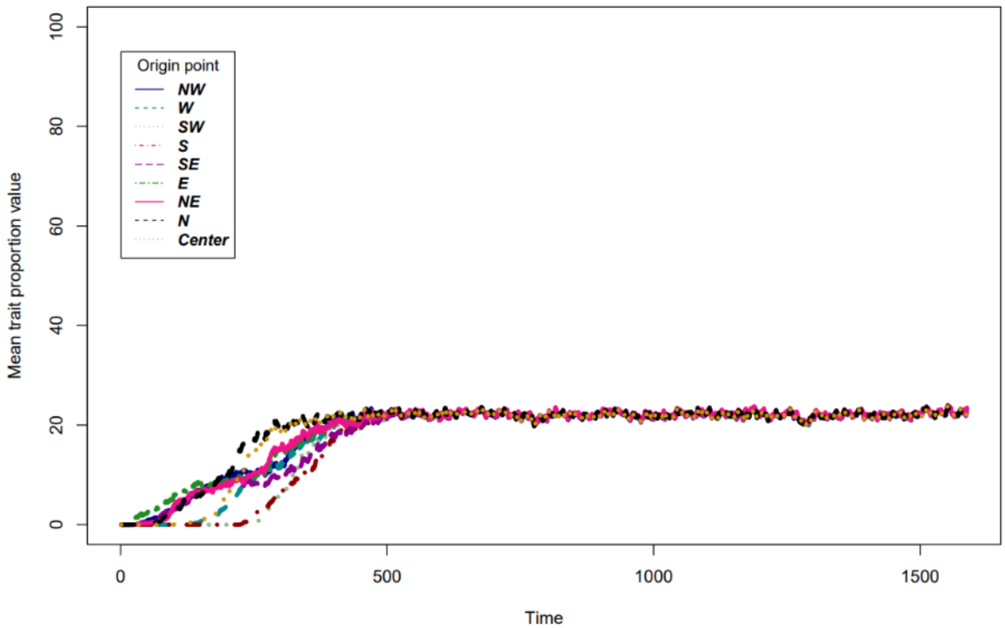

Figure 68: The mean trait proportion value trends over time for 10 simulations using the same instantiation of nodes and parameter values as the simulation in Figure? and different origin points (the location of each origin point is shown in Figure 43.

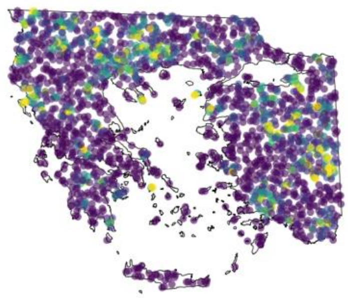

69.1.

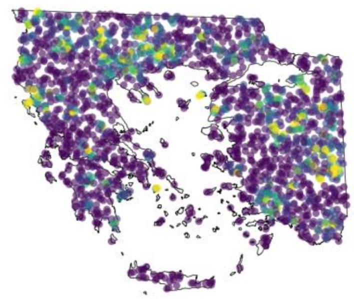

69.2.

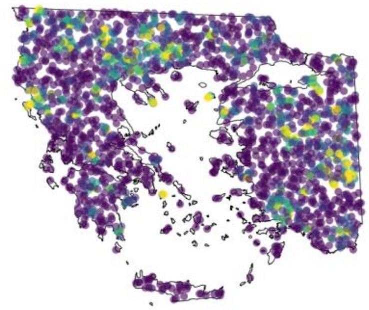

69.3.

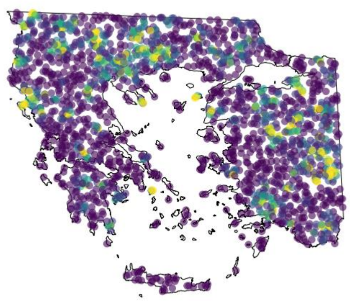

69.4.

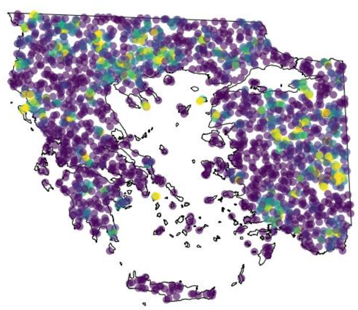

69.5.

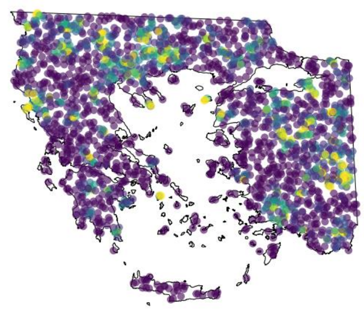

69.6.

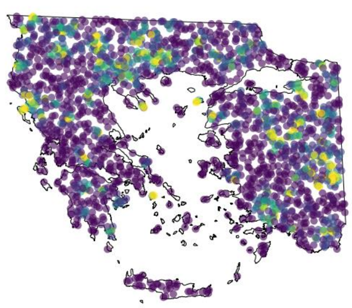

69.7.

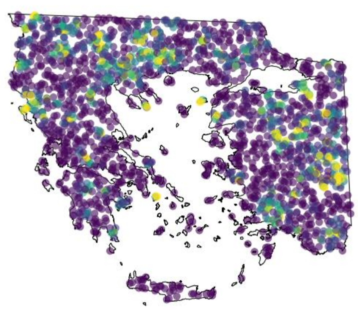

69.8.

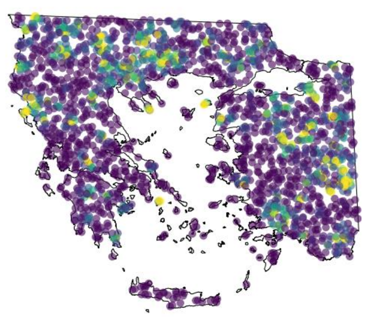

69.9.

Figure 69: The system shown at Timestep 1500, after the system has reached the equilibrium for simulations run with the same parameter values (N = 2000, k = 0.2, a = 1.2, origin = 0.1) and the same instantiation of communities/nodes using different origin points – 1 using the NW origin, 2 using the W, 3 using the SW, 4 using the S, 5 using the SE, 6 using the E, 7 using the NE, 8 using the N, and 9 using the Central origin. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

8.3.3. Experiment 3: Different landscape types

Simulations with the same parameter values (N = 2000, a =1.2, k = 0.2) were performed on the landscapes defined in section 8.1.3. As described previously, the landscapes are made up of zones that determine habitation probabilities. In landscape 1 (Figure 70.1), the probability of habitation is 0.45 in the coastal flat zone, 0.15 on the coastal steep zone, 0.30 in the inland flat zone, and 0.10 for the inland steep zone while in landscape 2 (Figure 31.1), the probability of habitation is 0.70 in the coastal flat zone, 0.05 on the coastal steep zone, 0.20 in the inland flat zone, and 0.05 for the inland steep zone.

The simulations running on landscape 1 (Figure 70) produce a similar transmission process to experiments 1 and 2. The main difference is that the locations of some community clusters that emerged as self-sustaining clusters retaining the trait over time are slightly different: for example, in the central Aegean it is the island of Melos rather than the island of Kea that forms a self-sustaining cluster. The most striking difference, however, is the fact that the trait now becomes prominent in Crete, which has two self-sustaining clusters that preserve the trait at high values over time and after the system has reached an equilibrium. On the other hand, a self-sustaining cluster never emerges in Euboea.

Figure 70.2 shows how the mean trend over time is affected by stochasticity for simulations with the same origin point (northern western Anatolia). It is evident that while there is some variability in the values over time, the shapes of the curves are qualitatively the same.

Looking at the shapes of the curves that are produced by varying the origin point (Figure 70.3), the picture that emerges is practically the same as in the previous examples, with the only difference being that the central Macedonian node now produces a double-s shaped curve rather than the regular logistic curve that it produced in previous examples.

The simulations running on landscape 2 (Figure 71), which reflects a more extreme tendency towards habitation preference on coastal zones, present a slightly different picture. The simulations with different origin points suggest that the clear distinction between double-s and single s-curves is less clear, and most origins generate a single S-curve. Taking the NW origin as an example, it is possible that the first phase of fast adoption (the first s in the double-s curve) for the model without trait loss is now slowed down by the effects of local trait loss. Additionally, for this origin point, the system reaches the equilibrium slightly slower when the process takes place in landscape 2 than in landscape 1, which represents a more extreme preference for coastal habitation and essentially halves the habitation probability for inland steep zones (0.05 rather than 0.10). Given that the NW origin is located in a very steep inland area, the density of communities there is lower than in other parts of the landscape, which means that transmission beyond that region should take more time compared to if inland steep zones were assigned a higher population probability. Based on the 10 simulations for the same origin point in Figure 71.2, stochasticity does not seem to produce major qualitative differences in the shape of the curves.

Figures 71.4 - 71.9 show that when coastal preference is so extreme, self-sustaining clusters with high trait proportion values are mostly located on the coasts, and Crete becomes even more prominent.

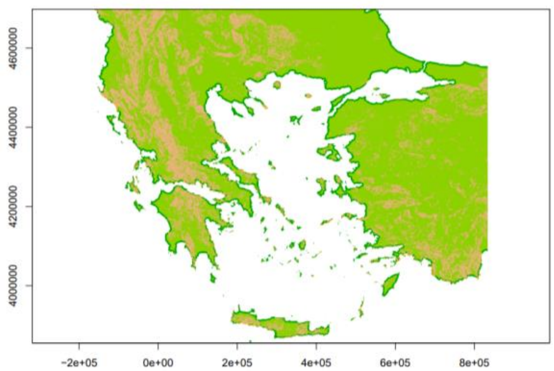

70.1.

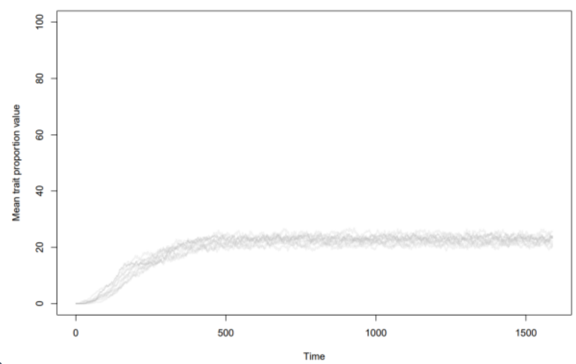

70.2.

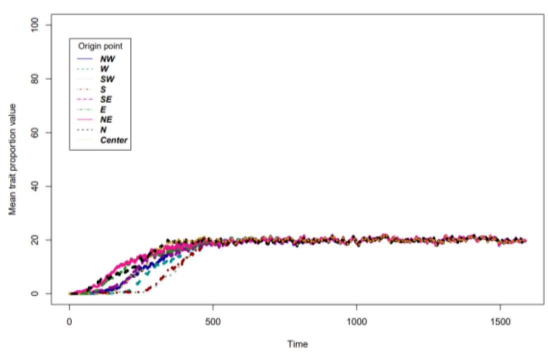

70.3.

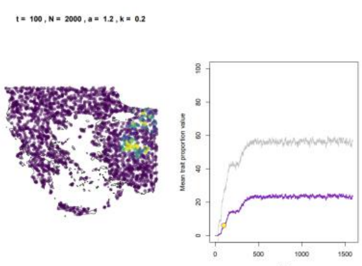

70.4

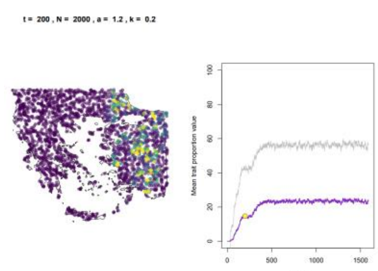

70.5

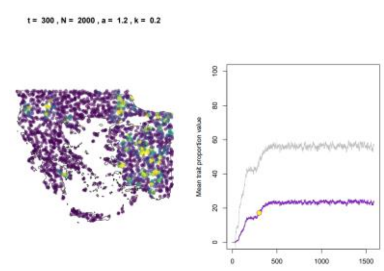

70.6.

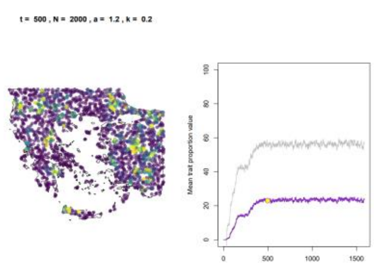

70.7.

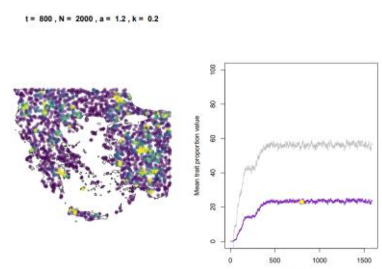

70.8.

Figure 70: Study area divided into an inland flat (light green), an inland steep (light orange), a coastal flat (dark green), and a coastal steep (orange) zone. Nodes have a 0.45 probability of falling on the coastal flat zone, a 0.15 probability of falling on the coastal steep zone, a 0.30 probability of falling on the inland flat zone, and a 0.10 probability of falling on the inland steep zone. Figure 70.2. shows the outputs of 10 simulations using the same parameter values and suggests that stochasticity does not cause significant changes in the shape of the curve. 70.3 shows the curves produced for different origin points. 70.4-70.8 shows the evolution of the system at different timesteps. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

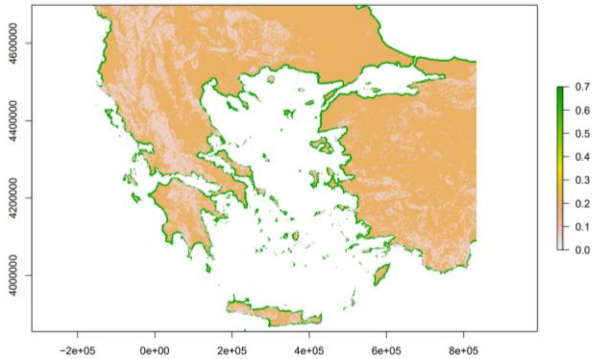

71.1.

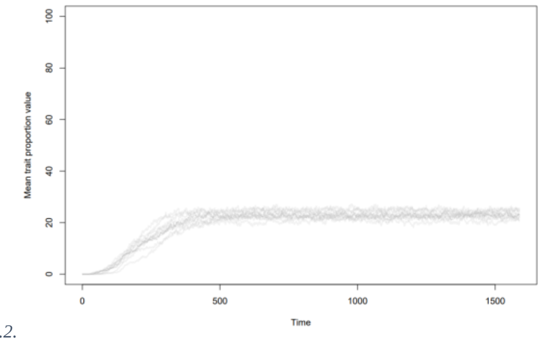

71.2.

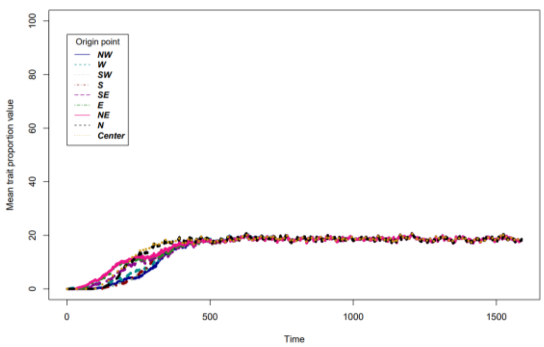

71.3.

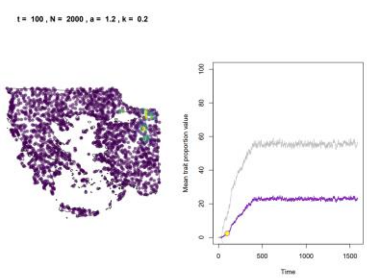

71.4.

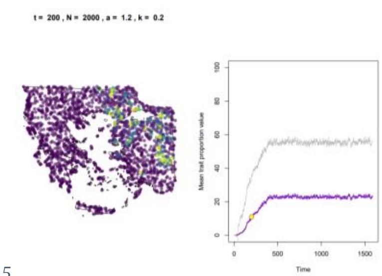

71.5.

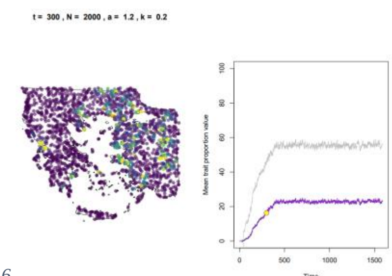

71.6.

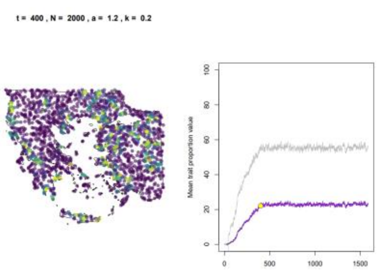

71.7.

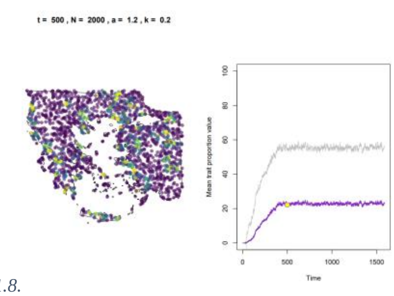

71.8.

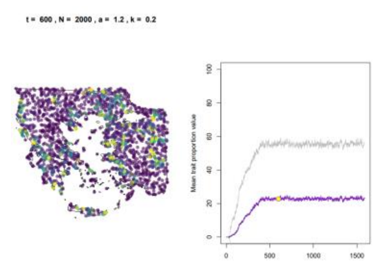

71.9.

Figure 71: Study area divided into an inland flat (light orange), an inland steep (orange), a coastal flat (dark green), and a coastal steep (orange) zone. Nodes have a 0.70 probability of falling on the coastal flat zone, a 0.05 probability of falling on the coastal steep zone, a 0.20 probability of falling on the inland flat zone, and a 0.05 probability of falling on the inland steep zone. Figure 71.2. shows the outputs of 10 simulations using the same parameter values and suggests that stochasticity does not cause significant changes in the shape of the curve. 71.3 shows the curves produced for different origin points. Figures 71.4-71.9 show the evolution of the system at different timesteps. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

8.3.4. Experiment 4: Parameter sweeps for *a* and *k*

As seen in section 8.3.1, comparing a simulation with $a = 1.2$ and $k = 0.6$ in an environment with N = 2000 communities was difficult because in the case of the model with trait loss, the spatial transmission process resulted in no observable outcome. In order to better identify the differences between the model with loss compared to the model without loss that was explored in section 8.1 and 8.2 in terms of their sensitivity to parameter values, I have performed parameter sweeps for the model with trait loss for the same range of values for *a* and *k* that was explored in section 8.2 for the model where trait loss does not occur.

Figure 72 shows that the parameter ranges used for the no loss simulations (Figure 72.1) are mostly uninformative for the model with trait loss (Figure 72.2). Successful transmission occurs only for the 4 highest values for *a* and the lowest value for $k = 0.2$. This suggests that, overall, including trait loss in the model affects the transmission process significantly, slowing it down or stopping it entirely, and that the model is particularly sensitive to parameter *k* which reflects communication costs. This is intuitively easy to grasp - when it is possible for communities to lose the trait, easy access to information from neighbouring communities becomes even more important for the 'successful' transmission of the trait because it allows communities to regain the trait quickly after losing it. In the presence of trait loss, only high values for the rate of local adoption *a* ensure trait transmission. Seen in the context of elite-imposed limits on innovation use, this could suggest that traits that are more prone to being lost due to stochastic reasons (for example, songs) would be more likely to disappear or not be transmitted beyond the origin community than traits that are copied in a more stable manner (e.g., forming techniques or decorative motifs on pottery).

More informative ranges for *k* are explored in Figure 73.1 and Figure 73.2. The former explores a wider range of values, and the latter focuses on a region of the parameter space that focuses on small variations in *k*. Figure 73.1 shows that the multiplicative relationship between *k* and *a* which was discussed previously, where diffusion is possible for both high *k* and high *a* simulations and low *k* and low *a* simulations is still prevalent in this region of the parameter space — visually, this is indicated by the diagonal pattern of intermediate values for the scaling parameter (cells with yellow-orange colours) on the top of the parameter space map. It appears as though even in the simulations with random trait loss, low values for the rate of local adoption e.g., $a = 0.6$, can be counterbalanced by low costs for crossing large distances, e.g., $k = 0.16$, having similar effects to the opposite case, where the rate of local adoption is high ($a = 1$) and the costs for crossing large distances is relatively higher ($k = 0.18$).

Figure 73.2 shows that looking at a very narrow range for *k* can also be informative in that it is possible to discern a non-linear pattern in the effects of *k* on the shape of the adoption curve. For example, looking at $a = 0.8$ suggests that increasing *k* linearly does not always result in increasing the scaling parameter: increasing *k* from to 0.164 to 0.165 slightly decreases its value, even though on average increasing *k* tends to increase the scaling parameter value.

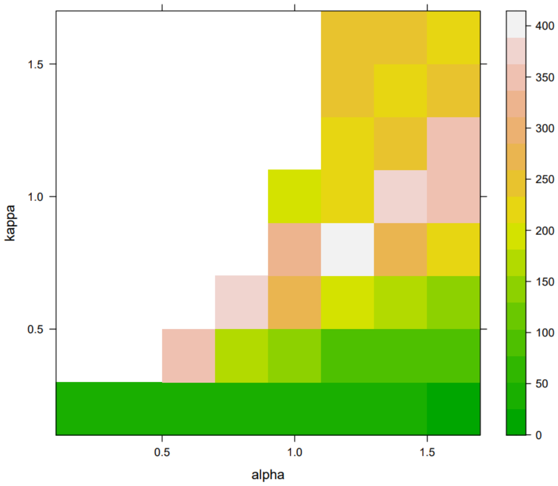

72.1.

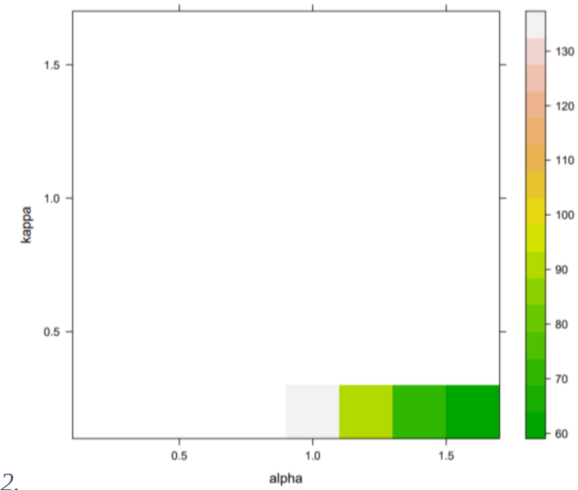

72.2.

Figure 72: Scaling parameter values for simulations of the model without trait loss (72.1) and simulations for the model with trait loss (72.2) using the same parameter ranges for a and k, N = 2000, and the value at the origin point = 0.1 (located in Western Anatolia).

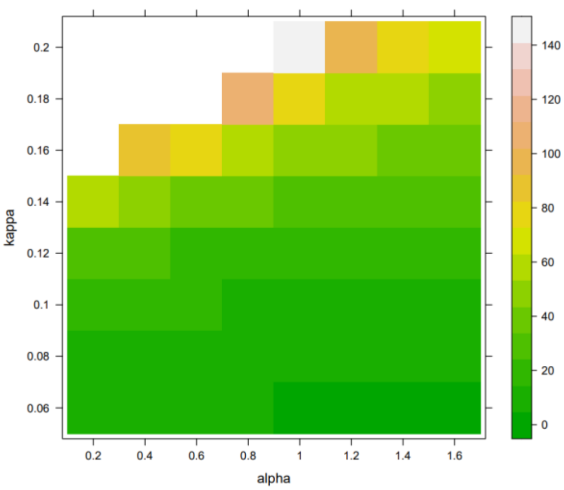

73.1.

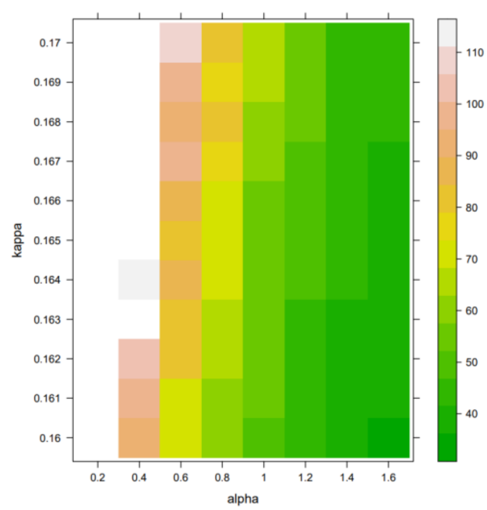

73.2.

Figure 73: Scaling parameter for simulations with different parameter ranges for the parameters a and k, N = 2000, and the value at the origin point = 0.1 (located in Western Anatolia).

Comparing the simulations with trait loss simulations to the empirical data (in the form a Monte Carlo simulation envelope) is challenging given that it has been impossible to identify a region in the parameter space where the system reaches a high enough equilibrium value and has a scaling parameter that is not essentially a step function rather than a logistic one.

Another issue is that there is a lot of variation around the mean trend even after it reaches the equilibrium value (which is not the case for the model without trait loss). This is reflected by the high value for the standard deviation in all simulation outputs. This in itself is not necessarily a problem and possibly parallels reality in that empirical data collection could be highly biased towards locations with resilience to trait loss when the cultural transmission process in the past was prone to stochastic loss. If the assumptions of the model hold in the real world, given that such locations are less prone to losing a trait as a result of their belonging in a 'self-sustaining' cluster, recovery bias would favour such areas as locations for data collection; note how the locations that consistently emerge in the simulations are near well-known sites (often near the Mycenaean palatial centres) with what appears to be continuous habitation and a rich archaeological record. Given that such locations emerge in the simulations as innovation 'nurseries' it makes sense that if the empirical mean was based only on samples drawn from those sites, the sample mean would tend to be higher than the actual population mean. The mean trend produced for the simulations is possibly not a good measure of central tendency when the model includes trait loss.

8.3.5. Experiment 5: Parameterising maximum trait loss

The simulations with trait loss explored in this chapter were designed with the assumption that, given that there is no additional information on how prone to loss the trait is, it is equally likely that a community will lose none, some, or all of its trait proportion value. In reality, however, it is safe to say that some traits are more prone to being lost than others, and that the potter's wheel in particular should be relatively less prone to being lost because of its mode of transmission — long apprenticeships and vertical transmission from the parental generation — which, as discussed in Chapter 2, is typically seen as a reliable transmission method.

The final experiment in this study involves varying the maximum amount of trait loss allowed at each timestep for each community. This means that at each timestep, rather than drawing a random loss amount from the range 0 to the community/node's current trait proportion value, I instead draw a random loss amount from the range 0 to a varying proportion of the community/node's current trait proportion value. This proportion is captured by a parameter μ, which represents the maximum trait loss proportion allowed at each timestep. For instance, a μ value of 0.1 sets the maximum trait loss allowed at each timestep to be 10 % of a node's current value.

Figure 74 explores the effect of varying this parameter on system-level outcomes while keeping the values for the parameters a and k constant at the values that achieved the best fit to the empirical trend (a =1, k =0.4). From a simple visual inspection, it appears that the value μ = 0.1 produces a better fit to the empirical trend than the case with no trait loss (μ = 0) and the cases where trait loss is allowed to take higher values. This is also confirmed by quantitatively comparing the simulation outputs to the empirical data using the method described in section 8.2.2. (Figure 75). This suggests that the pattern of spatial transmission of the potter's wheel over time that is empirically observed on the basis of the proportions of wheel-made over handmade pottery was due to a relatively low cost for communicating with other communities over long distances (either through long-distance communication or potter relocation since the model does not differentiate between the two), a lack of community level factors affecting the strength of transmission (such as local elites impeding or encouraging its adoption) and the fact that the technology was not very prone to being lost due to chance events.

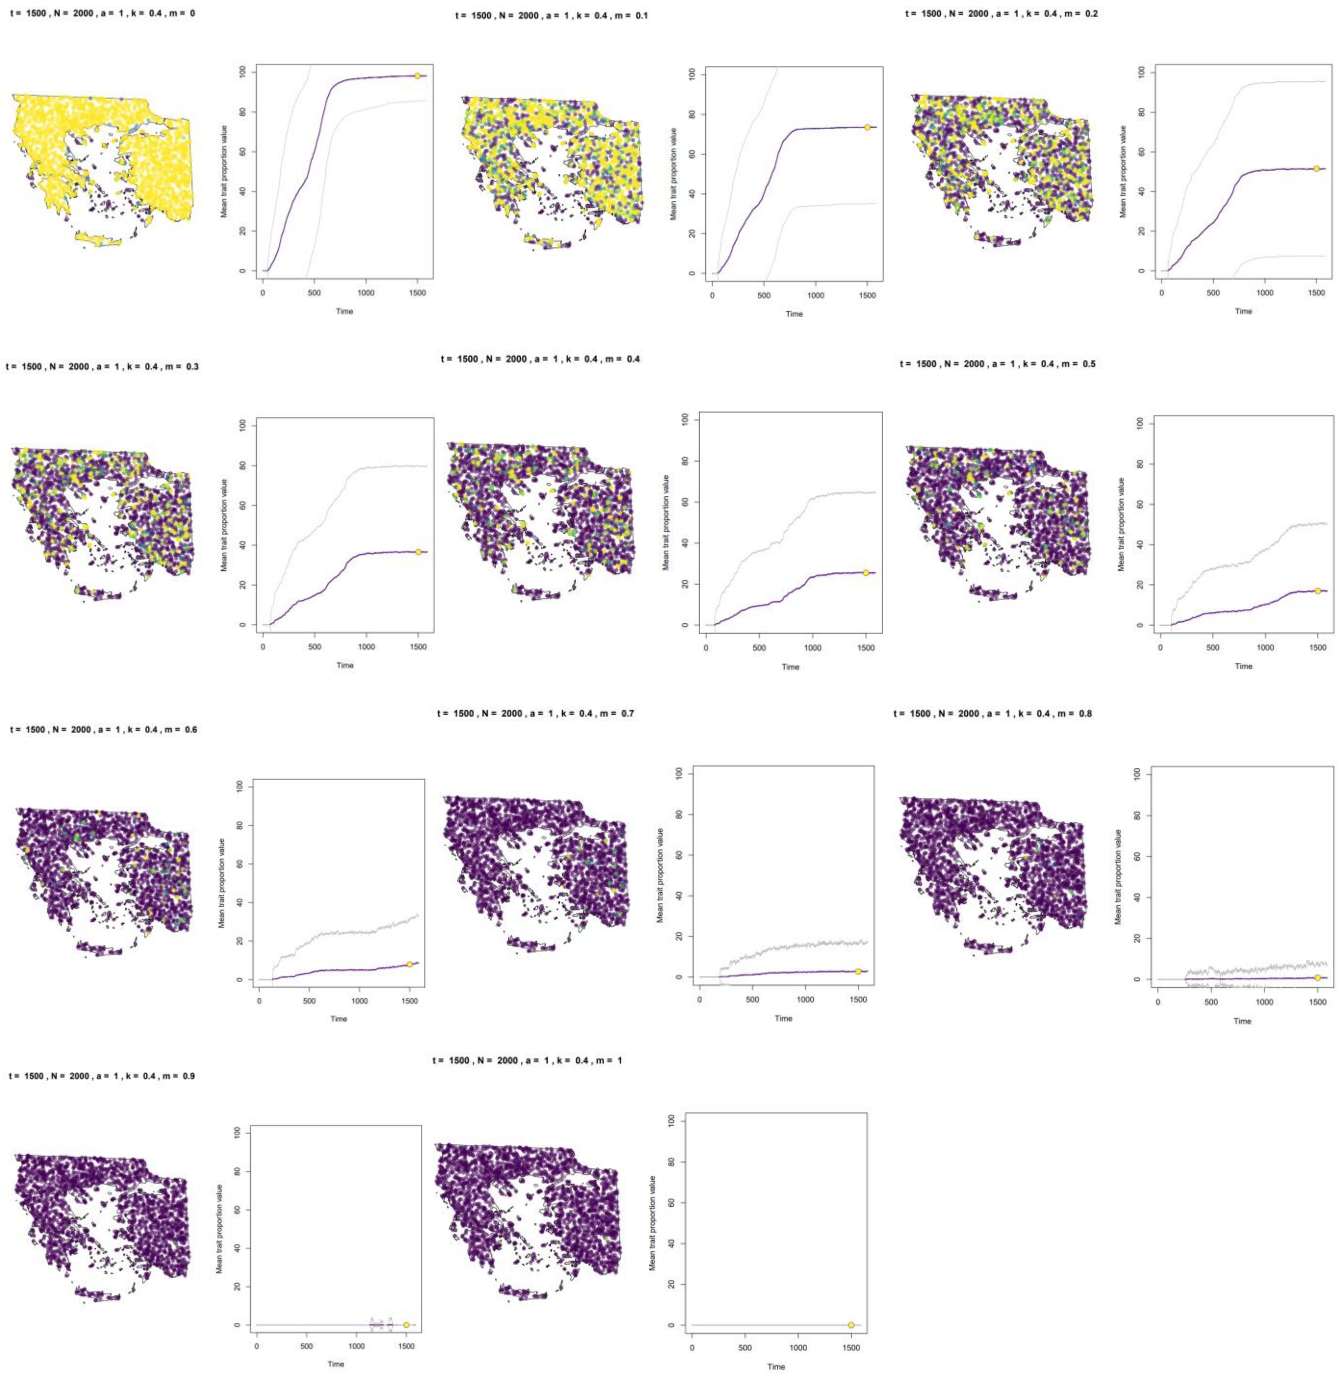

Figure 74: Simulations at timestep 1500 for the same instantiation of communities but with different values for the parameter for the maximum trait loss amount allowed at each timestep. Increasing the value of this parameter appears to lower the plateau that the system reaches at equilibrium as well as lowering the variability of trait proportion values at the level of the system – note the narrower standard deviation envelopes defined by grey lines for the simulation with $\mu = 0$ compared to $\mu = 0.1$. For the values of a and k that were explored here, values higher than $\mu = 0.7$ lead to no spatial transmission. Node colour indicates the trait proportion value of each node on a scale from purple (low) to yellow (high).

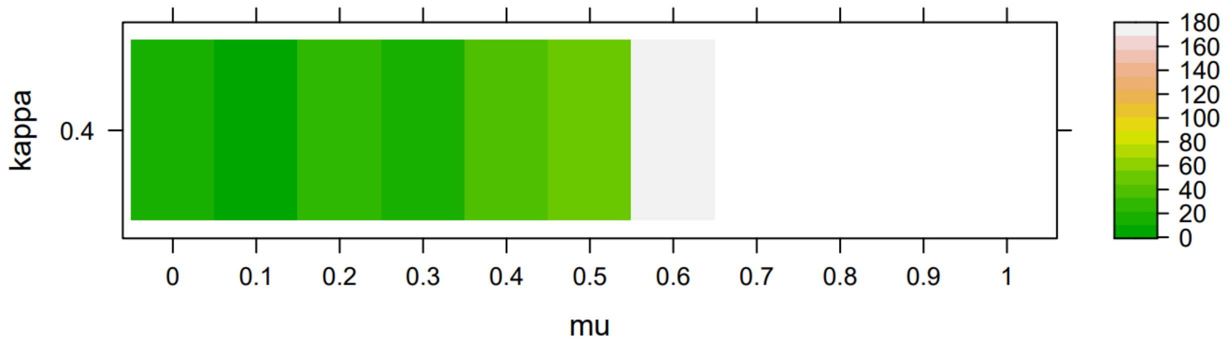

Figure 75: Sweep for the parameter μ which represents the maximum trait loss percentage allowed at each timestep for each community node for k = 0.4 and a = 1, which are the values that were associated with the simulation with the best fit to the empirical data for the model without trait loss (section 8.2.2.) – in this set of simulations the same effect (no trait loss) is captured by $\mu = 0$. The colour of each cell represents the mean difference between the value of the scaling parameter of a logistic function fit to the simulation and the scaling parameters fit to 1000 Monte Carlo simulations for the empirical data (the approach is described in detail in section 8.2.2.. The best fit to the empirical data is produced by the simulation with $\mu = 0.1$, which produces an even better fit compared to the simulations without trait loss.

8.4. A model with population fluctuations and no trait loss

As discussed in chapter 2, research in the field of cultural evolution has established through empirical observations and simulation, that larger population sizes are associated with higher probabilities of new traits randomly appearing through invention, and higher probabilities of them being transmitted through the population. There is a range of different scenarios that one could explore under this assumption; one option would be to keep the number of communities/population size N stable over time and run simulations for a range of different values for N. This was explored in section 8.1 (Figure 41), with outputs that showed that the speed at which a trait is transmitted over space significantly increases for a certain range of population values and then reaches a point at which increasing N does not significantly affect the shape of the diffusion curve.

Another option is not to assume that population size is at equilibrium, and to model the number of communities/population size as a function of time, with N increasing or decreasing linearly, exponentially or logistically — which means starting out exponential but levelling off when they reach a carrying capacity determined by environmental conditions or resource availability.

In this section I explore trait adoption trajectories for a population of communities whose size oscillates around an equilibrium using a non-monotonic function in order to capture the boom-and-bust patterns that empirical evidence — Weiberg et al. (2019) and my own analysis in section 5.4 — suggests occurred in the Bronze Age Aegean. As discussed in chapter 5, the overall picture that arises from an initial visual inspection of the time series trends associated with those proxies is one where population levels dropped at the end of the Early Bronze Age or the beginning of the Middle Bronze Age and then increased to their maximum levels for the Bronze Age during the Late Bronze Age, and dropped to a low again at the end of the Late Bronze Age. However, it is difficult to pinpoint the exact timing of the appearance and spread of novel traits in pottery in relation to those periods of population level boom-and-busts, since using different proxies results in differences in the timing of periods of population level increase and fall. Additionally, and as acknowledged by Weiberg et al. (2019: 755): “the amplitudes of the SPD should be treated cautiously and probably in relative rather than absolute terms given the overall sample size”. In the context of this modelling exercise, a relatively flexible approach is required to model population fluctuations over time.

A common approach to modelling population fluctuations in ecology is through a family of trigonometric functions like the sine and the cosine. Sinusoidal patterns are particularly common in the context of predator-prey models, where resource availability affects the growth rate of prey, which in turn affects the growth rate of predators and vice-versa, and both populations oscillate around an equilibrium (Benincà et al. 2009). In this section I will use a sinusoid to represent the number of communities in the simulation as a function of time. This sine curve will represent the fluctuations of the number of communities over time as oscillating around an environmentally determined equilibrium — which could be due to climatic conditions or resource availability.

The size N of a population oscillating around an equilibrium is represented by the following equation.

$$N(t) = Asin\ (b(t - c)) + \ d$$

In this equation, A is the amplitude of the sine wave, and it determines the strength of each oscillation. When it is set to be equal to 0, the function is a straight line equal to *d* and can represent a stable population size. The parameter *d* is the vertical shift of the midline of the graph which can represent the equilibrium or carrying capacity of the population. When it is set to be equal to 0 the value of the function will oscillate around 0 and have negative values on the bust phases, which is non-applicable when modelling populations. The parameter *b* is the frequency of the sine wave, which determines how quickly the population size rises and falls again. The frequency of the sine wave is related to its period, which is equal to $2\pi/b$ and represents the length of the portion of the function that is repeated.

The parameter *c* is the horizontal shift of the sine wave (sometimes also called a phase shift), which can "move" the wave to left or to the right. The ability to specify a horizontal shift can be very useful in the context of this modelling exercise, because of our ignorance of the exact timing of the transmission process in relation to local population fluctuations. It allows specifying whether the transmission process will start before, during, or after a boom or bust phase.

This function is so flexible that it enables modelling almost any population trajectory imaginable. With a small enough period and certain values of the horizontal shift, the population trajectories can be similar to those described by an exponential or logistic growth or decay function.

Fully exploring the range for the parameters of the sine function would allow potentially modelling any population trajectory possible — however, running the simulations for every possible combination of parameters is a very time-consuming process which would require vast computational resources. Furthermore, it is unnecessary given that we already have some information from the empirical record on the changes in population levels that happened during the Bronze Age. As we saw in Chapter 5, multiple lines of evidence suggest that population levels dropped from the end of the Early Bronze Age to sometime in the Middle Bronze Age and increased to their highest levels during the Late Bronze Age, followed by another population bust at the end of the Late Bronze Age.

8.4.1. Implementation

The implementation of the model remains mostly the same as the one described in Chapter 6 and at each timestep the trait proportion value for each community/node is calculated using the formula described in Section 6.4. What changes is the configuration of communities in space and their number over time. At each timestep, the number of communities is calculated using the equation described above. If the number of communities increases in the next timestep, I add a number of communities on random locations, and if the number of communities decreases, I choose a number of communities at random and remove them from the set of communities.

New communities are initialized with a trait proportion value equal to the trait proportion value equal to that of their nearest neighbour. This modeling choice is due to the assumption that the model does not deal with the introduction of new communities who are completely naïve and belong to a completely different cultural group. Instead, they represent new communities that split away from existing communities in the landscape — for example, a group of potters and consumers who were born in a neighbouring community and relocate and reproduce at a different location in the vicinity of their old community.

Due to computational limitations (each simulation becomes much more expensive to run when population fluctuations are introduced because community networks have to be processed at each timestep) I run the simulations for 160 timesteps instead of 1588 timesteps and the length of each timestep represents the passage of 10 years — so that the empirical window of analysis which lasts 1588 years is represented by having 160 10-year timesteps.

In the following experiment I use parameter value combinations that result in two population booms (based on the empirically observed population trajectories) by setting the period of the sine wave to be equal to 100 timesteps (setting the parameter *b* = 0.06283185). The parameter *d* which represents the midline of the wave and the equilibrium around which the population oscillates is set to *d* = 1000. The only parameters I vary in the following experiment are the horizontal shift which determines the number of communities at the beginning of each simulation and the amplitude of the sine wave. As discussed above, we know that population levels dropped at the end of the Early Bronze Age, but we do not know whether the introduction of the wheel was before or after this event. Varying the horizontal shift allows starting the transmission before and during the population boom and compare whether transmission trajectories are affected by this difference in initial conditions. Additionally, varying the amplitude allows comparing the effects of population fluctuations of different magnitude on transmission trajectories.

8.4.2. Results

Figure 76.1 shows the population level trajectories over time for a fluctuating population size where the first timestep coincides with the highest peak of a population boom — this allows exploring the effects of the transmission of the trait occurring during a population boom. I use *a* = 1, the value associated with the best fit to the empirical data for all previous models, and vary the parameter *k* (Figures 76.2-76.5). By varying the amplitude A to 0, 100, and 200 I explore cases where the population remains stable (A = 0) as well as vases where there are population fluctuations of different magnitude (A = 100 and A = 200). For all values of *k*, the simulations where there are population fluctuations reach the 1000 % equilibrium faster. The magnitude of the fluctuations does not seem to be as important since A = 100 and A = 200 lead to similar trajectories. However, differences in A become more important when the cost of crossing large distances *k* increases above 0.2 — this is reasonable given that the benefits of a higher population of communities in making transmission faster become more important when inter-community communication becomes more difficult.

None of the simulations where the number of communities fluctuates over time produce trajectories similar to the empirical data, at least on the basis of visual inspection. The only simulation that produces the distinctive double-S shaped curve that is empirically observed (see section 5.2.2.) is the simulation where the number of communities N remains stable over time at 1000 communities, and where $k = 0.1$.

A similar picture emerges from a set of simulations with the same parameter values but where initial conditions are different in that the transmission of the trait starts during a bust phase when population levels drop. This set of simulations results in trajectories where the amplitude of the population level sinusoidal function is even less important - note how close the simulations for A = 100 and A = 200 are for every value of k tested in Figure 77. As in the set of simulations in Figure 76, the best fit to the empirical data in the basis of visual inspection is produced by the simulation where the number of communities N is stable over time.

Overall, the simulations where the population is stable produce a better qualitative fit to the empirical data than the simulations with population fluctuations. This suggests that, either the fluctuations in population levels during the Bronze Age were not as extreme as the fluctuations in the simulations, or that the model with population level population fluctuations explored here is failing to capture something about how cultural transmission occurs between communities whose numbers fluctuate over time. It could be the case that in reality there was a spatial component to how new communities emerged/'died out' in the area that is not captured by the completely random spatial process that I used for this model — for instance, it is possible that new communities emerged in close proximity to existing ones with a large population size (members of the new generation exploring the vicinity of the old settlement in search of arable land). This shows that making model parameters a function of time in the simulations involves having additional information about the micro-level processes affecting them in space and time.

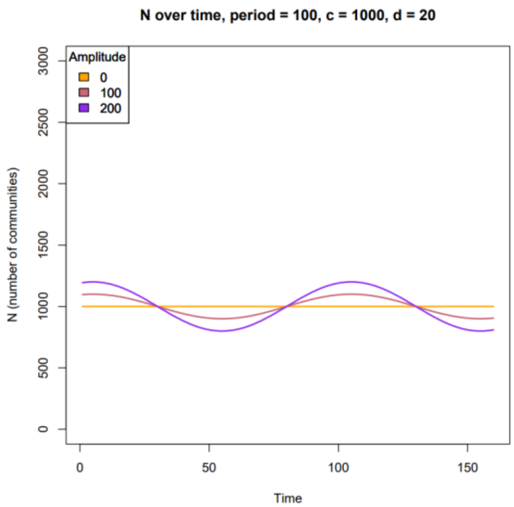

76.1.

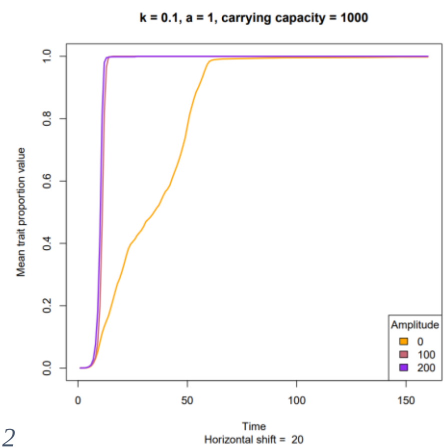

76.2

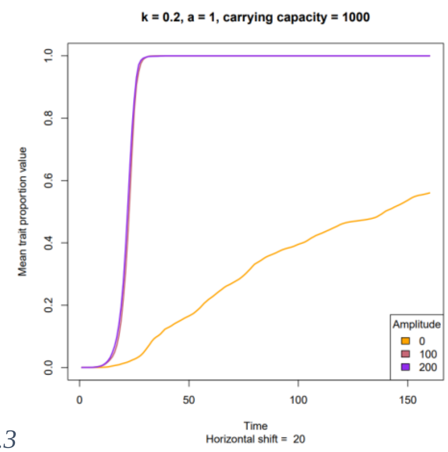

76.3

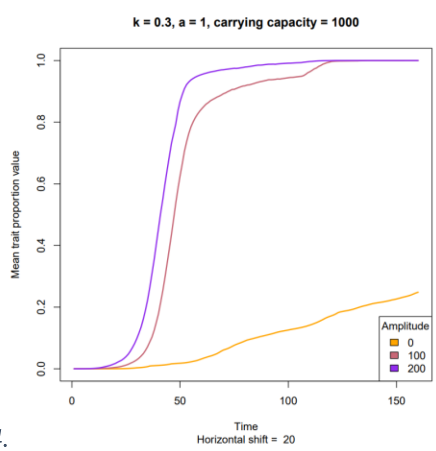

76.4.

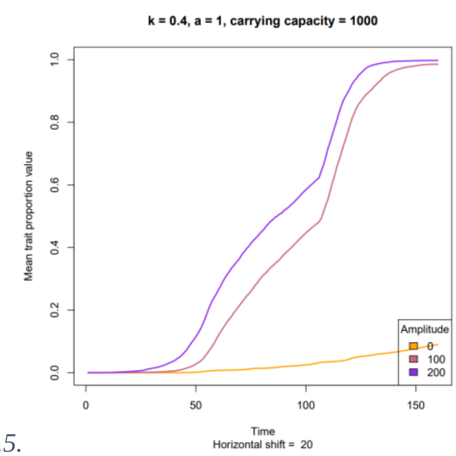

76.5.

Figure 76: Panel 76.1 shows the population level trajectories for three values of A (the amplitude of the sinusoid defining the number of communities over time) and when the horizontal shift d is equal to 20 (which makes the transmission process start during a population boom. Panels 76.2-76.5 show the outcomes of simulations with different values of k (0.1, 0.2, 0.3, and 0.4).

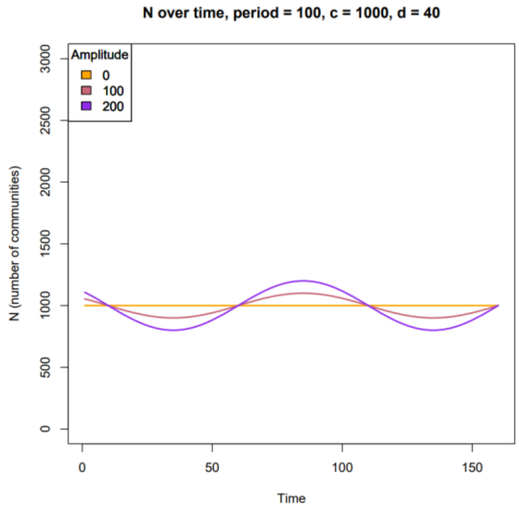

77.1.

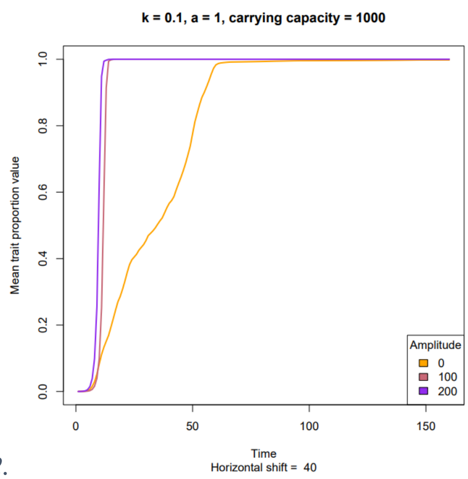

77.2.

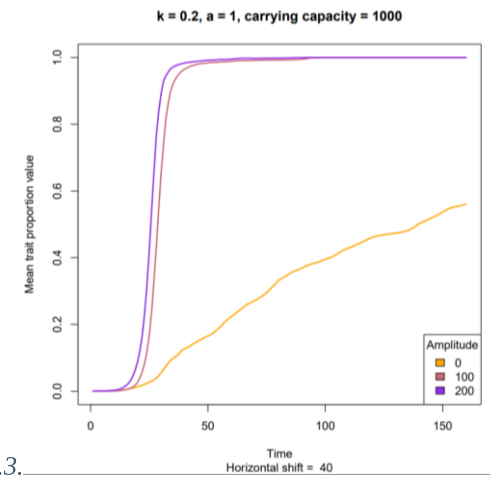

77.3.

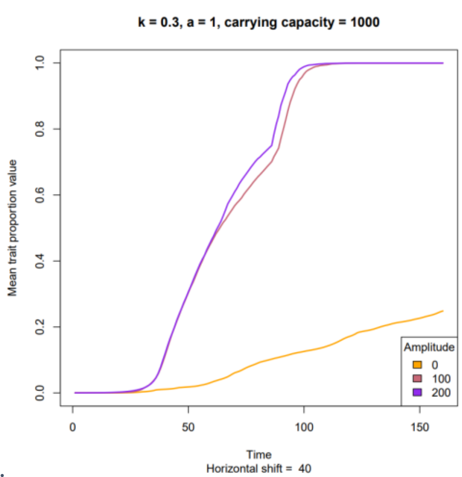

77.4.

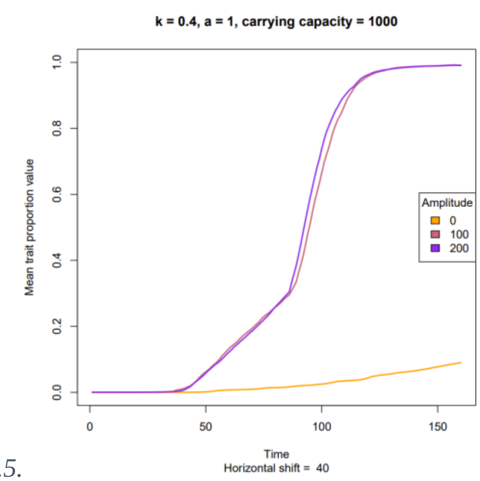

77.5.

Figure 77: Panel 77.1 shows the population level trajectories for three values of A (the amplitude of the sinusoid defining the number of communities over time) and when the horizontal shift d is equal to 40 (which makes the transmission process start during a population bust. Panels 77.2-77.5 show the outcomes of simulations with different values of k (0.1, 0.2, 0.3, and 0.4).

9. Insights gained and future directions

9.1. Insights into innovation transmission in the Aegean Bronze Age

Both the review of the literature on the potter's wheel, kilns, and IRT innovations and the simple and more 'realistic' simulations explored in this study suggest new ways of thinking about the transmission process of novel traits in prehistoric societies, where communication between communities was heavily affected by the existence of geographic boundaries and the capabilities of boating technologies. The Aegean has long attracted attention for its interesting geography and the nature of the communication networks that it enables.

The means of travel are, of course, particularly relevant in this respect. While the geography of the Aegean did not change much during the Holocene, the means of travel certainly did change, especially in the time period that is covered by the study. The results of the modeling sections of this study suggest that geography, varying degrees of effective long-distance travel, a relatively small degree of innovation loss due to chance factors and, to some extent, local response to innovation (e.g., a lack of elite-imposed limitations on the number of potters that could adopt the potter's wheel) can largely explain the observed system-level adoption patterns of the potter's wheel innovation (where the system is the entire study area). The modeling exercise was successful in explaining the most likely combination of those factors that explains the empirically observed adoption trend for the potter's wheel and the empirically observed timing of its arrival at different regions in the study area.

Of course, that is not to say that the model explains the adoption process at the level of the individual community or even the individual potter. The literature suggests that there was variation in the adoption of the same technology at the community level (see chapters 2 and 4) and explaining the reasons behind each potter's decision to adopt or reject it would require information that is, at best, difficult to acquire using only the methods employed in this study.

The cost of long-distance travel is particularly important. If pottery traits were transmitted through the movement of craftspeople or entire communities across the landscape, by extension this would mean that any change in sailing technology would also have affected the rate at which new traits spread between communities. In my model, the effects of improvements in sailing technology are explored by the decreasing the value of the parameter *k*, which represents the 'cost of travel' and penalises long-distance travel. Simulation outputs suggest that this parameter greatly affects the speed at which new traits are transmitted, the shape of the front in which they spread, and the relative importance of local connectivity — note how high communication costs make having a high local rate of adoption more important. This possibly means that the smaller costs in communication that sailing provided made the decisions of local elites have less of an impact on system-wide adoption overall.

In the case of the potter's wheel innovation, scholars have noted the slow adoption of the wheel in earlier phases before its association with Minoan and Mycenaean-driven production (Knappett 2016: 103-104, Hilditch, Jeffra, and Opgenhaffen 2021: 346) and, coincidentally before the advent of the sail, and the rapid increase in uptake after the wheel was incorporated in Minoan-driven production which coincides with the introduction of sailing in different parts of the Aegean (see Knappett 2016: 104, Gorogianni et al. 2016 for the islands of Kea and Aigina, and Jeffra 2013: 45 where Jeffra links the introduction of the potter's wheel in Crete to the "rise in long-distance exchange during the period of the first palaces"). While standardisation of production and craft specialisation linked to the palatial organisation of pottery production is a reasonable explanation for this — there seems to be a connection between wheel-made pottery and standardised production in the Aegean, as the introduction of the wheel-made pottery in sites like Phylakopi and Ayia Irini is associated with the 'Minoanising' conical cup, the fabric preparation and morphometric appear to be standardised (Berg 2004: 83) — the simulations suggest decreased travel costs as a result of Minoan sailing as an alternative factor that might have driven wheel adoption to the levels that it did after the first palace period. Of course, that is not to say that standardisation and specialisation should be excluded as potential explanations — they could even explain the improvement of sailing technology given that sailing boats are composite technologies made up of multiple 'innovative' traits and that making them possibly requires expert knowledge from multiple specialised craftspeople. It should be noted that the model is formulated on the basis of an equilibrium assumption in terms of the model parameters — sailing technology does not improve over time in the simulations, and communication costs remain the same for the duration of each run. Instead, I explore the effect of different travelling technologies as different runs with different parameter values. Future work in this direction could involve formulating a modified model with a non-equilibrium assumption that reflects technological improvement in travelling technologies by making the parameter k a function of time.

In the absence of empirical information, thinking about which came first – craft specialisation or improvement in sailing, becomes a chicken and egg situation. However, the simulations suggest that wheel technology can spread faster as a result of improved sailing technology, and that faster sailing can even counteract the system-wide effects of a community level factor such as elite-imposed local restrictions on the uptake of the wheel to some extent (but note that an elite-imposed low local rate of adoption, e.g., $a = 0.6$, would impose a threshold above which it is impossible for the trait to spread further at the community level). Elite-imposed restrictions on the community-level uptake of the wheel could take the form of the wheel only being used by specialists ‘attached’ to local elites using the technology to only make specific types of pottery that functioned as status symbols, as was probably the case in the Southern Levant and northern Mesopotamia according to Baldi and Roux (2016: 236)

The outputs of the model also suggest that geography matters when cultural traits are transmitted between pottery communities residing in permanent settlements, as they did during the Bronze Age. To be more specific, large scale geographic patterns matter more than micro-scale geography; what seems to matter more is where the origin of the trait is in relation to the East-West axis defined by Anatolia and mainland Greece (e.g., see Figure 44 in page 233). It is striking how small an effect the origin’s position in relation to the North-South axis has on system-wide adoption patterns. The tendency of local communities to cluster near coastlines and flat (potentially arable) land does not produce strong effects, either (see how similar the outputs of two simulations varying the landscape zone probabilities are in Figures 49-50, pages 241-242).

Including the potential for a trait to be lost as a result of stochastic sampling effects results in similar double-s shapes in simulated system-wide adoption trends but in these simulations the system mean oscillates around an equilibrium lower than the simulations with no trait loss. Such simulations also exhibit greater system-level variation in trait proportion values. In certain cases, it results in the trend becoming established in self-sustaining clusters of communities that are more resilient to innovation loss because their proximity allows them to get it back from their neighbours quickly. Of course, this effect is stronger when it becomes more expensive to travel over long distances (higher values for the parameter *k*). Interestingly, as was demonstrated by both the ‘toy’ and the ‘realistic’ simulations, in such cases the location of the trait’s origin does not matter. What matters is the structure of the network of communities and the cost of travel (edge weights) between them, as the network centrality measure known as node strength was shown to be positively correlated to higher community-level trait proportion values.

Certain areas in the Aegean emerge as regions in which such self-sustaining clusters appear in the simulations — and they are consistently located near existing large sites with continuous habitation — though it should be noted that in the simulations site location is random and the 'real' spatial distribution of potting communities (which remains unknown) might have given resulted in different regions becoming important under the same process. Notable examples of such regions include an area in central Euboea in the vicinity of Lefkandi, and in the Peloponnese, the gulf of the Argolid. The model generates the expectation that those locations should be less prone to losing a trait due to chance compared to other areas. In the case of the potter's wheel, the literature suggests that there is empirical evidence to suggest that Lefkandi was different to other areas regarding the intensity of wheel adoption. Spencer (2007) has attributed the high levels of adoption to aspects related to the organisation of pottery production which, based on the results of her petrographic analysis, was more standardised compared to other regions (Spencer 2007: 195). Knappett (2016: 102) acknowledged Spencer's assessment but also wondered: "Why does Lefkandi stand apart? This does then beg the question of why production should be organised differently in central Greece in the first place", moving on to suggest that "the Euboean case is quite hard to explain". The outputs of the model suggest that this can be explained by the location of Lefkandi in space in relation to other communities at the system-level of the entire Aegean.

In the model, communities in central Euboea like Lefkandi that have high node strength values are those that receive a high information content from the rest of the network (since node strength is defined as the sum of the weights of all edges that connect a node to other nodes in the network). The communities/nodes with the highest strength values belong in local clusters, which means that they receive a large sum of edge weights from the other nodes in the cluster (since small distances are afforded high edge weights in the model). Such communities are frequently located on coastal locations on or near the landmasses of Greece and Anatolia, particularly on gulfs like the Argolid. This suggests that gulfs can ensure the resilience of traits in the face of random loss by imposing a specific spatial structure on the networks of communities that develop around them, which allows fast inter-community transmission at the scale of the cluster.

This finding can throw light on some of aspects of geography that have perplexed archaeologists working in the area. In the context of Late Bronze Age transmission in Mycenaean pottery, Choleva et al (2020) empirically observed certain particularities about the transmission of the potter's wheel in the Argolid (using data from Tiryns and Mycenae) as opposed to other palatial centers like Ayios Vasileios in Laconia. For example, they note that in the Argolid there were two different finishing techniques used to produce kylikes compared to Ayios Vasileios which only had one, at least until the postpalatial period (1450-1100 BC), when the second technique disappeared from the Argolid (Choleva et al 2020: 273). Considering the outputs of the models with trait loss, it is possible that the spatial structure of the Argolid is such that, all else being equal with regard to the cultural transmission process, it can sustain cultural traits for longer periods of time in the face of trait loss. More importantly, they suggest that the Argolid was the region where "a variant of wheel-coiling method 3 — if not method 4 —, which on present knowledge was not available to the workshops at Ayios Vasileios" (Choleva et al 2020: 274) was developed. The authors also suggest that this variation enabled potters in those workshops to produce higher quantities of pottery which found fertile ground in the context of the palatial export economy (Choleva et al 2020: 274). Finally, in their words "future research will show whether or not this remarkable geographical coincidence of specialized production for export and specific wheel-coiling method is an example of cause and effect" (Choleva et al. 2020: 274). The simulations explored here show that, given that the assumptions informing the model hold, there is indeed something 'special' about the Argolid, which is not mere coincidence. If a new trait emerged anywhere in the study region (and, again, given that the assumptions of the model hold in the real world) it would be more likely for it to survive in the Argolid rather than in Laconia. Therefore, the development of a 'beneficial' new variety of wheel-coiling — I use the word beneficial here to describe a trait that enhances its 'bearer's' faster replication compared to alternatives, i.e., the fact that the new wheel-coiling variant (trait) made it easier to produce greater quantities of pottery ('bearer') with the trait — is more likely to take place in the Argolid rather than in Laconia, because the geography of the Argolid would protect it against stochastic loss at the crucial first stages of its transmission.

Knappett's position (2016) in his account on the spread of the potter's wheel in the Eastern Mediterranean (mostly comparing adoption trajectories in the regions of Euboea, Cilicia, and Cyprus) is against the notion of an East to West diffusion of the potter's wheel because according to his view, the adoption in those regions was "piecemeal" and "fragile" and depended on community-level factors. Furthermore, he suggests that elite involvement led to the adoption process becoming more fragile and that this explains the differences between adoption in other regions compared to Crete where he suggests adoption was more successful less dependent on elites due to 'scaffolding' in the social learning of the technique (Knappett 2016: 107).

The simulation models developed in this study incorporate most of the main factors discussed by Knappett — spatial transmission is captured by the parameter *k*, while local factors such as elite control over pottery production (a <1) are captured by the parameter *a* (though it should be noted that the model remains agnostic as to exactly which factor leads to lower or higher values for the parameter a and it is impossible to distinguish between elite involvement or other factors like the presence of social boundaries). They suggest that, if one is interested in the macro scale process of adoption across the entire Aegean, it is more likely that the observed process was the result of a spatially driven cultural transmission process that was largely dependent on the cost of crossing large distances (parameter k) than the effect of local community level factors (parameter a). This does not rule out that there might have been variability in the adoption processes for different regions — note that the empirical trend for the potter's wheel innovation is produced by integrating quantitative data by a very small number of studies across the Aegean (as discussed in Chapter 5). However, simulation is a useful tool to ask questions about the relative effects of different factors in such technological adoption processes and it can suggest ways in which certain factors can interact in space over time which can be hard to capture with a verbal model. In the case of Knappett's verbal model which suggests elite involvement is to blame for the low adoption of the technique in certain parts of the Aegean, the simulations suggest that community-level factors such as elite-imposed restrictions on the relative adoption of a technology at the level of the community can be offset by efficient and cheap long-distance communication between communities. Crucially, this effect was also observed in simulated outputs for scenarios where it is possible for a community to lose some (or all) of the trait due to chance — but this requires making communication costs even lower (Figure 73.1 in Chapter 8).

In this respect, the fact that the simulations point to low costs for inter-community contact as the main factor behind the observed adoption patterns in the Aegean is in agreement with Jeffra's (2013) view on the transmission of RKE-based techniques in Crete. Jeffra argues that "the evidence from the potting communities corresponds best with open, cooperative effort, prioritising group cohesion (perhaps extending beyond the boundaries of individual archaeological sites) during the period of technological change" (Jeffra 2013: 45). It should be noted that the simulation model explored in this study does not incorporate a distinction between different RKE-based techniques (I use a simple division between wheel-made and handmade pottery without differentiating between different methods of combining the use of RKE with coiling techniques), while Jeffra's analysis does. However, Jeffra (2013: 43) confirms that over time most potters in Crete converged to using a single RKE method (method 3), so this distinction is possibly not very important if one adopts a more deep-time perspective into the transmission of the potter's wheel in the area (as is the case for this study).

The models explored here can also throw light on some of the spatial transmission aspects of Minoanization, a cultural affiliation process linked to the transmission of the potter's wheel during the Middle Bronze Age onwards. In the model, when the origin of transmission is Crete, the curve is single S-curve which does not follow the characteristic double-S seen in cases where the origin is in one of the mainlands (Western Anatolia, Mainland Greece). This means that any process originating from Crete would appear as a relatively sudden sharp increase in the proportion of the new trait if measurements were taken from the entire Aegean region, since the relatively long 'incubation period' in the beginning of the process is the episode where the trait becomes established in the island of Crete. While Minoanization was a process involving the transmission of multiple cultural traits and the simulations only explore a case where a single trait gets transmitted, this has implications for how we interpret the relatively sharp increase in minoanizing traits in other regions of the Aegean. The models show that it does not necessarily have to be linked to a sudden colonization episode – it could simply be the result of an isolation-by-distance process which first spreads in Crete from an origin inside the island and then spreads quickly to the rest of the Aegean because it is so well-established in Crete by the time it propagates across the sea to spread beyond Crete. However, it should be noted that in the simulations where Crete is the origin of the trait I have used low trait proportion values for the origin node (0.1) and I'm therefore not exploring a scenario where the trait is already fully adopted in the origin node in Crete before it spreads to the rest of the Aegean, which would probably result in a shorter "slow adoption" episode (which would last for as long as the trait reaches the other regions) but would still show up as a sharp increase in the global trend when the trait reached the other regions. Additionally, the model shows that some locations like Kea are less prone to losing a trait after they receive it and can achieve higher proportions of the trait simply by virtue of their location in local clusters. This reflects what is observed empirically, as Gorogianni et al. (2016: 216) argue that the adoption of wheel-made pottery was faster in Ayia Irini in Kea than in Phylakopi in Melos.

The simulated transmission process that resulted in the best fit to the empirical data agrees in some respects with the verbal model proposed by Choleva (2020: 96) for the transmission of the potter's wheel across the Aegean: " Specialist potters seem to have been dispersed across the Aegean and created sporadic networks of apprenticeship that ensured the reproduction of their craft, therefore establishing the potter's wheel in an enduring transmission process sustained throughout the late EBII. The insight into the micro-history of the wheel phenomenon at each settlement actually reveals diversity regarding the patterns of appropriating and integrating the tool into the technological milieus, depending on the local and regional socio-cultural and economic dynamics". When one incorporates random trait loss into the simulation models, similar outcomes to the ones described here by Choleva (2020) emerge. The transmission process does not stop at the global scale, but certain areas and micro-regions lose the trait either permanently or for shorter periods of time. Dynamics at different local and regional scales such as clustering patterns lead to very different outcomes at the micro scale, but the transmission process is sustained throughout the Bronze Age at the global scale.

9.2. More general insights into innovation transmission

Apart from the insights that were specific to the specific case study of interest — the transmission of ceramic innovations in the Aegean — the modelling part of this study that used an abstract ‘toy’ environment produced valuable insights regarding spatial patterns and mechanisms that apply to spatial transmission processes that depend on the strength of intercommunity communication and have a local community level factor (a local adoption rate) that affects the transmission process.

The ‘toy’ simulations suggest that different combinations of the cost of crossing large distances and rate of local adoption result in transmission waves that propagate in different ways and with different speeds. For instance, when costs are cheap and the rate of local adoption is high, transmission waves are radial and fast. On the other hand, when they become more expensive, the transmission waves propagate along paths where there is a high density of communities, and they are slower compared to radial waves. Finally, very expensive long-distance travel and very low rates of local adoption result in transmission waves that propagate through clusters — this is the slowest propagation mechanism. Another useful insight is that the spatial structure of the network of communities and the existence of clusters is more important for fast and successful system-wide transmission than the location of the origin community in the environment and with respect to the location of the clusters.

Additionally, the ‘toy’ simulations with trait loss produced insights that have implications on how we infer the origin of a cultural transmission process in space using empirical data. When trait loss is possible, it is likely that the origin community/node will lose the trait permanently after some time has passed, particularly if the origin community is not located in a cluster/region with a high density of communities, even if the trait successfully spreads to the rest of the system. In such cases at the end of the transmission process the communities with the highest value for the trait are those located in well-connected clusters, while the origin community/node has trait proportion values near 0 %. Therefore, when attempting to trace the earliest origin of a cultural trait, it is important to assess to what extent to which the trait was prone to being lost due to drift. In this scenario, a researcher is more likely to 'miss' the earliest origin location if they attempt to infer the earliest origin solely on the basis of which location has the highest trait proportion values/highest incidence of the trait. If data of finer temporal resolution are available and the origin is inferred on the basis of early dating as well as high incidence of the trait, this would make the inference more robust in the face of this effect — however, such inferences would still fail to identify the earliest origin if the trait characterizes material that is prone to taphonomic loss (e.g., organic material).

Finally, the trend lines of the mean for all simulations with random trait loss are characterised by high variability around the mean, as a large number of communities lose the trait and never regain it (their trait proportion values are around 0 %). In such cases the trait becomes well-established in central locations that are near sites where research intensity is higher than in other areas. Therefore, if the trend of the empirical mean over time is produced using a sample coming solely from those sites, it is likely that it fails to capture the variability in the true population mean and the effect of communities that had very low incidence of the trait. This is another example where archaeology can fail to identify and account for the effect of absences in the empirical record.

9.3. Are simulation models an appropriate method for addressing archaeological questions with Aegean data?

In retrospect, the process of integrating legacy data to design and validate the simulation models for this study produced certain methodological insights on the suitability of simulation modelling to address archaeological questions using data from the Aegean Bronze Age that could benefit other archaeologists interested in using similar methods. Such methods are rarely used in Aegean archaeology and, at least to my knowledge, have never been used to study innovation diffusion in Aegean Bronze Age pottery. When thinking about why this is the case, the quantity and quality of data emerge as the elephant in the room. Excavations in the region are time-consuming, excavated material is published with significant delays if at all, and access to excavated material is limited by bureaucratic red tape. Formulating questions and models that incorporate the temporal and spatial dimension (as was the case for this study) ideally requires data collection from sites with deep stratigraphies for multiple sites across the Aegean.

Published studies by pottery specialists who generate the primary datasets are often characterised by a degree of methodological individualism. A relevant example that emerged from my review of the literature is the varying ways of describing the IRT trait (Chapter 4), which some archaeologists describe as a set of traits (lustrous, wheel-made, fine wares), while others describe it in terms of the alternating oxidizing-reducing-oxidating atmosphere used to produce the bichrome effect in the slip.

At the same time, a lack of unifying theoretical frameworks means that, on the whole, different researchers have different research questions, which is not a bad thing in itself, but it means that the process of “secondary retrieval” and “integration” of such legacy data (Chapman & Wylie

2016: 97-98) becomes more challenging. The use of a common data collecting method, as is the case for the recent tradition of using Roux and Courty's (1998) set of diagnostic criteria significantly alleviates problems in data integration. Unfortunately, sometimes some of the fine-grained and potentially valuable information emerging from such studies must be sacrificed in order to render it compatible with older legacy data. In this study, this meant 'sacrificing' the quantitative data on the prevalence of different wheel-coiling methods and only using quantitative data that use a two-fold distinction between wheel-made and hand-made pottery, since this was the reporting standard used in older publications (which used 'wheel-marks' and 'horizontal striations' as diagnostic criteria to identify wheel-made pottery macroscopically). The reliance on regional pottery typologies is another complicating factor. Researchers working in different regions date material using regional typologies, making it difficult to align the temporal dimensions associated with different regions into a single comparative framework. To make matters more complicated, typological chronologies are sometimes defined by the presence of the very traits that one is interested in dating, for example the use of the potter's wheel defining the Lefkandi I-Kastri phase in Cycladic pottery.

Finally, using quantitative data published by multiple researchers that specialize in specific time periods to trace the evolution of a phenomenon like the adoption of an innovation over time sometimes results in quantitative trends with regions of higher and lower uncertainty caused by certain periods being less intensely studied using quantitative methods compared to others in the same temporal window of analysis (resulting in smaller sample sizes for some sections of the time series). In the case of the potter's wheel innovation, for example, there is a relatively larger sample size of studies with available quantitative data for the periods coinciding with the introduction of the new technology (the late Early Bronze Age) compared to later periods where there is a general consensus that most pottery was wheel-made (like the Late Bronze Age). This means that, perhaps counterintuitively, disciplinary certainty in qualitative assessments results in higher quantitative uncertainty for such time series data.

However, there are reasons to remain optimistic regarding the potential of simulation models for explaining cultural change in the Bronze Age Aegean. In fact, it can be argued that Aegean archaeology is fertile ground for producing explanations with simulation models. It is one of the oldest branches of archaeological inquiry and as such has generated a wealth of verbal models by generations of archaeologists interested in tackling big questions such about cultural change (e.g. Renfrew 2011), social complexity (e.g., van Andel and Runnels 1988), urbanization (e.g., Whitelaw 2001), trading systems (e.g., Sherratt 1999), population movement and migration patterns, and long-distance travel (Broodbank 1989). There are large numbers of sites with long habitation sequences which enable looking at patterns of change over time roughly at the level of the community.

A good way of dealing with the problematic aspects of the available primary data is through selecting higher scales of abstraction for the formulation of the research questions and for the explanations one seeks to gain. Initially, one could start by formulating abstract and simple models to generate predictions about hypothesized processes, or compare simulation outputs with the empirical record solely on a qualitative basis (as was done here in Chapter 8 for the simulations incorporating trait loss) or with the aim to generate expectations of the empirical record to drive future data-collecting efforts (e.g. in this study the prominence of the Argolid as a self-sustaining cluster protecting innovation on an Aegean system-level suggests that this area is likely to preserve cultural traits that are prone to being lost due to chance). If there are enough available quantitative data (as was the case of the potter's wheel data in this study), it is recommended that any attempts to quantitatively compare simulation outputs to the empirical integrated data be done with an awareness of the uncertainty inherent in the empirical data and that they communicate this uncertainty to audiences in the most transparent way possible. The introduction of Monte Carlo methods can also help alleviate problems associated with dating uncertainty in the empirical data.

Archaeological data from other regions have recently been used in conjunction with bayesian methods such as Approximate Bayesian Computation with the aim to infer the parameter values that can generate the phenomenon of interest under the specified generative simulation model (e.g., Crema et al. 2016). Such studies typically focus on macroevolutionary change, examining trait diversity over time rather than focusing on the transmission of a single trait, and investigate a single community/location which is more likely to be published by a single research team and therefore less prone to data issues associated with the secondary integration of incompatible primary sources. An ideal dataset to use such methods would come from a site with a long habitation sequence, where the pottery specialist and the modeler would have good knowledge of the excavation practices and sampling methods used by the original excavators and where the data collection and reporting methods used by the pottery specialist would be known by the modeler. In this context, relying on presence/absence data rather than quantifying the site-level proportions of each trait is possibly a less time-consuming approach.

The use of more fine-grained models that actively incorporate the micro-level behaviour of agents as a set of behaviour rules (as is frequently the case for agent-based models), was not attempted in this study. A more micro-level model that incorporates the effects of the movement of individuals in inter-community transmission should have more secure empirical foundations for the movement/relocation/migration patterns prevalent in the study area during the time of interest. In the case of this study, the information available in the literature on the 'mobility' patterns of potters in the Bronze Age Aegean was limited and mostly comprised inferences made from the analysis of pottery. A more promising avenue for future research could be through combining such information with different lines of evidence, such as the analysis of biomolecular data through ancient DNA and strontium isotope techniques in order to establish the most common ways that individuals and/or entire communities moved from community to community or from region to region. For instance, such an approach could offer empirical support for ethnographically derived models of patrilocality and matrilocality, where the introduction of new forming techniques happens through the relocation of female/male potters to their spouse's community.

9.4. Avenues for future research

Apart from the use of biomolecular data to lend empirical support to agent behaviour rules, there are also system-level aspects affecting cultural transmission processes that could be addressed through simple modifications of the simulation model explored in this study. Future modelling efforts could focus on exploring the system under non-equilibrium conditions — in this study, communication costs and population size were treated as remaining stable over time, but, as mentioned earlier, there were improvements in maritime travelling technology during the Bronze Age.

Population levels also experienced fluctuations during the Bronze Age. In this study I explored a simple model with population fluctuations in section 8.4, and its outputs did not produce a good fit to the empirical record compared to the models where the population size remains stable over time. There are ways in which the model could be extended by incorporating additional archaeological information regarding the mechanism driving the introduction of the new traits (and especially the wheel) in the Aegean in order to get a better fit to the empirical data.

The most informative addition to the model would be a mechanism incorporating demic diffusion — in this case the potter's wheel would be introduced in the Aegean by new communities migrating from other regions, and the increase in population size would be associated with such migratory episodes. There is also empirical evidence to support this assumption; the literature on the potter's wheel suggests that the potter's wheel was first introduced in the Aegean islands and mainland by groups of potters migrating from Western Anatolia. However, incorporating demic diffusion into the model would require having a better picture of what drove the movement of the new communities in space that what is now available in the archaeological record — though research using biomolecular techniques should shed more light on migration patterns in the future. Such a model would also explore the effects of having an episode of migration where the new communities settling in the Aegean would have no knowledge of the technology/trait. This would allow us to explore what would happen if, following the population bust at the end of the Late Bronze Age, the Aegean was repopulated by new communities who had no knowledge of the potter's wheel and only used hand-building techniques. Incorporating this mechanism into the model is as simple as setting the trait proportion value of the new communities to 0 rather than the value of their nearest neighbour (which was the approach I took for the implementation of the model in Section 8.4).

Note that this approach makes population size as the number of communities independent of the transmission of the trait and therefore only allows modelling the transmission of selectively neutral innovations that do not improve or worsen the adopters' reproductive rates. Modifying the model in a way that couples the equation describing the population level over time to the equation describing the transmission of the trait would allow exploring the transmission trajectories of adaptive traits like agricultural practices.

Another promising future direction is exploring the adoption of innovations whose local adoption at the community level depends to some degree on the availability of resources from the environment, for example the availability of different types of geological resources, minerals, metals, etc. This is possibly relevant for the transmission of the IRT as discussed in chapter 4, as some archaeologists attribute the 'success' of the lustrous decorated slip in Crete to the availability of clays rich in calcium and illite (Betancourt et al. 2021). In this case we should expect two processes producing empirical signals of positive autocorrelation: the spatial transmission of the trait through inter-community contacts and spatial autocorrelation due to a common environmental variable (local geology). A simple way to explore such processes by modifying the main model would be to localise the parameter *a*, but rather than assigning each community with a random *a* value, this value could instead be drawn from a probability distribution by sampling each community's local geology. For example, if the innovation under investigation was a marble-working technique, a community whose local geology was rich in marble deposits would be assigned an *a* value equal to 1.

Finally, and specifically regarding the potter's wheel innovation, the model did not explore the transmission of different Rotative Kinetic Energy-based wheel-coiling methods and how they evolved over time, due to the lack of available empirical data. Since the use of Roux and Courty's (1998) identification framework is becoming increasingly popular, it is possible that in the future a larger dataset could be built in order to trace their adoption over time quantitatively.

Formulating a simulation model that incorporates different Rotative Kinetic Energy-based wheel-coiling methods would require a clearer theory on how the methods relate to one another. The literature suggests that there is an element of increasing difficulty linked to the adoption of these methods — each method uses an increasing amount of Rotative Kinetic Energy for different shaping operations (Roux and Courty 1998: 750). However, Roux and Courty argue that the methods "are not constrained by technical factors" and that "they may represent cultural behaviour" (Roux and Courty 1998: 747), which suggests that different methods are associated with use by different cultural groups. An agent-based model of a single community with different groups of potters and different degrees of cultural boundaries affecting the cultural transmission between them would be a good place to start to explore the mechanisms associated with the transmission of different methods at the level of the individual community/site.

9.5. Conclusion

This study has shown that there is much to be gained by adopting simulation modelling as a method to explain the transmission patterns of technological traits using data from the Aegean Bronze Age. While the integration of primary data to inform model building and build secondary datasets to validate such models is often time-consuming and laborious, the results of such efforts can be very rewarding in that they allow insights into the role of system-level factors which are often not captured by the empirical data. In the case of this study, simulation allowed experimenting with aspects of the Aegean's geography, the effects of travel costs afforded by different technologies enabling long-distance travel, the effects of trait loss due to chance, and empirically unobserved social factors (the rate of local adoption). After performing statistical analysis to incorporate the effects of temporal uncertainty on the empirical data, it was also possible to compare the empirical trend for the potter's wheel innovation to simulation outputs and suggest that the observed pattern of transmission and adoption at the level of the system (Aegean) had probably less to do with the local rate of adoption (which could be affected by decisions of local elites or the presence of social boundaries within communities) and more to do with the costs of communication between potters located at different communities.

The selection of an appropriate theoretical framework for addressing the research questions was very important in terms of formulating the model on the basis of reasonable assumptions about the mechanisms enabling inter-community cultural transmission. Cultural evolution combined with insights from ethnoarchaeology is a productive framework for investigating processes of change over time in the material culture of the Bronze Age Aegean.

References

Agresti. A. 2018. *Statistical Methods for the Social Sciences, Global Edition.* Pearson Education, Limited.

Alcock S.E. & Cherry J.F. 2004. 'Introduction'. In S.E. Alcock & J.F. Cherry (eds). *Side-by-side survey: Comparative regional studies in the Mediterranean world*. Oxford: Oxbow Books, 1-9.

Ammerman, A. J. & L. L. Cavalli-Sforza. 1971. Measuring the Rate of Spread of Early Farming in Europe. *Man* 6: 674-688.

Ammerman, A. J., & Cavalli-Sforza, L. L. 1979. The wave of advance model for the spread of agriculture in Europe. In C. Renfrew & K. L. Cooke (Eds.), *Transformations: Mathematical Approaches to Culture Change* (pp. 275–294). New York: Academic Press.

Andreou, S., & Psaraki, K. 2007. Tradition and innovation in the Bronze Age pottery of the Thessaloniki Toumba. *The Struma/Strymon River Valley in Prehistory. Proceedings of the International Symposium 'Strymon Prehistoricus.'*

Angelopoulou, A., 2008. The 'Kastri Group': evidence from Korfari ton Amygdalion (Panormos) Naxos, Dhaskalio Keros and Akrotiri Thera, in N.J. Brodie, Doole, G. Gavalas & C. Renfrew (eds.) *Horizon: A Colloquium on the Prehistory of the Cyclades*. Cambridge: McDonald Institute for Archaeological Research, 149–64.

Alram-Stern, E., & Horejs, B. 2018. Pottery Technologies in the Aegean and Anatolia During the 3rd Millennium BC: An Introduction. In E. Alram-Stern & B. Horejs (eds.). *Pottery technologies and sociocultural connections between the Aegean and Anatolia during the 3rd millennium BC*. Austrian Academy of Sciences, Vienna, 9-22.

Argyriou, A. V., Teeuw, R. M., & Sarris, A. 2017. GIS-based landform classification of Bronze Age archaeological sites on Crete Island. *PLoS One*, *12*(2), e0170727.

Arnold, D. E. 2006. The threshold model for ceramic resources: a refinement. *BAR International Series*, *1553*, 3.

Baldi, J., & Roux, V. 2016. The innovation of the potter's wheel: a comparative perspective between Mesopotamia and the southern Levant. *Levant*, *48*(3), 236-253.

Ballut, C., Renard, J., Cavanagh, W., & Orgeolet, R. 2017. Pottery Firing Structures in the Early Mediterranean: Micromorphological Evidence and Archaeological Data from Middle Neolithic Kouphovouno (Southern Greece). *European Journal of Archaeology*, 20(1), 98-119. doi:10.1017/eaa.2016.4.

Barceló, J. A., & Castillo, F. D. 2016. Simulating the past for understanding the present. A critical review. *Simulating prehistoric and ancient worlds*, 1-140.

Barrett, B. J. 2019. Equifinality in empirical studies of cultural transmission. *Behavioural processes*, *161*, 129-138.

Bass, F. M. 1969. A new product growth for model consumer durables. *Management science*, *15*(5), 215-227.

Benincà, E., Jöhnk, K. D., Heerkloss, R., & Huisman, J. 2009. Coupled predator–prey oscillations in a chaotic food web. *Ecology letters*, *12*(12), 1367-1378.

Bentley, R. A., & Shennan, S. J. 2003. Cultural transmission and stochastic network growth. *American Antiquity*, *68*(3), 459-485.

Bentley, R. A., & O'Brien, M. J. 2012. Cultural evolutionary tipping points in the storage and transmission of information. *Frontiers in Psychology*, *3*, 569.

Berg, I. 2004. The meanings of standardisation: conical cups in the late Bronze Age Aegean. *Antiquity*, *78*(299), 74-85.

Berg, I. 2007. Meaning in the making: The potter's wheel at Phylakopi, Melos (Greece*). Journal of Anthropological Archaeology*, 26(2), 234-252.

Berg, I. 2009. X-radiography of Knossian Bronze Age vessels: assessing our knowledge of primary forming techniques. *Annual of the British School at Athens,* 104, 137-173.

Berg, I. 2011. What's in a forming technique? An Investigation into wheel-throwing and wheel-coiling in Bronze Age Crete. *The Old Potter's Almanack*, 16(2), 9-12.

Berg, I. 2015. Potting Skill and Learning Networks in Bronze Age Crete. In: W. Gauss, G. Klebinder-Gauss, C. von Rüden, eds. *The Transmission of Technical Knowledge in the Production of Ancient Mediterranean Pottery, Proceedings of the International Conference at the Austrian Archaeological Institute at Athens, 23rd – 25th November 2012.* Wien: Österreichisches Archäologisches Institut, pp. 17–34.

Betancourt, P. P., Ferrence, S. C., Apostolakou, V., Brogan, T. M., Nodarou, E., & Hsu, F. S. C. 2016. Origins of the Mycenaean Lustrous Dark-on-Light Pottery Technology. *Ra-pi-ne-u: Studies on the Mycenaean World Offered to Robert Laffineur for His 70th Birthday*, 63-8.

Betancourt, P. P., Apostolakou, V., Brogan, T. M., Betancour, P. P., Bonga, L., Chalikias, K., Dierckx, H. M. C., Gluckman, A. S., Hsu, F. S. C., Nodarou, E., Sarasin, S. R., & Stamos, A. 2021. Conclusions on the Clay Vessels and Their Fabrics. In P. P. Betancourt, V. Apostolakou, & T. M.

Brogan (Eds.), *Bramiana: Salvaging Information from a Destroyed Minoan Settlement in Southeast Crete* (Vol. 66, pp. 109–116). INSTAP Academic Press. https://doi.org/10.2307/j.ctv21wj5sj.25.

Bevan, A. 2020. Spatial point patterns and processes. In M. Gillings, P. Hacigüzeller, G. Lock (Eds.), *Archaeological Spatial Analysis*, Routledge pp. 60-76.

Bevan, A., C. Frederick and A. Krahtopoulou 2003. A digital Mediterranean countryside: GIS approaches to the spatial structure of the post-medieval landscape on Kythera (Greece), *Archeologia e Calcolatori* 14: 217-236.

Bevan, A., Conolly, J., & Tsaravopoulos, A. 2006. The fragile communities of Antikythera. Archaeology International, 10(1).

Bevan, A., & Conolly, J. 2013. *Mediterranean Islands, Fragile Communities and Persistent Landscapes: Antikythera in Long-Term Perspective* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139519748

Bevan, A., & Crema, E. R. 2021. Modifiable reporting unit problems and time series of long-term human activity. *Philosophical Transactions of the Royal Society B, 376*(1816), 20190726.

Bilgen, A. N., Coşkun, G., Bilgen, Z., Ünan, N., Silek, S., Çırakoğlu, S., ... & Dikmen, B. 2013. Seyitömer höyük 2012 yılı kazı raporu. *Seyitömer Höyük Kazısı Raporu (2011-2012)*, 1-106.

Blegen, C. W., Caskey, J. L., & Rawson, M. 1951. *Troy, the third, fourth and fifth settlements.* Princeton, NJ: Princeton University Press.

Blitzer, H. 1990. ΚΟΡΩ ΝΕΪΚΑ: Storage-Jar Production and Trade in the Traditional Aegean. *Hesperia*, 675-711.

Bokulich, A. 2014. How the tiger bush got its stripes:‘how possibly’vs.‘how actually’model explanations. *The Monist*, *97*(3), 321-338.

Bonada, N., Dolédec, S. and Statzner, B. 2012. Spatial autocorrelation patterns of stream invertebrates: exogenous and endogenous factors. *Journal of Biogeography*, 39: 56-68. https://doi.org/10.1111/j.1365-2699.2011.02562.x

Boyd, R., & Richerson, P. J. 1985. *Culture and the evolutionary process*. Chicago: The University of Chicago Press.

Brainard, L. 2020. How to Explain How-Possibly. *Philosophers*, *20*(13).

Brandon, R. 1990. *Adaptation and Environment*, Princeton: Princeton University Press.

Branigan, K. 1991. Kommos II. The Final Neolithic through Middle Minoan III Pottery by Philip P. Betancourt. Echos du monde classique: *Classical views*, 35(2), 246-247.

Branigan, K. 2001. *Urbanism in the Aegean Bronze Age* (Sheffield Studies in Aegean Archaeology). Sheffield Academic Press.

Brea, L. 1956. A Bronze Age House at Poliokhni (Lemnos). *Proceedings of the Prehistoric Society*, 21, 144-155. doi:10.1017/S0079497X00017503.

Broodbank, C. 1989. The longboat and society in the Cyclades in the Keros-Syros culture. *American Journal of Archaeology*, 93(3), 319-337.

Broodbank, C. 2002. *An island archaeology of the early Cyclades*. Cambridge University Press.

Broodbank, C. 2010. “Ships a-sail from over the rim of the sea”: Voyaging, sailing and the making of Mediterranean societies c. 3500-500 BC. In Anderson, A., Barrett, J. H., and Boyle, K., (eds). *The Global Origins of Seafaring, McDonald Institute Monographs.* McDonald Institute for Archaeological Research, Cambridge; pp. 255–256.

Broodbank, C. 2013. *The Making of the Middle Sea: A History of the Mediterranean from the Beginning to the Emergence of the Classical World*. Thames & Hudson.

Broodbank, C., & Kiriatzi, E. 2007. The first 'Minoans' of Kythera re-visited: Technology, demography and landscape in the Early Bronze Age Aegean. *American Journal of Archaeology*, 111 (2), 241–274.

Brughmans, T. 2013. Thinking through networks: a review of formal network methods in archaeology. *Journal of archaeological method and theory*, *20*(4), 623-662.

Brughmans, T., Hanson, J., Mandich, M., Romanowska, I., Rubio-Campillo, X., Carrignon, S., ... & Ritondale, M. 2019. Formal modelling approaches to complexity science in Roman studies: a manifesto. *Theoretical Roman Archaeology Journal*, *2*(1), 1-19.

Buskell, A. 2017. Cultural attractor theory and explanation. *Philosophy, Theory, and Practice in Biology*, *9*(13).

Cadogan, G. 2011. Myrtos: From Phournou Koryphi to Pyrgos. *Hesperia Supplements*, 44, 39-49.

Caloi, I. 2011. «Le innovazioni tecnologiche nella Messarà: dal wheel-fashioning al wheel-throwing», in F. Carinci, N. Cucuzza, P. Militello & O. Palio (eds.), *Kretes Minoidos. Tradizione e identità minoica tra produzione artigianale, pratiche cerimoniali e memoria del passato* (Studi di Archeologia Cretese 10), Padova, p. 81-96.

Caloi, I. 2020. Breaking with Tradition? The Adoption of the Wheel-Throwing Technique at Protopalatial Phaistos: Combining Macroscopic Analysis, Experimental Archaeology and

Contextual Information. *Annuario della Scuola Archeologica di Atene e delle Missioni Italiane in Oriente*, Volume 97, 2019, 97, 9.

Campbell, D. T. 1974. “Evolutionary Epistemology.” In: Schlipp, P. A. (ed.), *The Philosophy of Karl Popper* Vol. I, 413-459. Illinois: La Salle

Campbell, D. T. 1960. “Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes.” *Psychological Review*, *67*(6): 380-400.

Catling, H. W., Hughes-Brock, H., Brodie, N., Jones, G., Jones, R. E., & Tomlinson, J. E. 2009. Sparta: Menelaion I: The Bronze Age. Volume 1: Text [CD-ROM content]. *The British School at Athens. Supplementary Volumes*, *45*, CD1–CD179. http://www.jstor.org/stable/40856090.

Cavalli-Sforza, L. L., & Feldman, M. W. 1981. *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.

Chapman, R., Wylie, A. 2016. *Evidential Reasoning in Archaeology*; Bloomsbury Academic: London, UK.

Chatzi-Vallianou, D. 1997 The potters’ quarter in LM III Gouves, in: R. Laffineur, P. P. Betancourt (eds.), *TEXNH. Craftsmen, Craftswomen and Craftsman-ship in the Aegean Bronze Age, Proceedings of the 6th International Aegean Conference, University of Liège 18–21 April 1996*, Aegaeum 16, Liège, pp. 333–345.

Cherry, J. F., Davis, J. L., & Mantzourani, E. 1991. *Landscape Archaeology as long-term History: Northern Keos in the Cycladic Islands from earliest settlement until Modern times* (No. 16). Cotsen Institute of Archaeology.

Choleva, M. 2012. The first wheelmade pottery at Lerna: wheel-thrown or wheel-fashioned?. *Hesperia*, 81(3), 343-381.

Choleva, M. 2015. Περιστρέφοντας τον κεραμικό τροχό στο Παλαμάρι Σκύρου: νέες τεχνολογικές πρακτικές και κεραμικές παραδόσεις στο Αιγαίο της Πρώιμης Εποχής του Χαλκού. In L. Parlama, M.D. Theochari, C. Romanou and St. Bonatsos (eds). *Ο Οχυρωμένος Προϊστορικός Οικισμός στο Παλαμάρι Σκύρου. Διεπιστημονική Συνάντηση για το Έργο Έρευνας και Ανάδειξης*. (Athens), Αθήνα: ΥΠΕΠ: 143-166.

Choleva M. 2018. Craft behaviours during a period of transformations. The introduction and adoption of the potter's wheel in central Greece during Early Bronze Age, in: I. Caloi & C. Langohr (eds.*), Technology in Crisis. Technological changes in ceramic production during periods of trouble*, Louvain-la-Neuve: Presses Universitaires de Louvain: 45-74.

Choleva, M. 2020. Travelling with the potter's wheel in the Early Bronze Age Aegean. *Annual of the British School at Athens*, *115*, 59-104.

Choleva, M., Jung, R., & Kardamaki, E. 2020. Working on the potter's wheel: technological insights into Mycenaean pottery production. *Ägypten und Levante: Internationale Zeitschrift für ägyptische archäologie und deren nachbargebeite*, (30), 219-282.

Christakis K. S. 1996. Craft Specialisation in Minoan Crete: The Case for Itinerant Pithos Makers, in B. Rutkowski and K. Nowicki (eds.) *Studies and Monographs in Mediterranean Archaeology and Civilization, ser. II, vol. 4. Art and Archaeology, vol. 3*, Warsaw, 63-74.

Christakis, K. S. 2004. Palatial Economy and Storage in Late Bronze Age Knossos. *British School at Athens Studies*, 299-309.

Collar, A., Coward, F., Brughmans, T., & Mills, B. J. 2015. Networks in archaeology: phenomena, abstraction, representation. *Journal of Archaeological Method and Theory*, *22*, 1-32.

Crema, E. R. 2012. Modelling temporal uncertainty in archaeological analysis. *Journal of Archaeological Method and Theory*, *19*(3), 440-461.

Crema, E. R. 2015. Modelling settlement rank-size fluctuations. In *Agent-based Modeling and Simulation in Archaeology* (pp. 161-181). Springer, Cham.

Crema, E. R., Kandler, A., & Shennan, S. 2016. Revealing patterns of cultural transmission from frequency data: equilibrium and non-equilibrium assumptions. Scientific Reports, 6(1), 39122–39122.

Crutchfield, J. P. 2008. Is anything ever new? Considering emergence. In Bedau, M. A. and Humphreys, P., editors, *Emergence: Contemporary Readings in Philosophy and Science*, pages 269–286. MIT Press, Cambridge, MA.

Cullen, T. 1985. Social implications of ceramic style in the Neolithic Peloponnese. In:
Kingery, W.D. (Ed.), *Ancient Technology to Modern Science. American Ceramic Society, Columbus*, Ohio, pp. 77–100.

Cultraro, M. 2005. *Aegeans on Smoke-shrouded Lemnos : A Re-assessment of the Mycenean Evidence from Poliochni and Other Sites.* Université de Liège Histoire de l'art et archéologie de la Grèce antique.

Cultraro, M. 2007. Domestic architecture and public space in Early Bronze Age Poliochni (Lemnos). *British School at Athens Studies*, 55-64.

Dakouri-Hild, A. 2001. Plotting fragments : a preliminary assessment of the Middle Helladic settlement in Boeotian Thebes. *Urbanism in the Aegean Bronze Age*, 103-18.

Dalfes, H. N., Kukla, G., & Weiss, H. (Eds.). 2013. *Third millennium BC climate change and old-world collapse* (Vol. 49). Springer Science & Business Media.

Day, P. M., Relaki, M., & Faber, E. W. 2006. Pottery making and social reproduction in the Bronze Age Mesara. In *Pottery and society: the impact of recent studies in Minoan pottery. Gold medal colloquium in honor of Philip P. Betancourt* (pp. 22-72). Archaeological Institute of America Boston.

Doudalis, G. 2018. An age of transition: ceramic innovations and social changes in the Mirabello gulf during the Middle Minoan period. *Българско е-Списание за Археология*, 8(2), 163-178.

Dray, W. H. 1957. *Laws and explanation in history*. Oxford: Oxford University Press.

Dray, W. H. 1968. On explaining how-possibly. *The Monist*, *52*(3), 390-407.

Dyson-Hudson, R., & Smith, E. A. 1978. Human territoriality: an ecological reassessment. *American Anthropologist*, *80*(1), 21-41.

Eder, B. and Lemos, I. S. 2019. From the Collapse of the Mycenaean Palaces to the Emergence of Early Iron Age Communities. *A Companion to the Archaeology of Early Greece and the Mediterranean*, 133-160.

Eerkens, J. W. & Lipo, C. P. 2007. Cultural Transmission Theory and the Archaeological Record: Providing Context to Understanding Variation and Temporal Changes in Material Culture. *Journal of Archaeological Research*, *15*(3), 239–274.

Evely, D. 1988. The potters' wheel in Minoan Crete. *Annual of the British School at Athens,* 83, 83-126.

Forrester, J. W. 1971. Counterintuitive behavior of social systems. *Theory and decision*, *2*(2), 109-140.

Fort, J., & Méndez, V. 2002. Wavefronts in time-delayed reaction-diffusion systems. Theory and comparison to experiment. *Reports on Progress in Physics*, *65*(6), 895.

Fortunato, L., & Archetti, M. 2010. Evolution of monogamous marriage by maximization of inclusive fitness. *Journal of evolutionary biology*, *23*(1), 149-156.

Gandon, E., & Roux, V. 2019. Cost of motor skill adaptation to new craft traits: Experiments with expert potters facing unfamiliar vessel shapes and wheels. *Journal of Anthropological Archaeology*, *53*, 229-239.

Gauss W. and Kiriatzi E., 2011, Pottery Production and Supply at Bronze Age Kolonna, Aegina. *An Integrated Archaeological and Scientific Study of a Ceramic Landscape,* Wien.

Goldman, H. 1931. *Excavations at Eutresis in Boeotia*. Cambridge, Mass.: Harvard University Press.

Gomart, L., & Burnez-Lanotte, L. 2012. Techniques de façonnage, production céramique et identité de potiers: une approche technologique de la céramique de style non rubané du site du Staberg à Rosmeer (Limbourg, Belgique). *Bulletin de la Société préhistorique française*, 231-250.

Gorogianni, E., Abell, N., & Hilditch, J. 2016. Reconsidering technological transmission: the introduction of the potter’s wheel at Ayia Irini, Kea, Greece. *American Journal of Archaeology*, 120(2), 195-220.

Gosselain, O. P. 2000. Materializing identities: an African perspective. *Journal of Archaeological Method and Theory*, *7*(3), 187-217.

Gosselain, O. 2002. *Poteries du Cameroun méridional: styles techniques et rapports à l'identité*.

Gosselain, O.P. (2008). Mother Bella was not a Bella. Inherited and transformed traditions in Southwestern Niger. In Stark, M., Bowser, B. and Horne, L. (eds). *Cultural Transmission and Material Culture: Breaking Down Boundaries*. The University of Arizona Press, Tucson, pp. 150–177.

Gosselain, O. P. 2011. Fine if I do, fine if I don’t. Dynamics of technical knowledge in Sub-Saharan Africa. In B.W. Roberts and M. Vander Linden (eds.). *Investigating archaeological cultures: Material culture, variability and transmission*: 211-227. New York: Springer.

Gosselain, O.P. and Livingstone Smith, A. 2005. The Source. Clay Selection and Processing Practices in Sub-Saharan Africa. In A. Livingstone Smith, D. Bosquet, and R. Martineau. (eds). *Pottery Manufacturing Processes: Reconstruction and Interpretation*, British Archaeological Reports International Series 1349, Oxford, pp. 33–47.

Haggis, D. 1996. Archaeological Survey at Kavousi, East Crete: Preliminary Report. *Hesperia: The Journal of the American School of Classical Studies at Athens*, *65*(4), 373-432. Retrieved November 11, 2020, from http://www.jstor.org/stable/148436.

Hale, C. M. 2016. The Middle Helladic Fine Gray Burnished (Gray Minyan) Sequence at Mitrou, East Lokris. *Hesperia: The Journal of the American School of Classical Studies at Athens*, 85(2), 243-295.

Hasaki, E. 2002. *Ceramic Kilns in Ancient Greece: Technology and Organization of Ceramic Workshops* [Doctoral dissertation, University of Cincinnati]. OhioLINK Electronic Theses and Dissertations Center. http://rave.ohiolink.edu/etdc/view?acc_num=ucin1023219003.

Hayden, B., 2003. *Reports on the Vrokastro area, eastern Crete. Vol. 1, Catalogue of pottery from the Bronze and early Iron Age settlement of Vrokastro in the collections of the University of Pennsylvania Museum of Archaeology and Anthropology and the Archaeological Museum, Herakleion, Crete.* Philadelphia: University of Pennsylvania, Museum of Archaeology and Anthropology.

Hayden, B., 2004. *Reports on the Vrokastro area, eastern Crete. vol. 2, The settlement history of the Vrokastro area and related studies.* (University Museum Monograph 119) Philadelphia: University of Pennsylvania, Museum of Archaeology and Anthropology.

Hayden, B., 2005. *Reports on the Vrokastro area, eastern Crete. Vol. 3, The Vrokastro regional survey project: sites and pottery. (University Museum Monograph 123)* Philadelphia: University of Pennsylvania, Museum of Archaeology and Anthropology.

Henrich, J. 2001. Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist*, *103*(4), 992-1013.

Henrich, J. 2004. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses—the Tasmanian case. *American Antiquity*, *69*(2), 197-214.

Henrich, J., Boyd, R., & Richerson, P. J. 2008. Five misunderstandings about cultural evolution. *Human Nature*, *19*(2), 119-137.

Henrich J. 2016. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton, NJ: Princeton University Press.

Herbich, I. 1987. Learning patterns, potter interaction and ceramic style among the Luo of Kenya. *African Archaeological Review*, *5*(1), 193-204.

Hilditch, J. R. 2008. *Reconstruction of technological choice, social practice and networks of exchange from a ceramic perspective in the Middle Bronze Age Cyclades* (Doctoral dissertation, University of Exeter).

Hilditch, J., Jeffra, C., & Opgenhaffen, L. 2021. The Tracing the Potter's Wheel Project (TPW): An Integrated Archaeological Investigation of the Potter's Wheel in the Bronze Age Aegean. *Interdisciplinaria Archaeologica*, *12*(2), 345-355.

Hood, M. 1965. Minoan Sites in the Far West of Crete. *The Annual of the British School at Athens*, 60, 99-113. Retrieved May 11, 2020, from www.jstor.org/stable/30103148.

Hood, S., Clutton-Brock, J., & Bialor, P. 1981. EXCAVATIONS IN CHIOS 1938-1955: PREHISTORIC EMPORIO AND AYIO GALA. VOLUME I. *The British School at Athens. Supplementary Volumes*, (15), I-425. Retrieved March 3, 2021, from http://www.jstor.org/stable/40855998.

Jameson, M. N., Runnels, C. N., & van Andel, J. H. 1994. *A Greek Countryside. The Southern Argolid from Prehistory to the Present Day*. Stanford.

Jeffra, C. 2013. A re-examination of early wheel potting in Crete. *Annual of the British School at Athens*, 108, 31-49.

Johnson, I. 2004. Aoristic analysis: seeds of a new approach to mapping archaeological distributions through time. *BAR International Series*, *1227*, 448-452.

Jones, R. 2021. The Decoration and Firing of Ancient Greek Pottery: A Review of Recent Investigations. *Advances in Archaeomaterials*, *2*(2), 67-127.

Jordan, P., & Shennan, S. 2003. Cultural transmission, language, and basketry traditions amongst the California Indians. *Journal of Anthropological Archaeology*, *22*(1), 42-74.

Joukowsky, M., & Benjamin, A. 1986. Prehistoric Aphrodisias : An account of the excavations and artifact studies / by Martha Sharp Joukowsky ; with contributions by Anna S. Benjamin ... [et al.]. (Archaeologia transatlantica ; 3). Providence, R.I. : Louvain-la-Neuve, Belgique: Brown University, Center for Old World Archaeology and Art ; Institut supérieur d'archéologie et d'histoire de l'art, Collège Erasme.

Kalantzopoulou, T. 2022. Living on the Edge: Habitation on the Uplands of East Crete–Preliminary Results from an Extensive Survey. South by Southeast: The History and Archaeology of Southeast Crete from Myrtos to Kato Zakros, 25.

Kaliampakos, D., Mavrikos, A., & Menegaki, M. 2011. Construction industry and archaeology: a land-use conflict on the island of Andros, Greece. *International Journal of Mining, Reclamation and Environment*, *25*(2), 152-160.

Karkanas, P., Berna, F., Fallu, D., & Gauß, W. 2019. Microstratigraphic and mineralogical study of a Late Bronze Age updraft pottery kiln, Kolonna site, Aegina Island, Greece. *Archaeological and Anthropological Sciences*, *11*, 5763-5780.

Katsarou, S., & Schilardi, D. 2004. Emerging Neolithic and Early Cycladic Settlements in Paros: Koukounaries and Sklavouna. *The Annual of the British School at Athens,* 99, 23-48. Retrieved March 10, 2020, from www.jstor.org/stable/30071530.

Katsonopoulou, D., Iliopoulos, I., Katsarou, S., & Xanthopoulou, V. 2016. Craftsmanship of big storage pithoi in the Early Helladic settlement of Helike, Achaea. In *Proceedings of the□ th Symposium of the Hellenic Society for Archaeometry.*

Kandler, A., & Powell, A. 2018. Generative inference for cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1743), 20170056.

Kelly, R. L. 1992. Mobility/sedentism: concepts, archaeological measures, and effects. *Annual Review of Anthropology*, 43-66.

Kiamil, T. 1980. The Yorton Culture within the Early Bronze Age of Western Anatolia. Dissertation, University of London. URL: http://discovery.ucl.ac.uk/id/eprint/1317604.

Kim, J. 2001. Elite strategies and the spread of technological innovation: The spread of iron in the Bronze Age societies of Denmark and southern Korea. *Journal of Anthropological Archaeology*, *20*(4), 442-478.

Kiriatzi, E. 2010. 'Minoanising'Pottery Traditions in the Southwest Aegean during the Middle Bronze Age: Understanding the Social Context of Technological and Consumption

Practice. *Mesohelladica: La Grèce continentale au Bronze Moyen. Athens: École française d'Athènes*, 683-699.

Kiriatzi, E., & Andreou, S. 2016. Mycenaean and Mycenaeanising Pottery across the Mediterranean: A Multi-Scalar Approach to Technological Mobility, Transmission and Appropriation. In E. Kiriatzi & C. Knappett (eds.), *Human Mobility and Technological Transfer in the Prehistoric Mediterranean* (British School at Athens Studies in Greek Antiquity, pp. 128-153). Cambridge: Cambridge University Press. doi:10.1017/9781316536063.010.

Kiriatzi, E., & Knappett, C. eds. 2016. *Human mobility and technological transfer in the prehistoric Mediterranean*. Cambridge University Press.

Knappett, C., Evans, T., & Rivers, R. 2008. Modelling maritime interaction in the Aegean Bronze Age. *Antiquity*, *82*(318), 1009-1024.

Knappett, C., & Van Der Leeuw, S. 2014. A developmental approach to ancient innovation: the potter's wheel in the Bronze Age east Mediterranean. *Pragmatics & Cognition*, 22(1), 64-92.

Knappett, C. 2016. Resisting innovation? Learning, cultural evolution and the potter's wheel in the Mediterranean bronze age. *Cultural phylogenetics: Concepts and applications in archaeology*, 97-111.

Kohler, T. A., VanBuskirk, S., & Ruscavage-Barz, S. 2004. Vessels and villages: evidence for conformist transmission in early village aggregations on the Pajarito Plateau, New Mexico. *Journal of Anthropological Archaeology*, *23*(1), 100-118.

Kopaka, K., Theou, E. 2018. Gavdos or living on the southernmost Aegean island in the Neolithic cultural horizons. In S. Dietz, F. Mavridis, Ž. Tankosić & T. Takaoglu (eds.). *Communities in Transition: The Circum-Aegean Area During the 5th and 4th Millennia BC,* Oxford, 441-55.

Kramer-Hajos, M., & O'Neill, K. 2008. The Bronze Age site of Mitrou in East Lokris: Finds from the 1988-1989 surface survey. *Hesperia*, 163-250,

Kvamme, K. L. 1990. Spatial autocorrelation and the Classic Maya collapse revisited: Refined techniques and new conclusions. *Journal of Archaeological Science*, *17*(2), 197-207.

Lake, M.W. 2015. Explaining the Past with ABM: On Modelling Philosophy. In G. Wurzer, K. Kowarik, H. Reschreiter. (eds). *Agent-based Modeling and Simulation in Archaeology. Advances in Geographic Information Science.* Springer, Cham. https://doi.org/10.1007/978-3-319-00008-4_1

Laland, K., Matthews, B., & Feldman, M. W. 2016. An introduction to niche construction theory. *Evolutionary ecology*, *30*, 191–202. doi:10.1007/s10682-016-9821-z

Lewis, H. M., & Laland, K. N. 2012. Transmission fidelity is the key to the build-up of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1599), 2171-2180.

Lewontin, R. C. 1983. Gene, Organism and Environment. In: Bendall, editor. *Evolution from molecules to men*. Cambridge: Cambridge University Press.

Lis, B. 2010. Cooking pottery in the Late Bronze Age Aegean–an attempt at a methodological approach. In B. Horejs, R. Jung, P. Pavúk (eds.), *Analysing Pottery. Processing–Classification–Publication, Comenius University in Bratislava, Bratislava*, 235-244.

Livingstone-Smith, A. 2000. Processing clay for pottery in northern Cameroon: Social and technical requirements. *Archaeometry*, *42*(1), 21-42.

Livarda, A., Orengo, H. A., Canellas-Bolta, N., Riera-Mora, S., Picornell-Gelabert, L., Tzevelekidi, V., ... & Krahtopoulou, A. 2021. Mediterranean polyculture revisited: Olive, grape and subsistence strategies at Palaikastro, East Crete, between the Late Neolithic and Late Bronze Age. *Journal of Anthropological Archaeology*, *61*, 101271.

Maniatis, Y. 2009. The Emergence of Ceramic Technology and its Evolution as Revealed with the use of Scientific Techniques. In A.J., Freestone, I.C., Rehren, T. (Eds.) *Advances in the Study of Ancient Technology*; Shortland. Oxbow Books: Oxford, UK.

Manning, S. W. 2014. Radiocarbon dating and archaeology: history, progress and present status. In *Material Evidence* (pp. 148-178). Routledge.

Maran, J. 2007. *Emulation of Aeginetan pottery in the Middle Bronze Age of coastal Thessaly: regional context and social meaning (Vol. 42, pp. 167-182).* Verlag der Österreichischen Akademie der Wissenschaften.

Marangou, L., Renfrew, C., Doumas, C., & Gavalas, G. 2008. Markiani on Amorgos: an Early Bronze Age fortified settlement–overview of the 1985–91 Investigations. In *Horizon. A Colloquium on the Prehistory of the Cyclades*. Cambridge: McDonald Institute of Archaeological Research Horizon. Cambridge: McDonald Institute of Archaeological Research. p (pp. 97-105).

Marketou, T. 2012. Dodecanese. In E. H. Cline (ed.). *The Oxford Handbook of the Bronze Age Aegean* (2012; online edn, Oxford Academic, 18 Sept. 2012), https://doi.org/10.1093/oxfordhb/9780199873609.013.0057, accessed 1 Mar. 2023.

Marketou, T. 2012. Rhodes. in E. H. Cline (ed.), *The Oxford Handbook of the Bronze Age Aegean* (2012; online edn, Oxford Academic, 18 Sept. 2012), https://doi.org/10.1093/oxfordhb/9780199873609.013.0057, accessed 1 Mar. 2023.

McDonald, W. 1972. Excavations at Nichoria in Messenia: 1969-71. *Hesperia: The Journal of the American School of Classical Studies at Athens,* 41(2), 218-273. Retrieved March 27, 2020, from www.jstor.org/stable/147682.

Mee, C., & Forbes, H. A. (Eds.). 1997. *A Rough and Rocky Place: The Landscape and Settlement History of the Methana Peninsula, Greece: Results of the Methana Survey Project, Sponsored by the British School at Athens and the University of Liverpool*. Liverpool: Liverpool University Press.

Mellink, M. 1986. The Early Bronze Age in West Anatolia. *The end of the Early Bronze Age in the Aegean*, *6*, 139-87.

Menelaou, S., Kouka, O., & Day, P. M. 2016. Pottery production and exchange at the Heraion, Samos during the late third millennium BC: First steps in the study of technology and provenance. *Journal of Archaeological Science: Reports*, 7, 480-488.

Menelaou, S., & Kouka, O. 2021. “Cooking up” pottery recipes: A diachronic technological and provenance analysis of cooking ware ceramics from prehistoric Heraion on Samos, Greece (ca. 4500–1700 BCE). *Journal of Archaeological Science: Reports*, *35*, 102716.

Mentesana, R., Kilikoglou, V., Todaro, S., & Day, P. M. (2019). Reconstructing change in firing technology during the Final Neolithic–Early Bronze Age transition in Phaistos, Crete. Just the tip of the iceberg?. *Archaeological and Anthropological Sciences*, *11*(3), 871-894.

Mesoudi, A., & O'Brien, M. J. 2008. The cultural transmission of Great Basin projectile-point technology I: an experimental simulation. *American Antiquity*, *73*(1), 3-28.

Mesoudi, A. 2016. Cultural evolution: a review of theory, findings and controversies. *Evolutionary biology*, *43*(4), 481-497.

Mesoudi, A. 2021. Blind and incremental or directed and disruptive? On the nature of novel variation in human cultural evolution. *American Philosophical Quarterly*, *58*(1), 7-20.

Mohlenhoff, K. A., & Codding, B. F. 2017. When does it pay to invest in a patch? The evolution of intentional niche construction. *Evolutionary Anthropology: Issues, News, and Reviews*, *26*(5), 218-227.

Morin, O. 2016. Reasons to be fussy about cultural evolution. *Biology & philosophy*, *31*, 447-458.

Morin, O., & Miton, H. 2018. Detecting wholesale copying in cultural evolution. Evolution and *Human Behavior*, 39(4), 392-401.

Μπρούσκαρη, Ε. 1993. *Διατήρηση αρχαίων σε υπόγειο ή σε ακάλυπτο χώρο: Το παράδειγμα της Κω*. Επτάλοφος ΑΒΕΕ.

Müller Celka S., Kiriatzi E., Charalambidou X., Müller N. S. 2018. Early Helladic II-III Pottery Groups from Eretria (Euboea). In E. Alram-Stern & B. Horejs (eds). *Pottery technologies and sociocultural connections between the Aegean and Anatolia during the 3rd millennium BC.* Austrian Academy of Sciences, Vienna.

Mylona, D. 2020. Marine resources and coastal communities in the Late Bronze Age southern Aegean: a seascape approach. *American Journal of Archaeology*, *124*(2), 179-213.

Neiman, F. D. 1995. Stylistic variation in evolutionary perspective: inferences from decorative diversity and interassemblage distance in Illinois Woodland ceramic assemblages. *American antiquity*, *60*(1), 7-36.

Netting, R. M., & Netting, R. M. 1981. *Balancing on an Alp: ecological change and continuity in a Swiss mountain community*. CUP Archive.

Nicklin, K. 1979. The location of pottery manufacture. *Man*, 436-458.

Nikolakopoulou, I., & Knappett, C. 2016. Mobilities in the Neopalatial southern Aegean: the case of Minoanisation. *Human mobility and technological transfer in the prehistoric Mediterranean*, 102-115.

Nodarou, E. 2003. *Pottery production, distribution and consumption in early Minoan West Crete: an analytical perspective* (Doctoral dissertation, University of Sheffield).

Noll, W., & Heimann, R. B. 2016. *Ancient Old World pottery. Materials, technology, and decoration.*

North, D. C. 1990. *Institutions, institutional change and economic performance*. Cambridge University Press.

O'Brien, M., & Laland, K. N. 2012. Genes, Culture, and Agriculture An Example of Human Niche Construction. *Current Anthropology*, *53*(4), 434–470.

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. 2003. *Niche Construction: The Neglected Process in Evolution (MPB-37)*. In: Monographs in population biology, vol 37. Princeton: Princeton University Press.

Orton, C. Tyers, P. & Vince, A. 1993. *Pottery in Archaeology*. Cambridge University Press.

Orton, D. 2017. *archSeries: tools for chronological uncertainty in archaeology*. R package version 0.0.0.9003. https://github.com/davidcorton/archSeries (2017)

Orton, D., Morris, J., & Pipe, A. 2017. Catch per unit research effort: sampling intensity, chronological uncertainty, and the onset of marine fish consumption in historic London. *Open Quaternary*, *3*(1), 1-20.

Perreault, C. 2018. Time-averaging slows down rates of change in the archaeological record. *Journal of Archaeological Method and Theory*, *25*(3), 953-964.

Philippa-Touchais, A., Touchais, G., Abell, N., Balitsari, A., Choleva, M., & Hruby, J. 2020. Argos, fouilles de l'Aspis-2017, 2018. *Bulletin archéologique des Écoles françaises à l'étranger*.

Popper, N. & Pichler, P. 2015. Reproducibility. In: G. Wurzer, K. Kowarik, and H. Reschreiter (eds.). *Agent-based Modeling and Simulation in Archaeology*. Cham: Springer, pp. 77–98. DOI: https://doi.org/10.1007/978-3-319-00008-4_4.

Powell, A., Shennan, S., & Thomas, M. G. 2009. Late Pleistocene demography and the appearance of modern human behavior. *Science*, *324*(5932), 1298-1301.

Premo, L. S. 2005. Patchiness and Prosociality: An Agent-Based Model of Plio/Pleistocene Hominid Food Sharing. In P. Davidsson, K. Takadama, and B. Logan (eds). *MABS 2004, Lecture Notes in Artificial Intelligence*, vol 3415. Springer, Berlin, pp 210–224.

Premo, L. S. 2007. Exploratory Agent-Based Models: Towards an Experimental Ethnoarchaeology. In: J. T. Clark, E. M. Hagemeister (eds). *Digital Discovery: Exploring New Frontiers in Human Heritage.* CAA 2006. Computer Applications and Quantitative Methods in Archaeology, Archeolingua Press, Budapest, pp 29–36.

Premo, L. S. 2008. Exploring Behavioral Terra Incognita with Archaeological Agent-Based Models. In B. Frischer, A. Dakouri-Hild (eds). *Beyond Illustration: 2D and 3D Technologies as Tools of Discovery in Archaeology*. British Archaeological Reports International Series. ArchaeoPress, Oxford, pp 46–138.

Premo, L. S. 2014. Cultural transmission and diversity in time-averaged assemblages. *Current Anthropology*, *55*(1), 105-114.

Premo, L. S. 2016. Effective population size and the effects of demography on cultural diversity and technological complexity. *American Antiquity*, *81*(4), 605-622.

Premo, L. S., & Scholnick, J. B. 2011. The spatial scale of social learning affects cultural diversity. *American Antiquity*, *76*(1), 163-176.

Pullen, D. J. 2011. Redistribution in Aegean palatial societies. Before the palaces. Redistribution and chiefdoms in mainland Greece, *American Journal of Archaeology*, *115*(2), 185–195.

R Core Team 2021. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Rahmstorf, L. 2016. Emerging economic complexity in the Aegean and Western Anatolia during the earlier third millennium BC", in B. Molloy (ed.). *Of Odysseys and Oddities: Scales and modes* *of interaction between prehistoric Aegean societies and their neighbors. Sheffield Aegean Round Table 2013*. Oxford: Oxbow, 225–276.

Ratcliffe, J. H. 2000. Aoristic analysis: the spatial interpretation of unspecific temporal events. *International Journal of Geographical Information Science*, *14*(7), 669-679.

Ratter, B.M. 2012. Complexity and Emergence—key concepts in non-linear dynamic systems. In M. Glaser, G. Krause, B.M. Ratter, and M. Welp (eds.). *Human-Nature Interactions in the Anthropocene: Potentials of Social-Ecological Systems Analysis*. Routledge: New York, NY, USA; pp. 83–101.

Renfrew, C. 1972. *The Emergence of Civilisation: the Cyclades and the Aegean in the third millennium BC*. Taylor & Francis.

Renfrew, C. 2011. *The Emergence of Civilisation: The Cyclades and the Aegean in the Third Millennium BC*. Oxbow Books. https://doi.org/10.2307/j.ctt1vgwb5m.

Renfrew, C., Gimbutas, M., & Elster, E. 1986. *Excavations at Sitagroi, a prehistoric village in northeast Greece volume 1 (Monumenta Archeologica 13)*. Los Angeles (CA): Institute of Archaeology at UCLA.

Renfrew, E., Renfrew, C., & Elster, E. 2003. *Prehistoric Sitagroi: Excavations in northeast Greece, 1968-1970. Vol. 2, The final report* / edited by Ernestine S. Elster and Colin Renfrew. (Monumenta archaeologica (Los Angeles, Calif.) ; v. 20). Los Angeles, CA: Cotsen Institute of Archaeology, University of California, Los Angeles.

Reyes-García, V., Balbo, A. L., Gómez-Baggethun, E., Gueze, M., Mesoudi, A., Richerson, P., ... & Shennan, S. 2016) Multilevel processes and cultural adaptation: examples from past and present small-scale societies. *Ecology and Society: a Journal of Integrative Science for Resilience and Sustainability*, *21*(4).

Rice, P. M. 1987. *Pottery Analysis: A Sourcebook*. Chicago: The University of Chicago Press.

Richerson, P. J., & Boyd, R. 1984. Natural selection and culture. *Bioscience, 34*(7), 430–434.

Rogers, E., M. 2003. *Diffusion of Innovations*. 5th edition. Free Press (Division of Simon and Schuster): New York, United States of America.

Romanowska, I., Brughmans, T., Bes, P., Carrignon, S., Egelund, L., Lichtenberger, A., & Raja, R. 2022. A study of the centuries-long reliance on local ceramics in Jerash through full quantification and simulation. *Journal of Archaeological Method and Theory*, *29*(1), 31-49.

Ross, R. M., Greenhill, S. J., & Atkinson, Q. D. 2013. Population structure and cultural geography of a folktale in Europe. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1756), 20123065.

Roux, V. 1994. La Technique du tournage : définition et reconnais sance par les macrotraces, in D. Binder et J. Courtin (eds.), *Terre cuite et société, la céramique, document technique, économique, culturel*, actes des 14es Rencontres internationales d'archéologie et d'histoire (Antibes, 21-23 octobre 1993), Juan-les-Pins, APCDA, p. 45

Roux, V. 2003. A dynamic systems framework for studying technological change: application to the emergence of the potter's wheel in the southern Levant. *Journal of Archaeological Method and Theory*, *10*(1), 1-30.

Roux, V. 2008. Evolutionary trajectories of technological traits and cultural transmission: a qualitative approach to the emergence and disappearance of the ceramic wheel-fashioning technique in the southern Levant during the fifth to third millennia BC. In M. Stark, B. Bowser, and L. Horne (eds.). *Breaking Down Boundaries: Anthropological Approaches to Cultural Transmission and Material Culture in Memory of Carol Kramer*, 82–104. Tucson: University of Arizona Press.

Roux, V. 2013. Spreading of innovative technical traits and cumulative technical evolution: continuity or discontinuity? *Journal of Archaeological Method and Theory*, *20*, 312-330.

Roux, V. 2015. Standardization of ceramic assemblages: Transmission mechanisms and diffusion of morpho-functional traits across social boundaries. *Journal of Anthropological Archaeology*, *40*, 1-9.

Roux, V. and Corbetta, D. 1989. *The potter's wheel: craft specialization and technical competence*. New Delhi : Oxford & IBH Pub. Co.

Roux, V., Bril, B., & Karasik, A. 2018. Weak ties and expertise: Crossing technological boundaries. *Journal of Archaeological Method and Theory*, *25*(4), 1024-1050.

Roux, V., & Courty, M. A. 1998. Identification of Wheel-fashioning Methods: Technological Analysis of 4th–3rd Millennium BC Oriental Ceramics. *Journal of Archaeological Science*, 25(8), 747-763.

Roux, V., & Jeffra, C. 2015. The spreading of the potter’s wheel in the Ancient Mediterranean: a social context-dependent phenomenon. In Gauß, W., Klebinder-Gauß, G., Von Rüden, C. (eds.). *The Transmission of Technical Knowledge in the Production of Ancient Mediterranean Pottery*. Wien: Österreichisches Archäologisches Institut, 165-182.

Roux, V., & Lara, C. 2016. Why potters have not borrowed the kiln? Comparing narratives of Indian and Ecuadorian potters. *Americae*, http://www.mae.u-paris10.fr/articles-articulos/why-potters-have-notborrowed-the-kiln-comparing-narratives-of-indian-and-ecuadorian-potter

Roux, V. 2019. *Ceramics and society: a technological approach to archaeological assemblages.* Springer.

Rubio-Campillo, X., Montanier, J. M., Rull, G., Lorenzo, J. M. B., Díaz, J. M., González, J. P., & Rodríguez, J. R. 2018. The ecology of Roman trade. Reconstructing provincial connectivity with similarity measures. *Journal of Archaeological Science*, *92*, 37-47.

Rutter, J. 1995. *Lerna: A Preclassical Site in the Argolid III: The Pottery of Lerna IV*. Princeton.

Rutter, J. B. 2012. Migrant drinking assemblages in Aegean Bronze Age settings. In J. Maran and P.W. Stockhammer, (eds.). *Materiality and Social Practice: Transformative Capacities of Intercultural Encounters*, 73-88.

Sagan, C. 1977. *The Dragons of Eden. Speculations on the Evolution of Human Intelligence*. Ballantyne, New York, NY, USA.

Şahoğlu, V. 2005. The Anatolian Trade Network and the İzmir Region During the Early Bronze Age. *Oxford Journal of Archaeology*, 24, 339-361.

Şahoğlu, V., Çayir, Ü., Gündoğan, Ü., & Tuğcu, İ. 2018. Çeşme–Bağlararası: Batı Anadolu sahil kesiminde bir Tunç Çağı yerleşimi.

Scerri, E. M., Thomas, M. G., Manica, A., Gunz, P., Stock, J. T., Stringer, C., ... & Chikhi, L. 2018. Did our species evolve in subdivided populations across Africa, and why does it matter? *Trends in Ecology & Evolution*, *33*(8), 582-594.

Schelling, T. C. 1969. Models of Segregation. *The American Economic Review*, 59(2), 488-493.

Schoep, L., & Knappett, C. 2003. Le Quartier Nu (Malia, Crète). L'occupation du Minoen Moyen II. *Bulletin de correspondance hellénique*, *127*(1), 49-86.

Shennan, S. J., Crema, E. R., & Kerig, T. 2015. Isolation-by-distance, homophily, and "core" vs."package" cultural evolution models in Neolithic Europe. *Evolution and Human Behavior*, *36*(2), 103-109.

Shennan, S. J., & Wilkinson, J. R. 2001. Ceramic style change and neutral evolution: a case study from Neolithic Europe. *American Antiquity*, *66*(4), 577-593.

Shennan, S. 2012. Darwinian cultural evolution. *Archaeological Theory Today*, 15-36.

Shennan, S. J. 2013. Long-term trajectories of technological change. In Richerson, P.J., Christiansen, M.H. (eds.). *Cultural Evolution: Society, Technology, Language, and Religion*. (pp. 143-155). Cambridge US: MIT Press.

Sherratt, S. 1999. *E pur si muove*: Pots, Markets and Values in the Second Millennium Mediterranean. In J. P. Crielaard, V. Stissi, and G. J. van Wijngaarden (eds.). *The Complex Past of Pottery: Production, Circulation and Consumption of Mycenaean and Greek Pottery [Sixteenth to Early Fifth Centuries B.C.]*. 163-211.

Sillar, B., & Tite, M. S. 2000. The Challenge of 'Technological Choices' for Materials Science Approaches in Archaeology. *Archaeometry*, *42*(1), 2-20.

Skourtanioti, E., Ringbauer, H., Gnecchi Ruscone, G. A., Bianco, R. A., Burri, M., Freund, C., ... & Stockhammer, P. W. 2023. Ancient DNA reveals admixture history and endogamy in the prehistoric Aegean. *Nature Ecology & Evolution*, 1-14.

Smith, J. M., & Haigh, J. 1974. The hitch-hiking effect of a favourable gene. *Genetics Research*, *23*(1), 23-35.

Sober, E. 2015. *Ockham's razors*. Cambridge University Press.

Spencer, C., & Bevan, A. 2018. Settlement location models, archaeological survey data and social change in Bronze Age Crete. *Journal of Anthropological Archaeology*, 52, 71-86.

Spencer, L. C. 2007. *Pottery technology and socio-economic diversity on the Early Helladic III to Middle Helladic II Greek mainland* (Doctoral dissertation, University of London).

Sperber, D. 1985. Anthropology and psychology: towards an epidemiology of representations. *Man* (NS)20, 73–89.

Steele, J. 2009. Innovation diffusion and travelling waves. In Shennan, S.J. (ed.). *Pattern and Process in Cultural Evolution*. (pp. 163-174). Berkeley: University of California Press.

Stellatou, A. I. 2007. *Final Neolithic and early Bronze Age settlement in the southern Aegean: A comparative spatial analysis of three regional surveys*. University of London, University College London (United Kingdom).

Stocker, S. 2003. Pylos Regional Archaeological Project, Part V: Deriziotis Aloni: A Small Bronze Age Site in Messenia. *Hesperia: The Journal of the American School of Classical Studies at Athens*, 72(4), 341-404. doi:10.2307/3182011.

Summers, K. 2005. The evolutionary ecology of despotism. *Evolution and Human Behavior*, *26*(1), 106–135.

Thér, R. 2014. Identification of pottery firing structures using the thermal characteristics of firing. *Archaeometry*, *56*, 78-99.

Tite, M. S., Kilikoglou, V., & Vekinis, G. 2001. Strength, toughness and thermal shock resistance of ancient ceramics, and their influence on technological choice. *Archaeometry*, *43*(3), 301-324.

Tobler, W. R. 1970. A computer movie simulating urban growth in the Detroit region. *Economic geography*, *46*(sup1), 234-240.

Triantaphyllou, S., Richards, M. P., Zerner, C., & Voutsaki, S. 2008. Isotopic dietary reconstruction of humans from Middle Bronze age Lerna, Argolid, Greece. *Journal of Archaeological Science*, *35*(11), 3028-3034.

Türkteki, M. 2014. Early Bronze Age pottery manufacture in Western Anatolia: Identifying Hybrid Technologies through X-ray Analysis. *Anatoliaca XL*, 93-109.

Van Andel, T. H., & Runnels, C. N. 1988. An essay on the 'emergence of civilization' in the Aegean world. *Antiquity*, *62*(235), 234-247.

Van der Leeuw, S. 2020. An Illustration of the Invention Process and Its Implications for Societal Information Processing. In *Social Sustainability, Past and Future: Undoing Unintended Consequences for the Earth's Survival* (New Directions in Sustainability and Society, pp. 237-262). Cambridge: Cambridge University Press.

Van de Moortel, A. Zachou, E. & Rutter, J. 2018. Mitrou and Proskynas in Prehistoric Lokris. Chronology, Interconnections and Society, in *PYRRA: Dakoronia Festschrift* (2018) 169-193. 169-193.

Vitale, S. 2016. Cultural Entanglements on Kos during the Late Bronze Age: A Comparative Analysis of 'Minoanisation'and 'Mycenaeanisation'at the 'Serraglio', Eleona, and Langada. In

E. Gorogianni, P. Pavúk, and L. Girella (eds.). *Beyond Thalassocracies: Understanding Processes of Minoanisation and Mycenaeanisation in the Aegean*. Oxford: Oxbow: 75-93.

Vitelli, K.D., 1984. Greek Neolithic pottery by experiment. In P.M. Rice (ed.). *Pots and Potters: Current approaches in ceramic archaeology*. Institute of Archaeology Monograph 24, Los Angeles, pp. 113-131.

Vitelli, K. D., 1994. Experimental approaches to Thessalian Neolithic ceramics: gray ware and color. *La Thessalie: quinze années de recherches archéologiques, 1975-1990 Bilans et perspectives*. Athens (1994

Walsh, M. J., Riede, F., & O'Neill, S. 2019. Cultural Transmission and Innovation in Archaeology. In A. M. Prentiss (ed.). *Handbook of Evolutionary Research in Archaeology* (pp. 49-70). Springer, Cham.

Warner, J.L. 1994. *Elmalı-Karataş II, The Early Bronze Age Village of Karataş.* Bryn Mawr: Bryn Mawr College.

Watrous, L. V., Haggis, D., Nowicki, K., Vogeikoff-Brogan, N., & Schultz, M. 2012. *Archaeological Survey of the Gournia Landscape: A Regional History of the Mirabello Bay, Crete, in Antiquity* (Vol. 37). INSTAP Academic Press (Institute for Aegean Prehistory).

Weiberg, E., Bevan, A., Kouli, K., Katsianis, M., Woodbridge, J., Bonnier, A., ... & Shennan, S. 2019. Long-term trends of land use and demography in Greece: A comparative study. *The Holocene*, *29*(5), 742-760.

Whitelaw, T. M. 2000. Settlement Instability and Landscape Degradation in the Southern Aegean in the Third Millennium. In P. Halstead, and C. Frederic (eds.). *Landscape and Land Use in Postglacial Greece* (Sheffield, Sheffield Studies in Aegean Archaeology 3), 135–61.

Whitelaw, T. 2001. From Sites to Communities: Defining the Human Dimensions of Minoan Urbanisation, in K. Branigan (ed.) *Urbanism in the Aegean Bronze Age*. Sheffield: Sheffield Academic Press, 15-37.

Whitelaw, T. 2017. The development and character of urban communities in prehistoric Crete in their regional context. *Minoan Architecture and Urbanism: New Perspectives on an Ancient Built Environment*, 114-180.

Whitelaw, T., Vasilakis, A., & Bredaki, M. (2018). Prehistoric Knossos: Tracing its long-term history through its surface record. In *Proceedings of the 11th International Cretological Congress* (Vol. 11, pp. 293-307). Historical and Folklore Society of Rethymnon.

Wiener, M. H. 2014. The interaction of climate change and agency in the collapse of civilizations ca. 2300–2000 BC. *Radiocarbon*, *56*(4), S1-S16.

Wimsatt, W.C. & Griesemer, J. 2007. “Reproducing entrenchments to scaffold culture: The central role of development in cultural evolution”. In R. Sansom, and R. Brandon (eds). *Integrating Evolution and Development: From Theory to Practice*. Cambridge, MA: The MIT Press, 227–323.

Wright, S. 1943. Isolation by distance. *Genetics*, *28*(2), 114.

Yeh, D. J., Fogarty, L., & Kandler, A. 2019. Cultural linkage: the influence of package transmission on cultural dynamics. *Proceedings of the Royal Society B*, *286*(1916), 20191951.

Yubero-Gómez, M., Rubio-Campillo, X., & López-Cachero, J. 2016. The study of spatiotemporal patterns integrating temporal uncertainty in late prehistoric settlements in northeastern Spain. *Archaeological and Anthropological Sciences*, *8*(3), 477-490.

Zachou, E. 2004. Rachi Panagias. In E. Alram-Stern (ed.). *Die ägäische Frühzeit: Die Frühbronzezeit in Griechenland* 738. Vienna: Verlag der Österreichischen Akademieder Wissenschaften.

Appendix 1: Pottery datasets from site catalogues

The pottery in the site catalogue from Quartier Nu, Malia provided by Schoep & Knappett (2003) were tabulated in the table below. Names for shapes are reported in English and were translated from French.

Shape	Function	Wheelmade	Wheel Coiled	Coiled	Fabric	Complete	Provenance
three-footed bowl	serving	0	0	1	red semi-coarse	1/3	local
jug with ovoid beak	serving	1	0	0	fine well fired	only base	local
jug with erect beak	serving	0	0	1	red semi-coarse	1	local
tripod globular jug	serving	0	1	0	red semi-coarse	3/4	local
pitcher, with 2 horizontal handles	serving	0	0	1	red coarse	4 fragments	local
Messara-type amphora	transport	0	0	1	light semi-coarse	fragments	Messara
straight mug, Scottish style	serving	1	0	0	light fine	1/8	local
jug with bridge spout, carinated and grooved	serving	1	0	0	light fine	lip fragment	local
tapered cup	serving	1	0	0	light fine	fragment	local
pitcher with three-lobed mouth	serving	1	0	0	red semi-coarse	1	local
low frusto-conical basin, handles down	storage	0	0	1	red coarse	1	local
small amphora with oval neck	transport	0	1	0	red semi-coarse	1	local
amphora with oval neck	transport	0	0	1	red semi-coarse	1	local
tripod tray	NA	NA	NA	NA	red-grey coarse	1/2	local
pitcher with three-lobed mouth	serving	1	0	0	red semi-coarse	1	local

small pitcher with three-lobed mouth	serving	1	0	0	red semi-coarse	without lip	local
ceramic core of a stuccoed basin	NA	NA	NA	NA	red coarse	fragment	local
ceramic core of a stuccoed basin	NA	NA	NA	NA	red coarse	lip and foot fragment	local
shallow dish, flared edge	serving	1	0	0	red coarse	3/4	local
lid, handle and rim	NA	NA	NA	NA	red coarse	1	local
ovoid / globular jar, with two horizontal handles	serving	0	0	1	red semi-coarse	1	local
lid, handle and rim	NA	1	0	0	red coarse	1	local
amphora with oval neck	transport	0	0	1	red semi-coarse	lip	local
honey jar?	storage	1	0	0	red semi-coarse	fragments	local
lid, handle and rim	NA	0	0	1	red semi-coarse	1	local
lid, handle and rim	NA	1	0	0	red coarse	1	local
ovoid jar with 2 handles	serving	0	1	0	red-orange semi-coarse	2/3	local
ovoid jar with 2 handles	serving	0	1	0	red semi-coarse	1	local
tapered cup	serving	1	0	0	red semi-fine	3/4	local
jug with bridge spout, carinated and grooved	serving	1	0	0	light fine	3/4	local

pitcher, with two vertical handles	serving	0	0	1	red coarse	3/4	local
hemispherical cup	serving	1	0	0	light fine	11 fragments	local
pitcher, with two vertical handles	serving	1	0	0	red fine	1	local
pithos	storage	NA	NA	NA	light coarse	lip fragment	local
pithos	storage	0	0	1	red-orange coarse	lip and fragments	local
tapered cup	serving	1	0	0	light fine	2/3	local
crater jar	NA	NA	NA	NA	dark grey coarse	half lip	local
pitcher	serving	0	0	1	red semi-coarse	1/4	local
tapered cup	serving	1	0	0	red semi-fine	3/4	local
tapered cup	serving	1	0	0	red semi-fine	1/2	local
carinated cup	serving	1	0	0	grey-red fine	1	local
carinated cup	serving	1	0	0	light fine	1/2	local
carinated cup	serving	1	0	0	red fine	1/2	local
tapered goblet	serving	1	0	0	orange-grey semi-fine	1	local
tapered goblet	serving	1	0	0	light fine	1/2	local

cistern	storage	0	1	0	red coarse	1/2	local
tray	serving	NA	NA	NA	red coarse	1/6	local
amphora with oval neck	transport	0	1	0	red semi-coarse	3 tessons	local
shallow tapered basin	storage	0	0	1	red coarse	2/3	local
pitcher	serving	0	1	0	red-orange semi-coarse	1/2	local
pitcher, 2 horizontal handles	serving	0	0	1	red coarse	1 fragment	local
Jar	serving	NA	NA	NA	red coarse	base	local
shallow tapered basin	storage	0	0	1	red semi-coarse	1/2	local
plate	serving	1	0	0	orange semi-coarse	1	local
pitcher	serving	0	1	0	red semi-coarse	1/3	local
tapered cup	serving	1	0	0	red semi-fine	2/3	local
tapered cup	serving	1	0	0	grey-red fine	2/3	local
tapered cup	serving	1	0	0	red semi-fine	2/3	local
?	NA	NA	NA	NA	red coarse	lip fragment	local
pedestal, from pedestaled lamp	lighting	0	0	1	red-brown coarse	1	local

?	NA	NA	NA	NA	red coarse grey in centre	fragment	local

Table 9: Pottery from Quartier Nu, Malia. Data from Schoep & Knappett (2003)

Appendix 2: Model verification using a toy example of a simple network with two nodes, k = 0, a = 1

The easiest way to verify that the code for the simulation produces outputs that are compatible with the mathematical formula is to calculate the result of the formula manually for a small number of timesteps and the smallest number of nodes possible (2) and then run the simulation with the same population size N for an equal number of timesteps and the same parameter values and compare the mean of both nodes for each timestep. The simplest model arises from setting $k = 0$ and $a = 1$, where $k = 0$ results in all edge weights being equal to 1 -since e raised to the 0th power equals 1- and $a = 1$, which ensures that the trait proportion value at each community/node does not increase or decrease on its own. Finally, both for the manual calculation and for the simulation, one of the nodes, node (n1) is assigned a trait proportion value of 0.1 while the other node (n2) is equal to 0.

For example, this is what the calculation for the values of n1 and n2 at Timestep 2 (t2) looks like:

n1t2 = 0.1 * 1 + (0 * 1 + 0 * 0.1) = 0.1

n2t2 = 0 * 1 + (0.1 * 1 + 0 * 0) = 0.1

and for the next timestep t3:

n1t3 = 0.1 * 1 + (0.1 * 1 + 0 * 0.1) = 0.2

n2t3 = 0.1 * 1 + (0.1 * 1 + 0 * 0.1) = 0.2

for timestep t4:

n1t4 = 0.2*1 + (0.2*1 + 0 * 0.2) = 0.4

n2t4 = 0.2*1 + (0.2*1 + 0 * 0.2) = 0.4

and finally, for timestep t5:

n1t5 = 0.4*1 + (0.4*1 + 0 * 0.4) = 0.8

n2t5 = 0.4*1 + (0.4*1 + 0 * 0.4) = 0.8

so that the mean of all communities/nodes for each timestep is:

t1 = 0.05, t2 = 0.1, t3 = 0.2, t4 = 0.4, t5 = 0.8

After running the simulation with the same parameter values, it is possible to confirm that the simulation output is the same as the result of the manual calculations, and that the implementation of the model in code follows the mathematical formula defined previously (Figure 1). Obviously, the more complex the simulation models become, the more probable it is that coding errors (sometimes called “bugs”) will occur, and various stages of testing were required during the development of the more complex models in subsequent chapters to ensure that such errors were identified and corrected. This is only an illustrative example of the advantages conferred by using toy models in the development stage of a simulation model: by reducing the model to a small number of elements in a simple environment, it is possible to manually calculate what the simulation outputs should be. This would have been impractical or even impossible to do in cases with thousands or millions of communities/nodes with different edge weights.

Ultimately, verification using such toy models renders the simulation process and the entire modelling exercise less opaque, allowing one to place their trust in the outputs of the simulation and the understanding that they can provide in explaining the phenomenon of interest.

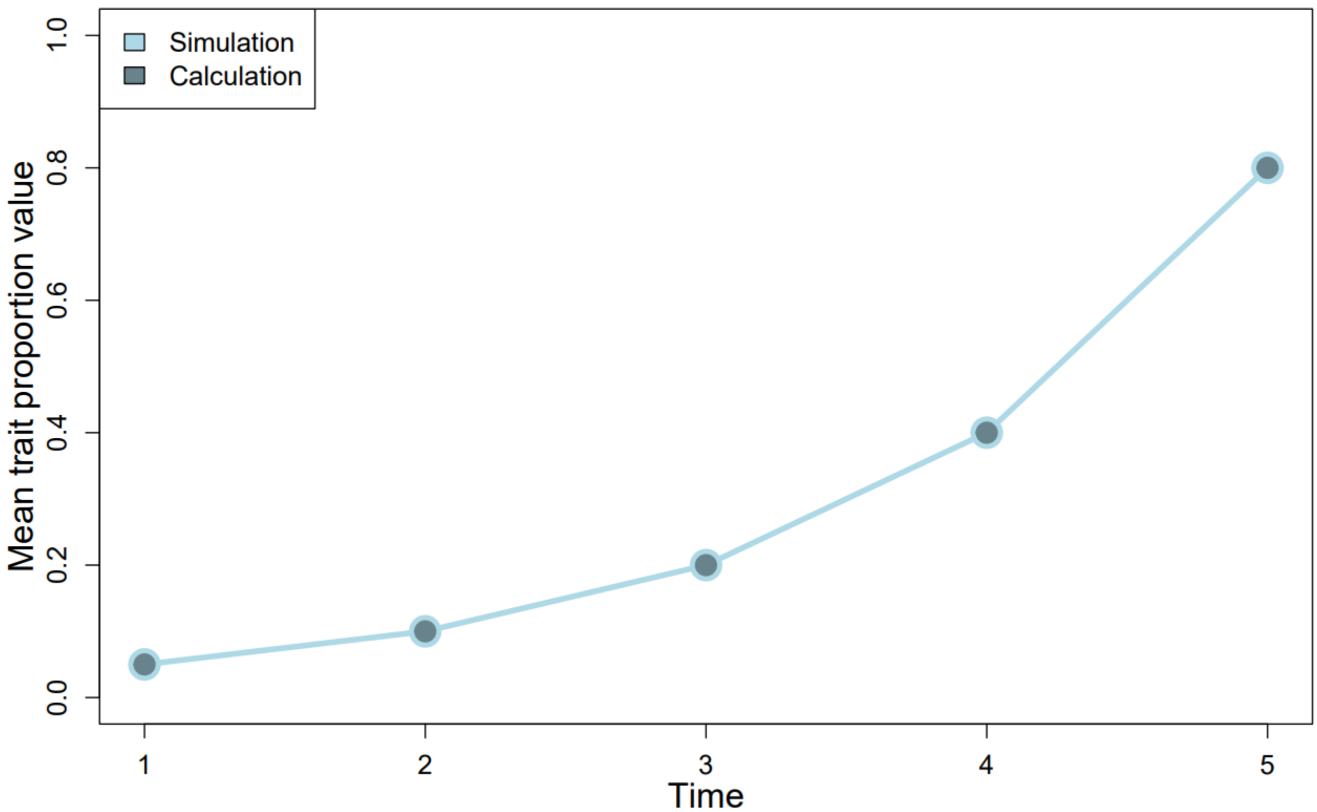

Figure 78. Mean trait proportion value per timestep as indicated by manual calculation (in grey) and by running the simulation model (in light blue). Plotting them on top of each other suggests that the simulation model produces the same results as the mathematical formula.

Appendix 3: Experiment using only rank weights

As discussed in the previous section, adding a new set of weights descriptive of neighbour rank to the model does not produce radically different trait adoption trajectories, but it does lead to the trait reaching certain regions (like Crete) relatively faster compared to other regions. However, combining both sets of weights makes assessing the effect of the rank weights on overall adoption trajectories difficult. This section includes simulations for a model that only uses rank weights, reformulated as such:

The trait value for each node is calculated for each timestep:

$$c_{it} = (c_{it-1})a + \sum_{i=j}^{N} c_{it-1}\, r_{i,j}$$

$$r_{i,j} = \begin{cases} \dfrac{1}{n_{i,j}}, i \neq j \\ 0, i = j \end{cases}$$

The results in Figure 79 suggest that, as would be expected, trait adoption increases exponentially, and the system reaches the 100 % equilibrium in 7 timesteps. This is due to the fact that this model essentially disregards the magnitude of geographic distances, taking only their ranking into account. As shown in Figure 80, which shows the weight values for node pairs including the arbitrarily chosen nodes 1 and 50, for each node there is one nearest node where their inter-node weight is equal to 1, a very high value that ensures nodes increase their trait value by their neighbours' value (plus the contribution of other nodes) each timestep. This could mean, for example, that in the cases where a pair of nodes has a roughly equal trait proportion value, nodes double their trait proportion value at each timestep.

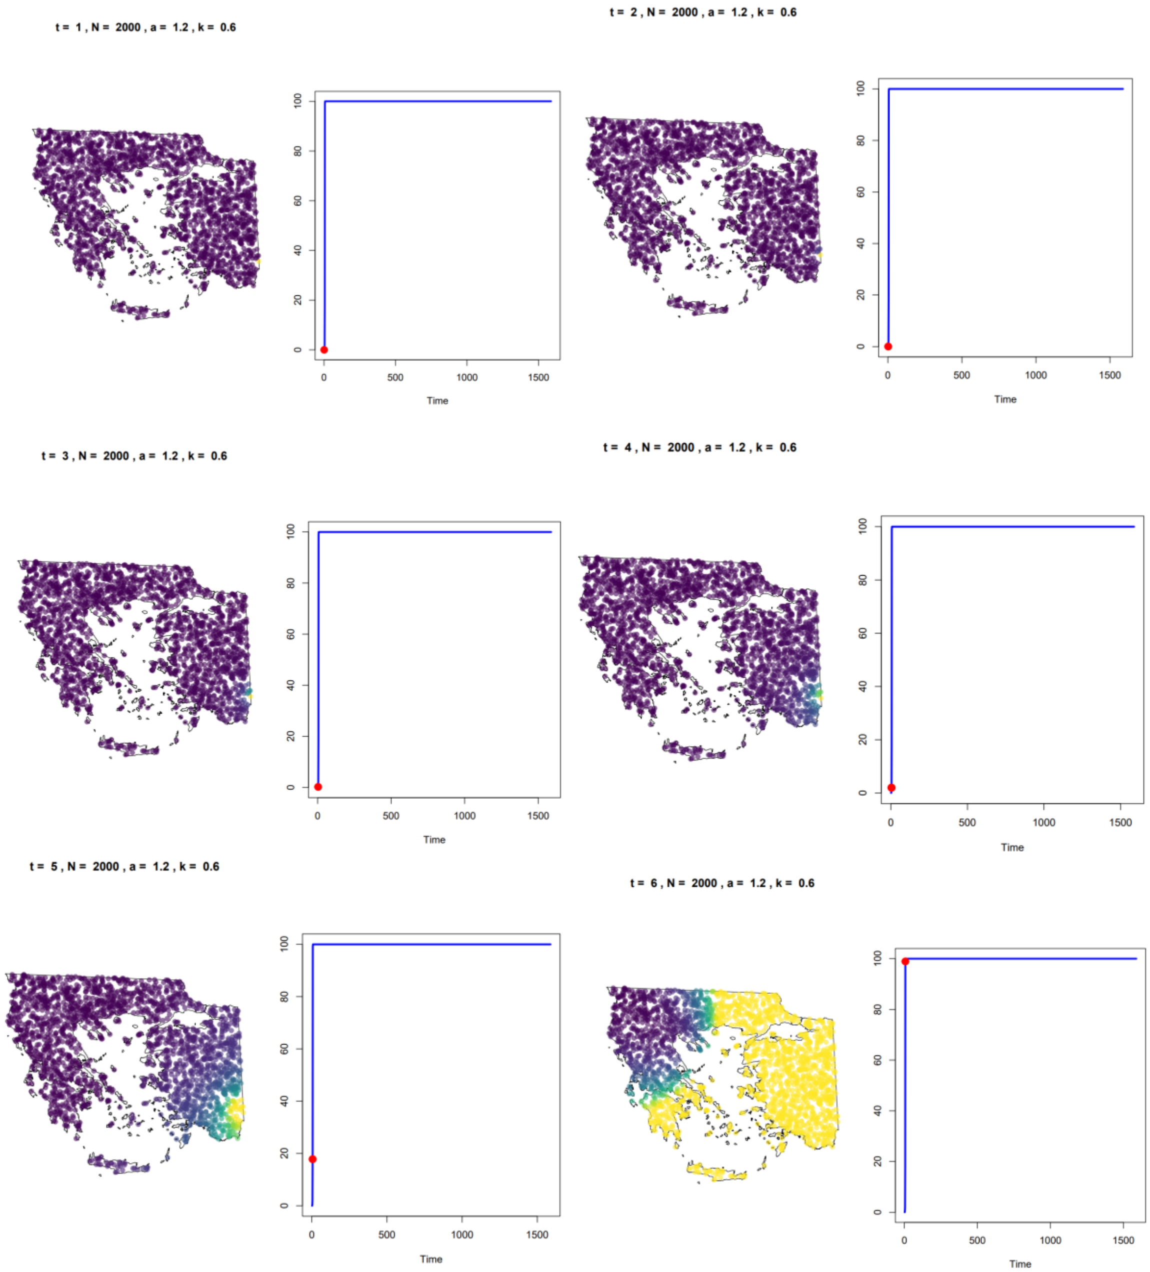

Figure 79: Simulation including a globally defined local adoption rate parameter a and a new set of weights determined by the reciprocal of the Euclidean distance ranking of nodes. Trait adoption at timestep 1, 2, 3, 4, 5, and 6. k = 0.6, a = 1.2, N = 2000. The easternmost node is the origin of the trait, with a starting value equal to 0.1.

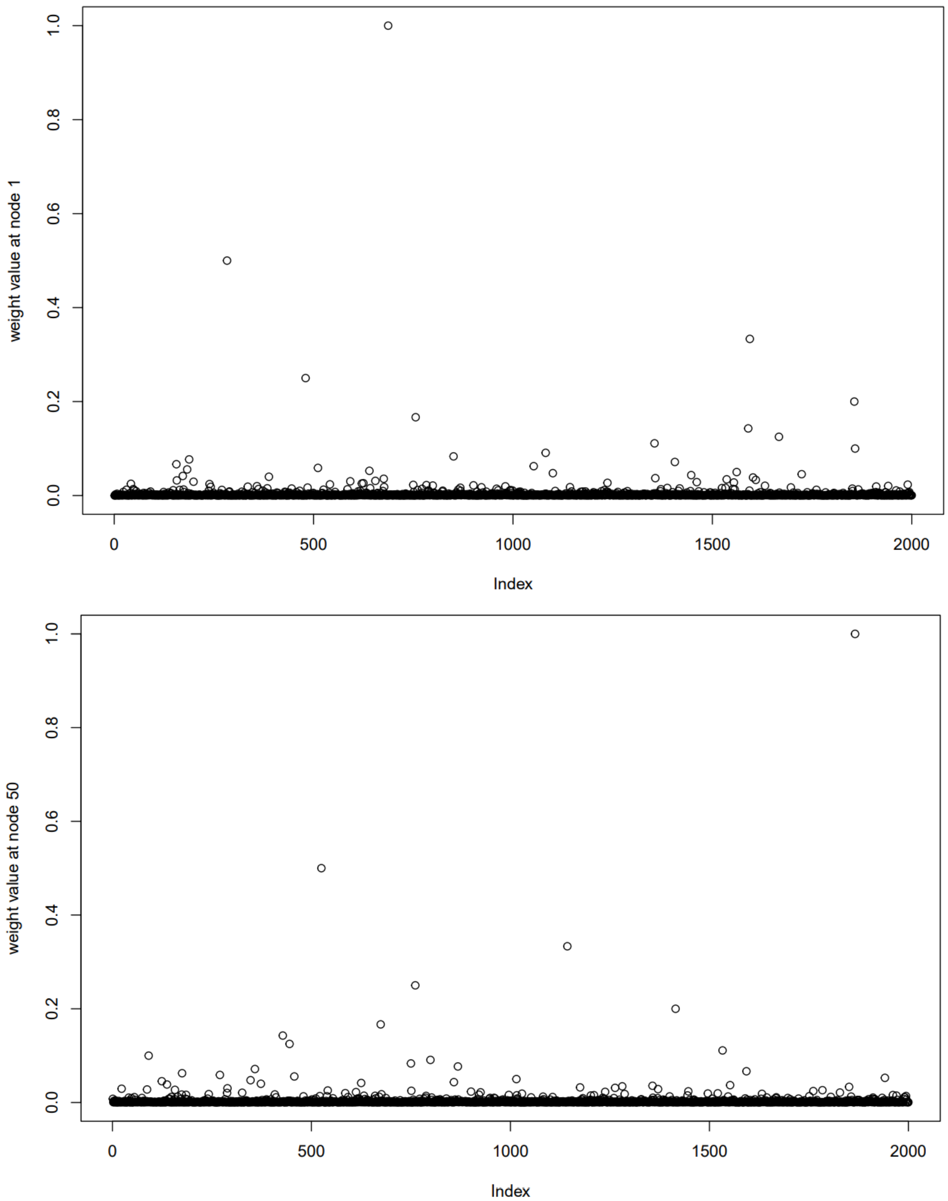

Figure 80: Weight values for all node pairs involving node 1 (left panel), and node 50 (right panel).