Content-Dependent Biases in Social Learning Strategies: a multiscale approach

Simon Carrignon

TESI DOCTORAL UPF / year 2019

THESIS SUPERVISOR
Sergi Valverde & Xavier Rubio-Campillo
Department of Experimental and Health Science

Universitat Pompeu Fabra
Barcelona

“Commune de caractère, équine, où se chevauchent les collines.”. BL

Acknowledgements

If I had to thank every people that made this work possible, I would have to come back so far in time that even the names of those people have fade out of my memory. I'll then keep it to the people who helped me during the past four years and a half. I cannot guarantee the language will stay the same all over these few lines, nor that it will change in a planned way.

First of all I want to thank my supervisors, Sergi Valverde and Xavier Rubio. Xavi who believed that I was the good candidate to do this and Sergi for accepting to follow us in this curious adventure that featured a Supercomputer, a Biomedicine school, the Roman Empire and a french student with a master in philosophy of science. I want to thank all my colleagues and friends at the BSC, Matías, Dani G., Paula, Mariña, Guido, Dani P., Miguel, Jordi, Hadrian,…The structures where science is done often have their own fate that pass over the individuals working within it without much care. Hopefully humans are still human, and you can count on them to care and share. I cannot forget the people that saved me so much time when I needed to travel (and they know how much I traveled): Luciana, Aina and Lisandra. People from the EPNET project, Jean-Marc, Luce, Ignacio and Iza. Also Tom, who brought me the case study I was missing and trusted me to explore it. I want to thank all students and postdocs I met in the Santa Fe Summer School. All those who stayed up all nights in Ariel's building to change the world again and again. Above all, I want to thank Jelena, Pinar, Philip, Devrim, the "boss" Gaetano, and Salva y Aina from the Barcelona's Complex System Lab. I have to specially thank Aina with whom I could share many doubts and difficulties: it often takes someone lost in the same dark corner of knowledge to be able to follow ahead and believe something can be done. One thing I learnt during those 4 years is that Academia is a place where you are supposed to know everything, every time, where you have to be the best and can't be wrong. Not sure if that's very good but what I'm sure is that I would have never been able to bear that without peers with whom I could be myself, mediocre, bad, doubting, failing over and over, unsure and to whom you can say "I haven't understand anything" or "This time I think I really fu…up". With all the people named before I was able to be that, and thanks to them I could make it through.

I have to thank all the friends I met in Barcelona, who accompanied me since I passed the door of the 40 carrer de la Torre Vélez. All the people from the "extend Guinardo's Family" and who continue to come to our terrace in La Sagrera after we moved: Adri & Alba (who will be cosmically linked for ever with this work, as they may have been the tiny touch that made the difference between staying in Paris or move), Arnaud mon gabache sûr, Jesus de Huelva (la version mucho mas simpatica que la 'oficial'), Juan-Pedro/Jota P/Jean-Pierre, Maria y Junao los "Bermianis", Ignacio otra vez, porque se lo merece, Lucas, one of the greatest mind and soul I had to meet and whose vitality should be an example for everyone. Mis dos hermanos de aqui, Jorge y Marcos, who lived and supported me since the beginning and until the very last and heavy months, where they couldn't enter the flat without seeing my dark hollow in the corner of the flat, every days and at every hours of the hottest days of Barcelona's summer. Gracias a estos dos he seguido apprendiando como vivir, disfrutar, tener nuevos sueños y "ademas", como lo dice el sors "trat[ar] de ser una buena persona". I thank also all my friends in Montreal (or who moved nearby) who welcome me each time I stop my flights in Quebec to say hello to Darcy. A special thank to Damien, for his help and those lengthy personalised podcasts that helped me to think, to Beuf la gouttiere, for the music and technical help, and of course the MHCC. I thank the 4dragz (both versions, David included), for their continuous help and to keep my intellectual standards high enough to be able to stay up to date with how the world goes (and with the Internet's finest content). J'ai la douce impression que ce n'est que le début. À tous mes jarrois, le brez le mlet, le steeve et tous les autres, on est plus trop au pays mais on continue de representer même à l'international. À Céline "Madame le Maire", toujours là, de la rue Benoît Duperrier à la Carrer d'Hondures en passant par la rue Samarcq.

The last thanks go obviously to my family, who have been always here, after 14 years at the University. Ce matin je mange le pain de Saint-Michel en Beaumont en écrivant ces dernières lignes et une semaine après il est aussi bon qu'au sortir du four (pas très sur de la portée profonde et métaphorique de cette phrase, elle devait simplement illustrer à quel point j'ai pu compter sur eux jusqu'aux derniers instants). Et ma deuxieme famille, Big-B, Bart, Nico, Pierre, Beufa, Groumf, ces 4 dernières années furent probablement celles durant lesquelles nous nous somme le moins vu en presque 20 ans, tout n'a pas été simple mais le moteur est toujours le même. En regardant les dernières photos on dirait presque qu'on est adulte ça m'a bien fait rire. En attendant j'ai reussi à finir mes études avant que l'un d'entre nous n'ai des gamins. Finally I thank Marìa who supported me every days, and stayed when most would have left.

Abstract

The content of what we learn shapes the evolution of human culture and society. In this thesis, we have quantified the influence of content-dependent biases in social learning strategies. Our theoretical framework combines agent-based models and Bayesian inference to measure content-dependent biases in large-scale social learning strategies. Our first empirical study measures the impact of social transmission biases in Twitter. The novelty of the second study is two-fold: ours is one of the rare uses of computational modelling in historical Roman Studies and one of the first tests of the impact of success bias across large spatial and temporal scales.

Resumen

El contenido de lo que aprendemos socialmente moldea la evolución de la cultura humana. En esta tesis hemos cuantificado la influencia de diferentes estrategias de aprendizaje social analizando procesos culturales en diferentes escalas. Se propone un marco teórico que combina los modelos basados en agentes y la inferencia bayesiana para detectar sesgos dependientes de contenido en la evolución cultural. El análisis se realizará sobre tres escenarios diferentes: un escenario teórico, que revela el potencial del sesgo de éxito, y dos casos de estudio empirico que representan distintas escalas espacio-temporales. En el primer caso, se estudia la influencia de transmisión social dependiendo del contenido de diferentes clases de noticias online, mientras que en el segundo se analiza la influencia de los sesgos de éxito en los cambios de distribución de vajillas en el este del Imperio Romano.

Preface

Non random mutation:

In 2010, while doing a master in Paris, I get back in touch with an old friend of mine: Jérémy Gardent. Five years before we lost sight of each other after we left the "Lycée des Portes de l'Oisan" where we first met in 2002. He went to the "ENS Lyon" and I started my undergrad at Grenoble's university (UJF at this time). When meeting again we quickly realize that during all those years, although we followed different path, our directions were the same. Within our respective schools, we had grown similar interests for similar topics. And so started days and nights of discussions about Science, Evolution and Life. Both of us had read Dawkins's book and its chapter about memes had a great impact on our imagination ; after some times a specific set of questions started to come back recurrently. What are the essential properties that allow entities to "evolve"? Do ideas, or cloud, or mountains, evolve, as living beings do? And how could we explore and test that? How can we be sure that this is evolving and this not? These questions didn't exactly pop up from nowhere: Jérémy was studying evolutionary biology and had various occasions to see how computational methods can link different empirical observations under a unified evolutionary history. On my side, while studying "Cognitive Sciences" I was impressed by how computer models were able to mimic life, moreover the beauty and elegance of the solution computers find when algorithm were built to replicate evolution. Random ideas on top of less random ideas, these discussions ended up leading us much farther than what we expected.

Occasionally exchanging random ideas was not enough for us: in order to follow our discussions in a more formal and systematic way, we decided, in March 2011, to apply for a funding to create a "Junior Lab" at Jérémy's school. After a first rejection, the project was accepted one year later. LaRéMI[1], a Junior Lab that aimed at studying music[2] using computers and evolutionary theories, was born. Young and naive, still believing our freshly acquired skills will solve every problems nature will put in our road, we were pretty sure to unveil the very essence of the evolution of Music.

Thanks to the few thousands euros the ENS gave to the junior lab, we organized several meetings to discuss *how and why* we expected to find this hidden essence of the evolution of Music. After heated debates, pizzas, homemade sandwiches and walk in the countryside, we ended up almost agreeing on this *how and why* (briefly summarized: we agreed on combining experimental work and computational models). But the fundings of the junior lab were (very) limited, and we needed to justify the need for this money by organising something institutions would find more "significant" than meetings and sandwiches. In 2013, with what was left on the budget, we organized an "international" conference, to present our findings, that didn't go much farther than "*how and why* we expected to find something". Mathieu Charbonneau and Olivier Morin, who I cannot thanks enough for that, kindly accepted to come to Lyon to discuss with us and to present their own work. As kindly as they came they tried (rightly) to temper our unbounded optimism, while at the same time they vigorously encouraged us to pursue on this track. At the end of the conferences we haven't exactly found the essence of the Evolution of Music, but we saw that we were not alone and more important, that the skillful academics who took the same road understood why we wanted to do what we were doing.

To perfect this unexpected journey in what we discovered on the way to be "Cultural Evolution", abstracts of our *how and why* were accepted in a conference on "Patterns of Evolution" in Lisbon. Then, a few month after the conference we organized in Lyon, Jérémy and I were on our way to our first *real* international conference, surrounded by hundreds of renown academics. Without knowing it at this time, we met there a great percentage of the scholars I will cite in this work. Again, we presented on a poster our "how and why", with reserve this time, and again people were kind and polite. But we get to visit Lisbon, paid by *our own* lab, and this is worth a thousand posters.

[1] This project succeed thanks to Pr Daniel Hromada, Dr Blaise Tymen, Christophe Chavatte and Manon Février, (though none was doctor at this time), who joined us in this adventure.

[2] We choosed to focus on music for various practical reason, but mostly because both J'er'emy and I were deeply involved in non-scientific activities revolving around music.

Two years later, the time for Jérémy to finish a Master degree in Musicology, and me one in History and Philosophy of Science, and both of us were starting PhDs to try to use quantitative methods to study cultural changes.

The tale of a third generation

Something was clear to me since the very beginning of this PhD: I did not want to decline "how and why this ingenious and novel way to articulate theory and data can help to understand social changes" in all its possible flavors. I wanted to *do* it. Implement the *how* and demonstrate the *why*.

This seemed even more than possible as since the first version of our Junior Lab project in 2011 and our enrollment in grad school in 2014-2015, the landscape of Cultural Evolution changed. In 2011 the main books cited to introduce the topic, were still the "First" books, written by "enlightened pioneers" 30 years ago (Boyd and Richerson, 1988; Cavalli-Sforza and Feldman, 1981; Durham, 1992). Those enlightened pioneers, when they started to study Culture using mathematical models, had to rely only on their personal curiosity and motivation to navigate those unknown seas[3]. Most of them were researchers already working for years in well defined disciplines, who decided to use their knowledge to study something else, something different.

Later on at the end of the 90s, early 00s, a second generation of younger researchers (Fiona Jordan, Alex Mesoudi, Claudio Tennie, Christine Caldwell, Alex Bentley, Rachel & Jeremy Kendal, to cite some of them), followed their path. This second generation still get their PhD in specific and well defined fields, such as Anthropology, Psychology, Biology, Ecology and so on. But as soon as they get finished, they followed the path drawn by the early pioneers and built there early career working on the questions they raised. The number of papers started to grow quicker and quicker, more and more book where written. Some of them were granted fundings and able to hire like minded postdocs, some could create their own research groups. This quick and impressive growth has been recently documented in a paper by Youngblood and Lahti (2018).

Now a great number of new books, centered on Cultural Evolution, can be cited (Mesoudi, 2011; Morin, 2015; Bentley and O'Brien, 2017; Henrich, 2015, to cite only a few of them). The Cultural Evolution Society has been created in

[3]This is obviously not exactly true but we decided to leave the history of cultural evolution and modeling for another work

2016 and organized its second conference in 2018, gathering together hundreds of researchers from all around the world. Cultural Evolution *is a thing*.

Alongside, a new, "third", generation of PhD students and young researchers is emerging. At the difference of the previous ones, these students *did* there PhDs within groups explicitly working on Cultural Evolution, composed by researchers focusing on such questions. To give an idea of the emergence of this third generation, I put here a list of these students, who I met or heard about during conferences, meeting, etc., and who explicitly mention that their thesis is a thesis about "Cultural Evolution", or working under the direction of researcher from the second generation[4]:

- Mathieu Charbonneau: *L'analogie de l'hérédité culturelle : fondements conceptuels de la théorie de la double hérédité* (2013)
- Maxime Derex: *Les mécanismes de l'évolution culturelle cumulative* (2013)
- Marius Kempe: *Experimental and theoretical models of cultural evolution* (2014)
- Sally Street: *Phylogenetic comparative investigations of sexual selection and cognitive evolution in primates* (2014)
- Marco Smolla: *Environmental effects on social learning and its feedback on individual and group level interactions* (2017)
- Charlotte O. Brand: *Sex differences in social learnning: exploring the links with risk aversion and confidence* (2017)
- Eva M. Reindl: *On the developmental origins of human material culture* (2017)
- Jérémy Gardent: *Mesurer les musiques pour parler du passé : la comparaison des musiques du Gabon comme source d'informations historiques* (2017)
- Rachel A. Harrison: *Experimental studies of behavioural flexibility and cultural transmission in chimpanzees and children* (2018)
- Oleg Sobchuk: *Charting Artistic Evolution: An Essay in Theory* (2018)
- Alice Williams: *Modelling the evolution of socio-political complexity* (2019)
- Maria Coto-Sarmiento: *Cuantificando el cambio cultural* (2019)

[4] It should be noted that this "story" about Cultural Evolution has no historical value. It is well known in History and Cultural Evolution that there is no such clear divisions between generations and most of the researchers cited and the numerous one who are missing don't fall right in the middle of one of the three generation proposed, but often fall somewhere in between. This preface is only a subjective interpretation of what happened and a real historical work should be done (also, 'the tale of a third generation' sounded good).

And again, they are only those I personally met or heard about and remember. I guess this list could easily be doubled if made exhaustive.

Hence, it looked very much like if time had come for research in Cultural Evolution to be done as "Normal Science", the way Kuhn theorized it. Without the need of justifying *how and why* it was possible and interesting to do so. The same way a PhD student in Biomedicine doesn't have to justify why he will use some algorithm to sort out meaningful differences between two DNA strands and why these differences can help him to infer an evolutionary distance between two species.

Given all that and my "multi-disciplinary" background mixing Computer Science, Evolutionary Biology and Cognition, I felt I could do "Normal Cultural Evolution Science" and work on some concrete questions, and solve a given problem, without having to care to much about justifying what I was suppose to do, how and why. The path was wide open for the third generation to start its journey through Cultural Evolution. I hope this thesis will give a good overview of this travel from the point of you of one of those third generation student trying to bind together the theory and the empirical observations using computer models.

Swarm Intelligence

This travel was impossible to do alone and for the rest of this document I will use the first person of the plural to underline this collective aspect. Here I will briefly details who is hidden behind this "we", and list some of the outcome we get from every chapter:

Chapter 2: Jean-Marc Montanier wrote a very first draft of the Agent Based Model structure from which I developed the framework used in the Chapter 2. Xavier Rubio wrote the pandora simulator which is used to run and parallelize this framework. Ignacio Morer designed the networks used in Section 2.5.1. I designed and implemented all models at both cultural and economic levels. I run and analysed all simulations. Selected outcome of this chapter:

- Peer reviewed Proceedings:
  - Carrignon, S., Montanier, J.-M., and Rubio-Campillo, X. (2015). Modelling the Co-evolution of Trade and Culture in Past Societies. In *Proceedings*

of the 2015 Winter Simulation Conference, WSC '15, pages 3949–3960, Piscataway, NJ, USA. IEEE Press

- Talks:
  - Carrignon, S., Montanier, J.-M., Michaud, J., and Rubio-Campillo, X. (2016). Co-evolution of culture and trade : impact of cultural network topology on economic dynamics. In *44th Computer Applications and Quantitative Methods in Archaeology Conference (CAA 2016)*
- Posters:
  - Morer, I., Carrignon, S., and Rubio-Campillo, X. (2016). Influence of the topology of cultural networks on the equilibrium of an exchange-based economy. In *7th Workshop on Complex Networks (CompleNet 2016)*
  - Carrignon, S. and Rubio-Campillo, X. (2017). Impact of different social learning mechanisms on the emergence of a Walrasian Equilibrium. In *European Human Behavior Evolution Association conference (EHBEA)*
- Code:
  - https://github.com/simoncarrignon/2015-WSC-source

Chapter 3: Theoretical formulation of the content model has been proposed by Alex Bentley. Original implementation of the neutral model and the Top Alberto model has been done by Damian Ruck and Alex Bentley. The re-implementation and implementation of all other models has been done by myself, as well as the analysis and processing of the original data, all experimental design, implementation of the this experimental setup, simulation and analyse of the results (including the whole Approximate Bayesian Computation). Selected outcome of this chapter:

- Talks:
  - Carrignon, S., Bentley, R. A., Ruck, D., and Gilchrist, M. (2018a). How the intrinsic value of information can change the spread of news in social media . In *2nd Conference of the Cultural Evolution Society*. Arizona State University
  - Part of it: Carrignon, S. (2018). Agent Based Modeling and Bayes Inference to learn about the past: the need for High Performance Computing. In *Conference on Complex System*

- Peer revied journal
  - Carrignon, S., Bentley, R. A., and Ruck, D. (2019). Modelling rapid online cultural transmission: Evaluating neutral models on twitter data with approximate bayesian computation. *Palgrave Communication*
- Code:
  - github.com/simoncarrignon/spreadrt

Chapter 4: The case study, the data set and the archaeological context and intrepretation has been brought by Tom Brughmans. Pre-processing of the data has been done by Iza Romanowska. All experimental design, implementation, simulation and analyse of the results (including the whole Approximate Bayesian Computation) has been done by myself. Selected outcome of this chapter:

- Talks:
  - Carrignon, S., Romanowska, I., and Brughmans, T. (2018b). An Agent-Based Model of Trade in the Roman East (25 BC – AD 150). In *28th Theoretical Roman Archaeology Conference (TRAC)*
  - Part of it in: Carrignon, S. (2018). Agent Based Modeling and Bayes Inference to learn about the past: the need for High Performance Computing. In *Conference on Complex System*
- Peer reviewed Journal:
  - Part of it in: Brughmans, T., Hanson, J., Mandich, M., Romanowska, I., Rubio-Campillo, X., Carrignon, S., Collins-Elliott, S., Crawford, K., Daems, D., Fulminante, F., de Haas, T., Kelly, P., Moreno Escobar, M. d. C., Paliou, E., Prignano, L., and Ritondale, M. (2019). Formal Modelling Approaches to Complexity Science in Roman Studies: A Manifesto. *Theoretical Roman Archaeology Journal*, 1:4
- Chapter in book:
  - Carrignon, S., Brughmans, T., and Romanowska, I. (Forthcoming). Transmission of cultural and economic strategies in inter-regional tableware trade. In Brughmans, T. and Wilson, A., editors, *Simulating Roman Economies. Theories, Methods and Computational Models*. Oxford: Oxford University Press

- Code:
  - https://framagit.org/sc/icrates_abc
  - https://framagit.org/sc/abcpandora

To end this preface, it should be noted that the full thesis, as well as the papers, talks and posters listed above, are only the final steps of this collective effort; a handful of selected snapshot from a wider picture. This whole picture is a huge patchwork of thousand of points and lines, random and less random variations. The results of unplanned meetings, debates and discussions, at the corner of a coffee break, during a journal club, a summer school or at the gala dinner of a conference. All those moments that flourish, out of sight of the relentless eyes of the reviewers and far away from the greedy races for fundings, when science really happens.

Despite their importance, such moments can hardly be translated in a written report. I will sum up some significant ones where I learnt these skills that are not in manuals nor on stackoverflow website, although they are so important when you sail interdisciplinary seas such as Cultural Evolution.

Summer schools and workshops:

- Modeling Complex Systems in Archaeology, 2018 DySoC Critical Workshop, Knoxville, USA.
- New England Complex System Institute Winter School, 2018 NECSI Winter School, Cambridge, US.
- UrbNet, Aarhus University, 2017 UrbNet, Aarhus, Denmark.
- Santa Fe Institute Complex System Summer School, 2016 SFI CSSS, Santa Fe, US.
- Data And Cities As Complex Adaptive Systems, 1st DACAS International Workshop, Manchester, England.
- The Computational turn: Simulation in Science., Scientific World Conception Summer School 2016, Vienna, Austria

Organisation and Edition:

- Assistant Editor (together with Maria Coto-Sarmiento), Since November 2017, Special Issue: Evolution of Cultural Complexity, Adaptive Behavior, Excepted publication date: September 2019.
- Evolution of Cultural Complexity III, Thessaloniki, Greece, Together with Sergi Valverde: Organisation and chair of the satellite "Evolution of Cultural Complexity" at the 2018 Conference on Complex System. Invited speakers: Peter Turchin & Anne Kandler.
- Evolution of Cultural Complexity II, Cancun, Mexico. Together with Maria Coto-Sarmiento, Sergi Valverde, Iza Romanowska and Xavier Rubio-Campillo: Organisation and chair of the satellite "Evolution of Cultural Complexity" at the 2017 Conference on Complex System. Invited speakers: Sergi Valverde, Tom Froese & Alex Bentley.

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

INTRODUCTION

Culture shapes the way we speak, the way we think, what we like and what we don't. We learn culture through social interactions: interactions with our parents, our friends and colleagues; interactions with people we read on the internet, people we see on the television or in videos shared in our favorite online social media. Flowing through social interactions, culture can spread and be transmitted from generation to generation. By understanding how these social interactions act, what factors impact them and how, we can understand why some traditions last millenniums while other fade quickly, why some technological inventions spread like wildfire while other disappear after a few months. On the other side, by looking at the records of cultural change throughout human history, one can infer what shaped those changes and shed light on the nature and the structure of social interactions in past human societies. The aim of this thesis is to contribute to the understanding of factors that impact cultural transmission, how they change the way culture evolve and how we can detect them and their importance throughout History.

1.1 Cultural Evolution and Social Learning Strategies

To do so, this thesis follows a relatively new approach known as "Cultural Evolution" and popularized by a growing number of researchers (Boyd and Richerson, 1988; Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 2005; Mesoudi, 2011; Youngblood and Lahti, 2018). Cultural Evolution proposes to study cultural change through the lens of the evolutionary and population thinking that biologists developed since Darwin's "Origin of Species" (Darwin, 1859). Thanks to this quantitative approach, Cultural Evolution aims at understanding how local, individual mechanisms can lead to global, long term change at the level of societies (Mesoudi, 2015).

Cultural Transmission is as central to Cultural Evolution as heritable variation and differential reproduction are central to Evolutionary Theory. By understanding how Culture is transmitted from generation to generation and why some cultural variants spread more than others, it becomes possible to understand the forces that drive culture changes all over Human History (Bettinger and Eerkens, 1999; Bentley and Shennan, 2003; Eerkens and Lipo, 2007; Premo, 2014; Mesoudi and O'Brien, 2008).

Researchers have defined Social Learning Strategies (SLS, see Laland, 2004) as the set of different process behind Cultural Transmission. SLS reflect the fact that humans can be influenced by various aspects of their environment or internal attributes and thus adopt different *strategies* when learning and transmitting cultural traits. These different strategies, applied by different kind of people, will change the outcome of Cultural Transmission, favoring or preventing the spread of certain cultural variants. In other words, this will *bias* Cultural Transmission toward or away from specific subset of cultural variants, generating what's is also called *Transmission Biases*[1]. Kendal et al. (2018) have classified SLS according to "what, who and when" is influencing cultural transmission, i.e. what is *the source* of the bias. In this thesis, we will focus on one particular source of bias, the *content* (or value) of cultural traits.

Biases impacted by content are crucial for the evolution of social learning in human. It has been shown that random social learning doesn't provide selective advantages (Rogers, 1988). Social learning becomes adaptive when there are ways for individuals to know the expected increase of fitness given by learned cultural traits (Boyd and Richerson, 1988; Henrich and McElreath, 2003; Morgan et al., 2012). In this context, we expect adaptive agents to select interactions increasing their fitness.

Different mechanisms can evaluate this expected outcome. Some are *direct*, where individual will manipulate or look at the artefact with a direct access to its intrinsic value, some are *indirect*, where individual cannot know this intrinsic

[1] In this thesis we will use *Social Learning Strategies*, *Cultural Transmission Bias* or simply *Bias* as strictly equivalent terms that describe any process of Cultural Transmission that modify the distribution of cultural variants in a non-random way.

value directly but will have access to cues or marks that are highly correlated with it. In Figure 1.1, we propose to classify SLS according to their dependence or not on the content (the intrinsic value) of the cultural artefact.

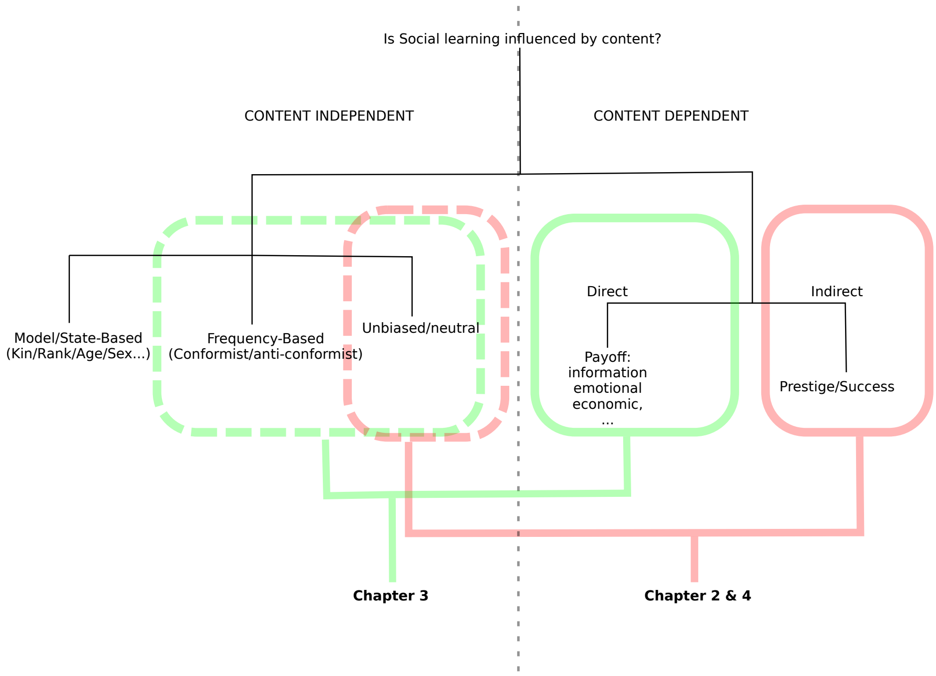

Figure 1.1: Classification of Social Learning Strategies (SLS) according to their dependence upon the content of cultural artefacts. The colored boxes indicates in which chapter the different bias will be studied. Dotted lines represent content-independent SLS, while full lines represent content-dependent SLS.

Thinking in terms of different social learning strategies allows to pull back theses processes at the individual level and to make hypotheses about their neural, psychological, anthropological and evolutionary bases (Kendal et al., 2018). In the meantime, while will enlarge the set of tools for detecting individual mechanisms behind social learning strategies and their impact in the historical record. This allows to test assumptions about social structure and the pace of social interactions at different places and times. In this context, previous studies have mainly focused in frequency-dependent social learning (Bentley et al., 2004; Ruck et al., 2017; Kandler et al., 2017; Crema et al., 2016; Pagel et al., 2019) studies where cultural artefacts have no intrinsic value. Thes works assume that cultural content is neutral and doesn't play any role in cultural transmission. In most of these studies the biases result from differences in the frequencies distribution of the variants. On the other hand, the studies that have focused on content-dependent bias (direct or indirect) have often limited their scope to experimental studies within the lab (Mesoudi, 2008), to the individual or small group scale (Henrich and Gil-White, 2001; Stubbersfield et al., 2019; Morgan et al., 2012) or to purely theoretical scenarios (Baldini, 2012; Kendal et al., 2009). Few attempts have been made to detect and quantify the impact of content-dependent biases in cultural changes at large temporal and spatial scale in real scenarios. This thesis seeks to fill that gap by developing and testing a method to detect and explore direct and indirect content bias in large data sets that describe cultural change in human society at the population level.

1.2 Methods

1.2.1 Agent Based Model

Exploring social learning strategies at large scales has a cost. It implies modelling the way human or group of humans process and evaluate information, directly from the artefact or by analysing and interacting with their environment, to attribute value to the cultural traits they will copy and transmit. Those mechanisms of evaluation are potentially complex and may involve different decision making processes acting at different levels. Such processes can hardly be describe as simple equations and are even harder to scale to describe populations interacting during decades, or even centuries.

To solve this issue, this thesis relies on Agent Based Modelling (ABM), a modelling approach that allows to keep mathematical formalisms while integrating a flexible and complete description of the social learning strategies we want to study. Its not a surprise if this method has gained more and more interest in disciplines where complex and heterogeneous individual interactions are central to understand the global properties. This has been the case in Sociology (Epstein and Axtell, 1996), Economy (Tesfatsion, 2006), Ecology (Grimm and Railsback, 2005) or Archaeology (Kohler et al., 2012).

In its general formulation, an Agent Based Model is a computer program where "agents" interacts together and with the environment given some predetermined "rules" (Axelrod, 1997). To mimic the real world where everything happen at the same time, agents' interactions are run in parallel. This is not trivial in a computer: processor can only interpret one command after the other in a sequential order. This is why the central core of an Agent Based Model is often a set of instructions that takes care that every action is run in the most parallel way possible (often by randomising the order at which agents are acting). If this is slowly changing with new computers with multiple nodes and cores, where in theory every agents could physically act in parallel, in practice a central way to orchestrate the interactions is always needed (Rubio-Campillo, 2014).

In these models, the agents and the environment represent more or less concrete real entities. Agents can be cells, animals, or totally abstract concepts, like a company, a group of people, a country or a generation of individuals. The environment can also represent a real physical one, like a room or a fields, or a totally abstract place, like an economic market, a world of idea or a mathematical construction. Both the agents and the environment can be specified with more or less precision: an animal can be represented by a few attributes that define its specie and age, or they can incorporate biological function like "eating", or complex cognitive properties that will guide their behavior. The environment can be a stylised place represented by a network that defines who see who, or it can integrate a complex physic, defining numerous different places with different properties and rules that will modulate and limit the behavior of the agents.

1.2.2 Approximate Bayesian Computation

These flexible formulations makes difficult to test and compare models one with each other or against real data, in a robust and consistent way. A new branch of techniques emerged to overcome this problem and calculate the likelihood of any model under different empirical knowledge. This method is called Approximate Bayesian Computation (ABC) and relies on the Bayes equation:

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)} \tag{1.1}$$

where θ is the vector of parameter for a given model, $P(\theta)$ is the prior distribution of the parameters, $P(D)$ the data and $P(D|\theta)$ the likelihood Gelman and Hill (2007).

This equation measures the degree of certainty about the observation we have. We can use Equation 1.1 to determine the distribution of parameters for which a given model can explain the data (so-called posterior distributions).Then, it is possible to compare different models through their posterior distributions and select the best options among competing hypotheses.

Finding the likelihood of Agent Based Models is not an easy task, as there is no straightforward way to do so. ABC approximates the likelihood of models by rejecting the parameter ranges that yield results too far from the empirical observations. The method was first described by Rubin (1984),but it was only during the past 15 years, with the increase of cheap computational power, that it started to be broadly used. The method was popularised by evolutionary biologists (Tavaré S et al., 1997; Beaumont et al., 2009; Csilléry et al., 2010; Toni et al., 2009) and is now attracting many scholars studying archaeology and cultural evolution (Rubio-Campillo, 2016; Crema et al., 2014; Kandler and Powell, 2018; Pagel et al., 2019). When applied to cultural evolution, ABC is often used to find the most probable social learning strategy generating the observed data. Here, we continue this trend and combine Agent Based Modelling with Approximate Bayesian Computation to quantify content-dependent biases in social learning.

1.3 Objectives

As we have seen, understanding content-dependent bias is a challenge for Cultural Evolution, for its potential adaptive role in the evolution of social learning as well as for the complex processes it involves. In this thesis we propose to push further this understanding by exploring the potential importance of content dependent social learning strategies:

- in the emergence of economic equilibrium,
- in how they can explain observed pattern of economics change and,
- in how they could be at the root of differences in spread of online information.

This will also gives us a great and unprecedented occasion to test how content-dependent biases are impacted by factors such as the nature of the cultural artefact that spreads, the pace and time scale at witch the cultural evolution take place, and if the bias involved in those changes is direct and indirect.

The methods developed in this thesis allow a formal comparison between content-dependent and independent Social Learning Strategies in different theoretical and empirical scenarios. We will do so with three case studies:

1. A purely theoretical scenario to quantify the expected outcome of success-biased Social Learning Strategies.

2. The spread of news in online social networks, where information propagates in a relatively short time scale (a few years).

3. The distribution of ceramic tablewares over inter-regional trade networks of the Roman East. This took place over large time scales (hundred years).

The theoretical scenario will allow us to develop the computational tools used all over the thesis while illustrating the potential importance of content-dependent biases. Both empirical scenarios imply different kind of cultural artefacts (tableware vs. online information), at two different time scales (9 vs. 400 years). Moreover, we expect those two different cultural changes to be driven by two different content-dependent biases: in online online information, the content is directly available. We expect *direct content bias* to be central in this scenario. In the other case, the intrinsic value of tableware isn’t directly accessible, thus an *indirect bias* (relying on the economic success of traders) sounds like a better candidate to drive the changes observed in the third scenario.

Thank to these two empirical scenarios we hope to explore how the nature of the cultural artefact and the geographical and time scale at which these artefacts spread can influence direct and indirect content bias.

1.4 Plan

In the Chapter 2 we will detail and theoretically explore how Success-Bias, an indirect bias toward content, can lead to complex economic scenario. We propose a computational framework that allows to explore it and compare to other social learning strategies, while in this we will solely compare it to an unbiased social learning strategies also call the neutral model (*cf.* the red box in the Figure 1.1).

In the Chapter 3 we will present a method that allows to formally test and compare different social learning strategies against empirical data. Direct content dependent bias and several frequency-dependent biases will be tested (*cf.* the green box in the Figure 1.1) to try to understand the spread of true and false information in the online social media Twitter.

In the Chapter 4 we build upon the framework developed in the Chapter 2 to solve some of the issues met in the Chapter 3 and test the ability of the neutral model, the Success Bias strategy and a independent learning strategy to reproduce observed economic change in the consumption of tableware in the Roman East from -300 to 300.

We will end the thesis by wrapping up the results of every chapter and discuss the importance played by the mechanisms explored in this thesis on the future of Humanity.

Chapter 2

A THEORETICAL FRAMEWORK TO QUANTIFY SUCCESS BIAS IN SOCIAL LEARNING STRATEGIES

2.1 Introduction

In this chapter we present a framework to explore social learning strategies. The design of this framework aims for a trade-off between the flexibility necessary for the implementation of multiple models of social learning and the structure necessary for the comparison between the models implemented. To achieve this flexibility we propose an Agent-Based Model implementing a general purpose evolutionary algorithm.

Within this framework we explore two strategies: one totally content agnostic, the unbiased or "neutral model"; and one which is indirectly dependent to content, the "success bias model". We will see how adding the indirect bias toward content (the success biased learning strategy) allows the emergence of an economic equilibrium without central planning. Our goal is to illustrate how social learning can be used to explore complex social systems such as trade and economic interactions with the help of Agent Based Modelling and Computer Simulation. This chapter makes clear how success bias can be the engine of complex emerging properties such as economic equilibrium, giving credit to its potential importance in human societies.

After introducing what we define as social learning and trade, this chapter will detail the architecture of the framework and explore the general properties of the interaction between trade and social learning. We then briefly show how it can be extended to explore different aspects of the system studied, by 1/ changing part of the network of social interactions and 2/ generalising the exploration of those properties using Bayesian Statistics.

2.2 Trade and Social Learning

Multiple social learning strategies drive cultural changes. One of the most studied, that we will use as a reference in the chapter and throughout the whole thesis, is the neutral model (Neiman, 1995). This model assumes that the nature of a trait does not bias the probability of an individual to acquire it. It is totally content-independent. It therefore means that no bias modifies the rate of transmission of the cultural traits, and that their success will depend only on their initial frequency in the population. Within analyses of real data, a neutral model produces a distinctive type of frequency distribution of cultural traits termed *power law*.

Neutral transmission predicts a power-law distribution of cultural (Bentley et al., 2004): an individual will copy the traits of a randomly chosen individual with a given probability. This copy can potentially introduce some errors in the acquired trait, which account for innovation processes. The individual will in turn continue to spread these cultural traits which will be further adopted by other individuals. This basic model can be enriched by several additional processes both in the innovation (Schillinger et al., 2014; Solé et al., 2013; Ziman, 2003) and the transmission (Heyes, 1994; Henrich and McElreath, 2003). Unbiased transmission works as a baseline for identifying frequency-dependent biases: if individual has higher tendency to copy the most common trait it is known as conformism, while the opposite is defined as anti-conformism.

This work explores the impact of a crucial element on the transmission of material culture: trade. Networks of good exchanges are being increasingly recognised as key elements that structured ancient societies (Temin, 2006; Remesal et al., 2014; Brughmans, 2010). The scenarios where this process emerges suggest a complex bias in the selection of cultural traits, which at the same time are identified as economic goods (Bentley et al., 2005; Macmillan and Huang, 2008). Transmission is not neutral anymore, as different prices for each good will introduce a dynamic content bias. This affects the frequency of the good within the population, which in turn modifies its price following a co-evolutionary dynamic.

These dynamics are studied using an Agent-Based Model (ABM) also often called Individual-Based Modelling by ecologists (Grimm and Railsback, 2005). This type of simulation is particularly useful for studying non-linear dynamics in heterogeneous environments within an evolutionary perspective (Lake, 2014) and is becoming a standard tool in the community of scholar studying social phenomenon and cultural evolution (Acerbi et al., 2009). More precisely we propose a framework that can be implemented in multiple ways depending on the model tested. Next section defines the framework, which is based on the basic processes found in evolutionary models of cultural change and social learning strategies. Next, we define the implementations used to explore the dynamics of the created framework. In the following section we analyse the results obtained with these two implementations. We then present two different ways this framework can be extended to explore particular features of the system. Finally, the concluding remarks discuss further possibilities of the presented framework.

2.3 Trade Model

To explore the co-evolution between trade and cultural change we have developed a framework where the different agents produce and trade goods to which they assign variable values. The model is composed of a population Pop of m agents, each defined by two vectors of size n. The first corresponds to the quantity of each good owned by the agent i:

$$\forall i \in Pop, \quad Q^i = (q^i_1, \cdots, q^i_n)$$

where Q^i is the total list of possessions of agent i, and q^i_j is the number of goods of type j that agent i possesses.

The second vector reflects the estimation of the value of a good made by an agent i:

$$\forall i \in Pop, \quad V^i = (v^i_1, \cdots, v^i_n)$$

where V^i is the total list of estimated values of agent i, and v^i_j is the value that agent i associates to one unit of a good of type j.

On top of these elements five processes are used: *production*, *consumption*, *cultural transmission*, *innovation* and *trade*. The *production* process describes the creation of goods by the agent. Once a good is produced by an agent i it is added to its quantity vector (Q^i). The consumption strategy of these goods is defined in the *consumption* process which decreases the number of goods in the vector (Q^i) and all goods are completely consumed at each iteration for all the models tested. The *trade* process models the exchange of goods between the agents which results in a modification of the quantity vectors (Q^i). The amount of goods exchanged is computed by the agents involved in the trade, within the *trade* process, based on their value vectors (V^i). The *cultural transmission* process implements the social learning strategies (i.e. who will copy who and how) that we want to explore in this chapter. During the *cultural transmission* an agent i can copy the entire value vector (V^j) of an agent j, where $j \neq i$. Finally, the *innovation* process also modifies the value vector V^i of an agent, but it differs from the *cultural transmission* process in that the modification is done without reference to the other agents.

The scheduling of the five processes is described in Algorithm 2.1 along with the vectors modified by each of these processes. On lines 3 and 4 all agents of the population are initialised with empty quantity vectors and random values. The code used to update the status of each agent at each iteration is presented between the lines 8 and 16. One can note that each of the five processes is executed synchronously by all agents. Moreover, the *trade* process is called at each iteration while the *cultural transmission* and *innovation* processes are executed only every $CulturalStep$. The idea behind this is to perform the *cultural transmission* based on a score that reflects the performance of the agent and not only one lucky or unlucky trading round. The timestep number used in the figures of this article refers to the number of times the *cultural transmission* and *innovation* processes are called.

In order to validate our model we first reproduce common results from the literature on cultural transmission. We then show that it is possible to transform our model to fit processes that are economically sound, i.e. the model should converge to optimal values as in Gintis (2006). To achieve these two goals, we have designed for each one a specific set of implementations of the five core processes (production, consumption, trade, cultural transmission and innovation).

Algorithm 2.1 General model of trade and social learning.

INITIALISATION:
for $i \in \#Pop$ **do** ▷ Initialize the agent with no goods and a random value vector
    $Q^i = (0, \cdots, 0)$
    $V^i = (v_0^i, \cdots, v_n^i)$ ▷ The values of v_j^i are selected randomly
SIMULATION:
loop $step \in TimeSteps$
    for $i \in Pop$ **do**
        $Production(Q^i)$
    for $i \in Pop$ **do**
        for $j \in Pop$ **do**
            $TradeProcess(V^i, Q^i, V^j, Q^j)$
    for $i \in Pop$ **do**
        $ConsumeGoods(Q^i)$ ▷ All goods are consumed
        if $(step \mod CulturalStep) = 0$ **then**
            $CulturalTransmission(V)$
            $Innovation(V^i)$

2.3.1 Neutral Scenario

The first scenario is designed to reproduce unbiased transmission, also known as the neutral or random model, where each good is a cultural trait without intrinsic positive or negative weight (Bentley et al., 2004, 2005; Mesoudi and Lycett, 2009). Under this hypothesis, the processes of *production* and *trade* are not relevant, and as a consequence, they do not modify the content of the quantity vectors of the agents.

Unbiased *cultural transmission* is implemented using "random copy": each agent has a probability to pick randomly one agent among all and copy its vector of values. The *innovation* process, termed "unbounded", is triggered with a low probability (μ) and draw a new random value to replace an element v_j^i. The probabilities to trigger *cultural transmission* and *innovation* are presented with other parameters in table 2.1.

The neutral hypothesis states that the "random copy" transmission and the "unbounded" innovation process used under a fixed population size leads to a distribution of frequency of cultural variants termed *power law*. This distribution is characterized by a small number of very frequent traits and a large number of rare traits. The main difference with similar distributions, such as exponential distribution, is that the rare traits are far from being absent of the distribution, i.e. the tail of the distribution is large. This distribution is formalised as:

$$P(v) = C/v^{\alpha} \tag{2.1}$$

where v is the number of times a variant has been repeated, $P(v)$ the probability to find that variant, C a constant, α a variable describing the slope of the curve obtained. We will therefore attempt to fit as well as possible the results obtained with this set of implementations to the "power law" distribution by modifying the α parameter.

2.3.2 Success Bias Scenario

In the second scenario, we are interested in the exchange of goods between agents in a barter process where each agent can choose his "prices" of exchange (*ie* the amount of good he expects to receive when exchange for an other good). We want to implement simple processes leading to the convergence of all prices to values acceptable by all agents, i.e. we would like to observe, at the end of an experiment, all the agents using a set of prices which allow them to trade efficiently.

Production: Each agent produces one good. The type of good produced by an agent i is assigned to it at the beginning of the simulation, does not change through the simulation, and is referred to as $produced^i$. At each time step, each agent, produces a number of units (of its production good) equal to the number of goods, which ensures that enough is owned to be traded for other goods. Moreover, when an agent consumes its own production good, it does not impact its inventory.

Cultural transmission: Social learning is biased towards the agents which are the best at trading, and is therefore termed success bias. To achieve this bias, the cultural transmission mechanism used takes into account the value vector of the other agents and relies on two new notions: *need* and *score*.

The *need* is a quantity of good that each agent tries to obtain. This quantity is different for each good but the need for a good is the same for all agents:

$$N = (n_1, \cdots, n_r)$$

The *score* s^i_j of an agent i reflects the ability of this agent to obtain the quantity of good j it needs. It is maximum when the quantity q^i_j that an agent i owns of the good j is equal to the need n_j for the good j and lowers proportionally to the distance between the need vector and the quantity vector. It is formally computed as follows for agent i and the good j:

$$s^i_j = \begin{cases} s_{max} = 1 & \text{if } q^i_j = n_j \\ 1 - \dfrac{|q^i_j - n_j|}{\sqrt{|(q^i_j)^2 - (n_j)^2|}} & \text{if } q^i_j \neq n_j \end{cases} \tag{2.2}$$

This function ensures that each good has the same weight in the final score, i.e. managing to get only the right amount of a good with a high "need" value will not give a better score to the agent. The Figure 2.1 depicts the score of an agent

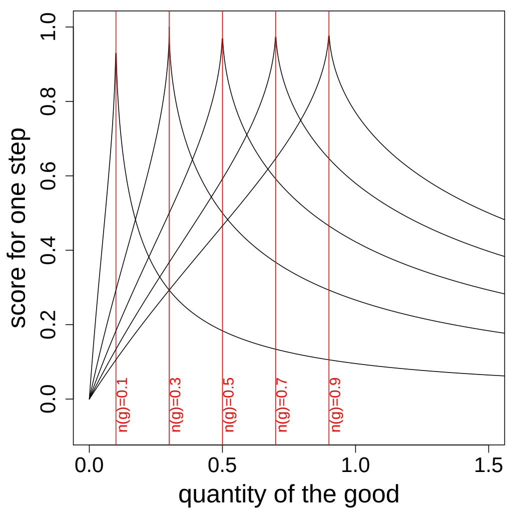

Figure 2.1: Theoretical payoff for one good g and one agent, given the quantity owned and different needs $n(g)$ (in red) for this good.

for different values of need, with regards to the quantity of the good possessed by the agent. The complete score of an agent i is termed s^i and corresponds to the sum of the s^i_j. An agent will choose from whom the price vector should be copied among the agents that produce the same good and have the highest score. In practice, the worst (in terms of score) twenty percent of the agents producing the same good will copy the prices of the best twenty percent producing the same good. This selection process is detailed in Algorithm 2.2.

Algorithm 2.2 Social Learning Process.

```
ToGet = 0.2 × #Pop / #Good
for g ∈ Good do
    ToReplace = {}
    while #ToReplace < ToGet do
        j = SelectRand(Pop, g)        ▷ Select randomly an agent j among the agents producing g
        X ∼ U([0, 1])                 ▷ Draw a random number from the uniform distribution between 0 and 1
        if X > ComputeScore(j) then   ▷ Select preferably the agents with the lowest scores
            ToReplace = {ToReplace, j}
    while #ToReplace > 0 do
        j = SelectRand(ToReplace)
        i = SelectRand(Pop, g)        ▷ Select randomly an agent i among the agents producing g
        X ∼ U([0, 1])
        if (X < ComputeScore(i)) then ▷ Select preferably an agent i with a high score
            if (ComputeScore(i) > ComputeScore(j)) then   ▷ Verify that agent i has a higher score than agent j
                CopyPrice(i, j)
                ToReplace = ToReplace − i
```

Trading: During the trading phase the value associated to a good by an agent corresponds to the subjective price of the good for this agent. Briefly summarised, for each good that it does not produce, an agent will trade with the first partner that offers an acceptable trade, i.e. an agent that proposes a satisfiable ratio between the other good and the good produced by the agent.

In more details, the trading phase starts by the agent looking at a first random agent producing another good. Let o be an agent producing g who proposes a trade and r an agent producing k that receives the proposition. As explained earlier, each has a quantity of good Q^o and Q^r. On the one side, o wants to exchange a quantity w_g^o of the good g for a quantity w_k^o of the good k. On the other side, r wants to exchange a quantity w_g^r of the good g for a quantity w_k^r. The tuples W^o and W^r describe the quantities of goods wanted by agent o and r for one trade proposition and are defined by:

$$W^o = (w_g^o = v_g^o, w_k^o = \frac{v_k^o}{v_g^o}), \quad W^r = (w_k^r = v_k^r, w_g^r = \frac{v_g^r}{v_k^r}) \tag{2.3}$$

Where v_j^i are the estimated value of the good j by the agent i as defined earlier. The requested quantity of the non produced good is simply the ratio between the estimated value of the requested good and the estimated value of the produced good.

Once the quantities are defined, the agents declare that the trade is possible if :

$$q_g^o >= w_g^o, \quad q_g^r <= w_g^o, \quad q_k^r >= w_k^o \tag{2.4}$$

$$w_g^o >= (q_g^r + w_g^r), \quad w_k^o <= w_k^r, \quad w_g^o <= w_g^r \tag{2.5}$$

The conditions 2.4 ensure that both agents have enough goods in their inventory to realise the trade while the conditions 2.5 ensure that the quantities of goods fit the will of both agents.

If a trade is possible, the two agents will exchange the agreed quantities. If the trade is not possible, the agent will continue to look at random partners for this good until either a partner is found or $TradeThreshold$ agents have been tried. At this point the agent will try to trade with agents producing another good. The process goes on until all goods have been tried. This trading process is described in Algorithm 2.3.

Algorithm 2.3 Trading Process for agent o

for $j \in Goods \ \& \ j \neq produced^o$ **do**
  $tradeAttempt = 0$
  for $r \in Pop \ \& \ produced^r = j \ \& \ tradeAttempt < TradeThreshold$ **do**
    if $acceptableTrade(W_o, W_r)$ **then**
      $trade(W_o, W_r)$
    else
      $tradeAttempt = tradeAttempt + 1$

Innovation: In a trading environment it seems unlikely that a price will change radically to a very different value. Therefore, a new and more realistic mechanism is proposed. The innovation process, coined "self referenced", is still triggered with probability μ but modifies the previous price by adding or subtracting a small amount taken randomly from a uniform distribution between $[0..\beta]$. In the rest of the text this will be the default innovation process used when the success bias model is mentioned, unless specified otherwise.

Expected outcome: Based on the set of implementations presented and given Equations 2.2 and 2.3, it is expected that the prices will converge to value allowing each agent to obtain quantities of resources exactly equivalent to the needs. The best possible price for each good satisfies the equations :

$$\begin{cases} \frac{v_k^o}{v_g^o} = n_k \\ v_g^o = n_g \end{cases} \quad => v_k^o = n_k \times n_g \tag{2.6}$$

for all $k \in Goods, o \in Pop$ and with $g = produced^o$ and $k \neq g$.
This means that for all $j \in Goods$ and $i \in Pop$:

$$\tilde{V}^i = \begin{cases} \tilde{v}_j^i = n_j & \text{if } j = produced^i \\ \tilde{v}_j^i = n_j \times n_{produced^i} & \text{else} \end{cases} \tag{2.7}$$

If such prices are reached, given the exchange rules defined in (2.3) and the exchange constraints (2.4) and (2.5), all exchanges will be optimally achieved, leading to a total score S for each agent of the population :

$$S = \sum_{i=0}^{CulturalStep} s^i(\tilde{Q}^i) \times ngoods \tag{2.8}$$

where $\tilde{Q}^i$ is the optimal quantity vector, i.e. the one for which $s^i(\tilde{Q}^i) = s_{max}$ and $ngoods$ the total number of goods. Remember that from equation 2.2, $s_{max} = 1$.

2.3.3 Experimental setups

The neutral scenario is tested through 15 experimental setups. The first six experimental setups are using one good, two population sizes (250 and 500 agents) and three values of μ (0.004, 0.016 and 0.064). The remaining experimental setups are using 500 agents, three number of goods (3, 6 and 9) and three values of μ (0.004,0.016 and 0.064). For each setup, we have performed 100 runs of $10,000$ timestep each. The success bias is tested on one unique experimental setup using 3 goods, 500 agents and μ equal to 0.004. Again 100 runs of $10,000$ timesteps are performed.

All parameters used to run the 16 experiments are summarized in Table 2.1. The full source code of the agent based model, which relies on the Pandora simulator (Rubio-Campillo, 2014), is available online: https://github.com/simoncarrignon/2015-WSC-source.

Table 2.1: Table of parameters used to run the experiments.

Parameter	Value
Number of goods	$\{1, 3, 6, 9\}$
Number of agents (N)	$\{250, 500\}$
Innovation probability (μ)	$\{0.004, 0.016, 0.064\}$
Innovation range (β)	0.005
Innovation process	$\{unbound, self\ referenced\}$
Cultural transmission probability	0.001
Social Learning Bias	$\{neutral, success\}$
CulturalStep	10
TradeThreshold	100
TimeSteps	$10,000$

2.4 Equilibrium

2.4.1 Neutral Scenario

We first analyse the results obtained in the neutral model with one good. The Figure 2.2 presents the results obtained for two population sizes N (250 and 500 agents) and μ varying through three values (0.004, 0.016, 0.064). The figure is a log-log plot of the average (across all the runs) of the distribution of variants obtained through all experiments. The y-axis shows the proportions of the variants of the prices used during the simulation, the x-axis shows how many variants achieve such proportions.

We observe that the lower the mutation rate, the closer to a line the result is. This line corresponds to the “power law” distribution explained in section 2.3.1, and is typical of the result obtained under the “neutral hypothesis”. In order to verify if the distribution is in fact a power-law, we follow the method proposed by Clauset et al. (2009) and the R implementation proposed by Gillespie (2015). Briefly, two values are returned: a) the estimation of the α parameter of the power law equation $P(v) = C/v^{\alpha}$; b) a p-value testing the null-hypothesis that our data could have been generated from a power law distribution.

The table 2.2 summarizes the results obtained on all setups. Each value shown in the table is the mean value for 100 simulations. We see that in almost

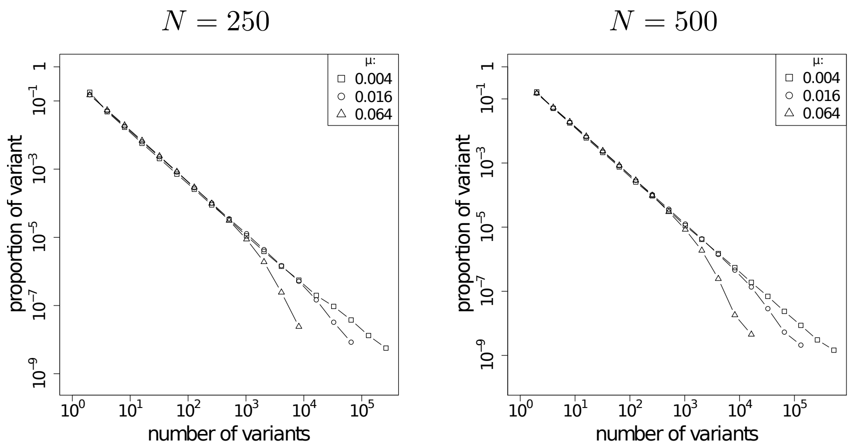

Figure 2.2: Distribution of proportions depending on the μ parameter with 250 agents (left) and 500 agents (right). The shape of the dot represent different μs, each line represents the mean obtained for 100 runs.

every case the null hypothesis cannot be rejected, which means that indeed the repartition of the price follows a power law. The only exception is for μ equal to 0.064, where the p-value is less than 0.05. In this last case the null hypothesis is rejected, and we therefore assume that the distribution does not follow a power law.

Table 2.2: Mean α & sd are calculated on 100 runs for our results, and 5 runs for Bentley et. al. 2004. pr is the percentage of run for which the p-value is less than 0.05, i.e. the percentage of runs for which we rejected the null-hypotheses stating that the distribution follows a power law.

		Our results		Bentley et. al. 2004
N	μ	α (sd)	p-value $(sd$ - $pr)$	α (sd)
250	0.004	1.53 (0.03)	0.58 (0.24 - .01)	1.54 (0.02)
	0.016	1.57 (0.02)	0.35 (0.28 - .05)	1.57 (0.01)
	0.064	1.66 (0.01)	0.0 (0.00 - 1)	1.67 (0.01)
500	0.004	1.50 (0.02)	0.59 (0.28 - .02)	1.53 (0.03)
	0.016	1.55 (0.03)	0.15 (0.17 - .10)	1.61 (0.04)
	0.064	1.78 (0.08)	0.0 (0.00 - 1)	1.81 (0.10)

For comparison purposes the results obtained by Bentley et al. (2004) (which tested the “neutral hypothesis” with the same methodology) are added in the last column of the Table 2.2. It appears that these results are highly similar to ours. We note also that our results match the ones presented in Mesoudi and Lycett (2009) where only one value of μ was tested ($\mu = 0.008$). However, it is difficult to know if the slight differences observed between our work and those previous studies are statistically significant as the two previous studies rely on only five runs for the computation of the mean of α (against 100 in our case).

Nonetheless, for high values of μ, previous works report that the distribution of variants follows a power law (Bentley et al., 2004). This claim is based on the fact that the estimated α (estimation based on linear regression on the log-log curve) has a high correlation coefficient. Recent works have shown that the use of correlation coefficient should be avoided (Clauset et al., 2009). Following these recent findings, our results preclude us to assume that the distribution of variant when μ is up to 0.064 follows a power law.

An additional series of experiments has been done to analyse how the system reacts when multiple goods are present. The mean values of α and p-value have been analysed for three innovation rates (0.004, 0.016, 0.064), four number of goods (1, 3, 6, 9), and 250 agents. The results are given in Table 2.3. The visual representation of those experiments is not given as it resulted in curves overlapping exactly the curves shown in Figure 2.2.

Innovation rate	Number of goods			
	1	3	6	9
0.004	1.53 (0.03)	1.53 (0.03)	1.53 (0.03)	1.53 (0.03)
0.016	1.57 (0.02)	1.57 (0.02)	1.56 (0.017)	1.57 (0.02)
0.064	1.66 (0.01)	1.66 (0.01)	1.66 (0.01)	1.66 (0.01)

Table 2.3: Mean α and standard deviation for all innovation rate and different of goods. The α given here is the mean α among all goods taken separately and computed over 100 runs with 250 agents.

The αs given in Table 2.3 as well as visual analysis, reveal that independently of the number of goods, the distributions obtained are exactly the same.

2.4.2 Distribution of variants

In order to understand the effect of introducing trading mechanisms, we compare first the distribution of values obtained in the success bias (with the self-referenced innovation mechanism) against the values obtained in the neutral model (with the unbounded innovation mechanism). The Figure 2.3.a) presents the results obtained from 100 runs for each model. All runs rely on the same experimental setup using three goods, 500 agents and μ equal to 0.004. In all following graphs a variant is one price of one given good. The distributions are first computed for each good independently and then averaged together.

On figure 2.3.a) it appears that the implementation of success bias model leads to a distribution of prices departing from the neutral hypothesis. In more details, the frequencies distribution has a plateau of common prices (a number of prices share similar and high proportions), which shows that, when trade is taken into account, the most common variants are more diverse.

In order to investigate which mechanism influences this departure from the neutral model, additional investigations have been performed. Here is presented the results of the analysis on the effect of the *innovation* process of the success bias model. To conduct this analysis the *innovation* process of the success model has been replaced by the *innovation* process of the neutral model, we will call it "success bias*". 100 runs have been performed with this model on the same experimental setup. The results obtained are compared to the neutral model in figure 2.3.b) and to the success bias model in figure 2.3.c).

On figure 2.3.b) it appears that the replacement of the innovation process leads to the creation of a distribution close to the one obtained with the neural model. On figure 2.3.c) we observe a strong reduction in the size of the plateau and an important difference between the two distributions. This analysis highlights the importance of the self referenced *innovation* process in the distribution of prices. This mechanism, by preventing the creation of totally random new prices, promotes the creation of few similar prices.

2.4.3 Study of scores

The results obtained with the success bias model are studied in more detail by investigating the ability of the population to find the price most suited for exchanges. This is done by first measuring the score of all agents in each of the two different models. The Figure 2.4 uses again the results obtained from 100

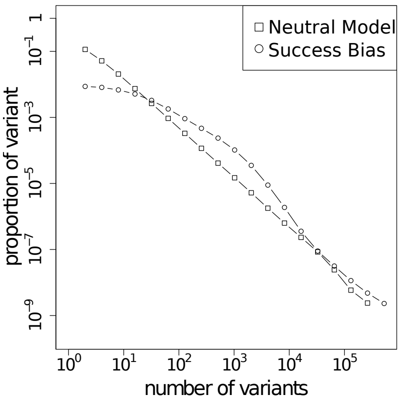

a)

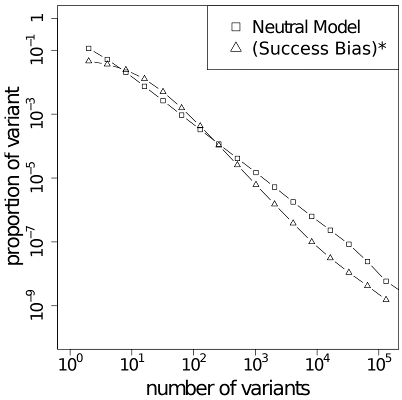

b)

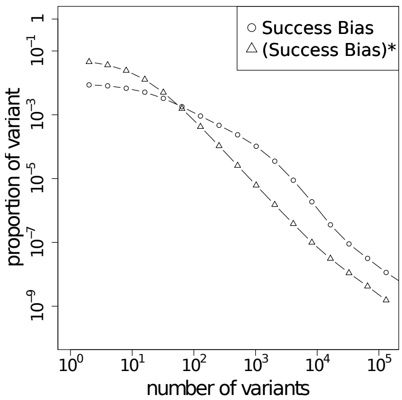

c)

Figure 2.3: Frequencies distributions, where each line represents the mean for 100 runs, for: a) the neutral model and the success bias models. b) the neutral model and the (success bias)* model (without the self-referenced innovation process). c) the success bias model and the (success bias)* model.

runs for each model where all runs rely on the same experimental setup using three goods, 500 agents and μ equal to 0.004. The figure represents as box plots the score computed thanks to equation (2.2) for all agents of all runs. The y-axis shows the score computed and the x-axis shows the timesteps. The left plot shows the results obtained in the neutral model and the right plot shows the results obtained in the success bias model.

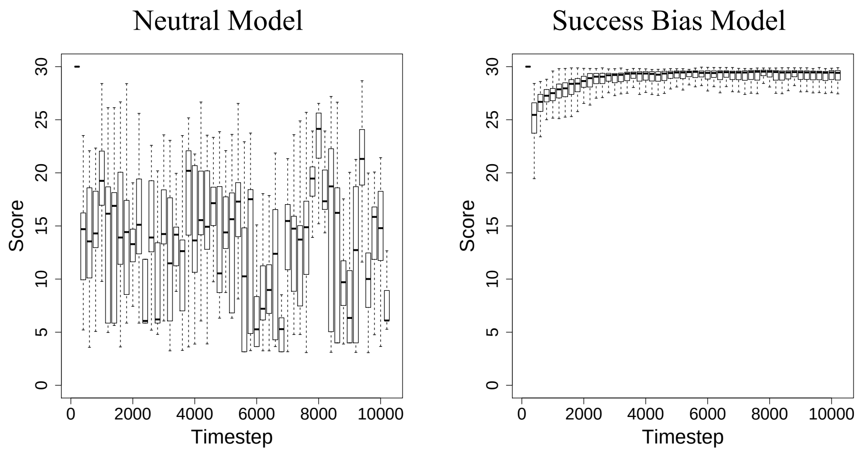

(a) (b)

Figure 2.4: Example of the evolution of the score of every agents for the two different models for two typical runs with 500 agents and three goods exchanged during $10,000$ timestep using the random model (a) and the success bias model (b) of cultural transmission.

As expected, the scores within the neutral model vary randomly. "Trends" may appear, where a bigger proportion of individuals adopt a better price that allow agents to reach better score (*cf.* iteration $8,000$ on Figure 2.4a) but such good score fall back as soon as another trend appears. However, with the success bias model, the score of all the agents increases. As the selection mechanism allows them to know who has found better vectors of prices, they will progressively adopt prices vectors that allow all of them to reach better scores.

The previous figure showed the capacity for the success biased model to increase its score but did not analyse the exact prices used. As explained in Section 2.3.2 we expect that the trade version of the *cultural transmission* and *in-*

novation processes will produce a convergence toward a set of prices for each good that will allow agents to exchange optimally the good they produce with the other goods. To verify this assumption we analyse the prices reached during the simulations. These are presented in figure 2.5 for the 100 runs relying to the experimental setup using three goods, 500 agents and μ equal to 0.004. For all runs, all agents and at each iteration we compute the difference between the prices used by the agent V_g and the optimal price $\tilde{V}_g$ (given by equation 2.7). The measures performed are presented as boxplots condensing the results for 100 runs.

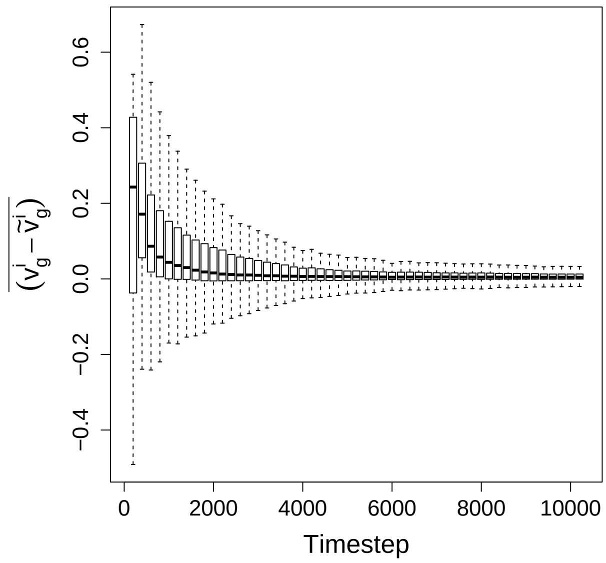

Figure 2.5: Evolution of the mean of the difference between the estimated value v_g^i and the optimal value $\tilde{v}_g^i$ (calculated with equation 2.7) for a good g and an agent i. As the optimal value $\tilde{v}_g^i$ depends on which good is produced by i, the mean of the difference between the estimated price and the optimal one is computed between all the agents that produce the same good. The figure represents this mean computed at each timestep for each good, for each group of agents and for 100 runs in a setup with 500 agents and three goods.

We observe that prices are indeed converging to the optimal prices which means that the agents within the success bias model are indeed improving their scores by reaching the optimal prices. Notably, a similar variation of prices was observed in the closely related economical model of Gintis (2006). This variation of the prices to the optimum offers an additional conformation of the validity of the success bias model.

2.5 Experimental Exploration

In the following two sections we propose to explore more in details some aspects of the framework we developed before. The first example tests different social networks, the second illustrates how we can use Bayesian probabilities to formally describe and test the properties of our models.

2.5.1 Social Network Topology

Throughout this chapter we developed a framework where dynamics at the economic level impact dynamics at the cultural level and vice versa. At both levels agents interact with each other: during the $TradeProcess$, they exchange at the economic level while during the $CulturalTransmission$ they exchange at the cultural level. In both cases agents interact within a network of other agents, in one level exchanging goods, in the other exchanging information. The two levels and the networks are represented in Figure 2.6.

In the previous sections the networks were fully connected: all agents interact with all other agents (under some condition detailed in Section 2.3.2), and no specific structure is imposed. Here we want to see if the shape of the cultural networks of the agents can impact the economic equilibrium. To do so we propose to change the topology of the cultural network by tweaking metrics that describe this topology and which are standard in network analysis: the average distance (L) and the average degree ($\langle k \rangle$) of the social network:

- L, the average distance, is the mean of all shortest paths between all nodes in the network and gives us an idea on how quick information goes from one point to another in the network.
- $\langle k \rangle$, the average degree, is the mean of all degree of all nodes in the network which gives us an idea of how many “contacts”, each agents have in the

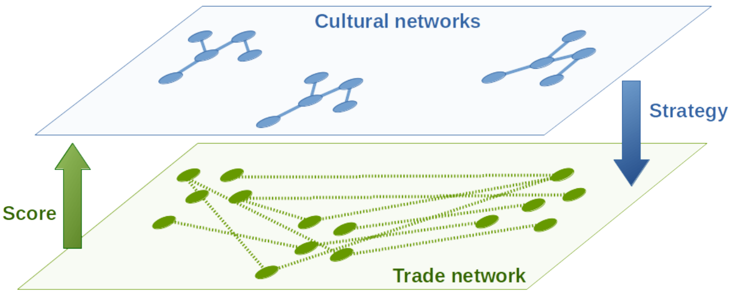

Figure 2.6: Illustration of the interaction between the Trade network (in green) and the Cultural networks. At the Cultural level, links in the networks represent Social Learning interactions, while at the Trade level, links represent Commercial (buy and sell) interactions.

	$\langle k \rangle_1$	...	$\langle k \rangle_n$
L_1	G_{11}		G_{1n}
...		...	
L_m	G_{m1}		G_{mn}

Table 2.4: Table summarizing how are combined different average distance L with different average degree $\langle k \rangle$ in our 28 networks G_{lv}.

contact list.

To test the influence of L and $\langle k \rangle$ we designed 28 different network topologies which combine different average distances L with different average degrees $\langle k \rangle$ as we show in Table 2.4. Where $L \in \{4, 6, 8, 10, 14, 18\}$ and $\langle k \rangle \in \{10, 20, 40\}$ (note that they are expected values, the real values only approximate them). The method followed to generate the topologies: we created a chain or ring lattices of v neighbours and then rewire other lattices with $v' < v$ until the former network with the average distance L is achieved. The Figure 2.7 illustrates some of the topologies designed using this method, for networks with 200 agents and 4 combination of $\{L, \langle k \rangle\}$.

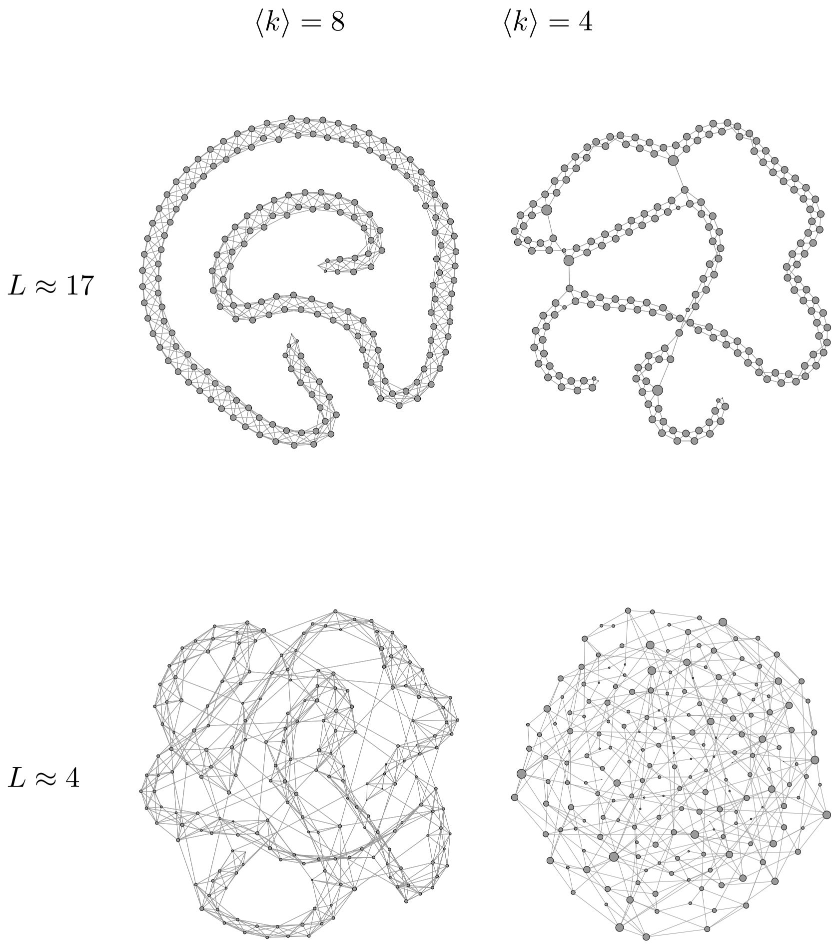

Figure 2.7: Visual representations of some of the topologies designed for chosen pairs of values $\{L, \langle k \rangle\}$

Simulation and Results

We run simulations with the success biased social learning and the 28 different topologies, each with a population of $m = 600$ agents and $n = 3$ goods. A penalty of 1 is given to the agents unable to exchange their good with one of the other goods. If the exchange is made, the penalty is reduced if the quantity gathered (Q^i) is close enough to an optimal value O^i shared by all the agents. During one timestep, the agents exchange their good 10 times before updating the values they attribute to prices. The score of the agents is given by the sum of the penalties. We show how the score is changing through time for all the simulations in Figure 2.8. As we can see in this figure, different topologies lead to different equilibrium at different speeds.

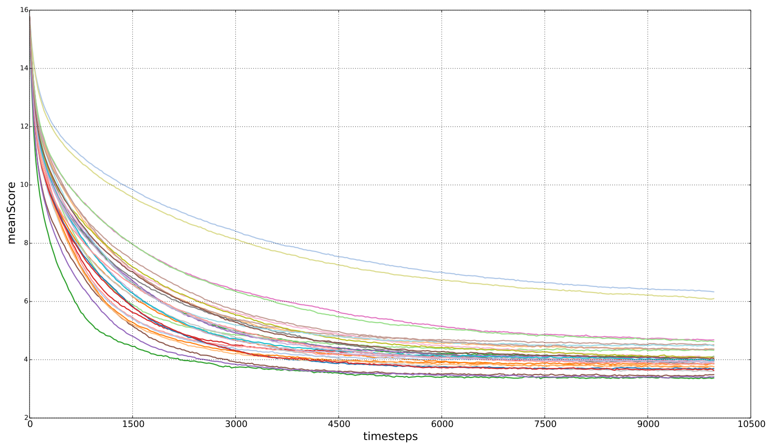

Figure 2.8: Evolution of the mean of the score of all agents during the whole simulation for all the different social network topologies.

To better explore this, we represent in Figure 2.9 the time taken for the overall mean score of all agents to reach 10% of the score expected under the ideal equilibrium. It clearly appears that the average distance L is highly correlated with the speed needed to reach the equilibrium. On the other hand, for equal values of L, there is no clear influence of the average degree ($\langle k \rangle$) in any aspects.

This simple example illustrates how crucial aspects of the network's topology impact the outcome of the model presented in this chapter. We have shown the capacity of the topology to modulate the speed to reach the equilibrium ; which suggests this to be an important factor regarding the resilience of economic systems against external shocks.

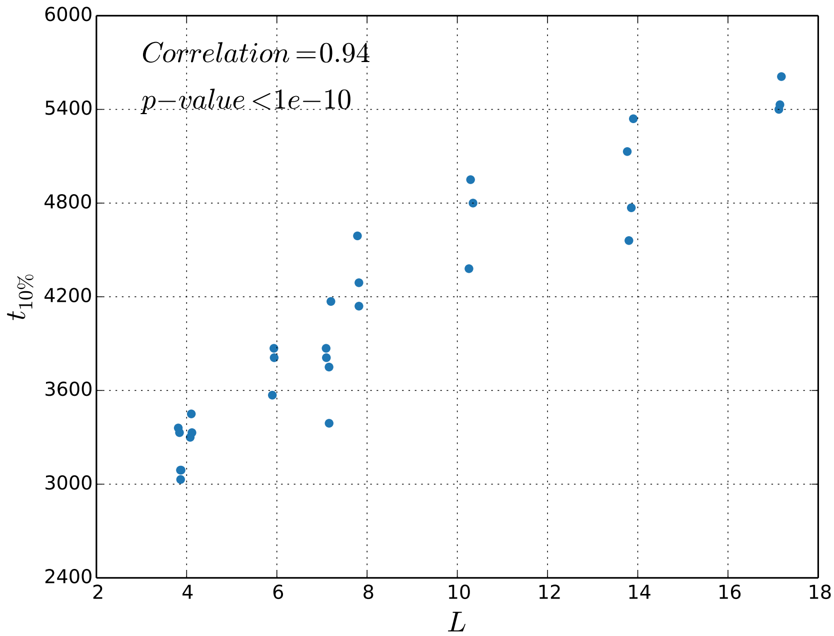

Figure 2.9: Time (in number of time step) needed to reach 10% of the score expected under the economic equilibrium, with regard to the average distance L of the social network.

Trade dynamics and cultural mechanisms are often studied separately, within the framework we presented they are brought together. This allows to perform various quantitative analysis of the influence of both side on the general dynamic of the system. Here we show how manipulating some properties of the *cultural* network can influence the speed of convergence of the trade mechanisms. In future works both cultural *and* trade networks could be explored at the same time while real networks could be used and compared.

2.5.2 Rate of Social Learning

In this section we propose to illustre Approximate Bayesian Computation methods described in Introduction.Our goal is to select among the scenarios envisaged in Section 2.1 (the neutral model versus the success bias model) the one that better reproduces the expected equilibrium and under which conditions, with the limitation that we haven’t any observation of this equilibrium to feed the Bayes equation shown in Equation 1.1, only theoretical expectations.

Fitting to Idealized Outcomes

This is why in this section we present a slightly different variation of ABC called: *Fitting to Idealized Outcomes* (FIO, *cf.* Gallagher et al., 2015), which compares the model to the output of known theoretical models, instead of empirical evidence. In the current setup the idealized output is the one reached when the prices follow the values predicted by the economic theory that describe the equilibrium we shown before. This theory is also known in economy as the General Equilibrium (GE).

The different steps of the FIO algorithm are given in Algorithm 2.4. To compute the posteriors, all simulations were run with 149 agents exchanging and producing three goods during $10,000$ time steps. As before, at each time step the agents realise 10 trades. Moreover, a way to calculate the distance to the idealize output has to be defined. This is done by the Equation 2.9. Additionally, we need to define a threshold under which a simulation is considered to be close enough to the Idealized Output. In the current setup a score of 0.25 for one agent means that this agents have succeeded on 15 exchanges other 20. Given the randomness of the model we considered this to be close enough to the general equilibrium and then set $\epsilon = 0.25$.

Algorithm 2.4 Algorithm to Fit the model to an Idealized Output (FIO).

1: sample of μ with $\mu \sim U(0, 1)$
2: run simulations with innovation rate = μ
3: compute distance S to idealized outcome (GE):

$$S = \frac{\sum_{i=1}^{n} s_i - s_{ge}}{n} \tag{2.9}$$

▷ n: total number of agents, s_i score of agent i, s_{ge}: ideal score
4: select 200 simulations with $S < \epsilon = .25$,
5: draw *posterior* distribution of μ for those simulations.

Model selection and posterior distributions

Figure 2.10 shows the distribution of S for the two different social learning strategies. For both models, 200 simulations are executed with value of μ randomly sampled between 0 and 1.

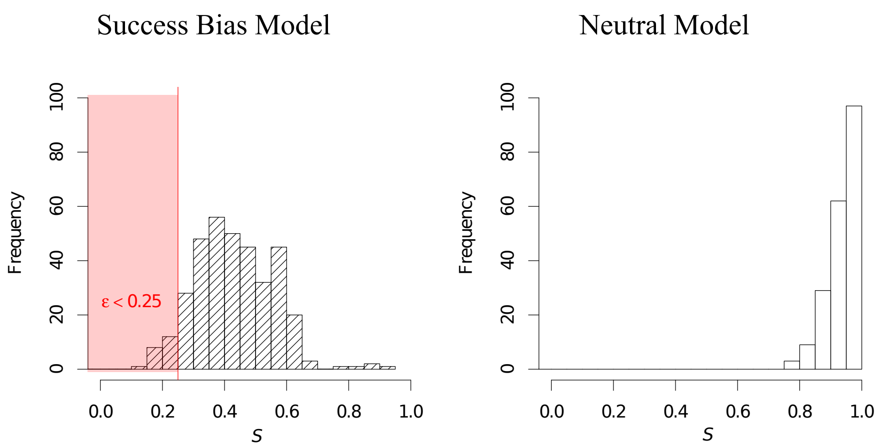

Figure 2.10: Distribution of all mean scores S for the simulations with success biased social learning (left) and unbiased social learning (right).

We see on Figure 2.10 (right) that the unbiased social learning isn't able to lead to General Equilibrium: there is no simulation under the 0.25 threshold ($\bar{S} =$

0.936), simulations are stuck in situations where almost all exchanges failed. On the other hand, a majority of the simulations that use success bias social learning (Figure 2.10 left) leads to a state where more than half of the trade are successful ($median(S) = 0.423$) with some of them where almost no trade fails ($S < \epsilon$, the red side on the left graph). From this we can say that the likelihood of the unbiased model to reproduce the General Equilibrium is null which is not the case with the success biased model. We can then conclude that success bias better reproduce the General Equilibrium than the random model.

This result isn't unexpected given what we presented in Section 2.3.2 and what can be seen in Figure 2.4; but we can now quantify precisely the values of innovation rate that have the highest probabilities to reproduce the theory by computing the marginal posterior distribution of μ. This can be done by looking at the distribution of μ in the simulations able to reproduce the GE. To have enough data to draw a precise distribution we run simulations with the success bias model until we get 200 simulations with $S < \epsilon$. Again these simulations use a μ randomly sampled between 0 and 1. The posterior distribution of μ is given by simply looking at the 200 ones that we selected, while the prior is simply the uniform distribution between 0 and 1 from which we initially sampled μ. Both distributions are represented in Figure 2.11.

Looking at the posterior distribution of μ in Figure 2.11 tells us that the innovation rate has to be lower than 0.5, which makes sense given that otherwise agents will be changing their prices almost once every two time steps, i.e.: randomly. Nonetheless, the mean of the posterior for the parameter μ is $\bar{\mu} = 0.175$, which correspond to around one change every six time steps. This is pretty high compared to, for example, the rate of genetic mutations observed in Biology. This may underline the difference of speed between biological and social processes, where social agents need to quickly adapt to complex and fast changing situation whereas the biological rhythm follows the slow pace of evolution. At the same time, this high rate of random changes would be drastically reduced if our agents were able to made non-random innovations and use self-adapting strategies which are processes closer to what is observed in real social systems.

This preliminary study illustrates how the use of simulation and model testing tools such as Fitting to Idealize Outcome, can help us to articulate our framework with other theories. In the next chapters we will see how we can go further and use empirical observations to test our models and how Bayesian inference can then be used to formally select between them and to find robust parameter intervals.

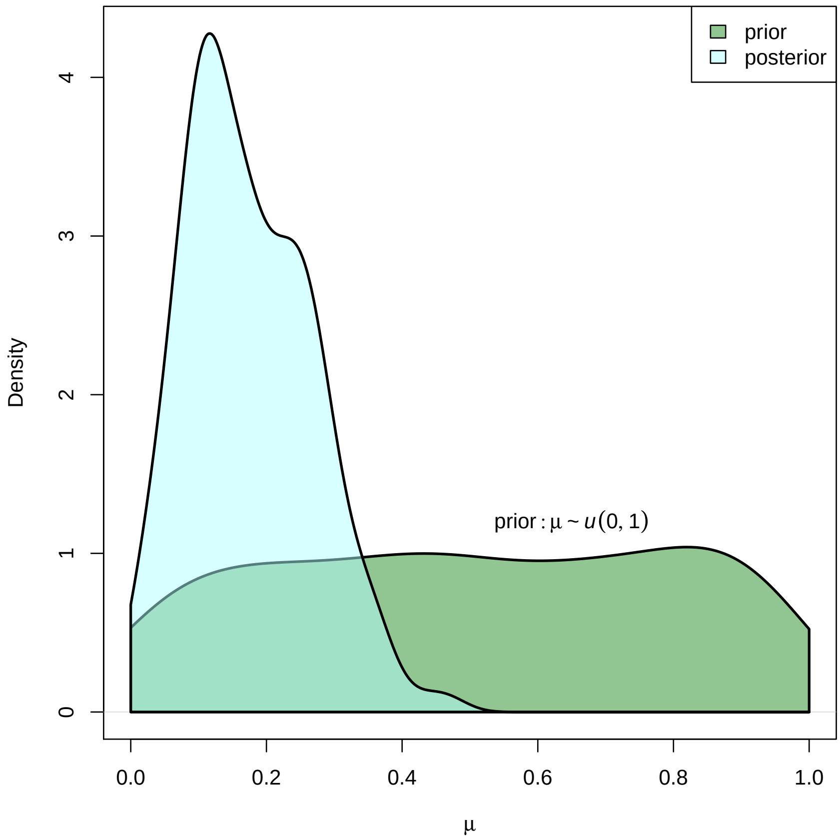

Figure 2.11: Prior (in green) and posterior (in light blue) distributions of the parameter μ for the success bias model. The posterior distribution is drawn using the distribution of μ for the 200 simulations selected by the FIO with $S < \epsilon$.

2.6 Conclusion

This chapter proposes a framework to simultaneously study cultural change and trade dynamics. The development of our framework was first aimed at simplicity which is achieved by the use of two vectors (quantity and value) and five processes (production, consumption, trade, cultural transmission and innovation). The second aim of the work conducted was to obtain a flexible framework which is possible since each of the processes can be implemented accordingly to the question studied. We have shown the validity of this approach by reproducing expected results on both the cultural and trade side. On the cultural transmission side we have shown that the implementation of a neutral model leads to the expected observations on the variants of the vector value: a power law. When implementing trading mechanisms we observe the convergence of prices to the expected values and the improvement of the score of the agents.

The implementation of a trading model within our framework successfully illustrates how economy can be viewed through the lens of cultural evolution and social learning. This aspect which has rarely been studied before, opens a new line of thinking on the weight of trade mechanisms within cultural change. Notably, we have demonstrated that the implementation of success bias leads to a distribution of prices departing from the neutral hypothesis, which is the reference in the study of cultural changes. With success biased social learning, the frequency distribution has a plateau of common prices (a number of prices share similar and high proportion), which shows that, when trade is taken into account, the most common variants are more diverse. Interestingly, we have not found references pointing at the ability of success biased social learning to keep a relatively large diversity in the frequency distribution. We suspect that this is mainly because it is not common to compare results of trading models to other cultural evolution models. It would therefore be interesting to compare the frequency distributions obtained by our model and the ones observed in current or past economies. On the side of cultural transmission, the results obtained can also be compared to the ones obtained when prestige biased cultural transmission is used. The idea behind this last comparison, is that trade could be interpreted as a particular prestige biased cultural transmission and therefore be fully integrated within the cultural evolutionary frameworks.

We ended this chapter by briefly illustrate some ways to explore this framework more cautiously. We gave an example on how the topology of the social network can be studied and another one that proposes a way to formally explore, compare and test against known theories the different social learning models we implemented in our framework, thanks to Bayes equation. In the next chapter this approach will be detailed more precisely and different social learning strategies will be tested against real data instead of idealized output.

It should be clear that the overall exploration proposed in this chapter only cover a tiny part of the properties of the model. Multiple implementations of the production mechanism can be studied so as to increase the number of goods per agent and include factors such as the meteorology or the various type of goods. More complex dynamics can be looked for. The trading mechanism can be naturally implemented in different ways, each reflecting a specific theory, and thus allowing their comparisons. Other cultural transmission mechanisms should also be tested. The trade network (which in this work can be interpreted as fully connected) can also be modified to study the effect of slow connections or the rupture of certain connections, following the exploration started with the social network. The agent themselves could become more complex and be endowed with the ability to learn behaviours which would again produce more realistic simulations. Moreover, various populations can be envisioned, each with their distinctive characteristics regarding cultural transmission and trade.

Finally, in terms of analysis of the results, multiple factors could be taken into account. The rapidity of the fluctuation of prices is a good starting point for the establishment of economical studies. For the study of a population of agent in general, it can be interesting to analyse the number of active agents and the composition of the population, or the number of goods agents accumulate during the simulation, as we will do in Chapter 4. But the rest of the thesis won’t go deeper in this theoretical exploration and will focus on proposing and testing ways to use the approach to understand and solve the problems raised by real case scenario.

Chapter 3

DETECTING THE IMPACT OF SOCIAL LEARNING BIAS IN ONLINE SOCIAL NETWORKS

3.1 Introduction

In this chapter we adapt the approach presented before to detect and quantify the importance of different social learning biases in the spread of news on Twitter, an online social platform. Recent studies have shown that false news spread quicker and farther than true news on this online platform. Our goal is to test if these differences can be attributed to different bias in social learning. Does the spread of false information speeded up by context bias when content bias slow down the spread of true information? In this chapter we propose different social learning biases to explore these dynamics of online information cascades and test if they are driven by the intrinsic content of the message, or by extrinsic value (e.g. popularity) while their intrinsic value is arbitrary.

Answering those questions will bring us to explore various aspect of the thesis: First we will present a method to articulate the theoretical framework presented in the previous chapter with empirical observations. Then, given the nature of the data, it will be the occasion to analyse how social learning bias can impact the spread of cultural information in a very special environment: online social media. This will give us a unique view of the pace and rhythm of cultural evolution on those platforms which are crucially changing the way we interact socially.

At the same time, the possibility to access the content of every messages makes this study a perfect candidate to explore the importance of direct content-dependent bias in cultural evolution at large scale. In practice, content dependent bias implies complex mechanisms at the individual levels. The complexity of those mechanisms makes the use of the tools proposed harder and we could not overcome all difficulties raised. Although we present an experiment to test the influence of content bias, the results cannot be interpreted with great confidence. Therefore, this chapter will mainly deal with the analysis of four shades of context-dependent biases and compare them together while the result of the content bias model will be discussed in a separate section.

The chapter starts by detailing the specificity of cultural evolution in online Social Media. Then we describe the five different models of social learning and the structure and pattern observed in the data we use. After this we present the method used to confront the models with the data and the problems this method raises. Then we present and discuss the result before concluding on their importance within the general scheme of the thesis.

3.2 Social Media and Cultural Evolution

Cultural evolution is undoubtedly altered by social media technologies, which impose new, often algorithmic, biases on social learning at an accelerated tempo on a vast virtual landscape of interaction. Unlike traditional societies that share in person (Danvers et al., 2019; Smith et al., 2019), sharing on social media is often not primarily kin or need-based. Important evolved psychologies, for example, such as shame and social exclusion (Robertson et al., 2018), or the visibility of social interactions involving others (Barakzai and Shaw, 2018), can be greatly altered in online social networks. One way they affect social learning is by the prominent display of social metrics (likes, shares, followers, etc.) that feed biases toward popularity and often novelty; digital social data far exceed what a human Social Brain can track without technological assistance (Crone and Konijn, 2018; Dunbar and Shultz, 2007; Falk and Bassett, 2017; Hidalgo, 2015).

Most social learning biases (Kendal et al., 2018) can be manipulated by a social media platform. Content-bias can be imposed, for example, by algorithmic filters that customize social media feeds to individual users. Context biases are also routinely imposed by social media platforms, which often prioritize popularity and recentness. Since strategies such as “copy recent success” are most competitive in fast-changing social landscapes (Rendell et al., 2010; Mesoudi et al., 2015), we might expect these context biases to flourish among social media (Bentley and O’Brien, 2017; Acerbi and Mesoudi, 2015; Kendal et al., 2018). Until recently, context-bias in online media was often underestimated. The hosts of Google Flu, for example, over-predicted influenza rates for 2013 (Lazer et al., 2014) by not accounting for context-biased learning about flu from other Internet users rather than individuals’ own symptoms (Bentley and Ormerod, 2010; Ormerod et al., 2014).

With over 300 million users worldwide, Twitter makes many social learning parameters explicit, the numbers of followers, re-tweets and likes of users and their messages. Aggregated Twitter content has previously been used for counting the frequencies of specific words across online populations, which can reveal mundane cycles of daily life (Golder and Macy, 2011), the risk of heart disease (Eichstaedt et al., 2015) and numerous other phenomena.

Subsequently, more work has been done on the dynamics of information flow online. Vosoughi et al. (2018) documented how false news travels “farther and faster” than true information among Twitter users. Here we begin to explore these dynamics by focusing on the sizes of Twitter “cascades” measured by Vosoughi et al. (2018), in terms of total number of re-tweets of messages that had been independently classified as true or false. Throughout the chapter we refer to “true” versus “false” rumours in their data, but it should be noticed it doesn’t correspond to an exact sample of the “true” and “false” rumour on Twitter. Instead, “true” and “false” here correspond to confirmed fact-checked rumours versus rumours that were debunked, respectively—see Vosoughi et al. (2018). We expect that a huge proportion of obviously true rumours are not confirmed[1] and thus doesn’t appear in the dataset, whereas obviously false rumours will tend to be more systematically debunked.

As we did on the previous model we will rely on the unbiased model as a null model, where Twitter messages are “neutral”, i.e. copied randomly with negligible bias, until proven otherwise. Models of unbiased (neutral) copying

[1] As Alberto Acerbi put it (I thank him for the example): “ Nobody would check […] whether Donald Trump has been elected president of the US”.

are well established in cultural evolution research (Acerbi and Bentley, 2014; Bentley et al., 2011; Neiman, 1995; Premo, 2014; Reali and Griffiths, 2010) and have also been applied to social media (Gleeson et al., 2014). In this approach, we assume that the probability of a message being re-tweeted depends only on the current frequency of the message and not on its content. Content bias would be identified when the model is falsified (Acerbi et al., 2009; Acerbi and Bentley, 2014; Bentley et al., 2004, 2007, 2014a; Kandler and Shennan, 2013; Mesoudi and Lycett, 2009; Neiman, 1995).

Unbiased copying models have been calibrated empirically against real data sets that represent easily-copied variants, such as ancient pottery designs (Bentley and Shennan, 2003; Crema et al., 2016; Eerkens and Lipo, 2007; Neiman, 1995; Premo and Scholnick, 2011; Shennan and Wilkinson, 2001; Steele et al., 2010), bird songs (Byers et al., 2010; Lachlan and Slater, 2003), English word frequencies since 1700 (Ruck et al., 2017), baby names (Hahn and Bentley, 2003), and Facebook app downloads (Gleeson et al., 2014). The time scales of these studies range from centuries to decades, months or days.

The two most important parameters of unbiased copying models are population size, N, and the probability, μ, of inventing a new variant (Hahn and Bentley, 2003; Neiman, 1995).

In order to compare different models in explaining the data, models need to be generalized, through multiple parameters, to generate as many outcomes as possible, taking into account sampling and different possible biases in cultural transmission (Kandler and Powell, 2018). The goal is to estimate probability distributions of those parameters to explain the set of posterior distributions.

To study the dynamics of Twitter cascades, we test several different models of Social Learning Strategies. These models can be compared based on their ability to replicate the data while minimizing the number of model parameters. Once we have determined the best model, we then estimate a probabilistic range of each parameter values to best fit each data set. The goal is to compare parameter ranges between different data sets.

Kandler and Powell (2018) advocate the use of Approximate Bayesian Computation (ABC), which can produce a probabilistic representation of parameter space that shows how likely the parameters are to explain the data. ABC allows models to be compared using Bayesian Inference to estimate the probability that a model explains the data (posteriors) given existing knowledge of the system (priors); models are often compared using likelihood ratios.

Using this approach on the Twitter data explored by Vosoughi et al. (2018) we can select the model of social transmission that best reproduce the observation. Moreover, as prior information on Twitter users is available, we hope to determine with precision the distribution of biases at the individual level in the population of Twitter users. This opens the possibility to precisely understand the process explaining *how* the observed differences appear (Vosoughi et al., 2018).

3.3 Models of re-tweeting

3.3.1 Neutral model

Here we consider the same Neutral model that we presented in Chapter 2, this time with a population of N Twitter users in a fully-connected network. The number of modelled Twitter users, N, is kept constant, as we will assume the modelled time period (days or weeks) is short compared to any growth in number of users.

Each Twitter user observes N randomly-selected other users in each time step. In this population, users either tweet something unique of their own or else re-tweet another message. At time t, each of the re-tweeting agents chooses randomly among the N agents, and either re-tweets that agent's message, with probability $(1 - \mu)$, or else composes an original new tweet, with probability μ. We run the model until reaching a steady state (for $\tau = 4\mu^{-1}$ time steps, see Evans and Giometto 2011). In this basic unbiased copying model, a re-tweet is chosen from among N other agents, as opposed to choosing from the different Tweet messages themselves. This means some/all of the N Twitter users may be re-Tweeting the same message. The number of *different* messages observed by each user is typically much less than N.

The implementation of this model is given by Algorithm 3.1. ,here:

1. $randSel(a)$ return an element of the vector a
2. $genNewTweets(n)$ return a vector with n new tweets that haven't been tweeted yet.

3.3.2 Context Biased models

Next, we modify this unbiased model to introduce context-biases through three different forms of popularity bias. The first is a frequency bias, where the prob-

Algorithm 3.1 Neutral model with random selection.

Input: N: number of users,
μ: probability of generating a new tweet,
$timestep$: n. of time step
Output: A table with the number of re-tweet for each tweet
function neutral($N, \mu, timesteps$)
$tweets \leftarrow matrix(timesteps \times N)$
$tweets[0,] \leftarrow genNewTweets(N)$
for t in $timesteps$ **do**
  for n in N **do**
    $X \sim U(0,1)$
    if $X < \mu$ **then**: $i \leftarrow genNewTweets(1)$
    else: $i \leftarrow randSel(tweets[t-1,]g)$
    $tweets[t,n] = i$

ability of a message being copied increases with frequency above the inherent frequency-dependent probability of the neutral model itself.

As social media feeds often highlight “trending” messages in some form, the other two versions represent “toplist” biases, in that Twitter users are biased towards the top y (where y is the size of a “trending” list) most popular messages (Acerbi and Bentley, 2014).

The first context-biased model derives from a more general model of discrete choice with social interactions (Brock and Durlauf, 2001; Bentley et al., 2014b; Brock et al., 2014; Caiado et al., 2016). A parameter β represents the overall magnitude of the biases. Another parameter, J, represents context-bias, specifically popularity bias here. For the population with global β and J, a simple representation (Bentley et al., 2014b; Brock et al., 2014; Caiado et al., 2016) for the probability, P_i, that a Twitter user re-tweets the message i is:

$$P_i = \frac{e^{\beta[U_i + Jp_i]}}{\sum_{i=1}^{k} e^{\beta[U_i + Jp_i]}} \tag{3.1}$$

Here the term U_i would denote the intrinsic payoff to choice i, which here could be the ’attraction’ (Acerbi, 2019) of message i for (re)Tweeting. Here we simply assume the messages have no intrinsic utility, i.e., $U_i = 0$ for all messages, i.

This yields:

$$P_i = \frac{e^{\beta J p_i}}{\sum_{i=1}^{k} e^{\beta J p_i}} \tag{3.2}$$

In this case the context-bias, Jp_i is based upon the popularity, p_i.

Note also that when $\beta = 1$ and/or $J = 0$, the model reduces to a random guess model, where each choice has equal probability regardless of its frequency , i.e., $P_i = 1/k$ for all choices, i. By contrast, under the neutral (a.k.a. random copying) model, the expected frequency of each future choice is predicted by its previous frequency. Equations 3.1 and 3.2 do not reduce to the basic neutral model in a simple way; the copying is neutral in this sense only for particular momentary combinations of β and $p_i(t)$.

The implementation of this conformist bias is given by Algorithm 3.2, with:

1. $biasselect(a, \beta)$ a function that returns an element i of the vector a of length n with a probability defined as:

$$\forall i \in a, P(i) = \frac{e^{\beta \times freq(i)}}{\sum_{j}^{n} e^{\beta \times freq(j)}}$$

Next, in our “Top threshold” model, Tweets are exhibited in a “top list”, such that a parameter C determines the fraction of individuals that will re-tweet a message from this list of the top y trending Tweets in the population (Acerbi and Bentley, 2014). The other $1 - C$ fraction of the population will re-tweets something else at random, per the neutral model.

The implementation of this Top threshold model is given by Algorithm 3.3 with:

1. $toptraits(a, b)$ a function that returns the bth more frequent elements of a

Our “Top Alberto” model, named for its inventor (Acerbi and Bentley, 2014), is a slight modification. At each time step in the Top Alberto model, a fraction, $C(1 - \mu)$, of the Twitter users compare the rank of their own tweet, to y most popular tweets, and the user only re-tweets something else if their tweet is not already on the top y list. The remaining fraction of individuals $(1 - C)(1 - \mu)$ re-tweet as in the basic neutral model. Note that a new set of ‘conformist’ agents, represented as a fraction $C(1 - \mu)$ of the population, are randomly selected at each time step.

The implementation of this Top Alberto model is given by Algorithm 3.4.

Algorithm 3.2 Random model where probabilities depend on tweets frequencies.

Input: N: number of users,
μ: probability of generating a new tweet,
$timestep$: n. of time step
β: weight of conformist bias,
Output: A table with the number of re-tweets for each tweet
function conformist($N, \mu, timesteps, \beta$)
$tweets \leftarrow matrix(timesteps \times N)$
$tweets[0,] \leftarrow genNewTweets(N)$
for t in $timesteps$ **do**
  for n in N **do**
    $X \sim U(0, 1)$
    if $X < \mu$ **then**: $i \leftarrow genNewTweets(1)$
    else:
      $i \leftarrow biasselect(tweets[t-1,]g, \beta)$
    $tweets[t, n] = i$

Algorithm 3.3 Top Threshold.

Input: N: number of users,
μ: probability of generating a new tweet,
$timestep$: n. of time step
C: percentage of conformists,
Y: n. of tweet in the top list
Output: A table with the number of re-tweet for each tweet
function conformist($N, \mu, timesteps, \beta$)
$tweets \leftarrow matrix(timesteps \times N)$
$tweets[0,] \leftarrow genNewTweets(N)$
for t in $timesteps$ **do**
  for n in N **do**
    $X, X^* \sim U(0, 1)$
    if $X < \mu$ **then**: $i \leftarrow genNewTweets(1)$
    else
      if $X^* < C$ **then**
        $i \leftarrow randomselect(toptraits(tweets[t-1,], Y))$
      else
        $i \leftarrow randomselect(tweets[t-1,])$
    $tweets[t, n] = i$

Algorithm 3.4 Top Alberto.

Input: N: number of users,
μ: probability of generating a new tweet,
$timestep$: n. of time step
C: percentage of conformists,
Y: n. of tweet in the top list
Output: A table with the number of re-tweet for each tweet
function conformist($N, \mu, timesteps, \beta$)
$tweets \leftarrow matrix(timesteps \times N)$
$tweets[0,] \leftarrow genNewTweets(N)$
for t in $timesteps$ **do**
  for n in N **do**
    $X \sim U(0, 1)$
    if $X < \mu$ **then**: $i \leftarrow genNewTweets(1)$
    else
      if $X < 1 - C$ **then**
        $i \leftarrow randomselect(tweets[t-1,]g)$
      else

        if $tweets[t-1, n] \in toptraits(tweets[t-1,], Y)$ **then**

          $i \leftarrow tweets[t-1, n]$
        else
          $i \leftarrow randomselect(tweets[t],)$

3.3.3 Content Biased model

Our last model, named "Cascades 3D" as it stores every cascades and captures all their three dimensions (*cf.* Vosoughi et al., 2018), implements the content bias part of Equation 3.1.

Algorithm 3.5 Cascades 3D.

Input: N: number of users,
μ: probability of generating a new tweet,
$timestep$: n. of time step
IC: Initial number of cascades
R: number of different unique rumours.
$\beta[N]$: an array defining the β value of each agents
$U[R]$: an array defining the utility for each rumour
Output: A 3D table with the metrics of each cascades
function cascade3D($N, \mu, timesteps, \beta$)
$poolCascades = generateCascades(IC, N, R)$
$tweets[0,] \leftarrow genNewTweets(N)$
for t in $timesteps$ **do**
  $alltweets = getAllTweets(poolCascades, t)$
  for n in N **do**
    $X \sim U(0, 1)$
    if $X < \mu$ **then**
      $newCascades \leftarrow genNewCascades(1, n, R)$
      $poolCascades.add(newCascades)$
    else
      $tl_n = samplePossible(alltweets, n)$
      $re - tweets = selectTweets(tl, n, \beta[n], u[tl])$
      $poolCascades.update(re - tweets)$

We already highlighted the main difference between the content-bias and the others context-dependent bias in the introduction: in the former, the intrinsic properties of the transmitted information will have an impact on the way this information is transmitted, whereas in the latter, only contextual, frequency-dependent aspects will affect this transmission. This will have a first important implication in a content-bias model: it has to encode this intrinsic information. This imply modelling the way this information is introduced and how it affect the transmission.

To do so we start again from Equation 3.1, with the term U_i that denotes the intrinsic payoff to choice i, and a term β, that denotes the bias of the agent toward this utility. Here we assume $Jp_i = 0$, meaning agent won't be biased by the context but only by the content U_i of the tweet. This yields:

$$P_i = \frac{e^{\beta U_i}}{\sum_{i=1}^{k} e^{\beta U_i}} \tag{3.3}$$

This probability is computed by the function $probaSelection()$ that implements Equation 3.3. This function is available in the package provided with this chapter and is called by the function $selectTweets()$ described at line 12 in Algorithm 3.5.

Given Equation 3.3, users with $\beta < 0$ will tend to re-tweet more tweets with $U_i < 0$, users with $\beta > 0$ will tend to re-tweet more tweets with $U_i > 0$ and users with $\beta = 0$ will be neutral and re-tweet all tweets with equal probabilities. We reproduce this in Figure 3.1. As one can see from this graph, users with a $\beta = -100$ (violet line), won't re-tweet any tweet with $U > -.75$ when faced with 10 different tweet of utility $U \in [-1, 1]$ while they will re-tweet with a probability of almost one a tweet with $U = -1$. On the other hand, a user with $\beta = 0$ (the green line on the graph) will have the same probability (0.1) to re-tweet any tweet independently of its utility.

This, implicitly, introduce two new important aspects absent from the previous models: heterogeneous populations of tweets with different utilities (U_i) and heterogeneous populations of Twitter users with different attraction toward those utilities (β). Those two things are represented respectively by the arrays $U[R]$ and $\beta[N]$ in Algorithm 3.5 and given *a priori*, as input to the model. If they could be informed by the data available, we will see how they can be explored using Bayesian Inference to find their shape. This will be detailed in section 3.5.3.

Aside those new elements, this model has been built to stick more closely to Twitter's underlying architecture. As we'll see, this adds layers of complexity that makes the extensive analysis of the model difficult. One thing is the potentially infinite Twitter memory: every Twitter user can re-tweet any tweet available during the whole Twitter's history, while he cannot re-tweet tweets he already re-twitted. Thus we need to store the activity of every Twitter users, as well as the whole history of tweet and re-tweet.

Even trickier: while scrolling through Twitter history, a user can be exposed to the exact same tweet multiple time. If one user follows multiple different users that re-tweeting the same initial tweet, this user will access to every "leaf" of the

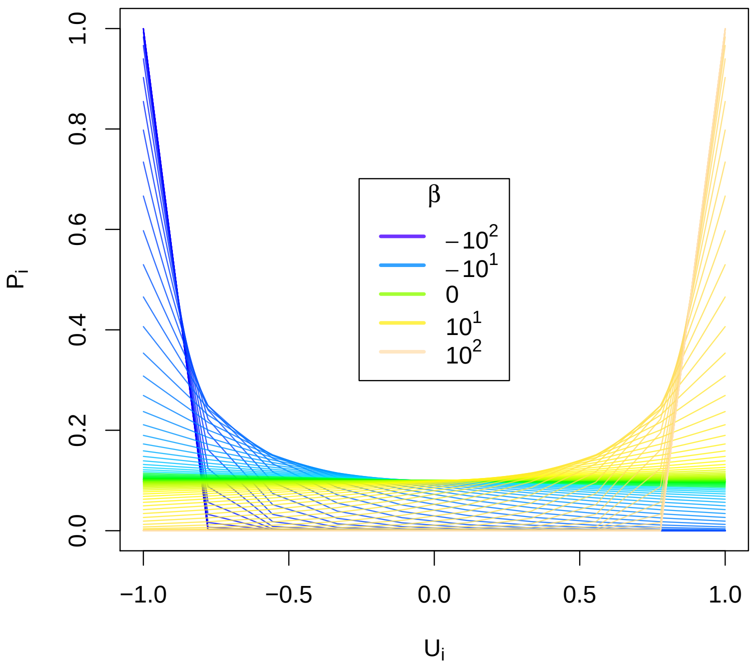

Figure 3.1: Illustration of the function 3.3 when agents with different β have to choose between 10 tweets with utilities regularly spaced between -1 and 1.

cascade of re-tweet[2] and thus being offered multiple time the opportunity to re-tweet the same original tweet. The order of apparition of those tweets will depend on the time when the other users re-tweeted the original one. Then the user can choose which one he will re-tweet in turn. This will add one re-tweet to one of the branch of the cascade, potentially broadening or deepening it, depending on which re-tweet he choose to re-tweet. This call for the need to store the full history of tweet and re-tweet, with a precise timing of who re-tweeted what, from who and when.

Those mechanisms greatly increase the computation time and the storage needed to execute the model. This makes the content-bias model order of magnum slower and heavier than the other models. If one can still reasonably run and test the content-bias model at the scale of one or two simulations, this quickly becomes a problem for the analysis we want to carry here. Indeed, as we will see in Section 3.5, Approximate Baysian Computation, the method we use in this chapter, relies on millions of simulation to compute the likelihood of the model. This makes the exploration of this model difficult. We will see in the next chapter a different way to implement ABC that optimize this process and that should solve the problems we will raise here.

In summary, our models for re-Tweeting activity are

- Unbiased copying
- Conformist copying
- Top Threshold copying
- Top Alberto copying
- Cascades 3D

All models implementations are available in the R-package here: github.com/simoncarrignon/spreadrt

While they do not span the space of all possible models, they require a rigorous means of discrimination when compared to the data. To do so, we use Approximate Bayesian Computation (Kandler and Powell, 2018).

[2] A cascade is a series of re-tweet started by a first original tweet, that can be describe as a tree, for more information about cascades of re-tweet and how they can be measured, see Vosoughi et al. (2018).

3.4 Data

Vosoughi et al. (2018) analyzed a set of 126,000 news stories distributed on Twitter from 2006 to 2017. These stories were (re)tweeted 4.5 million times by approximately 3 million Twitter users. These news stories were classified as true and false using multiple independent fact-checking organizations, which were over 95% consistent with each other and further confirmed by undergraduate students who examined a sample of these determinations (Vosoughi et al., 2018).

For each different message in this dataset, Vosoughi et al. (2018) measured how many Twitter users re-tweeted the message, called the “cascade size”, which we model here. Tweet cascades started by bots were not a significant factor in these data (Vosoughi et al., 2018). Vosoughi et al. (2018) also measured other dimensions of these Twitter cascades, included the “depth”(number of re-tweets by a unique new user from the origin tweet), breadth (maximum users involved in the cascade), and time elapsed over the cascade.

Since Vosoughi et al. (2018) counted multiple cascades for certain messages, we binned identical message cascades together, yielding a group of cascade distribution for each message (Figure 3.2). This reduces the dimension of the data set to one unique distribution and avoids the need to keep the full structure of each cascade. This is essential for the current study as it allows to apply crucial simplifications to the models in order to be able to directly compare them to the data collected by Vosoughi et al. (2018) (*cf.* the limitation we described in Section 3.3.3).

Figure 3.2a shows the distribution of cascade sizes when each cascade is taken separately, Figure 3.2b shows aggregate cascade sizes where we have aggregated the number of re-tweets for cascades of identical messages.

3.5 Approximate Bayesian Computation

Here we use Approximate Bayesian Computation (ABC) to calibrate our models against Twitter data. The aim is to find the distribution of parameters of each model knowing data distribution, i.e. the posterior distributions of the model and select the model with the highest probability to explain the data. We described this methods in Introduction and will here present one way to implement it.

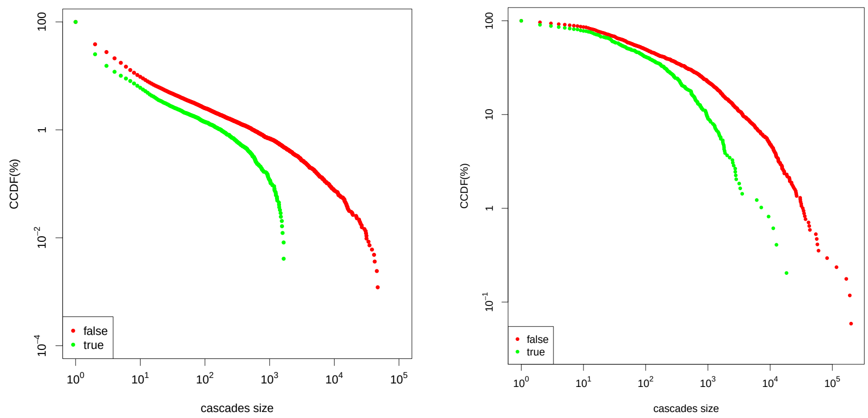

(a) (b)

Figure 3.2: Complementary Cumulative Distribution Functions (CCDFs) of (a) non aggregated cascades vs. (b) aggregated cascades with same rumours. Those CCDFs represent the percentage of respectively (a) cascades and (b) rumours that have reached a given number of re-tweets between 2007 and 2017.

3.5.1 Rejection Algorithm

For this chapter we adapted the ABC version used by Crema et al. (2016), the rejection algorithm, that can be summarized by the simple procedure described in Algorithm 3.6.

Algorithm 3.6 ABC: Rejection Algorithm.

```
sample of θ_i from the priors
run simulations : x_i = model(θ_i)
compute distance δ(x_i, d)
reject θ_i if δ(x_i, d) > ϵ
```

This algorithm sample parameters for the model from the prior, run simulations using those randomly chosen parameters and compare the result of the simulation with the data. This will be repeated until enough good simulations are found then the list of parameters that gave those good simulation is kept. This list of parameter is our approximation of the joint posterior distribution of the model given the data.

The prior will be the distribution used to draw the parameters. Before listing those priors, we will spend time to describe in details how we have chosen x_i and d, the summary statistics that represent respectively the results of a simulation and the data ; and the distance function $\delta(x_i, d)$ that will use those representation to compare them.

3.5.2 Summary statistics and distance to data

One central element of Approximate Bayesian Computation is the function that computes the distance between the simulation and the data (δ in Algorithm 3.6). It is this function that measures if a simulation and the parameters associated to it, are *close enough* to the data. Thus, this is what ultimately allows to draw a good approximation of the posteriors distribution (or bad if the distance is wrongly chosen). In order to calculate this distance, the function relies on a summary statistic of the empirical observation and of the result of the simulation (respectively d and x in Algorithm 3.6) that also needs to be carefully chosen.

In the previous chapter we presented a prototype of ABC where it was possible to calculate a simple distance between our simulations and an ideal, theoretical score. As this theoretical score is set by a theory, it is easy to formulate it in a way that perfectly align with the simulations. In the current study, finding the right distance isn’t as straightforward. The empirical observations, shown in Figures 3.2a & 3.2b, contain wide range of points not normally distributed and we want our simulations to reproduce the overall shape of this distribution, not just a summary or truncated version of it. We thus need to find the right way to compare this complex empirical distribution with the distributions generated by our simulations.

How to choose the right distance function?

Various methods are available to measure how two distributions are similar. An interesting one, borrowed from information theory, is the Kullback-Leibler Divergence (KL Divergence). This divergence measures how much information need to be added to one distribution to find the other. As information use the log of the frequencies, the KL divergence is less impacted by the non-normal shape of the data points. But as it is calculated as a ratio, the two distributions need to be defined on the same intervals (we cannot have a null probability). Here it means that for each possible cascade size, if the data or the simulation have more that one cascade of a given size s, the other distribution should also have at least one cascade of the same size s. If this is usually true for the smallest cascades (i.e. there is always more than one cascade of size 1, 2,3, or even 10, in the empirical data *and* in our simulations) it barely happens for bigger cascade sizes.

One solution is to calculate the distance only on intervals were both distributions are defined and remove the range where they are not. But this remove huge chunks of data where crucial differences may be. Another solution is to add one to all the possible sizes (this is often called Laplacian smoothing). But given the data we have this greatly bias the original distributions by introducing large cascades where no cascade existed before.

Another simple distance function can be designed by binning the size of cascades in a finite set of m categories. Given this we can calculate the euclidean distance between the two vectors of size m (one for the simulation and one for the data). The Kolmogorov-Smirnov statistic is also another common way to compute a distance between two distributions, without imposing much constraint on the distribution compared.

In the R-package provided with this chapter all those functions are implemented and scripts to test them are proposed in the vignette: abc_spreadrt.html. We give examples of how these functions impact the score of the simulations in Figures 3.3 and 3.4.

But the question remains: why prefer one method over the other? The answer obviously depends on what you want to do. And here we face one of the limit that transdisciplinarity imposed to this thesis. Ideally, the way to summarize and compare simulations with the data should enable us to answer the ultimate question of the study the ABC is used for. In this chapter we want to know what social learning biases can explain the differences observed in Figure 3.4. To know how different distance functions and summary statistics will change the answer to this question is a complex task and a central and well recognized problem in the ABC literature. Statisticians and Computer Scientists have started to propose solutions (Nunes and Balding, 2010; Fearnhead and Prangle, 2012) but these methods are often theoretical and haven't been tested extensively in real scenarios. Adapting them to different case studies isn't straightforward. In practice one still have to face a three sided problem:

1. Empirical: how to summarize and use the data without loosing important information that will be needed to answer the empirical questions.

2. Theoretical: how to articulate data and simulation with a distance function that enable the computation of the right posterior distribution.

3. Technical: how to implement and handle in a computationally and technically feasible way all the above, at scale that will involve complex hardware and software architecture.

The articulation of those three points is crucial and very problem dependent, making the whole process hardly generalisable. Moreover, each point mobilises very specific and advanced knowledges, from various scientific fields. From Psychology and Sociology for the point (1) to Probability and Statistics for the point (2), passing by different and distant subfields of Computer Science for the point (3). Ideally each point should be left to an expert of the domain, able to speak with the two others. Sadly it’s often not possible and one has to select the right trade-off relying only on her partial knowledge of each domain.

This is what we had to do here. Given that my formation didn’t allow me to deeply explore analytically the theoretical implications of different mathematical formulation on the probabilistic outcome of the study, I went for what I call an “empirical” test of the different approaches. The main idea is to start by the third point: implement a working solution and try different (1) and (2) until finding one that satisfies us. We will briefly outline how we did it and the problems it raise in the next paragraphs.

“Empirical” exploration of distance functions

Empirically exploring the impact of different functions has huge limitations: it is very slow and heavy. One has to run dozens of different experiments with different experimental setups to find the good candidate that will fit his needs.

The most naive way to do that is to run a full ABC, keeping as much information as possible during the simulation, and test different distance functions and summary statistics afterwards. Aside the fact that we may not know *a priori* what we will need and thus what we should keep and how, even the simplest approach became quickly impossible. Let see this with the current study. The most straightforward way to explore this scenario would be to keep all the cascades from all the simulations and then apply different distance functions to see how they modify the posterior obtained. The number of cascades for each simulation is given by $N + N \times \mu \times (t_{step} + \tau)$ which is $400,000$ given the prior used for parameters N, μ, t_{step} and τ (*cf.* next section, it can be much less or much more, but this is the expected value given the median of the priors). If we just store the size of the cascade this means to keep a vector of $400,000$ integers which is around 1.6Mo. If we want to store the full cascade we need to store the list of edge that represent who re-tweeted who and when. The number of edge is given by the total number of re-tweet during this simulation: $N + N \times (1 - \mu) \times (t_{step} + \tau)$ which is often $1,400,000$ given our prior. One edge of a cascade is represented by three integers: the id of the source, the id of the target and the time of the re-tweet. This leads to $3 \times 1,400,000 = 4,200,000$ integers (about 16Mo) to be stored if we want to keep track of all cascades of one simulation.

The space needed to keep the full cascades is at least 10 times bigger than the space needed to keep only the size. This is another important aspect that motivated the decision we took in Section 3.4 to explore only the size of the aggregated cascades. Still, even by keeping only the 1.6Mo needed to store the vector of integer with all the sizes this had to be kept for *all* simulations.

And this is the major problem of the rejection ABC presented before: it needs a huge number of those simulations (one million is often seen as “few” simulations). Even if, for the sake of the current example, we suppose that $100,000$ is enough to test the influence of the distance function, this means that around 160Go of data will be generated. 160Go that you will have to dig to recalculate different summary statistics and distances before finding the one you need. This is far from being something you can do at will, forget about loading it in memory to play around with on the fly. And this, let’s not forget it, does not produce the final results but simply to select the right distance function.

Nonetheless, given different constraint imposed by the project, this is the solution we chose to select the right tools; and we will use a similar approach in the next chapter. Here we present a first and simple example of such exploration. We ran 500 simulations using simple prior distribution chosen to run the model quickly. We stored all the distributions of the 500 simulations and applied various distance functions. The results are represented in Figure 3.3.

As we can see in Figure 3.3, some functions are more sensible to different part of the curve of the original data. From what we see here, the QS and KS functions look like the ones that better match the full shape of the distribution of the data. The QS function is an euclidean distance between the 100 percentil of both distributions and is detailed in equation 3.4. The KS function is the Kolmogorov-Smirnov distance as implemented by the R Core Team (2017) which uses Marsaglia et al. (2003)’s methods to calculate it.

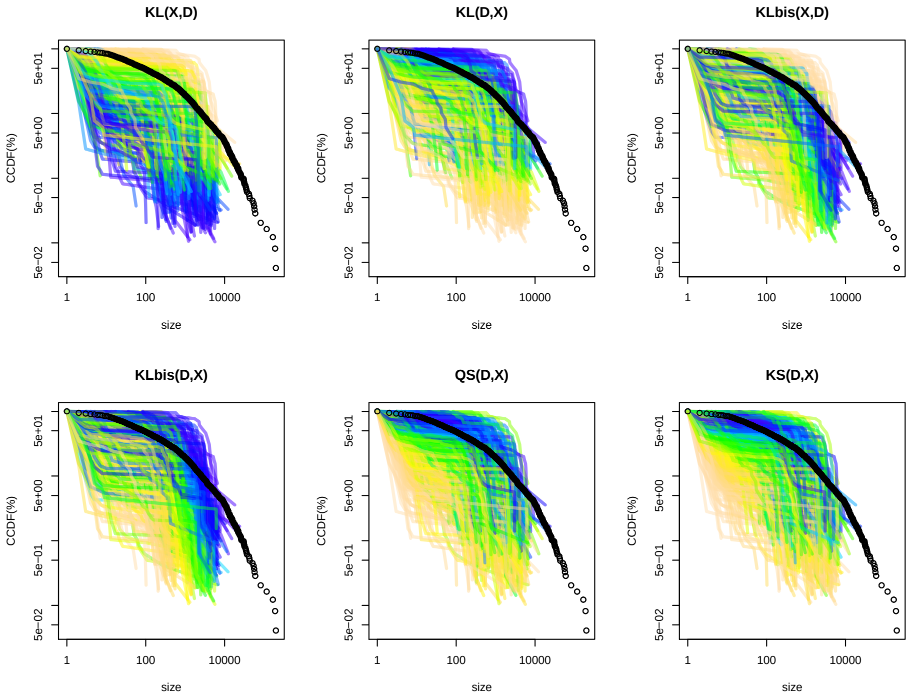

Figure 3.3: Comparing the effect of different distance metrics. Each plot represent the result of 500 simulations of the content bias model (the coloured lines) against the data from Vosoughi et al. (2018) (the black dots). The colours on each panel represent the distance to the data computed by different distance function. The colour scale used is a topographical scale, where the darker blues represent the distribution considered as the closest to the real data by the function used, while the paler yellows are measured as the farther.

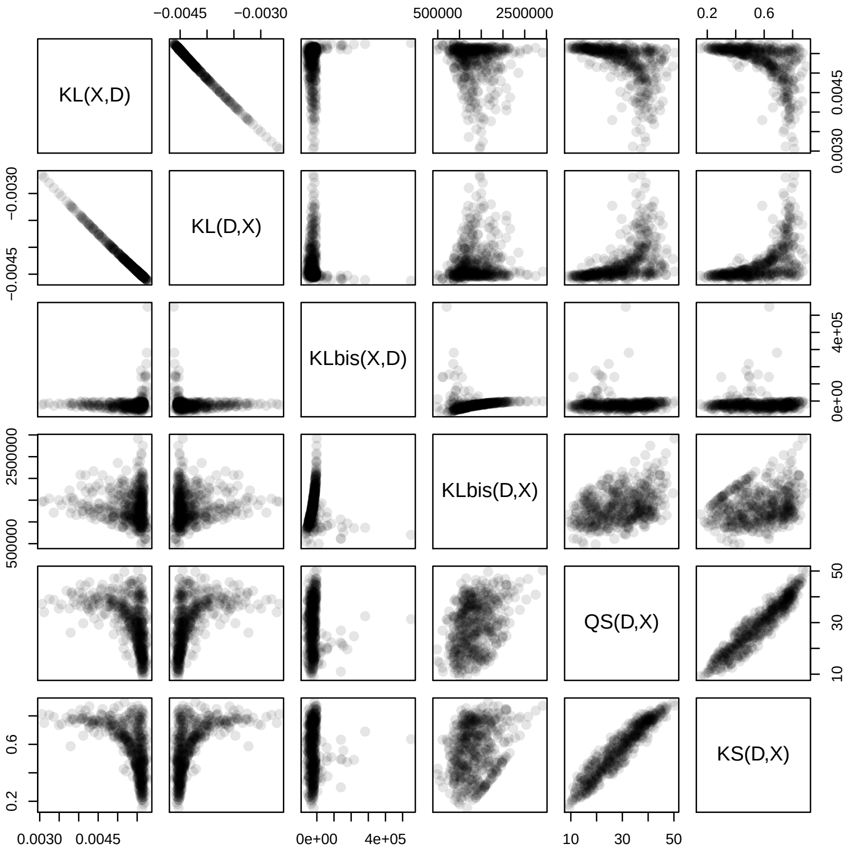

Figure 3.4: Comparing how different scores are correlated. In each panel the score of 500 simulations ran with the content-bias model is calculated given two different methods and compared.

parameters	priors
μ	$U(0,.3)$
N	$U(1000,10000)$
t_{step}	$U(10,300)$
Y	$U(0,1000)$
C	$U(0,1)$
β	$U(-2,2)$

Table 3.1: Value of priors for parameters inferred by the ABC, parameters Y and C only apply for the topfive models, beta for the conformist model.

To see how the different functions are related one with each other, we also plot in Figure 3.4 the joint distribution of all score two by two. It clearly shows that KS and QS are highly correlated. Thus choosing one or the other shouldn't impact that much our results. After some benchmarking it appears that the QS function run slightly quicker that the KS one (around 1.43 times quicker). All this led us to use this function QS as our distance function throughout the rest of this chapter. The function QS is defined as:

$$\sqrt{\sum_{i=1}^{100}\left(Q_i\left(log(S)\right)-Q_i\left(log(D)\right)\right)^2}, \tag{3.4}$$

where $Q_i(X)$ is the ith Percentile of the sample X. S is the sample generated by the simulation, D the data.

The implementation of all other scores is available in the R-package provided with this chapter and the full code used to run the empirical exploration is available in the vignette: abc_spreadrt.html.

3.5.3 Priors

We designed a list of priors for each parameters explored via ABC. This list is given in Table 3.1. The choice of the value has been motivated by knowledge from previous studies on neutral and frequency-dependent Social Learning Strategies (Bentley et al., 2004; Acerbi and Bentley, 2014) and the need to explore in details the parameter space of the various model analysed.

The content bias model involve a number of new mechanisms that are not present in the other models. Those new mechanisms need new parameters that

parameters	priors
μ	$U(0, .035)$
N	$U(15000, 30000)$
t_{step}	$U(50, 300)$
IC	$U(500, 2000)$
R	$U(1500, 2000)$
$\beta[N]$	$distrib([-100, -10, 0, 10, 100])$
$U[R]$	$distrib([-1, 0, 1])$

Table 3.2: Value of priors for parameters inferred by the ABC for the content model.

don't apply and cannot be translated within the other models. Moreover, as we mentioned it at various points, the computation and storage cost of the content model are much larger than the other. This prevents us to explore prior range as wide as we did for the other models, this would generate simulations lasting too long to be feasible. Thus, even for the parameters that can easily be translated within all models (e.g. the μ parameter), we defined a totally different and separate set of priors, used in a separate experiment designed to explore partially the content-biased model in a reasonable time.

One of the most interesting part of this model is the introduction of heterogeneous population of tweets and users. This has been done by introducing the parameters $\beta[N]$ and $U[R]$. Knowing the intrinsic attractiveness of all users toward different utilities is impossible. At the same time, and even if knowing the content of every tweet is in theory possible, it was not the case in the dataset we had access to so we could not use information from the data to initialise these priors. Moreover, it is not very clear what generates the attractive (or repulsive) character of a tweet among the information it transmits. In Vosoughi et al. (2018) they focus on the veracity of the information, but many authors argue that the emotional content of the tweet is the one that leads the spread of news online (Acerbi, 2019). Thus we decided to use priors as less informative as possible, in order to leave to the ABC the role of finding the right posterior distribution for those parameters.

To generate those uninformed priors we designed a function $distrib(T)$. This function takes a vector T of size c and associates to each element of T a proportion, in such a way that the sum of all proportion is equal to one. At the same time the function takes care that any category has the same probability to be represented in any proportion of the whole population. Phrasing it differently: if tweets can be split in three categories of $U \in \{-1, 0, 1\}$, after generating the prior we have the same probability to find an experiment where 90% of the tweets have $U = -1$, than to find an experiment where 90% of the tweets have $U = 0$ than to find an experiment where 90% of the tweets have $U = 1$. The same is true for the class of β at the user level. Given this function, the more likely distribution is the one when the population is evenly divided between the different classes (i.e.: $P(U = 0) = P(U = 1) = P(U = -1) = 0.33$).

The introduction of those new parameters allows us to explore which combinations of content and users may favors the spread of online news as Vosoughi et al. (2018) described it. Moreover, our model doesn’t assume any strict separation between true and false news but propose a more finely defined notion of attractiveness even within true and false tweet, which may be more inline with other theory of spread of online content (Acerbi, 2019).

3.5.4 Experiments

To perform the Approximate Bayesian Computation we ran the context-dependent models (Unbiased copying, Conformist copying, Top Threshold copying and Top Alberto copying) $8,770,560$ times. This leads to a total of $35,082,240$ simulations run on the Supercomputer Marenostrum of the Barcelona Supercomputing Centre. Execution of the simulation were parallelized thanks to the package Rmpi (Yu, 2002). The different step of the ABC were executed by functions greatly inspired by the package cTransmission (Crema and Kandler, 2019) developed for the paper by Crema et al. (2016). All the code used to implement, execute the models, run the ABC, parallelize the simulations and generate all visualizations is provided as a R-package available online here: github.com/simoncarrignon/spreadrt. The vignettes (spreadrt.html and abc_spreadrt.html) detail and give examples on how the code can be used to reproduce the experiments and figures.

The parameters of each simulation are randomly sampled from the priors described in Table 3.1.

In the next section we will use those $35,082,240$ simulations to select the best model able to reproduce the data and select for each model the $1,000$ best simulations to approximate their posterior distribution.

For the various reasons we mentioned in sections 3.3.3 and 3.5.2, we could not run as much simulations of the content model as we did for the others. Running 8 millions simulation was not technically possible so we ran only $405,888$ simulations that already took almost as much time as the full 35 millions simulations run in the first setup. This low number of simulation prevents us to strictly apply the Bayesian framework we described here. Nonetheless, we propose, in Section 3.7, a shorter and qualitative analysis of those $405,888$ simulations by comparing them to $405,888$ similar simulations randomly sampled from the previous analysis. The content bias simulations will be run by randomly sampling the parameters from the priors described in Table 3.2.

3.6 Results

3.6.1 Model Selection

To formally select between the different models we use the posterior distributions of the different models given the data. The Bayes equation described by Equation 1.1 then becomes:

$$P(m|D) = \frac{P(D|m)P(m)}{P(D)} \tag{3.5}$$

where $P(m)$ is the prior, the likelihood $P(D|m)$ is estimated through Approximate Bayesian Computation (Toni et al., 2009; Toni and Stumpf, 2010), and the probability of the data, $P(D)$, in the denominator, cancels out when we compare models to each other. To calculate $P(m, D)$ we define a level of acceptance t that determines the number of simulations we will accept (i.e. we will accept the tth best simulations). Then we calculate how many simulations of each model are below this acceptance level. Table 3.3 summarize this distribution of m for different level of $t \in [500, 5000, 50000]$.

We note that the Top Threshold is by far the model less likely to explain the data. The best models in Table 3.3 are the Unbiased and Top Alberto models. Since the Bayes factors do not change much even if we divide the level t by 100, the accepted simulations appear to be a good approximation of the real distribution.

To compare models more formally, having used uniform prior probability distributions for all models, we can compute the Bayes Factor K_{m_A,m_B} between pairs of models as follows:

$$K_{(m=A,m=B)} = \frac{P(m = A|D)}{P(m = B|D)} \tag{3.6}$$

For the smallest level of acceptance $t = 500$ and when modelling the true Twitter cascades, the Bayes factors show the Unbiased model to be slightly better than the Conformist model ($K = 1.18$) or the Top Alberto model ($K = 1.19$), while the Conformist and Top Alberto models are equivalent ($K = 1.01$) and the Top Threshold model is highly unlikely compared to the other three models ($K < 0.16$).

Similarly, when modelling the cascades of false tweets, the Top Threshold model is highly unlikely ($K < 0.02$) compared to any of the other three models. For false tweets, Unbiased and Top Alberto models are equally good ($K = 1.07$) and do better than the Conformist model ($K > 2.0$).

The number of parameters is, implicitly, taken into account in the Bayes factor: To approximate the likelihood while doing the ABC we randomly sample the same number of particles from the prior distribution, thus if the number of parameters for one model is higher, the parameter space is bigger and the sample size drawn from the prior will cover a smaller fraction of the total parameter space, yielding a lower probability to find good particles that we will not reject.

To calculate something comparable to AIC, we use the raw values from Table 3.3 divided by the total number of simulations. This would give us the approximated likelihood, L, for each model. Then AIC is $-2 \times lnL + 2p$ with L the likelihood and p the number of parameters. This gives a set of "corrected" Bayes factors as in Table 3.4, in which the basic neutral model (1) is the best for both sets of data.

3.6.2 Posterior Distributions

The ABC algorithm allows us not only to select between the models but also to look at the posterior distribution of the parameters that yield to simulations reproducing the data. The idea is then to explore the result of Equation 1.1, once the likelihood $P(D|\theta)$ has been approximated by the ABC.

We show the unbiased model is the most likely in section 3.6.1, so we present only its posterior probability distribution in Figure3.5. Each panel of Figure 3.5 compares the posteriors of the ABC done with the true tweets (in green) against

m	$P(m\|D, t = 500)$	$P(m\|D, t = 5000)$	$P(m\|D, t = 50000)$
Unbiased (1)	0.307	0.315	0.352
Conformist (2)	0.298	0.303	0.273
Top Threshold (3)	0.048	0.047	0.110
Top Alberto (4)	0.302	0.335	0.311
m	$P(m\|D, t = 500)$	$P(m\|D, t = 5000)$	$P(m\|D, t = 50000)$
Unbiased (1)	0.384	0.351	0.323
Conformist (2)	0.198	0.277	0.270
Top Threshold (3)	0.006	0.018	0.086
Top Alberto (4)	0.412	0.354	0.321

Table 3.3: Bayes Factor Table for different acceptance ratio. Top for distribution of true, bottom false.

Model	versus True	versus False
Neutral (1)	25.38	25.21
Conformist (2)	27.72	28.54
Top Threshold (3)	33.37	37.53
Top Alberto (4)	29.69	29.07

Table 3.4: “Corrected” Bayes Factor Table for different acceptance ratio.

the ABC done with the false tweets (in red) and the prior used in both case (in grey).

This allows us to quickly see the effect of the invention rate, μ, population size, N, and run-in time, τ. Notably, the most well-defined probability peak is for the invention rate μ in modelling the distribution of false tweet cascades. Also, there are fewer time steps required in modelling the true tweets compared to the false tweets, which is consistent with the finding that false tweet cascades persist longer (Vosoughi et al., 2018).

From those posterior distributions it is possible to determine the parameter range with the highest probability: the highest density region (HDR) (Hyndman, 1996) – often also called the Highest Posterior Density Region (HPD) when they are calculated on Posterior Distribution as it is the case here. We use the R package HDRCDE[3] (Hyndman, 2018) to calculate those HPD. The resulting intervals

[3]Most of the visual improvement developed for displaying the HDR presented in this thesis has been incorporated to the original package.

and modes are given in Tables 3.5 & 3.6.

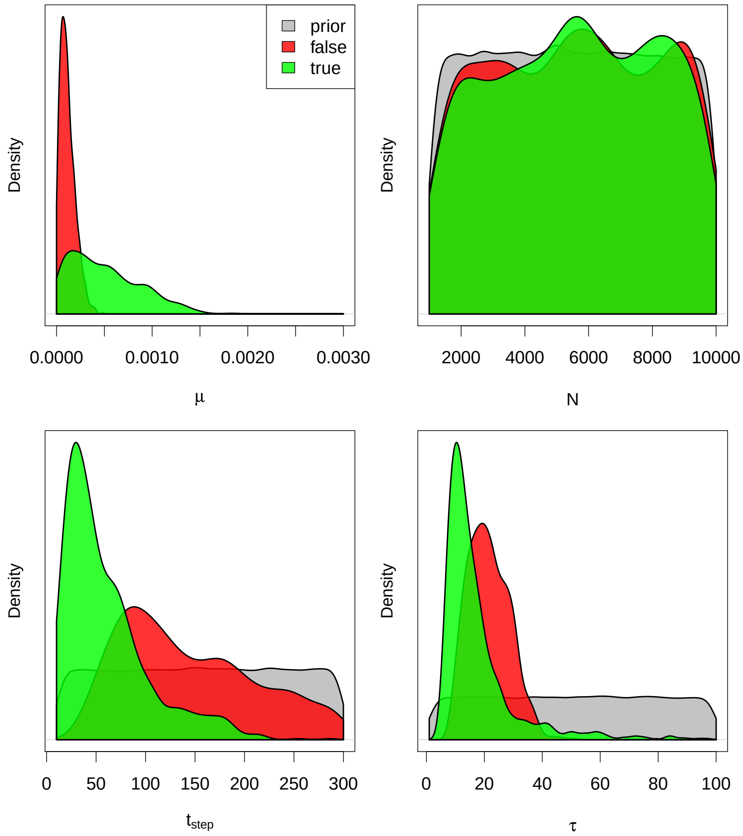

Figure 3.5: ABC posterior probability estimations for several parameters—invention rate μ, population size N, run time t_{step} and run-in time τ—of the unbiased (random) copying model, when the model is fit to the cascade size distributions of true (green) and false (red) tweets, respectively. The grey curve represents the prior distributions for each parameters.

3.6.3 Posterior Checks

As we explained it when we detailed the choice of the distance function, to store the full result of a simulation we need to store the size of all the cascades gener-

Parameter	mode	95% HDP
μ	2.48×10^{-4}	$[1.04 \times 10^{-5}; 0.12]$
N_{min}	5532	$[1050; 9579]$
t_{step}	23	$[5; 158]$
τ	9	$[2.7; 43]$

Table 3.5: mode and 95 % interval of the High Posterior Density region for the parameters of the unbiased model *wrt.* the distribution of true tweets.

Parameter	mode	95% HDP
μ	1.226×10^{-3}	$[6 \times 10^{-7}; 2.6 \times 10^{-2}]$
N_{min}	1528	$[1091; 9599]$
t_{step}	81	$[32; 274]$
τ	17	$[8; 34]$

Table 3.6: mode and 95 % interval of the High Posterior Density region for the parameters of the unbiased model *wrt.* the distribution of false tweets

ated. This, most often, leads to store an array of around $400,000$ integers (around 1.6Mo). Keeping this for the 9 million simulations ran for each model would have lead to about 72 TB of data. In order to save time and space while doing the ABC, we didn’t store it, but only the parameters used together with the distance to the data.

Thus, to check the adequacy between the simulations selected through ABC and the observed distribution, we could not just look back at the simulation selected by the ABC. In this context, a common way is to do what is called *posterior predictive check* (Gelman and Hill, 2007). The idea is to re-run a great amount of simulations by sampling the parameters from the selected posteriors distributions. We did this and re-run $10,000$ times each model, for both the posteriors obtained with true as well as false messages.

We present the results of the new simulations as distributions of cascade sizes. For a better visualization we binned the cascades with similar size within logarithmic bins. The High Density Regions for all bins and models are represented in Figures 3.6 to 3.9. The coloured dot represent the data from Vosoughi et al. (2018).

Figures 3.6 to 3.9 show how the respective models span a range of size distributions of aggregated cascade sizes. In each case, the simulation results are compared to the actual data of Vosoughi et al. (2018), for both the true messages and the false messages. The simulations and models fit well except underestimating the sizes of the top four or five largest tweet cascades, in both false and true categories (Figures 3.6 to 3.9). The models, particularly the random copying model (as expected given the Bayes Factors), otherwise predict the rest of the distribution of cascade sizes, albeit with different goodness of fit, which we consider below.

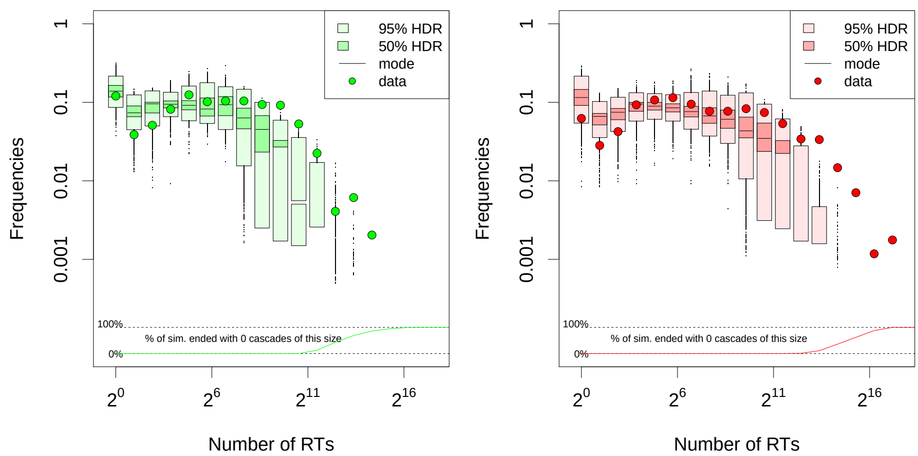

Figure 3.6: Posterior check of distributions of aggregated cascade sizes for random model. Both graph represent the percentage of rumours for which the accumulate number of RT falls within 18 bins of logarithmically growing size. For each bin the frequency of rumours that fall within it are represented by a coloured dots for the rumour from the original data set and by the mode and High Density Regions for the $10,000$ posterior checks. The percentage of simulations where no rumour felt within a given bin are represented as a curve at the bottom of the graph. The left graph shows data and posterior checks for the true rumours, the right graph for the false rumours.

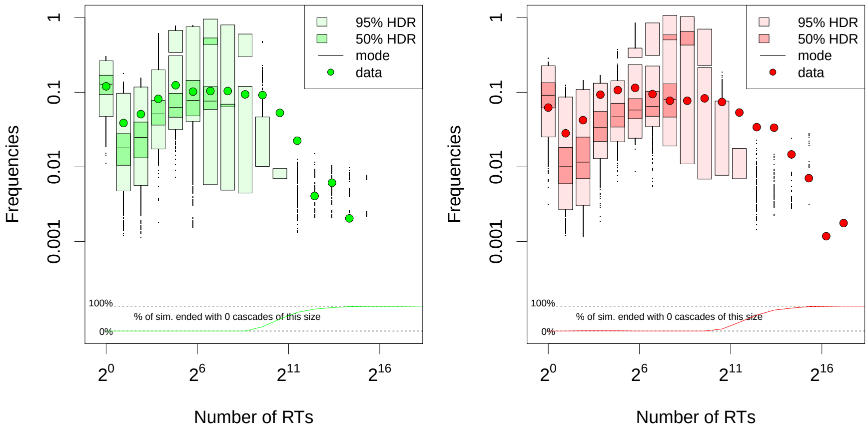

Figure 3.7: Posterior check of distributions of aggregated cascade sizes for the conformist model. Plot were generated as described in the caption of Figure 3.6.

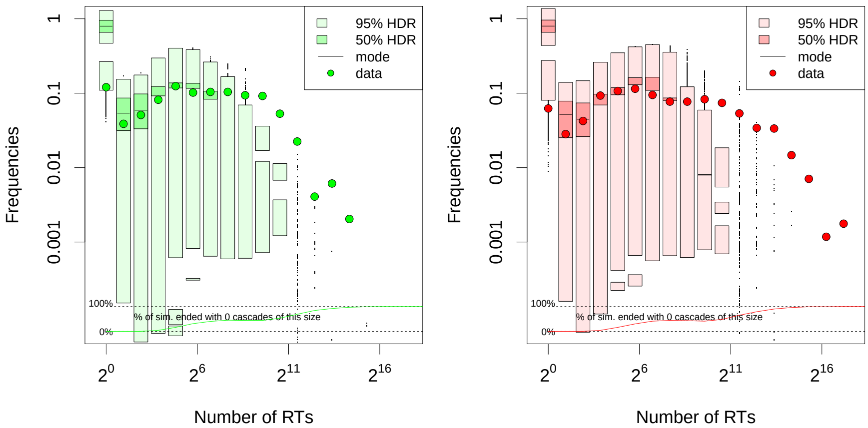

Figure 3.8: Posterior check of distributions of aggregated cascade sizes for the ”top threshold” model.Plot were generated as described in the caption of Figure 3.6.

3.7 Content Bias Exploration

Now that we reviewed in details the results of the ABC done with the context-dependent bias we propose a quick overview of the restricted ABC ran with the content-dependent bias.

3.7.1 Model Selection

To compare the content-bias model with all the others we use subsets of the simulations run in the previous experiment. But to apply the model selection procedure we used before, some adjustment need to be done. To run the content-biased model in a reasonable amount time we had to shorten the prior (*cf.* Tables 3.1 and 3.2). Though most of the parameters doesn't have a correspondence form one model to the other, we need to remove the simulations from the previous ABC with priors that obviously doesn't fall within the one used here (mainly, all the simulations with $\mu > 0.035$). We then sample $405,888$ simulations among the remaining one to match the sample size of all models.

From those $2,020,440$ simulations we can select the top 500, $5,000$, $50,000$ as we did before, and see how the five models are distributed. The results are represented in Table 3.7. More than 70% of the simulations present in the top 500

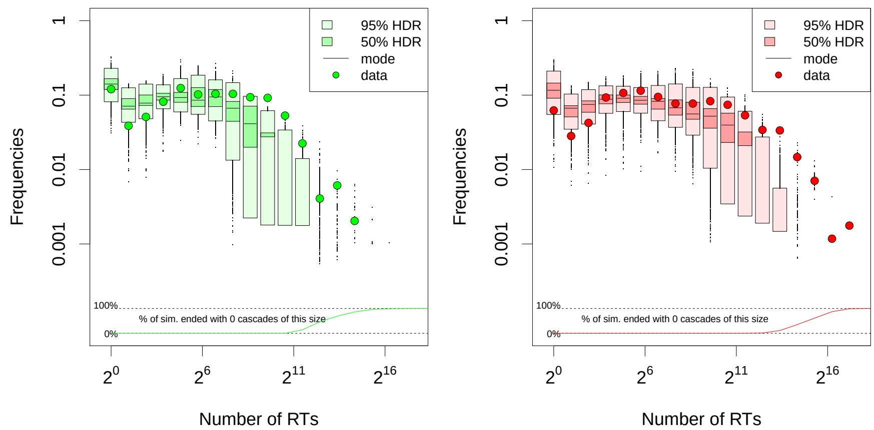

Figure 3.9: Posterior check of distributions of aggregated cascade sizes for "top Alberto" model. Plot were generated as described in the caption of Figure 3.6.

m	$P(m\|D, t = 500)$	$P(m\|D, t = 5000)$	$P(m\|D, t = 50000)$
Unbiased (1)	0.067	0.036	0.059
Conformist (2)	0.062	0.032	0.048
Top Threshold (3)	0.016	0.047	0.032
Top Alberto (4)	0.068	0.035	0.058
Content (5)	0.787	0.886	0.804
m	$P(m\|D, t = 500)$	$P(m\|D, t = 5000)$	$P(m\|D, t = 50000)$
Unbiased (1)	0.105	0.037	0.026
Conformist (2)	0.074	0.030	0.021
Top Threshold (3)	0.005	0.003	0.013
Top Alberto (4)	0.094	0.036	0.026
Content (5)	0.722	0.894	0.914

Table 3.7: Bayes Factor Table for different acceptance ratio. Top for distribution of true, bottom false.

are generated by the content model, both for the true and false distribution. We could be tempted to say with confidence that the content model as a much higher probability to explain the data than the five others. Nonetheless, it should be noticed that the 500 best simulations in the previous experiment represent 0.0014% of all the simulations. Given the limited amount of simulations available for this one, this top 500 represents more than 0.14% of the simulations of the current ABC. In order to select the same percentage of the best simulations we should look at the top 5, 50 and 500. If will look at those smaller ranks the results are much less robust and much more dependent on the sampling process.

To show this we repeated 50 times the process used to generate the Table 3.7, each time re-sampling new simulations among from the previous experiment. For each repetition we recorded the distribution of the five and 50 best models. The results are summarised in Figure 3.10. As we can see, the greater goodness of fit of the content bias is less clear when we are more restrictive. For the top 50 it seems that the three best models (Unbiased, Alberto and Content-biased) are generating similar amount of good simulations, whereas the content bias wasn't able to generate simulations good enough to enter in the top 5.

But given the small number of simulations analysed here, this could very well be an artefact due to the stochasticity of the process. This is why the main discussion of this chapter will not feature the results from the content-biased model. Nonetheless, given that they look promising, we will quickly review some inter-

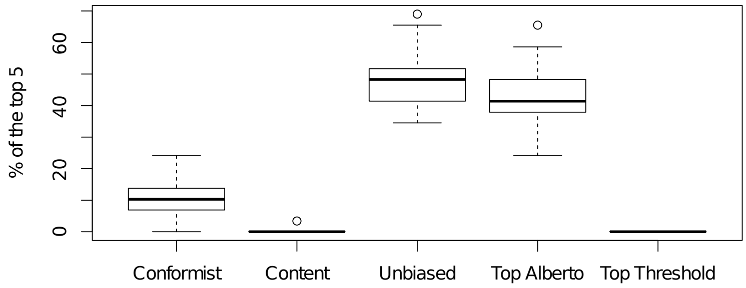

(a) Top Five

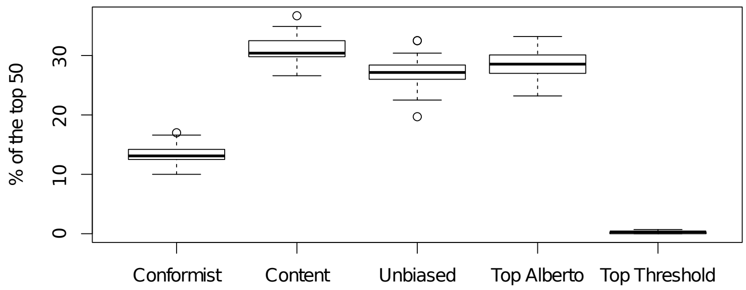

(b) Top Fifty

Figure 3.10: Robustness of model selection for $t = 5$ (a) and $t = 50$ (b) for 50 different sampling. Each box represent the distribution of the percentage of each model in the top 5 (a) and top 50 (b).

esting aspect of the posteriors in the next section, as we did before for the neutral model.

3.7.2 Posterior Distributions

The content biased model we described in Section 3.3.3 involves more processes with more parameters than the four other models. Interpreting its posteriors is not simple, as some parameters interact in subtle ways. Moreover, as we have shown previously, it’s very likely that those posteriors are not a good estimate of the real one, but an artefact generated by the low number of simulations. Thus, we will not review the marginal posteriors in detail as we did before, but will focus the interesting one as illustrative purpose.

To simplify the interpretation of the posterior, we will group together the users with $\beta > 0$ within the category “positively polarized”, the users with $\beta < 0$ within the category “negatively polarized” and both category as “polarized individual” whereas users with $\beta = 0$ are considered as “neutral individual”. We do the same for the different categories of tweet’s utility. In Figure 3.11 we represent the percentage of polarized users given the percentage of polarized rumours. The red area represents the posterior calculated with regard to the false tweets, the green area the posterior calculated with regard to true tweets and the grey area the priors generated by our function $distrib()$.

From Figure 3.11 we observe that polarization of users is correlated with the polarization of tweets, and this for both distributions (true or false tweet). In order to reproduce the observed data, the percentage of polarized tweets need to match the percentage of polarized users. The major difference between the fit to true and false tweets is the small shift away form the diagonal observed for the false tweet (the red area in Figure 3.11). This could illustrate that in order to have the bigger and wider cascades observed by Vosoughi et al. (2018), we need a slightly higher number of polarized individuals re-tweeting a fewer number of polarized tweets and not necessarily more polarized elements in general.

But this analysis should be taken with caution. The posterior distributions represented in Figure 3.11 may not be representative of the true posteriors as derived from too few simulations. Moreover, the priors use for the ABC and described in Table 3.2 are limited and may not include value that would allow to more closely reproduce the data. In the next chapter we will see a different implementation of ABC that allows to explore more efficiently model even if the prior aren’t well defined and thus could solve these problems.

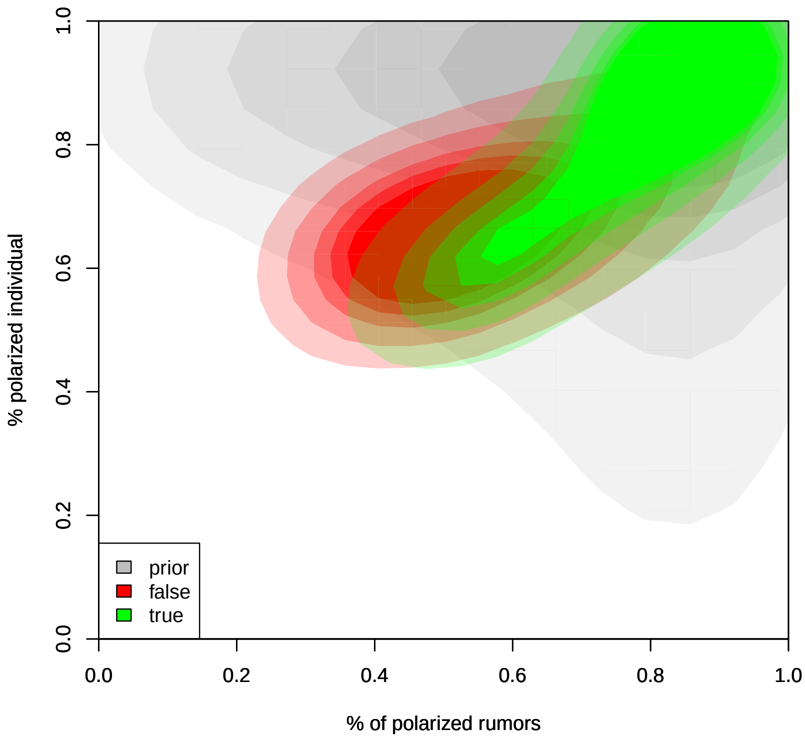

Figure 3.11: Joint posterior distribution of the percentage of polarized rumours and users. The grey area represent the prior distribution generated by the function $distrib()$. They red area represent the posterior distribution selected when fitting to the distribution of false tweet, the green while fitting the false tweets.

3.7.3 Posterior Checks

Using the method presented in Section 3.6.3, we run $10,000$ simulations to visually check how the parameters selected by the ABC allow the content-bias model to fit the data. We did so for both the true and false distributions. The results are represented in Figure 3.12 within 16 bins and High Density Regions as we did in Section 3.6.3.

As with the previous models, the content-bias predicts relatively well the overall distributions of tweets. This is a good surprise given the limited amount of simulations within which the bests were selected for this experiment in comparison with the other experiment. Nonetheless, and as observed in the previous checks, they all fail to generate the top four or five largest tweet cascades, in both false and true categories (Figure 3.12). This suggests that even with content bias and the polarized populations shown in Figure 3.11, it is still not enough to reproduce the spread and growth of the biggest cascades. But again, it may only reflect the fact that we haven’t explored the good prior or not enough precisely, which cannot be excluded given the small number of simulations ran with regard to the complexity of the model.

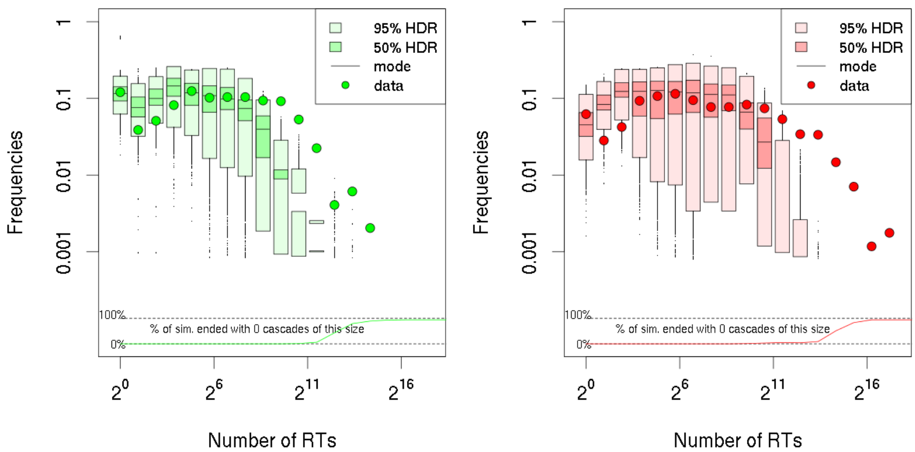

Figure 3.12: Posterior check of distributions of aggregated cascade sizes for the ”content” model. Plot were generated as described in Figure 3.6.

3.8 Discussion

In calibrating neutral model variations against distributions of Twitter cascade sizes, we find that the unbiased neutral model applies well to re-tweeting activity, and better than models with added conformity bias. Though preliminary results showed that the content-bias model could even do better, the technical obstacles that we presented throughout this chapter prevented us to fully demonstrate this. We have to keep the extended analysis and interpretation of the content model against the actual data for latter studies. Nonetheless we have successfully illustrated the use of ABC to resolve the joint probability distribution of all models presented here (Kandler and Powell, 2018).

Here we have applied the same set of models to both data sets, with the expectation that the different models, and the probability distributions on the parameter values for each model, would differ meaningfully between the false and true tweets. As Figure 3.5 and Tables 3.3-3.6 show, we find that the unbiased (and Top Alberto) model fits the distribution of both true and false Tweets, but the invention rate, μ, for the false Tweets (0.00123) is about five times higher than for the true Tweets (0.00025).

The modal value of the invention rate, μ, returned from our model runs was 0.00025 for the true tweets and 0.00123 for the false tweets (Tables 3.5 and 3.6). These values agree well with estimates we can make from the data from (Vosoughi et al., 2018, S2.3), which comprise 2,448 different rumours re-tweeted about 4.5M times, about two thirds of which ($\sim$ 3M) were re-tweeting the 1,699 false rumours and the other third ($\sim$ 1.5M) re-tweeting the 490 true rumours. This implies that about 0.00032 of the true tweets and 0.00057 of the false tweets were original, both of which are well within the High Posterior Density region for μ of the respective models (Tables 3.5 and 3.6).

In our tests, the most important parameter was the invention rate, μ, particularly in modelling the distribution of false tweet cascades. Another important parameter was the transmission bias, such that neutrality ($\beta = 0$) and positive-frequency bias ($\beta > 0$) can be evaluated in terms of likelihood of explaining the data (Kandler and Powell, 2018). The biased models performed worse than the unbiased ($\beta = 0$) model, as the biased models also failed to generate the largest cascades while implying more parameters (Figures 3.6-3.9). This is not due to limits on modelled population size, if it were, we would expect the posterior distribution for N to be more right-skewed. Indeed the value of N had little effect on the outcome.

Replicating the largest cascades was a challenge for all models. This suggests a few ways forward for future modifications. One is to allow for interdependence between the distribution of true and false messages, rather than as separate data sets since real Twitter users are exposed to both kind of messages. Another is to specify the network of connections among agents (Lieberman et al., 2005; Ormerod et al., 2012).

Heterogeneous networks may help resolve a discrepancy between our findings—that conformist biases were unhelpful in modelling Twitter data—and the highly-skewed nature of influence on social media. In a recentrecent study of social media by Grinberg et al. (2019), "Only 1% of individuals accounted for 80% of fake news source exposures, and 0.1% accounted for nearly 80% of fake news sources shared." The existence of "Influentials" is not unique to social media: in order to fit the neutral model to evolving English word frequencies over 300 years of books, Ruck et al. (2017) needed to assume that most of the copying was directed to a relatively small corpus of books, or "canon", within the larger population of millions of books.

Testing such models will require more granular data, including the content and word counts from tweeted messages, than we had access for this chapter. If counts of specific words through time are available, then additional diagnostic signatures include both the Zipf law of ranked word frequencies and turnover within "top y" lists of those words (Acerbi and Bentley, 2014; Bentley et al., 2007; Ruck et al., 2017). While we opted for parsimony here, more granular data would justify testing more complicated neutral model modifications, such as non-equilibrium assumptions (Crema et al., 2016), variable "memory" (Bentley et al., 2011; Gleeson et al., 2014) and isolation by distance effects (Bentley et al., 2014a).

3.9 Conclusion

Here we have tested variations of culture-evolutionary neutral models on aggregated Twitter data documenting the spread of true and false information. We use Approximate Bayesian Computation to resolve the full joint probability distribution of models with different social learning biases emphasizing context-biased versus content-biased learning.

This chapter begins to address how online social learning dynamics can be modelled through the tools of cultural evolutionary theory and how we can test hypothesis about those dynamic against real data. We have shown how models implementing those hypothesis can be precisely adjusted and compared to available record of online activity. Given the richness of the available data we expected to test a much wider range of more realistic model of social learning such as content biased social learning. We have seen throughout this chapter the problems that such methods raises when applied to more complex models. The realism of those models imply the modelling of more complex processes, that need more space to store more information and ultimately last much longer. These technical limits prevented us to fully explore the content model proposed.

However, far from illustrating the failure of such methods we think it opens the door for numerous fruitful future researches. The results presented here with the context-dependent models help to narrow he prior by excluding very unlikely ranges of parameters. In the meantime the difficulties faced with the content model highlighted the points that need to be improved, among them the computation cost of ABC and the difficulty to find the right distance function to bind together high dimension model and data. To improve the first point, we will see in the next chapter a solution proposed to minimize the computation cost of ABC.

Overall, this Chapter demonstrates how Culture Evolution can be used to understand social phenomenon in online social media. Finding tools to understand such phenomena is crucial, as social media are heavily reshaping the way we interact, its speed and breadth. Moreover when we know that these social media, by their digital nature, can easily be manipulate by algorithmic or semi algorithmic automation, which can ultimately lead to important changes in our societies (Ruck et al., 2019). On the other hand, the same way the nature of online social platforms exposes them to algorithmic and automatised manipulation, this can be used to collect and precisely monitor what is happening on such platforms. Thanks to that, online social media can be turned into the perfect experimental tools to test and explore hypothesis about Social Learning and Cultural Evolution and this chapter proposes and illustrates a method to do so.

Chapter 4

DETECTING THE IMPACT OF SOCIAL LEARNING BIAS IN LARGE-SCALE HISTORICAL TRADE NETWORKS

4.1 Introduction

In this chapter we combine the framework developed in Chapter 2 with the Approximate Bayesian Computation method introduced in Chapter 3, in order to study and detect the impact of content-dependent biases during 500 years of tableware trade in the Roman Empire. To do so we rely on the record of presence or absence of tablewares (plates, bowls, cups,…) in cities in the Eastern Roman Empire (Roman East) from 200 BC to AD 300. This record is of great interest for this thesis as they suggest changes led by large scale economic and social interactions.

Exploring Cultural Transmission at such large historical scale is the only way to test assumptions about the importance of such mechanisms in the evolution of Human Culture. By looking at long term changes we can quantify and detect processes and factors that have shaped and driven the political and economic history of our civilisations (Turchin et al., 2013). This is made possible thanks to the archaeological record, which allows to identify biases on cultural transmission over long-term trajectories (Lipo and Madsen, 2001; Shennan and Wilkinson, 2001; Bentley and Shennan, 2003; Steele et al., 2010; Crema et al., 2016; Kandler et al., 2017). However, material culture recovered from archaeological contexts is noisy and fragmented, and realising such study is a challenge (Porčić, 2014; Crema et al., 2014).

In the previous chapter thanks to the nature of the data we could explore direct content-dependent bias: it was possible to directly measure the content, or the intrinsic value, of a tweet. Here we don't have any clue about this intrinsic value: all tablewares are the same. We thus need to explore a different social learning strategy: success biased strategies as we described it in Chapter 2. Using the framework we developed to explore this strategy we will compare success bias to other content-independent models and test all of them against the data. Our goal is to find the social interactions that led cultural and economic changes in the Roman Empire, as well as to find empirical evidences of the process theoretically explored in Chapter 2.

We will first present the Archaeological and Historical context, the questions it raises and then translate those questions within the framework of cultural evolution and social learning. We will briefly describe a modified version of the algorithm presented in Chapter 2. This new version will allow us to implement the hypotheses we want to test while integrating the historical context studied here. To test these models against the data, we will propose another version of Approximate Bayesian Computation that solves some of the problems we have encountered in Chapter 3. After presenting the results of the ABC we will discuss their general and archaeological implications.

4.2 Tableware trade in Roman East

4.2.1 Archaeological Context

Vast quantities of foodstuffs, stones, minerals and craft products were traded over huge distances in Roman times, despite the significant limitations imposed by the then-current transport and communication technologies, and the uncertainties caused by climate and piracy. Seaborne commercial activity in particular facilitated long-distance trade flows throughout the entire Mediterranean region in late Hellenistic, late Roman Republican and Roman Imperial times. However, the extent to which the commercial actors involved in this inter-regional trade could depend on abundant reliable commercial information about the supplies and demands of goods from other parts of the Roman world is still uncertain. This issue lies at the heart of current debates on the functioning of the Roman economy and in particular its degree of economic integration, where availability of reliable commercial information is considered a condition for the Roman economy to be highly integrated (Wilson et al., 2012; Scheidel, 2012; Morris et al., 2007; Temin, 2013; Bang, 2008).

Ceramic tableware (thin-walled plates, cups and bowls) offers one of the most abundant sources for studying inter-regional trade, but also one of the only providing comparable and quantifiable information over centuries-long time periods (Wilson, 2009; Brughmans and Poblome, 2016). The recently aggregated tableware evidence from the eastern Mediterranean (Bes et al., 2018; Bes, 2015), which allows for the quantitative identification of centuries-long distribution patterns and can be used with computational models for formal hypothesis testing, reveals a particularly robust and well-studied distribution pattern (Hayes, 1997, p.14). Although a large number of fine ceramic tablewares were produced in the Eastern Mediterranean region during the late Hellenistic and Roman Early Imperial periods, only a handful achieved a commercial distribution that went beyond their region of production: ESA, ESB, ESC, and ESD. In addition to these four eastern-produced tablewares, a fifth ceramic tableware achieved a very wide distribution in the Roman East: Italian Sigillata (ITS). The identification of these wares across hundreds of sites and surveys in the Eastern Mediterranean allows us to observe trends in the width and overlaps of their distribution for a period of five centuries between 200 BC and AD 300 (for detailed discussions of this distribution pattern, see Bes 2015; Hayes 2008).

The archaeological record reveals in a number of ways that the imported tableware market in the east in this period was likely competitive. First, although each tableware has its own core region of distribution, there are large overlaps in the distribution patterns of tablewares in the Eastern Mediterranean. Second, the different tablewares also have strongly different distribution widths (we use the term distribution width to refer to the number of archaeological sites at which the ware is attested), which saw dramatic changes through time. ESA was for more than a century the most widely distributed ware, until the arrival of the western-produced ITS coincided with the sharp decrease in ESA distribution width and the increase in ESB distribution width. Third, the shapes of the eastern-produced tableware vessels show strong influences among themselves, but particularly striking is the widespread adoption of Italian Sigillata features upon its arrival in the Eastern Mediterranean markets in Augustan times, particularly in the period 10 BC - AD 15 (Hayes, 1997; Jones, 1950; Waagé, 1948). This could represent the adoption of the fashionable morphological features of a successful new imported product, which might have been motivated by an economic strategy to remain competitive on the eastern market.

Can these accumulated archaeological observations be explained by competition and can we specify the nature of this competition? Does the introduction in the east of new tablewares in general, and of the western-produced ITS in particular, lead to increased competition that reduces the dominance of ESA? Here, we explore the theory whether large-scale tableware buying decisions at markets throughout the Roman East were influenced by the buying strategies at other markets, and therefore also by the ability of traders to access reliable commercial information. In doing so we aim to leverage the potential of the existing tableware data to provide insights into the degree of Roman economic integration.

4.2.2 Tableware data

The data used in this chapter are evidenced in the ICRATES database, the largest collection of tableware sherds excavated and published from hundreds of sites in the eastern Mediterranean (Bes et al., 2018). From the entire database we here use a subset of 8730 chronologically datable entries of the five wares (ESA, ESB, ESC, ESD, ITS) from 178 eastern Mediterranean sites. We only take into account the main ware category (i.e. ESA, ESB, ESC, ESD, ITS) and do not distinguish between possible identification of different versions of the wares (e.g. ESB I-II, ITS-Arezzo) since this information is not reliable and comparable for all excavations included in this database. Similarly, we only take into account the presence of the absence of a ware at a site: the quantitative volume of that ware or the typological diversity recorded in the ICRATES database is not representative or comparable for all sites.

In order to determine the production regions (Table 4.1), out of the four eastern wares, only for ESC were actual ceramic production sites excavated, in Pergamon and the surrounding region (Meyer-Schlichtmann, 1988; Poblome et al., 2001). A combination of geochemical analyses and distributions of excavated pottery has allowed archaeologists to pinpoint the region of production of ESA in the Levantine coastal region between Latakia and Tarsos, of ESB in the Maeander valley in western Turkey and of ESD in (western) Cyprus (Hayes, 1972, 1985, 2008; Meyer-Schlichtmann, 1988). ITS was produced in a range of western workshops among others in Arezzo, Pisa, Lyon and in the Po Valley.

The provenance from the literature of the data used is summarized in Table 4.1. The geographical distribution of the data is represented in Figure 4.1.

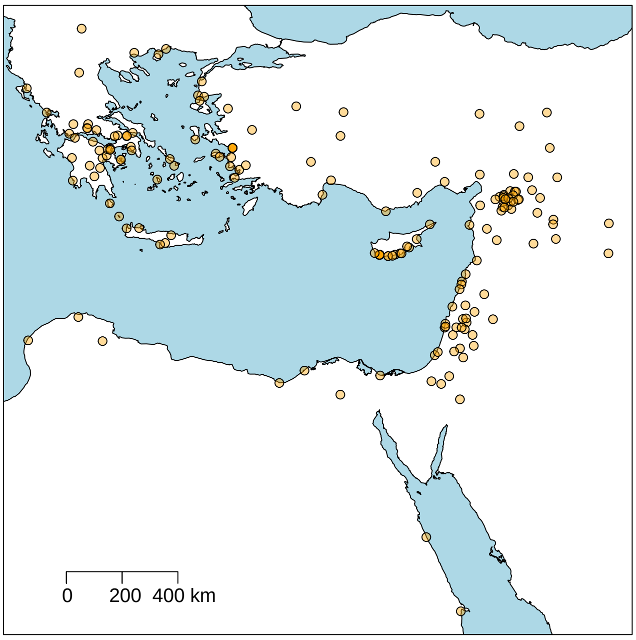

Figure 4.1: Geographical representation of the 178 eastern Mediterranean sites from the ICRATES database. Each point represents a site from the database.

To allow for identifying changes through time in distribution patterns, we draw on the standard dating ranges of the established typologies for these wares (Table 4.1). According to these, each morphological type has a different chronological date range. We count the number of sites at which each type of each tableware was found. We use cumulative probabilities to add up evidence at each site of the same ware but of different types with different dating ranges (some types have a narrow dating, but others can have very broad dating ranges). For each site/ware combination we calculate the probability that it existed in any given

Ware	Abbreviation	Typological and chronological standard	Region of production (Schneider 2000)
Eastern Sigillata A	ESA	Hayes (1985)	Coast between Tarsos (TUR) and Latakia (SYR)
Eastern Sigillata B	ESB	Hayes (1985)	Maeander Valley in western Asia Minor (TUR); possibly Aydin (ancient Tralleis)
Eastern Sigillata C	ESC	Hayes (1985, 1972) and Meyer-Schlichtmann (1988)	Pergamon and surrounding region
Eastern Sigillata D	ESD	Hayes (1985)	Cyprus (probably the western part)
Italian Sigillata	ITS	Ettlinger et al. 1990	Italy and Southern France

Table 4.1: Typological, chronological references and possible region of production for major tablewares studied in this paper

year, following a well-established approach in Roman archaeology (Bes, 2015; Fentress and Perkins, 1988; Willet, 2012) and assuming a uniform probability distribution. For example, a pottery find that is dated between AD 1 and AD 10 will add the value of $\frac{1}{10}$ for each year between AD 1 and AD 10, because the probability that it existed in any one of those years is 10 % when assuming a uniform probability distribution. A pottery find dated to AD 1-100 will add the value of $\frac{1}{100}$ to each year between AD 1 and AD 100 because the probability that it existed in any one of those years is only 1%. As we add these partial probabilities together we get the cumulative probability that a given tableware was used at any one time point. As we keep on adding more pottery finds we reveal a chronological overview of the pottery distribution in the region (Figure 4.2). In this paper we aim to understand whether these data patterns, illustrated in Figure 4.2, could reveal aspects of competition between tablewares on the eastern market and whether they can be explained by social learning mechanisms guiding the decisions of tableware traders. We will therefore briefly describe both data patterns.

The most striking tableware distribution pattern is the extremely wide distribution of ESA for more than a century and its significant decline shortly after the introduction and rapid expansion of the western-produced ITS (Figure 4.2, left panel). In the late first century AD, ESB reached a peak in its distribution width similar to that of ITS and the distribution of ITS contracts as rapidly as it increased more than a century before. An alternative representation of the same dataset is to explore the frequency distribution of the number of tablewares at each site (Figure 4.2, right panel). At the start and end of our period of study a majority of sites in the dataset have no tableware product because all wares' distributions are extremely limited. But the data is more interesting for the central period between 100 BC and AD 100, when many sites have evidence of only one ware and some have evidence of two or more wares.

These data patterns illustrate that the settlements in the Eastern Mediterranean market were not divided into exclusive market shares of specific products. Rather, each product had a core region of distribution close to its production area and there was a strong degree of overlap between the distributions of the tableware products (Bes, 2015; Brughmans, 2010). This overlap increases in particular between 100 BC and AD 140 as more wares come onto the eastern markets, most notably ESB and the western-produced ITS. For example, during this period ESA vessels were not uncommon in Ephesos which was in the core distribution region of ESB (Hayes, 2008, p.18).

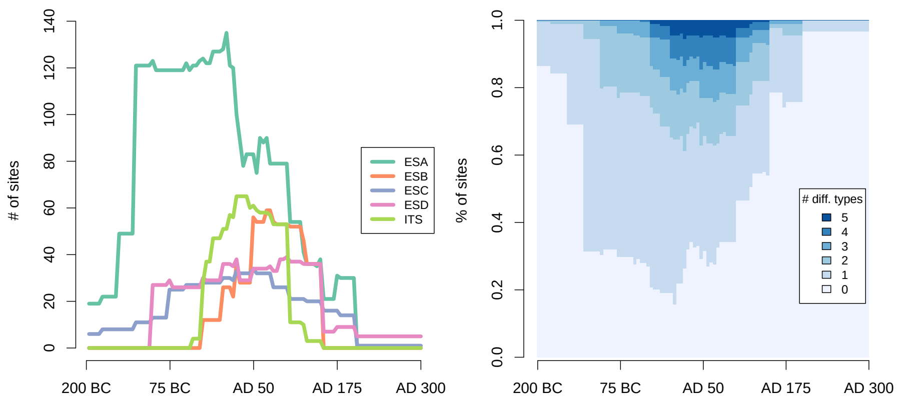

Figure 4.2: Data patterns derived from 8730 datable entries of 178 eastern Mediterranean sites from the ICRATES database. Left: number of sites for each pottery type; Note the disproportionate dominance of the oldest pottery type (ESA - blue-green line), and its decrease with the introduction of the western-produced ITS (green line). We will later refer to this metrics as "Pattern A". Right: distribution of distinctive pottery types at archaeological sites. Note the dominance of sites with none of the types or only one type. We will later refer to this metrics as "Pattern B".

4.2.3 Evidence for Imitation and hypothesis

Two further archaeological observations can also be interpreted as evidencing a competitive tableware market in the Roman East: the changes in ESA vessel shapes and the practice of stamping on ESB. In Augustan times (particularly between 10 BC and AD 15) the eastern-produced wares started widely imitating Italian wares (ITS) which were then relatively recently introduced on the eastern Mediterranean market (Hayes, 1997, p.52). ESA shapes were strongly inspired by ITS shapes, although they were not exact copies of ITS (Hayes, 1997; Jones, 1950; Waagé, 1948). A second phenomenon is the appearance of potter's name stamps impressed on the floor of the vessel on eastern-produced tablewares. Eastern wares originally did not have this practice, but it is widespread in ITS. Only a small percentage of ESA vessels was ever stamped (Hayes, 2008, p.17) and when stamps appear on eastern vessels they are typically written in Greek. The exception to this is ESB, which emerged around the same time as ITS and whose vessels have stamps written in Latin. Some ESB vessels show names known from the western ITS production, such as C. SENT from Arezzo and SERENI probably from Pozzuoli (Hayes, 1997, p.54-57). Other vessels have stamps including ARRETI and ARRETINA in reference to the Italian Sigillata production centre in Arezzo. These stamps are not an expression of the physical place of production, because it is clear from the vessel fabric they were not produced in western ITS workshops. Rather, they are an expression of influence of ITS on ESB. The practice of eastern wares adopting the habit of potter's name-stamps declined in the early 1st century AD, later stamps on eastern-made vessels are very rare.

Theories about the nature of the influence of ITS on eastern wares differ. The ESB stamps could evidence the existence of eastern workshop branches of Italian potters or even the actual physical presence of ITS potters in the east (Zabehlicky-Scheffenegger, 1995; Zahn, 1904), but they could equally reflect the copying by ESB producers of the ITS practice of potters' stamps. The morphological similarity between ESA and ITS could equally have been the result of active copying. This influence is not only illustrated for the case of ITS but also among eastern wares themselves. Some ESA vessels adopted 'motto'-type stamps from ESB, and some features of ESD show influence from ESA (Hayes, 1997, p.57-59). A further important theory is that this was an economic strategy to remain competitive on the eastern market, devised as a reaction to the rapidly growing distribution width of ITS.

The archaeological information offers clear evidence of influence among tablewares in the eastern Mediterranean, suggesting a competitive market where economic strategies of those active in tableware trade could have influenced each other. It suggests the emergence of ITS in particular made the eastern tableware market increasingly competitive and triggered the producers and traders of eastern wares to change their practices.

What was the role of competition between the traders who bought and sold tableware in giving rise to these data patterns? Can the data be explained by traders having access to each other's buying strategies, despite the significant distances involved and the logistical limitations for people in the ancient world to gather reliable information? If so, does the copying of the strategies of the most successful trader offer a good explanation? Or were traders not able to collect much reliable commercial information and instead changed their commercial strategies independently or through chance encounters with other traders ?

To explore those questions we translate them into three hypothesis about the Social Learning Strategies underlying the changes observed:

1. Independent Learning: Agent independently changes their tableware buying strategy (no access to reliable commercial information).

2. Unbiased Social Learning: Agent randomly copies the tableware buying strategy of another agent (limited access to reliable commercial information).

3. Success-Biased Social Learning: Agent copies the strategy of the most successful other agent (complete access to reliable commercial information).

4.3 Models and implementation

To implement these three hypotheses we use the agent-based model presented in Chapter 2. This model suits well the current case study as it articulates Social Learning Strategies with trade and economic activities, the two central elements of the system we described. In this model agents copy cultural traits from other agents and given those cultural traits they will calculate the price they use to trade. Cultural components are implemented in the model through social learning and innovation mechanisms, and economic components through production, trade and consumption mechanisms.

We have explored general properties of this model in Chapter 2. We have shown how the success biased social learning can lead to economic equilibrium.

In the last sections of Chapter 2 we have also started to detail how changing the configuration of this model changes its outcome. Here we will go one step further by testing it against the empirical data, using Approximate Bayesian Computation (ABC), the method shown in Chapter 3.

But before being able to use ABC, the theoretical model explored in Chapter 2 needs to be adapted to reflect the historical and archaeological context of this study. Thus a number of variables have been fixed and the following modifications were implemented:

- A period of 500 years is simulated, from 200 BC to AD 300. The order of appearance of different tablewares follows the standard chronology (Table 4.1; Figure 4.2):

  From 200 BC to 101 BC: ESA, ESC

  From 100 BC to 41 BC: ESA, ESC, ESD

  From 40 BC to 28 BC: ESA, ESC, ESD, ITS

  From 27 BC to AD 149: ESA, ESB, ESC, ESD, ITS

  From AD 150 to AD 199: ESA, ESC, ESD

  From AD 200 to AD 300: ESC, ESD

  To represent this within the model, the full length of the simulation is divided in periods that can be directly translated to the real chronology (*cf.* Section 4.4.2).

- We simulate 500 agents, where each agent representing the tableware buying and consuming behaviour of one urban settlement. This number was chosen to ensure reasonable computation time and because it roughly approximates the number of urban settlements in the Roman East estimated by Hanson et al. (2017); Wilson (2011) (*cf.* Section 4.4.2).

- The original model had an evenly distributed ratio of producers vs consumers, which in this implementation was modified to an uneven ratio of five producers (one for each type: ESA, ESB, ESC, ESD, ITS) vs 495 consumers.

- In addition to these five products, money was introduced as a sixth good produced by all agents and used to buy other products (how to introduce money and its impact on the original model was already proposed in the paper by Gintis 2006).

The original description of the implementation of this model is given in Algorithm 2.1. We reproduce it in Algorithm 4.1 with the function that introduces or removes goods in the tableware market following the sequence of appearance described before.

Algorithm 4.1 Model of social learning and economic trade with integration of accurate historical evidences.

```
 1: INITIALIZATION:
 2: for i ∈ #Pop do                ▷ Initialize the agent with no goods and a random value vector
 3:     Q^i = (0, ··· , 0)
 4:     V^i = (v^i_0, ··· , v^i_n)                    ▷ The values of v^i_j are selected randomly
 5: SIMULATION:
 6: loop step ∈ TimeSteps
 7:     if historicalChange() then
 8:         update(Q, V)  ▷ We add or remove a product given historical evidences and update associated
    value in consequences
 9:     for i ∈ Pop do
10:         Production(Q^i)
11:     for i ∈ Pop do
12:         for j ∈ Pop do
13:             TradeProcess(V^i, Q^i, V^j, Q^j)
14:     for i ∈ Pop do
15:         ConsumeGoods(Q^i)                                  ▷ All goods are consumed
16:         if (step mod CulturalStep) = 0 then
17:             CulturalTransmission(V)
18:             Innovation(V^i)
```

From this general model the three hypothesis described in Section 4.2.3 can be implemented:

1. To implement the independent learning hypothesis we simply remove the call to $CulturalTransmission$ at line 17.

2. Unbiased learning is achieved by switching $CulturalTransmission$ with a selection mechanism that randomly select cultural traits from the population with a probability μ.

3. To implement the success-bias mechanism we use the $CulturalTranmission$ process as we described it before in Algorithm 2.2.

parameter	description	initial value
EC	Total number of Economic Interaction	inferred
SC	Rate of Strategies change	inferred
μ	probability of individual update (i.e. innovation)	inferred
λ	probability social update (i.e. cultural transmission)	inferred
N	The total number of agents	500
μ_{max}	Amplitude of individual learning	inferred
λ_{str}	Strength of social learning	inferred
n_{good}	Number of goods produced and exchanged	from 3 to 6

Table 4.2: List of parameters tested in our different models. Note that some parameters are not used in all models.

We summarized in Table 4.2 all the parameters used in this model and their initial value if their is one. When the value is *inferred*, this means that the value will be set via the ABC. Note that some parameters are not used in all models (for example μ, the probability to copy another agent, is not used in the independent learning model).

4.4 Approximate Bayesian Computation

We described the Bayesian Inference approach in the Introduction and illustrate its use in Chapter 3, where we explored, tested and compared models of cultural transmission with the data from Twitter. Here, as we did in the previous chapter, we compare a content-dependent model of social learning with other content-independent biases. This time the content-dependent bias is an indirect one: success bias, that we presented in Chapter 2, while the content-independent biases are the unbiased and independent learning.

If the overall ABC process is the same and if the content-independent models can be straightforwardly translated from the models explored in the previous chapter, the success bias model implies new problematics and constraints that were absent in the direct content bias scenario and that we will quickly review here. We still need to introduce heterogeneous agents that react differently to different cultural inputs, as we did in Section 3.7, and we have shown how this adds a level of complexity. In the current scenario another process has to be added. As the utility isn't given anymore by a value directly observable from the cultural artefact, it has to be approximated through other indirect mechanisms. In this study this mechanism is *trade*, as we presented it in Chapter 2. Trade is complex: agents have to compute a certain number of things and interact with a huge number of other agents before finding a trade agreement (*cf.* Section 2.3.2, Algorithm 2.3 and Equations 2.2 to 2.5 for more precision about these processes). Moreover, this activity is highly stochastic: the order in which who meets who and when is randomly chosen and change every time step. This implies that various iterations of the economic interactions must happen to give a robust estimation of the underlying utilities and not just random noise.

Therefore, the indirect bias toward the most successful makes the model much slower than the direct content model we presented in Chapter 3. We have seen in Section 3.6 the computation cost of the ABC rejection algorithm (*cf.* Algorithm 3.6) and how it makes the exploration of heavy model difficult. This is worst with success bias model. If we don't have the space problems raised by the specific structure of Twitter cascades, the time taken to run one simulation in the current setup is simply too long. In Figure 4.3, we represent the distribution of time taken by 500 simulations from the unbiased and the success bias scenarios. As one can see, when looking to the pic around $1,200$ seconds on right panel, simulations run with the success bias model can often last 20 minutes. This makes impossible to run even a small illustrative experiment with $400,000$ simulations, as we did for the content bias model presented in Section 3.7.

4.4.1 Population Monte Carlo

Various solutions have been developed to optimize the number of simulations needed to approximate the posterior distribution since the popularization of ABC. Beaumont et al. (2009) proposed one, known as *Population Monte Carlo* (ABCPMC). This method accelerates the computation of the posteriors by updating the priors used throughout the process and dynamically lowering the ϵ-threshold used to consider a simulation as close enough to the data. Briefly summarized, a set of decreasing ϵs is defined and for each of those ϵ, the posteriors found at the previous one will be used as priors for the next one. This process greatly increases the speed of the algorithm by drastically decreasing the number of simulations needed to find the posteriors. Moreover it allows to find posteriors even if they are not included within the range of the initial priors, something the rejection algorithm presented before could not do.

A python implementation of this algorithm has been written by Akeret et al. (2015). Nonetheless, this implementation relies on `mpi4py` to parallelize the exe-

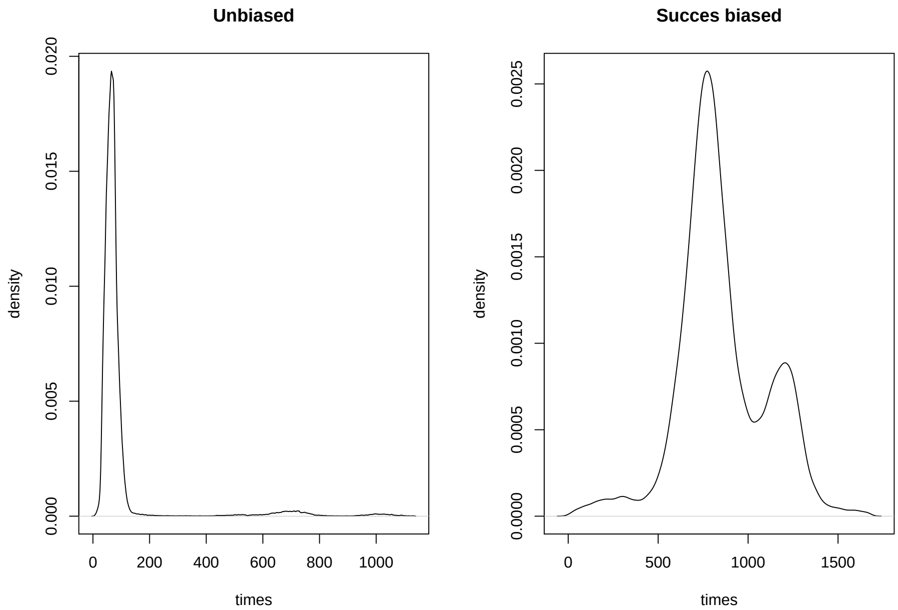

Figure 4.3: Distribution of time taken to run simulations for the unbiased (left) and the success biased (right) models. 500 simulations were run using parameters randomly sampled from the prior described in Table 4.3 and the time taken to finish each simulations recorded.

cution of the simulations. This is a good solution if the model is also implemented in python and needs to be run in simple environments, where all the resources needed are available at will and during an unlimited time. This is not our case, in this study the model is written in `C++` and have to be run on a shared supercomputing environment where resources are limited in space and time[1].

We thus had to reimplement almost entirely the `abcpmc` library, to interface it with the queuing systems we had access to in order to be able to launch our `C++` implementation of our model. An algorithmic description of the ABC is given in Algorithm 4.2 and the source code is available here: http://github.com/simoncarrignon/abcpandora.

Algorithm 4.2 ABC: The population Monte Carlo algorithm

INITIALIZATION:
$\epsilon = GenerateEpsilons()$ ▷ Generate a set of decreasing ϵs
$\theta_1 = GeneratePrior()$
RUN:
for ϵ_t in ϵ **do**
    while $pool.size < 500$ **do**
        $\theta_i = prior.genNewParam(\epsilon_t)$ ▷ Draw a vector of parameter from the prior
        $r = Model(\theta_i)$ ▷ Simulate the model
        $s = summary(r)$ ▷ Generate summary statistic
        if $\Delta(s, x) < \epsilon_i$ **then**
            $pool.add(r)$
        else
            $rm(\theta_i)$
    $prior = ModifyPriors(pool)$ ▷ Modify the prior using selected θs covariance matrix
return $pool$

[1] Thanks to BSC facilities the total amount of resources we had access to was almost unlimited. However, this access is limited in time and space by a queueing system. The ABC procedure has to handle this queuing system and respects the constraints it imposes. It should also be noted that the queue management system changed during the experiment ran for this chapter, passing from IBM platform LSF to SLURM workload manager.

4.4.2 Summary statistic and distance to data

As we have seen while exploring the Twitter data in Chapter 3, the key element to ABC approaches is the function used to calculate the distance to the data (the $\Delta(s, x)$ at Line 10 in Algorithm 4.2) and how the data (x) and simulations (s) are represented. In this chapter the nature of the data will raised different problems than the ones we faced in the previous study.

In the previous chapter we relied on a rich and detailed dataset. We could test our models against the full distribution of tweets and re-tweets with a great confidence that this distribution represented the real one. This is not the case with the ICRATES dataset: we know that the observations it contains represent only a tiny sample of what the real activity during Roman time. Thus, some assumptions have to be done when comparing our simulations with this part of what we know to be a bigger system.

Summary Statistic

Site sampling: The archaeological dataset we use was derived from 178 sites (*cf.* Figure 4.1). Because these are not assumed to represent the full population of urban settlements we simulate 500 agents. This number was chosen to ensure reasonable computation time and because it roughly approximates the estimated number of urban settlements in the Roman East (Wilson, 2011; Hanson, 2016; Hanson et al., 2017). To then compare this population of 500 agents with the 178 sites of the original dataset we normalize both metrics (patterns A and B in Figure 4.2) as a percentage of the total. Thus, instead of speaking about absolute *number* of sites or wares, we compare the *percentage* of sites and wares. This suppose that we assume the proportions described in the dataset as representative of the proportions we would find if we had access to the full urban settlements active during the period studied.

Time binning: To calculate the number of different types of ware present in one site at one time period, a duration (in year) has to be defined for this time period. Given the length in years of those time periods, the dataset can be divided in a finite number of periods that can be use to divide the results of the simulation in a similar way. If we choose to split the data in two periods, then the model will have to simulate two different periods and compare them to the data, if we decide to split the dataset in 100 periods, then the model will have to simulate 100 periods and the ABC will have to compare each one of them with the real data. On one side, having more periods makes the comparison between the model and the data longer, on the other hand, by splitting the data into a too small number of time periods, the patterns described in Section 4.2.2 disappear. The impact of binning the data within periods of different size is represented in Figure 4.4. As we can see on the left panel, the properties described before disappear to be replaced by straight lines, while binning the data within 50 or 200 periods doesn't change much the overall picture (while it greatly increases the computation time). Finding the right split means finding the good balance between loss of information and computation cost. We decided to divide the dataset into 50 periods (i.e. each period lasts 10 years long) as it was the smallest divisions that keeps as much information as possible. On the model side the periods will be represented by a different number of cultural interactions: every simulations will be split in 50 periods composed by an equal number of cultural interactions. The final number of cultural interactions per period is to be found by the ABC.

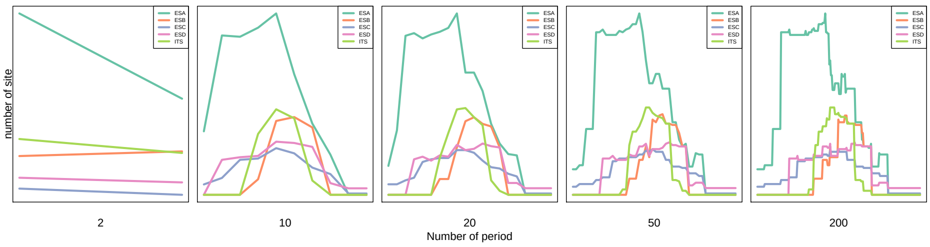

Figure 4.4: Impact of the size of the bin used to group data on the pattern observed.

Both dimensions, temporal and geographic, are crucial and can drastically change the way the distance to the data is calculated and what it measures exactly. Moreover, those two dimensions apply to both metric (patterns A and B, *cf.* Figure 4.2) independently and may not affect each one the same way. In order to simplify the code and the exploration we pre-process both patterns the same way (relative proportion and 50 periods).

In this seemingly harmless decision making process, a central aspect of the modelling approach we follow is hidden. Together with the choice of the distance function that we will describe later, this pre-processing of the observations bring back the three-sided problem we mentioned in Section 3.5.2. To choose the correct way to summarize the data, the precise implication of every aspect of the transformation of the empirical evidences has to be known (can we use the data sample as representative of the original population of urban settlement? can we bin together data from 10 different years?…). At the same time, one has to know the mathematical implications of these choices on the approximated probabilities (is it theoretically possible to approximate the posterior distribution of anything with $6 \times 5 \times 50$ dimensions using Monte Carlo methods?). Finally, one has to predict and handle the computation and storage cost that every trade-off will impose on the final stages of the ABC (e.g. while not binning the data may ensure that all essential information is kept, this will multiply by ten the time to compute the distance and by one thousand the space needed, a difference that can switch the exploration from slow to impossible).

This high level of transdisciplinarity is here again one of the major difficulty of this chapter. Without enough knowledge, trade-offs are done given what is possible more than what is the best. This is even more true in topics such as the one explored here, where the literature is scarce and scattered in wide and diverse scientific fields. We will see in the next paragraphs how we tried to reproduce the “empirical” exploration we proposed in the previous chapter in order to limit as much as possible arbitrary and sub-optimal decisions.

Distance to the data

While looking for the correct way to represent and summarize data and simulations, it is also important to choose a way to articulate them altogether within the right distance function. To do so we follow the same “empirical” procedure that we used in the previous chapter and presented in Section 3.5.2: we run thousands of simulations while keeping as much information as we can, and then apply various summary statistics and distance functions to see how they interact and which one suits best our needs.

Here we present results obtained with two different functions among the various we tried. The first one described by Equation 4.1 is the Euclidean distance between the data and the simulation at each period and for each measurement.

$$\delta(s,d) = \sqrt{\sum_{t=0}^{T}\sum_{i=0}^{W}(s_{i,t} - d_{i,t})^2} \times \frac{1}{T \times W} \tag{4.1}$$

The second one, described in Equation 4.2 uses a ’z-score’-like method to normalise the absolute difference between simulation and data by the mean and the standard deviation of this difference for each categories.

$$\delta^{'}(s,d) = \sum_{i=0}^{W} \frac{|s_{i,t} - d_{i,t}| - mean(|s_i - d_i|)}{sd(|s_i - d_i|)} \times \frac{1}{W} \tag{4.2}$$

In both functions, W is the total number of categories for the pattern observed: the type of ware for the pattern A, or the diversity for the pattern B. T is the number of periods between which the dataset and the simulation have been split and $mean(|s_i - d_i|)$ and $sd(|s_i - d_i|)$ are respectively the mean and standard deviation of the absolute difference between data and simulation for one category over all periods.

We calculate the distance using those two functions for both patterns and with different numbers of periods (from 10 to 100) for 2 500 simulations. An illustration of this exploration is given by Figure 4.5. For this figure we didn't used Equation 4.1 *as is*, but used a simpler variant which compute the mean over each period and category of the absolute difference (instead of the square root of the mean of the squared difference).

As we can see in this figure, the different distance functions are not capturing the same aspects of the data and are not affected equally by the way the period are binned together. Moreover, as shown on the right pane, it is hard to find simulation that are able to reproduce well both patterns at the same time.

This underlines the high dimensionality of the space we are trying to explore and illustrates well the difficulty we will have to capture with one simple metrics the whole picture we are trying to understand. Adding to this the time needed to run the model, and we have the perfect mix making our "empirical exploration" technique impossible to apply fruitfully. And indeed, finding the right function was not as simple as in the previous chapter and we ended choosing the one closest to the function used by Crema et al. (2016) and described by Equation 4.1. To take into account both patterns we simply took the mean between the two scores, as described in Equation 4.3.

$$\Delta(s,d) = \frac{\delta(s_a,d_a) + \delta(s_b,d_b)}{2} \tag{4.3}$$

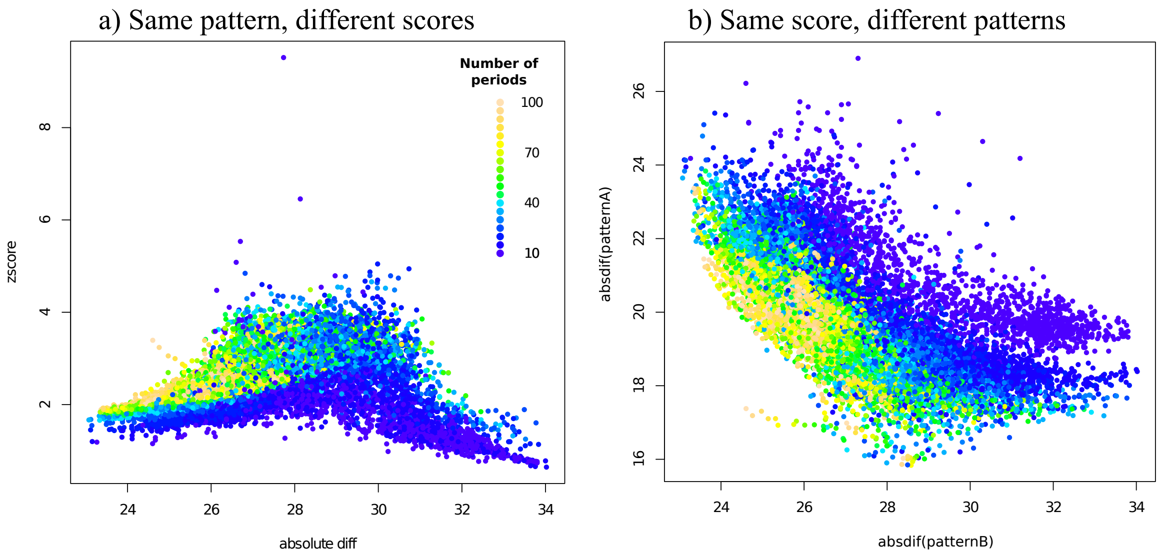

Figure 4.5: Empirical exploration of the score. The left panel (a) shows the influence of two difference distance functions applied to the same pattern (here the total number of site where each types were found). The right panel (b) shows how the two different patterns influence the result of the same distance function (absolute difference). The color of the points represents different time binning used to calculated the distance.

4.4.3 Parameters & Priors Distributions

Prior distributions were selected to cover wide but historically credible ranges given the technological constraints of the time: for the rates of innovation and social learning, distributions are uniform between zero and one; number of economic interactions (i.e. buying and selling) between one and three interactions per year; the prior for the number of cultural interactions (i.e. copying a trader's strategy) is set based on two constraints: at least two economic interactions take place between each two cultural interactions (to allow for information to be gathered), and there are a maximum of two cultural interactions per year.

Moreover, as we saw before, we should be able to divide the total number of cultural interactions in 50 periods that will be compared to the data. Thus, the total number of economic interactions and the total number of cultural interactions should respect the following inequality:

$$EC > 50 \times CI \tag{4.4}$$

The Table 4.3 describes every parameter that will be explored via ABC and as been built to follow the constraints listed before.

Parameters	Priors
μ	$U(0,1)$
μ_{max}	$U(0,10)$
λ	$U(0,1)$
λ_{str}	$U(0,10)$
EC	$U(50,1000)$
SC	$U(1,50)$

Table 4.3: Prior distributions for parameters inferred by the ABC. $U(X,Y)$ correspond to the uniform distribution between X and Y.

In addition to the parameters of the model, the ABC algorithm itself takes as input a decreasing sequences of ϵs that we give in Table 4.4.

4.5 Results

We ran Algorithm 4.2 for 13 different ϵs decreasing logarithmically (*cf.* Table 4.4), for the success-biased, unbiased and independent learning strategies

step	1	2	3	4	5	6	7
ϵ	0.13	0.011	0.0109	0.0108	0.0107	0.0106	0.0105
step	8	9	10	11	12	13	
ϵ	0.0104	0.0103	0.0102	0.0101	0.0100	0.0099	

Table 4.4: Value of epsilon for all step of the ABCPMC.

(*cf.* Algorithm 4.1). Each step of the ABCPMC was completed when 500 acceptable simulations (where acceptable means simulations that fall under the current threshold ϵ) were found. We therefore needed $6,500$ simulations to finish the whole ABC process (i.e. $13\epsilon \times 500$ acceptable simulations). To obtain $6,500$ acceptable simulations we had to run a total of $206\,902$ simulations for the independent learning model, $564\,211$ simulations for the unbiased model and $1\,267\,560$ simulations for the success-biased model.

We represent the percentage of simulations accepted for each model at every time step in Figure 4.6. As we can see, after the tenth step, the ratio of accepted simulations for the success biased social learning is already very low and just keeps going lower and lower. Given the time needed to run this model and the exponentially increase number of simulations need to go to the next step, we decided to stop the ABC at this thirteenth step, as shown in Figure 4.6.

4.5.1 Model selection

Using the approximation of the likelihood calculated by taking the simulation selected at the last time step of the ABC, we can compute the Bayes Factor of all models to formally select the more likely (Toni et al., 2009; Toni and Stumpf, 2010), as we did in Section 3.6.2.

We will note K the Bayes Factor between model m_1 and m_2 as:

$$K_{m_1,m_2} = \frac{P(D|m_1)}{P(D|m_2)} \tag{4.5}$$

Where $P(D|m_i)$ is the likelihood of m_i, as approximated by the ABC.

The Table 4.5 summarize all the Bayes factor between all models. For each line i, the Bayes ratio in each column j correspond to the ratio between m_i and m_j:

$$K_{m_i,m_j} = \frac{P(D|m_i)}{P(D|m_j)}$$

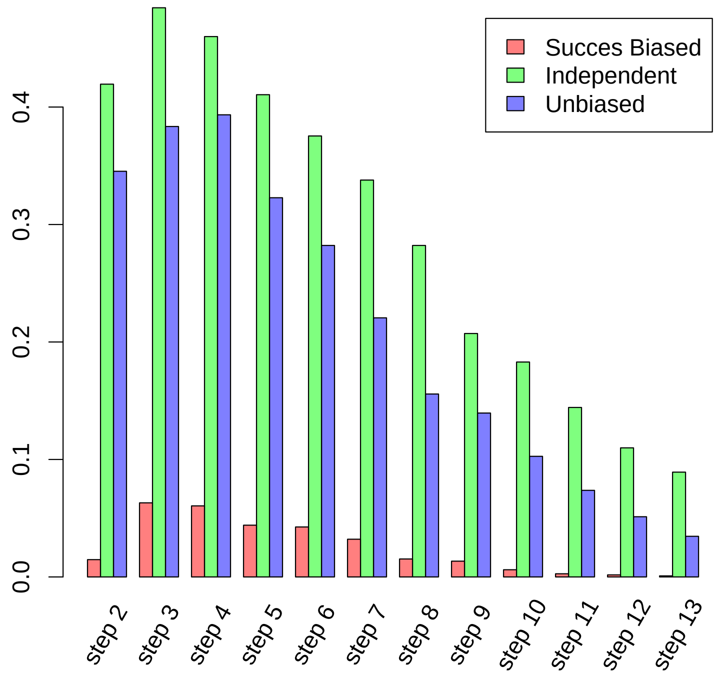

Figure 4.6: Evolution of the ratio of accepted simulations among the total number of simulations for each step of the ABC algorithm and for the three models. The first step has been removed as it represents an ϵ big enough to accept any simulation, then the three models present the same ratio of one for this step.

A common framework to interpret those ratios is given by Kass and Raftery (1995). Following their scale, independent learning is the more likely mechanism. It is twice most likely than the unbiased learning and 20 times more likely than the success bias.

The simulation and the population Monte Carlo version of approximate Bayesian computation (ABCPMC) results reveal that the independent learning model produces more good simulations and those good simulations are closer to the empirical dataset than both other models.

4.5.2 Posterior Distributions

The Approximate Bayesian Computation allows use to calculate an approximation of the posterior distribution of the model described by the parameters $\theta = [\theta_1, ..., \theta_p]$ given the data d: $P(\theta|d)$.

With this, we can represent the marginal posterior distribution of each parameters $\theta_1, ..., \theta_p$ for all models and compare them with the priors. Those priors, as well as each parameters studied in this experiment are described in Table 4.3. We represent all posteriors together with the priors of all models in Figure 4.7. The posterior distributions were drawn using the 500 accepted simulations from the last ABC step ($\epsilon = 0.0099$, *cf.* Table 4.4).

The posterior distributions for the independent learning and the unbiased models are very similar as we can see by comparing the green and blue curves in Figure 4.7. One parameter is obviously an exception to this: the rate of social learning, i.e. the probability of social update μ. It's not used in the independent learning as this strategies doesn't depend on social interaction thus it stays flat and equal to the prior distribution for the independent learning. The fact that the rest of posteriors for both model are similar is in line with the Bayes factors shown in Table 4.5 where those two model have a ratio close to one.

Nonetheless, as we have shown that independent learning is the hypothesis

	Unbiased	Independent	Success Biased
Unbiased	1	0.5	12
Independent	1.96	1	23
Success Biased	0.08	0.04	1

Table 4.5: Bayes factor for the three model

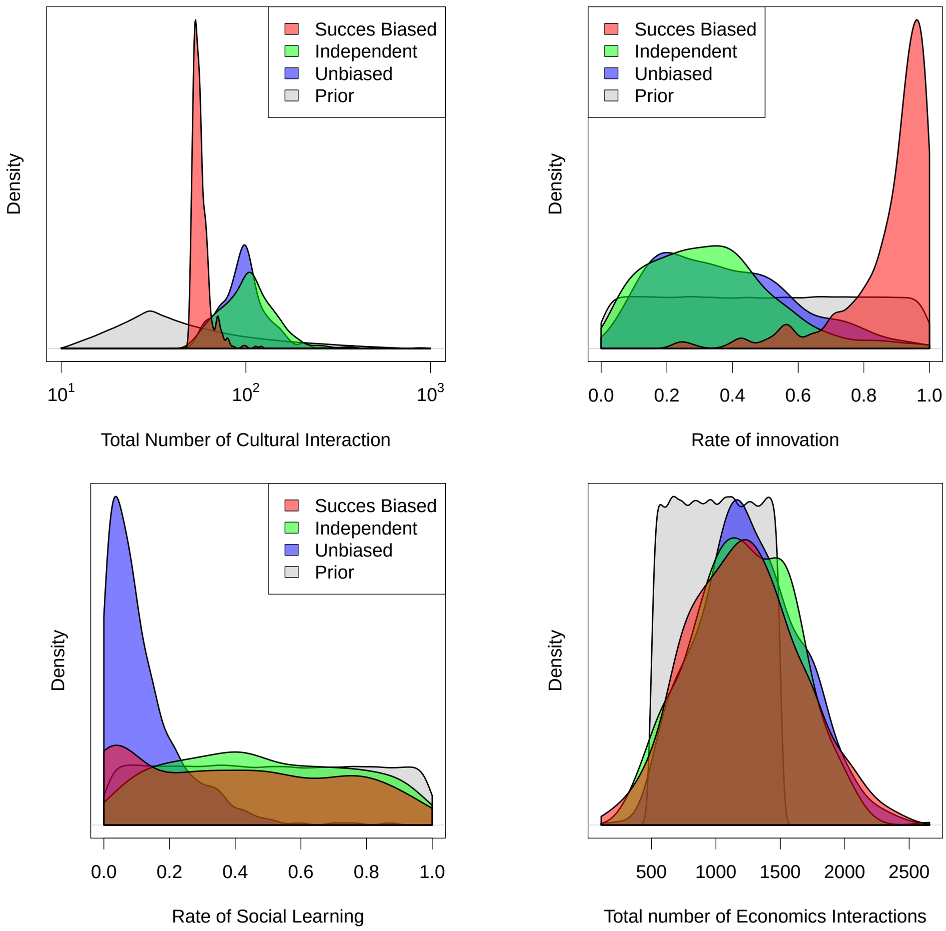

Figure 4.7: Comparison of the posterior distribution of four parameters of the three model.

better supported by the empirical data, we interpret more in detail the posterior distributions of the independent learning model only. To help the interpretation we determine the distribution's mode (the value with the highest probability) and the high density region (HDR) representing the parameter space within which 75% and the 95% of acceptable simulations lie (Figure 4.7). Were represent the most interesting posteriors distribution together with the mode and the 75% and 95% HDR for the independent model in Figure 4.8.

In the following technical overview of posterior distributions of key parameters for our hypotheses we focus on the results of the independent-learning model, because it offers simulations closer to the archaeological data:

Total number of economic interactions: the number of times agents go to the market to buy tableware during the whole simulation. The 75% HDR falls between 750 and 1700 economic interactions, i.e. 1.5 to 3.4 times per year.

Total number of cultural interactions: the number of times agents had the opportunity to copy strategies from other agents. The 75% HDR falls between 63 and 140 cultural interactions, i.e. once every 7.9 to 3.6 years.

Rate of innovation: the probability at each time step that an agent changes its strategy independent of other agents. The 75% HDR falls between probabilities of 0.088 and 0.51.

4.6 Discussion

The archaeological data presented in section 4.2.2 revealed clear evidence of influence of tablewares on each other (through distributions, stamps and morphologies), suggestive of a competitive market. Nonetheless the results of our ABM, in which cultural copying mechanisms were intertwined with economic mechanisms, and of the comparison of the three models' simulation results to the archaeological data using ABC, suggested that trader's innovation are independent with regards to each other's buying strategies. We observed that the model with the independent learning mechanism, in which commercial agents independently change their tableware buying strategies, was the more likely to generate simulations closer to the archaeological data patterns. The two other models, where agents updated their strategy by copying strategies from other agent, were not able to do so.

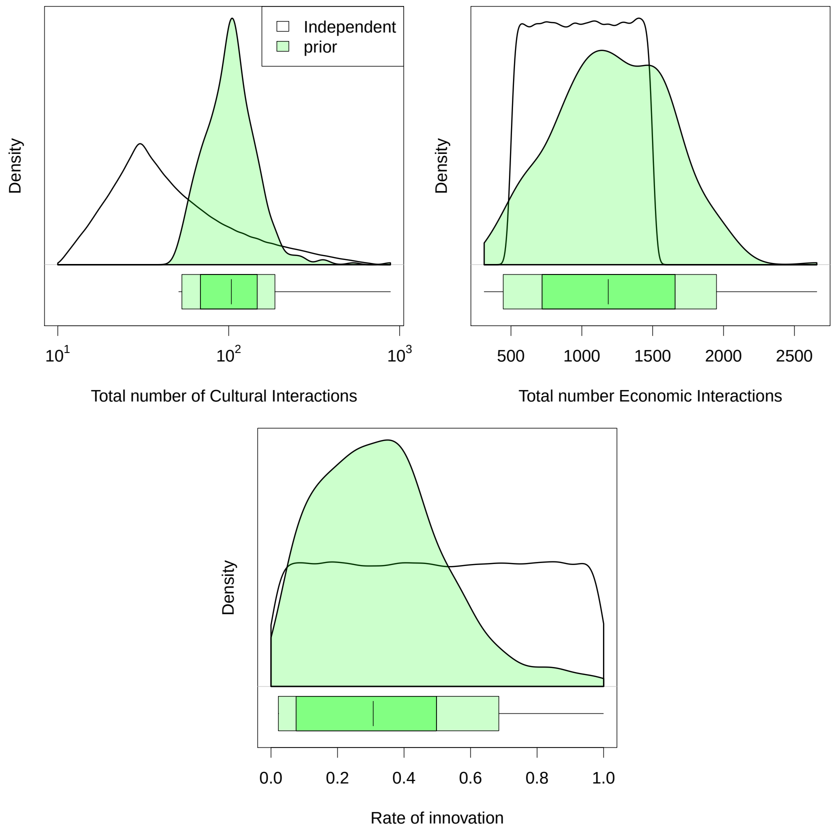

Figure 4.8: Comparison of the marginal posterior distribution of three parameters for the independent learning model. The green area represent the posterior distribution for the model while the white one represents the prior. At the bottom of each graph we represent the 75% HDR (darker green) and the 95% HDR (lighter green) as well as the mode of the distribution as a vertical line.

This result has important implications for the study of ancient inter-regional trade in the study area, and for Roman economy studies more broadly. It shows that copying the strategies of successful traders within the tableware market may not have played a central role in inter-regional tableware trade and whereas independent, uninformed and local adjustment seems to have leading the trade of this good. The reason could be that long-distance distribution of tableware was intertwined with major trade flows from foodstuff production (e.g. present-day Tunisia and Egypt) and mining and quarrying regions (e.g. present-day southern Spain and the Egyptian eastern desert), where tableware was one of the additional cargoes to fill up empty spaces in ships' hulls. Thus important mechanisms for explaining aspects of inter-regional trade in other goods such as foodstuffs, stones and metals may reveal other processes that could better explain the pattern observed.

These products made up the bulk of all long-distance trade in the Roman world, and their study is key to understanding the Roman economy. But large open access comprehensive datasets of centuries-long amphora container or stone distribution data for the entire Roman Empire or significant parts of it are currently few. Our findings, by quantifying credible mechanism for trade in craft products, present a prospect for contributing to this key topic in Roman archaeology and history by using tableware data. The precise nature of the link between trade in craft products and other goods and the use of cultural/economic transmission models with other evidence of inter-regional trade such as distributions of amphora containers and stones should be explored in future work.

Interpreting the posterior distributions resulting from the ABC allowed to precise those interpretation and identify credible parameter values for the modelled hypotheses. Although the credible parameter ranges are wide, this is not unexpected in a study of an economic system that functioned two millennia ago where very little information is available to fix tested parameter ranges as prior distributions. However, they still allow us to add an, for Roman economy studies, unprecedented level of specification to our theory of independent economic innovation. In acceptable simulations the traders bought tableware from other traders or producers around 1.5 to 3.4 times per year (75% HDR). Such limited frequencies are certainly historically and archaeologically supportable, given the significant limitations on the frequency of obtaining products from other parts of the empire posed by the then-current transport technologies and the financial requirements to organise inter-regional shipping. However, the suggested rate at which traders update there tableware buying strategies is far more limited, ranging from every 1.6 years to every 16 years (75% HDR).

To stick closer to the historical context future studies should focus on comparing these time estimates with what we know about ancient transport and communication infrastructure and technology, exploring the implications of this theory for the mobility and activity of commercial actors active in inter-regional trade in craft products in late Hellenistic and Roman times.

Our results revealed little about the role of ITS and we believe this western-produced ware should be the focus of future computational modelling research. Although we performed experiments to specifically explore how the presence of ITS might have stimulated competition on the eastern market, our approach did not succeed in identifying any effects other than those presented above. This should be further explored in future analysis, alongside exploring whether success-bias is a particularly viable theory for smaller time period, and more precisely the period 40 BC to AD 150, when ITS was distributed in the eastern Mediterranean. We also believe that the nature and processes of possible stylistic imitation of ITS features by eastern wares should be explored from a cultural transmission perspective using the methods we applied in this study.

At another level those results suggest that an indirect, success biased social learning isn't a good candidate to explain cultural changes linked to economic activity during large scale periods. Nonetheless some important methodological points should be observed more carefully and taken into account before discarding definitely the influence of such learning strategies. The first obstacle to the empirical detection of such bias is the computation time needed, which is even higher than the direct content bias we explored before. However we presented here an ABC algorithm that allowed to partially solve this problem. But in the meantime, the nature of the data raised another kind of problems. The high dimensionality, the noise and the low resolution of the record studied made difficult to find the right way to summarize and compare the data and the simulation in the right way, Trying to solve this issue we opted for a wrong option where we decided to keep as much information as possible while at the same time trying understand with a unique general model two very different patterns unfolding during long periods of time. To do this, the function we used (*cf.* Section 4.4.2) had to integrate at the same level very different periods that may have been driven by very different processes. The choice we made averaged under the same umbrella periods where four wares are competing altogether with periods where just two different types of ware are presents. This led the ABC to prefer a general random model able to cope partially with all part of the dataset, instead of a model that could better reproduce the data bu only at a small level.

A more careful exploration should look at both patterns and all periods as separate processes. ABC should be used at a lower level, to understand which model could generate the distribution of goods found at every time periods, as we suggest for the ITS.

4.7 Conclusion

When studying the impact of content-dependent learning on complex, long term interactions such as trade, Cultural Transmission studies usually tend to be theory driven. They analyse theoretically outcomes of different models and compare them together and rarely look back at the empirically record. In this chapter we present an approach that explores these theories in light of archaeological research questions raised by a robust observed archaeological data pattern. Such questions and data-driven selection of methodological and theoretical frameworks should become more widely applied in the historical disciplines and particularly in Roman Studies, where computational modelling is exceptionally rare. We believe this study demonstrates how computational modelling combined with ABC and a robust archaeological data pattern can make highly constructive contributions to theory evaluation, building and specification in Roman studies.

Social learning strategies such as success bias are expected to have played a crucial role in human evolution. Archaeological studies offers a unique occasion to test those hypotheses at large temporal and geographical scales. The overall study suggests that success bias is not the main driver of trade dynamics. The relevance of such bias cannot be fully discarded for various reasons. Throughout this chapter we have suggested that the broad range of data and the very diverse nature of the underlying processes require a careful articulation between data and theoretical model. Our modelling choices may have partly hidden the impact of success bias. Specifically, keeping every dimension of the dataset (i.e. the two patterns and the 500 years) in the distance function may have averaged out data fluctuations, which explains the rejection of the content-dependent hypothesis. Solving this issue requires additional empirical evidences.

This chapter provides a solid guide for future work to quantify social learning bias in archaeological studies and thus shows how Cultural Evolution can be a powerful tool to explore historical and archaeological questions. By phrasing archaeological and historical hypothesis in terms of Social Learning Strategies, it's possible to quantify and test which hypothesis is the most likely and reject the other ones. The framework proposed here allow to do so, even if the observations available are scarce and noisy and the hypotheses implies complex processes.

Chapter 5

CONCLUSION

5.1 His choice, my choice, our choices

Companies spend thousands of euros to be sure famous tennis men and women wear their clothes and exhibit their logos at every interview. We can trust companies that if they do so, they know they will get their money back and much, much more. And they know they will get their money back because they expect people to be influenced by what their idols do. If Roger Federer looks more classy than ever and as young as he was twelve years ago, it may be thanks to his training and his genetics. But it may also have to do with the white hat he always wears. Who really knows? The two first aspects aren't easy to fulfil, while buying a hat isn't such a big deal. With the approaching summer, you were going to buy a hat anyway, so why not buying the same than Roger? On the other hand, the day you will have to buy your next laptop and before spending almost a month of salary on it, you may think twice before choosing one simply because you saw Kylian Mbappé happily playing with the keyboard. You will spend hours checking and learning on Youtube about all specifications of every laptop on the market. You will look for their speed, memory size, battery charge, and so on and so forth, until you are sure to know exactly the "intrinsic value" of the laptop, and make the best choice[1].

[1] I know you could be more interested by the intrinsic value of a hat because you are a hardcore fisher and prefer to buy a macbook just because it's cool to do so, but as you went through this whole document I assume that you are more likely the kind of person who looks carefully at laptops' characteristics, sorry if I was wrong and in this case, just swap the two concepts.

Both scenarios will bias your behaviour at the moment you will take the final decision. But if both biases will for sure make companies richer, what if they had a deeper impact on how the world changes? What if instead of following your favourite sportsmen or sportswomen, you had spend hours of learning before buying the hat, as you did with the laptop? What if you had a deep look at how the hat was done, by who and where? To learn how much energy has been wasted to produce and transport it? And what if instead of focusing on the technical characteristics of your laptop you had done the same, looking at the conditions of its production, to be sure they limit their footprint on the environment and respectfully employ the workers producing it? This time the bias that would impact your final decision, multiplied by the millions of person it potentially influences, may have a great impact and change forever the future of earth as we know it. And if we jump one instant in this future: will our future selves be able to detect, while digging our fossilized garbage, that such changes were led by our knowledge of the intrinsic value of our cultural artefacts? Will they be able to detect the difference between the increase of a particular brand presence in our wardrobe due to Roger Federer's Grand Slam victories versus the increase led by our tendency to do as everyone else does?

Throughout this thesis we have shown examples of such biases and how all those questions can be answered. We proposed a computational framework that bind together data and theory to quantify and detect the role of the different social learning strategies that impact what we learn, how we learn it and from who. Thanks to this framework we have shown that yes, such bias *can* have a great impact at the population level, yes, it is possible to detect them by looking at empirical observations and yes, we can determine which bias was leading the changes observed.

5.2 Theoretical choices

We started in the Chapter 2 to explore the theoretical implications of such biases. Thank to a computational framework specially developed to this end, we compared two different biases: (1) neutral (or content-agnostic) and (2) success bias (or indirect content-dependent). We have shown that our model reproduces the well known fact that neutral bias leads to a power-law distribution of cultural traits when there is not any dependence on intrinsic value or explicit selection advantages (Bentley et al., 2004; Mesoudi and Lycett, 2009). On the other hand, we have shown that success-biased social learning keeps a relatively larger diversity of variants and exhibits a different signature than the neutral model.

In addition, we illustrated how Agent-Based model is a suitable approach to rigorously assess hypotheses in Social Learning and Cultural Evolution. The flexible and modular formulation of ABM allows to combine theories from different fields, as well as to predict and compare the outcome of different hypothesis. This can even be done probabilistically, by comparing the outcome of different models with the expected outcome of known theories, as we did with the Fitting to Idealized Outcome. Moreover, such modelling methods permit to define realistic individual decision making process and bring together micro and macro level of analysis, which revealed as promising approach for the study of Economy (Tesfatsion, 2003, 2006).

More interestingly, this chapter gave a great example on how Cultural Evolution and Economy can be integrated together trough computational modelling. While indicating how trade could be interpreted as a particular mode of prestige-biased cultural transmission, our study suggests a deeper link between trading models with cultural evolution models. The study appears as an interesting way to approach Economy as a special case of Culture.

We have also shown how our framework can be used to explore the underlying networks that structure cultural and economics interactions. This is a promising path to follow, as networks are know as a powerful tool to compare models and data, allowing a straightforward correspondence between observed and theoretical entities (Morer et al., 2018; Valverde et al., 2007; Brughmans, 2010). In this chapter we have been able to measure the effect of the network's topology in social learning dynamics and we show how the average degree of the node can greatly increase the speed to reach the equilibrium.

5.3 Online choices

In a third chapter we have tested variations of culture-evolutionary neutral models on aggregated Twitter data documenting the spread of true and false information. We used Approximate Bayesian Computation to resolve the full joint probability distribution of models with different social learning biases emphasizing context-biased versus content-biased learning. Doing so we find that the neutral model applies well to the observed re-tweeting activity, and better than models with added conformity bias.

In the meantime we he have shown how online social learning dynamics can be modelled through the tools of cultural evolutionary theory and how we can test hypothesis about those dynamic against real data. We have shown how models implementing these hypotheses can be precisely adjusted and compared to available record of online activity. We have exposed some problems ABC raises when it is applied to such dataset, where complex models reproduce with realism processes that need lot of space to store precise information and last longer.

Overall, the Chapter demonstrated how Culture Evolution can be used to understand social phenomenon in online social media, which is crucial given how they are heavily reshaping our social and cultural environment. On the other side, this chapter also demonstrated how thanks to the digital nature of online social media, they can be turned into great experimental tools to test and explore hypotheses on Social Learning and Cultural Evolution.

Here again, the use of the networks behind the spread of news could help to uncover properties that our model where not able to capture. Implementing and testing specific topology would help to reproduce the highly-skewed nature of influence on social media (Lieberman et al., 2005; Ormerod et al., 2012), following a fruitful tradition in studying cascading event (Watts, 2002).

5.4 Ancestral choices

In a last chapter we have demonstrated how computational modelling combined with robust data analysis is a powerful tool for theory evaluation in archaeology. The novelty of our approach is two-fold: (1) it provides a new example among the rare uses of computational modelling in historical Roman Studies and (2) we first tested the impact of success bias in large spatial and temporal scales. Both elements have great implication for the different fields implied. Computational modelling are crucial in archaeology: they allow to deal with the nature of the data, which is often noisy and missing while formalising verbally defined hypothesis (Costopoulos and Lake, 2010; Lake, 2014). Our approach shows how this can be done, and how different historical hypothesis can be quantitatively tested. On the other hand, exploring past society and archaeological data set is the only way to test the prediction made by cultural evolution (Shennan, 2002; Boyd and Richerson, 2005).

Our study suggests that success bias may not be the main driver of tableware trade in the Roman East. There are several possible explanations for this. Our model may have minimized the action of success bias even if it played a role in driving trade behaviour. In this case, additional data and different experimental scenarios are required to validate the hypothesis. More importantly, our fitting of the ABC model provides little evidence for social learning. Such pattern of independent economic behaviour is consistent with the significant technological and communication constraints and the lack of reliable information available to traders of the past.

This chapter gave a great example on how it is possible to track down complex social learning biases and their impact over economic changes at large scales by combining computational models, economic theories, cultural evolution and archaeological data. This approach has been shown as a good candidate to help us understand the economic crashes scattered throughout human history and to anticipate those to come.

5.5 Future choices, Evolving choices

Why and how social learning evolved is still a unsolved question. Is it just a happy by-product of the evolution of other cognitive capacities (Heyes, 1994, 2012) in an exaptation-like fashion (Gould and Vrba, 1982), was it a central element in our capacity to develop a cumulative culture where the technological advance of the past are used to develop more and more complex tools (Boyd et al., 2011), or did it evolve hand in hand with language to gossip about others’ reliability and punish the cheaters, in order to enhance human large groups cooperation (Dunbar, 1998; Dunbar and Shultz, 2007)? All those questions are still to be answered, but they all highlight the importance of the study of social learning to understand the Human specie. And about this importance of social learning one thing is sure: it has changed Human Societies and Human Societies have changed the way they learn socially (Henrich and Gil-White, 2001; Boyd and Richerson, 2005) and this thesis once again illustrates it. From epochs when social learning was bounded to local scales, limited by means of transportation and communication’s time between individual within different social groups, when inter-regional learning was almost absent in the Roman East where people had to rely on individual learning and small scale social interactions. To nowadays societies, trapped in a world wrapped up by social interactions, in which every one with a smartphone in his pocket can communicate and learn from others all over the world, at anytime. Between the two periods, the rate, the speed and the scale of social learning have incredibly increased. Where will this acceleration bring us isn’t easy to guess. We passed from societies where a few experts were interrogated by the limited members of their tribe about the plants that heal and the poisonous one (Henrich and Broesch, 2011) to a society where a few tweets from a few well known public figure are retweeted by millions of people worldwide and can lead to drastic political changes. In such society, Instagram’s “influencers” are paid proportionally to their number of followers and how they can weight social learning ; in this new society money is spend on bots and troll farms all around the world to manipulate cultural transmission (Ruck et al., 2019; Grinberg et al., 2019).

Some researchers argue that this acceleration of Cultural Change (Bentley and O’Brien, 2017), which is mainly powered by technological advances, will be handled thanks to other technologies (such as AI and big data mining). These other technologies should allow us to cope with the exponentially growing flow of information that we have to deal with (this thesis gives another example of such methods). But so far nothing is clear and the only thing we can advance with confidence is that the social learning strategies we use to learn the way we speak, the way we think, what we like and what we don’t, will continue to change and evolve, tied with the Human specie and its technologies. Cultural transmission will continue to transform, to be transformed, and to shape the future of Humanity, for the better or for the worst.

Future works should continue in that direction, testing at large scales predictions and models developed by experimental and theoretical studies. More complex and finely defined representation of individual choices making process and Social Learning Strategies could be implemented, such as the ones observed by Harrison (2018) or Brand (2017), and those mechanisms should be tested in a greater variety of empirical scenarios, such as those explored by Gardent (2017), Coto-Sarmiento (2019) or Sobchuk (2018). Such collaborations will be greatly eased by the common background that starts to grow with the birth of Cultural Evolution as a fields on itself.

Within the framework developed in this thesis, it becomes possible to integrate more realistic cognitive systems that would more faithfully reproduce the nature of social learning, which, as argue tenants of the Cultural Attractors Theory (Morin, 2015; Buskell, 2019; Scott-Phillips et al., 2014), is crucially missing in traditional Cultural Evolution works. This theory advances that perfect imitation isn’t important for the spread of culture while the cognitive process able to reconstruct the transmitted knowledge are central. They recently proposed theoretical models (Acerbi et al., 2019) that still need to be tested against the data and compared with other models of cultural evolution. The work proposed here could be used to do so.

Overall, by more precisely defining and increasing the number of processes and hypotheses studied in more heterogeneous geographic and temporal areas, our method will be a powerful device to unveil the common and specifics mechanisms that drove the evolution of Human Cultures toward their actual complexity and diversity.

Bibliography

Acerbi, A. (2019). Cognitive attraction and online misinformation. *Palgrave Communications*, 5(1):15.

Acerbi, A. and Bentley, R. A. (2014). Biases in cultural transmission shape the turnover of popular traits. *Evolution and Human Behavior*, 35(3):228–236.

Acerbi, A., Charbonneau, M., Miton, H., and Scott-Phillips, T. (2019). Cultural stability without copying. Url of the preprint: https://osf.io/vjcq3/.

Acerbi, A., Enquist, M., and Ghirlanda, S. (2009). Cultural evolution and individual development of openness and conservatism. *Proceedings of the National Academy of Sciences*, 106(45):18931–18935.

Acerbi, A. and Mesoudi, A. (2015). If we are all cultural Darwinians what's the fuss about? Clarifying recent disagreements in the field of cultural evolution. *Biology & Philosophy*, 30(4):481–503.

Akeret, J., Refregier, A., Amara, A., Seehars, S., and Hasner, C. (2015). Approximate Bayesian computation for forward modeling in cosmology. *Journal of Cosmology and Astroparticle Physics*, 2015(08):043.

Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*, volume 3. Princeton University Press.

Baldini, R. (2012). Success-biased social learning: Cultural and evolutionary dynamics. *Theoretical Population Biology*, 82(3):222–228.

Bang, P. F. (2008). *The Roman bazaar, a comparative study of trade and markets in a tributary empire*. Cambridge university press, Cambridge.

Barakzai, A. and Shaw, A. (2018). Friends without benefits: When we react negatively to helpful and generous friends. *Evolution and Human Behavior*, 39(5):529–537.

Beaumont, M. A., Cornuet, J.-M., Marin, J.-M., and Robert, C. P. (2009). Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990.

Bentley, R. A., Caiado, C. C., and Ormerod, P. (2014a). Effects of memory on spatial heterogeneity in neutrally transmitted culture. *Evolution and Human Behavior*, 35(4):257–263.

Bentley, R. A., Hahn, M. W., and Shennan, S. J. (2004). Random drift and culture change. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(1547):1443–1450.

Bentley, R. A., Lake, M. W., and Shennan, S. J. (2005). Specialisation and wealth inequality in a model of a clustered economic network. *Journal of Archaeological Science*, 32(9):1346–1356.

Bentley, R. A., Lipo, C. P., Herzog, H. A., and Hahn, M. W. (2007). Regular rates of popular culture change reflect random copying. *Evolution and Human Behavior*, 28(3):151–158.

Bentley, R. A. and O’Brien, M. J. (2017). *The acceleration of cultural change: from ancestors to algorithms*. MIT Press.

Bentley, R. A., O’Brien, M. J., and Brock, W. A. (2014b). Mapping collective behavior in the big-data era. *Behavioral and Brain Sciences*, 37(1):63–76.

Bentley, R. A. and Ormerod, P. (2010). A rapid method for assessing social versus independent interest in health issues: A case study of ‘bird flu’ and ‘swine flu’. *Social Science & Medicine*, 71(3):482–485.

Bentley, R. A., Ormerod, P., and Batty, M. (2011). Evolving social influence in large populations. *Behavioral Ecology and Sociobiology*, 65(3):537–546.

Bentley, R. A. and Shennan, S. J. (2003). Cultural Transmission and Stochastic Network Growth. *American Antiquity*, 68(3):459–485.

Bes, P. (2015). *Once upon a Time in the East. The Chronological and Geographical Distribution of Terra Sigillata and Red Slip Ware in the Roman East. Roman and Late Antique Mediterranean Pottery 6.* Archaeopress, Oxford.

Bes, P., Willet, R., Poblome, J., and Brughmans, T. (2018). Inventory of Crafts and Trade in the Roman East (ICRATES): database of tableware.

Bettinger, R. L. and Eerkens, J. (1999). Point Typologies, Cultural Transmission, and the Spread of Bow-and-Arrow Technology in the Prehistoric Great Basin. *American Antiquity*, 64(2):231–242.

Boyd, R. and Richerson, P. J. (1988). *Culture and the evolutionary process*. University of Chicago Press.

Boyd, R. and Richerson, P. J. (2005). *The origin and evolution of cultures*. Oxford University Press.

Boyd, R., Richerson, P. J., and Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(Supplement 2):10918–10925.

Brand, C. O. (2017). *Sex differences in social learnning:exploring the links with risk aversion and confidence*. PhD thesis, University of St Andrews.

Brock, W. A., Bentley, R. A., O'Brien, M. J., and Caiado, C. C. S. (2014). Estimating a Path through a Map of Decision Making. *PLOS ONE*, 9(11):1–10.

Brock, W. A. and Durlauf, S. N. (2001). Discrete Choice with Social Interactions. *The Review of Economic Studies*, 68(2):235–260.

Brughmans, T. (2010). Connecting the dots: towards archaeological network analysis. *Oxford Journal of Archaeology*, 29(3):277–303.

Brughmans, T., Hanson, J., Mandich, M., Romanowska, I., Rubio-Campillo, X., Carrignon, S., Collins-Elliott, S., Crawford, K., Daems, D., Fulminante, F., de Haas, T., Kelly, P., Moreno Escobar, M. d. C., Paliou, E., Prignano, L., and Ritondale, M. (2019). Formal Modelling Approaches to Complexity Science in Roman Studies: A Manifesto. *Theoretical Roman Archaeology Journal*, 1:4.

Brughmans, T. and Poblome, J. (2016). Roman bazaar or market economy? Explaining tableware distributions through computational modelling. *Antiquity*, 90(350):393–408.

Buskell, A. (2019). Looking for middle ground in cultural attraction theory. *Evolutionary Anthropology: Issues, News, and Reviews*, 0(0).

Byers, B. E., Belinsky, K. L., and Bentley, R. A. (2010). Independent Cultural Evolution of Two Song Traditions in the Chestnut-Sided Warbler. *The American Naturalist*, 176(4):476–489. PMID: 20712515.

Caiado, C. C., Brock, W. A., Bentley, R. A., and O'Brien, M. J. (2016). Fitness landscapes among many options under social influence. *Journal of Theoretical Biology*, 405:5–16. Advances in Modelling Biological Evolution: Linking Mathematical Theories with Empirical Realities.

Carrignon, S. (2018). Agent Based Modeling and Bayes Inference to learn about the past: the need for High Performance Computing. In *Conference on Complex System* .

Carrignon, S., Bentley, R. A., and Ruck, D. (2019). Modelling rapid online cultural transmission: Evaluating neutral models on twitter data with approximate bayesian computation. *Palgrave Communication*.

Carrignon, S., Bentley, R. A., Ruck, D., and Gilchrist, M. (2018a). How the intrinsic value of information can change the spread of news in social media . In *2nd Conference of the Cultural Evolution Society*. Arizona State University.

Carrignon, S., Brughmans, T., and Romanowska, I. (Forthcoming). Transmission of cultural and economic strategies in inter-regional tableware trade. In Brughmans, T. and Wilson, A., editors, *Simulating Roman Economies. Theories, Methods and Computational Models*. Oxford: Oxford University Press.

Carrignon, S., Montanier, J.-M., Michaud, J., and Rubio-Campillo, X. (2016). Co-evolution of culture and trade : impact of cultural network topology on economic dynamics. In *44th Computer Applications and Quantitative Methods in Archaeology Conference (CAA 2016)* .

Carrignon, S., Montanier, J.-M., and Rubio-Campillo, X. (2015). Modelling the Co-evolution of Trade and Culture in Past Societies. In *Proceedings of the*

2015 Winter Simulation Conference, WSC '15, pages 3949–3960, Piscataway, NJ, USA. IEEE Press.

Carrignon, S., Romanowska, I., and Brughmans, T. (2018b). An Agent-Based Model of Trade in the Roman East (25 BC – AD 150). In *28th Theoretical Roman Archaeology Conference (TRAC)*.

Carrignon, S. and Rubio-Campillo, X. (2017). Impact of different social learning mechanisms on the emergence of a Walrasian Equilibrium. In *European Human Behavior Evolution Association conference (EHBEA)* .

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

Charbonneau, M. (2013). *L'analogie de l'hérédité culturelle : fondements conceptuels de la théorie de la double hérédité*. PhD thesis, Université de Montréal.

Clauset, A., Shalizi, C. R., and Newman, M. E. (2009). Power-law distributions in empirical data. *SIAM review*, 51(4):661–703.

Costopoulos, A. and Lake, M. W. (2010). *Simulating change: archaeology into the twenty-first century*. University of Utah Press, Salt Lake City.

Coto-Sarmiento, M. (2019). *Cuantificando el cambio cultural*. PhD thesis, Universitat de Barcelona.

Crema, E., Edinborough, K., Kerig, T., and Shennan, S. (2014). An Approximate Bayesian Computation approach for inferring patterns of cultural evolutionary change. *Journal of Archaeological Science*, 50:160–170.

Crema, E. R. and Kandler, A. (2019). *cTransmission: An Approximate Bayesian Computation framework for inferring patterns of cultural transmission from frequency data*. R package version 0.1.0.

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

Crone, E. A. and Konijn, E. A. (2018). Media use and brain development during adolescence. *Nature Communications*, 9(1):588.

Csilléry, K., Blum, M. G. B., Gaggiotti, O. E., and François, O. (2010). Approximate Bayesian Computation (ABC) in practice. *Trends in Ecology & Evolution*, 25(7):410–418.

Danvers, A. F., Hackman, J. V., and Hruschka, D. J. (2019). The amplifying role of need in giving decisions. *Evolution and Human Behavior*, 40(2):188–193.

Darwin, C. (1859). *On the origin of species by means of natural selection, or, the preservation of favoured races in the struggle for life*. John Murray, London.

Derex, M. (2013). *Les mécanismes de l'évolution culturelle cumulative*. PhD thesis, Université Montpellier II.

Dunbar, R. I. M. (1998). *Grooming, gossip, and the evolution of language*. Harvard University Press.

Dunbar, R. I. M. and Shultz, S. (2007). Evolution in the Social Brain. *Science*, 317(5843):1344–1347.

Durham, W. H. (1992). *Coevolution: Genes, Culture, and Human Diversity*. Stanford University Press.

Eerkens, J. W. and 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.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., Weeg, C., Larson, E. E., Ungar, L. H., and Seligman, M. E. P. (2015). Psychological Language on Twitter Predicts County-Level Heart Disease Mortality. *Psychological Science*, 26(2):159–169. PMID: 25605707.

Epstein, J. M. and Axtell, R. (1996). *Growing artificial societies: social science from the bottom up*. Brookings Institution Press.

Evans, T. S. and Giometto, A. (2011). Turnover Rate of Popularity Charts in Neutral Models. *arXiv e-prints*, page arXiv:1105.4044.

Falk, E. B. and Bassett, D. S. (2017). Brain and Social Networks: Fundamental Building Blocks of Human Experience. *Trends in Cognitive Sciences*, 21(9):674–690.

Fearnhead, P. and Prangle, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):419–474.

Fentress, E. and Perkins, P. (1988). Counting African Red Slip Ware. In Mastino, A., editor, *L’Africa Romana: Atti del V Convegno di studio Sassari, 11-13 dicembre 1987*, pages 205–214. Sassari : Dipartimento di storia, Università degli studi di Sassari, Sassari.

Gallagher, E. M., Shennan, S. J., and Thomas, M. G. (2015). Transition to farming more likely for small, conservative groups with property rights, but increased productivity is not essential. *Proceedings of the National Academy of Sciences*, 112(46):14218–14223.

Gardent, J. (2017). *Mesurer les musiques pour parler du passé : la comparaison des musiques du Gabon comme source d’informations historiques*. PhD thesis, Muséum National d’Histoire Naturelle.

Gelman, A. and Hill, J. (2007). Data analysis using regression and hierarchical/multilevel models. *New York, NY: Cambridge*.

Gillespie, C. S. (2015). Fitting Heavy Tailed Distributions: The poweRlaw Package. *Journal of Statistical Software*, 64(2):1–16.

Gintis, H. (2006). The emergence of a price system from decentralized bilateral exchange. *Contributions in Theoretical Economics*, 6(1):1–15.

Gleeson, J. P., Cellai, D., Onnela, J.-P., Porter, M. A., and Reed-Tsochas, F. (2014). A simple generative model of collective online behavior. *Proceedings of the National Academy of Sciences*, 111(29):10411–10415.

Golder, S. A. and Macy, M. W. (2011). Diurnal and Seasonal Mood Vary with Work, Sleep, and Daylength Across Diverse Cultures. *Science*, 333(6051):1878–1881.

Gould, S. J. and Vrba, E. S. (1982). Exaptation–a missing term in the science of form. *Paleobiology*, 8(1):4–15.

Grimm, V. and Railsback, S. (2005). *Individual-based Modeling and Ecology*. Princeton Series in Theoretical and Computational Biology. Princeton University Press.

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., and Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, 363(6425):374–378.

Hahn, M. W. and Bentley, R. A. (2003). Drift as a mechanism for cultural change: an example from baby names. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(suppl_1):S120–S123.

Hanson, J. W. (2016). *An urban geography of the Roman world, 100 BC to AD 300*. Archaeopress, Oxford.

Hanson, J. W., Ortman, S. G., and Lobo, J. (2017). Urbanism and the division of labour in the Roman Empire. *Journal of The Royal Society Interface*, 14(136).

Harrison, R. A. (2018). *Experimental studies of behavioural flexibility and cultural transmission in chimpanzees and children*. PhD thesis, University of St Andrews.

Hayes, J. W. (1972). *Late Roman Pottery*. British School at Rome, London.

Hayes, J. W. (1985). Sigillate orientale. In Carratelli, G. P., editor, *Enciclopedia dell'arte antica. classica e orientale. Atlante della forme ceramiche II: ceramica fine romana nel bacino mediterraneo (tardo ellenismo e primo imperio)*, pages 1–96. Enciclopedia Italiana, Rome.

Hayes, J. W. (1997). *Handbook of Mediterranean Roman pottery*. British Museum Press, London.

Hayes, J. W. (2008). *Roman pottery, fine-ware imports. The Athenian Agora, results of excavations conducted by the American School of Classical Studies at Athens, volume XXXII*. The American School of Classical Studies at Athens, Princeton.

Henrich, J. (2015). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

Henrich, J. and Broesch, J. (2011). On the nature of cultural transmission networks: evidence from fijian villages for adaptive learning biases. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1567):1139–1148.

Henrich, J. and Gil-White, F. J. (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3):165–196.

Henrich, J. and McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3):123–135.

Heyes, C. M. (1994). Social learning in animals: categories and mechanisms. *Biological Reviews*, 69(2):207–231.

Heyes, C. M. (2012). What's social about social learning? *Journal of Comparative Psychology*, 126(2):193–202.

Hidalgo, C. (2015). *Why information grows: The evolution of order, from atoms to economies*. Basic Books.

Hyndman, R. J. (1996). Computing and Graphing Highest Density Regions. *The American Statistician*, 50(2):120–126.

Hyndman, R. J. (2018). *hdrcde: Highest Density Regions and Conditional Density Estimation*. R package version 3.3.

Jones, F. F. (1950). The pottery. In Goldman, H., editor, *Excavations at Gözlü Kule, Tarsus. Volume I text and plates. The Hellenistic and Roman periods*, pages 149–296. Princeton University Press, Princeton.

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

Kandler, A. and Shennan, S. (2013). A non-equilibrium neutral model for analysing cultural change. *Journal of Theoretical Biology*, 330:18–25.

Kandler, A., Wilder, B., and Fortunato, L. (2017). Inferring individual-level processes from population-level patterns in cultural evolution. *Royal Society Open Science*, 4(9).

Kass, R. E. and Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795.

Kempe, M. (2014). *Experimental and theoretical models of cultural evolution*. PhD thesis, Durham University.

Kendal, J., Giraldeau, L.-A., and Laland, K. N. (2009). The evolution of social learning rules: Payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology*, 260(2):210–219.

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., and Jones, P. L. (2018). Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences*. doi: 10.1016/j.tics.2018.04.003.

Kohler, T. A., Bocinsky, R. K., Cockburn, D., Crabtree, S. A., Varien, M. D., Kolm, K. E., Smith, S., Ortman, S. G., and Kobti, Z. (2012). Modelling prehispanic Pueblo societies in their ecosystems. *Ecological Modelling*, 241(0):30–41. Modeling Across Millennia: Interdisciplinary Paths to Ancient socio-ecological Systems.

Lachlan, R. and Slater, P. (2003). Song learning by chaffinches: how accurate, and from where? *Animal Behaviour*, 65(5):957–969.

Lake, M. W. (2014). Trends in Archaeological Simulation. *Journal of Archaeological Method and Theory*, 21(2):258–287.

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14.

Lazer, D., Kennedy, R., King, G., and Vespignani, A. (2014). The Parable of Google Flu: Traps in Big Data Analysis. *Science*, 343(6176):1203–1205.

Lieberman, E., Hauert, C., and Nowak, M. A. (2005). Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316.

Lipo, C. P. and Madsen, M. (2001). Neutrality,“style”, and drift: building methods for studying cultural transmission in the archaeological record. In Hurt, T. D. and Gordon F.M., R., editors, *Style and Function: Conceptual Issues in Evolutionary Archaeology*, pages 91–118. Bergin and Garvey.

Macmillan, W. and Huang, H. (2008). An agent-based simulation model of a primitive agricultural society. *Geoforum*, 39(2):643–658.

Marsaglia, G., Tsang, W. W., and Wang, J. (2003). Evaluating Kolmogorov’s Distribution. *Journal of Statistical Software, Articles*, 8(18):1–4.

Mesoudi, A. (2008). An experimental simulation of the “copy-successful-individuals” cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs. *Evolution and Human Behavior*, 29(5):350–363.

Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press.

Mesoudi, A. (2015). Cultural Evolution: A Review of Theory, Findings and Controversies. *Evolutionary Biology*, 43(4):481–497.

Mesoudi, A., Chang, L., Murray, K., and Lu, H. J. (2015). Higher frequency of social learning in China than in the West shows cultural variation in the dynamics of cultural evolution. *Proceedings of the Royal Society B: Biological Sciences*, 282(1798):20142209.

Mesoudi, A. and Lycett, S. J. (2009). Random copying, frequency-dependent copying and culture change. *Evolution and Human Behavior*, 30(1):41–48.

Mesoudi, A. and O’Brien, M. J. (2008). The Cultural Transmission of Great Basin Projectile-Point Technology I: An Experimental Simulation. *American Antiquity*, 73(1):3–28.

Meyer-Schlichtmann, C. (1988). *Die pergamenischen Sigillata aus der Stadtgrabung von Pergamon. Mitte 2. JH v. Chr. – Mitte 2. JH n. Chr.* de Gruyter, Berlin/New York.

Morer, I., Cardillo, A., Diaz-Guilera, A., Prignano, L., and Lozano, S. (2018). Comparing spatial networks: A’one size fits all’efficiency-driven approach. *arXiv preprint arXiv:1807.00565*.

Morer, I., Carrignon, S., and Rubio-Campillo, X. (2016). Influence of the topology of cultural networks on the equilibrium of an exchange-based economy. In *7th Workshop on Complex Networks (CompleNet 2016)* .

Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., and Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*, 279(1729):653–662.

Morin, O. (2015). *How Traditions Live and Die*. Foundations of Human Interaction. Oxford University Press, 1 edition.

Morris, I., Saller, R. P., and Scheidel, W. (2007). Introduction. In Scheidel, W., Morris, I., and Saller, R. P., editors, *The Cambridge economic history of the Greco-Roman world*, pages 1–12. Cambridge University Press, Cambridge.

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.

Nunes, M. A. and Balding, D. J. (2010). On Optimal Selection of Summary Statistics for Approximate Bayesian Computation. *Statistical Applications in Genetics and Molecular Biology*, 9.

Ormerod, P., Nyman, R., and Bentley, R. A. (2014). Nowcasting economic and social data: when and why search engine data fails, an illustration using Google Flu Trends. *arXiv e-prints*, page arXiv:1408.0699.

Ormerod, P., Tarbush, B., and Bentley, R. A. (2012). Social network markets: the influence of network structure when consumers face decisions over many similar choices. *arXiv e-prints*, page arXiv:1210.1646.

Pagel, M., Beaumont, M., Meade, A., Verkerk, A., and Calude, A. (2019). Dominant words rise to the top by positive frequency-dependent selection. *Proceedings of the National Academy of Sciences*, 116(15):7397–7402.

Poblome, J., Bounegru, O., Degryse, P., Viaene, W., Waelkens, M., and Erdemgil, S. (2001). The sigillata manufactories of Pergamon and Sagalassos. *Journal of Roman Archaeology*, 14:143–165.

Porčić, M. (2014). Exploring the Effects of Assemblage Accumulation on Diversity and Innovation Rate Estimates in Neutral, Conformist, and Anti-Conformist Models of Cultural Transmission. *Journal of Archaeological Method and Theory*, pages 1–22.

Premo, L. S. (2014). Cultural Transmission and Diversity in Time-Averaged Assemblages. *Current Anthropology*, 55(1):105–114.

Premo, L. S. and Scholnick, J. B. (2011). The Spatial Scale of Social Learning Affects Cultural Diversity. *American Antiquity*, 76(1):163–176.

R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Reali, F. and Griffiths, T. L. (2010). Words as alleles: connecting language evolution with Bayesian learners to models of genetic drift. *Proceedings of the Royal Society B: Biological Sciences*, 277(1680):429–436.

Reindl, E. M. (2017). *On the developmental origins of human material culture*. PhD thesis, University of Birmingham.

Remesal, J., Díaz-Guilera, A., Rondelli, B., Rubio-Campillo, X., Aguilera, A., Martín-Arroyo, D., Mosca, A., and Rull, G. (2014). The EPNet Project. Production and distribution of food during the Roman Empire: Economics and Political Dynamics. In *Information Technologies for Epigraphy and Cultural Heritage: Proceedings of the First EAGLE International Conference*, pages 455–464. Sapienza Universitá, Paris, France.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science*, 328(5975):208–213.

Robertson, T. E., Sznycer, D., Delton, A. W., Tooby, J., and Cosmides, L. (2018). The true trigger of shame: social devaluation is sufficient, wrongdoing is unnecessary. *Evolution and Human Behavior*, 39(5):566–573.

Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90(4):819–831.

Rubin, D. B. (1984). Bayesianly Justifiable and Relevant Frequency Calculations for the Applies Statistician. *The Annals of Statistics*, 12(4):1151–1172.

Rubio-Campillo, X. (2014). Pandora: A Versatile Agent-Based Modelling Platform for Social Simulation. In *Proceedings of SIMUL 2014, The Sixth International Conference on Advances in System Simulation*, pages 29–34. IARIA Publishing.

Rubio-Campillo, X. (2016). Model Selection in Historical Research Using Approximate Bayesian Computation. *PLoS ONE*, 11(1):1–18.

Ruck, D., Alexander Bentley, R., Acerbi, A., Garnett, P., and Hruschka, D. J. (2017). Role of neutral evolution in word turnover during centuries of english word popularity. *Advances in Complex Systems*, 20(06n07):1750012.

Ruck, D., Rice, N., Borycz, J., and Bentley, R. (2019). Internet research agency twitter activity predicted 2016 u.s. election polls. *First Monday*, 24(7).

Scheidel, W. (2012). Approaching the Roman economy. In Scheidel, W., editor, *The Cambridge companion to the Roman economy*, pages 1–21. Cambridge University Press, Cambridge.

Schillinger, K., Mesoudi, A., and Lycett, S. (2014). Copying error and the cultural evolution of “additive” vs.“reductive” material traditions: an experimental assessment. *American Antiquity*, 79(1):128–143.

Scott-Phillips, T. C., Laland, K. N., Shuker, D. M., Dickins, T. E., and West, S. A. (2014). The niche construction perspective: a critical appraisal. *Evolution*, 68(5):1231–1243.

Shennan, S. J. (2002). *Genes, Memes, and Human History: Darwinian Archaeology and Human Evolution*. Thames and Hudson, London.

Shennan, S. J. and Wilkinson, J. R. (2001). Ceramic Style Change and Neutral Evolution: A Case Study from Neolithic Europe. *American Antiquity*, 66(4):pp. 577–593.

Smith, D., Dyble, M., Major, K., Page, A. E., Chaudhary, N., Salali, G. D., Thompson, J., Vinicius, L., Migliano, A. B., and Mace, R. (2019). A friend in need is a friend indeed: Need-based sharing, rather than cooperative assortment, predicts experimental resource transfers among Agta hunter-gatherers. *Evolution and Human Behavior*, 40(1):82–89.

Smolla, M. (2017). *Environmental effects on social learning and its feedback on individual and group level interactions(2017)*. PhD thesis, University of Manchester.

Sobchuk, O. (2018). *Charting Artistic Evolution: An Essay in Theory*. PhD thesis, Universitas Tartuensis.

Solé, R., Valverde, S., Casals, M. R., Kauffman, S. A., Farmer, D., and Eldredge, N. (2013). The evolutionary ecology of technological innovations. *Complexity*, 18(4):15–27.

Steele, J., Glatz, C., and Kandler, A. (2010). Ceramic diversity, random copying, and tests for selectivity in ceramic production. *Journal of Archaeological Science*, 37(6):1348–1358.

Street, S. (2014). *Phylogenetic comparative investigations of sexual selection and cognitive evolution in primates*. PhD thesis, University of St Andrews.

Stubbersfield, J. M., Dean, L. G., Sheikh, S., Laland, K. N., and Cross, C. P. (2019). Social transmission favours the 'morally good' over the 'merely arousing'. *Palgrave Communications*, 5(1):70.

Tavaré S, Balding David J, Griffiths R C, and Donnelly P (1997). Inferring coalescence times from DNA sequence data. *Genetics*, 145(2):505–518.

Temin, P. (2006). The Economy of the Early Roman Empire. *The Journal of Economic Perspectives*, 20(1):pp. 133–151.

Temin, P. (2013). *The Roman Market Economy*. Princeton economic history of the Western world. Princeton University Press.

Tesfatsion, L. (2003). Agent-based computational economics: modeling economies as complex adaptive systems. *Information Sciences*, 149(4):262–268.

Tesfatsion, L. (2006). *Agent-based computational economics: A constructive approach to economic theory*, volume 2, pages 831–880. Elsevier.

Toni, T. and Stumpf, M. P. H. (2010). Simulation-based model selection for dynamical systems in systems and population biology. *Bioinformatics*, 26(1):104–110.

Toni, T., Welch, D., Strelkowa, N., Ipsen, A., and Stumpf, M. P. H. (2009). Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of The Royal Society Interface*, 6(31):187–202.

Turchin, P., Currie, T. E., Turner, E. A. L., and Gavrilets, S. (2013). War, space, and the evolution of Old World complex societies. *Proceedings of the National Academy of Sciences*, 110(41):16384–16389.

Valverde, S., Solé, R. V., Bedau, M. A., and Packard, N. (2007). Topology and evolution of technology innovation networks. *Phys. Rev. E*, 76:056118.

Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.

Waagé, F. O. (1948). Hellenistic and Roman tableware of north Syria. In *Antioch on-the-Orontes 4: part one, ceramics and Islamic coins*, pages 1–60. Princeton University Press, Princeton.

Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771.

Willet, R. (2012). *Red slipped complexity. The socio-cultural context of the concept and use of tableware in the Roman East (second century BC - seventh century AD). Unpublished PhD thesis*. PhD thesis, Katholieke Universiteit Leuven.

Williams, A. (2019). *Modelling the evolution of socio-political complexity*. PhD thesis, University of Exeter.

Wilson, A. (2009). Approaches to quantifying Roman trade. In Bowman, A. and Wilson, A., editors, *Quantifying the Roman economy. Methods and problems*, pages 213–249. Oxford University Press, Oxford.

Wilson, A. (2011). City Sizes and Urbanization in the Roman Empire. In Bowman, A. and Wilson, A., editors, *Settlement, Urbanization, and Population.*

Oxford Studies in the Roman Economy, pages 161–195. Oxford University Press, Oxford.

Wilson, A., Silver, M., Bang, P. F., Erdkamp, P., and Morley, N. (2012). A forum on trade. In Scheidel, W., editor, *The Cambridge companion to the Roman economy*, pages 287–317. Cambridge University Press, Cambridge.

Youngblood, M. and Lahti, D. (2018). A bibliometric analysis of the interdisciplinary field of cultural evolution. *Palgrave Communications*, 4(1):120.

Yu, H. (2002). Rmpi: Parallel Statistical Computing in R. *R News*, 2(2):10–14.

Zabehlicky-Scheffenegger, S. (1995). Subsidiary factories of Italian sigillata potters: the Ephesian evidence. In Koester, H., editor, *Ephesos. Metropolis of Asia. An interdisciplinary approach to its archaeology, religion and culture (Harvard Theological Studies 41)*, pages 217–228. Trinity Press International, Valley Forge.

Zahn, R. (1904). Thongeschirr (Scherben von Sigillata gefässen). In Wiegand, T. and Schrader, H., editors, *Priene: Ergebnisse der Ausgrabungen und Untersuchungen in den Jahren 1895-1898*, pages 430–434. Reimer, Berlin.

Ziman, J. (2003). *Technological innovation as an evolutionary process*. Cambridge University Press.