



UNIVERSITÉ DE STRASBOURG

ÉCOLE DOCTORALE AUGUSTIN COURNOT ED 221

BUREAU D'ÉCONOMIE THÉORIQUE ET APPLIQUÉE UMR 7522

THÈSE

pour l'obtention du titre de Docteur en Sciences Économiques

Présentée et soutenue le 15 Octobre 2024 par

Louis DALPRA

NAVIGATING COMPLEXITY: ESSAYS ON LEARNING AND DECISION MAKING IN THE NETWORKED ERA

Préparée sous la direction de Robin COWAN

Membres du jury :

Nicolas JONARD	Professeur, Université du Luxembourg	Président et Rapporteur
Alan KIRMAN	Directeur de Recherche, EHESS	Rapporteur
Annick VIGNES	Directrice de recherche INRAE, EHESS	Examineur externe
Moritz MÜLLER	Professeur Associé, BETA Strasbourg	Examineur interne
Robin COWAN	Professeur, Université de Strasbourg	Directeur

*L'Université de Strasbourg n'entend donner aucune approbation,
ni improbation aux opinions émises dans cette thèse ;
elles doivent être considérées comme propres à leur auteur.*

J'affirme être informé que le plagiat est une faute grave susceptible de mener à des sanctions administratives et disciplinaires pouvant aller jusqu'au renvoi de l'Université de Strasbourg et passible de poursuites devant les tribunaux de la République Française.

Je suis conscient que l'absence de citation claire et transparente d'une source empruntée à un tiers (texte, idée, raisonnement ou autre création) est constitutive de plagiat.

Au vu de ce qui précède, j'atteste sur l'honneur que le travail décrit dans mon manuscrit de thèse est un travail original et que je n'ai pas eu recours au plagiat ou à toute autre forme de fraude.

I affirm that I am aware that plagiarism is a serious misconduct that may lead to administrative and disciplinary sanctions up to dismissal from the University of Strasbourg and liable to prosecution in the courts of the French Republic.

I am aware that the absence of a clear and transparent citation of a source borrowed from a third party (text, idea, reasoning or other creation) is constitutive of plagiarism.

In view of the foregoing, I hereby certify that the work described in my thesis manuscript is original work and that I have not resorted to plagiarism or any other form of fraud.

Nom Prénom : Dalpra Louis

Ecole doctorale : Augustin Cournot ED 221

Laboratoire : BETA, UMR 7522

Acknowledgments

First, I would like to express my deepest gratitude to my thesis supervisor, Robin Cowan. His course on network economics was a pivotal moment in my academic journey, sparking my passion for research. His guidance throughout these PhD years has been invaluable. He consistently offered direction while giving me the freedom to explore my ideas and grow as an independent researcher. Working under his supervision has been an enriching experience, both academically and personally.

I would also like to thank the members of my jury: Nicolas Jonard and Alan Kirman, my two referees, as well as Annick Vignes and Moritz Müller, for agreeing to be part of this process. I am grateful for their time in reading my work and honored by their presence at my defense.

Special thanks to Moritz and Nicolas for being part of my thesis committee. Our meetings were always productive and insightful.

Although not on this jury, my deepest thanks go to Patrick Llerena. From the moment I decided to pursue a PhD, you have been instrumental in making it happen—first by hiring me as a research assistant, then by consistently offering your help and guidance throughout the process.

I would also like to thank the administrative teams at BETA and the doctoral school for their support, as well as the professors I've had the privilege to work with over the years.

A heartfelt thanks to my fellow PhD candidates, some of whom have become dear friends. Romane and Théo, for the many nights spent away from the thesis, trying every pizza and nightclub in Strasbourg, and for the visit of the hospital services..

At Bureau 133, thanks to Kenza and Pierre for being such great desk neighbors; your energy brings life and joy to the office. Also, thank you to Jibs, Laura, and Sofiene for your camaraderie. And to Jérôme, Agathe, Guillaume, and Anne-Gaëlle for all the good times we've shared.

Outside of BETA, thank you to my lifelong friends Thiebaud, Théo, and Victoria, as well as the friends I've made at the faculty: Alexandre, Biver, Malek, and Laurent. Thanks to Ghita and Boubker for their support during the early stages of this thesis.

Finally, this thesis would not have been possible without the constant love and care of my family. Thanks to my uncle and aunt for the Tuesday evening dinners that made your home feel like a second one to me. To my cousins Emile and Arthur for all the stimulating conversations and FIFA games, and to Rose for keeping me young by watching Les Marseillais.

Thank you to my grandparents for looking after me when I first arrived in Strasbourg and for always believing in me. To my grandmother Denise, who will not be here to read these lines but who has provided me with all the love I could ask. I'll always cherish the times spend with

you.

Thank you to my sister MS for your constant cheer and support. And finally, my deepest gratitude to my parents. None of this would have been possible without you, and I owe everything to the love and opportunities you've given me. This achievement is as much yours as it is mine.

Contents

0.1	Complex Economics	21
0.2	Social Networks to Understand Social Heterogeneity	23
0.3	Learning in Social Networks	26
0.4	Internet and the Networked Era	31
0.5	Outline of the Thesis	32
0.6	Chapter 2	33
0.7	Chapter 3	34
0.8	Chapter 4	35
0.9	Chapter 5	36
0.10	Methodology	37
0.11	Économie de la Complexité	41
0.12	Les Réseaux Sociaux pour Comprendre l'Hétérogénéité Sociale	43
0.13	Apprendre au Sein des Réseaux Sociaux	47
0.14	Internet et l'Ère des Réseaux	52
0.15	Plan de la thèse	54
0.16	Chapitre 2	54
0.17	Chapitre 3	55
0.18	Chapitre 4	56
0.19	Chapitre 5	57
0.20	Méthodologie	59
1	Experience goods, Reinforcement Learning & social networks	63
1.1	Introduction	63
1.2	The Model	67
1.2.1	Agents & Products	67
1.2.2	Interaction	68
1.2.3	Learning	68
1.2.4	Dynamics & Learning methods	70

1.3	Results	71
1.3.1	The role of Interaction	71
1.3.2	Social and Individual Learners	73
1.3.3	Curiosity	75
1.3.4	Trust-based social interactions	77
1.4	Conclusion	82
2	Internet and experience goods: The role of Influencers and rating-scale systems	87
2.1	Introduction	87
2.2	Model	89
2.2.1	Reinforcement Learning	89
2.2.2	Network Models	91
2.3	Results	92
2.3.1	Influencers & Star Networks	93
2.3.2	Influencers & Realistic Network Structures	95
2.3.3	Influencers and Heterogeneous Advertising	98
2.3.4	Review system	101
2.3.5	Review system with heterogeneous tastes	105
2.4	Conclusion	108
3	Opinions Polarization, Influencers, and Endogenous Online Social Networks	113
3.1	Model	117
3.1.1	Casual users	118
3.1.2	Influencers	119
3.1.3	Influencers-Casual Users Interaction	120
3.1.4	Discovering Influencers and the Platform's Role	121
3.1.5	Model flow	122
3.2	Results	123
3.2.1	Opinion dynamics	123
3.2.2	Distribution of Influencer's degree	130
3.3	Conclusion	135
4	Online Social Network protocols	139
4.1	Introduction	139
4.2	Modeling Networks	144
4.2.1	Creating a "Real-life" network	144

4.2.2	Cultural Diversity	146
4.2.3	Recommending Influencers	147
4.2.4	The Online Social Networks	148
4.3	Diffusion of Information	151
4.3.1	Model	151
4.3.2	Results	152
4.4	Collective Action	155
4.4.1	Model	156
4.4.2	Results	156
4.5	Job Market	159
4.5.1	Model	160
4.5.2	Results on Simple Networks	162
4.5.3	Results on Online Social Networks	164
4.6	Conclusion	165

*“All that is gold does not glitter,
Not all those who wander are lost;”*

– J.R.R. Tolkien

General Introduction

0.1 Complex Economics

Whenever I introduce first-year undergraduates to classical consumer theory in microeconomics, I stumble upon generalized perplexity in the classroom. Questions arise: Why are consumers depicted as always wanting more goods? In the work-leisure model, why is there an assumption that agents freely choose their optimal amount of work? And why is work evaluated solely through its monetary benefits, ignoring jobs that offer intrinsic value beyond financial gain? Of course there are reasons for these models — whether or not they’re satisfactory—rooted in their objectives or simplifications for analytical tractability.

But interestingly, by the time students reach their final undergraduate year, presenting them with models where firms opt out of production or models incorporating knowledge of a competitor’s reaction function in a duopoly scarcely prompts questions. Does this indicate a deeper understanding of economics, or a resignation to the abstraction from reality inherent in these models?

A peculiar shift seems to occur during economics education. As highlighted by Etzioni (2015), economics students develop a ”debased” moral compass. What factors contribute to this transformation?

A growing segment of the field argues that the root issue lies in the foundational assumptions of economics (Kirman (1989, 1992); Arthur (2009)). Since Herbert Simon’s critique in 1955 of the rational, utility-maximizing agent model (Simon (1955)), a significant debate has arisen over the realism of these assumptions. Simon pointed out the fallacy of assuming perfect rationality, a critique that resonates through to macroeconomics where the notion of equilibrium can obscure significant issues, as evidenced by the 2008 financial crisis shortly after Olivier Blanchard’s optimistic assessment about the state macroeconomic theory (Blanchard (2009)).

This acknowledgment of the limitations inherent in traditional economic models has spurred interest in alternative approaches. These approaches reconsider the behavior of economic agents, acknowledging the bounded rationality that characterizes human decision-making, as Simon

suggested, and further explored in the realm of behavioral economics (Gigerenzer and Selten (2002)). Agents, influenced by cognitive biases and without the computational prowess assumed in classical models, often opt for "satisficing" strategies over optimizing ones (Kahneman (2003)).

Additionally, economic behavior is not uniform but varies according to cultural, socioeconomic, and other factors, underscoring the heterogeneity of agents. Moreover, agents learn and adapt based on their experiences, with their preferences and behaviors evolving over time. This dynamism contradicts the static nature of utility functions or demand curves in classical models.

Interactions among agents further complicate the picture. An individual's behavior is often influenced by their social network, leading to changes in preferences and behaviors that classical models, with assumptions about independent consumers, fail to capture. This social dimension can lead to behaviors like consumption shifts based on peer influence or differentiation efforts. An agent might pick a good not only for its intrinsic value, but sometimes because it became desirable seen consumed by others, or sometimes because its consumption will permit to differentiate from others. Consumption is a social act that extends our identity.

The implication is that macroeconomic phenomena cannot be fully explained by aggregating the behaviors of isolated agents; instead, the interactions and adaptations among agents play a crucial role. Schelling's segregation model vividly illustrates this point (Schelling (1969)).

The author envisions a population living on a grid, divided into two classes of agents. Each agent observes the class of its neighbors and has a certain desire to be surrounded by agents of the same class. If the number of similar neighbors is too low, the agent relocates randomly to an empty space on the grid. The model and its result became famous because Schelling shows that even with a population that has a weak desire to be surrounded by peers — only moving if more than two thirds of his neighbors are of a different type — the global outcome ends-up being a segregated grid, with large neighborhoods containing only one class of agents.

This does not imply that the ghettos or highly segregated neighborhoods witnessed throughout history and today can be completely explained by this model and the individual drive of local agents, as numerous variables are at play, including significant political actions. Nonetheless, it shows that even with agents being fairly open to diversity, segregation can arise.

The core idea of complex economics, illuminated by Schelling's segregation model and Simon's adage that "the whole is different from the sum of its parts," emphasizes that the behavior of individual agents, when combined through interactions, leads to collective outcomes that may not be intuitive based on individual actions alone. This principle reveals that the aggregate dynamics of a system, such as the segregation levels within a city, cannot be accurately inferred just from analyzing the preferences of its residents in isolation. The unexpected patterns of

segregation arise from the nuanced interplay of individual choices, highlighting the need for economic modeling to incorporate the complexities of human behavior and social interaction more deeply. It underscores that understanding the macroscopic phenomena of economies or societies requires a careful consideration of the micro-level interactions that drive these systems, challenging us to rethink traditional approaches to economic theory.

0.2 Social Networks to Understand Social Heterogeneity

As stated earlier, to achieve the enriched understanding of economic agents advocated by proponents of the complex economics school of thought, we must acknowledge that while individual behavior — such as heterogeneous decision-making relying on heuristics, for example — is one piece of the puzzle, there is at least one equally important piece concerning social interactions.

In what follows, we define social interactions among economic agents as all the processes that cause an agent to behave differently due to its interactions with others. This broad definition can account for many different situations and intentions, whether the influence is exerted by individuals or groups, or received by them. It is important to acknowledge that the social influence that results from interactions can either be conscious or unconscious.

At the core of the idea of social influence is the tendency of humans to imitate and integrate what others do. This can be driven by the desire to fit into a group and belong, to differentiate by doing the opposite, to demonstrate contrasting values towards a group, or simply to achieve success. It can seem rational to imitate successful agents in order to become successful.

Social influences are evident in the consumption of many types of goods. We refer to "fashion" as the ongoing process of trends, which can be observed in clothing, furniture, architecture, cultural goods, and which essentially represents a process of social validation or the highlighting of certain features and characteristics that are currently deemed desirable by the mass market.

Most of us, whether intentionally or not, participate in this type of consumption. We may strive to be trendsetters, deliberate followers of the trend, or even unconsciously choose goods that are in vogue.

Furthermore, we may seek to signal our belonging to a certain social class. We purchase high-end goods to showcase our financial success, or at least to demonstrate our aspiration to join that particular social group (Veblen (2017); Leibenstein (1950); Bourdieu (1979); Johnson et al. (2018)).

While this type of social influence may seem uncoordinated, albeit partially driven by brands, in other domains, social influence is much more explicit. In his famous book, "The Protestant Ethic and the Spirit of Capitalism," Max Weber directly links Protestantism, a branch of Christianity that embodies a set of values, to the work behavior of its believers, thereby contributing

to the emergence of capitalist Western societies in Europe in the 19th century (Weber and Kalberg (2013)).

Such institutions, which are partially defined by written guidance and rules, as well as hierarchies, also cultivate close-knit communities among their members. These communities can serve as either pressure for individuals who may not wish to comply with all the rules, or as aspiration for those who strive to adhere more closely (Carvalho (2019)).

Furthermore, recent advances in science reveal that some behaviors are influenced by acquaintances, although we may not fully understand the mechanisms through which this occurs. Health behaviors such as obesity or smoking, for instance, have been shown to positively correlate with similar behaviors among peers (Christakis and Fowler (2007, 2008); Centola (2010)).

Therefore, to think that we can analyze the behavior of economic actors such as workers, consumers, or entrepreneurs without considering the social milieu they inhabit and the social interactions they engage in is to only see half of the picture.

Modeling or quantifying social interactions is easier said than done, which likely explains why the field of economics was slow to deeply engage with it until recently. This change occurred as network science began to gain greater recognition and visibility within the scientific community.

Networks formalize and allow for the visualization of the connections that bind various nodes together. In the context of social interactions, nodes can represent individuals or groups, while ties can represent various types of connections between two nodes, from family ties to formal work relationships, for example.

Constructing a network is therefore the explicit enumeration of all the agents we wish to incorporate, along with all the existing connections among them, based on the significance we attribute to those links.

The existence and study of networks—also known as graphs in the mathematical field from which they originated—can be traced back at least to the 18th century and Euler’s famous resolution of the “Seven Bridges of Königsberg” problem (Euler (1953)).

But the emergence of its existence as a field in itself - ‘Network Science’ - is much more contemporary, and is explained by at least two main factors according to Barabasi:

1. To establish the complex maps of the networks we want to study, a lot of data is needed. A social network requires accurate information about your friends, your friends’ friends, and so on. If we want to build the map of the internet, we need to know every web page and every link on web pages that redirects to others. It’s precisely the internet revolution that has facilitated the creation of these maps, with effective and fast data-sharing methods and cheap digital storage. Researchers are now able to collect, assemble, and analyze real and complex networks.

2. While networks were partially studied in different fields such as biology, chemistry, or social sciences, it's the universality of network properties that made the emergence of the field possible. As Barabasi writes (Barabási (2016)): "The architectures of networks emerging in various domains of science, nature, and technology are similar to each other, a consequence of being governed by the same organizing principles. Consequently, we can use a common set of mathematical tools to explore these systems."

On the eve of the 21st century, two important models began to explain some of the universal characteristics of network structures. The first one, proposed by Barabasi and Albert, identifies the 'scale-free' nature of many different networks such as the Internet, power grids, or social interactions. Here, the characteristic highlighted is the distribution of connections for the nodes of the graph. Specifically, while the majority of nodes will have a low and similar number of connections, a small group will have a significantly higher number of connections Barabási and Albert (1999).

The second model proposes to demonstrate that most of our interaction schemes occur in 'Small-Worlds', where even in systems with many nodes, reaching any of these nodes can be accomplished through small 'leaps' on the network, while the structure of the network remains tightly knit, implying clustering among small groups of nodes (Watts and Strogatz (1998)).

This model explains phenomena such as Milgram's famous experiment, where he asked individuals to reach others (who were unknown and geographically far) (Milgram (1967)). The results showed that any agent can reach any other in about six steps, coining the idea of the 'Small-World' in which we live.

Beyond demonstrating the shared structural properties of a wide diversity of networks, from technological to biological and social ones, these models also shed light on the measurement tools that any scientist can use to analyze a network.

Scale-Free structures highlight the importance of the degree distribution at the global scale, but also of degree centrality at the local level. A network is therefore useful for understanding the dynamics of global phenomena, such as the speed of diffusion, for example, but also for understanding the role of specific actors in these structures. Does the most central agent in a network have the same opportunities or knowledge as the least central, for example?

The same can be said about the Small-World model. It shows that we can measure graphs at the global level using concepts such as path length, whether we want to measure its average or shortest, and with the concept of embeddedness, which can measure the existence and strength of communities in networks. But it also applies at the individual level: Small-Worlds emphasize the idea of betweenness centrality and the importance of nodes that can act as bridges among different and remote communities.

While models emerging from mathematics and physics have allowed Network Science to emerge as a distinct field, some of the intuitions of network science were already present and the subject of works in other disciplines.

The notion of social capital — which could be described as the ability to mobilize resources from social connections — has existed in sociology since the 19th century and has been progressively used and developed throughout the 20th century.

Many authors have utilized this concept, such as Bourdieu (2021), who analyzes its effects and interdependence with other forms of capital; Putnam (2015), who shows its importance for altruism and democratic institutions; and Coleman (1998), who hypothesizes that social capital acts as a vector of conformity within societies.

Although none of these authors use networks as a tool to study social capital directly, the very idea of social networks is already present and useful in their analysis of society and its individuals.

More recently, network objects have become central in the works of Granovetter (1973) in the seventies and Burt (2018) in the nineties. Both authors highlight the importance of the structure of the graph at both global and individual levels, whether it is to generate 'good ideas' or to find jobs.

With the tools formalized in mathematics and physics, and the concept increasingly used in social science, Networks naturally became a tool and an object of study in Economics in the early 2000s with the development of models studying its impact on the job market matching performance, how it can be used to study production networks or how it helps to study the diffusion of knowledge, technology or behaviour (Jackson (2010)).

The present thesis will also utilize networks, but to underscore their significance in the learning process of individuals, within frameworks where interactions and the structure of those interactions have important implications at both individual and global levels.

0.3 Learning in Social Networks

Networks became particularly important in the study of diffusion. Whether we talk about the diffusion of knowledge, information, viruses, or electricity, using networks to describe the process is convenient as it generates a map describing all the roads a given object can use to reach the existing nodes, hence the points of interest of this map.

Having a map is a departure from the unknown. It allows us to understand, and sometimes predict, the scope or speed of a diffusion process because it explicitly shows the existing routes or the non-reachable nodes of the system. It can also serve as a method to identify the important

nodes, the ones that are essential for diffusion, and the ones that can be disconnected without impeding the diffusion process.

In turn, a diffusion process takes place only if the network on which it happens is non-empty (at least two nodes are connected), and being able to represent this network is crucial to understanding the dynamics of diffusion.

As mentioned above, we can study the diffusion of "real" objects such as people moving from one subway station to another, viruses from one body to another, or Wi-Fi signals from operators to houses and computers. Still, we can also study the diffusion of beliefs, emotions, information, and knowledge among the individuals of a society. As these factors are at the heart of many of our decisions and behaviors, they are likely to explain (some aspects of) economic processes.

The case of diffusion among individuals can be categorized as social learning, wherein the way agents send, accept, and internalize information or ideas responds to different rules depending on the nature of what is exchanged, on the identity of the people or group that sends it, but also on the nature of the individual recipient and on the culture of the society from which the recipient originates.

We don't react the same way to information about the evolution of a company's stock value as we do to the latest presidential poll. Similarly, we value information differently depending on who gave it to us. Was it an informed friend, or a partisan media source? Lastly, our personal background shapes how we deal with the information received. Will it lead us to invest in the company or to demonstrate against the rise of an extreme political party?

Social learning is therefore a complex mechanism, and its outcome, its efficiency, can depend on many cognitive and emotional variables such as trust, belief, or fear - to name just a few - depending on the scenario being studied. However complex and difficult to understand exactly how it works, social learning can be described as a diffusion process, on a network where nodes are individuals, and ties represent possible routes for pieces of information, knowledge, but also beliefs to be exchanged.

Foundational models that aim to represent social learning, or at least observational learning, are mostly based on two mechanisms: using weighted averages or using Bayes' rule (See Goyal (2011) for an extensive review).

In the first scenario, originating from DeGroot (1974), agents placed on a network structure can observe the behavior or beliefs of their direct neighbors (the ones with whom they share a tie), and adapt their own behavior or beliefs accordingly. In this framework, an agent gives fixed weights to its neighbors, expressing the value it gives to each neighbor. Then, its behavior over time is the weighted average of what its neighbors are believing/choosing, which are themselves following the same process with their own neighbors.

This type of observational learning can be useful to represent opinion dynamics where agents' beliefs evolve as they consider the opinions of others, whether it is through communication or observation. Several questions can be asked by studying this kind of model. Do we observe convergence of opinions? Under what conditions on the initial distribution of beliefs? What is the role of the network structure?

Although it is a useful model, it also has several limitations. For example, it assumes that at each round, an agent collects information about all of its neighbors' opinions. This kind of ubiquitous knowledge is contestable in itself.

The second major limitation arises from the updating mechanism. The scenario in which agents compile opinions and adopt the weighted average from their circle of influence is questionable. This approach may not be suitable or applicable in various contexts.

For example, we study opinion dynamics partially to understand democratic elections. Suppose that an agent has two opposed neighbors on the political spectrum. This agent gives equal weights to those two friends, making itself a centrist voter. Now, when the election arrives and no centrist candidate is on the line, to whom will the agent's vote go? What it chooses will reveal the actual weights it puts on its two friends because the outcome is binary; it won't be the weighted average of those two opinions. If the goal was to understand election outcomes, the model was ill-suited because the weighted average does not fit the final decisions agents would have to make during the election (although they can still have some predictive power).

Of course, while there are many cases where the model is limited, there are just as many scenarios where this kind of modeling can be useful, since not all processes end with a binary decision.

The second scenario, where agents use Bayes' rule, like in Acemoglu et al. (2011), mostly uses a model where agents are also placed on a network structure that limits their interactions to their neighborhoods and allows them to observe the actions and payoffs only of those neighbors. This then allows them to update their prior beliefs on the outcomes of a given event. In these kinds of models, agents are trying to find the "truth" or the maximizing behavior, which can be viewed as an action to undertake or not, but also an opinion to have or not.

In both streams of literature, and in most of the economic literature that takes an interest in social learning, there is mostly a common vision that agents learn from others by observing - and not through one-to-one communication, for example - the entirety of a set of neighbors at each period. Moreover, those agents use complex methods to determine their actions, whether it is computing a weighted average of their friends' decisions or updating a probability distribution based on the outcomes of those friends. Of course, we can choose to believe that those models work not because we're behaving this way - doing complex calculations - but because whatever we do, this can be approximated by those calculations.

Other papers have tried to use a perhaps more realistic approach by endowing agents in their models with "rules of thumb," which are simple heuristics. This is the case in Ellison and Fudenberg in the context of adopting different technologies (Ellison and Fudenberg (1995, 1993)). Their rules of thumb also involve a population observing the choices and payoffs of all neighbors, but now only a fraction of them revise their decisions at each round, and instead of complex calculations, they use simple rules such as: Use what was the most rewarding in the last round, or use what was the most popular in the last round.

In the same spirit, W.B. Arthur builds an interesting model, the "El-Farol Bar" problem, where agents want to avoid going to a bar if it's too crowded. Each time they go, they can observe if it's indeed too crowded or not and adapt their behavior accordingly for the next time they must decide to go or not (Arthur (1994)).

While not using network tools, each agent of this model is learning by observing the behavior of its geographical neighbors - the ones who also go to this particular bar. Arthur proposes that while agents are using rules of thumb to decide what to do, they are heterogeneous in which rule of thumb to use, therefore creating an "ecology" of rules.

Outside of economics, some authors have developed models of social learning. Notably in sociology, Granovetter (1978) develops a model of collective action where agents must decide whether they join or not a movement, depending on the decision of their neighbors that they can observe. It highlights an interesting rule of thumb where agents, instead of doing as the majority or doing the average of neighbors, decide to engage only if a precise proportion of their neighborhood has already joined the collective action.

In his paper, Granovetter use the example of a riot. Say that every agent in a crowd must decide whether to engage or not in a riot. Its a risky behavior that can have repercussions, thus, most of the people that would be ready to engage in it also need to observe that they're not alone in the movement. Depending on how strongly they believe in the cause for which they're ready to riot, they need to observe a different proportion of the population already engaged in it. Thus the riot is started by the most convinced agents, which spurs the involvement of the people that needed to witness those few rioters to engage themselves, and so on until every agent that had even the slightest willingness to riot — the ones that need to observe a very large proportion of the population participating — engage in the riot. There is therefore a cascading dynamics from the strongest believers to the weakest ones. The model is quite different than ones where agents engage if they observe a majority of agents already engaged. With the same crowd and the same distribution of willingness to engage in the riot, only the most convinced will participate, because every other agent in the crowd will observe a minority of rioters (its important to note that this configuration makes the assumption that every agent observe the behavior of every other one in the crowd, which is not so common).

Overall, understanding and accurately depicting social learning is complex and highly context-dependent. There is not one canonical model that would suit all kinds of information exchanges among individuals.

Our ambition with this thesis is to deepen some of the existing models by adding variables that can make them more suitable for specific contexts, but also to propose our own mechanisms of interactions that can lead to social learning.

For example, most of the existing models about social learning cast away the possibility for individuals not to use the actions of their peers. But in reality, agents have different reactions given their observations of their neighbors' actions. Even further, they might be learning about what type of behavior rewards them the most in the long run. And if they choose to be influenced by their neighbors, perhaps they are also learning about the best neighbors to listen to. These kinds of issues will be at the heart of our first chapter.

Another important idea that we use throughout the thesis is endogenously built networks. As they represent interactions among agents, a link between two nodes is likely to represent more than just a channel of communication. Instead, the existence of those links can be constructed using the processes that we know are at the basis of social interactions. Homophily, for example, can be used to represent agents connecting out of similarity; triadic closure for friends of friends becoming friends themselves; or preferential attachment when individuals choose to connect to popular agents on the network.

These concepts can be used to allow agents to form and break connections, leading to realistic network structures such as Scale-free or Small-World graphs, which will be used as the architecture of interactions of our populations. As we will show, interactions can lead to very different outcomes depending on which structure individuals are placed in.

Altogether, our models of social learning encapsulate many human processes — although they are not always at the center of the model, and thus discussed heavily — depending on the context. These include trust, attraction, repulsion, imitation, as well as honesty and dishonesty. Additionally, interactions may be bound by memory or driven by motives that are either interested or disinterested.

In addition to being centered around the notion of social learning, this thesis aims to explore the changes that the Internet and social media have brought with them. Most of the models developed here are set in the context of social media, encompassing the changes it brings to the way people communicate, but also to the new actors it creates, from Influencers to Online Social Network (OSN) platforms.

0.4 Internet and the Networked Era

A recent report from Pew Research indicates that more than 40% of TikTok users are using the platform to get news (Matsa (2023)). With around 150 millions American monthly active users (Statista (2023)), its thus safe to say that a significant part of the American society is partially informed by an online social network which algorithms are unknown and determined by its Chinese company.

As online social media becomes ubiquitous in our lives, it is important that our models adapt to this new paradigm. Specifically, social interactions are profoundly changed by OSNs. Firstly, they modify the number of interactions one can have. They allow us to send messages to any neighbor of that network, and even to communicate with a whole neighborhood at once by sharing posts. This illustrates that not only are online platforms modifying the scope of our interactions, but also their nature. Posting a status on Facebook or Twitter, visible to everyone connected to you, has no equivalent in the pre-Internet era, except perhaps at very specific events such as weddings or funerals, during which the content of the speech is very controlled. This radical departure from traditional communication necessarily impacts our understanding of global opinions or trends because we were never confronted with such forms of communication before.

Additionally, OSNs are not neutral places; they are firms that compete among themselves for your attention, for the time that you will spend using them. This leads them to develop algorithms that can distort your vision of the world.

Online social platforms also become the playground of Influencers, opinion leaders who seek revenues from their audience.

This new landscape is particularly suited to be studied through the lens of Network Science because OSNs are essentially network objects. What the platforms propose to their users is simply to create a profile (a node) online and to be able to connect to other individuals present on the OSN (thus creating a tie between two nodes). The value of a particular OSN then stems from two primary elements: the people that are already using it - how many people, but also what kind of people - and the rules, or policies (that we define later in the thesis as the protocols), that the OSN undertakes to provide the best experience possible, the one that will keep its users on its website for longer. The most common example is the algorithms that dictate what publications appear on the user's feed.

The first source of competitive advantage is commonly called Network Effects in the literature and has been the subject of research by many authors since the mid-eighties (Katz and Shapiro (1985); Arthur (1989); Cowan (1991)). This type of effect is not restricted to OSNs and can be found in many industries, particularly in markets where the value of an object or

a technology relies on its existing user base. This can occur because the value stems from the fact that others are also using it - like the telephone - or from people's reluctance to incur transition costs. For example, you might prefer to use Excel, Python, or a QWERTY keyboard not because they're the best options on the market, but mostly because they have become standard in society, and if you were to move from one job to another, there's a high probability that the new company will also ask you to work with those products (David (1985)). This second case mostly implies goods where there's a learning phase - such as learning to code in Python, learning to use shortcuts and functions, or learning your muscle memory to know where the keys are on the keyboard.

By contrast, the second explanation for the success of an OSN, or at least the levers it can actually use to differentiate itself from its competitors, is less studied in economics, mainly because it is a relatively new method of differentiation that emerged with OSNs in the late 2000s. We will show that most of the choices an OSN makes to define the rules of its platform can be translated into the rules that generate the network object itself. Therefore, we can study many policies the firms are taking by showing how they impact the network structure of the OSN, and then studying the impacts of the network structure on processes such as information diffusion or coordination to engage in a collective action. This, in turn, helps understand which OSNs could be more attractive because, from an objective viewpoint, they are more efficient for certain tasks. This work will be the subject of the last chapter of this thesis.

Lastly, it's important to recognize that some OSNs have specialized in specific purposes. LinkedIn is used to share about work-life, while Instagram focuses on pictures. These specificities give rise to different behaviors from individuals depending on which OSN they are using. Our models will aim to adapt accordingly.

As social learning adapts to new frameworks, with people communicating and observing others online as much - if not more - than offline, we aim at adapting accordingly, by proposing a nuanced and detailed vision of online social learning, mostly by allowing social networks to form endogenously under the constraints of platforms, and by acknowledging the existence of powerful forces such as recommendation algorithms or influencers.

0.5 Outline of the Thesis

Chapter 2 begins the thesis by asking whether social learning is the best behavior available in a context of uncertainty where agents are consuming experience goods. Chapter 3 adapts the model to online problems such as the existence of advertising from influencers or collaborative rating scale systems. Chapter 4 attempts to capture the role of influencers and algorithms in

OSNs for the dynamics of opinions, with a focus on polarization. Finally, Chapter 5 demonstrates how network protocols modify network structures and, consequently, social learning dynamics on those structures.

The subsequent sections will provide a comprehensive overview of each chapter and the methodologies employed throughout this thesis.

0.6 Chapter 2

This chapter explores the challenges of modeling consumer behavior, particularly when choosing among unfamiliar options through repeated trials. We argue that individuals employ diverse learning methods, and a single "optimal" strategy may not exist.

Focusing on experiential goods, where consumption itself provides learning, we propose a multi-dimensional reinforcement learning model. This model captures how agents not only learn to choose effective products but also to refine their decision-making processes over time.

Thus agents have the choice between learning about the best product through individual reinforcement learning, or through social learning, by using their social networks to get recommendations, and choosing accordingly.

We therefore propose a different vision of social learning. Here the idea is similar to a simple rule of thumb: agents increase the probability to choose again high-rewarding options more than low-rewarding ones. But to choose among the alternatives, we allow individuals to choose either between individual or social learning. The influence of peers is thus not an obligation for our agents, who will decide to use recommendations from their neighborhoods only if they come to believe that it is a more efficient way to get the best goods. Moreover, our agents are not observing the entirety of their neighbors decisions or outcomes, instead they choose only one.

Our findings reveal that a model incorporating both individual and social learning significantly outperforms a model limited to individual learning. This highlights the positive impact of social learning, even when individuals lack prior knowledge about the available options.

The benefit of social learning appears to stem from recommendations carrying implicit information about preferred options. However, individual learning also holds advantages, allowing some agents to discover good choices quickly. Nonetheless, individual learning comes with higher risk, as individuals may get stuck on sub-optimal options.

Therefore, while social learning may not lead to perfect choices for all agents, it ensures consistent performance at a high level. We further explore the role of optimal curiosity levels, which allow agents to sufficiently explore before committing to a choice.

Finally, the model incorporates a trust element, where agents learn to identify reliable

sources of advice within their social network. This element enhances overall performance and leads to interesting network structures, with certain agents becoming central figures for advice.

0.7 Chapter 3

This chapter investigates the impact of the internet on decision-making, particularly when choosing experience goods – products whose value is uncertain until after purchase. We focus on how online interactions, including online review systems and influencers, influence individual choices and outcomes in such situations. Here we use the same model as in chapter 2, with individuals having the opportunity of learning about choosing the best goods either through social or individual learning. We extend it by adding the presence of influencers, by modifying the network structure, and by comparing it to a model with aggregated online reviews.

The research explores the potential advantages of social learning, where individuals learn from the experiences of others, compared to individual learning through trial and error. The internet facilitates social learning by fostering online communities and platforms like review systems. However, the effectiveness of these systems can be limited by diverse user preferences. Aggregated reviews may not accurately reflect individual needs, highlighting the importance of considering the structure of online interactions and the potential limitations of relying solely on aggregated information.

Furthermore, this chapter delves into the role of online social networks and online influencers. These networks can connect individuals with similar tastes, potentially leading to more informed decisions. However, the action of online influencers can be detrimental, particularly for average-quality goods. Individuals may be swayed by endorsements without fully exploring alternative options, potentially missing out on superior choices. This effect is most pronounced in networks with structures exhibiting preferential attachment, whether in extreme cases (star networks) or more realistic setting such as scale-free networks (common in online platforms). We also found that some structure are much less impacted by influencers. This is the case for Small-World networks, structures exhibiting low average path length and high cliquishness.

Overall, this chapter suggests that while the internet offers promising tools for improving decision-making, critical evaluation and a cautious approach towards online influence are crucial for navigating the complexities of choosing experience goods in the digital age. Ultimately, social interactions with trusted individuals might be the most effective strategy for making informed decisions in such contexts. However, the effectiveness of these interactions, and the Internet’s influence as a whole, depends on factors like user preference distribution and the ethical conduct of online influencers, which may not always be ideal in real-world scenarios. This research contributes to a deeper understanding of these challenges and the potential benefits

and drawbacks of online interactions in the realm of decision-making for experience goods.

0.8 Chapter 4

This chapter examines the complex interactions between different user groups on online social networks (OSNs), particularly focusing on the relationship between aspiring influencers and the general user population. We reveal a dynamic system where both groups influence each other's behavior.

Regular users primarily share and discuss their opinions, potentially leading to increased polarization or greater consensus on specific topics. Aspiring influencers, on the other hand, strategically share content to maximize their online success. Interestingly, their presence seems to counteract the formation of extreme viewpoints, hindering both complete consensus and complete polarization.

We move beyond the traditional understanding of success in online networks by identifying key factors that drive success for aspiring influencers. The most significant factor appears to be a combination of two elements: strong word-of-mouth effects among regular users, where they share information about the influencers they follow, and a limited number of influencers that each user follows. This unique combination creates a "path dependence" for early success.

Using our model, we explored the potential effects of various platform algorithms on user engagement. In the ongoing debate regarding the role of OSNs in societal polarization, some argue that algorithms intentionally curate content aligned with users' existing opinions. However, our model demonstrates that such algorithms could potentially reduce polarization and even lead to a more homogeneous distribution of beliefs within the population, assuming users remain open to diverse viewpoints.

This finding aligns with recent studies examining polarization on platforms like Facebook during the 2020 U.S. election, which showed minimal algorithmic influence on polarization. This suggests the need to explore alternative explanations for the observed rise in societal polarization. Studies by Bakshy et al. (2015) and Andris et al. (2015), for example, highlight the role of user behavior itself in limiting exposure to divergent viewpoints, regardless of the platforms' algorithms.

We acknowledge potential areas for future research, including incorporating elements related to the spread of misinformation and allowing for dynamic connections between regular users to increase the model's complexity and reflect real-world scenarios more accurately. This research contributes significantly to our understanding of the multifaceted interactions within online social settings and provides valuable insights for further investigations.

0.9 Chapter 5

In this final chapter, we aim to elucidate the correlation between the policies and rules set by online social networks - termed network protocols - and users' experiences. We contend that, beyond the network effects that partially determine a platform's success, the protocols shaping the network architecture play a significant role in various user-driven processes on those networks.

To test this hypothesis, we begin by defining the network protocols. We focus on three different aspects: the nature of the ties allowed by the platform, which can be either directed or undirected; the users' ability to create and utilize sub-communities on the platform; and the type of recommendation algorithms that determine connections (real friends, individuals with similar interests, or highly connected agents). We then generate the various combinations of these rules, which translate into network-generating rules.

We observe significant variations when comparing these structures in terms of average path length or clustering - as expected from the generating rules. The next step is to test whether these variations lead to different user experiences. To do so, we focus on experiences likely to occur on an OSN and for which users can have an objective preference function. Specifically, we examine the efficiency of the network for information diffusion, the dynamics of collective action, and the effectiveness of the network in connecting job seekers with job providers.

Overall, the hypothesis that the network protocols, which influence the network structure of OSNs, also have an impact on common processes occurring in these media, is theoretically validated by our models and simulations.

There is high heterogeneity in the most and least performing structures, whether we examine the diffusion of information, the leverage of collective actions, or the performance of a job market. What's even more interesting is that an enhancing characteristic for one process might not be for another.

This research highlights an opportunity for OSN platforms to re-evaluate their goals and revise network protocols accordingly. Users benefit even more from this approach, as they can then choose the right platform for their specific needs. Strategic selection can significantly increase success rates. For instance, our models suggest that platforms like Instagram or Twitter, with structures closer to DSN, might not be ideal for job searches. Conversely, platforms like Facebook or Reddit, which facilitate the formation of sub-communities, can be powerful tools for mobilizing collective action.

0.10 Methodology

Apart from Network analysis, which is both an object of study, and a set of tools to understand the structure of interactions, all four chapters use numerical simulations, also called Agent-Based Models (ABMs) to study the models we imagine.

An agent-based model defines a population of autonomous agents who are given a set of rules defining their behavior, thus their responses to the possible situations they might be confronted with. Agents can represent individuals, institutions, firms, or any entity that will act independently. Furthermore, an ABM is not restricted to one type of agent, but allows study the co-existence of different type of entities, such as firms and states, or consumers and firms for example.

Moreover, ABMs allow our agents to interact among themselves, and thus we can study the implications of those interactions for individuals, but also for population-level variables. The final important ingredient for an ABM is its allowance for the simulation of time, which is crucial if we wish to understand dynamics.

An ABM is therefore a definition of a population, a set of rules that defines the goals, behaviors, reactions and consequences of interactions of each agent of this population — with each agent's set of rules potentially different than the ones of the others — but also of the existing (or potential) interactions among the agents, and of the time-sequence in which our population evolves.

Those type of models are very convenient to establish a bridge from the individual to the aggregate behavior. This interaction from the micro to the macro-level can lead ABMs to be constructed with different objectives in mind.

We can use them either to make sense of the world we witness — our data — by looking for the fitting behaviors and interactions of our individual agents, or we can build them to observe the world that arises when we hypothesize about the likely behaviors and interactions of agents in a given context.

Imagine that you can either have a collection of bones or the image of a body in motion and must recreate the working skeleton of that body. In the first philosophy, you pick the image of the body in motion and deduce the right place and interactions of the bones in a manner that would fit your image. This mirrors the approach of creating ABMs based on empirical observations, where each 'bone' or agent behavior is positioned to match the dynamics we observe in the real world. In the second one, you pick the collection of bones and arrange them in the most coherent way without having the image of the body in motion in mind. This is akin to developing ABMs grounded in theoretical frameworks, arranging agent behaviors and interactions logically to explore potential emergent phenomena.

In both methods, you're playing with bones. In the first one, so they fit the body. In the second one, so they articulate logically. Both methods are useful, and both methods have perils. You might want so hard for your bones to fit your picture that you're not thinking enough about the logical articulations of them, and your body might not walk for long. But with the other philosophy, you might exclude too much of the real body for the sake of the respect of logic, and the body you create might walk, but it might be out of this world, not representing the real body.

The models of this thesis were mostly built by picking bones and articulating them logically, aiming to strike a balance between theoretical coherence and empirical fidelity, endeavoring to reflect the familiar yet uncovering insights that might not be immediately apparent from observation alone, hopefully representing a body that is familiar to us.

Introduction Générale

0.11 Économie de la Complexité

Chaque fois que je présente aux étudiants de première année en économie la théorie classique du consommateur en microéconomie, je rencontre une perplexité généralisée dans la salle de classe. Des questions surgissent : Pourquoi les consommateurs sont-ils décrits comme désirant toujours plus de biens ? Dans le modèle travail-loisir, pourquoi suppose-t-on que les agents choisissent librement leur quantité optimale de travail ? Et pourquoi le travail est-il évalué uniquement à travers ses bénéfices monétaires, en ignorant les emplois qui offrent une valeur intrinsèque au-delà du gain financier ? Bien sûr, il y a des raisons pour ces modèles — qu’elles soient satisfaisantes ou non — enracinées dans leurs objectifs ou leurs simplifications pour la facilité analytique.

Mais, chose intéressante, lorsque les étudiants atteignent leur dernière année de licence, leur présenter des modèles où les entreprises choisissent de ne pas produire ou des modèles intégrant la connaissance de la fonction de réaction d’un concurrent dans un duopole suscite à peine des questions. Cela indique-t-il une compréhension plus profonde de l’économie, ou une résignation à l’abstraction de la réalité inhérente à ces modèles ?

Un changement particulier semble se produire au cours de l’enseignement de l’économie. Comme l’a souligné Etzioni en 2015 Etzioni (2015), les étudiants en économie développent une “boussole morale dégradée”. Quels facteurs contribuent à cette transformation ?

Un segment croissant du domaine soutient que le problème fondamental réside dans les hypothèses de base de l’économie Kirman (1989, 1992); Arthur (2009). Depuis la critique d’Herbert Simon en 1955 du modèle de l’agent rationnel maximisant l’utilité Simon (1955), un débat important a surgi sur le réalisme de ces hypothèses. Simon a souligné l’erreur de supposer une rationalité parfaite, une critique qui résonne jusqu’à la macroéconomie où la notion d’équilibre peut masquer des problèmes significatifs, comme en témoigne la crise financière de 2008 peu de temps après l’évaluation optimiste d’Olivier Blanchard sur l’état de la théorie macroéconomique Blanchard (2009).

Cette reconnaissance des limites inhérentes aux modèles économiques traditionnels a suscité un intérêt pour des approches alternatives. Ces approches reconsidèrent le comportement des agents économiques, en reconnaissant la rationalité limitée qui caractérise la prise de décision humaine, comme l’a suggéré Simon, et explorée plus avant dans le domaine de l’économie comportementale Gigerenzer and Selten (2002). Les agents, influencés par des biais cognitifs et sans la puissance de calcul supposée dans les modèles classiques, optent souvent pour des stratégies de “satisficing” plutôt que d’optimisation Kahneman (2003).

De plus, le comportement économique n’est pas uniforme mais varie selon des facteurs culturels, socio-économiques, et autres, soulignant l’hétérogénéité des agents. En outre, les agents apprennent et s’adaptent en fonction de leurs expériences, avec des préférences et des comportements évoluant au fil du temps. Cette dynamique contredit la nature statique des fonctions d’utilité ou des courbes de demande dans les modèles classiques.

Les interactions entre les agents compliquent encore le tableau. Le comportement d’un individu est souvent influencé par son réseau social, ce qui conduit à des changements de préférences et de comportements que les modèles classiques, avec leurs hypothèses sur des consommateurs indépendants, ne parviennent pas à capturer. Cette dimension sociale peut conduire à des comportements tels que des changements de consommation basés sur l’influence des pairs ou des efforts de différenciation. Un agent peut choisir un bien non seulement pour sa valeur intrinsèque, mais parfois parce qu’il est devenu désirable en le voyant consommé par d’autres, ou parfois parce que sa consommation permettra de se différencier des autres. La consommation est un acte social qui étend notre identité.

L’implication est que les phénomènes macroéconomiques ne peuvent pas être entièrement expliqués en agréant les comportements des agents isolés ; au contraire, les interactions et les adaptations entre les agents jouent un rôle crucial. Le modèle de ségrégation de Schelling illustre vivement ce point Schelling (1969).

L’auteur imagine une population vivant sur une grille, divisée en deux classes d’agents. Chaque agent observe la classe de ses voisins et a un certain désir d’être entouré d’agents de la même classe. Si le nombre de voisins similaires est trop faible, l’agent se déplace au hasard vers un espace vide sur la grille. Le modèle et son résultat sont devenus célèbres car Schelling montre que même avec une population ayant un faible désir d’être entourée de pairs — ne se déplaçant que si plus des deux tiers de ses voisins sont d’un type différent — le résultat global finit par être une grille ségréguée, avec de grands quartiers ne contenant qu’une seule classe d’agents.

Cela n’implique pas que les ghettos ou les quartiers fortement ségrégués observés à travers l’histoire et aujourd’hui peuvent être entièrement expliqués par ce modèle et la volonté individuelle des agents locaux, car de nombreuses variables sont en jeu, y compris des actions politiques

significatives. Néanmoins, cela montre que même avec des agents assez ouverts à la diversité, la ségrégation peut survenir.

L'idée centrale de l'économie complexe, mise en lumière par le modèle de ségrégation de Schelling et l'adage de Simon selon lequel "le tout est différent de la somme de ses parties", souligne que le comportement des agents individuels, lorsqu'il est combiné à travers des interactions, conduit à des résultats collectifs qui peuvent ne pas être intuitifs sur la base des seules actions individuelles. Ce principe révèle que les dynamiques globales d'un système, comme les niveaux de ségrégation au sein d'une ville, ne peuvent pas être précisément déduites simplement en analysant les préférences de ses résidents isolément. Les schémas de ségrégation inattendus résultent de l'interaction nuancée des choix individuels, soulignant la nécessité pour la modélisation économique d'incorporer plus profondément les complexités du comportement humain et de l'interaction sociale. Cela souligne que comprendre les phénomènes macroéconomiques ou sociétaux nécessite une considération minutieuse des interactions au niveau micro qui conduisent ces systèmes, nous défiant de repenser les approches traditionnelles de la théorie économique.

0.12 Les Réseaux Sociaux pour Comprendre l'Hétérogénéité Sociale

Comme mentionné précédemment, pour parvenir à une compréhension enrichie des agents économiques préconisée par les partisans de l'école de pensée de l'économie complexe, nous devons reconnaître que, bien que le comportement individuel - tel que la prise de décision hétérogène reposant sur des heuristiques, par exemple - soit une pièce du puzzle, il y a au moins une pièce tout aussi importante concernant les interactions sociales.

Dans ce qui suit, nous définissons les interactions sociales entre agents économiques comme tous les processus qui poussent un agent à se comporter différemment en raison de ses interactions avec les autres. Cette définition large peut rendre compte de nombreuses situations et intentions différentes, que l'influence soit exercée par des individus ou des groupes, ou reçue par eux. Il est important de reconnaître que l'influence sociale qui résulte des interactions peut être consciente ou inconsciente.

Au cœur de l'idée d'influence sociale se trouve la tendance des humains à imiter et intégrer ce que font les autres. Cela peut être motivé par le désir de s'intégrer dans un groupe et d'appartenir, de se différencier en faisant le contraire, de démontrer des valeurs contrastantes envers un groupe, ou simplement de réussir. Il peut sembler rationnel d'imiter les agents à succès pour devenir soi-même couronné de succès.

Les influences sociales sont évidentes dans la consommation de nombreux types de biens. Nous appelons "mode" le processus continu des tendances, qui peut être observé dans les vêtements, les meubles, l'architecture, les biens culturels, et qui représente essentiellement un processus de validation sociale ou la mise en avant de certaines caractéristiques jugées actuellement désirables par le marché de masse.

La plupart d'entre nous, que ce soit intentionnellement ou non, participons à ce type de consommation. Nous pouvons chercher à être des pionniers de tendances, des suiveurs délibérés de la tendance, ou même choisir inconsciemment des biens qui sont à la mode.

En outre, nous pouvons chercher à signaler notre appartenance à une certaine classe sociale. Nous achetons des biens haut de gamme pour montrer notre succès financier, ou du moins pour démontrer notre aspiration à rejoindre ce groupe social particulier Veblen (2017); Leibenstein (1950); Bourdieu (1979); Johnson et al. (2018).

Bien que ce type d'influence sociale puisse sembler non coordonné, bien qu'en partie poussé par les marques, dans d'autres domaines, l'influence sociale est beaucoup plus explicite. Dans son célèbre livre, "L'éthique protestante et l'esprit du capitalisme", Max Weber lie directement le protestantisme, une branche du christianisme qui incarne un ensemble de valeurs, au comportement au travail de ses croyants, contribuant ainsi à l'émergence des sociétés capitalistes occidentales en Europe au XIXe siècle Weber and Kalberg (2013).

Ces institutions, qui sont partiellement définies par des directives écrites et des règles, ainsi que par des hiérarchies, cultivent également des communautés soudées parmi leurs membres. Ces communautés peuvent servir de pression pour les individus qui ne souhaitent pas se conformer à toutes les règles, ou d'aspiration pour ceux qui s'efforcent de s'y conformer plus étroitement Carvalho (2019).

De plus, les avancées récentes en science révèlent que certains comportements sont influencés par les connaissances, bien que nous ne comprenions peut-être pas entièrement les mécanismes par lesquels cela se produit. Les comportements de santé tels que l'obésité ou le tabagisme, par exemple, ont montré une corrélation positive avec des comportements similaires chez les pairs Christakis and Fowler (2007, 2008); Centola (2010).

Par conséquent, penser que nous pouvons analyser le comportement des acteurs économiques tels que les travailleurs, les consommateurs ou les entrepreneurs sans tenir compte du milieu social dans lequel ils évoluent et des interactions sociales qu'ils entretiennent revient à ne voir que la moitié du tableau.

Modéliser ou quantifier les interactions sociales est plus facile à dire qu'à faire, ce qui explique probablement pourquoi le domaine de l'économie a tardé à s'y engager profondément jusqu'à récemment. Ce changement s'est produit lorsque la science des réseaux a commencé à gagner

en reconnaissance et en visibilité au sein de la communauté scientifique.

Les réseaux formalisent et permettent la visualisation des connexions qui lient divers nœuds entre eux. Dans le contexte des interactions sociales, les nœuds peuvent représenter des individus ou des groupes, tandis que les liens peuvent représenter divers types de connexions entre deux nœuds, allant des liens familiaux aux relations de travail formelles, par exemple.

Construire un réseau est donc l'énumération explicite de tous les agents que nous souhaitons incorporer, ainsi que de toutes les connexions existantes entre eux, en fonction de l'importance que nous attribuons à ces liens.

L'existence et l'étude des réseaux - également connus sous le nom de graphes dans le domaine mathématique dont ils sont issus - remontent au moins au XVIII^e siècle et à la célèbre résolution par Euler du problème des "Sept ponts de Königsberg" Euler (1953).

Mais l'émergence de son existence en tant que champ à part entière - la "science des réseaux" - est beaucoup plus contemporaine et s'explique par au moins deux facteurs principaux selon Barabasi :

1. Pour établir les cartes complexes des réseaux que nous voulons étudier, beaucoup de données sont nécessaires. Un réseau social nécessite des informations précises sur vos amis, les amis de vos amis, et ainsi de suite. Si nous voulons construire la carte d'Internet, nous devons connaître chaque page web et chaque lien sur les pages web qui redirige vers d'autres. C'est précisément la révolution Internet qui a facilité la création de ces cartes, avec des méthodes de partage de données efficaces et rapides et un stockage numérique bon marché. Les chercheurs sont maintenant capables de collecter, d'assembler et d'analyser des réseaux réels et complexes.
2. Bien que les réseaux aient été partiellement étudiés dans différents domaines tels que la biologie, la chimie ou les sciences sociales, c'est l'universalité des propriétés des réseaux qui a rendu possible l'émergence de ce champ. Comme l'écrit Barabasi Barabási (2016) : "Les architectures des réseaux émergents dans divers domaines de la science, de la nature et de la technologie sont similaires entre elles, conséquence d'être régies par les mêmes principes d'organisation. Par conséquent, nous pouvons utiliser un ensemble commun d'outils mathématiques pour explorer ces systèmes."

À la veille du XXI^e siècle, deux modèles importants ont commencé à expliquer certaines des caractéristiques universelles des structures de réseau. Le premier, proposé par Barabasi et Albert, identifie la nature "sans échelle" de nombreux réseaux différents tels qu'Internet, les réseaux électriques ou les interactions sociales. Ici, la caractéristique mise en avant est la distribution des connexions pour les nœuds du graphe. Plus précisément, tandis que la majorité

des nœuds auront un nombre de connexions faible et similaire, un petit groupe aura un nombre de connexions significativement plus élevé Barabási and Albert (1999).

Le deuxième modèle propose de démontrer que la plupart de nos schémas d'interaction se produisent dans des "petits mondes", où même dans des systèmes avec de nombreux nœuds, atteindre l'un de ces nœuds peut être accompli par de petits "sauts" sur le réseau, tandis que la structure du réseau reste étroitement liée, impliquant un regroupement parmi de petits groupes de nœuds Watts and Strogatz (1998).

Ce modèle explique des phénomènes tels que la célèbre expérience de Milgram, où il demandait à des individus d'atteindre d'autres personnes (qui étaient inconnues et géographiquement éloignées) Milgram (1967). Les résultats ont montré que tout agent peut atteindre un autre en environ six étapes, ce qui a donné naissance à l'idée du "petit monde" dans lequel nous vivons.

Au-delà de la démonstration des propriétés structurelles partagées par une grande diversité de réseaux, des technologiques aux biologiques et sociaux, ces modèles éclairent également les outils de mesure que tout scientifique peut utiliser pour analyser un réseau.

Les structures sans échelle mettent en lumière l'importance de la distribution des degrés à l'échelle globale, mais aussi de la centralité des degrés à l'échelle locale. Un réseau est donc utile pour comprendre les dynamiques des phénomènes globaux, tels que la vitesse de diffusion, par exemple, mais aussi pour comprendre le rôle des acteurs spécifiques dans ces structures. Est-ce que l'agent le plus central dans un réseau a les mêmes opportunités ou connaissances que le moins central, par exemple ?

Il en va de même pour le modèle des petits mondes. Il montre que nous pouvons mesurer les graphes à l'échelle globale en utilisant des concepts tels que la longueur de chemin, que nous voulons mesurer sa moyenne ou son plus court, et avec le concept d'enracinement, qui peut mesurer l'existence et la force des communautés dans les réseaux. Mais il s'applique aussi au niveau individuel : les petits mondes mettent en avant l'idée de la centralité d'intermédiation et de l'importance des nœuds qui peuvent agir comme des ponts entre différentes communautés éloignées.

Bien que des modèles issus des mathématiques et de la physique aient permis à la science des réseaux d'émerger en tant que champ distinct, certaines des intuitions de la science des réseaux étaient déjà présentes et faisaient l'objet de travaux dans d'autres disciplines.

La notion de capital social — qui pourrait être décrite comme la capacité à mobiliser des ressources à partir des connexions sociales — existe en sociologie depuis le XIXe siècle et a été progressivement utilisée et développée tout au long du XXe siècle.

De nombreux auteurs ont utilisé ce concept, tels que Bourdieu Bourdieu (2021), qui analyse ses effets et son interdépendance avec d'autres formes de capital ; Putnam, qui montre son

importance pour l'altruisme et les institutions démocratiques Putnam (2015) ; et Coleman, qui hypothèse que le capital social agit comme un vecteur de conformité au sein des sociétés Coleman (1998).

Bien qu'aucun de ces auteurs n'utilise les réseaux comme outil pour étudier le capital social directement, l'idée même de réseaux sociaux est déjà présente et utile dans leur analyse de la société et de ses individus.

Plus récemment, les objets de réseau sont devenus centraux dans les travaux de Granovetter dans les années soixante-dix Granovetter (1973) et de Burt dans les années quatre-vingt-dix Burt (2018). Les deux auteurs soulignent l'importance de la structure du graphe aux niveaux global et individuel, que ce soit pour générer des "bonnes idées" ou pour trouver des emplois.

Avec les outils formalisés en mathématiques et en physique, et le concept de plus en plus utilisé en sciences sociales, les réseaux sont naturellement devenus un outil et un objet d'étude en économie au début des années 2000 avec le développement de modèles étudiant son impact sur la performance de correspondance sur le marché du travail, comment il peut être utilisé pour étudier les réseaux de production ou comment il aide à étudier la diffusion des connaissances, de la technologie ou des comportements Jackson (2010).

La présente thèse utilisera également des réseaux, mais pour souligner leur importance dans le processus d'apprentissage des individus, dans des cadres où les interactions et la structure de ces interactions ont des implications importantes aux niveaux individuel et global.

0.13 Apprendre au Sein des Réseaux Sociaux

Les réseaux sont devenus particulièrement importants dans l'étude de la diffusion. Que l'on parle de la diffusion des connaissances, de l'information, des virus ou de l'électricité, utiliser des réseaux pour décrire le processus est pratique car cela génère une carte décrivant toutes les routes qu'un objet donné peut utiliser pour atteindre les nœuds existants, d'où les points d'intérêt de cette carte.

Avoir une carte est un départ de l'inconnu. Cela nous permet de comprendre, et parfois de prédire, l'ampleur ou la vitesse d'un processus de diffusion car cela montre explicitement les routes existantes ou les nœuds non atteignables du système. Cela peut également servir de méthode pour identifier les nœuds importants, ceux qui sont essentiels pour la diffusion, et ceux qui peuvent être déconnectés sans entraver le processus de diffusion.

En retour, un processus de diffusion a lieu uniquement si le réseau sur lequel il se produit n'est pas vide (au moins deux nœuds sont connectés), et être capable de représenter ce réseau est crucial pour comprendre les dynamiques de la diffusion.

Comme mentionné ci-dessus, nous pouvons étudier la diffusion d'objets "réels" tels que les personnes se déplaçant d'une station de métro à une autre, les virus d'un corps à un autre, ou les signaux Wi-Fi des opérateurs vers les maisons et les ordinateurs. Cependant, nous pouvons aussi étudier la diffusion des croyances, des émotions, de l'information et des connaissances parmi les individus d'une société. Comme ces facteurs sont au cœur de nombreuses décisions et comportements, ils sont susceptibles d'expliquer (certains aspects) des processus économiques.

Le cas de la diffusion parmi les individus peut être catégorisé comme apprentissage social, où la façon dont les agents envoient, acceptent et intériorisent des informations ou des idées répond à différentes règles en fonction de la nature de ce qui est échangé, de l'identité des personnes ou du groupe qui l'envoie, mais aussi de la nature du destinataire individuel et de la culture de la société d'où il provient.

Nous ne réagissons pas de la même manière aux informations sur l'évolution de la valeur d'une action d'une entreprise qu'aux derniers sondages présidentiels. De même, nous valorisons différemment l'information selon qui nous l'a donnée. Était-ce un ami informé ou une source médiatique partisane ? Enfin, notre parcours personnel façonne notre manière de traiter les informations reçues. Cela nous amènera-t-il à investir dans l'entreprise ou à manifester contre la montée d'un parti politique extrême ?

L'apprentissage social est donc un mécanisme complexe, et son résultat, son efficacité, peut dépendre de nombreuses variables cognitives et émotionnelles telles que la confiance, la croyance ou la peur - pour n'en nommer que quelques-unes - en fonction du scénario étudié. Cependant complexe et difficile à comprendre exactement comment il fonctionne, l'apprentissage social peut être décrit comme un processus de diffusion, sur un réseau où les nœuds sont des individus, et les liens représentent des routes possibles pour échanger des morceaux d'information, de connaissances, mais aussi des croyances.

Les modèles fondamentaux visant à représenter l'apprentissage social, ou du moins l'apprentissage par observation, sont principalement basés sur deux mécanismes : l'utilisation de moyennes pondérées ou l'utilisation de la règle de Bayes (voir Goyal (2011) pour une revue approfondie).

Dans le premier scénario, provenant de DeGroot DeGroot (1974), les agents placés sur une structure de réseau peuvent observer le comportement ou les croyances de leurs voisins directs (ceux avec qui ils partagent un lien), et adapter leur propre comportement ou croyances en conséquence. Dans ce cadre, un agent donne des poids fixes à ses voisins, exprimant la valeur qu'il accorde à chaque voisin. Ensuite, son comportement au fil du temps est la moyenne pondérée de ce que croient/choisissent ses voisins, qui suivent eux-mêmes le même processus avec leurs propres voisins.

Ce type d'apprentissage par observation peut être utile pour représenter la dynamique des opinions où les croyances des agents évoluent à mesure qu'ils considèrent les opinions des autres,

que ce soit par communication ou par observation. Plusieurs questions peuvent être posées en étudiant ce genre de modèle. Observons-nous une convergence des opinions ? Sous quelles conditions sur la distribution initiale des croyances ? Quel est le rôle de la structure du réseau ?

Bien que ce soit un modèle utile, il présente également plusieurs limitations. Par exemple, il suppose qu'à chaque tour, un agent collecte des informations sur les opinions de tous ses voisins. Ce genre de connaissance ubiquitaire est contestable en soi.

La deuxième grande limitation réside dans le mécanisme de mise à jour. Le scénario dans lequel les agents compilent des opinions et adoptent la moyenne pondérée de leur cercle d'influence est discutable. Cette approche peut ne pas être adaptée ou applicable dans divers contextes.

Par exemple, nous étudions la dynamique des opinions en partie pour comprendre les élections démocratiques. Supposons qu'un agent ait deux voisins opposés sur le spectre politique. Cet agent donne des poids égaux à ces deux amis, se faisant ainsi un électeur centriste. Maintenant, lorsque l'élection arrive et qu'aucun candidat centriste n'est en lice, vers qui ira le vote de l'agent ? Ce qu'il choisira révélera les poids réels qu'il accorde à ses deux amis parce que le résultat est binaire ; ce ne sera pas la moyenne pondérée de ces deux opinions. Si le but était de comprendre les résultats des élections, le modèle était mal adapté car la moyenne pondérée ne correspond pas aux décisions finales que les agents devraient prendre lors de l'élection (bien qu'elles puissent toujours avoir un certain pouvoir prédictif).

Bien sûr, bien qu'il y ait de nombreux cas où le modèle est limité, il existe autant de scénarios où ce type de modélisation peut être utile, car tous les processus ne se terminent pas par une décision binaire.

Le deuxième scénario, où les agents utilisent la règle de Bayes, comme dans Acemoglu et al. (2011), utilise principalement un modèle où les agents sont également placés sur une structure de réseau qui limite leurs interactions à leurs quartiers et leur permet d'observer les actions et les résultats uniquement de ces voisins. Cela leur permet ensuite de mettre à jour leurs croyances antérieures sur les résultats d'un événement donné. Dans ce genre de modèles, les agents essaient de trouver la "vérité" ou le comportement optimal, ce qui peut être vu comme une action à entreprendre ou non, mais aussi une opinion à avoir ou non.

Dans les deux courants de la littérature, et dans la plupart de la littérature économique qui s'intéresse à l'apprentissage social, il y a principalement une vision commune selon laquelle les agents apprennent des autres en observant - et non par une communication en tête-à-tête, par exemple - l'ensemble de leurs voisins à chaque période. De plus, ces agents utilisent des méthodes complexes pour déterminer leurs actions, qu'il s'agisse de calculer une moyenne pondérée des décisions de leurs amis ou de mettre à jour une distribution de probabilité basée sur les résultats

de ces amis. Bien sûr, nous pouvons choisir de croire que ces modèles fonctionnent non pas parce que nous nous comportons de cette manière - en faisant des calculs complexes - mais parce que quoi que nous fassions, cela peut être approximé par ces calculs.

D'autres articles ont essayé d'utiliser une approche peut-être plus réaliste en dotant les agents de leurs modèles de "règles empiriques", qui sont des heuristiques simples. C'est le cas d'Ellison et Fudenberg dans le contexte de l'adoption de différentes technologies Ellison and Fudenberg (1995, 1993). Leurs règles empiriques impliquent également une population observant les choix et les résultats de tous les voisins, mais maintenant seule une fraction d'entre eux révisé leurs décisions à chaque tour, et au lieu de calculs complexes, ils utilisent des règles simples telles que : Utiliser ce qui était le plus rémunérateur au dernier tour, ou utiliser ce qui était le plus populaire au dernier tour.

Dans le même esprit, W.B. Arthur construit un modèle intéressant, le problème du "El-Farol Bar", où les agents veulent éviter de se rendre dans un bar s'il est trop bondé. Chaque fois qu'ils y vont, ils peuvent observer s'il est effectivement trop bondé ou non et adapter leur comportement en conséquence pour la prochaine fois qu'ils doivent décider d'y aller ou non Arthur (1994).

Bien que n'utilisant pas d'outils de réseau, chaque agent de ce modèle apprend en observant le comportement de ses voisins géographiques - ceux qui se rendent également dans ce bar particulier. Arthur propose que, bien que les agents utilisent des règles empiriques pour décider quoi faire, ils sont hétérogènes quant à la règle empirique à utiliser, créant ainsi une "écologie" de règles.

En dehors de l'économie, certains auteurs ont développé des modèles d'apprentissage social Granovetter (1978). Notamment en sociologie, Granovetter développe un modèle d'action collective où les agents doivent décider s'ils rejoignent ou non un mouvement, en fonction de la décision de leurs voisins qu'ils peuvent observer. Il met en avant une règle empirique intéressante où les agents, au lieu de faire comme la majorité ou de faire la moyenne des voisins, décident de s'engager uniquement si une proportion précise de leur quartier a déjà rejoint l'action collective.

Dans son article, Granovetter utilise l'exemple d'une émeute. Disons que chaque agent dans une foule doit décider s'il participe ou non à une émeute. C'est un comportement risqué qui peut avoir des répercussions, donc, la plupart des gens prêts à s'y engager ont également besoin de voir qu'ils ne sont pas seuls dans le mouvement. En fonction de la force de leur croyance en la cause pour laquelle ils sont prêts à se révolter, ils doivent observer une proportion différente de la population déjà engagée. Ainsi, l'émeute est déclenchée par les agents les plus convaincus, ce qui incite les personnes qui devaient voir ces quelques émeutiers à s'engager elles-mêmes, et ainsi de suite jusqu'à ce que chaque agent ayant même la moindre volonté de se révolter — ceux qui ont besoin d'observer une très grande proportion de la population participant —

s'engagent dans l'émeute. Il y a donc une dynamique de cascade des croyants les plus forts aux plus faibles. Le modèle est assez différent de ceux où les agents s'engagent s'ils observent une majorité d'agents déjà engagés. Avec la même foule et la même distribution de la volonté de s'engager dans l'émeute, seuls les plus convaincus participeront, car chaque autre agent de la foule observera une minorité d'émeutiers (il est important de noter que cette configuration suppose que chaque agent observe le comportement de chaque autre dans la foule, ce qui n'est pas si courant).

En général, comprendre et représenter avec précision l'apprentissage social est complexe et hautement dépendant du contexte. Il n'existe pas de modèle canonique qui conviendrait à tous les types d'échanges d'informations entre individus.

Notre ambition avec cette thèse est d'approfondir certains des modèles existants en ajoutant des variables qui peuvent les rendre plus adaptés à des contextes spécifiques, mais aussi de proposer nos propres mécanismes d'interactions qui peuvent conduire à l'apprentissage social.

Par exemple, la plupart des modèles existants sur l'apprentissage social écartent la possibilité pour les individus de ne pas utiliser les actions de leurs pairs. Mais en réalité, les agents réagissent différemment en fonction de leurs observations des actions de leurs voisins. De plus, ils pourraient apprendre quel type de comportement les récompense le plus à long terme. Et s'ils choisissent d'être influencés par leurs voisins, ils pourraient également apprendre quels sont les meilleurs voisins à écouter. Ces types de questions seront au cœur de notre premier chapitre.

Une autre idée importante que nous utilisons tout au long de la thèse est celle des réseaux construits de manière endogène. Comme ils représentent les interactions entre les agents, un lien entre deux nœuds est susceptible de représenter plus qu'un simple canal de communication. Au lieu de cela, l'existence de ces liens peut être construite en utilisant les processus que nous savons être à la base des interactions sociales. L'homophilie, par exemple, peut être utilisée pour représenter des agents se connectant par similitude ; la fermeture triadique pour que les amis d'amis deviennent eux-mêmes amis ; ou l'attachement préférentiel lorsque les individus choisissent de se connecter à des agents populaires sur le réseau.

Ces concepts peuvent être utilisés pour permettre aux agents de former et de rompre des connexions, conduisant à des structures de réseau réalistes telles que les graphes sans échelle ou les petits mondes, qui seront utilisés comme l'architecture des interactions de nos populations. Comme nous le montrerons, les interactions peuvent mener à des résultats très différents selon la structure dans laquelle les individus sont placés.

Dans l'ensemble, nos modèles d'apprentissage social encapsulent de nombreux processus humains - bien qu'ils ne soient pas toujours au centre du modèle, et donc lourdement discutés - en fonction du contexte. Cela inclut la confiance, l'attraction, la répulsion, l'imitation, ainsi

que l'honnêteté et la malhonnêteté. De plus, les interactions peuvent être liées à la mémoire ou motivées par des intérêts soit intéressés, soit désintéressés.

En plus d'être centrée sur la notion d'apprentissage social, cette thèse vise à explorer les changements apportés par Internet et les médias sociaux. La plupart des modèles développés ici se situent dans le contexte des médias sociaux, englobant les changements qu'ils apportent à la façon dont les gens communiquent, mais aussi aux nouveaux acteurs qu'ils créent, des influenceurs aux plateformes de réseaux sociaux en ligne (OSN)

0.14 Internet et l'Ère des Réseaux

Un récent rapport de Pew Research indique que plus de 40 % des utilisateurs de TikTok utilisent la plateforme pour s'informer Matsa (2023). Avec environ 150 millions d'utilisateurs actifs mensuels aux États-Unis Statista (2023), il est donc sûr de dire qu'une partie significative de la société américaine est partiellement informée par un réseau social en ligne dont les algorithmes sont inconnus et déterminés par son entreprise chinoise.

À mesure que les médias sociaux en ligne deviennent omniprésents dans nos vies, il est important que nos modèles s'adaptent à ce nouveau paradigme. Plus précisément, les interactions sociales sont profondément modifiées par les OSN (Online Social Networks). Premièrement, ils modifient le nombre d'interactions que l'on peut avoir. Ils nous permettent d'envoyer des messages à n'importe quel voisin de ce réseau, et même de communiquer avec tout un quartier en partageant des publications. Cela illustre que non seulement les plateformes en ligne modifient la portée de nos interactions, mais aussi leur nature. Publier un statut sur Facebook ou Twitter, visible par tous ceux qui sont connectés à vous, n'a pas d'équivalent à l'ère pré-Internet, sauf peut-être lors d'événements très spécifiques tels que les mariages ou les funérailles, où le contenu du discours est très contrôlé. Cette rupture radicale avec la communication traditionnelle impacte nécessairement notre compréhension des opinions ou des tendances globales car nous n'avons jamais été confrontés à de telles formes de communication auparavant.

De plus, les OSN ne sont pas des lieux neutres ; ce sont des entreprises qui se font concurrence pour attirer votre attention, pour le temps que vous passerez à les utiliser. Cela les pousse à développer des algorithmes qui peuvent déformer votre vision du monde.

Les plateformes sociales en ligne deviennent également le terrain de jeu des influenceurs, des leaders d'opinion qui cherchent à tirer des revenus de leur audience.

Ce nouveau paysage est particulièrement adapté pour être étudié à travers le prisme de la science des réseaux, car les OSN sont essentiellement des objets de réseau. Ce que les plateformes proposent à leurs utilisateurs est simplement de créer un profil (un nœud) en ligne et de pouvoir se connecter à d'autres individus présents sur l'OSN (créant ainsi un lien entre

deux nœuds). La valeur d'un OSN particulier découle alors de deux éléments principaux : les personnes qui l'utilisent déjà - combien de personnes, mais aussi quel type de personnes - et les règles, ou politiques (que nous définissons plus loin dans la thèse comme les protocoles), que l'OSN adopte pour fournir la meilleure expérience possible, celle qui retiendra ses utilisateurs plus longtemps sur son site. L'exemple le plus courant est les algorithmes qui dictent quelles publications apparaissent dans le fil d'actualités de l'utilisateur.

La première source d'avantage concurrentiel est communément appelée effets de réseau dans la littérature et a fait l'objet de recherches par de nombreux auteurs depuis le milieu des années 80 Katz and Shapiro (1985); Arthur (1989); Cowan (1991). Ce type d'effet n'est pas restreint aux OSN et peut être trouvé dans de nombreuses industries, en particulier sur les marchés où la valeur d'un objet ou d'une technologie dépend de sa base d'utilisateurs existante. Cela peut se produire parce que la valeur découle du fait que d'autres l'utilisent aussi - comme le téléphone - ou de la réticence des gens à supporter les coûts de transition. Par exemple, vous pourriez préférer utiliser Excel, Python ou un clavier QWERTY non pas parce qu'ils sont les meilleures options sur le marché, mais principalement parce qu'ils sont devenus des standards dans la société, et si vous deviez passer d'un emploi à un autre, il y a de fortes chances que la nouvelle entreprise vous demande également de travailler avec ces produits David (1985). Ce second cas implique principalement des biens où il y a une phase d'apprentissage - comme apprendre à coder en Python, apprendre à utiliser les raccourcis et les fonctions, ou apprendre la mémoire musculaire pour savoir où se trouvent les touches sur le clavier.

En revanche, la deuxième explication du succès d'un OSN, ou du moins les leviers qu'il peut effectivement utiliser pour se différencier de ses concurrents, est moins étudiée en économie, principalement parce qu'il s'agit d'une méthode de différenciation relativement nouvelle qui a émergé avec les OSN à la fin des années 2000. Nous montrerons que la plupart des choix qu'un OSN fait pour définir les règles de sa plateforme peuvent être traduits en règles qui génèrent l'objet réseau lui-même. Par conséquent, nous pouvons étudier de nombreuses politiques que les entreprises prennent en montrant comment elles impactent la structure du réseau de l'OSN, puis en étudiant les impacts de la structure du réseau sur des processus tels que la diffusion de l'information ou la coordination pour s'engager dans une action collective. Cela, à son tour, aide à comprendre quels OSN pourraient être plus attractifs parce que, d'un point de vue objectif, ils sont plus efficaces pour certaines tâches. Ce travail fera l'objet du dernier chapitre de cette thèse.

Enfin, il est important de reconnaître que certains OSN se sont spécialisés dans des objectifs spécifiques. LinkedIn est utilisé pour partager sur la vie professionnelle, tandis qu'Instagram se concentre sur les photos. Ces spécificités donnent lieu à des comportements différents de la part des individus en fonction de l'OSN qu'ils utilisent. Nos modèles viseront à s'adapter en

conséquence.

À mesure que l'apprentissage social s'adapte à de nouveaux cadres, avec des personnes communiquant et observant les autres en ligne autant - sinon plus - qu'hors ligne, nous visons à nous adapter en conséquence, en proposant une vision nuancée et détaillée de l'apprentissage social en ligne, principalement en permettant aux réseaux sociaux de se former de manière endogène sous les contraintes des plateformes, et en reconnaissant l'existence de forces puissantes telles que les algorithmes de recommandation ou les influenceurs.

0.15 Plan de la thèse

Le chapitre 2 commence la thèse en se demandant si l'apprentissage social est le meilleur comportement disponible dans un contexte d'incertitude où les agents consomment des biens d'expérience. Le chapitre 3 adapte le modèle aux problèmes en ligne tels que l'existence de publicités d'influenceurs ou les systèmes de notation collaborative. Le chapitre 4 tente de capturer le rôle des influenceurs et des algorithmes dans les OSN pour la dynamique des opinions, en mettant l'accent sur la polarisation. Enfin, le chapitre 5 démontre comment les protocoles de réseau modifient les structures de réseau et, par conséquent, les dynamiques d'apprentissage social sur ces structures.

Les sections suivantes fourniront un aperçu complet de chaque chapitre et des méthodologies employées tout au long de cette thèse.

0.16 Chapitre 2

Ce chapitre explore les défis de la modélisation du comportement des consommateurs, en particulier lorsqu'ils choisissent parmi des options inconnues à travers des essais répétés. Nous soutenons que les individus utilisent diverses méthodes d'apprentissage et qu'il peut ne pas exister de stratégie "optimale" unique.

En nous concentrant sur les biens d'expérience, où la consommation elle-même procure un apprentissage, nous proposons un modèle d'apprentissage par renforcement multidimensionnel. Ce modèle capture la manière dont les agents apprennent non seulement à choisir des produits efficaces, mais aussi à affiner leurs processus de prise de décision au fil du temps.

Ainsi, les agents ont le choix entre apprendre à propos du meilleur produit par apprentissage par renforcement individuel ou par apprentissage social, en utilisant leurs réseaux sociaux pour obtenir des recommandations et choisir en conséquence.

Nous proposons donc une vision différente de l'apprentissage social. Ici, l'idée est similaire à une simple règle empirique : les agents augmentent la probabilité de choisir à nouveau des

options fortement rémunératrices par rapport à celles faiblement rémunératrices. Mais pour choisir parmi les alternatives, nous permettons aux individus de choisir entre l'apprentissage individuel ou social. L'influence des pairs n'est donc pas une obligation pour nos agents, qui décideront d'utiliser les recommandations de leurs voisins uniquement s'ils en viennent à croire que c'est un moyen plus efficace d'obtenir les meilleurs biens. De plus, nos agents n'observent pas l'intégralité des décisions ou des résultats de leurs voisins, mais choisissent seulement l'un d'entre eux.

Nos résultats révèlent qu'un modèle intégrant à la fois l'apprentissage individuel et social surpasse de manière significative un modèle limité à l'apprentissage individuel. Cela met en évidence l'impact positif de l'apprentissage social, même lorsque les individus manquent de connaissances préalables sur les options disponibles.

Le bénéfice de l'apprentissage social semble provenir des recommandations portant des informations implicites sur les options préférées. Cependant, l'apprentissage individuel présente également des avantages, permettant à certains agents de découvrir rapidement de bons choix. Néanmoins, l'apprentissage individuel comporte un risque plus élevé, car les individus peuvent se retrouver coincés sur des options sous-optimales.

Ainsi, bien que l'apprentissage social ne conduise pas à des choix parfaits pour tous les agents, il assure une performance constante à un niveau élevé. Nous explorons également le rôle des niveaux de curiosité optimale, qui permettent aux agents d'explorer suffisamment avant de s'engager dans un choix.

Enfin, le modèle intègre un élément de confiance, où les agents apprennent à identifier les sources de conseils fiables au sein de leur réseau social. Cet élément améliore la performance globale et conduit à des structures de réseau intéressantes, certains agents devenant des figures centrales pour les conseils.

0.17 Chapitre 3

Ce chapitre examine l'impact d'Internet sur la prise de décision, en particulier lors du choix de biens d'expérience - des produits dont la valeur est incertaine jusqu'après l'achat. Nous nous concentrons sur la manière dont les interactions en ligne, y compris les systèmes de notation en ligne et les influenceurs, influencent les choix individuels et les résultats dans de telles situations. Ici, nous utilisons le même modèle que dans le chapitre 2, avec des individus ayant la possibilité d'apprendre à choisir les meilleurs biens soit par apprentissage social, soit par apprentissage individuel. Nous l'étendons en ajoutant la présence d'influenceurs, en modifiant la structure du réseau et en le comparant à un modèle avec des avis en ligne agrégés.

La recherche explore les avantages potentiels de l'apprentissage social, où les individus ap-

prennent des expériences des autres, par rapport à l'apprentissage individuel par essais et erreurs. Internet facilite l'apprentissage social en favorisant les communautés en ligne et les plateformes comme les systèmes de notation. Cependant, l'efficacité de ces systèmes peut être limitée par la diversité des préférences des utilisateurs. Les avis agrégés peuvent ne pas refléter fidèlement les besoins individuels, soulignant l'importance de considérer la structure des interactions en ligne et les limitations potentielles de s'appuyer uniquement sur des informations agrégées.

En outre, ce chapitre se penche sur le rôle des réseaux sociaux en ligne et des influenceurs en ligne. Ces réseaux peuvent connecter des individus ayant des goûts similaires, ce qui peut conduire à des décisions plus éclairées. Cependant, l'action des influenceurs en ligne peut être préjudiciable, en particulier pour les biens de qualité moyenne. Les individus peuvent être influencés par des endorsements sans explorer pleinement les options alternatives, risquant ainsi de passer à côté de choix supérieurs. Cet effet est le plus prononcé dans les réseaux avec des structures présentant une attache préférentielle, que ce soit dans des cas extrêmes (réseaux étoilés) ou dans des contextes plus réalistes comme les réseaux sans échelle (courants sur les plateformes en ligne). Nous avons également constaté que certaines structures sont beaucoup moins impactées par les influenceurs. C'est le cas des réseaux de type "Small-World", des structures présentant une faible longueur de chemin moyenne et une forte cohésion.

Dans l'ensemble, ce chapitre suggère que, bien qu'Internet offre des outils prometteurs pour améliorer la prise de décision, une évaluation critique et une approche prudente vis-à-vis de l'influence en ligne sont cruciales pour naviguer dans les complexités du choix des biens d'expérience à l'ère numérique. En fin de compte, les interactions sociales avec des individus de confiance pourraient être la stratégie la plus efficace pour prendre des décisions éclairées dans de tels contextes. Cependant, l'efficacité de ces interactions, et l'influence d'Internet dans son ensemble, dépend de facteurs tels que la distribution des préférences des utilisateurs et la conduite éthique des influenceurs en ligne, ce qui peut ne pas toujours être idéal dans des scénarios du monde réel. Cette recherche contribue à une compréhension plus approfondie de ces défis et des avantages et inconvénients potentiels des interactions en ligne dans le domaine de la prise de décision pour les biens d'expérience.

0.18 Chapitre 4

Ce chapitre examine les interactions complexes entre différents groupes d'utilisateurs sur les réseaux sociaux en ligne (OSN), en se concentrant particulièrement sur la relation entre les aspirants influenceurs et la population générale des utilisateurs. Nous révélons un système dynamique où les deux groupes influencent le comportement de l'autre.

Les utilisateurs réguliers partagent et discutent principalement de leurs opinions, ce qui peut potentiellement conduire à une polarisation accrue ou à un plus grand consensus sur des sujets spécifiques. Les aspirants influenceurs, quant à eux, partagent stratégiquement du contenu pour maximiser leur succès en ligne. Fait intéressant, leur présence semble contrer la formation de points de vue extrêmes, empêchant à la fois le consensus complet et la polarisation complète.

Nous allons au-delà de la compréhension traditionnelle du succès dans les réseaux en ligne en identifiant les facteurs clés qui le favorisent pour les aspirants influenceurs. Le facteur le plus significatif semble être une combinaison de deux éléments : des effets de bouche-à-oreille forts parmi les utilisateurs réguliers, où ils partagent des informations sur les influenceurs qu'ils suivent, et un nombre limité d'influenceurs que chaque utilisateur suit. Cette combinaison unique crée une "dépendance au sentier" pour le succès précoce.

En utilisant notre modèle, nous avons exploré les effets potentiels de divers algorithmes de plateforme sur l'engagement des utilisateurs. Dans le débat en cours sur le rôle des OSN dans la polarisation sociétale, certains soutiennent que les algorithmes sélectionnent intentionnellement du contenu aligné avec les opinions existantes des utilisateurs. Cependant, notre modèle démontre que de tels algorithmes pourraient potentiellement réduire la polarisation et même conduire à une distribution plus homogène des croyances au sein de la population, à condition que les utilisateurs restent ouverts à des points de vue diversifiés.

Cette découverte s'aligne avec des études récentes examinant la polarisation sur des plateformes comme Facebook lors de l'élection américaine de 2020, qui ont montré une influence algorithmique minimale sur la polarisation. Cela suggère la nécessité d'explorer des explications alternatives pour la montée observée de la polarisation sociétale. Des études de Bakshy et al. (2015) et Andris et al. (2015), par exemple, soulignent le rôle du comportement des utilisateurs eux-mêmes dans la limitation de l'exposition à des points de vue divergents, indépendamment des algorithmes des plateformes.

Nous reconnaissons les domaines potentiels pour de futures recherches, notamment en incorporant des éléments liés à la propagation de la désinformation et en permettant des connexions dynamiques entre les utilisateurs réguliers pour augmenter la complexité du modèle et refléter plus précisément les scénarios du monde réel. Cette recherche contribue de manière significative à notre compréhension des interactions multiples dans les environnements sociaux en ligne et fournit des informations précieuses pour des investigations ultérieures.

0.19 Chapitre 5

Dans ce dernier chapitre, nous visons à élucider la corrélation entre les politiques et règles établies par les réseaux sociaux en ligne - appelées protocoles de réseau - et les expériences des

utilisateurs. Nous soutenons qu'au-delà des effets de réseau qui déterminent en partie le succès d'une plateforme, les protocoles qui façonnent l'architecture du réseau jouent un rôle significatif dans divers processus pilotés par les utilisateurs sur ces réseaux.

Pour tester cette hypothèse, nous commençons par définir les protocoles de réseau. Nous nous concentrons sur trois aspects différents : la nature des liens autorisés par la plateforme, qui peuvent être dirigés ou non dirigés ; la capacité des utilisateurs à créer et utiliser des sous-communautés sur la plateforme ; et le type d'algorithmes de recommandation qui déterminent les connexions (amis réels, individus ayant des intérêts similaires ou agents très connectés). Nous générons ensuite les diverses combinaisons de ces règles, qui se traduisent par des règles de génération de réseau.

Nous observons des variations significatives en comparant ces structures en termes de longueur de chemin moyenne ou de regroupement - comme prévu par les règles de génération. L'étape suivante consiste à tester si ces variations conduisent à des expériences utilisateur différentes. Pour ce faire, nous nous concentrons sur des expériences susceptibles de se produire sur un OSN et pour lesquelles les utilisateurs peuvent avoir une fonction de préférence objective. Plus précisément, nous examinons l'efficacité du réseau pour la diffusion de l'information, les dynamiques d'action collective et l'efficacité du réseau à connecter les chercheurs d'emploi avec les offreurs d'emploi.

Dans l'ensemble, l'hypothèse selon laquelle les protocoles de réseau, qui influencent la structure du réseau des OSN, ont également un impact sur les processus courants se produisant dans ces médias, est théoriquement validée par nos modèles et simulations.

Il y a une forte hétérogénéité dans les structures les plus performantes et les moins performantes, que nous examinons la diffusion de l'information, l'effet de levier des actions collectives ou la performance d'un marché du travail. Ce qui est encore plus intéressant, c'est qu'une caractéristique améliorante pour un processus peut ne pas l'être pour un autre.

Cette recherche met en évidence une opportunité pour les plateformes OSN de réévaluer leurs objectifs et de réviser les protocoles de réseau en conséquence. Les utilisateurs en bénéficient encore plus, car ils peuvent alors choisir la bonne plateforme pour leurs besoins spécifiques. Une sélection stratégique peut augmenter considérablement les taux de réussite. Par exemple, nos modèles suggèrent que les plateformes comme Instagram ou Twitter, avec des structures plus proches des DSN, pourraient ne pas être idéales pour les recherches d'emploi. À l'inverse, des plateformes comme Facebook ou Reddit, qui facilitent la formation de sous-communautés, peuvent être des outils puissants pour mobiliser l'action collective.

0.20 Méthodologie

En plus de l'analyse de réseau, qui est à la fois un objet d'étude et un ensemble d'outils pour comprendre la structure des interactions, les quatre chapitres utilisent des simulations numériques, également appelées modèles basés sur les agents (ABM), pour étudier les modèles que nous imaginons.

Un modèle basé sur les agents définit une population d'agents autonomes à qui l'on donne un ensemble de règles définissant leur comportement, donc leurs réponses aux situations possibles auxquelles ils pourraient être confrontés. Les agents peuvent représenter des individus, des institutions, des entreprises ou toute entité agissant indépendamment. De plus, un ABM n'est pas limité à un type d'agent, mais permet d'étudier la coexistence de différents types d'entités, comme les entreprises et les États, ou les consommateurs et les entreprises, par exemple.

De plus, les ABM permettent à nos agents d'interagir entre eux, et ainsi nous pouvons étudier les implications de ces interactions pour les individus, mais aussi pour les variables au niveau de la population. L'ingrédient final important pour un ABM est la possibilité de simuler le temps, ce qui est crucial si nous voulons comprendre les dynamiques.

Un ABM est donc une définition d'une population, un ensemble de règles qui définissent les objectifs, les comportements, les réactions et les conséquences des interactions de chaque agent de cette population - avec l'ensemble des règles de chaque agent pouvant être différent de celui des autres - mais aussi des interactions existantes (ou potentielles) entre les agents, et de la séquence temporelle dans laquelle notre population évolue.

Ces types de modèles sont très pratiques pour établir un pont entre le comportement individuel et le comportement agrégé. Cette interaction du niveau micro au niveau macro peut amener les ABM à être construits avec différents objectifs en tête.

Nous pouvons les utiliser soit pour comprendre le monde que nous observons - nos données - en cherchant les comportements et interactions adéquats de nos agents individuels, soit pour les construire afin d'observer le monde qui se manifeste lorsque nous émettons des hypothèses sur les comportements et interactions probables des agents dans un contexte donné.

Imaginez que vous ayez soit une collection d'os, soit l'image d'un corps en mouvement et que vous deviez recréer le squelette fonctionnel de ce corps. Dans la première philosophie, vous choisissez l'image du corps en mouvement et déduisez la place et les interactions correctes des os de manière à correspondre à votre image. Cela reflète l'approche consistant à créer des ABM basés sur des observations empiriques, où chaque "os" ou comportement d'agent est positionné pour correspondre à la dynamique que nous observons dans le monde réel. Dans la seconde, vous choisissez la collection d'os et les arrangez de la manière la plus cohérente sans avoir l'image du corps en mouvement en tête. Cela revient à développer des ABM basés sur des

cadres théoriques, arrangeant les comportements et interactions des agents de manière logique pour explorer les phénomènes émergents potentiels.

Dans les deux méthodes, vous jouez avec les os. Dans la première, pour qu'ils correspondent au corps. Dans la seconde, pour qu'ils s'articulent logiquement. Les deux méthodes sont utiles, et les deux méthodes ont des dangers. Vous pourriez vouloir tellement que vos os correspondent à votre image que vous ne pensez pas assez aux articulations logiques de ceux-ci, et votre corps pourrait ne pas marcher longtemps. Mais avec l'autre philosophie, vous pourriez exclure trop du vrai corps pour le respect de la logique, et le corps que vous créez pourrait marcher, mais il pourrait être hors de ce monde, ne représentant pas le vrai corps.

Les modèles de cette thèse ont été principalement construits en choisissant les os et en les articulant logiquement, visant à trouver un équilibre entre la cohérence théorique et la fidélité empirique, s'efforçant de refléter le familier tout en découvrant des idées qui pourraient ne pas être immédiatement apparentes à partir de l'observation seule, espérant représenter un corps qui nous est familier.

Chapter 1

Experience goods, Reinforcement Learning & social networks

This paper explores decision-making processes in experience goods markets, emphasizing how agents learn to choose in situations where value is uncertain until after consumption. The study examines agents employing individual and social learning strategies within a multi-dimensional reinforcement learning framework, particularly in scenarios of repeated choices. Agents in our model gather insights from personal experiences and through recommendations within their social networks. The simulation results highlight the benefits of combining individual and social learning. While social learning yields consistent outcomes across agents, individual learning offers the possibility of higher rewards but also greater risks.

A notable finding of this research is the development of asymmetrical influence patterns in social networks. This phenomenon refers to a tendency where certain agents become disproportionately influential in guiding others' choices, leading to a centralization in how advice is sought and followed within the network. This aspect of the model sheds light on the nuances of social dynamics in decision-making processes. The study enhances our understanding of consumer behavior in markets for experience goods, providing insights into the complex interplay of individual experiences and social influences in shaping economic decisions.

1.1 Introduction

Economic environments often require agents to make decisions in the face of uncertainty. Choosing optimally in these circumstances poses a significant challenge, and agents may adopt various strategies to cope. In instances where the situation occurs only once, opportunities for learning from experience are minimal. However, in contexts where agents repeatedly encounter uncer-

tainty, it is reasonable to assume that they will engage in a learning process. This process encompasses not only the strategies employed to acquire knowledge — such as consulting peers or using personal experience — but also the progressive refinement of decision quality. As the agent learns, they become more adept at discerning the most effective actions to take.

To make our point clearer, let's consider the example of experience goods. These goods are unique because their intrinsic value cannot be assessed until after they have been consumed. Consequently, when an agent faces the decision to purchase an experience good, there is inherent uncertainty regarding the optimal choice. In a scenario involving repeated purchases, while the consumer may opt for the familiar item, thereby knowing its expected value, the agent remains unaware of potentially better alternatives that may be available. This situation is often referred to as a multi-armed bandit. However, as the details of our model will point out later, the situation described is slightly different in the way agents behave.

Therefore, there is a trade-off between choosing a known option and exploring unknown alternatives that might lead to greater utility than the familiar choice. The assumption of this work is that in such scenarios, individuals are learning not only to select the right option from a pool of alternatives but are also discerning which decision-making method yields the most utility. Consider the example of choosing a doctor. In countries with a universal healthcare system, all doctors charge the same consultation fee; thus, price cannot serve as an indicator of quality. There is no definitive method for making the best choice under these circumstances, and people might employ various approaches: selecting the nearest doctor, choosing the first one that appears in an internet search, or asking acquaintances for recommendations, to name a few. Then, once a doctor is chosen, the patient evaluates the medical consultation based on the utility received. This feedback prompts two questions: Am I satisfied enough to continue seeing this doctor? And do I believe that the method I used to choose the doctor was the most effective? These two questions initiate two learning processes: one about finding the best doctor, and the other about determining the most effective method to find the best doctor.

Three foundational assumptions underpin the model presented: interactions are without prices and without strategic games, and are conducted by individuals with bounded rationality (Simon (1955)).

As discussed, many economic scenarios cannot be resolved through pricing mechanisms alone, particularly when alternatives are priced similarly. This situation is common for many experience goods, such as cultural products (books, cinema tickets, records), education, health-care, or restaurants within the same market segment.

By 'interactions without games,' we mean situations where the agents, in this context consumers, have no incentive to deceive or engage in behavior that would diminish the utility of others. This scenario aligns with many experience goods due to the inherent nature of these

products. An interaction between two consumers typically involves one recommending a product to the other, based on their own experience. Consequently, there is little to no incentive to recommend anything but the best option, as the recommending agent has already consumed the good. Furthermore, in scenarios of repeated consumption, there usually exists a delay before the good is consumed again. Thus, consumers recommend doctors or restaurants because they have already patronized them and do not foresee an immediate need to do so again.

The bounded rationality of our agents is manifested in the way they make decisions and, consequently, how they learn within this model. We propose that agents utilize reinforcement learning in this environment. Reinforcement learning embodies the Law of Effect' (Thorndike (1927), which posits that behaviors yielding favorable outcomes are likely to be repeated in the future, whereas those with detrimental outcomes are avoided. In our scenario, agents begin with no prior knowledge about the available alternatives and gradually learn to identify the optimal behavior (in this case, a consumption choice) over time. They improve by being afforded opportunities to repeat choices, thereby exploring different alternatives. This form of learning has been extensively examined by the behaviorist school of psychology. The significance of reinforcement in non-conscious learning has been established, notably by Thorndike (1927) and Skinner (1965), for both animals and humans. Unlike conscious learning, where individuals are aware of the learning process, reinforcement learning occurs without conscious awareness that one's actions are being influenced by past reinforcements. While humans are capable of various conscious learning mechanisms, these processes are cognitively demanding, and many of our actions are still largely governed by reinforcement learning (Brenner (2006)).

In the field of economics, reinforcement learning has primarily been applied in the context of game theory (Erev and Roth (1998)) and to explore social dilemmas and cooperation (Macy and Flache (2002)). Significant contributions that incorporate this assumption about agents' economic behaviors include the work of Cross et al. (2008), which develops a model where firms are depicted as learners; Arthur (1991), who uses a similar principle to devise economic agents that act like human agents'; and Kirman and Vriend (2000), who models the Marseille fish market with both buyers and sellers engaging in reinforcement learning.

As previously mentioned, our objective extends beyond merely observing the outcomes when agents are endowed with a learning mechanism to select from a range of options. We are equally interested in exploring the consequences of agents learning about their decision-making methods.

To this end, we introduce an alternative method for selection: social learning. Social learning entails situating the agents within a network that captures their social interactions.

More specifically, agents will have the opportunity to select a good based on the recommendations of direct neighbors—essentially, their acquaintances—during each iteration.

Social learning in the literature mainly involves DeGroot learning or Bayesian learning (see Golub and Sadler (2017) and Bikhchandani et al. (2021) for extensive and recent literature reviews). Both methods are different than what we propose here. DeGroot learning involve summing information from plural neighbors, and Bayesian learning updating prior beliefs, which therefore imply having an initial belief to begin with. By contrast in the scenario we study here, we define social learning as following the recommendation of one neighbor at a time.

The influence of social networks on behavior has been a subject of interest in various fields starting from the middle of the 20th century with the work of Katz and Lazarsfeld (1977) exploring the role of influencers on voter's decisions. Since then much progress has been made on finding where these influences take place and on quantifying them. Notably, Christakis and Fowler (2007, 2008) have shown that health behaviors were strongly related to our social networks. The way these behaviours propagate through our social networks has then been the subject of interest of Centola (2010) who highlights the complexity of this diffusion process, often requiring more than one neighbor.

Regarding the ties between consumption and social networks, we can highlight the work on path dependence and increasing returns (Arthur (1989), David (1985), Cowan (1991), Katz and Shapiro (1985)) which explain how the adoption of some technologies can be closely related to the number of consumers already using the technology. This stream of literature directly involves a consumption behavior modified by other buyers, hence implying a network effect. It is nevertheless limited in the way consumers interact and are influenced by others. It does not address the role of individual position or global structure on technology adoption.

Otherwise mainstream economics has generally assumed consumers with fixed tastes uncorrelated to other's consumption. Some authors have however noticed it might not be totally true. Leibenstein (1950) states that some activities are more desirable when shared with others. Smith (1937) and Veblen (1917) observe that wealthy agents enjoy a consumption behavior that makes them recognizable as such and Marshall (1890) points out that agents in general do wish for a consumption behavior that makes them different than others.

We should also notice works made on conformity and fad behavior (See, for example Banerjee (1992), Bernheim (1994)) and on the evolution of patterns of consumption (Granovetter and Soong (1983)). It led Cowan et al. (1997, 2004) to develop a model of demand with interdependence among consumers where we can distinguish three groups of reference for consumers: a peer group, a contrast group, and an aspirational group. The idea that consumers make choices to differentiate or to send signals to others has been confirmed by Gasana (2009) for consumption of clothes and automobiles in the US market and by Burgiel et al. (2017) with Polish consumers. Overall, there is more and more empirical evidence for network effects on consumption (see Bailey et al. (2022), De Giorgi et al. (2020), Agarwal et al. (2021)).

All of these different approaches involve one agent's decisions, behaviour or utility being affected by the actions of others.

Our model will then approach the case of repeated choice in experience markets when agents are allowed to use either their own learning mechanism, which will be modelled as reinforcement learning, or through the use of their social networks. They then also learn, again through a reinforcement mechanism, about which method yields the most utility. There are thus two layers of reinforcement learning.

Therefore, we place agents on a network structure and allow them to interact with their neighbors to make repeated choices in experience goods market. Through repeated choices and learning from experience, agents learn which products are the best to pick. Our model also allows agents to choose between individual and social learning, allowing them to explore the benefits of both approaches and learn about how to learn in general. By combining these two types of learning, our model offers a novel approach to studying consumer decision making in markets with multiple similarly priced options.

The simulations of the model described below give results about whether individual or social learning is preferable and about the naturally emerging proportion of each type of learners. We also explore the role of curiosity and trust for optimal decisions both at individual and population levels.

1.2 The Model

1.2.1 Agents & Products

We treat a fixed population of N agents each of which is repeatedly choosing which of 10 products to consume. Consumption lasts one period so at each period every agent repeats the decision. While the 10 alternatives all answer the same need, they don't all produce the same payoff/amount of satisfaction. Each product is associated with a given payoff, which remains the same for the whole simulation. Payoffs vary on a discrete scale from 1 to 10 and we call products by the level of payoffs they reward (so product 6 gives a payoff of 6 to the agent for example). Tastes are homogeneous so each product delivers the same payoff for everyone in the population and everyone prefers product i to product $i - 1$. Initially, agents have no information on the utility each product delivers; they only know that there are 10 products. Consequently, at the start of the simulation, agents have no way of discerning which product could give the best payoffs.

1.2.2 Interaction

Agents of our population will be able to interact among themselves. By interaction, we mean asking or giving advice on which product to choose for the current stage of the simulation. When giving advice, the agent will simply tell which product it would have chosen for the current round. Agents always answer solicitations from other agents. An agent who receive advice always follows it.

The interaction process takes place on a network structure that we define as follows. We use an undirected graph $G(V, N)$ that represents the ties between individuals. The nature of these ties captures the possibility for any of two connected agents to advise on buying a good or service. So this type of tie can involve friendship, family, or colleagues but it might be harder to exist between an employer and its employee for example.

Here $V = \{1, \dots, n\}$ is the set of agents and $N = \{N_i, i \in S\}$ the correspondence specifying, for each $i \in V$, the neighborhood N_i of i . Formally $N_i = \{j \in V - \{i\} | d(i, j) = 1\}$ and $d(i, j)$ is the length of the shortest path between node i and j . A direct connection implies $d(i, j) = 1$ and is the unique configuration that allows interaction between the two nodes.

To keep a coherent interaction scheme, advice always comes from someone who is, at least for the current round, using its own experience to choose a good. We want to avoid a situation where i asks neighbor j about which doctor to visit, j answers doctor A , but itself asks k and, following k 's advice, visits doctor B . Instead, if i asks j , and j itself asks k , and k is an individual learner, then k 's advice flows back to both j and i . The rare but problematic case where i asks j , who asks k , who itself asks i (creating a cycle) is resolved by forcing one among the three to become an individual learner for the round.

1.2.3 Learning

As agents initially have no information and will have to repeat consumption choices, they need to be able to learn through time. To this end, we implement a version of the Bush & Mosteller Bush and Mosteller (1955) model of reinforcement learning and its specification in Brenner Brenner (2006).

We write $A = \{1, \dots, 10\}$ the set of available products from which agents will have to choose. At any time an agent's opinions are represented by a vector of probabilities upon the set of alternative products. Formally we give a probability $p(a, t, i) (0 \leq p(a, t, i) \leq 1)$ to each alternative $a \in A$ at each time step t for every agent $i \in N$. We write the vector $\mathbf{p}(t, i) (= (p(a, t, i))_{a \in A, i \in N})$ with $\sum_{a \in A} p(a, t, i) = 1$.

The probability vector will evolve through the simulation as the agent experiences different alternatives and revises the probability of being picked of each product. The way it evolves

is through positive reinforcement so payoffs increase the chance of choosing the product again in the future, with increases proportional to the value of payoffs. Formally the change in probability $p(a, t, i)$ to choose the same product as in t at time-step $t + 1$ is given by :

$$p(a, t + 1, i) = p(a, t, i) + \begin{cases} \nu(\Pi(a, t, i)) * (1 - p(a, t, i)) & \text{if } a = a(t) \\ -\nu(\Pi(a, t, i)) * p(a, t, i) & \text{if } a \neq a(t) \end{cases}$$

with $a(t)$ the product selected in t , $\Pi(a, t, i)$ the payoff of product a and ν a parameter that controls how fast agents respond to new experience. Therefore, the idea is that each product chosen by the agent has a positive impact on the probability to choose it again¹, and since the sum of all probabilities is equal to 1, it also hurts the probability of choosing every other product. Information about past choices are only contained in the probability vector which is both the only element of “memory” and the only element used to make decisions¹.

We use the model with $\nu(\Pi) = \nu * \Pi$ and only positive values of Π . Alternatives with low payoffs will only marginally modify the probability vector in contrast to high payoffs alternatives which can greatly modify the vector. The level of impact of each alternative depends on the level of the ν parameter. ν is a critical parameter in this model as it controls the strength of reinforcement at each time step. Very high levels of ν will imply high modification in the probability vector and choosing an alternative only a few times can produce a lock-in effect to this alternative. By contrast very low levels of ν will never allow the agent to learn. It should also be noticed that in order to keep probabilities between 0 and 1, ν must not be set higher than $1/(\text{maximum payoff})$.

For example if we have four alternatives, each equally considered by the agent i at $t = 0$. Each rewarding 10, and $\nu = 0.08$. The reinforcement for any of the four options if picked after the first round will be of : $p(a, t + 1, i) = 0.25 + 0.08 * 10 * (1 - 0.25) = 0.85$ and every other options that was not picked : $p(a, t + 1, i) = 0.25 - 0.08 * 10 * 0.25 = 0.05$. In only one time step, the chance to pick an other alternative than the firstly picked is already very small. If picked a second time, the probability would increase to 0.97. Hence we can see that lock-in can happen quite fast if the level of ν is too high. At the other end, setting ν to a very low level has the opposite effect : the probability vector is only slightly impacted when an alternative is picked, and the learning process either doesn't happen or happens too slowly for agents to

¹We know that negative reinforcement could exist in a consumption perspective. In this model, we assume that all alternatives at least minimally answer the need motivating the buying act from the consumer. One way to understand it is that even if you're buying something and you're not totally satisfied, you still don't know what the other alternatives could give you. So the probability to choose it again increase because you're now familiar with this product compared to the others. We did test the model with negative reinforcement learning, and it only slightly slowed down the process that we describe later.

¹The learning mechanism is therefore different than in the multi-armed bandit where the prior of each product is only affected when the given product is selected and experienced. With reinforcement learning, each time you pick an alternative, the probability to choose again this good is modified, but so as the probabilities to choose every other goods.

benefit. Figure 1 an illustration of the learning process for both very high and very low levels of ν . When its too low (in blue), agents keep moving from one alternative to an other. By contrast when its too high (in red), the agent settles on one option too fast to learn about better options. These opposite effects are both cases we want to avoid in our model.




Figure 1.1: evolution of product choice for one agent other two runs. A run with a very low level of ν (in blue), and one with a very high level of ν (in red).

1.2.4 Dynamics & Learning methods

At first, agents are placed on a network structure that defines their neighborhoods. Each agent must choose a product at each round. To do so, in the general model, they can either use their product choice probability vector and so rely on their own experience (which will be limited in the early phase of the simulation) or ask advice at random from one of their neighbors. When they receive advice, they follow it blindly to choose.

Once everyone has picked a product, they get the associated payoff and can learn and gain experience following the reinforcement learning process described above. They learn about products, (and slightly or strongly increase the chance of choosing them again in the future depending on the level of payoff) but they also learn about their social behaviors.

The idea is that because agents have chosen either to rely on their experience or to ask a neighbor, they can now learn about which behavior is the most rewarding. If each time I took a decision on my own I was badly rewarded, I might start to ask my neighborhood about what they would pick. If my payoffs then increase, I will reinforce the pro-social behavior instead of

remaining an individual learner.

We then have a set $B_i = \{\text{individual, network}\}$ of possible social behaviors to adopt for agent i . As for the products, we give a probability $p(b, t, i)$ ($0 \leq p(b, t, i) \leq 1$) to each behavior $b \in B_i$ at each time step t . We also write the vector $\mathbf{p}(t, i) = (p(b, t, i))_{b \in B, i \in N}$ with $\sum_{b \in B} p(b, t, i) = 1$. Finally, the change in probability $p(b, t, i)$ to pick the same behavior as in t in $t + 1$ is given by :

$$p(b, t + 1, i) = p(b, t, i) + \begin{cases} \nu(\Pi(b, t, i)) * (1 - p(b, t, i)) & \text{if } b = b(t, i) \\ -\nu(\Pi(b, t, i)) * p(b, t, i) & \text{if } b \neq b(t, i) \end{cases}$$

With $\Pi(b, t, i)$ the reward associated with the social behavior adopted in t , which is equal to the payoff of the product picked at that round: $\Pi(a, t, i)$. As for the product reinforcement process, $\nu(\Pi(b, t, i)) = \nu * \Pi(b, t, i)$ ¹.

Once all agents have updated both product-choice and social behavior vectors, a new time-step starts and the model runs until the desired number of rounds.

1.3 Results

1.3.1 The role of Interaction

The initial outcome we examine is the contrast between two scenarios: one where agents can take their decisions based on social or individual learning and another where they operate in isolation, basing decisions solely on individual experience.

In the interactive scenario, agents are embedded within a stochastic network structured according to the Erdős-Rényi model Erdős et al. (1960) with a connectivity density of 10 percent, meaning each agent has a 10 percent chance of being connected to any other agent² For both the interactive and non-interactive models, our simulations incorporate a total of 500 agents. The learning parameter ν is consistently set at 0.02 for both models corresponding to 20% of its maximal possible strength.

¹ ν is not necessarily the same in the two learning processes. We could imagine setting a high level of ν in the product choice process but a low level of ν for the social learning process for example.

²We only display results for random network but changing the structure to ones with fat-tail degree distributions, or to ones with high cliquishness and low average path length has no effect on the dynamics observed.




Figure 1.2: Left: Distribution of product choices at last round (150) over 10 simulations of the model. Right: Evolution of mean aggregate payoffs for each round, over 10 simulations. Yellow represent results for the model with social and individual learning and red the model with only individual learning.

In the comparative analysis, the model permitting both social and individual learning exhibits a marked advantage, as a substantial proportion of the population learns to select the highest-reward options, outperforming the model that restricts agents to individual learning. Dynamically, both models display analogous trajectories: a rapid and effective learning phase initially, which then tapers off to a slower learning rate. Despite these similarities, the interaction-enabled model achieves higher levels of aggregate payoffs from the outset, specifically within the first 20 rounds. This advantage is evident in the final round, where over $\frac{3}{5}$ of the agents successfully identify the optimal choice, compared to approximately $\frac{2}{5}$ in the model without interaction. Consequently, a larger fraction of the population, when interaction is not present, settles for less optimal selections.

In the model without interaction, agents initially make random product selections due to a lack of information to guide their choices. In the early stages, they explore different options and can quickly dismiss the five least rewarding products because the reinforcement received is insufficient for any one of them to become a dominant choice. However, for the five most rewarding options, the reinforcement is significant enough that some agents might start to consistently choose less optimal products like the fourth or fifth most rewarding ones, due to stochastic repetition. This phenomenon, known as the ‘lock-in effect,’ varies in its onset with the level of the ν parameter. With $\nu = 0.02$, the model balances exploration and exploitation: agents explore at the start but are still likely to settle on a particular choice eventually. Consequently, a significant number of agents fail to identify the most rewarding product, settling

instead for sub-optimal choices.

The interactive framework markedly promotes the spread of recommendations for the best products, allowing a significant number of agents to circumvent lock-in on sub-optimal goods. The reason recommendations for the best products propagate first warrants a closer look. For instance, if agent i fortuitously selects A , the top doctor in town, on their first try, the rewarding experience increases the likelihood that A will be i 's recommendation to agent j in the next round. This is because A 's high reward has substantially raised its standing in i 's probability vector, indicating a recognition of A 's quality. In contrast, if i 's initial choice had been B , an average doctor, the slight increase in B 's vector probability is reflective of an unremarkable experience, making it less likely for i to recommend B again.

This creates a contamination effect,⁷ where exposure to the best goods substantially increases their promotion over others. As a result, more agents are introduced to these high-quality options, reducing the chances of settling for less. Thus, the outcomes at an aggregate level are considerably better than in models without interaction.

The reinforcement mechanism that operates within the recommendation system is crucial for enabling agents to convey information about high-quality goods. When an agent makes a recommendation to a neighbor, the suggestion is not based on their most recent selection or even the one they have chosen most frequently in the past n rounds. Instead, the recommendation reflects the choice they would likely make at that moment, which is determined by a probability vector. Thus when an agent i is asked to recommend a good to j , it draws a good using the probabilities in its vector and gives the result to j . This vector is continually adjusted through reinforcement, ensuring that it represents the agent's up-to-date preferences. It is this dynamic aspect of the recommendation system that facilitates the spread of the most advantageous alternatives.

1.3.2 Social and Individual Learners

Once we've established that the social interaction model was performing better than the no-interaction one, we wish to understand the effects of being either a social or individual learner within the interaction model. Mainly we want to understand if one behavior is more rewarding than the other, and whether one behavior gains more popularity than the other through the simulation. To address this issue we stick with a random network structure, inside which 500 agents are allowed to communicate among their neighborhoods. The ν parameter is still at a 20% level of its maximum potential.

Regarding the dynamics of the two behaviors, the population is split in two at the start, with half starting as individual learners and half as social learners. This is logical as agents don't have any information at the beginning and choose randomly between the two. There is

then a bit of variation within the first time-steps but only to, on average (for a few hundreds of simulations), return to a 50-50% split of the population between the two types of learning. The minimum number of social learners in the final round we found was 230 and a maximum of 270. On average, agents switch learning methods 15 times during a simulation before locking-in one. The mean standard deviation is around 10, indicating a fairly high heterogeneity in switching behaviors from the agents.




Figure 1.3: Aggregate individual payoffs for the total simulation, 150 rounds in each simulation. the greatest achievable payoff being choosing product 10 at every round, yielding an aggregate payoff of 1500. In red the density for individual learners, in yellow the one for social learners ³. The plot display the distributions for one representative simulation of the model

There's a more interesting fact when we look at the distribution of individual aggregate payoffs over a whole simulation for both types of learners.¹ As Figure 3 shows, the two distributions are quite different and this is confirmed with a Kolmogorov-Smirnov test ($D = 0.496, p < 2.2e - 16$).

What we see is that potentially, being an individual learner can permit an agent to reach higher aggregate payoffs, but also lower ones than with the social learning method. For individual learners, this behavior is easily explained by the lock-in effect of the reinforcement

¹We define social and individual learners with the behavior they picked in the last round. As people lock-in in one of both early in the simulation, this is equivalent to using the behavior they have chosen the most over time.

mechanism. If you're able to find product 10 early in the simulation and lock in to it, you will be able to reach almost the highest possible level of payoffs during the simulation. But if you lock in to least-rewarding alternatives, your payoffs will be harmed for the rest of the simulation. By contrast, social learners are all reaching very similar levels of payoffs, but not as high as the best performing individual learners. This can be explained by the process of diffusion of the best alternatives that take a bit of time to propagate to all social learners. Meanwhile, social learners are sometimes being recommended low-rewarding goods.

It explains why, while the two sub-populations reach equivalent average aggregate payoffs (around 1340 for both), individual learning can be seen as a riskier behavior than social learning from an individual point of view.

1.3.3 Curiosity

We have so far assumed a certain level of 'curiosity' in our agents. The issue they face—navigating a pool of alternatives with unknown values—necessarily involves a trade-off between exploration and exploitation. This means agents must balance the search for better options against the choice to settle for an alternative that appears to be satisfactory.

In this framework, 'curiosity' refers to the duration dedicated to exploring unfamiliar alternatives as opposed to exploiting those already known. This includes selecting previously untried goods and reconsidering goods that didn't initially seem optimal, even after encountering better options.

To model varying degrees of curiosity, we manipulate the ν parameter. As detailed in the model description, the magnitude of ν influences the reinforcement strength, thereby controlling the exploration phase's length before an agent commits to a specific good.

Thus, a low ν value equates to weaker reinforcement, prompting agents to explore longer before settling on a particular good. Conversely, a high ν value results in stronger reinforcement, leading to a quicker transition from exploration to exploitation of a chosen alternative.

Considering the reinforcement learning algorithm and the valuation of goods which span from 1 to 10, the upper limit for the parameter ν is set at 0.1. This represents the maximum reinforcement strength. On the opposite end of the spectrum, the minimum level of curiosity is equivalent to an absence of exploration, which would correspond to $\nu = 0$. In such a case, agents would make their choices randomly in every round, essentially eliminating the learning component of the model.




Figure 1.4: Population Payoff levels at last round when we vary the level of ν . Maximum possible being 5000. We do it for both the interaction and no-interaction model. $N = 500$, $R = 100$.

The results show the clear benefits of learning and a certain level of curiosity. When we move from 0 to 0.005, there is a steep increase in payoffs, highlighting the importance of learning over random choices. The peak value of ν is relatively low, at around 0.015 for the model without interaction and 0.025 for the model with interaction. Both models exhibit similar reactions to variations of the strength parameter. Once the peak is reached, there is a gradual decline in population-level payoffs in the last round as the level of ν increases.

There is no level at which the model without interaction outperforms the interaction model.

This graph also suggests the hypothesis that the interaction model is not simply a means to ‘open’ social learners to new alternatives through interactions, which would imply a kind of forced’ curiosity that distinguishes it from the no-interaction model. Rather, it involves the elicitation of the best alternatives that propagate through the network of agents, enabled by the reinforcement mechanism.

If interaction were merely adding a bonus level of curiosity to agents, then at the optimum level of curiosity in this model, both interaction and no-interaction models should roughly reach the same amount of payoffs, but at different levels of ν . Instead, the peak performance

reached by the interaction model is never attained by the no-interaction one, clearly indicating a different mechanism at play.

1.3.4 Trust-based social interactions

So far we've only considered unweighted network interactions. Thus, a social learner will ask any of its neighbor with an equal probability. Since our model makes the assumption of learning agents, both about which alternative to pick and which method of learning to use, it would be natural to also enable them to reinforce the probability to ask neighbors that have been giving good advice in the past.

Introducing this third layer of reinforcement will raises two primary questions: Will it allow social learners to perform better? And how will the network of interactions evolve once it's established that some connections are more utilized than others?

The method by which we enable agents to identify which neighbors to consult is consistent with the approach taken for the other two layers of reinforcement learning. Specifically, agent i increases its probability of consulting neighbor j based on the reward for that action, denoted by the value of the good (ranging from 1 to 10). As a result, poor advice has minimal impact on the probability vector representing the likelihood of choosing any direct neighbor. In contrast, good advice significantly boosts the probability of selecting j in subsequent rounds.

While we don't allow for agents to create or delete ties, this mechanism modifies the initial network by leading it to become a directed and weighted one. In this network, the weight of each connection can vary between 0 and 1, with these values indicating either a nonexistent or maximum probability of the tie being activated, from the originating agent to the receiving one.

This approach to curating social connections for guidance in uninformed scenarios can be conceptualized as a form of trust formation, grounded in experience and feedback. As previously discussed, the reinforcement learning equation from Bush & Mosteller has the potential to result in a lock-in to a single alternative. In this context, such a process can be interpreted as trust formation. Once an agent solidifies its choice on a neighbor j , it will unconditionally follow j 's recommendations, even if j alters the advice it provides to i .

Performance

We first evaluate the performance of this model regarding aggregate payoffs at last rounds, but also from an individual perspective, mainly asking if social learners, now able to choose among their neighbors, perform better than before.




Figure 1.5: At each round we plot the aggregate payoffs of the population for a representative run. We show the evolution for 150 rounds in 3 different models: When there is no social learning (in yellow), when there is social learning without weighting neighbors' ties (in red) and when we include the reinforcement of neighbors mechanism (in blue). $\nu = 0.02$, $N = 500$.

Figure 5 gives the aggregate payoffs at each round, for 150 rounds, for 3 different models. We mainly want to compare a model with (in blue) and without (in red) the third layer of reinforcement. We don't see a remarkable improvement from the simpler model both regarding the payoffs levels or dynamics of learning inside populations.

What's more interesting is the evolution of density from the initial model (with 2 layers of reinforcement) and the one with trust (3 layers).




Figure 1.6: We show the density of individual aggregate payoffs for 150 rounds and differentiate between agents mostly choosing to be individual learners or social learners. We compare the initial model (in dashed lines) with the one with trust (solid lines). $\nu = 0.02$, $N = 500$.

Comparison between the two models shows the distribution for social learners moves to the right, implying an improved average total payoff for a social learner. The mean total payoff of a social learner increases by around 50 points. There would therefore be an utility for a social learner to choose a neighbor based on a reinforcement mechanism instead of choosing at random. There is however an increased dispersion of total individual payoffs for social learners, reflecting the same process as for individual learners: reinforcement learning can lead to lock-in on sub-optimal options, here on sub-optimal neighbors.

Regarding the proportions of each type of learners, there is only a slight increase in the proportion of social learners from the 2-layer to 3-layer model. The mean number of social

learners over 10 simulations for each model only shows a 1% increase.

Network structure

By allowing agents to modify the probability to choose a neighbor over an other, the network evolves to become weighted and directed. In the initial model, the network is random which imply that the distribution of degree is normally distributed, far from real social network evidences.

We want to understand the in-degree distribution difference with the initial model. Mainly we're interested at whether some agents become more central than others, which would imply that they are more solicited than the rest of the population to give advice. Or, if the distribution remains normal.




Figure 1.7: We show the density of in-degree centrality for both model with two and three layers of reinforcement. $\nu = 0.02$, $N = 500$.

We define in-degree centrality of an agent i as the sum of the weights that each of its neighbors j give to the tie (from j to i). Say that agent i has n neighbors. The maximum value for in-degree is n , which would imply that each neighbors always ask advice from i rather than their other connections. At the other extreme if the in-degree is equal to 0, all neighbors have made the choice to never ask i about which product to choose. Results from figure 3 show a very different distribution for the model with and without reinforcement of neighbors. While the distribution is symmetric for the two layer model, adding neighbors' reinforcement leads to a fat right tail distribution. There are therefore a few agents that will be highly seek for by their neighbors. We don't find correlation between having a higher number of neighbors on the

initial random network and being central in the final one. It seems that there is a correlation between being successful and being highly central.

This might tell something interesting about the emergence of influence and opinion leaders in social networks. Here there is no preferential attachment mechanism since agents don't know the weights other agents use.




Figure 1.8: Scatter plot for individual aggregate payoffs over 150 rounds and in-degree centrality at round 150. Corr 0.28, p-value < 0.01. $\nu = 0.02$, $N = 500$.

1.4 Conclusion

Modeling agents' behaviors when they are learning to choose among multiple alternatives without prior information, and through repeated choice sequences, can be challenging. This is because individuals might employ different methods, none of which can be definitively assessed as the optimal one by observers.

We posit that in the context of consuming experiential goods, much of the learning is unconscious, implying that agents are not necessarily building elaborate strategies to choose. Furthermore, the methods they employ for learning are also subject to unconscious adaptation.

This has led us to model consumer behavior as a multi-dimensional reinforcement learning

process. In this framework, agents are not only learning to select the right products over time but are also refining their decision-making methods to yield better outcomes.

We discovered that a model facilitating both social and individual learning significantly outperforms a model limited to individual learning at the population level. Consequently, there is a discernible positive impact from social learning, even in the absence of superior prior knowledge among the population members.

The advantage of social learning appears to stem from the reinforcement mechanism inherent in the recommendation system. When an agent seeks advice on which good to select, the response is informed by the recommending agent's own probability vector, which effectively disseminates information about options with higher utility.

Nonetheless, there are advantages to individual learning as well. Fortunate agents may rapidly converge on choices that yield higher rewards, while less fortunate ones may become fixated on less rewarding options.

Therefore, individual learning presents a higher level of risk compared to social learning, which does not allow agents to achieve perfect scores but ensures that all social learners consistently perform at a high level with more uniform payoffs.

We also demonstrate the critical role of optimal curiosity levels, which enables agents to sufficiently explore different alternatives before settling on the most suitable one for exploitation.

Finally, by integrating a third dimension of reinforcement at the level of neighbor selection, we simulate a form of trust development wherein social learners hone their ability to identify the most reliable network peers for guidance. Incorporating this trust element allows social learners to enhance their overall performance. Additionally, we observe interesting emergent patterns within the social network's structure. Predominantly, certain agents become pivotal, evidenced by their disproportionate frequency in being consulted for advice, thus indicating a form of centralization in the advice-seeking process. It thus creates fat-tail degree distributions which are commonly found in networks of all kinds Barabási and Albert (1999).

There is therefore great benefits from learning through social interactions, and so even when no one is better informed than the others at the beginning.

This work has concentrated on the conventional methods by which agents select goods, yet the advent of the internet has transformed the ways in which individuals access information for decision-making in markets for experience goods. This transformation is especially pronounced with the emergence of online social networks and rating systems, which aggregate and disseminate information widely.

Future research could, therefore, extend the model to address these modern dynamics to more accurately capture the nuances of learning and decision-making in contemporary experience good markets.

Chapter 2

Internet and experience goods: The role of Influencers and rating-scale systems

With the advent of the Internet, consumers of experience goods have gained new methods for determining which product to choose. Primarily, we identify Influencers and Review Systems as predominant in this new landscape. We investigate whether these new methods of choice can enhance consumer utility, both at the individual and population levels. To this end, we constructed several models and analyzed them using computer simulations. Notably, we find that influencers negatively impact users' utility, especially when they exert their influence within scale-free network structures and promote average products rather than the lowest-quality options. As for review systems, our findings suggest that they are significantly less effective than learning from private information, particularly when tastes are heterogeneous, which is often the case for many experience goods.

2.1 Introduction

In 2023, France enacted a landmark law to regulate social media influencers, effective from June 1 (Mondaq (2023); Foley (2023); Euronews (2023); Enigma (2023)). This first-of-its-kind legislation in Europe aims to protect younger consumers from the promotion of harmful products and trends. It includes prohibitions on advertising cosmetic surgery and sports betting, and mandates clear labeling of altered images and videos.

These kind of laws are needed because online influencers do influence their followers, and

without regulations, it is a double-edge sword. While they can have positive impacts when promoting healthy behaviors like getting vaccinated for the flu (Bonnevie et al. (2020)), there is much more concern about influencers promoting dangerous practices such as gambling, as the new french law now prohibits.

Javed et al. (2022) research demonstrated the significant impact of fashion influencers on Instagram in guiding consumer decisions and extending content reach. Megane et al. (2019), found that influencers mold consumer purchasing intentions through their endorsements, though their authenticity is sometimes questioned by followers. Wang et al. (2022) study delved into the effects of influencer-led live streaming marketing, revealing that factors like influencers' expertise, negotiation skills, and after-sales services influence consumer trust and impulsiveness, ultimately swaying purchase decisions. Lu and Seah (2018) underscored the importance of social media influencers in managing consumer online engagement, noting that different influencer types elicit varying engagement levels based on social proximity and product traits. Collectively, these studies underscore the pivotal role of online influencers in directing consumer behaviors and shaping consumption trends.

Influential figures have always played a significant role in societies, but the advent of the internet and online social networks has dramatically amplified their reach and influence. This modern landscape offers more individuals the opportunity to exert influence over larger audiences and for extended periods.

When considering how consumers choose experience goods, it's essential to acknowledge influencers' role, as research indicates that consumers often rely on them for making decisions. Our study aims to explore the dynamics of this influence in a population that can either independently research their choices or rely on their social networks, which are increasingly populated with influencers—these highly connected nodes in the network. Importantly, influencers' impact extends beyond mere network connectivity, as their motivations often differ from those of the general population.

In particular, influencers may use their platforms to promote products, sometimes incentivized by financial gains, potentially leading them to endorse lower-quality alternatives. Our model aims to demonstrate how, within both simple and complex network structures, individuals are adversely affected by the promotion of subpar products. Intriguingly, the model suggests a heightened vulnerability to average-quality products, as consumers can more easily identify very low-quality options but may struggle to distinguish among average ones, which they are more likely to adopt.

This research also endeavors to unravel the complexities surrounding the review systems employed by numerous websites and platforms. These systems, essentially aggregating consumer reviews, are intended to guide users through uncertainty, particularly when selecting experience

goods.

While these review systems initially appear highly effective in mitigating uncertainty through the accumulation of past customer experiences, they are not without their drawbacks. A significant concern is the presence of fake reviews, which have been shown to harm consumer welfare (He et al. (2022); Akesson et al. (2023)). Moreover, even in the absence of fake reviews, these systems face challenges when addressing a population with diverse tastes.

Therefore, the objective of this chapter is to comprehend the significance of the public information that has become accessible recently with the advent of the internet. While Chapter One explored the private options available to individuals for making choices among a set of unknown alternatives, this chapter identifies and analyzes the role of two phenomena: Influencers and Review Systems.

While influencers are proactive actors online, driven by economic self-interest, review systems are supposedly a public asset shaped by the truthfulness of its users.

Our findings indicate that review systems are highly effective in populations with uniform tastes. However, this efficiency diminishes when considering a population with varied preferences, where individual learning becomes a more effective strategy.

A particularly intriguing aspect of our study is the exploration of social learning. Our simulations reveal that social learning is effective in populations with heterogeneous tastes, irrespective of their distribution. This effectiveness holds true in networks structured around the homophily principle (McPherson et al. (2001),) as well as in networks that do not exhibit homophilous tie formation.

We begin our analysis by outlining the baseline model, which closely aligns with the framework established in Chapter 1. This will set the foundation for our exploration. Subsequently, we delve into the construction of various network structures employed in our study, detailing the methodologies and rationale behind each. Following this, we present and discuss our key findings, focusing on two main aspects: firstly, the impact of influencers within these networks and how their presence shapes consumer decisions; and secondly, we critically examine the limitations and effectiveness of review systems in guiding consumer choices, particularly in the context of diverse consumer preferences and behaviors.

2.2 Model

2.2.1 Reinforcement Learning

The basic model, as established in Chapter 1, introduces a population connected through a network structure. This population faces a repeated decision-making process, choosing among

10 distinct products, labeled from Product 1 to Product 10. Each product provides a utility value equivalent to its numerical label.

At each decision round, each participant must select from these ten options.

In this model, the number of products and their respective values remain unchanged throughout the duration of the simulation. The primary interest in this model arises from the fact that the products are designed as experience goods. This means that no agent in the population can ascertain the value of a product prior to consuming it.

Therefore, agents with no prior information will must learn both which is the best product to pick, but also about which method is the most useful to learn.

Agents have two primary decision-making strategies available to them. The first strategy is social learning, where an agent randomly selects a direct neighbor within the network and emulates their choice at the current time step.

The second strategy involves individual learning, where agents make choices based on a reinforcement learning mechanism.

Reinforcement learning operates on the principle of strengthening the likelihood of choosing alternatives that have previously yielded higher payoffs. This is mathematically represented by a probability vector, encompassing the probabilities of selecting each of the ten alternatives, with the vector sum always equaling one. At each time step t , the selected alternative provides a reward, which then informs adjustments to the probability vector for future selections. The probability vector is dynamically updated at each step following the equation:

the change in probability $p(a, t, i)$ to choose the same product as in t at time-step $t + 1$ is given by :

$$p(a, t + 1, i) = p(a, t, i) + \begin{cases} \nu(\Pi(a, t, i)) * (1 - p(a, t, i)) & \text{if } a = a(t) \\ -\nu(\Pi(a, t, i)) * p(a, t, i) & \text{if } a \neq a(t) \end{cases}$$

with $a(t)$ the product selected in t , $\Pi(a, t, i)$ the payoff of product a and ν a parameter that controls how fast agents respond to new experience.¹

Agents are characterized as either social learners or individual learners. To determine their learning strategy at each round, we incorporate a second layer of reinforcement learning. This layer dynamically adjusts the probabilities of an agent adopting either a social or individual learning strategy in any given round, based on the payoffs received from previous choices.

The payoffs from chosen alternatives directly influence these probabilities. For instance, if an agent, employing social learning, selects a product with a high payoff (such as product 10), this will significantly reinforce their tendency towards social learning in future rounds, consequently

¹See chapter 1 for more details

reducing their inclination towards individual learning. Conversely, a lower payoff (e.g., a payoff of 1) will result in only a minor adjustment to the probability vector, subtly altering the likelihood of choosing between social and individual learning strategies in subsequent rounds.

Therefore at each round, each agent has two decisions to make: being a social or individual learner, which is decided through a reinforcement learning process, and once it is decided, what product to choose.

A social learner will simply mimic the pick of a neighbor taken at random.

An individual learner will pick a product using its probability vector associated to the pool of products, which is evolving through a reinforcement learning process.

finally, the probability vectors associated with the learning method and the pool of products are updated according to the learning process.

2.2.2 Network Models

We use two models of network structures that well characterize many social interactions processes. One is the the scale-free model Barabási and Albert (1999) and the other the small-world model Watts and Strogatz (1998).

The Barabási-Albert (BA) model, is a method for generating scale-free networks. It begins with a small initial network of m_0 nodes. New nodes are then added one at a time, each creating $m \leq m_0$ edges to existing nodes. The key aspect of the model is the "preferential attachment," where the probability $P(i)$ that a new node will connect to an existing node i increases with the degree k_i of node i , given by $P(i) = \frac{k_i}{\sum_j k_j}$. Here, the sum in the denominator is over all existing nodes j . This process results in a network where a few nodes become hubs with high degrees, while most have a low degree. The degree distribution of the network, $P(k)$, follows a power law, specifically $P(k) \propto k^{-\gamma}$ with γ typically in the range of 2 to 3, consistent with many real-world networks.

The Watts-Strogatz model is a process for generating networks that exhibit small-world properties, namely high clustering and short average path lengths. The model starts by creating a regular lattice, a ring of N nodes where each node is connected to k nearest neighbors (k is an even number). Then, each edge in the network is rewired with a probability p . During this rewiring process, one end of the edge is fixed, while the other end is reconnected to a node chosen uniformly at random over the entire network, avoiding self-loops and duplicate edges.

For $p = 0$, the network is a regular lattice, and for $p = 1$, it is a random network. The interesting behavior occurs for intermediate values of p , where the network exhibits both high clustering (like a regular lattice) and short average path lengths (like a random graph). This combination of properties is characteristic of many real-world networks and is what defines a small-world network.

Finally, for testing our model as designed in chapter 1 and compare it to review system, our population is placed on a random network following Erdos-Renyi model Erdős et al. (1960)

The Erdős-Rényi model, denoted as $G(N, p)$, is a simple and widely used model for generating random graphs. This model starts with N isolated vertices and adds edges between them randomly. Each possible edge between a pair of vertices is included in the graph with a probability p , independent of other edges.

This process results in a graph where the presence of each edge is a Bernoulli trial with probability p . The main parameter of interest, p , determines the graph's density. For low values of p , the resulting graph is likely to be sparse with many isolated vertices, while for high values of p , the graph becomes denser and more connected.

2.3 Results

While Chapter One delved into the dynamics of private learning, encompassing both individual and social learning, it's crucial to recognize that a significant portion of our learning also transpires through public channels.

Prior to the advent of the internet, public learning was predominantly facilitated by advertisements in various media. In the current digital age, new methodologies for product selection have emerged, notably through Influencers and Review systems. Influencers, active on social media platforms, are compensated by companies to promote products to their extensive follower base, thereby providing these products with significant exposure. On the other hand, review systems serve as collaborative platforms that aggregate user ratings and feedback for products post-consumption, effectively addressing the challenges posed by experience goods.

In today's market, consumers frequently rely on both channels to make informed decisions about experience goods. Influencers cover a wide range of cultural products, including movies, TV shows, books, restaurants, and travel destinations, most of which are also subject to review systems.

Thus, consumers are presented with an array of methods for choosing products, and the forthcoming simulations aim to discern the most effective among these.

We begin by examining the impact of influencers within a consumer population learning to select the optimal product from a variety of alternatives serving the same purpose.

Specifically, we investigate the influence of influencers on consumer utility, considering both the nature of the products they promote and the network structure through which their influence is disseminated.

2.3.1 Influencers & Star Networks

The phenomena induced by the internet we look at is the emergence of influencers on online social networks. Influencers are highly connected people on online social networks, where they can share a large variety of content which make them rise on these networks. They gather followers and as they grow, brands might solicit them to advertise their products to their communities as they can reach a lot of people at the same time with short publication.

Thus, influencers can lead their followers to sub-optimal alternatives because they can be paid to do so by companies. It will be up to the star to decide whether a product is worth advertising or not.

We analyze this situation with our model by setting up a star network structure, also called a hub and spoke, an extreme situation of preferential attachment. In this model, there is one star agent, the influencer, and the rest of the population which are followers. All followers are only connected to the star agent, so they have only one neighbor. In our scenario, the star is neither using the network nor its own experience. Rather, it is forced to choose one of the ten alternatives. This asymmetric relationship would translate into a directed network. Followers, as in the chosen-interaction framework, have the choice between relying on their experience or seeking advice from the influencer.




Figure 2.1: Left: Evolution of Aggregate Payoffs at the last round, for each product advertised. Right: Evolution of the number of individual learners at last round, for each product advertised.

Above two graphs shows the relationship between payoffs and users of the "personal experience" behavior and the star agent moving from product 1 to product 10. When the star is advertising the lowest-rewarding alternatives, the population is "robust" and everyone chooses to rely on their own experience instead of following the star agent. However, as the star is advertising products in the highest-rewarding range, more and more people start to listen to

him, and so to choose sub-optimal products. When the star is advertising the highest-rewarding products, around $3/5$ of the population have chosen to rely on its choice and so are able to make high rewards.

Interestingly, the worst-case scenario for the population is not when the influencer advertises the lowest-rewarding range of products, but rather when it is advertising average-rewarding ones. In this situation, many agents (around $1/2$ of the population) are following its advice, thus choosing a sub-optimal alternative thereafter.

The pattern we observe in the figures can be explained as follows: When the star advertises products 1 to 3, almost everyone in the population tries following influencer's advice, and realizes that individual learning generates better payoffs. Indeed an individual choice during the first rounds has a payoff probability of around 5 while following the influencer leads to inferior payoffs.

The situation gets trickier when the star advertises products near 5 since our agents will not easily realize that they could get better payoffs through individual learning. In the early phase of the simulation, the alternative you draw from individual learning will be crucial to determine whether you will follow the influencer or not. This is why around half the population follow the star while they could do at least as much with individual learning.

Finally, the situation gets better for our population when the influencer advertises products of the high-rewarding range. In this case, $2/3$ of the population follow the influencer, and it leads to higher aggregate payoffs at the last round than when everyone is an individual learner.

2.3.2 Influencers & Realistic Network Structures

We finally want to test the impact of Influencers on more realistic network structures. To do so we will use Scale-Free (SF) and Small-World (SW) network structures. As defined by Watts & Strogatz (1998), SW are useful to depict real-life observed phenomena such as the six degrees of separation experiment conducted by Milgram (1967). The idea of SW is to combine high clustering: the people I'm connected to are also mainly connected between themselves, and short average path length: every agent of the network is only a few connections away, so potentially easy to reach. Also a good approximation of some real-life structures, SF structures are defined by Barabasi & Albert (1999) as graphs where the degree distribution follows a power law, which implies that while the vast majority of nodes on the network have a similar, low number of connections, a few have a lot more links. These kinds of structures are generally related to "Mathew effect" kind of phenomena.

To test the impact of Influencers advertising goods, we use a model with 500 regular agents and 50 influencers, advertising to them. At first, we compare the aggregate payoffs in the last time step when there are influencers and when there are not. We do this both for SF

and SW graphs. Hence we will compare, for each good being advertised, the results with and without influencers (and so without advertising when there's no influencer). We also make sure the results of influencers are not interfering in the analysis. Hence in the no-star model, we create 500 agents while in the star model, we create 550 but then subtract the payoffs of these influencers from the aggregate payoff in the last round. This leaves us with both models showing payoffs for 500 regular agents.

To define influencers in both network structures, we use betweenness centrality as an indicator of central individuals. In the SF, it roughly coincides to degree centrality. In the SW graph, where everyone has a similar number of neighbors, betweenness centrality describes people connecting remote places of the network.

The graph below shows results as each good is advertised (x-axis). On the y-axis, results are expressed as a ratio from the star model to the no-star model (for the same network structure). A value of 1 then indicates that both models are performing equally for the given advertised good and network structure considered. Values below 1 indicate that the star model performs less than the no-star one, and values above 1 that the star model over-perform the no-star one. For each situation considered, we run 10 simulations of the model and show the mean over these 10 simulations. In each simulation, we run 100 time steps.

In a small world, influencers have little impact. They reduce payoffs slightly, relative to a world with no influencers, and they only improve payoffs when they are advertising the really top goods. By contrast, in a scale free world, agents greatly suffer from the presence of influencers. As shown earlier, critical situations happen when the goods being advertised are average-rewarding. This is when a lot of agents are deceived and start being social learners when they should remain individual ones. It seems that being an influencer on an SF structure (like most of our online social networks...) entails way more power than in an SW one, at least for the diffusion of recommendations.




Figure 2.3: Evolution of the number of individual learners as the product being advertised varies in both SF and SW network structures.

In the graph, we observe a distinct relationship between the proportion of social learners in the population and their reaction to various products advertised by influencers, highlighting different adaptations in SW (Small-World) and SF (Scale-Free) models.

In the SW model, regardless of the product advertised, a significant portion of the population consists of social learners. This proportion exhibits a mild increase, ranging from about forty percent for the least valued good to roughly fifty percent for the most rewarding one.

Conversely, the SF model shows a more dramatic change in the proportion of social learners. This figure starts at zero percent when the least valued product is advertised, escalating to over sixty percent for the most desirable product.

A key observation is that for product seven, the proportion of social learners is approximately half the population in both models, yet the aggregate payoffs are markedly different. This disparity suggests that influencers wield more power in the SF network. Their promotion of product seven significantly deters more individuals from choosing the optimal product. In contrast, the SW network demonstrates resilience to influencer advertising. The minor variation in the proportion of social learners across different products indicates a lower tendency to shift towards individual learning when faced with lower-end products. Moreover, the stability of the SW structure's results, even in the absence of influencers, further validates its robustness.

2.3.3 Influencers and Heterogeneous Advertising

The last thing we test with our model is when influencers are not all advertising the same product. Rather we use various distributions to test their impact on population payoffs. Mainly we test 5 different distributions: normal, uniform, polarized and transitional distributions from polarized to uniform and from normal to uniform. We show the frequency distributions of these functions in figure 4 (due to the small size of the number of influencers, we manually shape those distributions instead of using a defined function).




Figure 2.4: Distributions of goods advertised in the sub-population of influencers ($n = 50$).

We use the Scale-Free structure, 500 agents, 50 influencers, 100 simulations and 100 time-steps in each simulation. As before, the payoffs of influencers are subtracted from global aggregate payoffs so we easily compare the model with the setting without influencers.




Figure 2.5: Evolution of aggregate payoffs at each time-step for various distributions of advertised goods.




Figure 2.6: Evolution of the mean number of social learners at each time-step for the different distributions of product advertised tested.

Figure 5 shows that all distributions of advertised goods are detrimental to the global payoffs of the population. From the beginning of the model, the setting without influencers ranks highest in terms of payoffs and speed of learning. Among the various distributions we've tested, the one that leads to the highest payoffs is the polarized one, where roughly half of the influencers are advertising the lowest products, and the other half are advertising the highest rewarding ones.

By contrast, the lowest rewarding distribution is the bell-shaped one, where most influencers are advertising average products. This is coherent with our previous results, which showed that the most detrimental situation for the population was when the influencers were advertising products 6 and 7 on a Scale-Free structure.

Figure 6 provides some insights into the behaviors that lead to these results. It shows the evolution of the number of social learners over a hundred time steps for each distribution of products advertised. There is a symmetrical relationship between payoffs and the number of social learners. While the highest rewarding distributions are the ones with the fewest social learners, the lowest are the ones with the most social learners.

Most likely, in situations where average products are advertised, it's difficult for consumers to discern that better options exist, and thus, they're likely to trust social learning as the best method to choose a product at each round.

Instead, in situations where the lowest options are advertised, consumers can easily find better options (by testing individual learning at the beginning of the simulation), and therefore won't use social learning thereafter. This lead to better payoffs for those individual learners that won't listen to influencers bad recommendations. The consumers that keep being social learners are probably the ones that are connected to an influencer that promotes one of the highest rewarding option, which is then not an issue.

Overall we find that for all but the two best products, influencers action is always detrimental to the population whom would get higher aggregate payoffs without their presence.

Important finding reveals that under our assumptions of reinforcement learning, influencers do the most harm to consumers' payoffs when they are advertising average products instead of the lowest options.

They are minimally detrimental when their presence is mitigated by a small-world network structure, where no agent has a significantly higher degree than others. We also find a limited impact when the distribution of the products they advertise is polarized (with either ads about very low/high products).

The next part finally analyze the potential for review systems, but also show their limitations, by modeling heterogeneous tastes of agents.

2.3.4 Review system

With the advent of the internet, markets for experience goods have increasingly adopted review systems to guide consumers. These systems, typically implemented by search engines like Google, specialized websites such as TripAdvisor and TheFork, or two-sided market platforms like Airbnb, Uber, and Deliveroo, aim to mitigate the inherent uncertainties associated with experience goods by providing value indicators.

How these systems work is simple: all goods are available to be reviewed online, most of the time simply by grading it on a scale. The current trend is using stars, with 5 the highest number possible, and half a star the minimum.

It follows that when consuming an experience good, each consumer can review the good to inform the rest of the population, and everyone can use the reviews to choose a good. While we'll be considering a very basic model, the very concept of grading goods can be altered by many variables. For example:

- The willingness to rate honestly : The one leaving the review might be biased either positively or negatively by personal relationships to the company providing the good. While this is true in a review model and it might alter the quality of the grade, this is equally true in a social learning context, and as we chose not to include it in the former model, we won't do it either in the later.
- The conversion from the experience to the grade : while it can be easy to rank options, it can be harder to give accurate grades. We'll make the assumption that they are entirely capable of translating experience to reviews here.
- In reality, goods are not delivering the same value at each time. This could be verified for many experience goods, especially when they include a human-delivered service like restaurants or healthcare. Coupled with the hypothesis that agents would be more willing to leave reviews for particularly good or bad experiences, that might spur from an incident for example, could lead to distorted grades. However as for the learning model, we make the simplifying assumption that a good always delivers its true value and that agents review all the goods they consume.

Therefore, the model we use to represent the review system is straightforward: an agent chooses a good based on the existing reviews which are averaged to give a single grade. Naturally, the agent chooses the best existing options.

If two products have the same grade, the agent pick any of the two at random.

A simulation consists of several rounds. In each, all agents pick a good and rate it. Then the average grades are adjusted after each new review.

The only problem is the initialization phase. Because every good is unrated at the beginning, people must choose at random. Say that all goods are rated the average grade of 5 (on a scale from 1 to 10) initially.

If an agent selects a good whose true value is below 5, it is not a problem because the agent will rate it according to its true value. This action will lower the good's grade to below 5, leading the remaining agents to choose from the higher-graded options that are still available.

But if the next agent pick an above-average good, then it will grade it its true value and therefore, the good will be the highest graded, even if its not the good with value 10. Thereafter, all remaining agents are locked to this option because they use a simple maximum function.

To solve this issue we choose to update the grades only starting at round 2. It allows for the five hundred agents to pick a good at random and rate them their true value. It means that every agent chooses at random once. then the grades are updated once every agent choose once. It ensures (with a large enough population) that all goods are picked at least once, and that no lock-in effect appears.

It follows that this version of the review model is the most efficient possible. It allows agents to perform the best score possible from round 3, and never deviate thereafter.




Figure 2.7: Comparison of models where agents can only use the review system, or only use either individual or social learning, or when they can use the three strategies. We display mean (among 10 simulations) aggregate payoffs at population level at each round.

Our next point of interest is the situation where individuals have access to reviews, alongside private learning methods like individual and social learning. Does the presence of a review

system cause everyone to forsake private signals in favor of public reviews?

While many of us have access to review systems, and in theory, they offer the most effective way of gathering information, it seems that we do not solely depend on them in our everyday lives—at least, not universally.

Hence, our goal is to investigate whether our model can replicate this observed behavior in the population.

Figure 7 shows a comparison of the evolution of aggregate payoffs across three different models. Although the model that exclusively uses the review system is the most efficient, as previously mentioned, comparing a model that incorporates both public (review system) and private (individual and social learning) information with a model that depends solely on private signals uncovers interesting dynamics.

In the initial hundred rounds, the model that integrates both public and private information gathering is more efficient than the model relying only on private information, achieving higher payoffs more quickly. However, this efficiency diminishes afterwards.

There is a slight decrease in aggregate payoffs after the initial rounds, which is not attributed to the average of multiple simulations depicted in the graph but represents a consistent trend observed in each simulation.

In the long term, both models yield similar levels of aggregate payoffs, likely due to the behavioral choices of individuals. In the model where both public and private learning are available, the majority (approximately sixty percent of the population) predominantly opts for the review system as their most frequent behavior, and thus, likely as the behavior they commit to in the long run. Meanwhile, twenty percent continue as individual learners, and another twenty percent as social learners.

These numbers emerge and stabilize after approximately a hundred rounds, coinciding with the point at which the two models begin to converge in terms of the levels of payoffs they reward to the population.

Figure 8 illustrates the individual payoffs associated with the most frequent behaviors. It shows that the average payoffs for individual and social learners are lower than those for users of the review system.

Thus, at the population level, there seems to be no advantage for the existence of the review system if people can also learn to use other methods to make choices among alternatives (and thereby fail to learn to use the review system). However, the users who do utilize the review system significantly benefit from it.




Figure 2.8: box-plots of total aggregate payoffs by sub-population. Each dot is the aggregate payoff of one agent over one simulation. A blue dot indicates the aggregate payoff of an agent that has chosen the review-system method the most over the 100 rounds of the simulation.

We then question whether the review system retains its usefulness for its users when we discard the assumption that users have similar tastes. The following section presents results from a model incorporating heterogeneous tastes.

2.3.5 Review system with heterogeneous tastes

While the reviewing system is indeed a better option in a context of homogeneous tastes, as in the models we've studied before, it may not be the case if we consider heterogeneous tastes.

While so far we considered that the goods would reward the same payoffs to all users, indicating a shared ordering of preferences over the 10 alternatives, some experience goods

might not exhibit such an universal ranking. For example the value of cultural experience goods will greatly vary from one individual to another depending on tastes.

In this situation, reviews aggregated from the whole population, hence from agents with potentially different tastes might not give a useful information to a given individual .

Contrarily, it is known that social networks are inherently connecting people with similar characteristics and tastes McPherson et al. (2001); Puetz (2015); Duricic et al. (2021).

We thus adapt the model to allow for heterogeneous tastes and networks built through homophily. We keep the 10 products set, but give a different “best option” to each agents following an uniform law. This means that if an agent’s best alternative given its taste is product 2, then picking it would reward them the maximum payoff possible of 10.

To make sure that each agent, depending on its best option, still has 9 inferior alternatives, we use a circular representation of tastes. Therefore if agent’s preference is product 3 its second best options are product 2 and 4, third best options products 1 and 5, fourth best option products 10 and 6, fifth best options products 9 and 7 and its worst option is product 8.

Regarding the network formation, once everyone has a preferred choice, we simply repeat a procedure where each agent can pick the closest neighbor (in terms of tastes) among a pool of agents selected at random from the population. Controlling the size of the pool allows to control the strength of the homophily. We repeat the procedure until we have the desired graph density.

We can measure the homophily inside our network with the following measure:

$$\text{Network Homophily Index} = \frac{1}{N} \sum_{i=1}^N \left(\frac{\sum_{j \in \text{Neighbors}(i)} \text{Proximity}(i, j)}{\text{Max Difference} \times \text{Number of Neighbors}(i)} \right)$$

Where: N is the total number of agents in the network, $\text{Neighbors}(i)$ represents the set of direct neighbors of agent i , $\text{Proximity}(i, j)$ is the measure of taste proximity between agent i and their neighbor j , and Max Difference is the maximum possible difference in tastes, considering the circular representation.

The measure vary from 0, which would indicate the highest level possible of homophily: every agent is only connected to agents with the exact same taste, and 1 which indicates the total absence of homophily: all agents are only connected to individuals with the farthest taste then theirs. A random network, where agents connect to each other without taking homophily into account, will score 0.5 on the homophily measure for example.




Figure 2.9: We compare various levels of homophily in the network of interactions and the model using global reviews. plot show mean aggregate results at each round for 10 simulations of each model. $N = 500$, $\nu = 0.02$.

We evaluate various versions of the network, constructed with differing degrees of homophily. The levels range from 0.5, which corresponds to no homophily in the construction of the network, to 0.09, the highest level of homophily we were able to implement. In each network we keep the same level of density.

We then employ each of these networks to simulate a model of private information, where agents decide which product to choose through either individual or social learning. In these configurations, there are no review systems. This version of the model corresponds to the one explored in Chapter One, except now agents have heterogeneous tastes, and the people they are connected to somewhat reflect these preferences.

We compare these various simulations, where we vary the level of homophily, to a scenario where agents exclusively use the review system. The results are presented in Figure 9.

There is a clear advantage for the model with social and individual learning over relying on a review system when tastes are heterogeneous and that the network is built through homophily. Even a network without homophily, which is represented by the one with a level $H = 0.5$ which is build by reducing the size of the pool to only one agent (chosen at random in the population), and which corresponds to a random structure, scores better than the review model. And even in this case, there are social learners in the population that score lower than individual ones, but still above the average payoff they could expect if choosing a product at random.

A potential explanation would be that the social learners in this model are agents for which their neighborhood is composed of more agents that have similar tastes than dissimilar ones. Therefore they can consistently score better than if choosing at random by asking neighbors. At early stages of the learning process, this can be an advantage over the individual process.

Subsequently, a distinct positive correlation emerges between homophily and performance. The beneficial impacts of local social learning, as evidenced in these results, could be amplified by incorporating the probable mechanism of identifying which neighbor's advice to follow. In the model being examined, while agents currently select a neighbor randomly for social learning, the approach's efficiency could be substantially improved if agents had the ability to choose a neighbor based on whose recommendations align more closely with their own preferences.

2.4 Conclusion

This study aimed to explore the impact of the internet on decision-making in high-uncertainty situations, particularly when selecting experience goods. We focused on examining the influence of online influencers and review systems on the decision-making process and outcomes for the population.

As demonstrated in the first chapter, social learning was more beneficial to individuals in this context than individual learning, represented by reinforcement learning. However the advent of the internet brought unprecedented changes in information aggregation and diffusion.

It enabled users to connect with others at a low cost in terms of both price and time. While agents previously communicated with a few acquaintances regularly and accessed global information through news media, they could now engage in one-on-one communication with anyone using technology, primarily through forums or online social networks.

The easy aggregation of information led to the development of useful tools. In the context of experience goods, where the value of a good cannot be assessed before consumption, gateway websites like Google and dedicated platforms created ways to inform users by allowing them to

grade goods and access others' ratings.

However, these systems perform best when everyone shares a common preference ordering. Their efficiency decreases when considering agents with diverse tastes and different preference orders. In such cases, aggregated reviews start losing their informational value. Simultaneously, social networks regain interest due to their structure reflecting a homophily process that binds together people with shared tastes.

Under the hypothesis that tastes are heterogeneous and uniformly distributed among the population, networks formed endogenously around the idea of homophily attachment enable social agents to learn and perform better than with the review system. Even networks not built around homophily (heterogeneous agents connects at random) perform better than random choices at each round, reinforcing the importance of social interactions in decision-making.

Online social networks also gave rise to influencers, agents who capitalize on their popularity by selling their audience's attention to brands.

We demonstrate that in a context where agents can either heed influencers' endorsements or learn independently which goods to choose, influencers are most detrimental for average options. This is because, for very poor-quality goods, agents learn not to trust influencers, and for high-quality options, they learn to heed influencers' ads. However, for average options, many agents struggle to discern that better options are available, and continue to rely on influencers instead of exploring alternatives independently.

While these results hold true for a network structure built on the most extreme case of preferential attachment (star network), they are also true on scale-free networks, commonly found in online social networks. Conversely, influencers do not pose a significant threat in populations evolving within small-world networks.

Overall, these simulations reveal that social interactions with local acquaintances might be the best solution for making informed choices in the context of experience goods. While the changes brought by the internet seemed promising for improving population utility, they hinge on assumptions about people's distribution of tastes and the ethical choices of influencers, which might prove unrealistic.




Figure 2.2: Comparison of models performance in terms of aggregate mean payoffs at final round with and without influencers for SF and SW structures as the good being advertised varies. 1 is the mean performance for the models without influencers on the network structure.

Chapter 3

Opinions Polarization, Influencers, and Endogenous Online Social Networks

We propose a model to improve our understanding of the escalating polarization of opinions in recent Western societies. This phenomenon is often attributed to the role of online social network platforms' algorithms and the influence of key individuals in creating echo chambers of opinions. Our model situates the population within a network, enabling interactions both among individuals and with these influencers. Distinctly, influencers, as a separate class of agents, learn to shape their opinions to maximize their utility by expanding their follower base. Concurrently, the platform can implement policies to control content visibility based on users' existing beliefs.

Computer simulations of our model reveal several crucial insights. Contrary to popular belief, influencers tend to moderate public opinion, while platform policies favoring content alignment with users' existing views contribute to reducing polarization. Furthermore, our findings explain the Pareto distribution of influencers' network degrees, attributing it to a combination of word-of-mouth dynamics within the population and a limit on the number of influencers an individual is likely to follow.

Introduction

The emergence of online social networks (OSNs) in the mid-2000s profoundly transformed how people communicate and interact. This change is evident in direct communication whether

for one-on-one interactions or for broader audiences with posts, but also for the sharing of beliefs and opinions (Peter and Muth (2023); Center (2014)), the formation and participation in interest groups (Howard et al. (2011)), and the way people access news (Shearer (2021)).

Primarily, online social media platforms have transformed various aspects of both our social circles and our awareness of events outside these circles.

Most OSNs operate using feeds. These feeds display posts and recent news from individuals within an agent's neighborhood, meaning all the people to whom the agent is connected. This connection grants access to their profiles, posts, and reactions to other content on the platform. In contrast to offline interactions, the breadth of friends and acquaintances one can now access and stay updated with online is vast. Traditionally, an individual would engage with a few distinct and close-knit circles daily, such as family, colleagues, or close friends. However, with online platforms, that same individual can maintain a sense of closeness and stay informed about the lives of high-school friends, colleagues from a decade ago, or distant cousins residing across the country.

While Dunbar empirically demonstrated and theorized that the size of our social circles could average around 150 people (Dunbar (1992); Dunbar et al. (2015)), the shift from offline to online interactions has undoubtedly influenced how individuals engage with this newly accessible group.

The manner in which people interact on OSNs is fundamentally different from traditional forms of communication, such as phone calls or written letters. While OSNs do offer a semblance of these traditional interactions through chat messaging, they also introduce novel modes of communication, including posts, re-posts, comments, and reaction buttons. These innovative methods foster distinct user behaviors. For instance, posts allow individuals to address their entire network simultaneously, while re-posts and comments enable users to endorse or share content they find compelling. Reaction buttons attempt to compensate for the lack of non-verbal cues inherent in online communication. In essence, these platforms alter not only the content people share but also the breadth of their audience in certain instances. Consequently, online communication diverges significantly from offline interactions.

Furthermore, OSNs often suggest connections to new individuals, regardless of whether you personally know them. They also display content in your feed from individuals outside of your immediate network.

The primary goal of this approach is to prevent users from leaving the platform. Once a user has caught up with all the activities of their connections, they might not have a reason to stay. However, much like traditional media, most OSNs generate revenue by monetizing the time users spend on their platforms, selling this engagement data to advertisers. Consequently, the longer an individual remains on the platform, the more beneficial it is for the OSN.

Algorithms employed by OSNs for recommending new connections and showcasing content from unfamiliar users are typically proprietary secrets of the respective companies. However, it's widely believed that these algorithms can drive various outcomes, contingent upon the platform's objectives.

For instance, if an OSN prioritizes introducing users to real-life acquaintances, it might suggest connections with many mutual friends on the platform. Such a strategy could enhance tie density within a confined user group. Alternatively, a platform might focus on connecting users with similar interests or beliefs, potentially bridging gaps and fostering interactions that might not occur offline. Another approach could prioritize promoting already-popular users, gauging popularity through their connection count or the engagement their posts receive. Most likely, leading OSNs like Facebook, X, or Instagram deploy a blend of these strategies to optimize user engagement and network growth.

In turn, these user retention methods employed by social networks are believed to potentially have concerning consequences on public discourse. They might lead to the formation of echo chambers, which could amplify both affective and ideological polarization, a serious concern for the well being of a society (Axelrod et al. (2021); Esteban and Schneider (2008)). Additionally, the rise of macro and micro-influencers on these platforms may result in asymmetric relationships among users.

Echo chambers emerge when OSN recommendation algorithms curate feeds based on users' presumed preferences. For instance, if the algorithm deduces a user leans politically right, it might predominantly showcase posts from right-leaning media outlets or influencers. The concern here is the absence of pluralism and heterogeneity in news sources and perspectives. Such a narrow viewpoint may not only reinforce existing biases but could also contribute to further radicalization within one's current political stance. Moreover, this lack of diverse exposure can stifle the evolution of more nuanced opinions that span across the political spectrum. In the end, echo chambers could lead to polarization. This issue is highly debated both in the news Chapin (2018) and inside the scientific community (Maes and Bischofberger (2015); Keijzer and Mäs (2022)), with papers highlighting both the detrimental effects of echo chambers (Cinelli et al. (2020); Cinus et al. (2021)) or its absence of effect to the matter (Bruns (2019)).

Recently, a series of research papers (Nyhan et al. (2023); Guess et al. (2023b,a)) had the unique opportunity to collaborate with Meta, the parent company of Facebook and Instagram, to conduct a series of experiments on their Online Social Networks (OSNs). These studies notably investigated the influence of echo chambers, network algorithms, and re-posting behaviors on political polarization during the 2020 American election. The collective findings from these three studies indicate that these phenomena had no discernible impact on political polarization.

While they are not typically studied in conjunction, a fundamental characteristic of many OSNs is their propensity to give rise to influencers—highly connected individuals on the network who garner attention and followership from their community. Such relationships especially thrive on directed OSNs like X or Instagram, where connections aren’t necessarily reciprocal, allowing for asymmetries to develop (Aparicio et al. (2015)).

These individuals have been identified as opinion leaders, as defined by Katz et al. (1955), in recent studies examining their impact on vaccine behavior (Bonnievie et al. (2020)) and commercial product promotion (Nandagiri and Philip (2018)). Recognizing their effectiveness in shaping behavior, politicians quickly adopted influencers for their campaigns. They have been effectively employed by the American government to promote pro-vaccine initiatives against the coronavirus and to campaign in the U.S. for both Biden and Trump (Lorenz (2021); Goodwin et al. (2020)).

Such practices pose challenges, especially when the commercial partnerships between the political party or candidate and the influencers are not transparently disclosed, potentially leading to audience manipulation. Moreover, political advertising by influencers can amplify extreme messages or give rise to conspiracy theories (Riedl et al. (2021)). Recent research from Gibson et al. (2023) indicates that exposure to these types of influencers can escalate extremist views, and that encountering paid political advertisements can erode trust in political institutions.

In turn, profound changes have occurred due to the global adoption of social media as a medium for interactions within local networks of acquaintances and global discussions on national or international topics. Previously, our communications were more direct, often private, and within closely-knit circles. Today, we express ourselves through posts intended for broad audiences, many of whom we might not know personally, using likes and reposts to share our stances. While we once sourced our information, knowledge, and opinions primarily from journalists and editorialists in newspapers and on television, we are now the audience of monetarily incentivized influencers who might promote any message without disclosures if paid sufficiently. While this shift isn’t entirely negative, given the ease of keeping in contact and the potential to foster powerful social movements, there are genuine concerns about how these transformative platforms are shaping global opinion dynamics.

In light of this discussion, the model to be presented in the next section aims to illustrate how opinions evolve on an OSN. Within this model, influencers are distinctly defined based on their incentives, and their differences from other platform users are clearly delineated. Furthermore, the model facilitates the study of the effects of various policies that OSNs might implement using algorithms.

The literature on opinion dynamics modeling is extensive. Comprehensive overviews of both theoretical models and empirical findings are available in Peralta et al. (2022). The contributions of these models in explaining polarization have also been explored by Keijzer and Mäs (2022).

The model developed in this work draws from the philosophy of Attraction and Repulsion mechanisms, as described in Axelrod et al. (2021). The fundamental concept is that our opinions evolve differently depending on whether we communicate with like-minded individuals or those with divergent views. In the former scenario, an individual might adjust their beliefs to align more closely with the interlocutor. In contrast, in the latter situation, the individual might adjust their beliefs to further differentiate themselves within the opinion space.

The novelty of our approach arises from the inclusion of influencers within the ecosystem. In our model, these agents actively pursue new connections, striving to amass as many followers as possible. Empirical evidence indicates that influencers can monetize their extensive networks, suggesting that such behaviors are likely to emerge online.

To expand their network connections, agents categorized as influencers will disseminate content. With the platform’s assistance, this content might reach audiences based on the engagement these influencers have previously garnered. Our model’s influencers operate as reinforcing learners; they adjust the content they share based on the rewards (or feedback) they receive.

This model juxtaposes ‘casual’ users, who primarily seek information or communication, with influencers who are motivated to maximize their influence. This dynamic enables us to observe the endogenous evolution of the OSN through simulations, which will be elaborated upon later.

Consequently, our model serves to integrate two pivotal types of agents that are instrumental in shaping the theoretical discourse on opinion dynamics: platforms (via algorithms) and influencers (via content dissemination).

Subsequent sections will provide a detailed explanation of the model and its dynamics. We will explore the evolving structural characteristics of networks that result from this model, evaluate the impact of influencers on the opinion landscape, and finally assess the effects of various platform algorithms.

3.1 Model

The framework is based on three categories of participants: casual users, who constitute the bulk of online social networks; a small number of influencers; and the social network platform. In terms of opinions, we focus on a single-topic opinion space with a continuous scale ranging

from 0 to 10. The distance between two opinions is gauged by their absolute difference.

3.1.1 Casual users

Casual users represent the majority of OSN participants. They engage with these networks primarily for entertainment, information, or other specific interests. Unlike influencers, casual users aren't actively pursuing an increase in connections. As a result, they have distinct motivations and behaviors, especially in terms of how their opinions evolve when interacting with others.

Initially, casual users are interconnected through a random undirected network, defined by the Erdős-Rényi model (Erdos et al. (1960)).

Initial beliefs are assigned independently, drawn from a random uniform distribution on the set $[0,10]$.

Importantly, we do not correlate the distribution of initial beliefs with the network structure. The distribution of any agent's neighbors beliefs is therefore uniform.

At each round of a simulation, agents will have the opportunity to interact among themselves.

When two agents do interact, the outcome is a revision of the belief they held before the interaction. The way we model the revision process of two interacting agents i and j is inspired by the Attraction-Repulsion Model (ARM) proposed by Axelrod et al. (2021). We consider two main effects:

- **Attraction:** When agents i and j hold beliefs that are sufficiently close in the opinion space, the initiating agent i updates its opinion to a position closer to j . Empirical evidence supporting this behavior can be found in Myers (1982) Myers (1982). Mathematically, the updated belief $b_{i,t}$ can be updated taking the mean of both beliefs:

$$b_{i,t} = \frac{1}{2} * (b_{j,t-1} + b_{i,t-1}) \quad (3.1)$$

- **Repulsion:** Conversely, if agents i and j have beliefs that are sufficiently distant in the opinion space, agent i adjusts its opinion to a position even further from j . Such behavior is also empirically observed, as noted by Flache et al. (2017) Flache et al. (2017). This can be mathematically represented as:

$$b_{i,t} = b_{i,t-1} + r * (b_{i,t-1} - b_{j,t-1}) \quad (3.2)$$

Considering how the two equations are related, for the attraction and repulsion effects to be equally strong, r should be equal to 0.5.

Notably, interactions in our model are not always reciprocal. Agent i might modify its belief due to influence from agent j , while agent j remains unaffected. This is analogous to user behaviors on OSNs like Twitter, where each individual has a profile showcasing their tweets, retweets, and replies. Thus, agent i might be influenced by visiting j 's profile without j reciprocating the visit or experiencing any change in belief.

Lastly, to determine the point at which two opinions diverge enough for repulsion to occur instead of attraction, we introduce a threshold T . This threshold T can take values between 0 and 10, representing the difference between two beliefs. A value of $T = 0$ indicates a scenario where even the slightest difference between two opinions results in repulsion, thereby eliminating attraction. Conversely, $T = 10$ signifies a scenario devoid of repulsion, as the maximum possible difference between any two opinions in our model cannot exceed 10.

3.1.2 Influencers

As previously discussed, influencers, as defined in this context, differ significantly from casual users. While casual users primarily engage with the platform for personal or social reasons, influencers are driven by the pursuit of network richness. They aim to convert this vast network reach into monetary benefits, such as through commercial partnerships and advertising opportunities. While the term 'influencer' often conjures images of individual personalities, in the context of an OSN—especially within political spheres—it can also encompass entities akin to traditional news companies.

These differing motivations lead to distinct behaviors exhibited by influencers on the platform. Their primary goal is to attract the maximum number of casual users to connect with them. Such connections allow these users to access and engage with the content shared by the influencer. In essence, influencers are in pursuit of followers.

To model this unique agent category on the platform, we initialize them without any immediate neighbors, contrasting the approach taken with casual users who are integrated into the random network structure from the outset.

Unlike casual users, influencers do not start with a predefined belief. Instead, over time, they experiment to determine which opinions yield the greatest success in terms of acquiring new followers and eliciting reactions. We propose modeling this learning process using reinforcement learning. A variety of reinforcement learning models have been introduced and utilized in academic literature. Key examples include works by Arthur (1993); Erev and Roth (1998); Bush and Mosteller (1955), and a justification for employing this learning paradigm can be found in Brenner (2006). Central to reinforcement learning is Thorndike's 'law of effect' (Thorndike (1927)), which posits that behaviors rewarded in the past are more likely to be repeated, while those that led to adverse outcomes are avoided.

We adopt the model proposed by Bush and Mosteller (1955), which provides each agent (in our context, the influencers) with a probability vector. This vector outlines the likelihood of selecting each available alternative—in our case, opinions ranging from 0 to 10—for subsequent interactions. Following every interaction or experience, this probability vector undergoes updates. For instance, when an influencer shares an opinion on the OSN, the reactions it garners and the net change in followers serve as feedback. This feedback is quantified into a reward function. A high reward amplifies the likelihood of the influencer expressing that opinion again, whereas an opinion leading to follower attrition would see its associated probability reduced for future iterations.

Mathematically, the vector $\mathbf{p}(t)$ represents the probabilities associated with the selection of various opinions a at time t . The update to this probability vector is influenced by the reward $\Pi(t)$ corresponding to the selected action $a(t)$ multiplied by a variable ν that can be used to control the impact of rewards on the probability vector, and thus to shape the learning process. This relationship is defined as:

$$p(a, t+1) = p(a, t) + \begin{cases} \nu \Pi(t)(1 - p(a, t)) & \text{if } a = a(t) \wedge \Pi(t) \geq 0 \\ \nu \Pi(t)p(a, t) & \text{if } a = a(t) \wedge \Pi(t) < 0 \\ -\nu \Pi(t)p(a, t) & \text{if } a \neq a(t) \wedge \Pi(t) \geq 0 \\ -\nu \Pi(t) \frac{p(a, t)p(a(t), t)}{1 - p(a(t), t)} & \text{if } a \neq a(t) \wedge \Pi(t) < 0 \end{cases} \quad (3.3)$$

This configuration of the reinforcement learning model accommodates both positive and negative rewards. This flexibility is beneficial as our reward modeling encompasses both types.

3.1.3 Influencers-Casual Users Interaction

Rewards for influencers are determined by the reactions of casual users to the influencer's most recent opinion post. These reactions are quantified in terms of new connections, disconnections, and “reactions,” which are designed to mimic comments or reaction buttons on OSNs.

When a casual user encounters an influencer's post in the current time-step, they first revise their opinion, just as they would when interacting with another casual user in the network through the Attraction-Repulsion model. Subsequently, they decide whether to follow or unfollow the influencer and whether to leave a comment on the post. Naturally, the option to unfollow is only available if the casual user had previously chosen to follow the influencer.

To model the probabilities of following, unfollowing, and reacting, we use the following functions:

- The probability for an individual to connect to an influencer (star) after encountering their content is a linearly decreasing function of the absolute difference between the individual's

opinion and that of the influencer:

$$p(\text{connect}) = 1 - \frac{|b_{i,t} - b_{s,t}|}{O} \quad (3.4)$$

where O represents the “maximum” opinion (in our case, 10). $b_{s,t}$ is the opinion posted by the influencer at time-step t , and $b_{i,t}$ is the opinion of the casual user.

- The probability for an individual to react is very high when the difference between the two opinions is either at its maximum or its minimum:

$$p(\text{react}) = \left(\frac{|b_{i,t} - b_{s,t}|}{2 \cdot O} - 1\right)^2 \quad (3.5)$$

- The probability to sever an existing connection from an individual to an influencer is described by a convex increasing function. We opted against a linearly increasing function for the connection probability, believing that individuals more readily form connections than they sever them:

$$p(\text{disconnect}) = \left(\frac{|b_{i,t} - b_{s,t}|}{O}\right)^2 \quad (3.6)$$

Thus, the reward an influencer receives at the end of a time-step is the change in connections from the previous time-step plus the number of comments garnered by their post.

3.1.4 Discovering Influencers and the Platform’s Role

The final aspect to address is how influencers are discovered in this model. Since they are not initially connected, we allow for several processes, each with varying influence, to enable influencers to reach casual users on the OSN. Mainly we distinguish the following effects:

- n : the “natural” rate of reach. At each round, the influencer reaches a fixed proportion of individuals at random in the population. This is used to initially introduce influencers in the early rounds where the other two processes are inactive. Furthermore, it represents a kind of random discovery process that might exist in reality, albeit to a limited extent.
- f : This parameter represents the proportion of agents an influencer reaches in the current round, determined by the engagement level of their most recent opinion publication. Specifically, the number of reactions and connections an influencer’s content garnered in the previous round directly influences the extent of their reach in the current round. For example, if due to its last publication, an influencer gained n new followers and the post itself gave rise to c comments, then at present, the influencer reaches $f * (n + c)$ agents.

This mechanism reflects a system where the platform gives precedence to content that has previously generated significant engagement from casual users. The underlying rationale is that content which has demonstrated appeal and has engaged users effectively in the past is more likely to encourage casual users to spend additional time on the platform in the present.

- w : the "word-of-mouth" parameter. Individuals recommend influencers to others. Consequently, the more agents already connected to an influencer, the more agents the influencer can potentially reach. This exhibits a clear Matthew effect (Perc (2014)).

It is evident from the model that the primary role of the platform is to curate which influencers are showcased to casual users. In the baseline model, we employ the f parameter, representing an OSN that is pragmatic in its recommendations. However, by adjusting this algorithm, we can introduce various recommendation strategies. For instance, the OSN could favor influencers who share similar views with users, those with opposing views, or even those presenting moderate opinions.

We will delve deeper into these scenarios in the subsequent section.

3.1.5 Model flow

To summarize, the initialization phase and a typical time-step execution of the model proceed as follows:

A population of casual users is placed within a neighborhood on a randomly connected, undirected network. Each agent is given an initial belief.

A small subset of influencers, relative to the overall population size, is introduced into the same network; however, they initially have no connections.

As the first round begins, each casual user randomly selects a neighbor to engage with, leading to the belief revision process.

Simultaneously, during this round, influencers choose an opinion to share. Based on the previously explained rules, they are presented to casual users. Consequently, influencers may gain and/or lose ties within the network and receive comments from users.

In response to these rewards, influencers adjust their probability vectors, which influence the opinions they will share in subsequent rounds.

3.2 Results

3.2.1 Opinion dynamics

We initially focus on the evolution of belief distributions within the population. Given the prevalent debate, as mentioned in the introduction, regarding the role of OSNs in intensifying polarization on political topics, we aim to investigate its impact within our model.

To quantify polarization, we adopt the definition provided by Esteban and Ray (1994). According to the authors, polarization is observed when:

1. There are a few large groups that dominate, in contrast to isolated individuals who have a minimal impact on the indicator.
2. There is a high degree of homogeneity within each group.
3. There is a high degree of heterogeneity between different groups.

In our framework, a group is defined as all the agent sharing the same belief. Homogeneity and heterogeneity are gauged based on the beliefs of the users. Here homogeneity is ensured since the group is composed only of people sharing the exact same belief. Heterogeneity is the distance is measured as the distance between two beliefs. The polarization measure is as follows:

$$P_{(\pi,y)} = K \sum_{i=1}^n \sum_{j=1}^n \pi_i^{1+\alpha} \pi_j |y_i - y_j|$$

Here, K is a constant. y_i denotes the class of i (in our context, a belief), and π_i represents its corresponding weight in the population, such that $\sum_{i=1}^n \pi_i$ constitutes the total population. For example y_i can be the belief 5, and if it's the belief of a third of the total population, then its weight $\pi_i = 1/3$.

α belongs to the interval $(0, \alpha^*]$ where $\alpha^* = 1.6$. “This parameter reflect the degree of ”polarization sensitivity” of the measure. The larger is its value, the greater is the departure from inequality measurement”. For more details see Esteban and Ray (1994).

Specifically, we are exploring the role of influencers on polarization. We perceive them, along with the enhanced interactions within personal social circles, as the most significant agents of change compared to offline interactions.

Thus, to assess their impact, we run our model with and without the presence of influencers within the population and examine the resultant variations in the distribution of opinions among the population over time.




Figure 3.1: Each panel represent a simulation lasting 100 time-steps of two distinct models: one without influencers and one with influencers. The evolution of polarization in the absence of influencers is represented in each panel by the red line. For the model with influencers, we separate the population representation. In blue we show the evolution of polarization for the casual users population, while in yellow the evolution of the indicator for influencers at the same round. We use a repulsion strength parameter of $1/2$. We test from $T = 1$ to $T = 9$. $N = 300$ for both casual users populations. $N = 30$ for the population of influencers.

Figure 1 depicts simulation results, illustrating the dynamics of polarization through a hundred time steps for two distinct scenarios. Each panel presents the dynamics at a different threshold level T . Lower levels of T signify high intolerance to differing opinions and a propensity for polarization, while higher levels of T indicate substantial tolerance to varying views. To recap, if the difference between two opinions is below T , attraction occurs; if it's above, repulsion occurs.

In the studied scenarios, maximum polarization is observed when the population is split into two groups of equal size ($N/2$), with one group adhering to opinion 0 and the other to opinion 10. This polarization level is approximately $P = 3.4$ with $\alpha = 1.6$.

In contrast, when the population uniformly concurs on an opinion, the polarization is $P = 0$. The simulations initiated from a state featuring a uniform distribution of opinions, corresponding to a polarization level of approximately $P = 1$.

When examining the model devoid of influencers (lines in red), it reveals that only two outcomes are plausible: absolute polarization or unanimous consensus. The transition between these states is observed when the threshold level T lies between 5 and 6.

Our best hypothesis explaining why the transition occurs at this level of T is related to the proportion of agents that cannot be repulsed away during the first round, which impacts subsequent rounds.

When $T = 4$, in the first round, all agents have a positive probability of meeting someone whose belief is such that it would cause repulsion, and therefore, a move towards polarization. This is because the threshold is low enough so that, for any agent across the belief spectrum (which spans from 0 to 10), there is at least one segment of the population (holding beliefs divergent enough to be outside the threshold) that will lead to repulsion. However, for agents holding belief 5, only meeting agents with beliefs 0 and 10 will lead to repulsion, which accounts for 20% of the population (since beliefs are randomly attributed).

But when we move to $T = 5$, those agents with belief 5 are not repulsed by meeting an agent who believes in 0 or 10, because it now falls within the threshold boundaries. Consequently, these agents can only converge towards the agent they meet in the first round, regardless of that agent's belief.

But for $T = 5$, our simulations show that the population keeps ending-up being polarized. We believe this is because the proportion of those agents that can't be repulsed by any divergent belief is not high enough (it's only 9% of the population at first round). Furthermore, by interacting with others at the first round, they will average their opinion and the opinion of the agent they just met. Thus they won't be holding belief 5 in the next round, unless they meet an other agent holding belief 5, which has a 1,6% chance of happening. The other way any agent holds belief 5 at this next round is if someone believing 4 meet someone believing 6

or if someone with belief 7 meet someone with belief 3. In total there is 4,8% of interactions that can lead to an agent holding belief 5 in the second round. This drastic fall of those type of agents might explain why we still witness polarization as a long-term dynamic when $T = 5$.

Then for $T = 6$, the proportion of agents that can only converge with someone else opinion grow to all the agents holding beliefs 4, 5 and 6, which represent 27% of the population. And our simulations shows that on average, at the next round, we still have around 21% of agents with beliefs in that range from 4 to 6 where agents can only get closer to other's opinions in the next round.

We believe that this sustainable proportion of agents who will not repulse from any other belief is the reason why starting at $T = 6$, the dynamics move from polarization to consensus.

This is a first important result from this model. To converge towards conformity, a society don't need to have only people ready to converge towards the opinion of others, but a proportion of 20% seems to be enough so that their action leads the whole population towards consensus.

However it's important to note that these results happen when the threshold is already high ($T = 6$), and the same for the whole population.

The introduction of influencers into the Online Social Network (OSN) significantly alters polarization at the population level. With identical model parameters, the dynamics of population opinion (depicted in red) diverge markedly. The dynamics of the influencers' opinions are also illustrated (in yellow).

The interaction between influencers and casual users constitutes an interdependent system where influence on opinions is bidirectional. Influencers, once they have gathered enough followers, can broadcast an opinion, thereby attracting or repulsing a large number of users simultaneously. However, what influencers choose to publish is shaped by their perceptions of user preferences, which are inferred from users' decisions to follow or unfollow them and from their reactions. Thus, the outcomes observed are inherently a result of this interdependence, with one group potentially exerting more influence over the other.

For low thresholds levels, from 1 to 5, the population confronted with influencers still rise towards polarization, but the process is slowed down, apparently due to the influencers who slowly polarize themselves, learning to diffuse what the population will likely engage with. It's interesting to note that there is a steeper curve of polarization in the first rounds in the model with influencers compared to the population without their presence. Our explanation for that phenomena is the increased number of interactions each agent has when there are influencers in the model. Indeed while the population without influencers only interact once by round, the presence of influencers gives the opportunity to make more interactions at each round. And because the population is initially confronted with very diverse opinions, each interaction

increases the chance for polarization for these thresholds levels (There are more chances to interact with someone with an opinion above the threshold).

Then the situation is also very interesting for threshold levels superior to 5. While the issue sparking this model is polarizing societies, there's likely an evenly strong case to be made about the perils of living in a society where total conformity—everyone holding the same exact opinion—prevails.

Because opinions are not truths, there is danger whenever an idea is no longer debated but taken for granted by the whole group. From a purely utilitarian perspective, it has been demonstrated that diversity is associated with higher levels of innovation (Hewlett et al. (2013)). Galton's renowned experiment on the wisdom of the crowd (Galton (1907)) also showed that diverse populations tend to make more accurate estimations than individual experts when asked to guess about something.

In this regard, it appears that a world with influencers permits societies that are slightly less uniform, allowing for polarization levels different from 0, although they remain low. For threshold levels from 6 to 9, the presence of influencers shifts the population's opinion distribution from a uniform one to what appears to be, around 0.5, a distribution more akin to a normal distribution.

In every simulation, the convergence is centered around the opinion 5.

An explanation for the asymmetrical impact of influencers - when transitioning from thresholds that induce polarization to those that induce total conformity - may be found in the initial rounds of the model.

The population swiftly gravitates towards polarization for $t < 6$ as they are exposed to uniform distributions of beliefs (both from their neighborhoods and from the few influencers who initially propose random opinions and are discovered with the natural rate of discovery n). This leads the influencers to rapidly adapt to this polarization by diffusing extreme opinions themselves. If the initial opinions encountered by people were more aligned with their beliefs, then polarization would not occur as quickly.

Conversely, for $T \geq 6$, encountering uniform distributions of beliefs is not problematic as it results in more attraction than repulsion; it becomes statistically less likely to encounter someone whose belief would induce repulsion. However, by disseminating random opinions and gradually learning the population's beliefs, influencers actually decelerate the unification process by presenting users with opinions far removed from their own. This effect is enduring, as these divergent opinions are introduced to new users in each round, either through word-of-mouth or platform algorithms.

As we've analyzed, the sub-population of influencers seems to act as an anchor for the

opinion's distribution of the population, both slowing down processes of polarization and total unification. We wish to validate this hypothesis by asking what would happen if those same influencers were to diffuse random opinions instead of learning what the population wishes. Figure 2 shows this experiment for every threshold level from 1 to 9 for one simulation of 50 rounds.

There is a clear anchor effect from the influencers when they're spreading random opinions. Indeed while it's not able to change the early dynamic of polarization for thresholds below 6, it is nevertheless acting as a regulator for those cases. For greater thresholds levels, the population closely align with the distribution of influencers, showing that influencers do have a great impact on the opinions of the population.




Figure 3.2: Each box represent a simulation lasting 50 time-steps of the model with influencers with different thresholds levels. Lines represent the evolution of polarization among the respective populations. blue is the whole population ($N = 300$), and yellow the sub-population of influencers ($N = 30$). We use a repulsion strength parameter of $1/2$. We test from $T = 1$ to $T = 9$.

3.2.2 Distribution of Influencer's degree

A common feature of OSNs with directed connections (where i can connect to j without the reverse being true) is the emergence of positively skewed distributions of in-going connections (such as the number of followers on Twitter, for example).

This process has been uncovered for many different types of networks in Barabási and Albert (1999) with a simple explanation for its existence: the 'preferential attachment' hypothesis. For the authors, the explanation for this network feature would lie in a Matthew effect, also sometimes called the 'rich-get-richer' effect. Basically, it explains that the high-degree nodes would have gained their numerous connections not from intrinsic qualities or performance, but rather from the opportunity of being on the network at its beginning. There would then be a natural tendency for new entrants to connect to people already there, combined with a simple heuristic of connecting to the already well-connected.

While there is certainly some truth to this explanation, we suppose that some other dynamics might be at play which could explain as well the asymmetrical success of Influencers. Because in this model the whole sub population of influencers arrive at the same time (at $t = 0$), we can explore those other hypotheses.

In the baseline model we presented earlier, all influencers succeeded in obtaining a high number of connections. This is expected since they all learn the same way. If we look at the degree distribution of the global population, including influencers, then yes, we have a scale-free distribution with a few highly connected agents and the rest of the population with far fewer. But certainly not all who wish to be influencers succeed, and so we will instead focus our attention on the distribution of degree inside the sub-population of influencers to try to understand what drives success and failure in this specific group of agents.

The two main factors likely to generate disparities are the word of mouth process of diffusing information about existing influencers and the bounded time and attention the population of casual users are willing to spend on the platform.

The word-of-mouth mechanism operates by expanding the potential pool of new followers through the existing follower base of an influencer.

This process is crucial for influencers as it allows them to expand the pool of potential new followers and the number of potential reactions to their posts, which in turn influences whether the platform is likely to highlight them in the upcoming rounds.

Specifically, if an agent i has x neighbors within the population of regular users and is connected to an influencer k , this agent can 'introduce' its neighbors to k by informing them about the influencer's existence.

The w parameter defined earlier then control the proportion of those x neighbors that will




Figure 3.3: Influencer's degree distribution when we vary the strength of the word-of-mouth mechanism coupled with variations on scarcity of users' capacity to follow. In blue the distributions for $T = 4$, in red for $T = 5$ and in yellow for $T = 6$. Each panel is a different combination of the maximum number of influencer any casual user can follow and of the level of word-of-mouth parameter. For example 13_0.2 is the simulation of the model when any casual user can follow at max 13 influencers, and the the word-of-mouth parameter is set to 20%.

be reached through the word-of-mouth mechanism.

To test whether the word-of-mouth mechanism can drive inequalities in the success of the influencer population, we run our model while varying w from $w = 0$ to levels of 20%, 40% and 60%.

What the word-of-mouth mechanism likely induces is a self-reinforcing process, wherein an influencer can expand their pool of potential followers in each round, provided the structure of the network of casual users remains homogeneous. This is also contingent on the influencer continually converting a similar proportion of potential to actual followers, implying the dissemination of opinions that resonate well with the beliefs of the population, or at least a significant portion of it.

This process is inherently inclusive, allowing all influencers the opportunity for growth through word-of-mouth diffusion. However, minor disparities in the ability to disseminate the 'right' opinions early in the simulation could lead to significant discrepancies. This is particularly true as success through word-of-mouth increases the likelihood of being spotlighted by the platform, thereby further expanding the pool of potential new followers.

Moreover, it will be crucial for influencers to achieve success early on, especially considering the variable of scarcity. The scarcity hypothesis posits that each casual user, having only a finite amount of time to spend on the platform, cannot connect and engage with every influencer. This is particularly plausible when considering that users might prefer to avoid redundancy in the opinions they are exposed to and may wish to prevent being overwhelmed by excessive information, which would hinder their ability to access their preferred content on the OSN. Thus, we argue that the scarcity hypothesis, which suggests a limit to the number of influencers a casual user is willing to connect to, is plausible. It's important to emphasize that scarcity does not lead to user lock-in as users always have the option to disconnect from any influencer if they find the diversity in the opinion space to be too broad.

Even though agents have the option to disconnect, the interplay of word-of-mouth and scarcity is likely to augment the path dependency of initial successes for any influencer. This can potentially expand their pool of prospective followers while concurrently undermining the success of other influencers. This is because existing followers might have already reached their connection limit and, therefore, are not open to forming connections with other influencers.

The results from Figure 3 confirms this hypothesis. By coupling the two variables and observing the distributions of in-degree for influencers for three different levels of the threshold we can see that the most skewed distributions arise for very high levels of scarcity and at least existing word-of-mouth mechanism. When agents can connect to many influencers, distributions are much more symmetrical.

These results show that the preferential attachment hypothesis can be further detailed, in

this context, as a combination of a word-of-mouth mechanism where success can breed further exposition and thus, more success, and a scarcity parameter the model the bounded time we spend on a social platform, or the limits to people's willingness to follow and engage with influencers.

Therefore it enrich the explanation of why some influencers will be more successful than others.

Platform's algorithms

Lastly, we employ the model to investigate the implications of potential policies that the platform could implement. These policies, manifested as algorithms, would regulate the visibility of various influencers to the casual user base.

We examine three distinct algorithms utilized by the platform for influencer recommendations: one that prioritizes influencers who have recently expressed opinions aligning with the user's current views; one that does the contrary, highlighting those with opposing views; and a third that exclusively presents influencers with moderate opinions to users.

We choose to only focus on the three threshold which represent most of the dynamics in this model. When $T = 4$ or below we're on a easily polarizing context, $T = 5$ act as a transitory state and $T = 6$ or above will most likely move opinion distribution toward consensus.

When $T = 4$ we witness diverse dynamics which highlight a clear impact from the algorithms on population's opinions. We display a benchmark dynamic in black that represent the opinion dynamic without algorithms (the platform simply shows engaging influencers without taking care of individuals opinions). Then in red is shown the so-called echo chamber where the platform highlight influencers that are close on the opinion space. In blue the opposite (platform display influencers distant on the opinion space to the user) and in yellow a moderate filter which would only show influencers with recent moderate opinions, typically in the range 4 to 6.

for this threshold level, there is a clear detrimental effects of showing opposite opinions to users, contrary to moderate and filter algorithms, which seem to act as decreasing polarization inside the population.

The effect is even more pronounced for $T = 5$. In this scenario, while showing opposite opinions through influencers increase polarization compared to the benchmark model, the two other policies lead to a population with converging opinions, in the range 0.5 which typically exhibits normal-law types distributions.

Finally, for $T = 6$, we don't discern any effects of the policies compared to the benchmark model.

These results gives theoretical foundations for recent empirical results (Guess et al. (2023b,a); Nyhan et al. (2023)) which found no detrimental effects of online social networks (and therefore




Figure 3.4: 3 Different algorithms tested for $T = 4, T = 5, T = 6$. In red an algorithm that propose close-opinion influencers, in yellow influencers with moderate opinions, in blue ones with distant opinions and in black a benchmark simulation without any algorithm.

of echo-chambers) on polarization. It seems that if people do indeed behave in an Attraction-Repulsion kind of way regarding their opinions, then displaying people close to you on the opinion space will likely lead to less polarization than if no algorithm is in place if people's threshold is below 5, and actually leads to more consensus otherwise.

This result is not necessarily expected from the model, because while the attraction effect will take place under an echo-chamber with the influencers, casual users are still connected to other casual users which are not necessarily close on the opinion space. So a typical user, in a given round, can both get closer to influencer's opinion, but also repulsed from a neighbor's one. In term, depending on the scope of these effects, we could also have witnessed more polarization.

3.3 Conclusion

This study has illuminated the diverse effects that occur online when various actors engage with differing motives. Specifically, distinguishing between individuals who aspire to be influencers and the general population reveals a compelling bi-influential system in which both groups reciprocally influence each other.

While regular users predominantly share and discuss opinions—leading them towards either polarization or consensus—influencers aim to propagate views that will maximize their success on the platform. Consequently, the presence of these influencers seems to decelerate the polarization process and reduce the likelihood of total consensus.

We were also able to gain deeper insights into the factors that drive success for influencers, moving beyond the conventional explanation of preferential attachment in scale-free structures. We discovered that the most pronounced skew in network success arises from the conjunction of two factors: firstly, potent word-of-mouth effects among regular users, indicating that they disseminate information about the influencers they follow to their peers. Secondly, a pronounced scarcity in the number of influencers that a regular user is willing to follow. This amalgamation fosters a path dependence for early success, even though individuals can unfollow influencers.

Utilizing the model, we could assess various algorithms that a platform could implement to optimize user engagement. There is a prevailing debate regarding the impact of OSNs on polarization, with some proposing that algorithms intentionally display content aligning with user opinions. Contrarily, our model demonstrates that such algorithms would actually mitigate polarization and could potentially result in a homogeneous distribution of beliefs within the population, contingent upon the population's receptivity to divergent opinions.

This finding notably aligns with recent empirical studies investigating polarization on Facebook during the 2020 American election campaign. Similarly, these studies did not observe an impact of algorithms on polarization.

This suggests alternative explanations for the increasing polarization observed in western societies. For instance, Bakshy et al. (2015) discovered that users themselves tend to limit their exposure to divergent news and opinions, irrespective of the algorithms employed by the platforms.

Furthermore, the study by Andris et al. (2015) on cross-partisan collaborations reveals a trend towards polarization that predates the creation of the first OSN, as illustrated in Figure 5 (sourced from their article).

Lastly, we recognize that numerous extensions could be integrated into this model to either explore issues closely related to Online Social Networks (OSNs) or enhance the model's realism. For instance, this framework could be instrumental in investigating the dissemination of misinformation online. Additionally, incorporating the ability for regular users to establish and sever connections among themselves users could increase coherence.




Figure 3.5: Each member of the U.S. House of Representatives from 1949–2012 is drawn as a single node. Republican (R) representatives are in red and Democrat (D) representatives are in blue. Edges are drawn between members who agree above the Congress’ threshold value of votes. The threshold value is the number of agreements where any pair exhibiting this number of agreements is equally likely to be comprised of two members of the same party (e.g. D-D or R-R), or a cross-party pair (e.g. D-R).

Chapter 4

Online Social Network protocols

In the competitive market of Online Social Networks (OSNs) used by the population, explaining why one platform outperforms another, or why users migrate, remains a complex challenge. While existing literature often emphasizes the competitive advantage created by network effects, our research proposes that network protocols - the foundational rules shaping the creation of OSNs and the interactions within them - play a crucial role in why users prefer one platform over another. To substantiate our argument, we employ computer simulations of different network structures, derived from various network protocols. Our findings reveal significant insights; for instance, directed networks can markedly impede the diffusion of information, and the presence of sub-communities is vital for enhancing collective actions. These simulations demonstrate that the nuances of network design can lead to vastly different outcomes, providing a deeper understanding of user behavior and platform dynamics in online social networks.

4.1 Introduction

The recent purchase of Twitter (now known as "X") by Elon Musk and its subsequent announcements regarding the future of the platform (such as the major changes to the Twitter Blue program) quickly fostered waves of worry among many of its users, encouraging them to look for and collectively migrate to an alternative Online Social Network (OSN) that would serve the same functions as Twitter. The favored network soon became Mastodon, an open-source alternative that essentially serves as a substitute for Twitter, allowing users to post and view short text messages — capturing the essence of the former OSN.

Figure 1 illustrates the evolution of active users on Mastodon over time, with November 2022 marking Elon Musk's acquisition of Twitter. Although the call to migrate from Twitter



Figure 4.1: Evolution of Mastodon's users base

to Mastodon had a discernible impact on Mastodon's user statistics, it's noteworthy to see that this surge in growth persisted for only a month. A rapid decline ensued, with over one and a half million users ceasing to use the OSN within a few months. Thus, six months after Twitter's acquisition, Mastodon can report a doubling of its active user base. In contrast, Twitter's latest data report (from the end of 2022) cites approximately 240 million daily users, and Musk recently announced a count of 250 million daily users. While the migration could be considered a win for Mastodon, it hardly appears to be a loss for Twitter.

These recent events provide a compelling illustration of the perplexing competition dynamics that platforms can sometimes exhibit. How can we account for the abrupt halt in migration after an impressive gain of 2 million users, with most presumably coming directly from Twitter?

Literature on platform competition typically cites network effects as the primary force shaping dynamics in such markets. Network effects, as seminally developed by Katz and Shapiro (1985) and Arthur (1989), detail how the value of certain goods is intrinsically tied to their existing user base. Put simply, owning a phone in a world where no one else has one provides no utility to its owner. Thus, network externalities are a crucial factor when considering online social networks. Because if no one is using Facebook, its utility becomes null for anyone that would consider joining.

The literature also makes evident that the winning competitor is not necessarily the superior one but rather the one that gained an advantage in the early days of the market. Therefore, understanding why one social network has more users than its competitors and becomes the market leader often hinges on the "historical events" that favored it.

In November 2022, such a historical event occurred. Elon Musk purchased Twitter, prompting many active users to reconsider which OSN they should use. Strictly speaking, we aren't

exactly in the situation described by Arthur’s model because the population isn’t deciding which good to adopt. Instead, they’re determining whether they’re willing to switch from one to the other, a situation Farrell & Saloner explored through theoretical models (Farrell and Saloner (1985, 1986)).

Overall, given such a dazzling start, the network externalities stream of literature would suggest that a snowball effect should be taking place. This rapid growth would only further fuel the subsequent adoption by millions of other users. Notably, cascades leading to the fall of an online social network have recently been studied by Török and Kertész (2017).

However, if this is not the trend we’re observing, there might be another factor explaining the sudden decrease in the Mastodon hype. While Mastodon is designed as an open-source alternative to Twitter, and using it can feel very much like Twitter, we argue that certain network *protocols* result in a notably different experience. This could explain Twitter’s resilience to the detriment of Mastodon following the buyout by Musk.

If you examine any major online social network (OSN) today, it’s evident that they largely offer similar functionalities to their users. Most are structured in the following way: Users can share “posts”, which can manifest in different formats, including text messages, audio clips, videos, or photos. In response, other users can engage with these posts through reactions (such as “likes”) or by leaving comments. Additionally, most of these platforms provide a private messaging feature that facilitates direct communication between users. However, not all platforms prioritize the same functionalities. For instance, Instagram emphasizes photo sharing, TikTok focuses on brief videos, and Twitter on concise text messages. Yet, in principle, one could use Twitter in much the same way as Instagram, or LinkedIn as an alternative to YouTube.

Then the distinct user experience of these OSNs might lie from the unique network protocols each employs. By ‘network protocols’, we refer to the strategic choices made by the platforms that directly influence the structure of the user network. In this paper, we highlight the primary protocols as: the nature of ties, the limits on the number of users an individual can connect to, the type of user recommendations, and the ability (or lack thereof) to establish sub-communities.

By ‘the nature of tie’, we mean the way in which an OSN allows users to establish connections with others. For instance, Facebook and LinkedIn operate on the basis of reciprocal ties. When you want to add another user to your network, you send them a request. The recipient can then accept, decline, or ignore this request. On the other hand, platforms like Twitter and Instagram use directed ties. This means that you can follow someone without needing their approval. However, this doesn’t guarantee that they will follow you back.

Most platforms also impose limits on the number of users one can connect with. For instance, Facebook allows up to 5,000 connections, while LinkedIn permits 30,000. On directed networks

like Instagram or Twitter, there's a cap on the number of people you can follow, but no ceiling on the number of followers you can have.

This interplay of the first two protocols explains why there aren't macro-influencers boasting hundreds of millions of followers on Facebook or LinkedIn like on Instagram — such colossal neighborhoods simply aren't permitted by the platform's constraints.

Furthermore, the ability for users to form and rally around sub-communities can transform the network's structure, making it more dense. The essence of most sub-communities is the capacity to engage with all of its members through public posts. The idea is to facilitate sharing around a common interest. So, even if you're not directly connected to everyone in the group, a single post within the community can effectively reach them all. Therefore a sub-community could be considered a complete sub-graph among all members, although the weight these ties carry might not be as strong as a 'real' connection on the OSN.

Lastly, a significant feature of an online network is its recommendation algorithm. This algorithm determines whom the platform will suggest you connect with. Recommendations can be explicit, such as "Here are 10 people on this network you might know," or they can be more subtle. For instance, the platform might highlight activities from people you aren't directly connected to in your feed, typically framed as "You might be interested in this." Most major OSNs employ both strategies. We can categorize the typical recommendations into a few types: individuals the platform believes you know in real life, individuals with similar interests or backgrounds, and those already well-connected on the platform, often termed "influencers".

In turn, Recommendation algorithms play a pivotal role in shaping network structures. For instance, they can amplify the reach of established 'influencers' or promote emerging ones, likely leading to the formation of scale-free structures, as described by Barabási and Albert (1999).

Given its status as open-source, non-profit software, Mastodon avoids the typical recommendation algorithms often found in other platforms. While it might display global trending topics, the platform doesn't offer the same level of personalized content as major OSNs do. While this approach holds several merits, such as preventing echo chambers and addressing data privacy concerns, this particular network protocol might adversely impact user experience, especially during the crucial initial phase of platform discovery. This could, in part, account for the observed decline in active users following a brief period of intensive growth.

In the same time frame, Meta's Threads, a direct competitor to X, has notably outperformed Mastodon. Garnering tens of millions of monthly users in the U.S. alone (Intelligence (2023)), Threads likely benefits from Meta's established use of recommendation algorithms, similar to those in Facebook and Instagram. This contrast in user engagement with Mastodon, which eschews such algorithms, might bolster the argument about the pivotal role of network protocols in shaping social platform success.

In light of this discussion, it becomes clear that while many platforms, ranging from streaming services and game consoles to newspapers, rely on third-party contributors—such as directors, game developers, and writers—to carve out their competitive advantage, OSNs stand apart in this landscape. The unique competitive strength of OSNs doesn't primarily derive from external creators, but from their intrinsic network protocols. It's the distinct architecture and rules set by these platforms that determine how users can establish connections, interact, and subsequently shape the network. This inherent design not only dictates the user experience but also results in a specific network structure, which becomes the platform's defining asset and source of competitive advantage.

In this chapter, we postulate a causal linkage between the network protocols instituted by firms and the utility consumers extract from their OSNs. This hypothesis is grounded in the notion that these protocols critically shape the interaction dynamics among users, which in turn impacts the efficacy of these networks in fulfilling the diverse expectations of their user base. Such a framework suggests that the architectural choices of network protocols are not merely technical decisions but also key determinants of consumer utility that can therefore influence the success of one OSN over another.

The first part of the chapter will therefore be devoted to generating the different networks, each following a different combination of protocols, and then to study how different they are in terms of various network characteristics.

Then the second part of our argumentation tests the statement that different network structures will result into different dynamics or outcomes for the users. While it is obvious that a star network will unfold different processes than a complete graph, we're interested to know if the realistic network protocols we model here will result into different enough network structures so that they do have an impact on the processes we will test. This second part then tests three different processes that frequently happen and that are partially the reason why people use OSNs.

We focus on how information spreads, how groups can come together for a common cause, and how job seekers connect with potential employers. These processes, by reflecting prototypical uses of social networks, bring part of the value to the users of the platform. We're focusing on quantifiable measures of success. For instance, most would agree that an OSN is more effective if it spreads information quickly, helps a community mobilize, or aids in job searches. Admittedly, our focus does not encompass the entire spectrum of platform utilities, notably those related to entertainment. Given the subjective nature of entertainment—what captivates one individual might bore another—it presents challenges in terms of precise modeling. Consequently, we have chosen to exclude it from our current analysis.

Combining the two parts of this chapter, we'll be able to link the performance of a particular network structure, on one or plural processes, to the network protocols that shaped it.

It can then be used either by firms to understand how they can improve or build new OSN that would fit their ideas, and the functionalities they believe users will want to use it for, or by consumers, to choose the right OSN according to their need.

The rest of this study is structured as follow: We start by generating different online social networks that would mimic the combinations of various network protocols and evaluate their network properties. Then we apply the three enumerated processes: the diffusion of information, the dynamics of collective action and a simulated job market on each of these network structures. We demonstrate that for each process, we identify significant effects of network protocols. This is an important finding for platforms seeking to enhance their performance and, more importantly, for users who should be aware of certain platforms' limitations and choose the right OSN based on their specific needs.

4.2 Modeling Networks

We begin by building the networks that will serve as our OSNs, each designed based on a unique set of protocols outlined in the introduction. Of the four protocols mentioned, we retain three: the nature of ties, the ability to form sub-communities, and the recommendation algorithm. We exclude the 'limits on individual degree' protocol, as it primarily poses a constraint for Facebook.

We categorize the nature of ties and the ability to form sub-communities as binary variables. Specifically, a tie can be either directed (D) or undirected (U), and the network either allows (C) or does not allow (N) the creation of sub-communities. As for recommendation algorithms, we consider the three previously discussed options: recommendations based on real-life acquaintances (F), those based on similar interests or characteristics (A), and recommendations of influencers, which we also refer to as 'Stars' (S). This categorization results in twelve distinct network structures. Throughout our discussion, we will refer to these structures using their respective initials. For instance, an undirected network that recommends real-life acquaintances and allows the formation of sub-communities will be designated as 'UFC'.

4.2.1 Creating a "Real-life" network

While modeling some protocols is straightforward, such as depicting the nature of ties, others require a bit more "background building". This is particularly true when considering the recommendation of real-life acquaintances. If an Online Social Network (OSN) aims to suggest

connections to people you're familiar with, it presupposes that the platform has access to some level of information about your real-life social circles, whether partial or complete. In practice, this data can be obtained through various techniques. For instance, Facebook encourages users to provide their phone numbers and grant access to their phone contacts. Should you consent, Facebook can then cross-reference your contacts with other users who have also shared their phone numbers. Additionally, if you disclose details such as your age, place of study, or workplace, the OSN can make educated guesses about other individuals you might know who have shared similar information.

We thus need a network that can replicate this real-life scenario. The aim is to represent a social network as it existed before the advent of Online Social Networks (OSNs). Such neighborhoods typically consisted of families, friends, colleagues, and so on. In this context, ties are not weighted.

Importantly, we want the network to embody two characteristics commonly observed in social networks: small-world and scale-free properties.

The small-world phenomenon in social networks has been well-understood since the studies of Milgram (1967); TRAVERS (1969). The structure and modeling of small-world networks were further elaborated upon by Watts and Strogatz (1998).

Similarly, the concept of scale-free networks, introduced by Barabasi & Albert, underscores the idea of asymmetric node connectivity within graphs. The emergence of scale-free structures in social networks can be attributed to factors such as popularity, hierarchies, or brokerage roles. To substantiate the presence of these scale-free properties, we examined two datasets: a mail communication network and ego networks from Facebook. These datasets are particularly appropriate proxies for real-life networks as they largely exclude connections to 'unknown people,' such as influencers who gain prominence through OSNs like Twitter or Instagram. Both datasets are sourced from the Stanford Large Network Dataset Collection¹. Our analysis indeed confirms both the scale-free and small-world nature of these networks.

We thus aim to generate a network that would allow N individuals to connect through two distinct processes: homophily and popularity. While homophily promotes connections with similar agents, leading to the formation of cliques, popularity fosters the emergence of scale-free-like structures.

This leads to two questions: how can an agent discern someone as more similar than another and how can it assert that one agent is more popular than another?

To deal with homophily in our model, we arrange every agent from our population along a circular segment. We position each agent at equal intervals from one another. This arrangement

¹<http://snap.stanford.edu/data/>

is designed to reflect the level of similarity between agents: the closer two agents are on the segment, the more similar they are to each other; conversely, the greater the distance between two agents on this circular segment, the more dissimilar they are. Thus if an agent wants to connect to someone based on homophily, it will pick an agent that is close on the segment.

Then to deal with popularity, we simply use degree centrality, meaning that an agent will consider another agent popular proportionally to its degree on the network. We artificially create slight inequalities in degree at the beginning of the model by selecting a few agents and making them connect at random. There is a few agents that will have a higher degree before the rest of the population join and they will potentially become the high-degree agents found in scale-free networks.

Once this initialization phase is done, agents sequentially pick an other agent to connect to. They all have the same probability p of connecting based on homophily and the probability $1 - p$ of connecting based on popularity. If the agent connects to someone similar, it will look into the closest not-already connected agent on the circular segment. If it connects to someone popular, it selects from the pool of the highest-degree agents.

After all agents have made their initial choices, this process is repeated for the entire population multiple times, until a specified network density is achieved.

We use the following parameters for the model: $N = 1000$, a probability of connection based on homophily $p = 0.9$, 3% of the agents initially connect using a random process with a density of 30%, and the connection process is repeated 10 times. This results in a connected network with a density of 2%, an average path length of approximately 3, and a clustering coefficient of 63%. Both the average path length and the clustering coefficient suggest that the graph possesses small-world properties. Below, we present the degree distribution on a log-log scale, which further indicates a scale-free-like behavior of this structure.

4.2.2 Cultural Diversity

The next step is to recommend individuals who share similarities across one or multiple dimensions. This is essential not only for the recommendation protocol of the Online Social Network (OSN) but also to ensure that, when agents form sub-communities, they do so based on shared interests or values.

To capture the richness of cultural heterogeneity within populations, we adopt the conceptualization of culture presented by Axelrod (1997) in his paper on cultural dissemination. In this framework, every agent possesses a cultural vector of length F . Each element within this vector represents a cultural feature. These features can be conceptualized as overarching cultural markers, such as religion, musical preference, sporting affiliations, political views, and so forth. Each feature can assume one of T possible traits. Taking the "religion" feature as an example,




Figure 4.2: Degree distribution of the Real-life network

its possible traits might include Christianity, Judaism, Islam, and others. By determining the values of F and T , we can calibrate the depth of cultural diversity within our population. For instance, with $F = 5$ and $T = 3$, the cultural profile of an agent i might appear as: $[2, 1, 2, 3, 2]$

These cultural vectors are then use for the recommendation of similar people. The OSN will typically propose to connect to someone that has a given proportion of similar traits. Similarly, agents will join sub-communities that are centered around a given trait of a specific feature.

4.2.3 Recommending Influencers

Lastly, some of our OSNs will make recommendations based on popularity. At the onset of the OSN's creation, there won't be any standout or popular agents. To kickstart this dynamic, we employ a method similar to the one used for generating the real-life network. The first batch of agents to join the OSN will be randomly interconnected. With a small enough initial group, and by establishing connections among them at a density of 30%, we ensure sufficient variability in their connectivity. As a result, while some agents might have only one connection, others could have as many as four. When subsequent agents join the OSN, and if they receive recommendations to connect with these early members, the initial agents will naturally emerge as more popular choices. Over time, as this dynamic intensifies, these early agents evolve into "stars"—individuals with a degree significantly higher than the majority of the population.

4.2.4 The Online Social Networks

We can now generate networks based on the twelve combinations of network protocols we discussed. These combinations arise from three main choices:

- The network’s structure, which can be either undirected or directed.
- The platform’s recommendation system, which suggests friends, people with similar traits, or star users.
- The OSN’s policy on sub-communities, either allowing or disallowing their creation.

All combinations exhibit a consistent adoption pattern resembling an S-curve, reflecting real-world adoption processes, as described by Rogers et al. (1962). We model this using a Susceptible-Infected approach with a 10% infection rate, applied to the real-life network.

Starting with an initial 2% of the population, each subsequent round sees new members join based on platform recommendations. If the platform suggests friends, there’s a chance of an erroneous recommendation. This simulates real-life scenarios where platforms may mistakenly suggest unfamiliar connections.

For similarity-based recommendations, cultural vectors use parameters $F = 10$ and $T = 5$. The platform then suggests users sharing at least half of the newcomer’s traits.

If sub-communities are allowed, new members can either form individual connections or join a sub-community, linking them to all its members. In directed networks, this affiliation generates both incoming and outgoing connections.

The connection process repeats until the desired network density is achieved. To ensure our model’s accuracy, we aim for low density, reflecting the sparse nature of real-world OSNs, as noted by Bhattacharya et al. (2020). Even with just a thousand agents, two iterations produce a 2% density, surpassing most real-world OSN densities.

Following table reports basic network statistics for each of the combination tested. We simulate each structure ten times and report the average for each measure:

Metric	DFC	DAC	DSC	DFN	DAN	DSN
Density	0.02	0.02	0.02	0.00	0.00	0.00
Average Path Length (In-going links)	3.05	3.04	2.89	11.92	6.09	2.26
Silent Nodes	159.20	236.60	241.90	176.20	785.10	980.70
Clustering Coefficient	0.61	0.60	0.60	0.06	0.02	0.01
KS-test	0.07	0.07	0.07	0.03	0.08	0.02
p-value	0.94	0.83	0.90	0.96	0.57	0.85

Metric	UFC	UAC	USC	UFN	UAN	USN
Density	0.02	0.02	0.02	0.00	0.00	0.00
Average Path Length	3.02	3.10	3.00	6.03	4.34	4.00
Clustering Coefficient	0.60	0.63	0.60	0.07	0.02	0.01
KS-test	0.06	0.06	0.06	0.03	0.08	0.03
p-value	0.87	0.98	0.86	0.95	0.39	0.72

Looking at the following metrics, we can see great variations from the different structures. First regarding density, it is clear that the community feature greatly improve the number of edges in the network. This is expected since communities are designed, as explained previously, as complete components.

Regarding Clustering, the role of communities is here again predominant. It is interesting to notice that the networks without communities doesn't all perform similarly. The structures that promote the connections to Influencers (USN and DSN) are the least performing structures in terms of clustering. The ones that promote tie building through real-life connections are the ones with the better scores, leading to clustering coefficients around 7.5%. Though this is still far from the clustering expected in small-world networks.

The next metric we look at is the average path length. For undirected networks, the calculus and interpretation is straightforward. Here again sub-communities drastically reduce the average path length. Then in networks without sub-communities, the recommendation of friends greatly increase the average length. Recommending influencers or similar people leads to around the same number of degree.

But in directed networks, we first need to clarify the context in which we compute the average path length. Indeed, our directed graphs are the representation of online social networks such as Twitter or Instagram, where someone can "follow" an other agent, while the reciprocity is not guaranteed, and remains the choice of the other agent. The way we modeled these networks, an out-going link from i to j imply that i follows j . This means that it becomes reachable by j , who can diffuse posts that will be seen by i , but not the other way around.

Thus, we are interested in path length through in-going links. The way we built our networks, everyone is at least following one agent of the OSN, so everyone is reachable by at least one person. But not everyone is followed by someone, which means that some agents will be

”silent” in the sense that what their posting on the network won’t reach anyone else in the OSN.

Consequently, we present two pieces of information: Average path length through in-going links, and the number of silent agents in the network. Strictly looking at average path length it looks like the DSN structure will be the most efficient structure to diffuse information. However, considering the number of silent nodes (980), it is evident that only a very few agents receive all the connections. In this context, only the information of those stars will spread through the network. Similarly, DAN is bad both at allowing everyone to diffuse information with an average of 785 silent agents, but also at reaching people fast, with an average path length of 6. DFN on the other hand, while having the highest average path length of all, is able to keep its number of silent agents quite low. Then all the structures allowing sub-communities have similar behavior: about a quarter of the population is silent, but the average path length is low, with a mean of 3 degree of separation.

Finally, we’re interested in the scale-free attribute of our networks. To analyze it we use the Kolmogorov-Smirnov test to assess whether the degree distribution could have been generated using a power law. Here all OSNs have high enough p-values (superior to 0.1) to assume that our data could have been generated the power law. However, upon examining the linear representation found in the appendix for each of these distributions, it becomes apparent that some do not completely align with the commonly held perceptions of scale-free networks.

While the structure without sub-communities do indeed seem to show a scale-free structure, with the vast majority at very low levels of degree and very few agents with significantly higher number of connections than this majority, it is important to notice that not all recommendations systems (friends, similar agents or influencers) leads to the same “power” acquired by the connection-rich agents. Both for undirected and directed networks, the highest power of the rich is found in the star recommendation system which would be expected. Then recommending similar agents leads to more power for the influencers than the friends protocol. This could be explained by the fact that while friends will tend to connect among themselves, and thus create small cliques, mimicking the “real-life” network generated at the beginning, connecting to similar people will not necessarily lead to those same clusters, but exhibit a more random pattern.

For the undirected networks, every agent that has joined a sub-community will have a greater number of connections, which explains the high frequency of agents with around 20 ties. Those networks, nevertheless, shows both large amount of people with almost no connections, and very few with high number of connections.

The detailed network analysis of the various combinations we tested unveils a great diversity

of structures. This diversity suggests that we will likely observe distinct dynamics when we apply and test various social processes on these structures.

That's what we do in the next section: On every of the twelve different OSN, we test the dynamics for the diffusion of information, the leverage of a collective action, and the efficiency of a job market.

4.3 Diffusion of Information

4.3.1 Model

Now that we have generated our networks, and assessed how different their structural properties could be with various protocols, we want to understand how these variations can lead to different outputs in terms of objective measures. To start we investigate the role of the network structure in the diffusion of information. What we're curious about is about the speed, but also about the breadth of diffusion in our networks.

Evaluating the speed of diffusion on OSNs can be of great importance, as a growing part of the population are using them as a source of information (Shearer and Mitchell (2021)). It is even more important to understand how the diffusion happens with the presence of fake, or unverified news that can have important political or economic impacts (Wang (2020); Kogan et al. (2019)).

To simulate the propagation of a given information piece, we use a Susceptible-Infected (SI) model where agents on the network have two possible states: either being Susceptible, meaning that they did not receive the information (they accept information with probability equal to 1 once reached), or Infected, meaning that they had receive and accepted the information. This very basic framework is enough to show great disparities, as we'll show later.


We add one element to this model regarding the sub-community propagation. As we've explained, sub-communities work as complete sub-graphs inside the OSNs. However it would not be realistic to model the interaction between two agents that have decided to connect in the same manner as two agents that belongs to the same sub-community. Indeed, as the sub-community grow, so will the number of posts on it. We hypothesize here that the probability to care and read a post coming from the sub-community is decreasing as the sub-community grows. This is because there will naturally be more posts, and agents might not care about reading so much of them, or because they're not interested to engage with the community at all times when they're on the platform. Platforms like Facebook design those pages that act as sub-communities as places you must willingly go, which represents a barrier. And while posts from the sub-community can be displayed on the feed of the user, not all posts are shown. this

is thus coherent with the hypothesis that as a sub-community will grow, the less posts will be seen by a member of the community.

Therefore, for networks that allow sub-communities, we introduce the action from users to decide whether or not they'd like to open the piece of information that a neighbor wishes to share with them. If the information comes from a neighbor the agent chose to connect to, then $p = 1$, else, if the post comes from a neighbor the agent connected to through the sub-community, then $p = 1/s_c$ with s_c the sub-community that connects the two agents.

On each network, the process then goes as follow: We infect 1% of the population at random to be the initial owners of information. At each round, they give the information to 10% of their neighborhoods.

Of course, in such a framework, for connected networks, with enough rounds, everyone would end up infected. What's interesting to look at is the speed of diffusion. But many directed networks won't be strongly connected, and thus, the steady state is not necessarily a full spread of the information. First because, as we've seen, some structures induce that a large proportion of the population is silent which means that even if they want to share information, they can't because no one's connected to them. Secondly loops or subsets of nodes that aren't able to communicate to the rest of the graph can exist in such structures.



Consider the network 1 shown above. If the agent having information is among nodes i, j or k , then the information won't be able to move further than in this subset since no one following them. Thus l and m will never get the information piece.

4.3.2 Results

Figure 3 shows the diffusion spread for each structure. Results are the average for 10 simulations on each OSN.




Figure 4.3: Diffusion of Information

From the plot we witness a great diversity of behaviors resulting from the different OSNs structures. This is a first result regarding the importance of network protocols for the user's experience. This implies that while some OSNs will be fitted to diffuse information, some will not be at all. Particularly, directed networks, that don't allow sub-communities (DSN, DAN) are very slow at diffusing information. This is mainly because of the asymmetric out-going degree distribution that renders a great part of the network unable to share information. It is important to note that while all the other structures almost don't show variation from one simulation to the other, DSN and DAN exhibit higher levels of variation. For the DSN structure, only 10% of the simulations are able to diffuse to the remaining populations. As the structure is very close to a star-network, the probability of diffusion depends on whether the star (in our case one or plural stars) is part of the initial agent with the information. If no star is infected in the beginning, the diffusion becomes impossible. In terms of diffusion of information, these kind of networks are close to mass-medias as only one or few sources are able to diffuse to a very large audience. These structures are close to OSNs like Twitter or Instagram, where ties are directed, there is no place for sub-communities, and the algorithm is at least partially trying to recommend influencers and content based on center of interests.

We also observe that while sub-communities should act as the greatest force towards fast transmission of information, the simple rules we introduced in order to weight these type of ties makes them slower than some undirected OSNs without the sub-community feature.

Overall, undirected OSNs will be faster than directed ones and the sub-community is an advantage for diffusion only for directed networks.

To our knowledge, there is no comparison in the literature regarding the diffusion of information in a directed social network as opposed to an undirected one, which appears to be a key factor influencing diffusion speed in our model. However, some literature does confirm other findings of ours. Notably, the speed of information diffusion is influenced by various network characteristics, such as modularity (the existence and number of communities), average degree, and the relative degree of social hubs, all of which play a role in determining the efficiency of information spreading (Peng et al. (2020); Peres (2014)). On the other hand, the clustering coefficient has a negative impact on diffusion (Peres (2014)). This aligns with our findings: Among undirected networks, the most effective are those that do not allow sub-communities.

Overall, this is a first compelling result supporting our hypothesis that different network protocols will yield different outcomes for users. In terms of information diffusion, building a directed or undirected network drastically modifies the probabilities that a piece of information will reach the entire population. This is due to the many silent nodes (individuals not followed by anyone) present in the network. The proportion of silent nodes even increases, thus rendering

diffusion slower when the platform pushes agents to follow influencers or when it doesn't allow for the formation of sub-communities.

We now keep building our argument by testing the influence of the network protocols on a collective action process.

4.4 Collective Action

In 1978, Granovetter introduced the concept of Thresholds in collective behaviors (Granovetter (1978)). The idea, is to consider that for some actions, people need that a sufficient proportion of the population has engaged in it already before they can themselves join. Granovetter list many of those cases such as political protests or the diffusion of innovations.

Consider for example, the case of products with high network externalities. Adoption will be highly dependent on the existing user base. While some will be early adopters, ready to advocate for the further spread of the innovation, the majority will need to see that many others have already adopted the product so that they can feel confident that the product will meet its purpose, given that it has established a sufficient user base. In his seminal paper, Granovetter consider the case where agents have perfect information about the choices among the whole population. But it could be argued that in many contexts, local information, meaning looking at what your neighbors are doing, is more realistic than having global information.

The process of a collective action taking form inside a population is quite different than the propagation of simple piece of information. While the population is aware of the option of adopting the product, or of joining a protest, they'll be inclined to do so only if their perceptions of the proportion of the population already doing it is high enough. Depending on the distribution of thresholds among the population, and on the accessible information - here about the structure of the network - collective actions can vary in scope and length to take form.

Studying these kind of processes on online social platforms can be particularly useful as they drastically modify the way people can access information and form beliefs about the willingness of the population to engage in the given collective action, compared to traditional mass medias, or offline social networks, which are way more constrained geographically for example.

In the early 2010s, we witnessed a compelling example of how OSNs were used to amplify and spread anti-government protests in the Arab world, an event now known as the Arab Spring (Howard et al. (2011)). More recently, these networks have been instrumental in the development and dissemination of the MeToo movement (Hosterman et al. (2018); Manikonda et al. (2018)), especially on Twitter. The movement is often associated with its hashtag, which serves as a key tool for spreading the message on the platform. Additionally, OSNs played a

prominent role in Hong Kong's Umbrella Movement (Shen et al. (2020)).

Given the critical importance of OSNs in the existence and effectiveness of these movements, they must be understood as strategic tools. This section aims to demonstrate how varying network protocols - such as different recommendation algorithms, community policies, and the nature of connections among users - can influence the spread and, consequently, the success of collective actions.

4.4.1 Model

As for the diffusion of information, we use one of the simplest modelling of thresholds models of collective action. The goal is not to perfectly depict the process of engagement in a precise context, but rather to establish important differences among the different structures that we've generated.

Let's denote the set of agents as A , such that $A = \{a_1, a_2, \dots, a_{1000}\}$. For each agent a_i in A , there is an associated threshold t_i . Each t_i is drawn from the uniform distribution over the interval $[0, 1]$.

Furthermore, the process, if it is to start at all, needs to have agents with a threshold level of 0. To make sure of their presence, we define a random subset S of A such that $|S| = 0.01 \times |A|$. Then for each a_i in S , $t_i = 0$. For each a_i , we also allow a 5% interval below t_i which will be interpreted as meeting the required threshold. Imagine an agent j , with $t_j = 0.68$ and a degree of 3. What we want to allow is that if 2 neighbors of j engage in the collective action, meaning that j calculates a threshold of 0.66, j should nevertheless be more inclined to join because its threshold is very close, than not joining until its third neighbor has also joined, which might never happen.

Once this is defined, the process goes as follow: we draw a random ordering of the population. Then sequentially, for each a_i in the random drawing, a_i computes its ratio of engaged neighbors at the given period s $r_{i,s}$ and compares it with its threshold. If it meets its requirement: $v_i - 0.05 * v_i > r_{i,s}$, then a_i joins the collective action and can be seen as such directly for the next agent and for all subsequent rounds. There is no disengagement in this version of the model.

4.4.2 Results

To generate results we proceed as follow: For each of the 12 structures we test, we run 50 simulations. All simulations share the same original real-life network. They also share the same original distribution of thresholds. What will change from a simulation to another is simply the generation of the OSN. This way we ensure that we have isolated the structure and the

underlying network protocols as the main cause of differences from one structure to another in terms of levels of engaged population in the collective action. For each simulation, we run 30 time-step, which means that every agent has 30 chances to re-consider its choice of joining or not the collective action.




Figure 4.4: Final levels of engagement after 30 rounds in a population of 1000 agents.

As for the diffusion of Information, our results suggest a high heterogeneity regarding the ability of a given network structure to foster a collective action.

Figure 4 shows the final level of engaged agents after 30 rounds, for each simulation performed on every structure. There is a clear performance enhancement when the network allows for sub-communities. On these structures, the collective action is able to reach 75% of the population on average. During the least effective simulations on those structures, the level of engagement still reach 60 % of the population.

Among the OSNs that don't allow sub-communities, DFN is the closest performing one, with a similar mean level of engagement but a significantly higher variation from one simulation to the others, as shown in the figure 5.

It is also interesting to note that the directed structures are performing better inside the group of sub-community OSNs. Among those networks, the recommendation algorithm doesn't seem to have a high impact on performance.




Figure 4.5: Box plots for the 12 OSNs, in red OSNs with sub-communities, in blue OSNs without sub-communities.

The least performing structure is DAN, a directed network without sub-communities, and where recommendations are based on common interests. In such a structure neither influence or community can serve as a vector to engagement, which result in collective actions gathering in average 20% of the population.

Finally one structure stands out by its behavior: DSN, a directed network without sub-communities, and where recommendations are based on degree (so it recommends influencers). This combination of network protocols creates a highly asymmetric network in terms of degree, with many silent agents (no one follows them) and a small group of macro-influencers (some get more than 25% of the whole population connected to them). As for the experiment about the diffusion of information, this structure leads to either very small or very high levels of engagement. Typically high levels of engagement will be achieved when one or many initial engaged people are those influencers. Say that one influencer has a threshold of 0. It will start by convincing all of its followers with low enough thresholds. Then since influencers will also tend to follow other influencers, if propagation succeeds in this cluster of influencers, then the whole population can become engaged.

By contrast, in this framework, 'silent nodes,' which are agents that follow others but are not followed themselves, cannot propagate or fuel any collective action due to their virtual invisibility within the network. If the majority of the initial agents initiating an action are silent nodes, the action is unlikely to progress in this network. Moreover, even if these silent nodes become willing to participate in a movement as a result of influence from their network, their shift does not impact others' decisions, as they are not being listened to, unlike in a

strongly connected network.

In a scenario where 99% of the population are silent nodes, willing to engage in a collective action but only following a few influencers, the initiation of the movement depends heavily on these influencers. If the influencers are unwilling to participate, the movement is unlikely to ignite. Conversely, if the influencers do choose to participate, the movement can gain momentum easily.

This particular structure is close to networks such as Instagram, TikTok or Twitter (now X) which are allowing directed nodes, are mostly recommending influencers, and don't help users to build sub-communities.

Again, network protocols, as expressed by their resulting network structures, do modify the dynamics of a collective action process.

We finally test the impact of those protocols for a job market, where both job suppliers and seekers signal their presence on a network of acquaintances.

4.5 Job Market

The development of the Internet, and soon after of online social networks, gave job market participants a new tool to look for and advertise job openings. As many papers highlight (Stevenson (2008); Kuhn (2014)), internet as a search tool has drastically been adopted by users both for employed and unemployed users over the years and it has proven successful to increase the chance of being hired compared to people not using internet to search for jobs (Choi (2023)).

We can distinguish two different ways in which Internet has created new ways to look for jobs. First the use of websites and forums have acted as global information places where jobs openings could be published. For example we saw the rise of websites such as Monster.com or Craigslist.org being used as such. While this practice doesn't radically change how people might find jobs, because it is close to job fairs or looking for job ads in newspapers, it nevertheless reduce time delays because of the instantaneous of access online and increase the scope of people reached because everyone can access those websites regardless of their geographic localization.

Later, as Online social networks emerged, their potential for conveying information about job seeking and openings was quickly recognised, and dedicated OSNs were soon created, with its most successful one being LinkedIn. The processes that allow LinkedIn to function are closer to what people were doing well before the internet: using their networks of personal relationships, whether being friends and families (which could be labeled as strong ties) or more distant acquaintances (weak ties) to get jobs. It is well known that people use these

networks to find jobs, and even that the weak ties might be more useful to that purpose than strong ones (Granovetter (1973)). Empirical work has, since then, refined this theory, showing a paradoxical relationship: while weak ties help more people to find jobs, the value of a single strong tie is more valuable at the margin. This is explained by the numerous number of weak ties that we have, compared to only a few strong ones (Gee et al. (2017)).

Then, what an OSN such as LinkedIn allows is for each user to be able to explicitly write their network down, by virtually connecting to other members. Then as for traditional OSNs, a feed allows for users to share and receive posts from their neighborhoods.

It can therefore be used as a similar way to channel jobs opportunities and job requests through the same network but online. It can also enhance the capability to maintain such a network of professional acquaintances and thus be more efficient than the offline traditional way of doing it.

As for the previous processes, we now ask whether there are better structures than others to be used as job markets. This is both in the interest if platforms such as LinkedIn, which can build and maintain a competitive advantage by providing the network protocols that will ensure the highest level of matches among job seekers and job providers, and naturally in the interest of users of those networks, as they will profit from the most efficient OSN.

4.5.1 Model

We build a labor market with N participants, represented by a graph $G = (V, E)$, where V is the set of vertices ($N = |V|$), each representing a market participant, and E is the set of edges representing the social connections between participants. The graph represents a social network used by its users, firms or workers, to communicate about professional information such as in LinkedIn. Here the information diffusion is purely local and can only move from one direct neighbor ($d_{i,j} = 1$) to another.

Each vertex $v \in V$ can take one of three states: “S” (Sleeping), “O” (Offering), or “L” (Looking). “S” indicates an employed participant without any job to offer, “O” indicates a participant with a job offer, and “L” indicates an unemployed participant seeking a job.

When an agent i moves from the “S” state to either the “O” or “L” state, it shares this new information with its direct neighbors $N(i) = \{j \in V : (i, j) \in E\}$. If a neighbor is compatible (i is “O” and j is “L” or the inverse) then a match is created, and both agents return to state “S”.

We hypothesis that at least two other processes can have a significant importance in the dynamics we wish to explore here. Mainly we want to enrich the behaviors of our agents by gifting them memory and a varying willingness to share information in this particular framework.

A memory variable $M_{i,t}$ depicts the state of knowledge of agent i about the other agents offering/looking for a job at time t . As stated, information diffuse locally and thus it is mainly

filled with information about direct neighbors looking/offering a job at the current period. The memory length m gives the number of rounds the information will stay in agent i 's vector. The shorter m , the harder it will be for two neighbors of i , j and k (that are not direct neighbors themselves) to match if they're not offering/looking for a job at the same time period. if $m = 1$, and j becomes "O" at $t = 1$ and k becomes "L" at $t = 3$, then the match is never made because of i 's short memory for example.

We allow agent i 's memory vector to be enriched by information about outside its direct neighborhood agents looking/offering a job if this information is transmitted by one of i 's direct neighbor.

A direct neighbor j can transmit information about an other agent k (that is not in i 's neighborhood) to i depending on j 's willingness to help k find a job.

For simplicity, we posit that the willingness to propagate information about an other agent's status (O or L) is decreasing with the distance (degree of separation on the network) to that agent. The idea is that we might be very inclined to help friends and share their CV, and we might also do it for friends of friends, but as the distance increase, there is certainly a tendency to ignore. Thus the probability P of information diffusion can be written as: $P(d) = (0.5)^d$ where d is the degree of separation from the source. The factor 0.5 is the decay rate, which means the probability is halved for each degree of separation. It follows that the chance to share information about a neighbor is 1/2 and about a neighbor of a direct neighbor is only 1/4 for example.

Therefore matches can be created between j and k even if they belong to different neighborhoods, as long as an agent i possess in its current memory the information that the two agents are compatible: one is "L" and one is "O" and that the information has been transmitted to i by the agents on the path from j to k .

To sum-up a round in this model goes as follow:

- A proportion of agents are selected to switch state from "S" to either "O" or "L" at random.
- These agents share their new states with their neighborhoods. If any neighbor is compatible or has information about someone that is in its recent memories, a match is created and both actors of the match return to state "S".
- If that's not the case, neighbors write the status of the agent in their memory.
- Then everyone has the opportunity to share information about offers and proposals they have in their last-round memory to their direct neighbors. They can't give a piece of information to the one that initially gave them. The probability to accept the new in-

formation is negatively correlated with the distance from the agent offering/proposing a job.

- A new layer of memories for the whole population is created so they can store and diffuse information in the next round.

The principal element we track for now is the number of matches made over time. By setting the number of rounds and the number of new entrants in the market at each period, we can know the maximum number of jobs created in a perfect information/diffusion world. It is expressed by: $M = \frac{p \cdot R}{2}$ with M the total number of matches made, p the number of new “O” or “L” posts at each round and R the number of rounds in the simulation.

4.5.2 Results on Simple Networks

Before translating this model into our ecology of online social networks, we can try to understand what elements are key for a network to perform well.

First if we keep aside memory and diffusion decay, the only determinant of performance will be average path length. That is, a complete graph will ensure the fastest and best structure for jobs suppliers and seekers. This is because any two nodes with matching statuses will be able to directly communicate and match in a given round. By contrast, a linear or ring structure where everyone has only 2 neighbors on average will perform the worst because it achieve the highest average path length possible.

It follows that small-world or scale-free networks, for example which have low average path length will necessarily show good performances. Of course, the importance of average path length is true if job suppliers and seekers appear at random “geographic” places on the network. If all appears on a small restricted region of a graph, then the average path length of the whole network become useless to predict performance. For now the emergence of suppliers and seekers is indeed random on the graph.

Below is a table showing matching performances for various networks, each with similar density values, when we omit memory and diffusion decay. Here omitting memory is giving perfect memory to all agents and omitting diffusion decay is allowing information to diffuse without decay.

To test models, we set the number of job seekers and offerers appearing at each round to one, the number of rounds to 20, and the size of the population to 500. Consequently, the maximum number of matches possible, denoted as T , is 20, assuming one possible match per round. For each network structure, we run 10 simulations and calculate the mean number of matches based on these simulations.

	Mean number of Matches	Average path length	Density
Line	10.5	167	0.004
<i>SW, Rewiring probability</i>			
0	14.4	32	0.016
0.01	17.8	7	0.016
0.05	18.4	4.5	0.016
0.1	18.6	4	0.016
0.5	19	3	0.016
1	18.7	3	0.016
SF	18.5	3	0.016
Complete	20	1	1

These results confirm the intuition: the lowest the average path length the better for matching. Neither cliquishness nor "stars" inside networks have a significant impact on results. Indeed random graphs (when $p=1$) perform as well as all small-world structures (from 0.01 to 0.1 rewiring probabilities) and the scale-free structure which does have a few highly connected agents does not perform better than a random one.

Adding bounded memory and diffusion decay will create two barriers: a temporal and a "geographical" one. With low memory size, our agents will not be able to recall much information from the past, and thus the condition for a successful match will mainly be that seekers and suppliers manifest themselves inside a short, coordinated time-frame. With high diffusion decay, information will have trouble moving further than direct neighborhoods of agents. Reduction of average path length and density of networks becomes even more important with this behavior.

We now simulate the same networks, with the same number of rounds and same number of new entrants at every round. But we now include bounded memory where agents only remember the last round, and diffusion decay with the probability to diffuse: $P(d) = 0.5^d$ with d the shortest path from i to j . Results are then:

	Mean number of Matches	Average path length	Density
Line	0.9	167	0.004
<i>SW, Rewiring probability</i>			
0	10.6	32	0.016
0.01	12.3	7	0.016
0.05	15.2	4.5	0.016
0.1	16.7	4	0.016
0.5	17.7	3	0.016
1	17.4	3	0.016
SF	18.8	3	0.016
Complete	20	1	1

As expected, all structures (except the connected graph naturally) are affected by the new barriers. The line network becomes extremely inefficient, and it seems that the passage from a ring network to a random one exhibits a different pattern than before in terms of mean Matches. The number of matches seems to be more sensitive to slight changes in the average path length. The most remarkable results comes from the performance of the scale-free structure which now is the most efficient structure but also the only one that didn't decrease its mean number of matches. It thus appear the most robust to low memory and low diffusion.

4.5.3 Results on Online Social Networks

As for the simulations in simpler structures, the maximum number of matches, that would happen in a complete graph is of 20 matches because we create, at each of the 20 rounds, one agent looking for a job, and one agent offering one. We test each structure 15 times and report the box-plots for each one in figure 7.

As for the 2 precedent processes, the structure of the network is of great importance for this simulated job market. The main lessons are that undirected networks perform significantly better than directed ones, and that allowing sub-communities can also increase the number of matches.

The worst structure is the DSN one, which creates structure similar as star-graphs. While in the two previous processes of collective action and of information diffusion, this structure was either the best or the least performing, here it's always the worst one. It is most likely due to the one-way flow of information, through the influencers of theses networks, which while working to diffuse information or engage people into doing something, is not working for a job




Figure 4.6: Statistics for each OSN, 15 simulations for each.

market, where information must flow in both ways to create matches.

The group of undirected, sub-community structures are performing the best with a level of 75% job matches among the total possible ones. This is the case even though we discounted the probability to read information when it came from a community-based link. Nevertheless, this higher clustering and the presence of influencers in these structures, which place them in the Small-World category of graphs, seems to be the most efficient one. A characteristic that we also observed for SW networks in the previous section.

4.6 Conclusion


Overall, the hypothesis that the network protocols, which influence the network structure of OSN, have also an impact on common processes happening in these medias is theoretically validated by our models and simulations.


There is a high heterogeneity in the most and least performing structures whether we look at the diffusion of information, the leverage of collective actions or the performance of a job market. What's even more interesting is that an enhancing characteristic for one process might not be for another.

There is thus room for platforms to look at what they aim at achieving with their OSN, and revising their network protocols accordingly. This is even more important for users, which can use the right network for the right purpose. Thinking about it might increase the chance of success. Our very simple models would tell, for example, not to use Instagram or Twitter which are close to DSN structures, if you're looking for a job. Or that platforms that allow

sub-communities, such as Facebook or Reddit can be very powerful to gain traction in collective actions.

Appendix








Figure 4.7: Dynamics for a single simulation on each structure for collective action

General conclusion

When I began this thesis, my main objective was to emphasize the dynamic interplay of influence between social networks and individual behaviors, acknowledging that causality is bidirectional. The likelihood that you smoke or abstain, drink or abstain, or lean politically left or right, often reflects the behaviors of those close to you. Beyond this broad influence, your specific position within a network impacts not only your actions but also the dynamics of the entire group. The rate and extent of behavioral diffusion are shaped by the network's pathways of influence. For example, smoking habits within a classroom can vary significantly based on the structure of student interactions. Similarly, an individual's number of connections and their centrality within a network critically influence the pattern of diffusion. Thus, networks significantly shape our behaviors.

However, these networks are not imposed upon us. We actively create, maintain, and sever the connections that shape our networks, often in response to observed behavioral changes. Consider the classroom scenario where smoking spreads. Depending on certain network architectural features, smoking might permeate the entire classroom. Yet, suppose other values are prevalent among the students—such as religious beliefs that strongly discourage smoking. These students may resist adopting this habit by severing ties with peers who smoke. Consequently, the network becomes less dense, potentially features longer average paths, and could even become disconnected.

Thus, the initial social structure among students outlines a set of potential pathways for behaviors to spread. When effective, this connectivity exposes some students to practices they prefer to avoid. Consequently, the social structure adapts to align with the individual beliefs of each student, with the network being reshaped by these behaviors.

Moreover, people are becoming increasingly aware of their presence within networks and recognize their positions within these structures. This awareness is especially pronounced on online social platforms, where the network is not only palpable but also integral to the platform's purpose.

These observations led us to further test these observations through the several models we

built. Overall we find several key findings:

- The network structure do influence the behaviors of agents. This is a result we find in several chapters. In chapter 3, we show that the outcomes for social learners, when they are exposed to influencers, are greatly deteriorated when they move from clustered structures such as small-worlds, to structures with highly skewed degree distributions, where a small number of agents are greatly connected compared to the rest of the population. These are commonly known as scale-free structures.

Testing the importance of structures is also the object of chapter 5. In it we test if the network rules of OSNs affect social processes that happen on them. We find great variations among the different structures tested for information diffusion, leveraging collective action but also for the efficiency of a job market. Depending on the purpose of the network, some structures will be better suited than others.

- The behaviors of agents can easily spur the emergence of skewed distributions of degree. We first identify this phenomena in chapter 2, where social learners also learn about which agents to listen to. We find that some agents become central in the sense that they gather more people asking them at each round than other ones. And it is not necessarily correlated with them being more informed than others.

In chapter 4 we also find a skewed degree distribution in the opinion model, where influencers are trying to obtain more connections from the population. We show that under realistic hypothesis some influencers will succeed far more than others. Creating both a skewed distribution in the overall population but also inside the sub-population of influencers.

- Social learning is efficient. This the main finding of chapter 2. In a population where no one has better information's than others, it is safer, and on average as efficient, than individual learning. So allowing people to interact easily among themselves like an online social network would do can have benefits. It can lead agents to better learn, to change their opinions, to join a collective movement or even to find jobs more efficiently.
- But when they are manipulated, those same networks can be detrimental for their users. This is the case in chapter 3 when we introduce influencers to the network of interactions, or when the platform that regulates the online social network elicit the wrong algorithm in chapter 4 and chapter 5.

Overall, our works validates the idea of co-influence between networks and behaviors, and give more depths into the precise mechanisms that are at play. We focused on learning in the

context of social interactions, and showed how social interactions modify this process, with an emphasis on the role of online social networks, which, as we showed, greatly modify the dynamics of social learning.

Conclusion Générale

Lorsque j'ai commencé cette thèse, mon principal objectif était de mettre en lumière l'interaction dynamique d'influence entre les réseaux sociaux et les comportements individuels, en reconnaissant que la causalité est bidirectionnelle. La probabilité que vous fumiez ou vous absteniez, que vous buviez ou vous absteniez, ou que vous penchiez politiquement à gauche ou à droite, reflète souvent les comportements de ceux qui vous entourent. Au-delà de cette influence générale, votre position spécifique au sein d'un réseau impacte non seulement vos actions mais aussi la dynamique du groupe entier. Le taux et l'étendue de la diffusion des comportements sont façonnés par les chemins d'influence du réseau. Par exemple, les habitudes tabagiques au sein d'une classe peuvent varier considérablement en fonction de la structure des interactions entre les élèves. De même, le nombre de connexions d'un individu et sa centralité au sein d'un réseau influencent de manière cruciale le schéma de diffusion. Ainsi, les réseaux façonnent significativement nos comportements.

Cependant, ces réseaux ne nous sont pas imposés. Nous créons, maintenons et rompons activement les connexions qui façonnent nos réseaux, souvent en réponse à des changements de comportement observés. Prenons l'exemple de la classe où le tabagisme se propage. En fonction de certaines caractéristiques architecturales du réseau, le tabagisme pourrait imprégner toute la classe. Pourtant, supposons que d'autres valeurs prévalent parmi les élèves, comme des croyances religieuses décourageant fortement le tabagisme. Ces élèves peuvent résister à cette habitude en coupant les liens avec leurs pairs qui fument. Par conséquent, le réseau devient moins dense, présente potentiellement des chemins moyens plus longs et pourrait même se déconnecter.

Ainsi, la structure sociale initiale parmi les élèves délimite un ensemble de chemins potentiels pour la propagation des comportements. Lorsque cette connectivité est efficace, elle expose certains élèves à des pratiques qu'ils préfèrent éviter. Par conséquent, la structure sociale s'adapte pour s'aligner sur les croyances individuelles de chaque élève, le réseau étant remodelé par ces comportements.

De plus, les gens prennent de plus en plus conscience de leur présence au sein des réseaux et

reconnaissent leur position dans ces structures. Cette prise de conscience est particulièrement prononcée sur les plateformes sociales en ligne, où le réseau est non seulement palpable mais aussi essentiel à la finalité de la plateforme.

Ces observations nous ont conduit à tester davantage ces observations à travers les différents modèles que nous avons construits. Dans l'ensemble, nous trouvons plusieurs résultats clés :

- La structure du réseau influence les comportements des agents. C'est un résultat que nous trouvons dans plusieurs chapitres. Au chapitre 3, nous montrons que les résultats pour les apprenants sociaux, lorsqu'ils sont exposés à des influenceurs, se détériorent considérablement lorsqu'ils passent de structures regroupées comme les petits mondes à des structures avec des distributions de degré très biaisées, où un petit nombre d'agents sont beaucoup plus connectés que le reste de la population. Ce sont communément des structures dites sans échelle.

Tester l'importance des structures est également l'objet du chapitre 5. Nous y testons si les règles du réseau des OSN affectent les processus sociaux qui s'y produisent. Nous trouvons de grandes variations parmi les différentes structures testées pour la diffusion de l'information, la promotion de l'action collective, mais aussi pour l'efficacité d'un marché du travail. En fonction de l'objectif du réseau, certaines structures seront mieux adaptées que d'autres.

- Les comportements des agents peuvent facilement favoriser l'émergence de distributions de degré biaisées. Nous identifions ce phénomène pour la première fois au chapitre 2, où les apprenants sociaux apprennent également à identifier les agents auxquels ils doivent prêter attention. Nous trouvons que certains agents deviennent centraux dans le sens où ils rassemblent plus de personnes les sollicitant à chaque tour par rapport à d'autres. Et cela n'est pas nécessairement corrélé avec le fait qu'ils soient mieux informés que les autres.

Au chapitre 4, nous trouvons également une distribution de degré biaisée dans le modèle d'opinion, où les influenceurs essaient d'obtenir plus de connexions de la part de la population. Nous montrons que, sous des hypothèses réalistes, certains influenceurs réussiront beaucoup plus que d'autres. Créant à la fois une distribution biaisée dans l'ensemble de la population mais aussi à l'intérieur de la sous-population des influenceurs.

- L'apprentissage social est efficace. C'est la principale conclusion du chapitre 2. Dans une population où personne n'a de meilleures informations que les autres, il est plus sûr, et en moyenne aussi efficace, que l'apprentissage individuel. Ainsi, permettre aux gens d'interagir facilement entre eux, comme le ferait un réseau social en ligne, peut avoir

des avantages. Cela peut conduire les agents à mieux apprendre, à changer d'opinion, à rejoindre un mouvement collectif ou même à trouver des emplois plus efficacement.

- Mais lorsqu'ils sont manipulés, ces mêmes réseaux peuvent être préjudiciables à leurs utilisateurs. C'est le cas au chapitre 3 lorsque nous introduisons des influenceurs dans le réseau d'interactions, ou lorsque la plateforme qui régule le réseau social en ligne utilise le mauvais algorithme aux chapitres 4 et 5.

Dans l'ensemble, nos travaux valident l'idée de co-influence entre réseaux et comportements, et approfondissent les mécanismes précis en jeu. Nous nous sommes concentrés sur l'apprentissage dans le contexte des interactions sociales et avons montré comment les interactions sociales modifient ce processus, en mettant l'accent sur le rôle des réseaux sociaux en ligne qui, comme nous l'avons montré, modifient grandement la dynamique de l'apprentissage social.

Bibliography

- Acemoglu, D., Dahleh, M. A., Lobel, I., and Ozdaglar, A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236.
- Agarwal, S., Qian, W., and Zou, X. (2021). Thy neighbor’s misfortune: Peer effect on consumption. *American Economic Journal: Economic Policy*, 13(2):1–25.
- Akesson, J., Hahn, R. W., Metcalfe, R. D., and Monti-Nussbaum, M. (2023). The impact of fake reviews on demand and welfare. Technical report, National Bureau of Economic Research.
- Andris, C., Lee, D., Hamilton, M. J., Martino, M., Gunning, C. E., and Selden, J. A. (2015). The rise of partisanship and super-cooperators in the us house of representatives. *PloS one*, 10(4):e0123507.
- Aparicio, S., Villazón-Terrazas, J., and Álvarez, G. (2015). A model for scale-free networks: application to twitter. *Entropy*, 17(8):5848–5867.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394):116–131.
- Arthur, W. B. (1991). Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *The American economic review*, 81(2):353–359.
- Arthur, W. B. (1993). On designing economic agents that behave like human agents. *Journal of Evolutionary Economics*, 3(1):1–22.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406–411.
- Arthur, W. B. (2009). Complexity and the economy. In *Handbook of Research on Complexity*. Edward Elgar Publishing.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of conflict resolution*, 41(2):203–226.

- Axelrod, R., Daymude, J. J., and Forrest, S. (2021). Preventing extreme polarization of political attitudes. *Proceedings of the National Academy of Sciences*, 118(50).
- Bailey, M., Johnston, D., Kuchler, T., Stroebel, J., and Wong, A. (2022). Peer effects in product adoption. *American Economic Journal: Applied Economics*, 14(3):488–526.
- Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The quarterly journal of economics*, 107(3):797–817.
- Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *science*, 286(5439):509–512.
- Barabási, A.-L. (2016). *Network Science*. Cambridge University Press.
- Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.
- Bernheim, B. D. (1994). A theory of conformity. *Journal of political Economy*, 102(5):841–877.
- Bhattacharya, S., Sinha, S., and Roy, S. (2020). Impact of structural properties on network structure for online social networks. *Procedia Computer Science*, 167:1200–1209.
- Bikhchandani, S., Hirshleifer, D., Tamuz, O., and Welch, I. (2021). Information cascades and social learning. Technical report, National Bureau of Economic Research.
- Blanchard, O. (2009). The state of macro. *Annu. Rev. Econ.*, 1(1):209–228.
- Bonnevie, E., Rosenberg, S. D., Kummeth, C., Goldbarg, J., Wartella, E., and Smyser, J. (2020). Using social media influencers to increase knowledge and positive attitudes toward the flu vaccine. *Plos one*, 15(10):e0240828.
- Bourdieu, P. (1979). *La Distinction*. Les Editions de Minuit.
- Bourdieu, P. (2021). *Forms of Capital: General Sociology*. Polity Press.
- Brenner, T. (2006). Agent learning representation: advice on modelling economic learning. *Handbook of computational economics*, 2:895–947.
- Bruns, A. (2019). *Are filter bubbles real?* John Wiley & Sons.
- Burgiel, A. et al. (2017). Social comparisons and their role in shaping contemporary consumers’ behaviors. *Annales Universitatis Mariae Curie-Skłodowska, Sectio H Oeconomia*, 51(2):47–57.

- Burt, R. S. (2018). Structural holes. In *Social stratification*, pages 659–663. Routledge.
- Bush, R. R. and Mosteller, F. (1955). *Stochastic models for learning*. John Wiley & Sons, Inc.
- Carvalho, J.-P. (2019). Religious clubs: The strategic role of religious identity. *Advances in the Economics of Religion*, pages 25–41.
- Center, P. R. (2014). Religion and electronic media: One-in-five americans share their faith online.
- Centola, D. (2010). The spread of behavior in an online social network experiment. *science*, 329(5996):1194–1197.
- Chapin, S. (2018). Who’s living in a ‘bubble’? *The New York Times*.
- Choi, E. J. (2023). Does the internet help the unemployed find jobs? *Information Economics and Policy*, 62:101017.
- Christakis, N. A. and Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4):370–379.
- Christakis, N. A. and Fowler, J. H. (2008). The collective dynamics of smoking in a large social network. *New England journal of medicine*, 358(21):2249–2258.
- Cinelli, M., Morales, G. D. F., Galeazzi, A., Quattrociocchi, W., and Starnini, M. (2020). Echo chambers on social media: A comparative analysis. *arXiv preprint arXiv:2004.09603*.
- Cinus, F., Minici, M., Monti, C., and Bonchi, F. (2021). The effect of people recommenders on echo chambers and polarization. *arXiv preprint arXiv:2112.00626*.
- Coleman, J. (1998). *Foundations of Social Theory*. The Belknap Press.
- Cowan, R. (1991). Tortoises and hares: choice among technologies of unknown merit. *The economic journal*, 101(407):801–814.
- Cowan, R., Cowan, W., and Peter Swann, G. (2004). Waves in consumption with interdependence among consumers. *Canadian Journal of Economics/Revue canadienne d’économique*, 37(1):149–177.
- Cowan, R., Cowan, W., and Swann, P. (1997). A model of demand with interactions among consumers. *International Journal of Industrial Organization*, 15(6):711–732.
- Cross, J. G. et al. (2008). A theory of adaptive economic behavior. *Cambridge Books*.
- David, P. A. (1985). Clio and the economics of qwerty. *The American economic review*, 75(2):332–337.

- De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2020). Consumption network effects. *The Review of Economic Studies*, 87(1):130–163.
- DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical association*, 69(345):118–121.
- Dunbar, R. I. (1992). Neocortex size as a constraint on group size in primates. *Journal of human evolution*, 22(6):469–493.
- Dunbar, R. I., Arnaboldi, V., Conti, M., and Passarella, A. (2015). The structure of online social networks mirrors those in the offline world. *Social networks*, 43:39–47.
- Duricic, T., Kowald, D., Schedl, M., and Lex, E. (2021). My friends also prefer diverse music: homophily and link prediction with user preferences for mainstream, novelty, and diversity in music. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 447–454.
- Ellison, G. and Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of political Economy*, 101(4):612–643.
- Ellison, G. and Fudenberg, D. (1995). Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, 110(1):93–125.
- Enigma (2023). France adopts pioneering law for influencers.
- Erdős, P., Rényi, A., et al. (1960). On the evolution of random graphs. *Publ. math. inst. hung. acad. sci*, 5(1):17–60.
- Erdos, P., Rényi, A., et al. (1960). On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60.
- Erev, I. and Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881.
- Esteban, J. and Schneider, G. (2008). Polarization and conflict: Theoretical and empirical issues.
- Esteban, J.-M. and Ray, D. (1994). On the measurement of polarization. *Econometrica: Journal of the Econometric Society*, pages 819–851.
- Etzioni, A. (2015). The moral effects of economic teaching. In *Sociological Forum*, volume 30, pages 228–233. Wiley Online Library.

- Euler, L. (1953). Leonhard euler and the königsberg bridges. *Scientific American*, 189(1):66–72.
- Euronews (2023). France has approved a law that targets influencers.
- Farrell, J. and Saloner, G. (1985). Standardization, compatibility, and innovation. *the RAND Journal of Economics*, pages 70–83.
- Farrell, J. and Saloner, G. (1986). Installed base and compatibility: Innovation, product pre-announcements, and predation. *The American economic review*, pages 940–955.
- Flache, A., Mäs, M., Feliciani, T., Chattoe-Brown, E., Deffuant, G., Huet, S., and Lorenz, J. (2017). Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, 20(4).
- Foley (2023). France’s new influencer law adopted june 1, 2023.
- Galton, F. (1907). Vox populi. *Nature*, 75(1949):450–451.
- Gasana, P. U. (2009). Relative status and interdependent effects in consumer behavior. *The Journal of Socio-Economics*, 38(1):52–59.
- Gee, L. K., Jones, J. J., Fariss, C. J., Burke, M., and Fowler, J. H. (2017). The paradox of weak ties in 55 countries. *Journal of Economic Behavior & Organization*, 133:362–372.
- Gibson, R., Bon, E., Darius, P., and Smyth, P. (2023). Are online political influencers accelerating democratic deconsolidation? *Media and Communication*, 11(3):175–186.
- Gigerenzer, G. and Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. MIT press.
- Golub, B. and Sadler, E. (2017). Learning in social networks. *Available at SSRN 2919146*.
- Goodwin, A., Joseff, K., and Woolley, S. C. (2020). Social media influencers and the 2020 u.s. election: Paying ‘regular people’ for digital campaign communication.
- Goyal, S. (2011). Learning in networks. In *Handbook of social economics*, volume 1, pages 679–727. Elsevier.
- Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443.
- Granovetter, M. and Soong, R. (1983). Threshold models of diffusion and collective behavior. *Journal of Mathematical sociology*, 9(3):165–179.
- Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, 78(6):1360–1380.

- Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., et al. (2023a). How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science*, 381(6656):398–404.
- Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., et al. (2023b). Reshares on social media amplify political news but do not detectably affect beliefs or opinions. *Science*, 381(6656):404–408.
- He, S., Hollenbeck, B., and Proserpio, D. (2022). The market for fake reviews. *Marketing Science*, 41(5):896–921.
- Hewlett, S. A., Marshall, M., and Sherbin, L. (2013). How diversity can drive innovation. *Harvard business review*, 91(12):30–30.
- Hosterman, A. R., Johnson, N. R., Stouffer, R., and Herring, S. (2018). Twitter, social support messages, and the #metoo movement. *The Journal of Social Media in Society*, 7(2):69–91.
- Howard, P. N., Duffy, A., Freelon, D., Hussain, M. M., Mari, W., and Maziad, M. (2011). Opening closed regimes: what was the role of social media during the arab spring? *Available at SSRN 2595096*.
- Intelligence, I. (2023). Threads is struggling to retain users - but it could still catch up to X. *The Verge*.
- Jackson, M. O. (2010). *Social and economic networks*. Princeton University Press.
- Javed, S., Rashidin, M. S., and Xiao, Y. (2022). Investigating the impact of digital influencers on consumer decision-making and content outreach: using dual aisas model. *Economic research-Ekonomska istraživanja*, 35(1):1183–1210.
- Johnson, C. M., Tariq, A., and Baker, T. L. (2018). From gucci to green bags: Conspicuous consumption as a signal for pro-social behavior. *Journal of Marketing Theory and Practice*, 26(4):339–356.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American economic review*, 93(5):1449–1475.
- Katz, E. and Lazarsfeld, P. F. (2017). *Personal influence: The part played by people in the flow of mass communications*. Routledge.
- Katz, E., Lazarsfeld, P. F., and Roper, E. (1955). *Personal influence: The part played by people in the flow of mass communications*. Routledge.

- Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *The American economic review*, 75(3):424–440.
- Keijzer, M. A. and Mäs, M. (2022). The complex link between filter bubbles and opinion polarization. *Data Science*, 5(2):139–166.
- Kirman, A. (1989). The intrinsic limits of modern economic theory: the emperor has no clothes. *The Economic Journal*, 99(395):126–139.
- Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of economic perspectives*, 6(2):117–136.
- Kirman, A. P. and Vriend, N. J. (2000). Learning to be loyal. a study of the marseille fish market. In *Interaction and Market structure*, pages 33–56. Springer.
- Kogan, S., Moskowitz, T. J., and Niessner, M. (2019). Fake news: Evidence from financial markets. *Available at SSRN*, 3237763.
- Kuhn, P. J. (2014). The internet as a labor market matchmaker. *IZA World of Labor*.
- Leibenstein, H. (1950). Bandwagon, snob, and veblen effects in the theory of consumers’ demand. *The quarterly journal of economics*, 64(2):183–207.
- Lorenz, T. (2021). To fight vaccine lies, authorities recruit an ‘influencer army’. *The New York Times*.
- Lu, Q. S. and Seah, Z. Y. (2018). Social media influencers and consumer online engagement management. In *Digital Marketing and Consumer Engagement: Concepts, Methodologies, Tools, and Applications*, pages 1398–1406. IGI Global.
- Macy, M. W. and Flache, A. (2002). Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences*, 99(suppl.3):7229–7236.
- Maes, M. and Bischofberger, L. (2015). Will the personalization of online social networks foster opinion polarization? *Available at SSRN* 2553436.
- Manikonda, L., Beigi, G., Liu, H., and Kambhampati, S. (2018). Twitter for sparking a movement, reddit for sharing the moment: #metoo through the lens of social media. *arXiv preprint arXiv:1803.08022*.
- Marshall, A. (1890). Principles of economics macmillan. *London (8th ed. Published in 1920)*.
- Matsa, K. E. (2023). More americans are getting news on tiktok, bucking the trend seen on most other social media sites. <https://pewresearch.org>.

- McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444.
- Megane, L., Pauline, M., and Amitabh, A. (2019). The online influencers strategies and their impact on consumers decision process. - . , 18(3):419–447.
- Milgram, S. (1967). The small world problem. *Psychology today*, 2(1):60–67.
- Mondaq (2023). France’s new law on influencer marketing.
- Myers, D. G. (1982). Polarizing effects of social interaction. *Group decision making*, 125:137–138.
- Nandagiri, V. and Philip, L. (2018). Impact of influencers from instagram and youtube on their followers. *International Journal of Multidisciplinary Research and Modern Education*, 4(1):61–65.
- Nyhan, B., Settle, J., Thorson, E., Wojcieszak, M., Barberá, P., Chen, A. Y., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., et al. (2023). Like-minded sources on facebook are prevalent but not polarizing. *Nature*, pages 1–8.
- Peng, H., Nematzadeh, A., Romero, D. M., and Ferrara, E. (2020). Network modularity controls the speed of information diffusion. *Physical Review E*, 102(5):052316.
- Peralta, A. F., Kertész, J., and Iñiguez, G. (2022). Opinion dynamics in social networks: From models to data. *arXiv preprint arXiv:2201.01322*.
- Perc, M. (2014). The matthew effect in empirical data. *Journal of The Royal Society Interface*, 11(98):20140378.
- Peres, R. (2014). The impact of network characteristics on the diffusion of innovations. *Physica A: Statistical Mechanics and Its Applications*, 402:330–343.
- Peter, C. and Muth, L. (2023). Social media influencers’ role in shaping political opinions and actions of young audiences. *Media and Communication*, 11(3).
- Puetz, K. (2015). Consumer culture, taste preferences, and social network formation. *Sociology Compass*, 9(6):438–449.
- Putnam, R. D. (2015). Bowling alone: America’s declining social capital. In *The city reader*, pages 188–196. Routledge.
- Riedl, M., Schwemmer, C., Ziewiecki, S., and Ross, L. M. (2021). The rise of political influencers—perspectives on a trend towards meaningful content. *Frontiers in Communication*, 6:752656.

- Rogers, E. M. et al. (1962). Diffusion of innovations. *Diffusion of innovations*.
- Schelling, T. C. (1969). Models of segregation. *The American economic review*, 59(2):488–493.
- Shearer, E. (2021). More than eight-in-ten americans get news from digital devices. *Pew Research Center*, 12.
- Shearer, E. and Mitchell, A. (2021). News use across social media platforms in 2020 pew research center’s journalism project.
- Shen, F., Xia, C., and Skoric, M. (2020). Examining the roles of social media and alternative media in social movement participation: A study of hong kong’s umbrella movement. *Telecommunications and Informatics*, 47:101303.
- Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118.
- Skinner, B. F. (1965). *Science and human behavior*. Number 92904. Simon and Schuster.
- Smith, A. (1937). *The wealth of nations [1776]*, volume 11937. na.
- Statista (2023). U.s. tiktok monthly active users 2023. <https://www.statista.com>.
- Stevenson, B. (2008). The internet and job search. Technical report, National Bureau of Economic Research.
- Thorndike, E. L. (1927). The law of effect. *The American journal of psychology*, 39(1/4):212–222.
- Török, J. and Kertész, J. (2017). Cascading collapse of online social networks. *Scientific reports*, 7(1):16743.
- TRAVERS, J. (1969). An experimental study of the small world problem. *Sociometry*, 32(4):425–443.
- Veblen, T. (2017). *The theory of the leisure class*. Routledge.
- Wang, T.-L. (2020). Does fake news matter to election outcomes?: The case study of taiwan’s 2018 local elections. *Asian Journal for Public Opinion Research*, 8(2):67–104.
- Wang, X., Aisihaer, N., and Aihemaiti, A. (2022). Research on the impact of live streaming marketing by online influencers on consumer purchasing intentions. *Frontiers in psychology*, 13:1021256.
- Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442.

Weber, M. and Kalberg, S. (2013). *The Protestant ethic and the spirit of capitalism*. Routledge.