



Grant Agreement number: 101132559

Project acronym: ST4TE

Project title: Strategies for just and equitable transitions in Europe

Type of action: RIA

D2.4 Report on Algorithmic Fairness and Inequality

Deliverable leader	CLIMAFIN
Authors	Antoine Mandel, Marco Duenas
Due date	01.04.2025
Actual submission date	17.03.2025, revised 07/04/2025
Dissemination Level	Public



ARISTOTLE
UNIVERSITY
OF THESSALONIKI



EFISCENTRE



University
of Ferrara



Maastricht
University



G S
S I

GRAN SASSO
SCIENCE INSTITUTE
SCHOOL OF ADVANCED STUDIES
Scuola Universitaria Superiore



CLIMAFIN



Utrecht
University



This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101132559.

Document Revision History

Date	Version	Author/Editor/Contributor	Summary of main changes / Status
17.03.2025	0.1	Antoine Mandel, Marco Duenas	
07.04.2025	0.9	Antoine Mandel, Marco Duenas	Refactored layout. Introduced a new section on the impacts on socio-economic dynamics in order to clarify relationships with WP4.
15.04.2025	1.0	Antoine Mandel, Marco Duenas	Improved layout, submitted version

Disclaimer

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

Copyright

This document may not be copied, reproduced, or modified in whole or in part for any purpose without written permission from the ST4TE Consortium. In addition, an acknowledgement of the authors of the document and all applicable portions of the copyright notice must be clearly referenced.

All rights reserved.

This document may change without notice.



This project has received funding from the European Union's Horizon Europe under grant agreement No 101132559.

Executive summary

This deliverable contributes to work-package 2 “*Assessing the impact of green, digital and twin transition on inequalities*”. It focuses on objective 3 “*identifying new forms of (social) inequality that could emerge from the twin transition*” and more precisely on inequalities in terms of social capital that could be induced by recommendation algorithms.

Specifically, it reports on the work performed in task 2.4 “Algorithmic fairness and inequality” and how it will feed in task 4.5. The core motivation for task 2.4 is the remark that, in the context of the digital transition, recommendation algorithms have become a key determinant of the formation of socio-economic relationships and, more broadly, of the social structure. This naturally raises the question of the impacts of these algorithms on inequalities in terms of social capital/status/centrality. To address this question, task 2.4 has developed analytical and computational models of social network formation processes induced by recommendation algorithms. These models allow to characterize how recommendation algorithms determine the structure of online social networks and the associated distribution of centrality.

Therefrom, one can investigate the impact of recommendation algorithms on inequalities in terms of access to information social capital, and more broadly on social interactions. The models also allow to define systemic measures of algorithmic fairness and to investigate the potential trade-offs between fairness and efficiency for online platforms. We find a strong relationship between the structure of users’ preferences and the properties of the algorithms. When preferences are strongly correlated, the more hierarchical algorithms perform better as they rapidly identify the objects/followees collectively preferred by users. When preferences are negatively correlated, uniform recommendation algorithms perform as well (or even better) than hierarchical algorithms as they allow for an efficient exploration of the network. In terms of fairness, we find that the uniform recommendation algorithm is always perfectly fair. For low levels of correlation between preferences, hierarchical algorithms do not entail substantial divergence from this fair benchmark because no object/followee can sustain a high degree. Likewise, for high levels of correlation, unfairness is limited because the set of efficient networks is small and rather homogeneous as most agents seek to form links with the few “star” objects. Unfairness reaches a maximum for intermediary level of correlation because this is a domain where (i) central nodes can emerge and be identified by hierarchical algorithms, (ii) the set of efficient networks is sufficiently large and heterogeneous for divergence with the uniform benchmark to materialize.

Our results hence highlight the presence of a trade-off between efficiency and fairness (in particular for intermediate level of preference correlation), as the hierarchical recommendation algorithms that ensure fast convergence to efficient networks are also those that lead to high level of unfairness. However, our analysis also hints at a simple solution to mitigate this trade-off. The designer of an online platform can adapt the recommendation algorithm to the different phases of the network formation process. Hierarchical algorithms can be used in the early phase of the process to reach rapidly the set of efficient networks and ensure user satisfaction and adoption. Once the user base is established, more uniform recommendation algorithms can be used to further explore the set of potential networks and guarantee fairness without paying an efficiency cost. There are evidences of platform behaviour being aligned with our conclusions in the sense that recommendation algorithms have been updated to foster a more extensive exploration of the network and the recommendation of more diverse content. Notably, a major wave of such algorithmic updates has occurred in 2024.

Beyond fairness “per se”, recommendation algorithms can also have substantial impacts on socio-economic dynamics. First, they induce changes on the structure of social networks. Uneven distributions of centralities in social networks can increase inequality in terms of access

to information, social capital, or symbolic power. Second, network structure is a key determinant of technological diffusion and adoption. It is thus key for the success of the twin transition to have network structures aligned for rapid and efficient technological diffusion. Third, social networks are key in the formation of public opinion and in its potential polarization. A successful twin transition requires social network structures that are conducive to consensus formation rather than polarization. Overall, it appears that some form of monitoring and/or regulation of recommendation algorithms is warranted. A first step in this direction is the requirement to make algorithms available to regulators (and ideally to users as well) in order to evaluate their performance in terms of fairness and their potential systemic bias. A second step would be the actual regulation of the algorithms. Our results suggest that a potential route would be to impose to recommendation algorithms specifications that express fairness requirements. This approach would also avoid the perils of uniformization by regulation.

Contents

1	Introduction	4
2	Mathematical and computational analysis of recommendation algorithms	9
3	Measuring efficiency of recommendation algorithms	14
4	Invariant distributions of networks	17
5	Measuring fairness of recommendation algorithms	22
6	Discussion and Conclusion	26

1 Introduction

1.1 Context and motivation

With the rise of electronic commerce and online social networks, recommender systems have become key drivers of socio-economic interactions in a wide range of domains such as the formation of commercial relationships [Dinerstein et al., 2018, Johnson et al., 2023], the organisation of the labor market [Komiya and Noda, 2024], or the dynamics of opinions [Santos et al., 2021, Cinus et al., 2022]. This structural role of recommendation algorithms leads to the emergence of ethical and social standards for companies deploying these algorithms [Redman, 2019] and raises the question of algorithmic regulation in the public debate [Makridis and Teodorescu, 2024]. Specifically, algorithmic fairness has emerged as a key instance of the debate on global inequalities in the online realm [Kleinberg et al., 2018, Barocas et al., 2023].

The existing literature has mostly focused on fairness in the context of decisions based on the predictions of statistical and machine learning models [Mitchell et al., 2021, Wang et al., 2023]. Specifically, it has investigated *what fairness might mean in the context of decisions based on the predictions of statistical and machine learning models* [Mitchell et al., 2021]. Yet, existing analysis provide an incomplete perspective by focusing mostly on the outcomes of single instances of recommendation and by adopting a purely statistical approach that has difficulties shading light on the algorithmic drivers of these discriminations.

Against this background, there is increasing concern about systemic bias in recommendation algorithms. First, an emerging issue in the public debate is the fact that the recommendation algorithms might hinder access to information (see Figure 1). Second, major episodes of manipulations of recommender systems have been recorded recently, in particular during the last Romanian presidential election. In the regulation sphere, the European Commission has recently sent requests for information to YouTube, Snapchat, and TikTok on recommender systems under the Digital Services Act¹. The commission is concerned by the potential role of recommendation algorithms in amplifying systemic risks “including those related to the electoral process and civic discourse, users’ mental well-being (e.g. addictive behaviour and content “rabbit holes”), and the protection of minors”. Rather than statistical information, the commission is asking the platforms about the parametrization of recommendation algorithms. This highlights the shift from a regulation based on the statistical properties of algorithms towards a more proactive perspective focusing on the design and the functioning of the algorithms.

Accordingly, our analysis focuses on the impacts of recommendation algorithms on the formation of social networks. This allows first to analyse algorithmic fairness from a systemic perspective. Beyond fairness “per se”, recommendation algorithms can also have substantial impacts on socio-economic dynamics. First, they induce changes on the structure of social networks. Uneven distributions of centralities in social networks can increase inequality in terms of access to information, social capital, or symbolic power. Second, network structure is a key determinant of technological diffusion and adoption. It is thus key for the success of the twin transition to have network structures aligned for rapid and efficient technological diffusion. Third, social networks are key in the formation of public opinion and in its potential polarization. A successful twin transition requires social network structures that are conducive to consensus formation rather than polarization.

¹<https://digital-strategy.ec.europa.eu/en/news/commission-sends-requests-information-youtube-snapchat-and-tiktok-recommender-systems-under-digital>

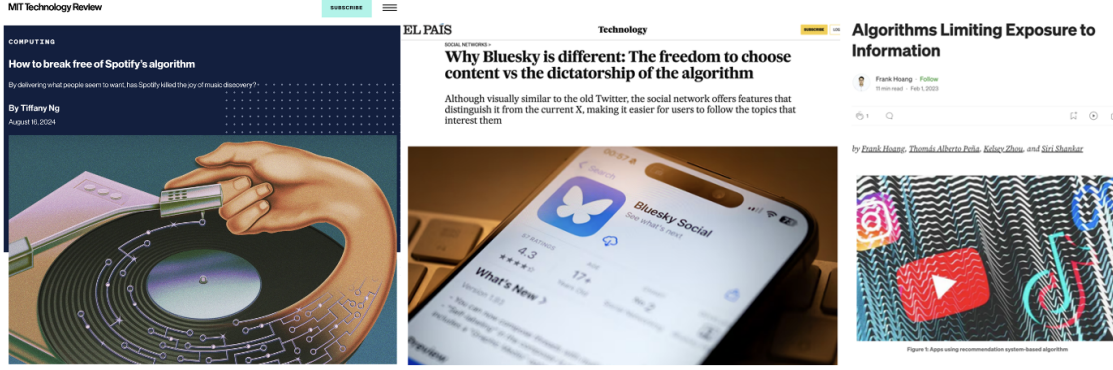


Figure 1: Illustration, through three recent media pieces, of the raising concern about recommendation algorithms hindering access to information. See respectively <https://medium.com/@frankhoang161/algorithms-limiting-exposure-to-information-16af4cfe76c>, <https://www.technologyreview.com/2024/08/16/1096276/spotify-algorithms-music-discovery-ux/> and <https://english.elpais.com/technology/2024-11-23/why-bluesky-is-different-the-freedom-to-choose-content-vs-the-dictatorship-of-the-algorithm.html>

1.2 Our approach

In order to contribute to the understanding of the systemic properties of recommendation algorithms, with a specific focus on their properties in terms of fairness and, we aim to develop a model of the long-term evolution of economic networks governed by recommendation algorithms, define measures of fairness that account for long-term and structural impacts, identify fairness-efficiency trade-offs in this context, and provide algorithmic solutions to overcome them.

More precisely, we introduce a dynamic model of network formation driven by link recommendation algorithms and utility-based individual choices. We consider a bipartite network where agents link towards certain objects. This framework encompasses the case of a social networks where the set of objects is the set of agents. Each agent has a preference over potential links, as well as a capacity of attention corresponding to the maximal number of links that it can form. The network is controlled by a platform which, at every time step, recommends a link to an agent. Assuming that link acceptance is utility-based, this recommendation algorithm induces a Markov process over the set of networks. If every link has a positive probability to be recommended, the process converges towards the class of efficient networks where each agent maximizes its utility, i.e., is linked to its preferred followees.

In this setting, we measure the efficiency of a recommendation algorithm, from the point of view of the platform and its users, via the speed at which it reaches the recurrent class of efficient networks. We then measure fairness, from the point of view of objects/followees, via the properties of the stationary distribution of networks. Namely, by considering the link allocation problem as a bankruptcy problem [see Aumann and Maschler, 1985], we provide micro-foundations for measuring the fairness of a recommendation algorithm through the entropy of its invariant distribution, or more generally through the Kullback-Leibler divergence between this distribution and a fair benchmark.

In order to operationalise these definitions, we provide analytical characterizations of the hitting time to the recurrent class and of the invariant distribution of the Markov chain associated to a recommendation algorithm. We further provide an exact algorithm for the computation of the invariant distribution when the associated Markov chain is time-reversible. These results

allow us to investigate precisely the fairness efficiency trade-off for link recommendation algorithms. We characterize algorithms in terms of the strength of their hierarchical features, i.e., strength of the attachment parameter for preferential attachment type of algorithms or length of the sequence of recommendations for follower of follower type of algorithms. We find a strong relationship between the structure of users' preferences and the properties of the algorithms. When preferences are strongly correlated, the more hierarchical algorithms perform better as they rapidly identify the objects/followees collectively preferred by users. When preferences are negatively correlated, uniform recommendation algorithms perform as well (or even better) than hierarchical algorithms as they allow for an efficient exploration of the network. In terms of fairness, we find that the uniform recommendation algorithm is always perfectly fair. For low levels of correlation between preferences, hierarchical algorithms do not entail substantial divergence from this fair benchmark because no object/follower can sustain a high degree. Likewise, for high levels of correlation, unfairness is limited because the set of efficient networks is small and rather homogeneous as most agents seek to form links with the few "star" objects. Unfairness reaches a maximum for intermediary level of correlation because this is a domain where (i) central nodes can emerge and be identified by hierarchical algorithms, (ii) the set of efficient networks is sufficiently large and heterogeneous for divergence with the uniform benchmark to materialize.

Our results hence highlight the presence of a trade-off between efficiency and fairness (in particular for intermediate level of preference correlation), as the hierarchical recommendation algorithms that ensure fast convergence to efficient networks are also those that lead to high level of unfairness. Our analysis also hints at a simple solution to mitigate this trade-off. It suffices to update the recommendation algorithm once the class of efficient networks is reached. More precisely, the designer of an online platform can adapt the recommendation algorithm to the different phases of the network formation process. Hierarchical algorithms can be used in the early phase of the process to reach rapidly the set of efficient networks and ensure user satisfaction and adoption. Once the user base is established, more uniform recommendation algorithms can be used to further explore the set of potential networks and guarantee fairness without paying an efficiency cost. This can notably be implemented through the Metropolis process, which will have the effect to obtain a uniform stationary distribution, hence ensuring perfect fairness.

We believe that our results and methodologies can have a strong impact on the design of platforms recommending links in a social network. Apart from giving analytical results on the convergence and stationary distribution on the states of the networks, the paper suggests a general methodology to achieve both rapid convergence and fairness.

1.3 Relation with the literature

Our results relate to two main streams of literature. On the one hand, the expanding literature that aims to quantify the fairness of algorithms used in socio-economic applications. On the other hand, the theoretical literature that investigates network formation processes and their social efficiency.

As for the algorithmic fairness literature, it has mostly focused on "*what fairness might mean in the context of decisions based on the predictions of statistical and machine learning models.*" [Mitchell et al., 2021]. Numerous surveys [see e.g. Caton and Haas, 2020, Barocas et al., 2023] are available on the matter. A first strand of literature has focused on the potential emergence of algorithmic biases against certain groups in relation to attributes such as gender, age, or racial origin [Kleinberg et al., 2018]. This has led to a number of domain-specific analysis, investigating algorithmic discrimination in contexts such as labor markets [Lambrecht and Tucker,

2019, Komiyama and Noda, 2024], online matching markets [Ma et al., 2021, 2023], credit scoring [Bono et al., 2021, Hurlin et al., 2024], criminal justice [Berk et al., 2017], or healthcare decisions [Yang et al., 2023]. Our approach differs from these analysis as the source of discrimination we consider is endogenous, being related to network characteristics, and as we analyze fairness from the point of view of followers/suppliers rather than from that of users/applicants.

From this latter perspective, our contribution is closely related to a number of contributions that have analyzed algorithmic fairness from the point of view of sellers in online advertising markets and marketplaces [see e.g. Cao et al., 2024, Bateni et al., 2022]. Two common objectives in this thread of literature are the definition of fairness metrics and the design of algorithms that overcome the fairness-efficiency trade-off. In particular, Lejeune and Turner [2019] investigate a Gini index based measure and study an optimization problem to maximize the spread of impressions across targeted audience segments. Balseiro et al. [2021] define a fairness measure based on a nonlinear regularizer and propose an online resource-allocation algorithm, with the aim of maximizing efficiency and fairness subject to a resource constraint. Bateni et al. [2022] study a setting where an online platform dynamically allocates a collection of goods (e.g. advertising slots) to budgeted buyers (e.g. advertisers) and considers a weighted proportional fairness metric. Li et al. [2024] investigate how the value from advertising in a dynamic allocation problem can be adjusted by a general fairness metric. Our approach shares the focus of the online advertising literature on the measure of fairness and the mitigation of the fairness vs efficiency trade-off. However, our representation of recommendation algorithms as stochastic processes provides a different perspective. It leads to a representation of the evolution of the system as a Markov chain, to a measure of efficiency through stopping time, and to a measure of fairness through the deviation from a benchmark distribution. Furthermore, we approach the problem from a systemic perspective where the utility of agents/followers, the efficiency of the platform, and the fairness to the objects/followers are linked through a network structure.

With respect to the network formation literature [see e.g. Jackson, 2008, Bramoullé et al., 2016], the main innovation of our work is to consider a recommendation algorithm as the central institution coordinating the formation of links, whereas the existing literature mostly focuses on fully decentralised interactions. From this perspective, our approach relates to the matching literature [Roth, 2015], although the latter focuses on the design of optimal algorithms rather than on the analysis of empirical ones. Also, our algorithmic approach allows to represent explicitly the network formation process, and to analyze its efficiency, whereas the existing literature focuses mostly on equilibrium networks in a game-theoretic sense. Furthermore, most of the existing contributions focus on social welfare, measured as the sum of individual utilities at equilibrium, rather than on fairness or equity. A notable exception is Navarro [2014] that puts forward a network formation processes that is both stable in the sense of Jackson and Wolinsky [1996] and fair in the sense of Myerson [1977].

Finally, its network perspective relates our work to the literature on link recommendation algorithms [e.g. Ferrara et al., 2022, Santos et al., 2021, Cinus et al., 2022]. In this setting, unfairness has been linked to the bias of algorithms towards extreme content/actors and its impact on the polarization of users and of the public debate [Vaidhyanathan, 2018, Eisenstat, 2021, Grabisch et al., 2023]. Our work draws inspiration from the regulatory measures that have been proposed in this framework [Ghosh, 2021] such as modifying the structure of interactions [Fagan, 2018], increasing platform liability for user-generated content, requiring algorithmic transparency, or fostering platform self-regulation [Ghosh, 2020].

1.4 Structure of the report

The rest of this report is structured as follows. Sections 2 to 5 describes our analytical model (that is also reported in the companion working paper [see Grabisch et al., 2024]). Section 3 introduces measures of efficiency for recommendation algorithms. Section 4 investigates the structural properties of a set of standard recommendation algorithms. Section 5 provides a micro-founded measure of fairness for recommendation algorithms and analyze the trade-off between efficiency and fairness. Section 6 discusses the policy implications of our results and concludes.

2 Mathematical and computational analysis of recommendation algorithms

2.1 Structure of interactions and recommendation algorithm

We consider the formation of bipartite networks between a set N of n agents and a set M of m objects. The set of potential links between agents and objects is $L = N \times M$. Accordingly, networks are represented by adjacency matrices $W \in \{0, 1\}^{n \times m}$ where $W_{i,j} = 1$ if there is a directed link between agent i and object j . We denote by $\mathcal{W}$ the set of all such adjacency matrices. Given a network W , for each agent $i \in N$, we denote by $M_i(W) := \{j \in M \mid W_{i,j} = 1\}$ the set of objects to whom it is connected, also referred to as its set of “followees”. Accordingly, for each object $j \in M$, we denote by $N_j(W) := \{i \in N \mid W_{i,j} = 1\}$ its set of “followers”. The cardinal of the set of followees of i is denoted by $d_i^+(W)$ and referred to as the out-degree of agent i while the cardinal of the set of followers of j is denoted by $d_j^-(W)$ and referred to as the in-degree of object j .

Remark 1. A particular case of interest subsumed in our framework is that of social networks where the objects are the agents themselves, i.e., $N = M$.

In this setting, we consider network formation processes driven by the interplay between individual choices and recommendation algorithms. Individual characteristics are given by a budget/capacity of attention and a preference ordering over objects. The capacity of attention of agent i is defined as the maximal number of active links $\bar{d}_i$ that it can maintain. The ordering over objects is represented by a utility function $u_i : M \rightarrow \mathbb{R}_+$ where for $j, k \in M$ one has $u_i(j) \geq u_i(k)$ if i prefers object j to object k . Assuming additivity, this utility over objects induces a utility over networks for i , $U_i : \mathcal{W} \rightarrow \mathbb{R}_+$, defined as the sum of the utility of the followees of i , i.e.,

$$U_i(W) = \sum_{j \in M} W_{i,j} u_i(j). \quad (1)$$

The recommendation algorithm characterizes the organization/institution (typically an online platform) that steers network formation. It suggests to agents with some probability a link in $L = N \times M$. Formally, a recommendation algorithm is a transition probability from the set of networks $\mathcal{W}$ to the set of links L , i.e., a matrix $Q \in \mathbb{R}^{\mathcal{W} \times L}$ so that, given the state of the network $W \in \mathcal{W}$, $Q(W, (i, j))$ is the probability that the link (i, j) is recommended. In particular, one has $\sum_{(i,j) \in L} Q(W, (i, j)) = 1$ for all $W \in \mathcal{W}$ and $Q(W, (i, j)) \geq 0$ for all $(i, j) \in L$. The evolution of the network is then determined by a sequential process of link recommendation by the algorithm and link formation by the agents. We assume that this process is characterized by the following principles.

Definition 1. A network formation process governed by a recommendation algorithm is such that:

- (1) Only recommended links can be formed.
- (2) Agents only accept to form links that (weakly) increase their utility.
- (3) The number of links that an agent can form is bounded by its capacity of attention.

Formally, the network formation process associated to a recommendation algorithm Q and a link capacity $\bar{d} = (\bar{d}_1, \dots, \bar{d}_n)$ is a Markov chain on $\mathcal{W}$ with transition matrix $P_{(Q, \bar{d})}$ such that for

$W' \neq W$:

$$P_{(Q, \bar{d})}(W, W') = \begin{cases} Q(W, (i, j)) & \text{if } W' = W + V_{i,j}, U_i(W') \geq U_i(W) \text{ and } d_i^+(W) < \bar{d}_i \\ \frac{Q(W, (i, j))}{|\{\ell \mid U_i(W + V_{i,j} - V_{i,\ell}) \geq U_i(W)\}|} & \text{if } W' = W + V_{i,j} - V_{i,k}, U_i(W') \geq U_i(W) \text{ and } d_i^+(W) = \bar{d}_i \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $V_{i,j}$ denotes the matrix for which (i, j) is the only non-zero coefficient, with value 1.

Hence, we assume that an agent who has not reached his maximal linkage capacity forms every desirable link that is recommended (first line of the definition of P_Q) whereas an agent that has reached its maximal linkage capacity forms a recommended link by substituting a lower utility link chosen uniformly at random among the existing ones (second line of the definition of P_Q). The probability that the network does not change following a recommendation is implicitly defined by complementarity as $P_Q(W, W) = 1 - \sum_{W' \neq W} P_Q(W, W')$. Our analytical specification of the network formation process also implicitly assumes that a single link is recommended (and potentially formed) every period.

In the following, for sake of simplicity, we will denote $P_{(Q, \bar{d})}$ by P_Q whenever there is no risk of ambiguity.

2.2 Examples of recommendation algorithms

Our modelling framework allows to represent all Markovian recommendation algorithms, i.e., algorithms where the recommendation only depends on the current state of the network. That is actually the case for most algorithms implemented in online platforms, which generally derive their recommendations on the basis of network metrics that are functions of the current state of the network. We highlight below a few algorithms of theoretical and/or of operational relevance. A more comprehensive list of algorithms can be found, e.g., in Ferrara et al. [2022].

Uniform recommendation We first consider, as a benchmark, the uniform recommendation algorithm where each link is recommended with an equal probability independently of the structure of the network. With L denoting the set of all possible links, the uniform recommendation algorithm is characterized by the transition matrix $U \in \mathbb{R}_+^{W \times L}$ such that for all $(W, (i, j)) \in \mathcal{W} \times L$, one has

$$U(W, (i, j)) = 1/|L| := 1/nm. \quad (3)$$

In case $N = M$, loops are excluded so that $|L| = n(n - 1)$. (todo really ?)

Preferential attachment We then consider a degree-based recommendation algorithm, where nodes with higher in-degree have a larger probability to be recommended. Namely, we consider the preferential attachment (PA) algorithm² defined for every $(W, (i, j)) \in \mathcal{W} \times L$ by:

$$PA(W, (i, j)) = \frac{\epsilon + d_j^-(W)}{m \sum_{h \in M} (\epsilon + d_h^-(W))}. \quad (4)$$

where $\epsilon > 0$ is a constant.

²The probability of recommendation is set proportional to $\epsilon + d_j^-(W)$ rather than to $d_j^-(W)$ to ensure that the algorithm has full support (see Assumption A below). We consider $\epsilon = 1$ by default in the following.

We further consider extensions of the preferential attachment algorithms where the strength of the degree-based recommendation is specified by a function f of the degree. Namely, we call “general preferential attachment” any recommendation algorithm (PA_f) of the following type:

$$\text{PA}_f(W, (i, j)) = \frac{f(d_j^-(W))}{m \sum_{h \in M} f(d_h^-(W))} \quad (5)$$

where $f : \mathbb{N} \rightarrow \mathbb{R}_{++}$ is an increasing function of the degree, which satisfies $f(d) > 0$ for every integer $d \geq 0$. The standard PA algorithm corresponds to $f(d) = 1 + d$, but other functions are natural, for example $f(d) = 1 + d^\alpha$ with $\alpha > 0$. A value greater (resp., smaller) than 1 for α indicates a stronger (resp., weaker) preferential attachment effect.

2-hops In the case of social networks, i.e., $M = N$, the followee of followee algorithm recommends to an agent a followee of one of its followees. It is thus similar to the 2-hops algorithm (see Ferrara et al. [2022]) in the sense that it recommends an agent that is two steps away from the agent receiving the recommendation. Formally, this corresponds to the transition probability given by:

$$\text{FF}(W, (i, j)) = \frac{W_{ij}^2}{\sum_{\substack{h, k \in N \\ h \neq k}} W_{hk}^2}. \quad (6)$$

We shall consider a slightly modified version in order to ensure that the algorithm has full support in the sense of Assumption A below:

$$\text{FF}(W, (i, j)) = \frac{\epsilon + W_{ij}^2}{\sum_{\substack{h, k \in N \\ h \neq k}} (\epsilon + W_{hk}^2)} \quad (7)$$

with $\epsilon > 0$.

We shall also consider n -hops variants of the algorithm that recommend the n -th followee, i.e., an agent at a distance of n in the followee network. Namely:

$$\text{FF}_n(W, (i, j)) = \frac{\epsilon + W_{ij}^n}{\sum_{\substack{h, k \in N \\ h \neq k}} (\epsilon + W_{hk}^n)} \quad (8)$$

with $\epsilon > 0$.

2.2.1 Asymptotic properties of the dynamics

The (random) network formation process induced by a recommendation algorithm Q is a Markov chain with set of states $\mathcal{W}_0 = \{W \in \mathcal{W} : d_i^+(W) \leq \bar{d}_i, \forall i \in N\}$ and transition matrix P_Q . Starting from an arbitrary (possibly empty) network, the algorithm sequentially proposes links to agents, which are accepted and included in the network if they represent a utility improvement for the agents. It is straightforward to remark from Definition 1 that the utility of agents is non-decreasing as the algorithm unfolds. Furthermore, if the algorithm is able to recommend the highest utility links, the network will eventually reach a maximal level of utility for each agent i , given its capacity $\bar{d}_i$ and its preferences. In other words, agent i will only form links with the set $E_i(u)$ of objects that yield maximal utility within capacity:

$$E_i(u) = \{j \in M : |\{k \in M : u_i(k) > u_i(j)\}| < \bar{d}_i\}. \quad (9)$$

We shall refer to $E_i(u)$ (E_i for short in absence of ambiguity about the preference profile) as the set of acceptable objects for agent i and call a link $\ell = (i, j)$ acceptable if $j \in E_i$. We shall denote by L_a the set of acceptable links and by $L_{na} := L \setminus L_a$ the set of non acceptable links. We define the set of efficient networks as those containing only acceptable links, that is

$$\mathcal{E} := \{W \in \mathcal{W}_0 : W_{i,j} > 0 \Rightarrow (i, j) \in L_a\}.$$

Given that L_a only contains links yielding maximal utility, the set of efficient networks can equivalently be defined as the set of networks maximizing total utility, that is:

$$\mathcal{E} = \operatorname{argmax}_{\mathcal{W}_0} \sum_{i \in N} U_i = \{W \in \mathcal{W}_0 : \sum_{i \in N} U_i(W) = \max_{W' \in \mathcal{W}_0} \sum_{i \in N} U_i(W')\}.$$

As mentioned above, the fact that the link acceptance process is utility based implies that the utility of agents is non-decreasing as the algorithm unfolds. If the algorithm furthermore recommends efficient links with some probability, the set of efficient networks will be reached, eventually. More precisely, let us introduce the following assumption:

Assumption A. For all $W \in \mathcal{W}_0$ and $\ell \in L_a$, one has $Q(W, \ell) > 0$.

Remark 2. Assumption A holds in particular if Q has full support in the sense that for all $W \in \mathcal{W}_0$ and each $\ell \in L$ one has $Q(W, \ell) > 0$.

Assumption A implies that $\mathcal{E}$ is the unique recurrent class of P_Q .

Lemma 1. If Q satisfies assumption A, then $\mathcal{E}$ is the unique recurrent class of P_Q , i.e., all the recurrent states of P_Q are such that each agent has maximal utility. In addition, $\mathcal{E}$ is aperiodic.

Standard results from the theory of Markov chains imply that the Markov chains induced by P_Q will reach $\mathcal{E}$ in a finite expected time and that the average time they spend in the network configuration W is given by the invariant distribution of P_Q , which has support in $\mathcal{E}$. Namely, one has the following proposition.

Proposition 1. Let $(X_t^Q)_{t \in \mathbb{N}}$ be a Markov chain with transition matrix P_Q , where Q is such that Assumption A holds. Then, for the underlying probability measure $\mathbb{P}$:

1. $\mathbb{E}(\inf\{t \in \mathbb{N} \mid X_t \in \mathcal{E}\}) < +\infty$ and in particular $\lim_{T \rightarrow +\infty} \mathbb{P}(X_T \in \mathcal{E}) = 1$.
2. There exists a unique probability distribution π_Q on $\mathcal{E}$ such that $P_Q \pi_Q = \pi_Q$ (*stationary or invariant* distribution). Furthermore, π_Q has full support in $\mathcal{E}$.
3. For all $W \in \mathcal{E}$, one has $\lim_{T \rightarrow +\infty} \frac{\mathbb{P}(\sum_{t=1}^T X_t^Q = W)}{T} = \pi_Q(W)$.

In other words, the random network formation process induced by the recommendation algorithm Q will lead almost surely to the set of efficient networks $\mathcal{E}$. Each agent's utility will increase monotonically during this convergence process. Furthermore, once the set of efficient networks $\mathcal{E}$ is reached, the algorithm alternates between networks configurations in $\mathcal{E}$, and the frequency of occurrence of a network configuration W is given by the invariant distribution π_Q , which is determined by the characteristics of the recommendation algorithm Q .

2.3 Structure of the class of efficient networks

The structure of the class of efficient networks depends on the preferences of agents and their linkage capacity. If $|E_i| = \bar{d}_i$, e.g., if the preference ordering of i is strict, the set of efficient connections for i is uniquely determined, and there exists a unique efficient network. In this setting, all recommendation algorithms are essentially equivalent and the analysis of their asymptotic properties is not relevant.

The situation of interest for us is rather when there are multiple efficient networks, i.e., where $|E_i| > \bar{d}_i$ for at least some i . In this setting, the set of acceptable objects/links can be partitioned between objects at and above minimal utility, respectively, i.e., $C_i := \{j \in M \mid u_i(j) = \min_{E_i} u_i\}$ and $E_i \setminus C_i = \{j \in M \mid u_i(j) > \min_{E_i} u_i\}$. On the one hand, as underlined in Example 1 below, links to objects in $E_i \setminus C_i$ are “undisputed”: they are present in every efficient network and their prevalence is thus independent of the recommendation algorithm. On the other hand, links in C_i are the actual object of choice influenced by the recommendation algorithm. If $|E_i| > \bar{d}_i$, all the links in C_i cannot simultaneously be present in an efficient network and the recommendation algorithm will determine their relative frequency of occurrence. Accordingly, the set of efficient networks can be described in terms of the sets C_i as follows.

$$\mathcal{E} = \{W \in \mathcal{W}_0 : \forall i \in N, M_i(W) \subseteq (E_i \setminus C_i) \cup K, K \subseteq C_i, |M_i(W)| = \bar{d}_i\}.$$

Example 1. Let us consider a network with 10 agents, and suppose that agent 1 can have 5 followees ($\bar{d}_1 = 5$). The utilities for agent 1 are given below:

$$u_1(2) = u_1(3) = 10, u_1(4) = u_1(5) = 9, u_1(6) = u_1(7) = u_1(8) = 8, u_1(9) = u_1(10) = 7.$$

Then $E_1 = \{2, 3, 4, 5, 6, 7, 8\}$. Observe that anyway links $(1, 2), (1, 3), (1, 4), (1, 5)$ will be formed, but then for the 5th link there is a choice between agents 6, 7 and 8, who have the same utility. Therefore, $C_1 = \{6, 7, 8\}$ and agents 9 and 10 will never be chosen. Remark that as agents 2, 3, 4 and 5 are chosen anyway, the value of their utility is not important for our analysis and can be assumed to be equal.

Lemma 1 tells us that $\mathcal{E}$ is the unique recurrent class of the Markov chain, which is acyclic. Let us focus on this class and restrict the transition matrix P_Q to $\mathcal{E}$, studying its essential properties. They are summarized in the following proposition.

Proposition 2. The unique recurrent class $\mathcal{E}$ has the following properties:

1. The total number of efficient networks is

$$|\mathcal{E}| = \prod_{i \in N} \binom{|C_i|}{\bar{d}_i - |E_i| + |C_i|}. \quad (10)$$

2. Let $W \in \mathcal{E}$ and consider the transition from W to $W + V_{i,j} - V_{i,k} \in \mathcal{E}$. It has probability:

$$P_Q(W, W + V_{i,j} - V_{i,k}) = \frac{Q(W, (i, j))}{|C_i \cap M_i(W)|}. \quad (11)$$

3. In (11), the denominator $|C_i \cap M_i(W)|$ does not depend on W .

In the following, our key concern is the selection properties of recommendation algorithms among efficient networks. Undisputed links are mostly irrelevant in this respect. In order to focus our analysis on actual choice situations, we shall consider in the following a stylized setting described by the following assumption³.

³The results in subsection 4 are nevertheless valid without Assumption B

Assumption B. The utility functions u_i are binary with values in $\{-1, 1\}$ and for all $i \in N$, one has $|\{j \in M \mid u_i(j) = 1\}| \geq \bar{d}_i$.

Under Assumption B, the set of objects is partitioned in a binary manner between satisfactory ones ($u_i(j) = 1$) and non-satisfactory ones ($u_i(j) = -1$). Consequently, one has for all $i \in N$, $E_i = \{j \in M \mid u_i(j) = 1\}$ and $C_i = E_i$. Hence, there are no undisputed links and recommendation algorithms fully control the formation of the network. It thus seems a natural framework to investigate the properties of these algorithms.

3 Measuring efficiency of recommendation algorithms

Following Proposition 1, any recommendation algorithm satisfying Assumption A will eventually converge towards the set of efficient networks. From this perspective, any recommendation algorithm is “efficient”. However, for impatient users, the waiting time before receiving suitable recommendations ought to be a key determinant of choice. Hence, in a setting of competition between online platforms [Zhang and Sarvary, 2011, Gelper et al., 2021], the speed at which the recommendation algorithm provides efficient recommendation is likely to be a key determinant of performance and comparative advantage. In our model, the speed at which an algorithm Q provides efficient recommendation can be measured at the system’s level via the expected time required to reach the recurrent class, i.e., the set of efficient networks. Indeed, this measures the expected time before the algorithm, or equivalently the platform, is able to provide acceptable recommendations to all its users.

In the following, we characterize the expected hitting time under the binary utility assumption (Assumption B). In this setting, the network formation process is fully controlled by the interplay between the recommendation algorithm and individual preferences. We further consider that the initial network is empty, which is a natural benchmark to ensure comparability between recommendation algorithm and preference profiles.

Assumption C. The initial network is empty, i.e., $\mathbb{P}(W(0) = \emptyset) = 1$.

In this setting, the expected hitting time is formally defined as

$$T^{Q,u} = \mathbb{E}(\inf\{t \in \mathbb{N} \mid W(t) \in \mathcal{E}\}) \quad (12)$$

where $W(t)$ is the state at time t of the Markov chain with transition matrix P_Q , and $\mathbb{E}$ is the expectation for the underlying probability measure. It is important to note that $\mathcal{E}$ depends on the utility profile $u = (u_1, \dots, u_n)$, hence the superscript u in the expected hitting time.

One can determine $T^{Q,u}$ as follows. Under Assumption B, only acceptable links will be formed. More precisely, the trajectories of a Markov chain with transition matrix P_Q can be described by the sequence of links $(\ell_t := (i_t, j_t))_{t \in \mathbb{N}} \in (L \cup \{\emptyset\})^{\mathbb{N}}$ that gets formed, with the convention that $\ell_t = \emptyset$ indicates that no link gets formed at step t . Under assumption B, non-acceptable links have negative utility and are never formed (in other words, one has $\ell_t = \emptyset$ whenever a non-acceptable link is proposed). A trajectory reaches the set of efficient networks when each agent i has formed $\bar{d}_i$ satisfactory links. Accordingly, let us define a *minimal hitting path* as a sequence of acceptable links of minimal length leading to an efficient network, namely:

Definition 2. A minimal hitting path is a sequence of links $\ell = (\ell_1, \dots, \ell_s) = ((i_1, j_1), \dots, (i_s, j_s)) \in L^s$ such that:

1. For each $\sigma \in \{1, \dots, s\}$, $j_\sigma \in E_{i_\sigma}$.

2. For each $i \in N$, $|\{\ell_\sigma = (i_\sigma, j_\sigma) \in \ell \mid j_\sigma \in E_{i_\sigma}\}| \geq \bar{d}_i$.

3. For any subpath ℓ' of ℓ , there exists $i \in N$ such that $|\{\ell_\sigma = (i_\sigma, j_\sigma) \in \ell' \mid j_\sigma \in E_{i_\sigma}\}| < \bar{d}_i$.

Furthermore, we shall say that a trajectory $\lambda = (\lambda_t)_{t \in \mathbb{N}}$ follows a minimal hitting path ℓ if ℓ is the sequence of links in $(\lambda_t)_{t \in \mathbb{N}}$ leading to $\mathcal{E}$, i.e., ℓ_s is exactly the s th acceptable link in λ . Finally, we shall denote by $V_\sigma(\ell)$ (or V_σ in absence of ambiguity), the network formed by the successive formation of the links ℓ_1 to ℓ_σ (noting that if an agent i is involved in more than $\bar{d}_i$ links, only the last $\bar{d}_i$ links are part of V_σ).

We shall denote by $\mathcal{M}(u)$ the set of minimal hitting paths, for a given preference profile u . As emphasized above, a minimal hitting path is a sequence of satisfactory links of minimal length leading to the set of efficient networks $\mathcal{E}$. Conversely, any trajectory of the Markov chain leading to $\mathcal{E}$ must follow a minimal hitting path. Building on the ergodic properties put forward in Proposition 1, one can then partition the trajectories of the Markov chain in terms of the minimal path they follow to reach the set of efficient networks and accordingly compute the expected hitting time as the sum of hitting times conditional on the minimal path followed. Namely, one has:

Lemma 2. The expected hitting time to $\mathcal{E}$ associated to a recommendation algorithm Q and a preference profile u is given by:

$$T^{Q,u} = \sum_{\ell \in \mathcal{M}(u)} P_\ell^Q T_\ell^Q$$

where P_ℓ^Q is the probability that the Markov chain follows the minimal hitting path ℓ and T_ℓ^Q is the expected time for hitting $\mathcal{E}$, conditional on the Markov chain following the minimal hitting path ℓ .

The probability of a minimal hitting path and the associated conditional expected hitting time can then be computed as follows.

Lemma 3. One has:

$$P_\ell^Q = \frac{Q(\emptyset, \ell_1)}{Q(\emptyset, L_a)} \times \frac{Q(V_1, \ell_2)}{Q(V_1, L_a \setminus V_1)} \times \cdots \times \frac{Q(V_{s-1}, \ell_s)}{Q(V_{s-1}, L_a \setminus V_{s-1})}$$

$$T_\ell^Q = \frac{1}{[Q(\emptyset, L_a)]^2} + \frac{1}{[Q(V_1, L_a \setminus V_1)]^2} + \cdots + \frac{1}{[Q(V_{s-1}, L_a \setminus V_{s-1})]^2},$$

where $Q(W, L')$ is the probability that Algorithm Q recommends a link in L' when in state W .

Lemma 3 highlights the determinants of the hitting time through a given minimal path. The probability that the path is followed depends on the rate at which the algorithm recommends the corresponding links, relatively to the rate at which it recommends acceptable links, i.e., how prominent for the algorithm the links constituting the paths are, as compared to the broader set of acceptable links. The conditional hitting time depends on the probability that acceptable links are recommended along the path. More specifically, the conditional waiting time is inversely proportional to the square of this probability.

Finally, the expected hitting time is obtained by summing over the set of minimal paths, i.e., over the paths that the algorithm can follow to reach an efficient network.

Proposition 3. The expected hitting time to the class of efficient networks $\mathcal{E}$ for a recommendation algorithm Q and a preference profile u is given by

$$T^{Q,u} = \sum_{\ell \in \mathcal{M}(u)} \left(\frac{Q(\emptyset, \ell_1) Q(V_1, \ell_2) \cdots Q(V_{s-1}, \ell_s)}{Q(\emptyset, L_a) Q(V_1, L_a \setminus V_1) \cdots Q(V_{s-1}, L_a \setminus V_{s-1})} \right) \times$$

$$\left(\frac{1}{[Q(\emptyset, L_a)]^2} + \frac{1}{[Q(V_1, L_a \setminus V_1)]^2} + \cdots + \frac{1}{[Q(V_{s-1}, L_a \setminus V_{s-1})]^2} \right).$$

Proposition 3 highlights that the efficiency of the algorithm, i.e., the expected hitting time, essentially depends on the ability of the algorithm to identify acceptable links. In turn, the set of acceptable links is determined by individual preferences. From this perspective, the recommendation algorithm oughts to explore efficiently the preference profile. To do so, the ability of the recommendation algorithm might depend on the structure of preferences. We first highlight this dependency through two examples.

Example 2. Let us first consider a case where for all $i \in N$, $\bar{d}_i = 1$ and the preferences of each agent are distinct and thus negatively correlated. More precisely, we consider the preference profile $\underline{u}$ such that $C_i = \{\pi(i)\}$ where π is a permutation of N that leaves no element invariant ($\pi(i) \neq i$). In this setting, a minimal path is of the form $\ell := (\ell_1, \dots, \ell_n) = ((i_1, \pi(i_1)), \dots, (i_n, \pi(i_n)))$ where each $i \in N$ appears exactly once. By symmetry, each such minimal path has the same probability provided the recommendation algorithm is anonymous. Furthermore, one has for each $s \in \{1, \dots, n-1\}$:

$$\text{PA}(V_{s-1}, L_a \setminus V_{s-1}) \leq \frac{1}{|L_a \setminus V_{s-1}|} = \text{U}(V_{s-1}, L_a \setminus V_{s-1}). \quad (13)$$

Indeed, links towards objects for which there does not exist a link in V_{s-1} are less likely to be recommended by $\text{PA}(V_{s-1}, \cdot)$ because they are of degree zero. However these are the only remaining acceptable links given the structure of preferences. Following Equation (3), one then has for each $\ell \in \mathcal{M}(\underline{u})$, $T_\ell^{\text{PA}} > T_\ell^{\text{U}}$ and thus

$$T^{\text{PA}, \underline{u}} > T^{\text{U}, \underline{u}}. \quad (14)$$

Hence, the uniform recommendation algorithm is more performant than the preferential attachment algorithm for the uncorrelated preference profile $\underline{u}$. One can further remark that, in this setting, for a generalized preferential attachment algorithm PA_f , the faster f would increase, the worst would be the performance of the algorithm. On the contrary, a generalized preferential attachment algorithm PA_f based on a decreasing function of the degree would outperform uniform attachment.

Example 3. Let us then consider a case where for all $i \in N$, $\bar{d}_i = 1$ and agents have perfectly (positively) correlated preferences. More precisely, we consider the preference profile $\bar{u}$ such that for all $i \in N$, $C_i = \{1\}$ (where 1 is an arbitrarily chosen element in M). In this setting, the set of minimal paths is given by the possible permutations of the links $\{(i, 1) \mid i \in N\}$. By symmetry, each such minimal path has the same probability provided the recommendation algorithm is anonymous. The relative performance of algorithms is then inverse to that in Example 2. For each minimal path $\ell = (\ell_1, \dots, \ell_s)$, one clearly has for each subpath $(\ell_1, \dots, \ell_r)$ of ℓ

$$\text{PA}(V_{r-1}, L_a \setminus V_{r-1}) \geq \frac{1}{|L_a \setminus V_{r-1}|} = \text{U}(V_{r-1}, L_a \setminus V_{r-1}). \quad (15)$$

Indeed, links towards 1 are more likely to be recommended by $\text{PA}(V_{r-1}, \cdot)$ than by the uniform recommendation algorithm because 1 is the only object with positive degree in V_{r-1} . These are the only acceptable links given the structure of preferences.

Following Equation (3), one then has for each $\ell \in \mathcal{M}(\bar{u})$, $T_\ell^{\text{PA}} < T_\ell^{\text{U}}$ and thus

$$T^{\text{PA}, \bar{u}} < T^{\text{U}, \bar{u}}. \quad (16)$$

Hence, the uniform recommendation algorithm is less performant than the preferential attachment algorithm for the correlated preference profile $\bar{u}$. One can further remark that, in this setting, for a generalized preferential attachment algorithm PA_f , the faster f would increase, the best would be the performance of the algorithm.

Overall, Examples 2 3 highlight that the relative performance of recommendation algorithms depends on the structure of preferences. “Hierarchical” algorithms, such as preferential attachment perform well when preferences are correlated but uniform, “egalitarian”, algorithms might perform better when preferences are highly heterogeneous. In order to investigate more precisely the relationship between the structure of preferences and the efficiency of recommendation algorithms, we run a series of numerical simulations in which we investigate the performance of different algorithms for varying level of correlation between individual preferences. Namely, we generate samples of random preference profiles with varying level of correlation (following the procedure described in Appendix 6.5.8) and then simulate the network formation process for the uniform recommendation algorithm, the 2 hops algorithm, and the preferential attachment algorithm for which we vary the strength of the degree-based recommendation by considering variants PA_f with f of the form $f(x) = x^\alpha$ ($\alpha = 1$ corresponds to the standard PA algorithm and $\alpha = 0$ to the uniform recommendation case).

We report in Figure 2 the results of these numerical experiments for a range of parameter configurations. These numerical results generalize the conclusions drawn from Examples 2 and 3. For uncorrelated preference profiles, uniform recommendation algorithm performs systematically better than preferential attachment algorithm and the performance differential increases with the strength of the preferential attachment. As the level of preference correlation increases, the performance of preferential attachment algorithms gradually improves. The algorithms with weak preferential attachment (low α) are the most efficient for weakly correlated profiles and algorithms with strong preferential attachment (high α) progressively outperform them as the level of correlation increases. For fully correlated preference profiles, the standard preferential attachment algorithm ($\alpha = 1$) becomes the best performing algorithm and relative performance is determined by the strength of preferential attachment. One also observes a temporary decrease of performance for the preferential attachment algorithm for weakly correlated preference profile. These feature are robust to parameter change. They highlight clearly the relation between the structure of preferences and the performance of algorithms. More “hierarchical” algorithms of the preferential attachment type perform better for correlated/hierarchical preference profiles where some objects dominate the preference orderings. More “egalitarian” algorithms perform better for diverse/heterogeneous preference profiles.

Additional series of numerical experiments focusing on variants of the 2-hops algorithm further strengthen this conclusion about the relationships between performance of hierarchical recommendation algorithm and correlation of preferences. Namely, we have considered recommendation algorithms of the form n -hops (with $n = 2, 3, 5$). As n increases, the determinants of recommendation shift from local to global properties of the network. Indeed, neglecting the ϵ term, the n -hops algorithm recommends agent j to i as a function of the share of walks of length n leading from i to j . As n increases, the share of walks leading to j increasingly depends on the overall centrality of the agent. Specifically, as n tends towards infinity, the share of walks leading to j is the eigenvector centrality of j . Hence n -hops algorithms are increasingly hierarchical as they increasingly account for the global centrality features of recommended agents. The results reported in Figure 3 highlight that as the level of correlation of preferences increases, the relative performance of the n -hops algorithms improves, i.e., more hierarchical algorithms perform better for more correlated/hierarchical preference profiles.

4 Invariant distributions of networks

Once the Markov chain reaches the class of efficient networks, the recommendation algorithm governs the evolution of the network through this recurrent class and hence generates the




Figure 2: Hitting time to the recurrence class, averaged over 25 simulations, for a range of (preferential attachment) recommendation algorithms and parameter values: $n = 200, m = 200, a = 10$ in the upper-left panel, $n = 20, m = 20, a = 10$ in the upper-right panel, $n = 200, m = 400, a = 10$ in the lower-left panel, $n = 200, m = 200, a = 4$ in the lower-right panel

invariant distribution of the Markov chain over efficient networks. This invariant distribution determines in particular the systemic properties of the algorithm in terms of fairness, which we shall investigate in subsection 5. In this subsection, we provide a characterization of the invariant distribution for the uniform and preferential attachment algorithms. We then put forward an exact numerical method for the computation of these distributions.

4.1 Characterization of the invariant distribution

Uniform recommendation We first provide a characterization of the invariant distribution for the uniform recommendation algorithm. In that case, the transition probabilities between two efficient networks are perfectly symmetric. Accordingly, each efficient network is equally likely to emerge and the invariant distribution is uniform. Namely, one has:

Proposition 4. Under the uniform recommendation algorithm, the uniform distribution over $\mathcal{E}$ is the only stationary distribution of P_U .




Figure 3: Hitting time to the recurrence class, averaged over 25 simulations, for a range of (hops) recommendation algorithms and parameter values: $n = 200, m = 200, a = 10$ in the left panel, $n = 20, m = 20, a = 4$ in the right panel

Preferential attachment In the case of preferential attachment recommendation algorithms, there is a simple relationship between changes in the degree distribution and the transition probabilities between networks, which implies that the Markov chain P_{PA} is time-reversible and that the stationary distribution can be characterized simply by a detailed balanced condition.

Specifically, let us recall that a Markov chain with transition matrix P and set of states S is time-reversible if it has the same transition probabilities when running backwards in time, i.e., it satisfies the detailed balance equations (see, e.g., Levin et al. [2009]):

$$\pi_i p_{ij} = \pi_j p_{ji} \quad (s_i, s_j \in S). \quad (17)$$

Furthermore, an equivalent condition for time-reversibility is that the chain satisfies the Kolmogorov criterion [see Kelly, 2011], i.e., for any finite sequence of distinct states $s_{j_0}, \dots, s_{j_k}$, one has

$$p_{j_0 j_1} \cdots p_{j_{k-1} j_k} p_{j_k j_0} = p_{j_0 j_k} p_{j_k j_{k-1}} \cdots p_{j_1 j_0}. \quad (18)$$

In the case of preferential attachment, any such loop contains similar degree increments and decrements (and thus similar transition probabilities) in both temporal directions. This allows to prove that the Markov chain is time-reversible:

Proposition 5. Let f be a function defining a general preferential attachment algorithm PA_f .

1. The Markov chain P_{PA_f} is reversible, and its stationary distribution π_{PA_f} satisfies:

$$f(d_j^-(W)) \pi_{PA_f}(W) \sum_{h \in M} f(d_h^-(W + V_{i,j} - V_{i,k})) = f(d_k^-(W + V_{i,j} - V_{i,k})) \pi_{PA_f}(W + V_{i,j} - V_{i,k}) \sum_{h \in M} f(d_h^-(W)). \quad (19)$$

2. If $f(d) = 1 + d$ (classical PA), the expression simplifies to

$$f(d_j^-(W)) \pi_{PA_f}(W) = f(d_k^-(W + V_{i,j} - V_{i,k})) \pi_{PA_f}(W + V_{i,j} - V_{i,k}).$$

As exemplified in subsection 4.3 below, the reversibility of P_{PA_f} also allows to derive bounds on the eigenvalues and on the convergence time for the Markov chain [Diaconis and Stroock, 1991].

4.2 Computation of the invariant distribution

For the preferential attachment algorithm, and more generally for recommendation algorithms inducing a reversible Markov chain, the detailed balance condition (Equation (19)) can be used to compute step by step the stationary distribution provided one can find a path that passes through all the efficient networks. This latter condition amounts to ensure that the *adjacency graph* associated to the transition matrix is Hamiltonian. More precisely, the adjacency graph G_Q associated to the transition matrix P_Q (restricted to $\mathcal{E}$) is the directed graph whose set of nodes is $\mathcal{E}$ and such that there is an arc from state W to state W' if and only if $P_Q(W, W') > 0$. This graph is Hamiltonian if it admits a Hamiltonian path, i.e., a path passing through all nodes without repetition. It turns out that every recommendation algorithm, and preferential attachment in particular, induces the Hamiltonian property on the adjacency graph.

Proposition 6. For any recommendation algorithm, the adjacency graph G_Q of the transition matrix P_Q is Hamiltonian.

One can then compute the stationary distribution of networks for the preferential attachment algorithm, and more generally for recommendation algorithms inducing a reversible Markov chain, as follows.

1. Consider the adjacency graph G_{PA_f} of the transition matrix P_{PA_f} . We have proved in Proposition 6 that G_{PA_f} is Hamiltonian, i.e., there exists a path passing through all nodes without repetition.
2. Let us thus consider a Hamiltonian path for G_{PA_f} and denote the sequence of states (networks) in this path by $W_1, \dots, W_q$, where q is the number of efficient networks.
3. By construction, between W_i and W_{i+1} , only one link has been deleted and a new link has been added, for $i = 1, \dots, q - 1$. Therefore, $\pi_{PA_f}(W_{i+1})$ can be expressed in terms of $\pi_{PA_f}(W_i)$ through (19).
4. Letting $\pi'_{PA_f}(W_1) = 1$, we can compute $\pi'_{PA_f}(W_2), \dots, \pi'_{PA_f}(W_q)$ in this order.
5. The stationary distribution π_{PA_f} is then obtained by normalization:

$$\pi_{PA_f}(W_i) = \frac{\pi'_{PA_f}(W_i)}{\sum_{j=1}^q \pi'_{PA_f}(W_j)}.$$

4.3 A numerical example

We present in this subsection a detailed example of network with $N = M$, computing the transition matrices and find numerically the stationary distributions. Another similar and more complex example is shown in Appendix 6.6.2.

Example 4. We consider a setting where $M = N = 5$. Let us take 5 agents and assume that the following links of the network are those with maximal utility (i.e., they remain fixed): (1,2), (2,3), (3,5), (4,3), (5,1) Suppose that the sets C_i of possible choices are:

$$C_1 = \{3, 4\}, \quad C_2 = \{1, 4\}, \quad C_3 = \{1, 2\}, \quad C_4 = \{2, 5\}, \quad C_5 = \{2, 3\}$$

and the out-degree of each agent is 2. Then, each agent has two choices. We suppose that the initial state W_1 is the one where the 1st choice is taken for each agent. There are $2^5 = 32$ states.

The states (together with the way they are coded) and the transition matrix for the preferential attachment are given in Appendix 6.6.1.

There are 5 possible transitions from each state (each agent has one choice). Observe that the diagonal has strong values (i.e., the process is slow to converge). The stationary distribution is given in Figure 4. There are 4 states more probable than the others: states 5, 17, 21 and 25.




Figure 4: Stationary distribution Ex. 4 with PA

The maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.956459.

The transition matrix for the modified 2-hops (follower of follower) with $\epsilon = 0.1$ is given in Appendix 6.6.1. The diagonal is less strong than for PA. The stationary distribution is given in Figure 5. There are several states more probable than the others, which are the states 1, 5, 9,




Figure 5: Stationary distribution Ex. 4 with FF

13, 17, 21, 25 and 29 (they include those obtained with PA).

The maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.9713848, indicating a process slightly slower than preferential attachment.

It is possible to explain why the above states are more probable than the others. Looking at the graph of fixed choices, one can see that the in-degree of agent 4 is 0. Now, agents 1 and

2 have as 2nd choice the possibility to create a link to 4. In the PA model as well as in the FF model, these agents have no incentive to choose 4, therefore they have a strong tendency to choose their first choice. This choice is coded as 0 in the table of states (see Appendix 6.6.1). Therefore, states where agents 1 and 2 take value 0 are more probable. These are exactly the states $1 + 4k$, $k = 0, 1, 2, \dots, 7$. In addition, for PA, agent 3 is the most attractive since it has the maximum in-degree (equal to 2). Agent 3 can be chosen by agent 5 as 2nd choice. Therefore, states where agent 5 takes the value 1 are more probable among the previous ones. These are states 17, 21, 25, 29 (see table of states in Appendix 6.6.1).

5 Measuring fairness of recommendation algorithms

In this subsection, we focus on the measure of algorithmic fairness from the point of view of objects or followees, i.e., we ask whether the algorithm provides fair opportunities to be followed.

5.1 Micro-founded measures of unfairness

From the point of view of an individual object/followee, the number of links received is a natural measure of success. In this respect an object/followee j can “claim” a link from each agent i for which it is acceptable, i.e., such that $j \in E_i$. However, there are $e_i := |E_i|$ objects that have a similar claim whereas agent i can only form $\bar{d}_i$ links. One can consider the situation as a bankruptcy problem where an estate of $\bar{d}_i$ links is to be divided among $|E_i|$ objects, each having a claim of 1 on the estate (see Aumann and Maschler [1985]). In this simple setting, it is straightforward that any solution rule for the bankruptcy game (e.g., Shapley value, nucleolus, proportional rule) will yield the same outcome, namely that each object shall receive $\bar{d}_i/|E_i|$ from the estate. Summing over agents, this suggests that the benchmark/fair allocation of links across agents shall be such that each object/followee j receives (in expectation) a total of b_j links where

$$b_j := \sum_{\{i|j \in E_i\}} \frac{\bar{d}_i}{|E_i|}. \quad (20)$$

If $(b_j)_{j \in M}$ is a fair allocation of incoming links over objects, it appears natural to measure the unfairness of a recommendation algorithm by the extent to which the induced distribution of in-degrees for objects at convergence entails deviation from this benchmark allocation. More precisely, let us denote by e^Q the expected allocation of in-degrees associated to the stationary network distribution induced by the recommendation algorithm Q . That is, for all $j \in M$:

$$e_j^Q = \sum_{W \in \mathcal{E}} \pi^Q(W) \sum_{i \in N} W_{i,j}. \quad (21)$$

We shall define the unfairness of the recommendation algorithm Q as the Kullback-Leibler divergence⁴ between its expected allocation of in-degrees and the fair benchmark. Namely, the unfairness $\mathcal{U}(Q)$ of the recommendation algorithm Q is given by:

$$\mathcal{U}(Q) = \sum_{j \in M} e_j^Q \log \frac{e_j^Q}{b_j}. \quad (22)$$

Hence, the benchmark allocation has an unfairness of zero, i.e., it is completely fair, and unfairness increases with the distance to the benchmark.

⁴This is a standard choice when comparing distributions. Yet our results are qualitatively similar for all standard metrics on $\mathbb{R}^M$.

Remark 3. A notable particular case is that where each agent has the same linkage capacity and number of acceptable links, i.e., if $\bar{d}_i = \bar{d}_j = \bar{d}$ and $|E_i| = |E_j|$ for all i, j . One then has $b_i = b_j = \bar{b}$ for all i, j , and the unfairness of the algorithm Q is given by:

$$\mathcal{U}(Q) = \sum_{j \in M} e_j^Q \log(e_j^Q) - n\bar{d} \log(\bar{b}) \quad (23)$$

where $n\bar{d}$ is the total number of links. Hence, the unfairness of the expected allocation of in-degrees associated to an algorithm is equal, up to a constant, to the opposite of its Shannon entropy.

5.2 Entropy as a measure of fairness

In relation to the results of the preceding subsection, a key remark is that the uniform recommendation algorithm always leads to a perfectly fair distribution. Indeed, following Proposition 4, the uniform recommendation algorithm induces the uniform distribution over the set of efficient networks $\mathcal{E}$. In this setting, the expected number of links received by object j is exactly b_j . Indeed, under Assumption B, it is straightforward to remark that efficient networks are those such that each agent i has exactly $\bar{d}_i$ links towards objects in E_i . Accordingly, drawing a network according to the uniform distribution amounts to draw independently for each i , $\bar{d}_i$ links in E_i . Each acceptable object j thus has a probability $\bar{d}_i/|E_i|$ to be linked to agent i or equivalently receives $\bar{d}_i/|E_i|$ expected links from agent i . By independence, the total expected number of links towards object j is the sum of expected links received from each agent. That is, one has:

Remark 4. If assumption B holds, then the expected number of links to object j for the uniform distribution on $\mathcal{E}$ is given by $\sum_{\{i|j \in E_i\}} \frac{\bar{d}_i}{e_i} = b_j$.

Hence, the uniform distribution on $\mathcal{E}$ yields the fair distribution of links. In view of Proposition 4, an immediate corollary of this remark is that the uniform recommendation algorithm systematically leads to the fair distribution of links.

Proposition 7. Under Assumption B, one has $e_j^U = b_j$ for all $j \in M$, and accordingly $\mathcal{U}(U) = 0$.

Another consequence of Remark 4 is that unfairness can also be measured at the level of distribution over networks (rather than through the expectation of in-degrees). Indeed, given that the uniform distribution over networks is perfectly fair, the fairness of a distribution can be measured by the divergence with respect to the uniform distribution, that is through the entropy. In other words, we consider as a macro-level measure of fairness the entropy

$$\mathcal{F}(Q) = - \sum_{W \in \mathcal{E}} \pi^Q(W) \log(\pi^Q(W)). \quad (24)$$

Hence, the measure of unfairness $\mathcal{U}$ and the measure of fairness $\mathcal{F}$ are consistent in the sense that the uniform algorithm yields maximal fairness and minimal unfairness. The (un)-fairness of other algorithms may differ because one measures the divergence to the uniform benchmark in terms of the expected allocation of in-degrees whereas the other measures this divergence directly in terms of the invariant distribution over networks.

5.3 Numerical analysis of fairness

In order to investigate the fairness of an extended range of algorithms, we repeat the numerical experiments of subsection 3 and simulate the evolution of the Markov chain over 500000 periods (convergence almost always occurs within the first 100000 periods) in order to approximate the invariant distribution and the fairness of the corresponding recommendation algorithms. The results are reported in Figures 6 and 7. The results for preferential attachment and n -hops type of algorithms are qualitatively similar. The level of unfairness first increases with the level of correlation of preferences, reaches a maximum, and then decreases as the preferences become fully correlated. In quantitative terms, the level of unfairness is higher for the more hierarchical algorithms, i.e., preferential attachment with larger α parameter and n -hops algorithms with larger n . These results can be explained as follows. For low levels of correlation, the preferential attachment mechanism does not really “bite” because no object can sustain a high in-degree. For high level of correlation, the set of efficient networks is small and rather homogeneous as most agents seek to form links with the few “star” objects (in the limit of perfect correlation, there is even a single efficient network). Accordingly, the divergence between the uniform distribution and distributions induced by hierarchical recommendation algorithms is limited. Unfairness reaches a maximum for intermediary level of correlation because this is a domain where (i) central nodes can emerge and be identified by hierarchical algorithms, (ii) the set of efficient networks is sufficiently large and heterogeneous for divergence with the uniform benchmark to materialize.

5.4 Fairness-efficiency trade-off

The previous results highlight the presence of a trade-off between efficiency and fairness (in particular for intermediate level of preference correlation), as the recommendation algorithms that ensure fast convergence to efficient networks are also those that lead to high level of unfairness. However, our analysis also hints at a simple solution to mitigate this trade-off. It suffices to update the recommendation algorithm once the class of efficient network is reached. More precisely, the designer of an online platform can adapt the recommendation algorithm to the different phases of the network formation process. Hierarchical algorithms can be used in the early phase of the process to reach rapidly the set of efficient networks and ensure user satisfaction and adoption. Once the user base is established, more uniform recommendation algorithms can be used to further explore the set of potential networks and guarantee fairness without paying an efficiency cost.

This shift towards a uniform distribution can also be achieved through the modification of the recommendation algorithm via the Metropolis (Monte-Carlo Markov Chain, MCMC) process. Indeed, it can be used to deviate the algorithm towards the uniform distribution (see, e.g., Levin et al. [2009]). More precisely, suppose $\Psi = [\psi_{ij}]$ is an arbitrary transition matrix, and π is an arbitrary distribution (e.g. the uniform distribution over efficient networks in our context). Then the Metropolis chain is defined by

$$p_{ij} = \begin{cases} \psi_{ij} \left(1 \wedge \frac{\pi_j \psi_{ji}}{\pi_i \psi_{ij}}\right), & \text{if } j \neq i \\ 1 - \sum_{k:k \neq i} \psi_{ik} \left(1 \wedge \frac{\pi_k \psi_{ki}}{\pi_i \psi_{ik}}\right), & \text{if } j = i, \end{cases}$$

and has stationary distribution π . The approach is to reject with some probability transitions given by Ψ . Hence, P converges more slowly than Ψ , but the Metropolis process is the fastest among all processes based on this principle.




Figure 6: Unfairness (divergence with respect to the benchmark distribution), averaged over 25 simulations, for a range of (preferential attachment) recommendation algorithms and parameter values: $n = 200, m = 200, a = 10$ in the upper-left panel, $n = 20, m = 20, a = 10$ in the upper-right panel, $n = 200, m = 200, a = 10$ in the lower-left panel, $n = 200, m = 200, a = 4$ in the lower-right panel. Stationary distributions are approximated by the empirical frequency observed between hitting time of the recurrence class and $T = 500000$.




Figure 7: Unfairness (divergence with respect to the benchmark distribution), averaged over 25 simulations, for a range of (hops) recommendation algorithms and parameter values: $n = 200, m = 200, a = 10$ in the left panel, $n = 20, m = 20, a = 4$ in the right panel. Stationary distributions are approximated by the empirical frequency observed between hitting time of the recurrence class and $T = 500000$.

Example 5. We have applied the Metropolis process of Example 4, to the preferential attachment and the 2-hops algorithms, which yields indeed a uniform distribution. For preferential attachment, the maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.967487, slightly above the original value (0.956459). Similarly, for 2-hops, the maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.984847, slightly above the original one (0.9713848).

6 Discussion and Conclusion

6.1 Summary of results

This report presents the model of network formation governed by recommendation algorithms developed in task 2.4. The model is built as a finite Markov chain in which the algorithm sequentially proposes links to agents who decide which links they form on the basis of a preference profile over target objects/followees. This simple structure allows for both mathematical and computational analysis of the model. We analyze the convergence properties and the structure of the stationary distribution of the chain that we relate, respectively, to the efficiency and to the fairness of the recommendation algorithm.

We first show that the utility-based nature of the link acceptance process implies that the system converges to an absorbing class of efficient networks, provided the recommendation algorithm has a sufficiently large support. In this sense, any recommendation algorithm is “efficient”. However, for impatient users, the waiting time before receiving suitable recommendations is a key element in comparing platforms and recommendation algorithms. Accordingly, we define the efficiency of an algorithm as its expected hitting time to the absorbing class of efficient networks. We provide an analytical formula for this hitting time that underlines that the performance of a recommendation algorithm is tightly linked to the structure of the preference profiles. When preferences are uncorrelated, uniform recommendation algorithms perform as well (or even better) than hierarchical algorithms whose recommendations favour more central

nodes, e.g., in terms of degree for preferential attachment. When preferences are correlated, hierarchical algorithms perform better as they rapidly identify the objects/followees collectively preferred by users.

We then provide an analytical characterization of the asymptotic distribution of (efficient) networks for uniform and preferential recommendation algorithms. We further provide an exact algorithm for the computation of this distribution for the class of recommendation algorithms that induce reversible Markov chains. The invariant distribution characterizes the long-term behavior of the recommendation algorithm as it determines the set of networks generated and their frequency of occurrence. We thus measure the fairness of a recommendation algorithm in function of the associated invariant distribution. Specifically, we define the fairness of a recommendation algorithm as its divergence from a benchmark where each object/followee receives a number of links proportional to the number of agents for whom it is a priori desirable. This approach provides micro-foundations for using the entropy of the invariant distribution as a measure of fairness for a recommendation algorithm. The uniform recommendation algorithm is always perfectly fair. The fairness of more hierarchical algorithms, such as preferential attachment and n -hops, depends on the structure of preferences. They are relatively fair when the correlation between preferences is very high or very low but can be very unfair for intermediate level of correlation as the recommendation algorithms focus on central nodes whereas a number of alternative network configurations would be equally desirable.

6.2 Evolution of recommendation algorithms

Our results put forward the existence of a trade-off between fairness and efficiency for recommendation algorithms. However, our analysis also hints at a simple solution to mitigate this trade-off. The designer of an online platform can adapt the recommendation algorithm to the different phases of the network formation process. Hierarchical algorithms can be used in the early phase of the process to reach rapidly the set of efficient networks and ensure user satisfaction and adoption. Once the user base is established, more uniform recommendation algorithms can be used to further explore the set of potential networks and guarantee fairness without paying an efficiency cost. There are evidences of platform behavior being aligned with our conclusions in the sense that recommendation algorithms have been updated to foster a more extensive exploration of the network and the recommendation of more diverse content. A major wave of such algorithmic updates has occurred in 2024 concerning notably Instagram and Snapchat (see Figure 8)

It appears that these algorithmic updates stem from both competitive and regulatory pressure. In our setting, the impact of competition on algorithmic fairness and openness seems ambiguous. On the one hand, competition implies a strong efficiency requirement for platforms that must thus identify rapidly and recommend broadly the content that is likely to be appreciated by the largest number of possible users. This reduces diversity in recommendation and has a negative impact on fairness. On the other hand, competition implies a need for diversification that can foster the adoption of more exploratory and fair algorithms.

6.3 Impacts on socio-economic dynamics

Beyond fairness “per se”, recommendation algorithms can also have substantial impacts on socio-economic dynamics. First, they induce changes on the structure of social networks. Indeed, link-recommendation algorithms lead to more hierarchical social structures than decentralized/offline social network formation process based on local interactions, which shall be comparable to recommendation algorithms with low-levels of preferential attachment or n -hops



Figure 8: Illustration, through two recent media pieces, of the redesign of recommendation algorithms to foster more extensive exploration of social networks focusing on the case of Snapchat and Instagram. See respectively <https://www.bloomberg.com/news/articles/2024-09-17/snap-updates-video-feed-in-competition-with-meta-tiktok?embedded-checkout=true> and <https://www.techtimes.com/articles/304135/20240430/instagram-algorithm-will-focus-more-original-content-straying-away-reposted.htm>

with small n . This implies more uneven distribution of centralities in social networks and thus increased inequality in terms of access to information, social capital, or symbolic power. These structural inequalities can have substantial impacts on the dynamics processes unfolding on social networks.

First, in the context of structural economic change, networks play a crucial role in the diffusion of new technologies Robertson et al. [1996], Deroian [2002], Hanaki et al. [2010], Vega and Mandel [2018]. Hence, it is key for the success of the twin transition to have network structures aligned for the rapid and efficient diffusion of technological innovations. Hierarchical recommendation algorithms might however lead to inefficiencies in this respect. Indeed, innovative firms, jobs, skills, or people might have their access to information and/or their social (media) visibility reduced because backward-looking recommendation algorithms will, by default, discriminate against newly created/innovative entities. This can slow down matching processes on the labor and funding markets for emerging industries and sectors. This can also hamper the diffusion of new products, unless they are adopted by already influential users.

More specifically, in the context of opinion formation, it is well-known that network centrality, in particular eigenvector centrality, is a key determinant of influence on opinion dynamics [see

e.g. DeGroot, 1974, Deffuant et al., 2001, Hegselmann and Krause, 2002, Golub and Jackson, 2010]. Accordingly, hierarchical recommendation algorithms will generate highly skewed distributions of social influence where a few actors can have extremely large and rapid impacts on a large share of the population as exemplified by the role of social-media “influencers” in the 2024 Romanian presidential election campaign.

Additionally, hierarchical recommendation algorithms will lead to a very heterogeneous density of links with a high concentration of links around high-centrality nodes and very sparse links in the rest of the networks. Such unbalanced distribution of links make social networks highly vulnerable to polarization [see e.g. Grabisch et al., 2023]. Indeed, weak links between moderate/central agents can hardly resist to centrifuge forces in the presence of highly connected nodes with extreme/polar opinions. In such situations, social network can disconnect, leading to further amplification of polarization mechanisms. The polarization of opinions can be a major obstacle to twin transitions as exemplified by the case of the French yellow vests [see e.g. Douenne and Fabre, 2022].

6.4 Policy recommendations

Overall, it appears that some form of monitoring and/or regulation of recommendation algorithms is warranted. A first step in this direction is the requirement to make algorithms available to regulators (and ideally to users as well) in order to evaluate their performance in terms of fairness and their potential systemic bias. Some online platforms allow users to select the recommendation algorithms they use but substantial opacity remain as underlined by the recent request to some online platforms by the European Commission under the Digital Services Act (see footnote 1 above) seems well aligned with these objectives. A second step would be the actual regulation of the algorithms. Our results highlight that, in theory, an automatic form of regulation could be possible, e.g. by including desired algorithmic features through the Metropolis algorithm. This is however likely to raise major efficiency concerns. A more appropriate route could be to impose to recommendation algorithms specifications that express fairness requirements. This approach would also avoid the perils of uniformization by regulation. Indeed, our results also hint that no recommendation algorithm is perfect for all situations. Both the efficiency and the fairness of an algorithm in fact depend on the structure of preferences of users.

References

- R. Aumann and M. Maschler. Game theoretic analysis of a bankruptcy problem from the Talmud. Journal of Economic Theory, 36(2):195–213, 1985.
- S. Balseiro, H. Lu, and V. Mirrokni. Regularized online allocation problems: Fairness and beyond. In Proceedings of the 38th International Conference on Machine Learning, pages 630–639. PMLR, New York, 2021.
- S. Barocas, M. Hardt, and A. Narayanan. Fairness and Machine Learning: Limitations and Opportunities. MIT Press, 2023.
- M. H. Bateni, Y. Chen, D. F. Ciocan, and V. Mirrokni. Fair resource allocation in a volatile marketplace. Operations Research, 70(1):288–309, 2022.
- R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth. Fairness in criminal justice risk assessments: The state of the art. arXiv:1703.09207, 2017.
- T. Bono, K. Croxson, and A. Giles. Algorithmic fairness in credit scoring. Oxford Review of Economic Policy, 37(3):585–617, 2021.
- Y. Bramoullé, A. Galeotti, and B. W. Rogers, editors. The Oxford Handbook of the Economics of Networks. Oxford University Press, 2016.
- H. Cao, L. Ma, Z. E. Ning, and B. Sun. How does competition affect exploration vs. exploitation? A tale of two recommendation algorithms. Management Science, 70(2):1029–1051, 2024.
- S. Caton and C. Haas. Fairness in machine learning: A survey. arXiv preprint arXiv:2010.04053, 2020.
- F. Cinus, M. Minici, C. Monti, and F. Bonchi. The effect of people recommenders on echo chambers and polarization. In Proceedings of the International AAAI Conference on Web and Social Media, volume 16, pages 90–101, 2022.
- G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. Advances in Complex Systems, 3(01n04):87–98, 2001.
- M. H. DeGroot. Reaching a consensus. Journal of the American Statistical Association, 69:118–121, 1974.
- Frédéric Deroian. Formation of social networks and diffusion of innovations. Research policy, 31(5):835–846, 2002.
- P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of Markov chains. Annals of Applied Probability, 1:36–61, 1991.
- M. Dinerstein, L. Einav, J. Levin, and N. Sundaresan. Consumer price search and platform design in internet commerce. American Economic Review, 108(7):1820–1859, 2018.
- Thomas Douenne and Adrien Fabre. Yellow vests, pessimistic beliefs, and carbon tax aversion. American Economic Journal: Economic Policy, 14(1):81–110, 2022.
- Y. Eisenstat. How to hold social media accountable for undermining democracy. Harvard Business Review, January 2021.

- F. Fagan. Systemic social media regulation. Duke Law & Technology Review, 16(1):393–439, 2018.
- A. Ferrara, L. Espín-Noboa, F. Karimi, and C. Wagner. Link recommendations: Their impact on network structure and minorities. arXiv preprint arXiv:2205.06048, 2022.
- S. Gelper, R. van der Lans, and G. van Bruggen. Competition for attention in online social networks: Implications for seeding strategies. Management Science, 67(2):1026–1047, 2021.
- D. Ghosh. What will tech regulation look like in the Biden era? Harvard Business Review, December 2020.
- D. Ghosh. Are we entering a new era of social media regulation? Harvard Business Review, January 2021.
- B. Golub and M. O. Jackson. Naïve learning in social networks and the wisdom of crowds. American Economic Journal: Microeconomics, 2(1):112–149, 2010.
- M. Grabisch, A. Mandel, and A. Rusinowska. On the design of public debate in social networks. Operations Research, 71(2):626–648, 2023.
- Michel Grabisch, Antoine Mandel, and Agnieszka Rusinowska. Efficiency versus fairness in link recommendation algorithms. Documents de travail du Centre d'Économie de la Sorbonne, 2024.
- Nobuyuki Hanaki, Ryo Nakajima, and Yoshiaki Ogura. The dynamics of R&D network in the IT industry. Research Policy, 39(3):386–399, 2010.
- R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence, models, analysis and simulation. Journal of Artificial Societies and Social Simulation, 5(3):1–33, 2002.
- C. Hurlin, C. Pérignon, and S. Saurin. The fairness of credit scoring models. Management Science, 2024.
- M. O. Jackson. Social and Economic Networks. Princeton University Press, 2008.
- M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. Journal of Economic Theory, 71:44–74, 1996.
- J. P. Johnson, A. Rhodes, and M. Wildenbeest. Platform design when sellers use pricing algorithms. Econometrica, 91(5):1841–1879, 2023.
- F. P. Kelly. Reversibility and Stochastic Networks. Cambridge University Press, 2011.
- J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan. Algorithmic fairness. American Economic Association Papers and Proceedings, 108:22–27, 2018.
- J. Komiyama and S. Noda. On statistical discrimination as a failure of social learning: A multi-armed bandit approach. Management Science, 2024.
- A. Lambrecht and C. Tucker. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. Management Science, 65(7):2966–2981, 2019.
- M. A. Lejeune and J. Turner. Planning online advertising using Gini indices. Operations Research, 67(5):1222–1245, 2019.

- D. A. Levin, Y. Peres, and E. L. Wilmer. Markov Chains and Mixing Times. American Mathematical Society, 2009.
- X. Li, Y. Rong, R. Zhang, and H. Zheng. Online advertisement allocation under customer choices and algorithmic fairness. Management Science, 2024.
- W. Ma, P. Xu, and Y. Xu. Group-level fairness maximization in online bipartite matching. Preprint, <https://arxiv.org/abs/2011.13908>, 2021.
- W. Ma, P. Xu, and Y. Xu. Fairness maximization among offline agents in online-matching markets. ACM Transactions on Economics and Computation, 10(4):1–27, 2023.
- C. Makridis and M. H. M. Teodorescu. Fairness in machine learning: Regulation or standards? Brookings Institution, 2024.
- S. Mitchell, E. Potash, S. Barocas, A. D’Amour, and K. Lum. Algorithmic fairness: Choices, assumptions, and definitions. Annual Review of Statistics and Its Application, 8:141–163, 2021.
- R. B. Myerson. Graphs and cooperation in games. Mathematics of Operations Research, 2: 225–229, 1977.
- N. Navarro. Expected fair allocation in farsighted network formation. Social Choice and Welfare, 43(2):287–308, 2014.
- T. Redman. Are your algorithms upholding your standards of fairness? Harvard Business Review, November 2019. Accessed: 2024-12-23.
- Maxine Robertson, Jacky Swan, and Sue Newell. The role of networks in the diffusion of technological innovation. Journal of management studies, 33(3):333–359, 1996.
- A. E. Roth. Who gets what—and why: The new economics of matchmaking and market design. Houghton Mifflin Harcourt, 2015.
- F. P. Santos, Y. Lelkes, and S. A. Levin. Link recommendation algorithms and dynamics of polarization in online social networks. Proceedings of the National Academy of Sciences, 118(50):e2102141118, 2021.
- S. Vaidhyanathan. Antisocial Media: How Facebook Disconnects us and Undermines Democracy. Oxford University Press, 2018.
- Solmaria Halleck Vega and Antoine Mandel. Technology diffusion and climate policy: a network approach and its application to wind energy. Ecological Economics, 145:461–471, 2018.
- Y. Wang, W. Ma, M. Zhang, Y. Liu, and S. Ma. A survey on the fairness of recommender systems. ACM Transactions on Information Systems, 41(3):1–43, 2023.
- J. Yang, A. A. S. Soltan, D. W. Eyre, and D. A. Cifton. Algorithmic fairness and bias mitigation for clinical machine learning with deep reinforcement learning. Nature Machine Intelligence, 5(11):884–894, 2023.
- K. Zhang and M. Sarvary. Social media competition: Differentiation with user generated content. Marketing Science, 47:48, 2011.

Appendix

6.5 Proofs

6.5.1 Proof of Lemma 1

Given a link $\ell := (i, j) \in N \times M$, we denote by V_ℓ the matrix for which (i, j) is the only non-zero coefficient. With a slight abuse of notation, we also denote by $V_\emptyset$ the matrix 0 in $\mathbb{R}^{N \times M}$ and let $L' := L \cup \{\emptyset\}$.

1. Let us first remark that following Definition 1, one has:

$$P_{(Q, \bar{d})}(W, W') \geq 0 \Rightarrow \forall i \in N, U_i(W') \geq U_i(W).$$

This implies in particular that a network W' is accessible from W only if $\sum_{i \in N} U_i(W') \geq \sum_{i \in N} U_i(W)$.

2. Let us first show that any network $W \notin \mathcal{E}$ is not recurrent. Indeed, if $W \notin \mathcal{E}$, there exists $\ell \in L_a$ and $\ell' \in L'$ such that $U_i(W + V_\ell - V_{\ell'}) > U_i(W)$. Let $W' = W + V_\ell - V_{\ell'}$. According to Assumption A, one has $Q(W, \ell) > 0$ and thus, following Definition 1, $P_{(Q, \bar{d})}(W, W') > 0$ and W' is accessible from W . Furthermore, following 1. above, W is not accessible from W' because $\sum_{i \in N} U_i(W') > \sum_{i \in N} U_i(W)$. This implies that W is not recurrent.

3. Let us then show that any network $W \in \mathcal{E}$ is recurrent. It is straightforward that for any $W' \in \mathcal{W}_0$, there exists a finite sequence of link addition/deletion $(\ell_1, \ell'_1), \dots, (\ell_S, \ell'_S)$ such that

(i) For all $\sigma = 1, \dots, S$, $\ell'_\sigma \in L_a$,

(ii) $W = W' + \sum_{s=1}^S V_{\ell_s} - V_{\ell'_s}$

(iii) For all $i \in N$ and all $\sigma = 1, \dots, S$:

$$U_i(W' + \sum_{s=1}^{\sigma} V_{\ell_s} - V_{\ell'_s}) \geq U_i(W' + \sum_{s=1}^{\sigma-1} V_{\ell_s} - V_{\ell'_s}).$$

This implies, following Assumption A that for all $s = 1, \dots, S$,

$$P_{(Q, \bar{d})}(W' + \sum_{s=1}^{\sigma-1} V_{\ell_s} - V_{\ell'_s}, W' + \sum_{s=1}^{\sigma} V_{\ell_s} - V_{\ell'_s}) \geq 0.$$

Thus W is accessible from W' . As this holds for all $W' \in \mathcal{W}_0$ and all $W \in \mathcal{E}$, this implies that all $W \in \mathcal{E}$ are recurrent and that there is a unique recurrent class.

4. Finally, as for all $W \in \mathcal{E}$, one has $P_{(Q, \bar{d})}(W, W) > 0$, the Markov chain is aperiodic.

6.5.2 Proof of Proposition 2

1. Observe that, given that the current state W is in $\mathcal{E}$, if i is selected by the recommendation algorithm, there are $|C_i \setminus M_i(W)| \cdot |C_i \cap M_i(W)|$ possible moves (adding a new link and removing an old one). This yields $\prod_{i \in N} |C_i \setminus M_i(W)| \cdot |C_i \cap M_i(W)|$ possible moves in total from W . Therefore the total number of efficient W is

$$|\mathcal{E}| = \prod_{i \in N} \binom{|C_i|}{\bar{d}_i - |E_i| + |C_i|}.$$

2. As we restrict to $\mathcal{E}$, (2) reduces to the second case. The denominator takes into account the different (considered as equiprobable) choices for deleting a link, which yields the formula.
3. $C_i \cap M_i(W)$ is the set of choices for agent i to delete a link. In $\mathcal{E}$, we have $|M_i(W)| = \bar{d}_i$. Also, $M_i(W) \setminus C_i$ is the set of indisputable links, which therefore does not depend on W . We have then

$$|C_i \cap M_i(W)| = \underbrace{|M_i(W)|}_{\bar{d}_i} - |M_i(W) \setminus C_i|,$$

which proves the claim.

6.5.3 Proof of Lemma 2

Let us denote by $\mathcal{T}$ the set of trajectories for the Markov chain with transition matrix P_Q and an initially empty network. The set $\mathcal{T}$ can be partitioned between the subset $\mathcal{T}_{nc}$ of trajectories that do not reach the recurrent class $\mathcal{E}$ and the subset of trajectories $\mathcal{T}_c$ that reach the recurrent class. Following Proposition 1, the set $\mathcal{T}_{nc}$ has probability zero. Thus, without loss of generality, the expected hitting time can be computed conditional on the trajectory being in $\mathcal{T}_c$. Now, by definition, each trajectory in $\mathcal{T}_c$ follows a minimal path. This implies that:

$$T^{Q,u} = \sum_{\ell \in \mathcal{M}(u)} P_\ell^Q T_\ell^Q$$

where P_ℓ^Q is the probability that the Markov chain follows the minimal hitting path ℓ and T_ℓ^Q is the expected time for hitting $\mathcal{E}$ conditional on the Markov chain following the minimal hitting path ℓ .

6.5.4 Proof of Lemma 3

Let us consider a minimal hitting path $\ell = (\ell_1, \dots, \ell_s)$. Starting from the empty network, we have:

- Step 1: from $\emptyset$ to the network $\{\ell_1\}$, with $\ell_1 = (i_1, j_1)$. The probability that ℓ_1 is the first acceptable link recommended by Q from the empty network is given by

$$P_{\emptyset, \ell_1}^Q = \sum_{t \in \mathbb{N}} [1 - Q(\emptyset, L_a)]^t Q(\emptyset, \ell_1) = \frac{Q(\emptyset, \ell_1)}{Q(\emptyset, L_a)}$$

where $Q(\emptyset, \cdot)$ is the probability of recommendation from the empty network. Hence, $[1 - Q(\emptyset, L_a)]^t Q(\emptyset, \ell_1)$ corresponds to the probability of the recommendation of t unsatisfactory links (in the complement of L_a) before the recommendation of ℓ_1 . The second equality comes from the summation of the geometric series.

Accordingly, conditional on ℓ being followed, the expected waiting time until ℓ_1 is formed is given by

$$T_{\emptyset, \ell_1}^Q = \sum_{t \in \mathbb{N}} (t+1) [1 - Q(\emptyset, L_a)]^t = \frac{1}{[Q(\emptyset, L_a)]^2}.$$

- Step k : from $V_{k-1} := \{\ell_1, \dots, \ell_{k-1}\}$ to $V_k := V_{k-1} \cup \{\ell_k\}$, with $\ell_k = (i_k, j_k)$. The probability that ℓ_k is the first new acceptable link formed from V_{k-1} is given by

$$P_{V_{k-1}, \ell_k}^Q = \sum_{t \in \mathbb{N}} [1 - Q(V_{k-1}, L_a \setminus V_{k-1})]^t Q(V_{k-1}, \ell_k) = \frac{Q(V_{k-1}, \ell_k)}{Q(V_{k-1}, L_a \setminus V_{k-1})}.$$

Accordingly, conditional on ℓ being followed, the expected waiting time until ℓ_k is formed from V_{k-1} is given by

$$T_{V_{k-1}, \ell_k}^Q = \sum_{t \in \mathbb{N}} (t+1) [1 - Q(V_{k-1}, L_a \setminus V_{k-1})]^t = \frac{1}{[Q(V_{k-1}, L_a \setminus V_{k-1})]^2}.$$

Finally, we obtain

$$P_\ell^Q = \frac{Q(\emptyset, \ell_1)}{Q(\emptyset, L_a)} \times \frac{Q(V_1, \ell_2)}{Q(V_1, L_a \setminus V_1)} \times \cdots \times \frac{Q(V_{s-1}, \ell_s)}{Q(V_{s-1}, L_a \setminus V_{s-1})}$$

$$T_\ell^Q = \frac{1}{[Q(\emptyset, L_a)]^2} + \frac{1}{[Q(V_1, L_a \setminus V_1)]^2} + \cdots + \frac{1}{[Q(V_{s-1}, L_a \setminus V_{s-1})]^2}.$$

6.5.5 Proof of Proposition 4

It suffices to prove that the transition matrix P_U is symmetric, i.e., for any two states W, W' , $P_U(W, W') = P_U(W', W)$. Let us put $W' = W + V_{i,j} - V_{i,k}$. Using (11) and the definition of U , we get:

$$P_U(W, W + V_{i,j} - V_{i,k}) = \frac{1}{nm|C_i \cap M_i(W)|}$$

$$P_U(W + V_{i,j} - V_{i,k}, W) = \frac{1}{nm|C_i \cap M_i(W + V_{i,j} - V_{i,k})|}.$$

However, by Proposition 2 (3), $|C_i \cap M_i(W)|$ does not depend on W . Therefore, the two expressions are equal.

6.5.6 Proof of Proposition 5

1. We first prove that the matrix P_{PA_f} is reversible. For this, it suffices to show that it satisfies the Kolmogorov criterion [see Kelly, 2011], i.e., that for every finite sequence of distinct recurrent networks $\mathcal{W} := W_0, \dots, W_T$, one has:

$$P_{PA_f}(W_0, W_1)P_{PA_f}(W_1, W_2) \cdots P_{PA_f}(W_T, W_0) = P_{PA_f}(W_0, W_T)P_{PA_f}(W_T, W_{T-1}) \cdots P_{PA_f}(W_1, W_0) \quad (25)$$

(see (18)). Let us set accordingly $\vec{\mathcal{W}} := W_0, \dots, W_T, W_0$ and $\overleftarrow{\mathcal{W}} := W_0, W_T, \dots, W_1, W_0$.

Consider in $\vec{\mathcal{W}}$ the transition from W_t to W_{t+1} , and suppose that $W_{t+1} = W_t - V_{i_t, j_t} + V_{i_t, k_t}$. Using (11) and (5), the probability of transition is

$$P_{PA_f}(W_t, W_{t+1}) = \frac{PA_f(W_t, (i_t, k_t))}{|C_{i_t} \cap M_{i_t}(W_t)|} = \frac{f(d_{k_t}^-(W_t))}{m \sum_h f(d_h^-(W_t)) |C_{i_t} \cap M_{i_t}(W_t)|}.$$

Observe that that term $m \sum_h f(d_h^-(W_t))$ depends only on W_t (let us abbreviate it by $\gamma(W_t)$), while the term $|C_{i_t} \cap M_{i_t}(W_t)|$ depends only on i_t (let us denote it by $\kappa(i_t)$). Hence, computing the left hand-side in (25) yields:

$$LHS = \frac{\prod_{t=0}^T f(d_{k_t}^-(W_t))}{\prod_{t=0}^T \gamma(W_t) \kappa(i_t)}.$$

Consider now the reverse sequence $\overleftarrow{\mathcal{W}}$. Observe that in the transition W_{t+1} to W_t , the link (i_t, k_t) is deleted, while the link (i_t, j_t) is created. Therefore,

$$P_{\text{PA}_f}(W_{t+1}, W_t) = \frac{\text{PA}_f(W_{t+1}, (i_t, j_t))}{|C_{i_t} \cap M_{i_t}(W_{t+1})|} = \frac{f(d_{j_t}^-(W_{t+1}))}{\gamma(W_{t+1})\kappa(i_t)}.$$

Therefore, the right hand-side of (25) reads:

$$RHS = \frac{\prod_{t=T}^0 f(d_{j_t}^-(W_{t+1}))}{\prod_{t=T}^0 \gamma(W_{t+1})\kappa(i_t)}$$

with $W_{T+1} := W_0$. It follows that

$$\frac{LHS}{RHS} = \frac{\prod_{t=0}^T f(d_{k_t}^-(W_t))}{\prod_{t=T}^0 f(d_{j_t}^-(W_{t+1}))}.$$

It remains to prove that the above ratio reduces to 1. We will show that each term in the numerator has a corresponding term in the denominator.

Case 1: Transition W_0 to W_1 . We have that $W_1 = W_0 + V_{i_0, k_0} - V_{i_0, j_0}$. As the sequence $\overrightarrow{\mathcal{W}}$ returns to W_0 , there must exist t such that the in-degree of k_0 in W_t is $d_{k_0}^-(W_0)$. Denote by t_0 the first such time step, i.e., the in-degree decreases from W_{t_0-1} to W_{t_0} . Consequently, in the reverse sequence $\overleftarrow{\mathcal{W}}$, there is an increment of one unit of the in-degree of k_0 in the transition $W_{t_0} \rightarrow W_{t_0-1}$. Therefore, $d_{k_0}^-(W_0) = d_{k_0}^-(W_{t_0})$, and the term $f(d_{k_0}^-(W_0))$ in RHS cancels $f(d_{k_0}^-(W_t))$ in LHS.

Case 2: Transition from W_t to W_{t+1} , $t > 0$. The in-degree of k_t increases by 1 during this transition.

Case 2a: Suppose that the previous change in the in-degree of k_t is a negative increment (-1) , say during the transition $W_s \rightarrow W_{s+1}$, $s < t$. Then in the reverse sequence $\overleftarrow{\mathcal{W}}$, in the transition $W_{s+1} \rightarrow W_s$, the degree of k_t will be incremented by 1, and $d_{k_t}^-(W_t) = d_{k_t}^-(W_{s+1})$. Therefore, the term $f(d_{k_t}^-(W_t))$ in RHS cancels $f(d_{k_t}^-(W_{s+1}))$ in LHS.

Case 2b: Suppose that the previous change in the in-degree of k_t is a positive increment $(+1)$. Suppose first that $d_{k_t}^-(W_t) \geq d_{k_t}^-(W_0)$. As the sequence returns to W_0 , there must exist a time step s where in the transition $W_s \rightarrow W_{s+1}$ the in-degree of k_t decreases and $d_{k_t}^-(W_{s+1}) = d_{k_t}^-(W_t)$. Therefore, in the reverse sequence $\overleftarrow{\mathcal{W}}$, there is an increment of $d_{k_t}^-$ in the transition $W_{s+1} \rightarrow W_s$, and $d_{k_t}^-(W_t) = d_{k_t}^-(W_{s+1})$, and the corresponding terms in LHS and RHS cancel each other. Suppose now that $d_{k_t}^-(W_t) < d_{k_t}^-(W_0)$. As the sequence starts from W_0 , there must exist a time step s where in the transition $W_s \rightarrow W_{s+1}$ the in-degree of k_t decreases and $d_{k_t}^-(W_{s+1}) = d_{k_t}^-(W_t)$. Then, the reasoning is much the same as in the previous case.

2. It remains to prove (19). Considering the two states W and $W + V_{i,j} - V_{i,k}$, we get from (17)

$$\pi_{\text{PA}_f}(W) \frac{f(d_j^-(W))}{m \sum_h f(d_h^-(W)) |C_i \cap M_i(W)|} = \pi_{\text{PA}_f}(W + V_{i,j} - V_{i,k}) \frac{f(d_k^-(W + V_{i,j} - V_{i,k}))}{m \sum_h f(d_h^-(W + V_{i,j} - V_{i,k})) |C_i \cap M_i(W + V_{i,j} - V_{i,k})|}.$$

As $|C_i \cap M_i(W)|$ does not depend on W by Proposition 2 (3), the result follows.

6.5.7 Proof of Proposition 6

Using the notation introduced above, the set of possible choices for agent i is C_i , among which the agent has to choose $\bar{d}_i - |E_i \setminus C_i| =: k_i$ links. Therefore, the number of choice subsets is $\binom{c_i}{k_i}$, with $c_i = |C_i|$. We order these subsets by the following procedure. Let for simplicity $C_i = \{1, 2, \dots, c_i\}$. Let us define the following order on subsets (omitting braces, commas and subindex i for simplicity and ordering the elements of subsets in numerical increasing order):

$$\begin{aligned} &12\dots(k-1)k, 12\dots(k-1)(k+1), \dots, 12\dots(k-1)m, 12\dots km, 12\dots k(m-1), \dots, 12\dots k(k+1), \\ &12\dots(k+1)(k+2), \dots, 12\dots(k+1)m, \dots, 12\dots(k-2)(m-1)m, 12\dots(k-1)(m-1)m, \\ &\dots, 12\dots(k-1)k(k+1), \dots, \dots, (m-k+1)\dots m. \end{aligned}$$

The rationale is that each “position” is alternatively increasing and decreasing, starting from the last position. For example, taking $c_i = 7$ and $k_i = 3$, the following sequence is obtained (reading the columns from top to bottom and from left to right):

123	135	167	234	356
124	134	267	235	357
125	145	256	236	367
126	146	257	237	467
127	147	247	347	457
137	157	246	346	456
136	156	245	345	567

Doing so, it can be checked that between two consecutive subsets, there is only one element in the set difference. Let us number the subsets in this sequence from 1 to $\ell_i := \binom{c_i}{k_i}$.

Once the previous operation has been done for each agent, a particular state is coded by a chain (or word) of n numbers $s_1 \dots s_n$, with $1 \leq s_i \leq \ell_i$ corresponding to the s_i th subset in the sequence of choices for agent i . It remains to order these words in such a way that between two consecutive words, only one agent has changed its choice subset, taking the next or preceding choice subset in the sequence. Due to the above construction of the sequence, it follows that between two consecutive states, only one link has been deleted, and one added, i.e., a Hamiltonian path in G_Q has been constructed. The principle of the ordering of the words resembles the previous one and is as follows: alternatively increase from 1 to ℓ_i then decrease from ℓ_i to 1 each “letter” in the word, starting from the last letter. This gives for example, taking $n = 3$, $\ell_1 = 2$, $\ell_2 = 3$, $\ell_3 = 6$, the following order (reading the columns from top to bottom and from left to right):

111	126	131	236	221	216
112	125	132	235	222	215
113	124	133	234	223	214
114	123	134	233	224	213
115	122	135	232	225	212
116	121	136	231	226	211

As it can be checked, between two consecutive words, only one letter has changed, and the increment is ± 1 .

6.5.8 Generation of random preference profiles

For sake of simplicity, we assume that each agent has the same maximal number of links $\bar{d}$ and the same number of acceptable links $\bar{a}$. For each level of correlation $\alpha \in [0, 1]$, a random

preference profile is generated by repeating the following operation for each agent. We consider two “urns”, urn 1 initially contains objects 1 to $\bar{a}$, urn 2 initially contains all objects. We then consider a binomial distribution with n trials and probability of “success” α . In case of success, we draw an object in the first urn, remove it from both urns, and assign it as acceptable to the agent (i.e., we set $u_i(j) = 1$ where i and j respectively are the agent and the object under consideration). In case of failure, a similar procedure is followed but the object is drawn from the second urn. This procedure generates a preference profile for each agent. One shall remark that if $\alpha = 0$, the preferences of all agents are independent, while for $\alpha = 1$ the preferences are fully correlated as for all i , $E_i = C_i = \{1, \dots, \bar{a}\}$.

6.6 States and transition matrices for numerical examples

6.6.1 States and transition matrices for Example 4

There are $2^5 = 32$ states, which are ordered in the following way (for each agent, 0 indicates 1st choice, 1 indicates 2nd choice):

state number					
1	0	0	0	0	0
2	1	0	0	0	0
3	0	1	0	0	0
4	1	1	0	0	0
5	0	0	1	0	0
6	1	0	1	0	0
7	0	1	1	0	0
8	1	1	1	0	0
9	0	0	0	1	0
10	1	0	0	1	0
11	0	1	0	1	0
12	1	1	0	1	0
13	0	0	1	1	0
14	1	0	1	1	0
15	0	1	1	1	0
16	1	1	1	1	0
17	0	0	0	0	1
18	1	0	0	0	1
19	0	1	0	0	1
20	1	1	0	0	1
21	0	0	1	0	1
22	1	0	1	0	1
23	0	1	1	0	1
24	1	1	1	0	1
25	0	0	0	1	1
26	1	0	0	1	1
27	0	1	0	1	1
28	1	1	0	1	1
29	0	0	1	1	1
30	1	0	1	1	1
31	0	1	1	1	1
32	1	1	1	1	1

The transition matrix for the preferential attachment and the transition matrix for the modified 2-hops (follower of follower) with $\epsilon = 0.1$ are given below:

6.6.2 A second numerical example

We study a more sophisticated situation with 6 agents, where the fixed choices are (1,2), (2,4), (3,5), (4,1), (5,2), (6,1) (see figure below). Suppose that the sets C_i of possible choices are:

$$C_1 = \{3\}, \quad C_2 = \{1, 6\}, \quad C_3 = \{2, 4, 6\}, \quad C_4 = \{2, 3, 5, 6\}, \quad C_5 = \{3, 4\}, \quad C_6 = \{2, 4, 5\}$$

and the out-degrees for agents 1 to 6 are respectively 2, 2, 3, 3, 3, 4. This means that agent 1 has only one choice (to form link (1,3)), agent 2 has 2 choices (agents 1 and 6), agent 3 has 3 choices (to choose 2 agents among agents 2, 4, 6), agent 4 has 6 choices (choose 2 agents among agents 2, 3, 5, 6), and agents 5 and 6 have only 1 choice, namely to take all their possible links. Therefore, the total number of states is $1 \times 2 \times 3 \times 6 \times 1 \times 1 = 36$, listed in the table below.

state number						
1	0	0	0	0	0	0
2	0	1	0	0	0	0
3	0	0	1	0	0	0
4	0	1	1	0	0	0
5	0	0	2	0	0	0
6	0	1	2	0	0	0
7	0	0	0	1	0	0
8	0	1	0	1	0	0
9	0	0	1	1	0	0
10	0	1	1	1	0	0
11	0	0	2	1	0	0
12	0	1	2	1	0	0
13	0	0	0	2	0	0
14	0	1	0	2	0	0
15	0	0	1	2	0	0
16	0	1	1	2	0	0
17	0	0	2	2	0	0
18	0	1	2	2	0	0
19	0	0	0	3	0	0
20	0	1	0	3	0	0
21	0	0	1	3	0	0
22	0	1	1	3	0	0
23	0	0	2	3	0	0
24	0	1	2	3	0	0
25	0	0	0	4	0	0
26	0	1	0	4	0	0
27	0	0	1	4	0	0
28	0	1	1	4	0	0
29	0	0	2	4	0	0
30	0	1	2	4	0	0
31	0	0	0	5	0	0
32	0	1	0	5	0	0
33	0	0	1	5	0	0
34	0	1	1	5	0	0
35	0	0	2	5	0	0
36	0	1	2	5	0	0

For PA, the transition matrix has values on the diagonal which are above 0.9, implying a slow process. The maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.977688, and when applying the Metropolis process (see subsection 5.4), this value changes to 0.98159. The stationary distribution is given on Figure 9. We can see




Figure 9: Stationary distribution Ex. 6.6.2 with PA

that the states 1, 7 and 25 are much more probable.

With FF, the diagonal in the transition matrix is less strong, implying a faster convergence. The maximum between the modulus of 2nd largest eigenvalue and modulus of smallest value is 0.973091, and when applying the Metropolis process, this value changes to 0.991725. The stationary distribution is given on Fig. 10. As before and even more strongly, states 1, 7 and 25




Figure 10: Stationary distribution Ex. 6.6.2 with FF

are much more probable. This can be explained as follows: considering fixed links, agents 3 and 6 have in-degree 0. Observe that agent 6 appears as a possible choice for agents 2, 3 and 4. As they have no incentive to choose agent 6, choices corresponding to avoid 6 (for agent 2: value 0; for agent 3: value 0; for agent 4: values 0, 1, 4) are more probable. As it can be checked in the above table, these correspond to states 1, 7 and 25.