



HAL
open science

The closed loop between opinion formation and personalised recommendations

Wilbert Samuel Rossi, Jan Willem Polderman, Paolo Frasca

► **To cite this version:**

Wilbert Samuel Rossi, Jan Willem Polderman, Paolo Frasca. The closed loop between opinion formation and personalised recommendations. *IEEE Transactions on Control of Network Systems*, 2022, 9 (3), pp.1092 - 1103. <10.1109/TCNS.2021.3105616>. <hal-03400149>

HAL Id: hal-03400149

<https://hal.science/hal-03400149v1>

Submitted on 24 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

The closed loop between opinion formation and personalised recommendations

Wilbert Samuel Rossi, Jan Willem Polderman and Paolo Frasca

Abstract—In online platforms, recommender systems are responsible for directing users to relevant content. In order to enhance the users’ engagement, recommender systems adapt their output to the reactions of the users, who are in turn affected by the recommended content. In this work, we study a tractable analytical model of a user that interacts with an online news aggregator, with the purpose of making explicit the feedback loop between the evolution of the user’s opinion and the personalised recommendation of content. More specifically, we assume that the user is endowed with a scalar opinion about a certain issue and receives news about it from a news aggregator: her opinion is influenced by all received pieces of news, which are characterized by a binary position on the issue at hand. The user is affected by a confirmation bias, that is, a preference for news that confirm her current opinion. The news aggregator recommends items with the goal of maximizing the number of user’s clicks (as a measure of her engagement): in order to fulfil its goal, the recommender has to compromise between exploring the user’s preferences and exploiting what it has learned so far. After defining suitable metrics for the effectiveness of the recommender systems (such as the click-through rate) and for its impact on the opinion, we perform both extensive numerical simulations and a mathematical analysis of the model. We find that personalised recommendations markedly affect the evolution of opinions and favor the emergence of more extreme ones: the intensity of these effects is inherently related to the effectiveness of the recommender. We also show that by tuning the amount of randomness in the recommendation algorithm, one can seek a balance between the effectiveness of the recommendation system and its impact on the opinions.

I. INTRODUCTION

Recommendation systems are ubiquitous in web services like social networking service and e-commerce platforms. Their purpose is sieving the available information and direct the user to the most relevant content. Recommendation systems leverage a wide array of machine learning techniques, which allow not only to quantify the absolute relevance of the items but also to tailor the recommendations to the expected tastes of the users, whose online behaviors are suitably recorded. Besides being monumental achievements of computer science, recommendation systems are essential to the

W.S. Rossi is with the Department of Sciences, University College Groningen, University of Groningen, 9718 BG Groningen, The Netherlands, w.s.rossi@rug.nl; J.W. Polderman is with the Department of Applied Mathematics, University of Twente, 7500 AE Enschede, The Netherlands, j.w.polderman@utwente.nl; P. Frasca is with Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, GIPSA-lab, 38000 Grenoble, France, paolo.frasca@gipsa-lab.fr.

This work has been partly supported by CNRS through grants S2IH PEPS “MOB” and 80|PRIME “DOOM” and by ANR through MIAI@Grenoble Alpes (ANR-19-P3IA-0003). The authors would like to thank all the colleagues who have provided useful comments on this research, including Doina Bucur, Alessandro Panconesi, Chiara Ravazzi, Tommaso Venturini, and the participants to the 2019 DBAED conference and the Polarisation seminar of the IAS-UvA.

user experience and access to information, news or purchase opportunities. As online activities become more and more prominent in people’s lives, questions are being asked about the effects (if any) of recommendation systems on the online and offline behaviors of the users. Even though the details of commercially deployed recommendation systems are kept secret and their operation is typically opaque to users, some of their key features are well known, such as being based on popularity and personalization. Our investigation specifically questions the role of personalization.

Concerns on the social impact of recommendation systems are specially perceived as relevant when it comes to the access to news. Indeed, also the news market shows a clear trend towards personalised information access. A number of personalised news aggregators like Google News and Apple News have emerged: these services collect news articles from different media outlets, rank them according to the estimated preferences of the user and propose her the curated collection [1], [2]. Users can explicitly set their interests, but the recommender systems also automatically construct user profiles by recording previous reading activities [3], [4]. In recommendation systems, personalisation enhances user experience, but political activists and scholars have raised concerns that excessive personalisation might narrow down the positions available to users about specific issues, effectively enclosing users into so-called “filter bubbles” that favour the emergence of opinion polarisation and radicalisation [5], [6]. Even though this concern has been downplayed by subsequent research [7], it is clear that personalization has at least the potential to reinforce the user’s idiosyncrasies and biases. Indeed, extensive research has shown that individuals are prone to *confirmation bias*. By this term we mean the unintentional tendency to acquire and process evidence that confirms one’s preconceptions and beliefs, possibly leading to an unconscious one-sided case-building process [8], [9], [10]. Since empirical evidence supports the idea that confirmation bias is extensive, strong and multiform, its effects may be amplified by curation algorithms.

The main purpose of this paper is to propose a tractable mathematical model of the interplay between a user and a learning system that produces personalized recommendations. This model will allow us to quantify the reciprocal reinforcement of confirmation bias and personalized curation. More specifically, we mathematically model the opinion formation process of a user that reads news from a personalized news aggregator. We restrict ourselves to pieces of news that bear implications for one specific issue, say, highlighting the benefits/drawbacks of immigration or arguing in favor/against US sanctions against Iran. News are characterized by a binary attribute that defines their positive or negative *position* on the

given issue. The opinion of the user evolves as an affine system that integrates the received news (actually, their positions) along time. Owing to the confirmation bias, news items are clicked upon with a probability that is larger when their position is closer to the user opinion. The recommender system has the objective of improving the engagement of the users, measured as the number of clicks. In order to achieve this purpose, the recommender follows a randomized strategy that balances “exploration”, that is, identifying which position is more appreciated by the user, and “exploitation”, that is, proposing contents that are most likely to be clicked on.

In view of our intention to obtain a tractable model, both user and recommender have closed-form mathematical descriptions that are very simple but at the same time include the key features and phenomena that are found in reality: users assimilate the information they receive and are prone to follow their confirmation biases, while recommendation systems record the actions of the users and personalize their recommendations to increase the users’ engagement. Further discussion on the soundness of our assumptions is provided in Section II after describing the model.

By combining extensive simulations and analytical results, we are able to obtain a clear picture of the interaction between user’s opinion and recommendation system. Our findings can be summarized by the following claims, which will be detailed and discussed in Section III: (i) Recommendations make opinions more extreme, due to the combination of personalisation and confirmation bias; (ii) More extreme opinions contribute positively to user engagement, and therefore to the benefit of the recommender system; (iii) The performance of the recommender, measured by improving the click-through rate, is entangled with its impact on the opinions; (iv) The effects of the recommender system can be mitigated by increasing its level of randomness, but this choice also reduces its performance, thereby introducing a relevant trade-off.

a) Literature review: Our paper is related to several recent works that have addressed recommendation systems or opinion dynamics by mathematical models. We survey some of the literature that has inspired our work, trying to emphasize similarities and differences.

A large literature has developed and studied mathematical models of opinion evolution. In this literature, opinions are cognitive orientation towards some objects, displayed attitudes, or subjective certainties of beliefs, and as such they can be quantified [11], [12], [13], [14], [15]. Individuals revise their opinions following social interactions [16] or after obtaining new information, which might confirm or challenge their views. Nowadays much social interactions and information seeking takes place online: experimental studies have demonstrated that online activities influence feelings and offline behaviours [17], [18], [19].

Incorporating models of online platforms in models of opinion dynamics is therefore of paramount importance [20], [13]. The paper [21] models opinion evolution on a social media platform that is able to favour the circulation of certain opinions over others, by selectively tuning their diffusion probabilities. Platform effects have also been explored by agent-based models in relation to filter bubbles [22] and

polarization [23]. In the filter bubble perspective, the platform could restrict the interactions of the users to a limited number of most similar individuals: our recent paper [24] investigates a simple opinion model based on this idea.

Coming to papers with a more defined focus on recommendation systems, [25], [26] look at systems that recommend products for purchase to a population of users. Both works study the evolution of the popularities of the products to identify potential distortions due to the recommender systems, but bear some relevant differences. First, [26] includes social ties between the users: these ties are shown to mitigate the effects of the recommender. Second, in [26] the probability that a product is recommended is proportional to its popularity, whereas in [25] that probability is non-linearly increasing in the popularity. As a consequence, only in the latter model the recommender system is able to distort the market and create artificial hits. This scenario is consistent with what we observe in our model, which is also non-linear in nature. Similar models of recommenders and socially connected users are studied in the recent paper [27] by means of simulations.

The paper [28] investigates the polarising effect on user opinions of collaborative recommender systems used to provide personalised suggestions of items (books in their example): the items have a binary attribute of which the recommender is completely agnostic. The paper compares three popular recommendation algorithms and analytically computes the probability that the next recommended item holds a specific attribute. The authors interpret the attribute share of the items owned by the user as her opinion and define as polarising an algorithm that makes the attribute share more uneven. The paper [29] adopts the same setting as [28], but investigates numerically the co-evolution of attribute shares and recommendations over a sequence of time steps. Different recommendation systems are also compared in [30], where the authors highlight that including randomness and growing the pool of items available for recommendation are important to mitigate the recommender’s negative effects.

More generally, the reciprocal influence between user and recommender is increasingly recognized as a key aspect in the analysis and design of recommendation systems [31], in evaluating both their impact on the users and their performance, which should be broadly defined to include objectives such as diversity, serendipity, novelty, or coverage. Compared to this growing literature, spanning from early studies like [32] (which was motivated by then-popular blog recommendations) to current research about filter bubbles in e-commerce [30], our distinguishing feature is the ability to analytically study the dynamics of user-recommender interconnection, as opposed to relying on simulations and data analysis.

b) Outline: In Section II, we describe our mathematical model that includes both the user and the recommender systems. We also discuss the experimental and theoretical backgrounds of our modeling choices. Section III contains the detailed analysis of the model, which exploits both numerical simulations and mathematical analysis. In Section IV, we summarize our main results and indicate some directions for future research.




Fig. 1. The closed loop between the user and the news aggregator. The diagram includes the variables exchanged by the two interacting dynamical systems, as well as their internal state variables.

II. MODEL: USER & RECOMMENDER IN INTERACTION

Our purpose is to mathematically model a user that interacts in closed loop with an online news aggregator, see Figure 1. The model that we are going to present is therefore made of two components: the *user model*, which includes the opinion dynamics and the confirmation bias, and the *news aggregator model* with the idealised recommender system. The user is endowed with a scalar signed opinion about a specific issue and receives news regarding the issue from an online news aggregator. The news aggregator proposes articles to the user, distinguishing between two antithetic positions (positive vs negative). The news aggregator adopts a recommender system to choose the articles to propose in a personalised way: the system tracks the clicks on the different headlines to understand user's preference and maximise her engagement, i.e., the number of clicks.

A. User model: opinion dynamics & click model

The user is endowed with a scalar *opinion* that evolves in discrete time

$$o_{\text{usr}} : \mathbb{N} \rightarrow [-1, 1],$$

and that represents her inclination about a given issue, as well as with a time-independent *prejudice* $o_{\text{usr}}^0 \in [-1, 1]$ that coincides with her initial opinion about the issue, i.e., $o_{\text{usr}}(0) = o_{\text{usr}}^0$. The prejudice encodes influences that are external to the recommendation system, such as the prior and continuing exposure to the opinions of relatives or trusted parties [33].

The user receives at each time step a news item, that is, an article, that supports a definite *position* about the issue at hand: the position can therefore take on opposite binary values

$$p_{\text{art}} : \mathbb{N} \rightarrow \{-1, +1\}.$$

At each time t , the user receives an article with position $p_{\text{art}}(t)$ and, upon receiving the recommended item, updates her opinion $o_{\text{usr}}(t)$ to

$$o_{\text{usr}}(t+1) = \alpha o_{\text{usr}}^0 + \beta o_{\text{usr}}(t) + \gamma p_{\text{art}}(t), \quad (1)$$

where α, β, γ are non-negative real scalars and $\alpha + \beta + \gamma = 1$ (that is, $o_{\text{usr}}(t+1)$ is a convex combination of o_{usr}^0 , $o_{\text{usr}}(t)$ and $p_{\text{art}}(t)$). The weights α, β and γ describe the relative importance of the prejudice, of the previous opinion (memory) and of the new information, respectively, in shaping the user's new opinion.

Beside updating her opinion, the user decides whether to read the recommended article or not, i.e., to click on its headline or not, following her interest. We model the $\{0, 1\}$ *click decision*, meaning $\{\text{no click}, \text{click}\}$, as a stochastic process such that at each time the decision is a Bernoulli random variable with opinion-dependent parameter $\theta(o_{\text{usr}}, p_{\text{art}})$, i.e.,

$$\text{clk}(t) \sim \text{Bernoulli}(\theta(o_{\text{usr}}(t), p_{\text{art}}(t))).$$

Function $\theta : [-1, 1] \times \{-1, +1\} \rightarrow [0, 1]$ quantifies the subjective interest $\theta(o_{\text{usr}}, p_{\text{art}})$ of a user with opinion o_{usr} in an item with headline of position p_{art} . The user is subject to a confirmation bias [8] and prefers contents that are consistent with her opinion o_{usr} : we model this fact by choosing function

$$\theta(o_{\text{usr}}, p_{\text{art}}) = \frac{1}{2} + \frac{1}{2} o_{\text{usr}} p_{\text{art}}. \quad (2)$$

After making her click choice (and possibly reading the article), the user moves on to the next recommended item, that is, the dynamics moves on to the next time step.

a) *On the justification of the user model:* In our model, opinions belong to the interval $[-1, 1]$, which are sometimes referred to as *polar* opinions and indicate the degree of proclivity towards one or two competing alternatives [34]. The opinion dynamics assimilate the information that the user receives. The specific linear model (1) is very close to various models proposed both by sociologists and economists in the context of naive learning [35], [36], [37] and of opinion formation under social influence [33]. Assuming that the evolution of opinions follows a convex (and therefore, positive) combination is consistent with observations in experimental social psychology: for both live [38] and online interactions [39], that after the interaction opinions get closer than before. In our case, the stream of news items is the source of information and is assimilated by the user [35].

In modeling the effect of the news item, we make a twofold assumption: we assume that the article's position about the issue is binary in nature and that the user revises her opinion upon being recommended the item, that is, without the need to read the full article. In fact, we assume that the article's headline is sufficient to influence the user to revise her opinion. The binary nature of the news items, which is also featured in similar recent models [37], corresponds well to situations where news outlets take partisan perspectives [9]. The assumption that receiving the headline is sufficient to be influenced is consistent with the heuristic model of persuasion by [40], based on the observation that opinion changes are often the outcome of minimal amounts of information and superficial judgements. There is however another, perhaps more compelling, reason to make this assumption: this assumption allows us to focus more precisely our analysis on the effect of closed-loop recommendations. Since the user is influenced by all received news items, she effectively has no confirmation bias as per her process of information assimilation and opinion update, whereas her bias is made apparent to the recommender system through her clicking history. Informally speaking, we may say that in our model all bias in the dynamics of the user opinion is only due to the effect of *biased recommendations*: unbiased recommendations

would result in an unbiased opinion and the confirmation bias only bears effects through the recommendation system. Our assumption is therefore conservative in nature: assuming the user to be influenced only after reading the articles can only result in a stronger bias on the opinion.

Finally, the specific expression (2) is supported by the literature as it translates into our framework the definition of *biased user* by [28, Def. 9]. Moreover, we observe that it satisfies several common-sense properties:

- $\theta(o_{\text{usr}}, +1) \geq \theta(o_{\text{usr}}, -1)$ for $o_{\text{usr}} > 0$ while $\theta(o_{\text{usr}}, +1) \leq \theta(o_{\text{usr}}, -1)$ for $o_{\text{usr}} < 0$: the interest of the user is higher for articles that have a position closer to her opinion;
- $\theta(o_{\text{usr}}, +1)$ is non-decreasing in o_{usr} while $\theta(o_{\text{usr}}, -1)$ is non-increasing in o_{usr} : if the match between the opinion of the user and the position of the article increases, the interest of the user does not decrease;
- $\theta(o_{\text{usr}}, +1) = \theta(-o_{\text{usr}}, -1)$: the interest is symmetric in the opinion-position match;
- $\theta(+1, +1) = \theta(-1, -1) = 1$: complete alignment between opinion and position makes the click almost certain.

B. News aggregator model: 2ϵ -greedy recommender system

The news aggregator has the purpose of *maximising the user engagement*, measured by the *ratio of clicks* on the suggested content. At each time step the news aggregator has to choose whether to recommend an article with “positive” or “negative” position. The online service tracks the user’s activities by logging clicks and estimates the user’s interest in a position as the empirical probability that the user clicks on an article bearing that position. The estimated user interests are then used to generate the recommendation [3]. In making the recommendation, the recommender system faces the *exploration-exploitation dilemma* of sequential decision problems [41], which arises between staying with the most successful option so far and exploring the other option, which might turn better in the future: this dilemma has been specifically identified in the context of news recommendations [42]. Moreover, as users’ interests change over time, the system needs to incrementally update the user’s profile to reflect such changes. We adopt for the recommender system an *2ϵ -greedy algorithm*: with probability 2ϵ , the recommender system randomly explores the binary options; with probability $1 - 2\epsilon$, the recommender system recommends the most successful option so far. This approach, although not optimal, is supported by the literature for time-varying settings like ours [43].

Let us now describe in details how the system keeps track of past clicks to learn the most successful choice. The sets $T_+(t)$ and $T_-(t)$ collect the time steps until t at which an article with position $+1$ or -1 , respectively, has been recommended:

$$\begin{aligned} T_+(t) &= \{s : 0 \leq s \leq t-1 \text{ and } p_{\text{art}}(s) = +1\}, \\ T_-(t) &= \{s : 0 \leq s \leq t-1 \text{ and } p_{\text{art}}(s) = -1\}. \end{aligned}$$

The cardinalities of the above sets

$$r_+(t) = \#T_+(t), \quad r_-(t) = \#T_-(t),$$

count how many times an article with position $+1$ or -1 has been proposed until time t . The counters $a_+(t)$ and $a_-(t)$ record how many times until time t the user has accepted a recommendation with positions $+1$ and -1 , respectively,

$$a_+(t) = \sum_{s \in T_+(t)} \text{clk}(s), \quad a_-(t) = \sum_{s \in T_-(t)} \text{clk}(s).$$

According to this discussion, we shall evaluate the performance of the recommender systems by the *click-through rate*

$$\text{ctr}(t) = \frac{a_+(t) + a_-(t)}{t}. \quad (3)$$

We observe that, by their definition, the integer sequences $r_+(t)$, $r_-(t)$, $a_+(t)$ and $a_-(t)$ have null initial value and evolve according to

$$\begin{aligned} r_+(t+1) &= r_+(t) + 1 & a_+(t+1) &= a_+(t) + \text{clk}(t) \\ r_-(t+1) &= r_-(t) & a_-(t+1) &= a_-(t) \end{aligned}$$

if $p_{\text{art}}(t) = +1$ and

$$\begin{aligned} r_+(t+1) &= r_+(t) & a_+(t+1) &= a_+(t) \\ r_-(t+1) &= r_-(t) + 1 & a_-(t+1) &= a_-(t) + \text{clk}(t) \end{aligned}$$

otherwise.

At any time $t \geq 2$ the recommender system computes the ratios $\frac{a_+(t)}{r_+(t)}$ and $\frac{a_-(t)}{r_-(t)}$, i.e., the proportions of clicks collected by the positions $+1$ and -1 , respectively, to estimate the future odds of collecting a click upon proposing the positions $+1$ and -1 . Then the recommender system proposes with probability $1 - \epsilon$ an article representing the position that has received the higher proportion of clicks and with the complementary probability ϵ the other position. Should a tie arise, both positions receive equal probabilities. In formulas, the decision rule reads

$$\mathbb{P}(p_{\text{art}}(t) = +1) = \begin{cases} 1 - \epsilon & \text{if } \frac{a_+(t)}{r_+(t)} > \frac{a_-(t)}{r_-(t)} \\ \frac{1}{2} & \text{if } \frac{a_+(t)}{r_+(t)} = \frac{a_-(t)}{r_-(t)} \\ \epsilon & \text{if } \frac{a_+(t)}{r_+(t)} < \frac{a_-(t)}{r_-(t)} \end{cases} \quad (4)$$

The parameter ϵ controls the trade-off between exploration and exploitation and is typically small; in most numerical simulations we shall take $\epsilon = 0.05$. At times 0 and $+1$ the recommender system follows an initialisation procedure by which positions are proposed in a random order, i.e., $\mathbb{P}(p_{\text{art}}(0) = +1, p_{\text{art}}(1) = -1) = \mathbb{P}(p_{\text{art}}(0) = -1, p_{\text{art}}(1) = +1) = 0.5$, so that $r_+(2) = r_-(2) = 1$.

III. ANALYSIS OF THE CLOSED-LOOP DYNAMICS

We now move on to present our findings about the dynamical model described in the previous section. Our presentation is divided into two main parts. In the first part, we observe that typical trajectories of the dynamical model are characterized by a definite majority of either $+1$ or -1 recommendations. This observation supports the study of the expected dynamics *conditioned upon a given majority*: these conditional expectations can be calculated in closed form and turn out to describe the stochastic dynamics very accurately. In the second part, we build on these formal derivations and on extended simulations

to discover the effects of recommendations on the evolution of the opinions.

We begin by writing the closed-loop dynamics that summarizes the model in the previous section. By defining the state vector $\mathbf{x}(t) := [r_+(t), r_-(t), a_+(t), a_-(t), o_{\text{usr}}(t)]^\top$, the dynamics¹ can indeed be written as

$$\begin{cases} \mathbf{x}(t+1) = A\mathbf{x}(t) + \mathbf{f}(\mathbf{x}(t)) & t \geq 2 \\ \mathbf{x}(2) \text{ discrete random variable,} \end{cases} \quad (5)$$

where A is the update matrix

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \beta \end{bmatrix}$$

and the vector $\mathbf{f}(\mathbf{x}(t))$ is a random variable that can take on four values, corresponding to the cases “position +1, no click”, “position +1, get click”, “position -1, no click” and “position -1, get click”. In the recommender rule (4), the probability of each of these cases depends non-linearly on the state $\mathbf{x}(t)$. To encode these probabilities, we define the difference

$$\Delta(\mathbf{x}(t)) := \frac{a_+(t)}{r_+(t)} - \frac{a_-(t)}{r_-(t)}. \quad (6)$$

Clearly, $\Delta(\mathbf{x}(t)) > 0$ if and only if +1 is the most successful recommendation so far. Using the modified step function $h_\epsilon: \mathbb{R} \rightarrow \{\epsilon, \frac{1}{2}, 1 - \epsilon\}$

$$h_\epsilon(s) := \begin{cases} 1 - \epsilon & \text{if } s > 0 \\ \frac{1}{2} & \text{if } s = 0 \\ \epsilon & \text{if } s < 0, \end{cases}$$

the probabilistic model of $\mathbf{f}(\mathbf{x})$ is $\mathbf{f}(\mathbf{x}) =$

$$\begin{cases} [1, 0, 1, 0, \alpha o_{\text{usr}}^0 + \gamma]^\top & \text{with prob. } h_\epsilon(\Delta) \theta(o_{\text{usr}}, +1) \\ [1, 0, 0, 0, \alpha o_{\text{usr}}^0 + \gamma]^\top & \text{with prob. } h_\epsilon(\Delta) (1 - \theta(o_{\text{usr}}, +1)) \\ [0, 1, 0, 1, \alpha o_{\text{usr}}^0 - \gamma]^\top & \text{with prob. } (1 - h_\epsilon(\Delta)) \theta(o_{\text{usr}}, -1) \\ [0, 1, 0, 0, \alpha o_{\text{usr}}^0 - \gamma]^\top & \text{with prob. } (1 - h_\epsilon(\Delta)) (1 - \theta(o_{\text{usr}}, -1)) \end{cases}$$

A. Majority trajectories: numerical evidences and analysis

To showcase the evolution of the closed-loop dynamics (5), we present in detail a set of simulations in Figures 2, 3 and 4. An immediate observation is that the dynamics are shaped by a prevalence of +1 or -1 recommendations. For instance, the simulation in Figure 2 is clearly characterized by having $p_{\text{art}}(t) = +1$ most of the times. Correspondingly, the ratio $a_+(t)/r_+(t)$ is larger and more stable than the ratio $a_-(t)/r_-(t)$. These facts are consistent with the opinion $o_{\text{usr}}(t)$ being always positive (most often around 0.65). We refer to this feature as being a +1-majority trajectory. Conversely, Figure 3 shows a -1-majority trajectory (left) and a trajectory that takes some time to define itself as +1-majority (right). To highlight the effects of the recommender system,

¹For times $t = 0, 1$ the system is driven by the stochastic initialization of the recommender system. The specific form of $\mathbf{x}(2)$ can be easily derived but is not relevant to our analysis that focuses on long-term behaviors. We therefore avoid to report it.




Fig. 2. Simulation with $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $o_{\text{usr}}^0 = 0.30$ and recommended choice of $p_{\text{art}}(t)$ with $\epsilon = 0.05$. The most recommended position is +1.

these dynamics should be compared against Figure 4, where the sequence of recommended positions is purely random, that is, $\epsilon = \frac{1}{2}$. We can observe that in this case the click-through rate $\text{ctr}(t)$ is much smaller than in Figure 2.

With the support of these initial insights from the simulations, we proceed to develop a mathematical analysis of the stochastic dynamics (5). To this purpose, it is natural to calculate the conditional expectation

$$\begin{aligned} \mathbb{E}[\mathbf{x}(t+1)|\mathbf{x}(t)] &= \mathbb{E}[A\mathbf{x}(t) + \mathbf{f}(\mathbf{x}(t))|\mathbf{x}(t)] \\ &= \begin{bmatrix} r_+(t) \\ r_-(t) \\ a_+(t) \\ a_-(t) \\ \beta o_{\text{usr}}(t) \end{bmatrix} + \begin{bmatrix} h_\epsilon(\Delta(\mathbf{x}(t))) \\ 1 - h_\epsilon(\Delta(\mathbf{x}(t))) \\ h_\epsilon(\Delta(\mathbf{x}(t))) \theta(o_{\text{usr}}(t), +1) \\ (1 - h_\epsilon(\Delta(\mathbf{x}(t)))) \theta(o_{\text{usr}}(t), -1) \\ \alpha o_{\text{usr}}^0 + \gamma (2h_\epsilon(\Delta(\mathbf{x}(t))) - 1) \end{bmatrix} \end{aligned} \quad (7)$$

for $t \geq 2$. In principle, one could now derive the evolution of $\mathbb{E}[\mathbf{x}(t)]$ by taking the expectation of both sides: however, we refrain from this operation for two reasons. Not only the calculation is made difficult by the function $h_\epsilon(\Delta(\mathbf{x}))$ being non-linear, but is also unlikely to bring significant insight, as we explain next.

Simulations suggest that every trajectory is shaped by an overwhelming majority of either +1 or -1 recommendations: either choice yields a definite behavior. These two distinct groups of solutions would not be described well by an average dynamics: we should instead try to distinguish these two groups. Motivated by this observation, we want to define a suitable pair of dynamics and, to this purpose, we study the process conditional on having either $\Delta(\mathbf{x}(t)) > 0$ or $\Delta(\mathbf{x}(t)) < 0$ for $t \geq 2$. We thus introduce the notation

$$\begin{aligned} \mathbb{E}^+[\mathbf{x}(t)] &:= \mathbb{E}[\mathbf{x}(t) | \Delta(\mathbf{x}(s)) > 0 \text{ for } 2 \leq s < t], \\ \mathbb{E}^-[\mathbf{x}(t)] &:= \mathbb{E}[\mathbf{x}(t) | \Delta(\mathbf{x}(s)) < 0 \text{ for } 2 \leq s < t] \end{aligned}$$

and we proceed with the calculations for case $\Delta(\mathbf{x}(t)) > 0$, which means $h_\epsilon(\Delta(\mathbf{x}(t))) = 1 - \epsilon$. We recall that in this

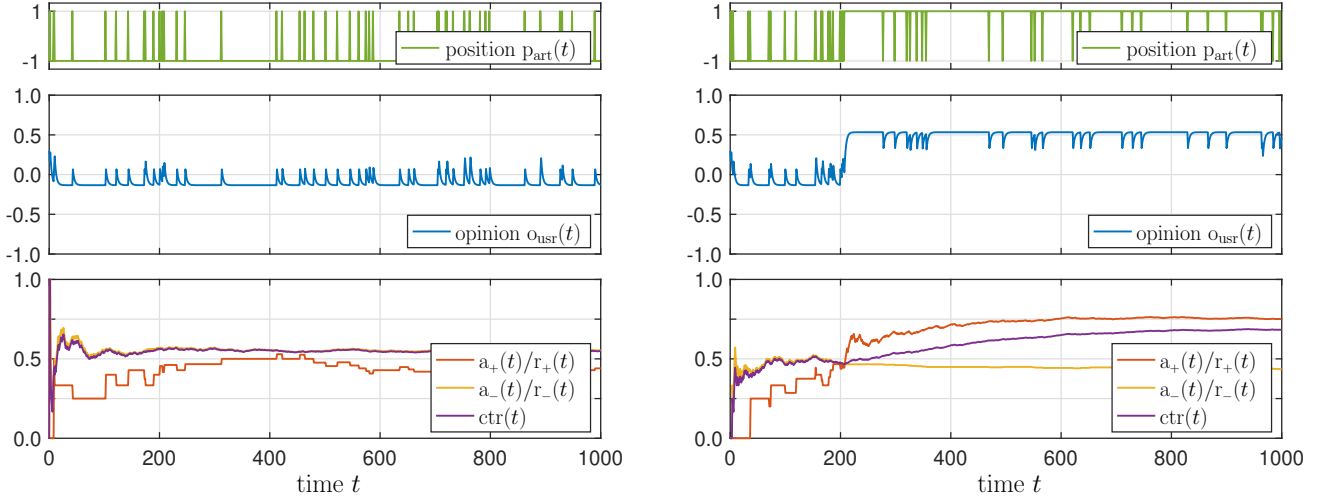


Fig. 3. Simulations with $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $o_{\text{usr}}^0 = 0.30$ and recommended choice of $p_{\text{art}}(t)$ with $\epsilon = 0.05$. Left: a -1 -majority trajectory. Right: a $+1$ -majority trajectory that features a majority of -1 s in its initial phases.




Fig. 4. Simulation with $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $o_{\text{usr}}^0 = 0.30$ and random choice of $p_{\text{art}}(t)$ (i.e $\epsilon = 0.50$).

scenario the most favourable option is recommending articles with position $+1$. Equations (7) and (2) give

$$\begin{aligned} \mathbb{E}^+[\mathbf{x}(t+1)|\mathbf{x}(t)] &= \begin{bmatrix} r_+(t) \\ r_-(t) \\ a_+(t) \\ a_-(t) \\ \beta o_{\text{usr}}(t) \end{bmatrix} + \begin{bmatrix} 1 - \epsilon \\ \epsilon \\ (1 - \epsilon)\theta(o_{\text{usr}}(t), +1) \\ \epsilon\theta(o_{\text{usr}}(t), -1) \\ \alpha o_{\text{usr}}^0 + \gamma(1 - 2\epsilon) \end{bmatrix} \\ &= \begin{bmatrix} r_+(t) + 1 - \epsilon \\ r_-(t) + \epsilon \\ a_+(t) + \frac{1}{2}(1 - \epsilon)o_{\text{usr}}(t) + \frac{1}{2}(1 - \epsilon) \\ a_-(t) - \frac{1}{2}\epsilon o_{\text{usr}}(t) + \frac{1}{2}\epsilon \\ \beta o_{\text{usr}}(t) + \alpha o_{\text{usr}}^0 + \gamma(1 - 2\epsilon) \end{bmatrix} \end{aligned}$$

We take the expectation a second time and get

$$\mathbb{E}^+[\mathbf{x}(t+1)] = B_{\text{R}}^+ \mathbb{E}^+[\mathbf{x}(t)] + \mathbf{b}_{\text{R}}^+ \quad (8)$$

where

$$B_{\text{R}}^+ = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & \frac{1}{2}(1 - \epsilon) \\ 0 & 0 & 0 & 1 & -\frac{1}{2}\epsilon \\ 0 & 0 & 0 & 0 & \frac{1}{\beta} \end{bmatrix} \quad \text{and} \quad \mathbf{b}_{\text{R}}^+ = \begin{bmatrix} 1 - \epsilon \\ \epsilon \\ \frac{1}{2}(1 - \epsilon) \\ \frac{1}{2}\epsilon \\ \alpha o_{\text{usr}}^0 + \gamma(1 - 2\epsilon) \end{bmatrix}.$$

The quantity $\mathbb{E}^+[o_{\text{usr}}(t)]$ has an autonomous dynamics and converges to

$$\lim_{t \rightarrow \infty} \mathbb{E}^+[o_{\text{usr}}(t)] = \frac{\alpha o_{\text{usr}}^0 + \gamma(1 - 2\epsilon)}{\alpha + \gamma}. \quad (9)$$

The analogous calculations for $\Delta(\mathbf{x}(t)) < 0$ yield

$$\lim_{t \rightarrow \infty} \mathbb{E}^-[o_{\text{usr}}(t)] = \frac{\alpha o_{\text{usr}}^0 - \gamma(1 - 2\epsilon)}{\alpha + \gamma}; \quad (10)$$

these two limit values differ by

$$\lim_{t \rightarrow \infty} \mathbb{E}^+[o_{\text{usr}}(t)] - \lim_{t \rightarrow \infty} \mathbb{E}^-[o_{\text{usr}}(t)] = 2 \frac{\gamma}{\alpha + \gamma} (1 - 2\epsilon), \quad (11)$$

which does not depend on o_{usr}^0 . This quantity, which we call *discrepancy*, measures the effect of recommendations on opinions and corresponds to the vertical gap between the green and magenta lines in Figure 5 (right). This figure confirms that trajectories split into $+1$ -trajectories and -1 -trajectories, concentrating around the conditional expectations.

B. Recommender dynamics and click-through rate

Having understood the evolution of the user dynamics, we can proceed to study the internal variables of the recommender, which follow

$$\begin{aligned} \mathbb{E}^+[r_+(t+1)] - \mathbb{E}^+[r_+(t)] &= 1 - \epsilon \\ \mathbb{E}^+[r_-(t+1)] - \mathbb{E}^+[r_-(t)] &= \epsilon \\ \mathbb{E}^+[a_+(t+1)] - \mathbb{E}^+[a_+(t)] &= \frac{1}{2}(1 - \epsilon) + \frac{1}{2}(1 - \epsilon)\mathbb{E}^+[o_{\text{usr}}(t)] \\ \mathbb{E}^+[a_-(t+1)] - \mathbb{E}^+[a_-(t)] &= \frac{1}{2}\epsilon - \frac{1}{2}\epsilon\mathbb{E}^+[o_{\text{usr}}(t)]. \end{aligned}$$




Fig. 5. The average opinion $\overline{o_{\text{usr}}}(t_{\text{max}})$ with $t_{\text{max}} = 5000$ corresponding to simulations with $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$ and $o_{\text{usr}}^0 \in \{-1.00, -0.90, \dots, 1.00\}$. There are 1000 simulations for each prejudice o_{usr}^0 . Left: random choice of $p_{\text{part}}(t)$ (i.e. $\epsilon = 0.50$). Right: recommended choice of $p_{\text{part}}(t)$ with $\epsilon = 0.05$; +1-majority and -1-majority simulations are distinguished; the equations of the magenta and green line are (9) and (10), respectively.

Therefore, for large t we have

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}^+[r_+(t)]}{t} = 1 - \epsilon \quad (12a)$$

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}^+[r_-(t)]}{t} = \epsilon \quad (12b)$$

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}^+[a_+(t)]}{t} = \frac{1}{2}(1 - \epsilon)(1 + \mathbb{E}^+[o_{\text{usr}}(\infty)]) \quad (12c)$$

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}^+[a_-(t)]}{t} = \frac{1}{2}\epsilon(1 - \mathbb{E}^+[o_{\text{usr}}(\infty)]) \quad (12d)$$

and consequently, the expected click-through rate (3) becomes

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}^+[\text{ctr}(t)] &= \frac{1}{2} + \frac{1}{2}(1 - 2\epsilon)\mathbb{E}^+[o_{\text{usr}}(\infty)] \quad (13) \\ &= \frac{1}{2} + \frac{1}{2}(1 - 2\epsilon)\frac{\alpha o_{\text{usr}}^0 + \gamma(1 - 2\epsilon)}{\alpha + \gamma}. \end{aligned}$$

An analogous calculation gives

$$\lim_{t \rightarrow \infty} \mathbb{E}^-[\text{ctr}(t)] = \frac{1}{2} - \frac{1}{2}(1 - 2\epsilon)\frac{\alpha o_{\text{usr}}^0 - \gamma(1 - 2\epsilon)}{\alpha + \gamma}. \quad (14)$$

Observe that in both cases the click-through rate achieved by the recommender system is larger than the click-through rate with the random choice $\epsilon = 0.5$. The difference between the asymptotic click-through rates is

$$\mathbb{E}^+[\text{ctr}(\infty)] - \mathbb{E}^-[\text{ctr}(\infty)] = \frac{\alpha}{\alpha + \gamma} o_{\text{usr}}^0 (1 - 2\epsilon).$$

This quantity is proportional to $1 - 2\epsilon$ and to o_{usr}^0 : intuitively, recommending +1 is more rewarding for larger o_{usr}^0 . This last comment leads us to look more carefully at the link between prejudices and recommendations.

For $\Delta(\mathbf{x}(t))$ to be positive, one expects $o_{\text{usr}}(t)$ to be positive: indeed, by (12) we have that $\lim_{t \rightarrow \infty} \mathbb{E}^+[\Delta(\mathbf{x}(t))] = \lim_{t \rightarrow \infty} \frac{\mathbb{E}^+[a_+(t)]}{\mathbb{E}^+[r_+(t)]} - \frac{\mathbb{E}^+[a_-(t)]}{\mathbb{E}^+[r_-(t)]} = \mathbb{E}^+[o_{\text{usr}}(\infty)]$. Using (9), the condition $\mathbb{E}^+[o_{\text{usr}}(\infty)] > 0$ reads

$$o_{\text{usr}}^0 > -\frac{\gamma}{\alpha}(1 - 2\epsilon), \quad (15)$$

i.e., o_{usr}^0 should not be too negative. Any o_{usr}^0 that satisfies this condition is compatible with $\Delta(\mathbf{x}(t)) > 0$ and +1-majority trajectories. In the scenario $\Delta(\mathbf{x}(t)) < 0$, analogous calculations lead to

$$o_{\text{usr}}^0 < \frac{\gamma}{\alpha}(1 - 2\epsilon). \quad (16)$$

Together, conditions (15) (16) split the interval $[-1, 1]$ of possible prejudices o_{usr}^0 into three parts:

- A $o_{\text{usr}}^0 < -\frac{\gamma}{\alpha}(1 - 2\epsilon)$: only $\Delta(\mathbf{x}(t)) < 0$ seems possible and we should observe only -1-majority trajectories;
- B $-\frac{\gamma}{\alpha}(1 - 2\epsilon) \leq o_{\text{usr}}^0 \leq \frac{\gamma}{\alpha}(1 - 2\epsilon)$: both $\Delta(\mathbf{x}(t)) < 0$ and $\Delta(\mathbf{x}(t)) > 0$ are possible, allowing both +1-majority and -1-majority trajectories;
- C $o_{\text{usr}}^0 > \frac{\gamma}{\alpha}(1 - 2\epsilon)$: only $\Delta(\mathbf{x}(t)) > 0$ seems possible and we should observe only +1-majority trajectories.

This reasoning explains simulations in Figures 5 (right) and 6 and confirms that *strong prejudices lead to recommendations that are consistent with the prejudices*.

C. The impact of recommendations on opinions

We are now ready to make some important observations about the effects of the recommendations on the evolution of the opinions. Statistically, recommendations have a significant polarizing effect and this effect is closely related to their effectiveness in increasing the click-through rate.

Our first observation is that *most trajectories produce opinions that are more extreme than the prejudices*, that is, they exhibit some polarization effect. In Figure 7 the shaded areas correspond to opinions that got milder during their evolution, i.e., $|\overline{o_{\text{usr}}}(t_{\text{max}})| \leq |o_{\text{usr}}^0|$. It is clear that a large majority of the trajectories falls in the non-shaded areas, thereby implying that for most choices of the parameters, recommendations make opinions more extreme, that is, contribute to polarize opinions towards either -1 or +1.

Our second observation is that *recommendations are more effective when opinions are more extreme*. Figure 8 represents




Fig. 6. The empirical probability of +1-majority and -1-majority trajectory, with the majority character identified by $\bar{p}_{\text{part}}(t_{\text{max}})$. Dashed blue lines have abscissas (15) (16). Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $\epsilon = 0.05$. For each prejudice o_{usr}^0 there are 1000 simulations with $t_{\text{max}} = 5000$.

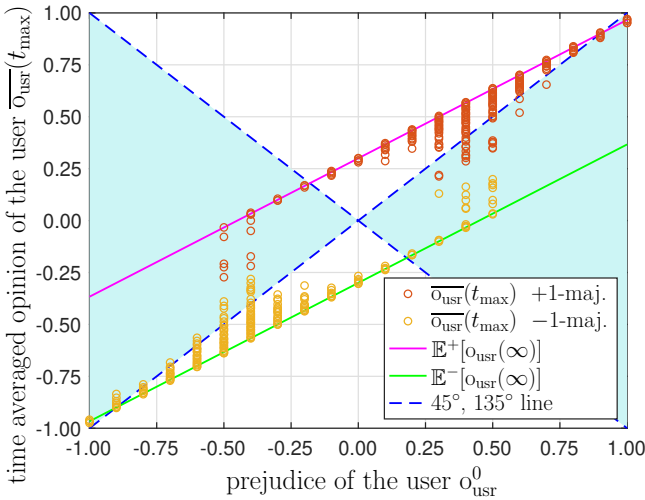


Fig. 7. Plot of final opinions against prejudices (initial opinions). In shaded areas, the time averaged opinion $\bar{o}_{\text{usr}}(t_{\text{max}})$ is *less extreme* than the prejudice o_{usr}^0 , i.e., $|\bar{o}_{\text{usr}}(t_{\text{max}})| \leq |o_{\text{usr}}^0|$; in non-shaded areas, it is *more extreme*. Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $\epsilon = 0.05$, 1000 simulations.

the click-through rate $\text{ctr}(t_{\text{max}})$ versus the time averaged opinion $\bar{o}_{\text{usr}}(t_{\text{max}})$. We observe that the dashed lines in the plot, with expressions (13) and (14), correspond to the two possible weighted averages of (2), i.e. $(1-\epsilon)\theta(o_{\text{usr}}, +1) + \epsilon\theta(o_{\text{usr}}, -1)$ and $\epsilon\theta(o_{\text{usr}}, +1) + (1-\epsilon)\theta(o_{\text{usr}}, -1)$ and that the realizations concentrate on theoretical predictions in their portions above 0.50. These simulations confirm that the recommender system increases the ratio of collected clicks and that this ability is enhanced by extreme opinions.

Our third observation sheds more light on the connection between recommender systems and polarizing effect. Not only we have already observed that recommendations are more effective when opinions are extreme, but actually we find an explicit *correlation between the effectiveness of recommendations and their impact on opinions*. Indeed, the average click-




Fig. 8. Click-through rate $\text{ctr}(t_{\text{max}})$ with respect to time-averaged opinion $\bar{o}_{\text{usr}}(t_{\text{max}})$, with the simulations distinguished by their +1-majority and -1-majority character. The equations of the shaded black and cyan line are (13) and (14), respectively. Parameters: $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $\epsilon = 0.05$, $o_{\text{usr}}^0 \in \{-1.00, -0.90, \dots, 1.00\}$; for each prejudice o_{usr}^0 we show 1000 realizations with $t_{\text{max}} = 5000$

through rate is such that

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{2} (\mathbb{E}^+[\text{ctr}(t)] + \mathbb{E}^-[\text{ctr}(t)]) &= \frac{1}{2} + \frac{1}{2} \frac{\gamma}{\alpha + \gamma} (1 - 2\epsilon)^2 \\ &= \frac{1}{2} + \frac{1}{4} (1 - 2\epsilon) \lim_{t \rightarrow \infty} (\mathbb{E}^+[o_{\text{usr}}(t)] - \mathbb{E}^-[o_{\text{usr}}(t)]) \end{aligned} \quad (17)$$

where we used (11) to recognise the *discrepancy*. This average is relevant if the prejudice o_{usr}^0 is zero, as in this case the trajectories are equally likely to be +1-majority and -1-majority. Figure 9 plots the sample average of the click-through rates $\text{ctr}(t_{\text{max}})$ against the *discrepancy*. The simulations with $o_{\text{usr}}^0 = 0.00$ match expression (17), while the simulations with $o_{\text{usr}}^0 = 0.33$ are distributed above the former and confirm that the *click-through rate is monotonically increasing with the discrepancy*.

Our fourth and final observation is that the tuning parameter ϵ has a crucial influence on both the effectiveness of the recommender systems and on its side effects, and can be used to mitigate them. Indeed, ϵ controls the amount of randomness injected in the system. If $\epsilon = 0.5$, we have random recommendations that propose +1 and -1 with equal probabilities. For a value of ϵ in $(0, \frac{1}{2})$, the algorithm favors the option that stimulated a larger interest in the user; this bias gets larger for the smaller ϵ . According to our analysis, this recommendation bias produces an average user opinion that differs from its unbiased counterpart ($\epsilon = 0.5$) by an *opinion distortion*

$$\begin{aligned} \Delta_{\text{usr}}^{\pm} &:= \mathbb{E}^{\pm}[o_{\text{usr}}(\infty); \epsilon] - \mathbb{E}^{\pm}[o_{\text{usr}}(\infty); \epsilon = 0.5] \\ &= \pm \frac{\gamma}{\alpha + \gamma} (1 - 2\epsilon), \end{aligned} \quad (18)$$

which turns out to be half of the *discrepancy*. At the same time, the recommendation bias produces a click-through rate




Fig. 9. Sample average of the click-through rate $\text{ctr}(t_{\max})$ against the discrepancy. The different values of the discrepancy, defined in (11), have been obtained with 116 different combinations of the parameters α, β, γ , while keeping $\epsilon = 0.05$. Each cross represents 1000 simulations with $t_{\max} = 1000$.

larger than in the unbiased case by a *click-through rate gain*

$$\begin{aligned} \Gamma_{\text{ctr}}^{\pm} &:= \mathbb{E}^{\pm}[\text{ctr}(\infty); \epsilon] - \mathbb{E}^{\pm}[\text{ctr}(\infty); \epsilon = 0.5] \\ &= \pm \frac{1}{2} \frac{\alpha}{\alpha + \gamma} o_{\text{usr}}^0 (1 - 2\epsilon) + \frac{1}{2} \frac{\gamma}{\alpha + \gamma} (1 - 2\epsilon)^2. \end{aligned} \quad (19)$$

Both these quantities are monotonically increasing in $(1 - 2\epsilon)$. Hence, our analysis suggests that *mitigating the impact of the recommender system on the opinions has a price in terms of the achievable click-through rate*. To make this connection clearer, we can combine (18) and (19) and deduce

$$\Gamma_{\text{ctr}}^{\pm} = \frac{1}{2} \frac{\alpha}{\gamma} o_{\text{usr}}^0 \Delta_{\text{usr}}^{\pm} + \frac{1}{2} \frac{\alpha + \gamma}{\gamma} (\Delta_{\text{usr}}^{\pm})^2, \quad (20)$$

thereby showing that the relation between the the opinion distortion $\Delta_{\text{usr}}^{\pm}$ and the click-through rate gain $\Gamma_{\text{ctr}}^{\pm}$ is independent of ϵ and therefore depends only on the characteristics of the user. This relation is well matched by the simulations in Figure 10.

IV. CONCLUSION

In this paper, we have proposed a mathematical model of the interaction between a user and an online service providing personalized recommendations. The model is simple enough to allow for its analytical treatment, while being rich enough to include the main features of the user-platform interaction, which we identified on the user side as information assimilation and confirmation bias, and on the recommender side as measuring the user’s engagement and learning through exploration-exploitation. Our results evaluate the effects of personalization on the evolution of user’s opinion, showing that personalized recommendations typically drive users towards more extreme opinions.

In its simplicity, our model is not meant to quantitatively simulate the actual evolution of user opinions. Instead, its purpose is to qualitatively show which consequences can originate from certain features of the recommendation systems.




Fig. 10. Empirical click-through rate gain against empirical opinion distortion, with the simulations distinguished by their +1-majority and -1-majority character. The empirical “ctr gain” is $[\text{ctr}(t_{\max}); \epsilon]$ minus the sample average of $[\text{ctr}(t_{\max}); \epsilon = 0.5]$. The empirical “op. distortion” is $[\overline{o}_{\text{usr}}(t_{\max}); \epsilon]$ minus the sample average of $[\overline{o}_{\text{usr}}(t_{\max}); \epsilon = 0.5]$. In the simulations, $\alpha = 0.20$, $\beta = 0.70$, $\gamma = 0.10$, $o_{\text{usr}}^0 = 0.33$ and $t_{\max} = 5000$.

While we believe that our model is informative enough to make it relevant in the heating debate on the impact of social media platforms on our societies, we are well aware of its limitations. In a nutshell, our model describes the behavior of a *single user that receives one recommendation at the time from a binary set of possible items*. Instead, real recommendation systems propose multiple recommendations from large sets of possible items to a large number of users that can have social ties and shared interests. Therefore, several extensions would be useful to bring our model closer to reality: the recommendation system could return a list of items, which could take values on some discrete set or continuous interval, so to account for nuanced positions.

As per the users, the linear user model that we assume is also a strong simplification: more complex nonlinear models could account for bounded confidence or antagonistic reactions. Moreover, the user opinion could take values on a higher-dimensional space and could be influenced by external information sources. Perhaps, the most important extension would be the inclusion of a *social network of interactions* between multiple users. Indeed, our focus on a single user has allowed us to highlight the feedback loop between the user’s opinion and the recommendations, but has limited the scope of our work in two ways. First, our recommender system was not allowed to exploit either social ties or shared interests to provide its recommendations. Instead, real recommender systems are *collaborative* and effectively take advantage of these features [44]: this fact has been included in some mathematical models [28], [26]. Second, recommendations were the only drive to the opinion dynamics in our model. Instead, opinion dynamics are also driven by social interactions (both directly and through the collaborative elements of the recommender system), creating a complex entanglement of effects. On this matter, we note that some experimental

studies on Facebook have reported that ideological content is primarily filtered by user’s social connections rather than by the curation algorithms, suggesting that user preferences may have stronger impact than algorithmic personalisation [7]. A future model that includes both social and recommendation effects could shed more light on this issue.

APPENDIX

A. Dependence on opinion model’s parameters

Along the paper, we have made a running choice of α , β and γ , but it is clear that different users can be characterised by different parameters. We therefore explore with simulations the dependence of the results on these parameters. In this exploration we keep $\epsilon = 0.05$ and take two non-negative values of the prejudice, i.e., $o_{\text{usr}}^0 = 0.00$ (see Figure 11) and 0.33 (see Figure 12): for each value we repeat the simulations with different α , β and γ . For the simulations we select 116 points on the 2-simplex $\{(\alpha, \beta, \gamma) : \alpha, \beta, \gamma \geq 0; \alpha + \beta + \gamma = 1\}$: 66 points lie on the grid with spacing 0.10 while 50 points are randomly chosen. For each combination of the parameters, we run 1000 simulations till $t_{\text{max}} = 1000$ and compute two variables of interest, namely the empirical probability of obtaining a +1-majority trajectory and the sample average of the click-through rate $\text{ctr}(t_{\text{max}})$. We use triangular colormap plots to represent these two variables of interest with respect to α, β, γ . The triangles represent the 2-simplex and the colours encode the values of the variables of interest.

Our exploration of the parameter space confirms the validity of the simulations shown in the rest of the paper and indicates that the values of the parameters influence the evolution in a rather intuitive way: namely, larger weights on the prejudice make easier for the recommender system to identify whether +1 or -1 is the best recommendation, but reduce the recommender’s effectiveness in reaping clicks.

REFERENCES

- [1] D. Billsus and M. J. Pazzani, “User modeling for adaptive news access,” *User Modeling and User-Adapted Interaction*, vol. 10, no. 2, pp. 147–180, 2000.
- [2] M. Karimi, D. Jannach, and M. Jugovac, “News recommender systems – Survey and roads ahead,” *Information Processing & Management*, vol. 54, no. 6, pp. 1203–1227, 2018.
- [3] J. Liu, P. Dolan, and E. R. Pedersen, “Personalized news recommendation based on click behavior,” in *Proceedings of the 15th International Conference on Intelligent User Interfaces*, ser. IUI ’10. New York, NY, USA: ACM, 2010, pp. 31–40.
- [4] M. A. Beam, “Automating the news: How personalized news recommender system design choices impact news reception,” *Communication Research*, vol. 41, no. 8, pp. 1019–1041, 2014.
- [5] E. Pariser, *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.
- [6] D. Lazer, “The rise of the social algorithm,” *Science*, vol. 348, no. 6239, pp. 1090–1091, 2015.
- [7] E. Bakshy, S. Messing, and L. A. Adamic, “Exposure to ideologically diverse news and opinion on Facebook,” *Science*, vol. 348, no. 6239, pp. 1130–1132, 2015.
- [8] R. S. Nickerson, “Confirmation bias: A ubiquitous phenomenon in many guises,” *Review of General Psychology*, vol. 2, no. 2, pp. 175–220, 1998.
- [9] S. Mullainathan and A. Shleifer, “The market for news,” *American Economic Review*, vol. 95, no. 4, pp. 1031–1053, 2005.
- [10] M. Del Vicario, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi, “Modeling confirmation bias and polarization,” *Scientific reports*, vol. 7, p. 40391, 2017.
- [11] N. E. Friedkin and E. C. Johnsen, *Social influence network theory: A sociological examination of small group dynamics*. Cambridge University Press, 2011.
- [12] C. Castellano, S. Fortunato, and V. Loreto, “Statistical physics of social dynamics,” *Reviews of Modern Physics*, vol. 81, pp. 591–646, 2009.
- [13] A. Flache, M. Mäs, T. Feliciani, E. Chattoe-Brown, G. Deffuant, S. Huet, and J. Lorenz, “Models of social influence: Towards the next frontiers,” *Journal of Artificial Societies and Social Simulation*, vol. 20, no. 4, p. 2, 2017.
- [14] A. V. Proskurnikov and R. Tempo, “A tutorial on modeling and analysis of dynamic social networks. Part I,” *Annual Reviews in Control*, vol. 43, pp. 65–79, 2017.
- [15] —, “A tutorial on modeling and analysis of dynamic social networks. Part II,” *Annual Reviews in Control*, vol. 45, pp. 166–190, 2018.
- [16] M. Moussaïd, J. E. Kämmer, P. P. Analytis, and H. Neth, “Social influence and the collective dynamics of opinion formation,” *PLoS one*, vol. 8, no. 11, p. e78433, 2013.
- [17] R. M. Bond, C. J. Fariss, J. J. Jones, A. D. I. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler, “A 61-million-person experiment in social influence and political mobilization,” *Nature*, vol. 489, no. 7415, pp. 295–298, 09 2012.
- [18] S. Aral, “Social science: Poked to vote,” *Nature*, vol. 489, no. 7415, pp. 212–214, 09 2012.
- [19] A. D. I. Kramer, J. E. Guillory, and J. T. Hancock, “Experimental evidence of massive-scale emotional contagion through social networks,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 24, pp. 8788–8790, 2014.
- [20] M. Maes and L. Bischofberger, “Will the personalization of online social networks foster opinion polarization?” Available at SSRN 2553436, 2015. [Online]. Available: <http://dx.doi.org/10.2139/ssrn.2553436>
- [21] P. Bolzern, P. Colaneri, and G. D. Nicolao, “Opinion influence and evolution in social networks: A Markovian agents model,” *Automatica*, vol. 100, pp. 219–230, 2019.
- [22] D. Geschke, J. Lorenz, and P. Holtz, “The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers,” *British Journal of Social Psychology*, vol. 58, no. 1, pp. 129–149, 2019.
- [23] A. Sirbu, D. Pedreschi, F. Giannotti, and J. Kertész, “Algorithmic bias amplifies opinion fragmentation and polarization: A bounded confidence model,” *PLoS ONE*, vol. 14, no. 3, pp. 1–20, 2019.
- [24] W. S. Rossi and P. Frasca, “Opinion dynamics with topological gossiping: Asynchronous updates under limited attention,” *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 566–571, 2020.
- [25] D. Fleder and K. Hosanagar, “Blockbuster culture’s next rise or fall: The impact of recommender systems on sales diversity,” *Management Science*, vol. 55, no. 5, pp. 697–712, 2009.
- [26] M. Bressan, S. Leucci, A. Panconesi, P. Raghavan, and E. Terolli, “The limits of popularity-based recommendations, and the role of social ties,” in *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’16, New York, NY, USA, 2016, pp. 745–754.
- [27] N. Perra and L. E. Rocha, “Modelling opinion dynamics in the age of algorithmic personalisation,” *Scientific reports*, vol. 9, no. 1, p. 7261, 2019.
- [28] P. Dandekar, A. Goel, and D. Lee, “Biased assimilation, homophily, and the dynamics of polarization,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 15, pp. 5791–5796, 2013.
- [29] L. Spinelli and M. Crovella, “Closed-loop opinion formation,” in *Proceedings of the 2017 ACM on Web Science Conference*, ser. WebSci ’17. New York, NY, USA: ACM, 2017, pp. 73–82.
- [30] R. Jiang, S. Chiappa, T. Lattimore, A. György, and P. Kohli, “Degenerate feedback loops in recommender systems,” in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 2019, pp. 383–390.
- [31] A. J. B. Chaney, B. M. Stewart, and B. E. Engelhardt, “How algorithmic confounding in recommendation systems increases homogeneity and decreases utility,” in *Proceedings of the 12th ACM Conference on Recommender Systems*, ser. RecSys ’18, 2018, pp. 224–232.
- [32] K. El-Arini, G. Veda, D. Shahaf, and C. Guestrin, “Turning down the noise in the blogosphere,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2009, pp. 289–298.
- [33] N. E. Friedkin and E. C. Johnsen, “Social influence and opinions,” *Journal of Mathematical Sociology*, vol. 15, no. 3–4, pp. 193–206, 1990.
- [34] V. Amelkin, F. Bullo, and A. K. Singh, “Polar opinion dynamics in social networks,” *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5650–5665, 2017.




Fig. 11. Simulations with prejudice $o_{\text{usr}}^0 = 0.00$ and parameters α, β, γ covering the 2-simplex. For each choice of parameters, we run 1000 simulations up to $t_{\text{max}} = 1000$, with position $\text{part}(t)$ recommended using $\epsilon = 0.05$. Left: the empirical probability of +1-majority trajectory, that is about 0.50. Right: the sample average of the click-through rate values $\text{ctr}(t_{\text{max}})$.




Fig. 12. Simulations with prejudice $o_{\text{usr}}^0 = 0.33$ and parameters α, β, γ covering the 2-simplex. For each choice of parameters, we run 1000 simulations up to $t_{\text{max}} = 1000$, with position $\text{part}(t)$ recommended using $\epsilon = 0.05$. Left: the empirical probability of +1-majority trajectory. Right: the sample average of the click-through rate values $\text{ctr}(t_{\text{max}})$.

- [35] B. Golub and M. O. Jackson, “Naive learning in social networks and the wisdom of crowds,” *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–49, 2010.
- [36] D. Acemoglu and A. Ozdaglar, “Opinion dynamics and learning in social networks,” *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.
- [37] F. Germano and F. Sobbrino, “Opinion dynamics via search engines (and other algorithmic gatekeepers),” *Journal of Public Economics*, vol. 187, p. 104188, 2020.
- [38] N. E. Friedkin, “The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem,” *IEEE Control Systems*, vol. 35, no. 3, pp. 40–51, 2015.
- [39] B. Clemm von Hohenberg, M. Maes, and B. Pradelski, “Micro influence and macro dynamics of opinion formation,” *Available at SSRN 2974413*, 2017. [Online]. Available: <http://dx.doi.org/10.2139/ssrn.2974413>
- [40] S. Chaiken, “The heuristic model of persuasion,” in *Social influence: the Ontario symposium*. Psychology Press, New York, US, 1987, vol. 5, pp. 3–39.
- [41] S. Bubeck and N. Cesa-Bianchi, “Regret analysis of stochastic and nonstochastic multi-armed bandit problems,” *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [42] W. Li, L. and Chu, J. Langford, and R. E. Schapire, “A contextual-bandit approach to personalized news article recommendation,” in *Proceedings of the 19th International Conference on World Wide Web*, ser. WWW’10, New York, NY, USA, 2010, pp. 661–670.
- [43] D. E. Koulouriotis and A. Xanthopoulos, “Reinforcement learning and evolutionary algorithms for non-stationary multi-armed bandit problems,” *Applied Mathematics and Computation*, vol. 196, no. 2, pp. 913–922, 2008.
- [44] D. Jannach, M. Zanker, A. Felfernig, and G. Friedrich, *Recommender systems: an introduction*. Cambridge University Press, 2010.



Wilbert Samuel Rossi received the PhD degree in Mathematical Engineering from Politecnico di Torino, Turin, Italy, in 2015. In 2013 he was visiting PhD student at the Department of Automatic Control, Lund University, Lund, Sweden. He was a post-doctoral researcher at the Department of Applied Mathematics, University of Twente, Enschede, The Netherlands from 2015 to 2018. Since September 2018, he has been an Assistant Professor at the University College Groningen, University of Groningen, Groningen, The Netherlands. His research interests include dynamics and control in (large-scale) networks, cascading behaviours, (algorithmically-mediated) opinion dynamics.



Jan Willem Polderman was born in Kruiningen, The Netherlands. He received the M.Sc. degree in mathematics from the University of Groningen, Groningen, The Netherlands in 1983, and the Ph.D. degree in mathematics, also from the University of Groningen, in 1987. From 1983 to 1987 he was a Research Assistant at the Center for Mathematics and Computer Science in Amsterdam, where he worked under the supervision of J.H. van Schuppen and J.C. Willems. The title of his dissertation was “Adaptive Control and Identification: Conflict or Conflux?”

Since 1987 he has been with the Department of Applied Mathematics of the University of Twente, Enschede, The Netherlands, first as a Lecturer, and since 2003, as an Associate Professor. He has made contributions to adaptive control, the behavioral approach to systems theory, coding theory, and hybrid systems. With I. Mareels he wrote the book *Adaptive Control: An Introduction* (Basel, Switzerland: Birkhauser 1996). With J.C. Willems he published the book *Mathematical Systems Theory: A Behavioral Approach* (Berlin, Germany: Springer Verlag, 1998). He has been an Associate Editor of *Systems & Control Letters* and *Automatica*. From 2008 until 2019 he was the Programme Director of the bachelor’s and master’s program in applied Mathematics and the master’s program *Systems & Control*. Since 2020 he is Programme Director of the bachelor’s and master’s *Industrial Design Engineering*.



Paolo Frasca (M’13-SM’18) received the Ph.D. degree from Politecnico di Torino, Turin, Italy, in 2009. After Postdoctoral appointments with the CNR-IAC (Rome) and in Turin, he has been an Assistant Professor with the University of Twente (Enschede, The Netherlands) from 2013 to 2016. Since October 2016, he has been a tenured CNRS Researcher with GIPSA-lab, Grenoble, France.

His research interests include the theory of network systems and cyber-physical systems, with applications to infrastructural and social networks. On

these topics, he has (co)authored more than 90 journal and conference papers and the book *Introduction to Averaging Dynamics Over Networks* (Springer). He has been an Associate Editor for several conferences and journals, including *Automatica*, *IEEE Control Systems Letters*, *International Journal of Robust and Nonlinear Control*, *Asian Journal of Control*, and the Conference Editorial Boards of the European Control Association and of the IEEE Control Systems Society.