



Delayed and Indirect Impacts of Link Recommendations

Han Zhang

pariszhang@berkeley.edu
University of California, Berkeley
USA

Yixin Wang

yixinw@umich.edu
University of Michigan
USA

Shangen Lu

shangen@wharton.upenn.edu
Wharton School, University of Pennsylvania
USA

Mihaela Curmei

mcurmei@berkeley.edu
University of California, Berkeley
USA

ABSTRACT

The impacts of link recommendations on social networks are challenging to evaluate, due to feedback loops between algorithmic recommendations and underlying network dynamics. Observational studies have limitations in answering causal questions; naive A/B experiments often result in biased evaluations due to unaccounted network interference and finally, existing simulations primarily employ static network models that do not take into account dynamics. Departing from existing approaches, we employ simulations to study *dynamic* impacts of link recommendations. Specifically, we propose an extension to the Jackson-Rogers network evolution model and investigate how link recommendations affect network evolution *over time*. Our experiments demonstrate that link recommendations can have surprising *delayed* and *indirect* effects on the structural properties of networks. Effects of recommendations vary in the short-term and long-term, such as the immediate reduction in degree inequality but eventual increase in degree inequality through friend-of-friend recommendations. Furthermore, even after recommendations are discontinued, their impacts can persist in the network, in part by altering natural network evolution dynamics. These results provide valuable insights into the interplay between algorithmic interventions and natural network dynamics and highlight the limitations of current evaluation paradigms.

CCS CONCEPTS

• Computing methodologies → Network science; • Theory of computation → Random network models.

KEYWORDS

recommender systems, link recommendation, dynamic graphs

ACM Reference Format:

Han Zhang, Shangen Lu, Yixin Wang, and Mihaela Curmei. 2023. Delayed and Indirect Impacts of Link Recommendations. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FACCT '23)*, June 12–15, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3593013.3594021>



This work is licensed under a Creative Commons Attribution International 4.0 License.

FACCT '23, June 12–15, 2023, Chicago, IL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0192-4/23/06.

<https://doi.org/10.1145/3593013.3594021>

1 INTRODUCTION

Link recommendation algorithms such as Facebook's "People You May Know", Twitter's "Who to Follow" and LinkedIn's "Recommended for You" have an ever-increasing influence on the evolution of social networks, with some accounts crediting over 50% of links to algorithmic recommendations [28]. This can cause downstream effects on information flow, opinion dynamics, and resource allocation. For instance, recommendations can be a polarizing force increasing network segregation, which in turn may re-inforce opinion echo-chambers [7] or restrict access to information and resources for less connected communities [6, 10, 25]. At the same time, recommendation systems can also surface "long-range" connections between nodes that would not otherwise be exposed to each other, and thus, increase network integration by promoting and maintaining diverse links [11, 23]. Given the ubiquity of algorithmic recommendations on social media, studying their impacts is seeing increased academic and regulatory interest. However, such studies are challenging for a variety of reasons, such as lack of normative framing [9] and limited access to large-scale platforms. In this work, we explore a more foundational evaluation challenge stemming from the fact that social networks have underlying dynamics that interfere with algorithmic recommendations.

Existing real-world evaluations of link recommendation algorithms rely on A/B tests (experimental) or longitudinal data (observational). However, both experimental and observational evaluations can yield misleading conclusions. The validity of A/B tests relies on the Stable Unit Treatment Value Assumption (SUTVA) [17, 18], which is violated in case of network interference. Observational studies may fail to assess causal impacts, as longitudinal evaluation lacks counterfactual measurements on what the network evolution would have been without algorithmic recommendations. These challenges have motivated the use of simulation-based evaluations of link recommendations. In simulation studies, one can make explicit network modeling assumptions and then evaluate the impact of recommenders. Despite a growing number of works in this space [7, 13, 15, 16, 25, 26], existing simulation-based evaluation falls short of providing insights into the mechanisms through which recommendations impact social networks. Existing simulation studies primarily consider static networks, and thus do not take into account feedback loops between link recommendation and organic network evolution.

In dynamic settings, the main challenge is to measure impacts relative to a "baseline" network. Given this, existing works often measure the effects relative to the initial networks; they are rarely

measured relative to a plausible counterfactual based on the natural evolution of the graph without link recommendations. Such evaluations can lead to qualitatively wrong conclusions. For instance, Abebe et al. [1] shows that triadic closure – the most common type of friend-of-friend recommendations – can reduce segregation with respect to the initial network before intervention. However, triadic closure can, at the same time, increase network segregation in relative terms with respect to a natural evolution dynamic which assumes the addition of random edges.

In this work, we initiate a study to explore *dynamic* impacts of link recommendations through simulations. We find that link recommendations can have surprising delayed and indirect effects on the structural properties of networks. For instance, we find that the short-term effects of friend-of-friend recommendations can be qualitatively different from long-term effects. In the short term they can alleviate the degree inequality, but in the long term they will increase the degree inequality. Moreover, we demonstrate that the effects of link recommendations can persist in the network even after recommendations have been discontinued. This phenomenon is due to the fact that recommendations impact the network in two ways: directly through the creation of algorithmic edges, and indirectly by influencing the natural growth dynamics. Indirect effects amplify the direct effects, contributing to the persistent impact of link recommendation algorithms. A stylized illustration of the indirect and delayed effects can be found in Figure 1.

Contributions. We enumerate our contribution along modeling, evaluation, and experimental findings.

- (1) *Modeling:* We propose a dynamic network formation model which extends upon the Jackson-Rogers model [19] to incorporate algorithmic recommendations. Unlike the classic model, our dynamic model includes not only the original two phases – dubbed “meeting strangers” and “meeting friends” – but also a third phase of “meeting recommendations.” Additionally, we consider latent node representations, enabling us to model community structure and node activity levels in a flexible manner.
- (2) *Evaluation:* We monitor the progression of network metrics over different intervention windows. We compare the immediate impacts observed during intervention with the delayed impacts observed after the intervention has ended. Further, we measure the indirect impacts that recommendations have on network properties. To tease out the indirect impact, we compare the observed network evolution with a counterfactual baseline network that discounts the influence of recommendations on natural growth.
- (3) *Experimental findings:* Our study reveals diverse qualitative patterns for delayed and indirect effects; we find that different duration of the intervention and/or different times of measurement can lead to drastically different conclusions about “How do recommendations impact networks?” Furthermore, we find that indirect effects can be substantial; they can significantly amplify the impact of recommendations and persist even after the intervention has ended.

2 RELATED WORK

Simulation studies. Simulation studies typically analyze the impact of link recommendations on static networks. Some works [12, 13, 20,

26] focus on a single round of recommendations and examine observed, finding that homophily – the preference for within-group links – leads to exposure bias toward the more homophilous group, even if it’s a minority. Similarly, [15, 16] find that homophily and node degree are strong indicators of which nodes will receive disproportionate visibility. Other works [15, 16] make explicit behavioral assumptions on how nodes accept link recommendations and consider cumulative effects of link recommendations over multiple rounds. They reveal algorithmic amplification of biases via “rich get richer” effects and increases in observed homophily over time. A shortcoming of existing simulation-based evaluations is that they implicitly assume that the addition of algorithmic edges is the only change in the network. Conversely, in our work, we make explicit modeling assumptions about the underlying network dynamics. This evaluation setup allows us to measure the impact of link recommendations with respect to counterfactual natural evolution. Furthermore, by emphasizing underlying temporal dynamics we can pose more subtle evaluation questions such as: “How does the effect of algorithmic intervention fade over time once recommendations are stopped?” or “How do algorithmic recommendations bias the underlying network growth?”

Platform studies. There is limited publicly available research evaluating link recommendation algorithms on real social networking platforms. In one experimental study [9], several recommendation algorithms were compared on IBM’s SocialBlue network and found to reduce group homophily. The study also revealed that friend-of-friend recommendations had the highest rate of acceptance but the lowest level of edge activity. On the other hand, [23] found that recommending more distant connections or “weak ties” through LinkedIn’s “People You May Know” algorithm led to higher transmission of job opportunities. In observational settings, a longitudinal study comparing the Twitter network before and after the introduction of the “Who To Follow” recommender in 2010 [27] found that while recommendations increased the number of connections for all users, the highest gains were achieved by the most popular nodes. Conversely, a study comparing links formed naturally and links formed via recommendations on Flickr and Tumblr [3] found that the recommended links were more diverse and less biased towards popular users. So far, existing observational platform studies provide a limited understanding of the underlying mechanisms and lack access to counterfactual network evolution. A/B tests, on the other hand, may produce wrong estimates of the effects due to network interference. Our counterfactual simulations allow us to articulate in stylized settings the source of bias in both longitudinal and A/B evaluations.

Theoretical investigations. The impacts of friend-of-friend recommendations in homophilous networks were studied theoretically in a number of works. Under choice homophily, which captures the setting when nodes preferentially accept recommendations to within-group nodes, [5, 26] show that recommendations lead to more exposure gains for the homophilous group which further exacerbates homophily. In contrast, [1] shows that, when the closure of triangles via friend-of-friend recommendations is not biased in favor of in-group edges, recommendations can in fact improve network integration. In our work, we investigate friend-of-friend recommendations and their impacts on network segregation and




Figure 1: Delayed and indirect effects. The image displays two counterfactual evolution of the network. The solid blue line represents the trajectory for an intervention interval $[t, \bar{t}]$ in which the full network receives algorithmic recommendations. The dashed black trajectory corresponds to a counterfactual network evolution without recommendations. The total causal effect of recommendation at time t , Effect_t , can be computed as the difference between the counterfactual trajectories at time T (solid blue and dashed black line). The delayed effect at some time $T \geq \bar{t}$ is the difference: $\text{Effect}_T - \text{Effect}_{\bar{t}}$. The solid purple trajectory illustrates the counterfactual evolution of the network in which the indirect effects are removed. The difference between the purple and dashed black curve captures the direct effects and the difference between the blue and purple line captures the indirect effects.

show that their effects further depend on the length of intervention as well as the time of measurement.

Delayed algorithmic impacts. Temporal dynamics associated with algorithmic interventions have been previously studied in the context of fairness in Machine Learning. These works showcase broad settings where algorithmic interventions designed to improve fairness [4, 9, 21], robustness [22] or diversity [8] in the short term, lead to the opposite effect in the long run. Our work uncovers similar surprising temporal dynamics in the case of link recommendations.

3 METHODOLOGY FOR EVALUATING DELAYED AND INDIRECT IMPACTS

To illustrate temporal complexities of evaluating link recommendation algorithms in a dynamic setting, we consider a stylized network evolution model. We propose an extension to the classic Jackson-Rogers (JR) network evolution model [19]. The JR model has been validated empirically and shown to display properties of real network such as decreasing diameter over time, increased edge densification and emergence of community structure [19, 24]. Our extension adds an optional recommendation phase to model the feedback loop between natural network evolution and algorithmic recommendations.

We analyze two simple baseline recommendation algorithms: a neighborhood-based and an affinity-based algorithm. Recommendation algorithms deployed in practice use both neighborhood and affinity information, with potentially additional learning components. Here, however, to understand the underlying mechanisms behind evaluation phenomena in dynamic networks, we opt for simplicity in the choice of recommendation algorithms. Algorithm 3.1 summarizes the simulation and evaluation procedure.

3.1 Dynamic network model

Let $G^t = (V^t, E^t)$ be an undirected network where V^t and E^t are the set of nodes and edges at time t . Nodes are characterized by their group identity g_i and latent representation v_i . We assume group identities are static over time and latent representations are sampled from a group-specific multivariate normal distribution $v_i \sim \mathcal{D}_{g_i}$. At each time step, the network evolves via new nodes which first connect "organically" the existing nodes. Further at each step, the network evolves "algorithmically" via connections mediated by a link recommender.

Natural growth. Similar to the classic JR model, upon the arrival of a new node, the network evolves in two phases: "Meeting Strangers" and "Meeting Friends". In the "Meeting Strangers" phase, the new node makes connections with existing nodes at random. Unlike the classic model where connections are made with a fixed probability, we model connections probabilities based on latent representation of the nodes v_i . This additional modeling assumption allows us to consider various community structures. In the "Meeting Strangers" phase, N_s candidate nodes are sampled uniformly from the existing network, among which the arriving node might make zero, one or more connections. Specifically, the arriving node i connects with each candidate node j with probability proportional to the inner product of their respective latent embeddings: $p_{i,j} = \sigma(\langle v_i, v_j \rangle)$ where $\sigma(\cdot)$ is a scaled and translated sigmoid function¹.

In the subsequent "Meeting Friends" phase, the incoming node i makes additional connections based on neighborhood proximity. Here N_f candidate nodes are sampled from the set of nodes at

¹We consider $\sigma(x) = \frac{1}{1+e^{-ax+b}}$ where we set a and b to match desired average linkage probabilities.

distance 2 (neighbors of neighbors) from node i . Node i connects with each candidate node j with constant probability. Optionally, we model attrition effects by considering node departures from the network with a hazard function² that increases with the age of the node.

Algorithmic intervention. We introduce algorithmic intervention as a third "meeting recommendations" phase, whereby nodes in the network receive link recommendation. This phase applies not only to incoming nodes but also to existing nodes in the network. Upon receiving a recommendation, nodes accept it according to a behavioral model. We consider neighborhood and affinity-based recommendations. The prototypical neighborhood recommendation is the friend-of-friend (FoF) recommendation where candidate nodes are selected uniformly from the set of nodes at distance 2: $\mathbb{P}(\text{rec}_i = j) = \frac{\mathbf{1}(\text{dist}(i,j)=2)}{\sum_{j'} \mathbf{1}(\text{dist}(i,j')=2)}$. Conversely, affinity-based (Latent) recommendations utilize the latent structure rather than local neighborhood structure to make recommendations. Specifically, the Latent recommender computes affinity scores for all the nodes in the network in terms of inner products and recommends candidates with probability proportional to the scores: $\mathbb{P}(\text{rec}_i = j) \propto e^{\beta s_j}$, where $s_j = \langle v_i, v_j \rangle$ is the similarity between node i and node j 's embeddings. High values of softmax temperature parameter β correspond to more deterministic recommenders. Lower values of β lead to more random recommendations and qualitatively capture the effects of estimation noise for learning-based recommenders.

Behavioral models. We consider two behavioral models for accepting link recommendations: the constant probability baseline and the embedding-based probability where nodes accept new links based on proximity in latent space. Embedding-based probabilities model more granular notions of choice homophily [5]. We consider an additional behavioral option in which upon acceptance of a new recommended link, an edge is removed at random from the set of existing edges. This option is in line with several works which model recommendations as a *rewiring* [5, 14, 25] process that does not impact the average degree.

3.2 Structural metrics

Clustering coefficient: Clustering coefficient of a node is the ratio of triangles it forms Δ_i and the maximum number of triangles it could have formed given its current degree d_i : $c_i = \frac{2\Delta_i}{d_i(d_i-1)}$. The clustering coefficient captures the micro-level cohesion in the network [16]. The average clustering coefficient is the metric averaged over all the nodes. Averaging the metric over communities measure clustering at the community level.

Gini coefficient: We measure the inequality in the degree distribution via Gini coefficient. An increase of the Gini coefficient resulting from the use of recommendations is often used to demonstrate biases in link recommendation [14–16]. The Gini coefficient can be computed for an ordered list of node degrees as: $G = \frac{2 \sum_{i=1}^n i d_i}{n \sum_{i=1}^n d_i} - \frac{n+1}{n}$. Similarly, this metric can be computed for the entire graph or restricted to communities.

² $h(a) = cd^a + k$ where a is the age of the node and c, d, k are tuned to model mean and variance of lifespans in the network.

Algorithm 1: Simulating network evolution

```

input : initial  $G_0$ ; time-steps  $T$ ; hazard function;
        communities: prevalence  $c_g$ , latent distribution  $\mathcal{D}_g$ ;
        natural growth:  $N_s, N_f$ , dist-2 connect prob  $p_2$ ;
        intervention: window  $[t, \bar{t}]$ , recommender, behavior;

for  $t$  in  $1 \dots T$  do
    Natural growth;
    sample incoming node  $i$  according to group prevalence
       $c_g$  and latent distribution  $\mathcal{D}_g$ ;
    strangers  $\leftarrow$  samples  $N_s$  nodes;
    for  $s$  in strangers do
        | add edge  $i - s$  with probability  $\propto \langle v_i, v_s \rangle$ 
    end
    friends  $\leftarrow$  samples  $N_f$  neighbors of neighbor nodes;
    for  $f$  in friends do
        | add edge  $i - f$  with probability  $p_2$ 
    end
    if  $t \in [t, \bar{t}]$  then
        Algorithmic intervention;
        for node  $j$  in treatment group  $G_t^{\text{treatment}}$  do
            candidate = recommender( $j, G_t$ );
            if behavior( $j$ , candidate) = accept then
                | add edge  $j$ -candidate
            end
        end
    end
    nodes to remove  $\leftarrow$  hazard function
end

```

Homophily: We define monochromatic and bichromatic edges to be edges that link two nodes from the same community and nodes from different communities, respectively. The homophily of a community measures the propensity of nodes to favor within-group connections compared to a non-preferential baseline. We compute the homophily of a community g as follows: $H_g = \frac{|E_{gg}|}{|E_g|} - \frac{n_g}{n}$, where $|E_{gg}|$ denotes the number of monochromatic edges within g , $|E_g|$ denotes the number of total edges that have at least one node in community g . Finally n_g denotes the size of g , and n denotes the total size of the network, the ratio $\frac{n_g}{n}$ is the baseline fraction of within-group links when nodes have no group-based preferences. The vast majority of existing simulation-based evaluations study recommendation-induced changes in homophily [1, 9, 13].

3.3 Evaluation

We denote by $\mathcal{G}$ the evolution of the network under natural dynamics and by $\mathcal{G}(\text{rec}, [t, \bar{t}])$ the evolution when the network receives link recommendation according to rec algorithm over the $[t, \bar{t}]$ intervention interval. The subscript T in $\mathcal{G}_T(\cdot)$ is used to denote the snapshot of the network at time T . For a structural metric m , we denote by $m(\mathcal{G}_t(\text{rec}, [t, \bar{t}]))$ the value of this metric evaluated at time t . Given a network evolution model, we simulate the counterfactual trajectories of the metrics for intervention intervals $[t, \bar{t}]$ of different lengths.

Total effects. The total effect of intervening in a network evolution $\mathcal{G}$ with recommender rec for intervention interval $[t, \bar{t}]$ on the metric m at time T is defined as the difference between two counterfactual measurements:

$$\text{Effect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) := m(\mathcal{G}_T(\text{rec}, [t, \bar{t}])) - m(\mathcal{G}_T).$$

This definition compares between the “treatment” universe where the whole network received an algorithmic intervention and a counterfactual “control” universe where the network evolved without algorithmic recommendations.

Delayed effects. When the measurement occurs during the intervention interval, i.e. $T \in [t, \bar{t}]$ we call the effect *immediate*. Conversely, when the measurement occurs after the intervention stops, i.e. $T > \bar{t}$, we can define the notion of *delayed* impact as:

$$\begin{aligned} \text{DelayedEffect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) = \\ \text{Effect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) - \text{Effect}_{\bar{t}}(\mathcal{G}, \text{rec}, [t, \bar{t}]) \end{aligned}$$

The notion of delayed impacts allows us to characterize the impact of link recommendations into three broad categories: diminishing, amplifying, and persistent based on the sign of the delayed effect ($\text{DelayedEffect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}])$). Note that the delayed impacts measure the difference in effect sizes between time $\bar{t}$ and time T , rather than the difference in the metric $m(\mathcal{G}_T(\text{rec}, [t, \bar{t}])) - m(\mathcal{G}_{\bar{t}}(\text{rec}, [t, \bar{t}]))$. These two notions are equivalent only when the metric remains constant from time $\bar{t}$ to time T under natural network evolution dynamics.

Indirect effects. The temporal evolution of the network in the presence of link recommendations is affected *directly* by the addition of algorithmic edges; but also *indirectly*, as the addition of algorithmic edges biases the natural growth dynamics. In our model, recommendations have an indirect impact on natural dynamics in the “Meeting Friends” phase of the network evolution. Upon arrival, a new node forms the initial neighborhood by “Meeting Strangers”. In the next phase, the arriving node connects with nodes at distance 2 from itself (neighbors of neighbors). In the presence of algorithmic edges, there can be nodes at distance 2 that require algorithmic edges in order to be reachable. For instance, in Figure 2, at time $t+2$ node F arrives and connects with B in the first phase of natural growth. In the second phase, nodes $\{A, C, D, E\}$ are reachable for the purposes of “Meeting Friends”, however node A is only reachable because of algorithmic edge $B - A$. If the incoming node F connects to a node such as A , the resulting edge $F - A$ is said to be *mediated* by recommendation.

To measure direct impacts, we design a counterfactual experimental procedure to remove the indirect influence of algorithmic edges on natural growth by discount mediated edges. Constructing a post-hoc counterfactual by simply removing mediated edges from the network would inadvertently decrease the edge density in the network and thus bias the analysis. To address this, the counterfactual procedure is defined as follows: upon arrival of a new node i , we modify the “Meeting Friends” phase to only consider node candidates that are at distance 2 from i via non-algorithmic edges. For this counterfactual, the network contains no mediated edges. Since organic link are created only when a new node arrives, this counterfactual removes all *indirect* effects of recommenders

on natural growth. Formally, we refer to this unmediated evolution as $\hat{\mathcal{G}}(\text{rec}, [t, \bar{t}])$. Finally we can define direct effects as the difference between the metric evaluated for the unmediated counterfactual and the natural growth trajectory. The indirect effects are the difference between total effect and direct effect:

$$\begin{aligned} \text{DirectEffect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) &= m(\hat{\mathcal{G}}_T(\text{rec}, [t, \bar{t}])) - m(\mathcal{G}_T), \\ \text{IndirectEffect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) &= \text{Effect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]) - \\ &\quad \text{DirectEffect}_T(\mathcal{G}, \text{rec}, [t, \bar{t}]). \end{aligned}$$

Longitudinal evaluation. In a simulated setting, one has access to both counterfactuals: the evolution of the network in the presence of recommendations and in their absence. In longitudinal evaluations (i.e. observational studies), effects are measured as the difference in a metric before and after the intervention:

$$\widehat{\text{Effect}}_T^{\text{obs}}(\mathcal{G}, \text{rec}, [t, \bar{t}]) := m(\mathcal{G}_T, \text{rec}, [t, \bar{t}]) - m(\mathcal{G}_t)$$

This measure is biased whenever $m(\mathcal{G}_T) \neq m(\mathcal{G}_t)$.

A/B evaluation. The lack of valid counterfactuals motivates the use of A/B tests in settings with interventional access. In an A/B test, nodes are divided into two groups: treatment and control; the treatment group receives recommendations, while the control group does not. However, in networks, the Stable Unit Treatment Value Assumption (SUTVA) [18] does not hold due to network interference. There are various methods for choosing treatment nodes, as well as methods to correct for network interference in estimation procedures.

Let a scheme for choosing treatment nodes exist and let $\mathcal{G}^{AB}(\text{rec}, [t, \bar{t}])$ be the network evolution where a group of nodes are assigned to the treatment group and receive recommendations. To estimate the impact of recommendations on a metric of interest, we compare the values of the metric on the treatment and control groups:

$$\begin{aligned} \widehat{\text{Effect}}_T^{AB}(\mathcal{G}, \text{rec}, [t, \bar{t}]) := \\ m^{\text{treatment}}(\mathcal{G}_T^{AB}, \text{rec}, [t, \bar{t}]) - m^{\text{control}}(\mathcal{G}_T^{AB}, \text{rec}, [t, \bar{t}]). \end{aligned}$$

In evaluating the impact of recommendations, we consider both naive estimators that do not correct for network interference, and more sophisticated ones that take network externalities into account.

4 COUNTERFACTUAL EVALUATION

We simulate counterfactual evolutions of the network by applying different recommendation interventions for various time periods. Although accessing counterfactuals in reality is infeasible, these simulations offer insights into dynamic and temporal phenomena.

4.1 Setup

The results in sections 4.2 and 4.3 consider a simple experimental setup to illustrate delayed and indirect effects of FoF and Latent³ recommenders. In the baseline setup, we consider two equally sized communities. We sample latent embeddings for the nodes independently from $\mathcal{N}(\mu_c, \sigma I)$, with $\mu_1 = [0, 1]$ and $\mu_2 = [1, 0]$; for both

³For Latent recommendations we use softmax temperature $\beta = 10$, which in our simulations corresponds to $\approx 80\%$ chance of recommending ‘closest’ node.




Figure 2: Mediated edges. Edge $A - B$ is a recommended edge and a direct effect of the recommendation. Edge $F - A$ is a mediated edge as node A would not be reachable from F in the absence of $A - B$ and thus an indirect effect of recommendation.

communities, the variance of the embeddings is set to $\sigma = 0.05$. We sample 50 nodes in each group and initialize both within-group and across-group edges by connecting pairs of nodes $i - j$ with probability proportional to the inner product of their embeddings. Then, for each node in the network, we consider their neighbors at distance 2 and connect to them with probability $p_1 = 0.05$. This results in a slightly homophilic initial network with homophily $h_1 = h_2 \approx 0.1$. Upon initialization, at each time step, 5 new nodes arrive. For natural growth, we consider $N_s = N_f = 100$ and the connection probability of connecting to a candidate node in the "Meeting Friends" phase: $p_2 = 0.05$ (which adds an average of 5 edges per node). Further, we assume a behavioral model where nodes accept recommended edges with a constant probability, 0.5. At each time step we measure structural metrics of the network such as clustering coefficient, Gini coefficient and homophily. We repeat each trajectory for 5 random seeds and report average results along with confidence bands.⁴

In section 4.4, we vary these assumptions by considering majority-minority structure, differentiated homophily, and within-group heterogeneity. Finally in Appendix A.1 and A.2, we consider additional behavioral assumptions, different variants of the underlying dynamics and modifications to the recommendation algorithms.

4.2 Delayed effects

We find that affinity and neighborhood-based recommenders have opposite long-term effects on homophily and clustering coefficients. Latent leads to increases in homophily as well as global clustering whereas FoF recommendations reduces clustering. Both recommenders have diminishing delayed impacts with respect to homophily and clustering, as upon the end of the intervention, the trajectory of the metric regresses to the counterfactual natural growth trajectory. Meanwhile, effects with respect to the Gini coefficient are qualitatively different. Latent recommendations lead to increases in degree inequality both in the short and in the long term. Conversely, FoF recommendations decrease the Gini coefficient in the short term but increase it in the long term. Figure 3 illustrates these findings.

Clustering and homophily. In the case of Latent recommendations, if nodes i , j , and k are similar in embedding space, it is likely they are from the same community and all three pairs of edges have been proposed as recommendations, leading to a strong bias for




Figure 3: Trajectories of homophily, global clustering and Gini coefficient for Latent and FoF recommendations applied over varying intervention intervals. The solid black lines represent the evolution of the metrics under natural growth dynamics. Left column corresponds to a behavioral model where nodes accept recommendations with constant probability and right column corresponds to acceptance probability proportional to embedding similarity. The blue lines represent trajectories for the Latent recommender whereas the orange lines correspond to FoF. Dashed and dotted lines correspond to various intervention intervals. The shaded area corresponds to the 95% confidence interval.

closing triangles within the community. This results in increased homophily and clustering, especially with the behavioral assumption that nodes accept links based on embedding similarity (see Fig. A1c and A1d). The bias is further exacerbated for high-temperature β . Conversely, with a low β value, which corresponds to a more stochastic recommender, it is less likely that all three edges of a

⁴Code for reproducing experimental results can be found here.

given triplet will close, thus slightly lowering the homophily value compared to the high- β case.

The FoF recommender intervention mimics the "Meeting Friends" phase of natural growth, which has a bias towards forming cross-community links, resulting in decreased homophily. Decreased clustering occurs because random FoF recommenders lack the bias of connecting nodes with a large common neighborhood, unlike neighborhood-based models such as Adamic-Adar [2], which favor links with nodes with a large number of common neighborhoods and thus increase clustering (see Fig. A2).

Gini coefficient. The Gini coefficient for recommendations shows amplifying delayed effects. The natural growth trend exhibits a slight increase in inequality over time. Latent recommendations exacerbate the wealth gap between popular and unpopular nodes through a bias towards nodes with high embedding norms, creating the "rich-get-richer" effect. FoF recommendations initially reduce degree inequality as they are less biased towards popular nodes. Surprisingly, once recommendations stop, the Gini coefficient increases dramatically for both recommenders, particularly for FoF. This is due to a "relative-rich-minority" effect caused by differences in edge density between "older" and "newer" nodes. Through recommendations, the existing nodes have accumulated a large number of connections. After recommendations stop, as new nodes are arriving and natural growth is continuing to take place, the set of "rich" nodes becomes relatively smaller, while the majority of nodes, who come after recommendations stop, have much fewer connections. This results in higher inequality as measured by the Gini coefficient. If a rewiring behavioral model is assumed (see Fig. A1b and A1d), where nodes receiving recommendations do not change their degree, the delayed impacts diminish, similar to the case of clustering and homophily.

4.3 Indirect effects

The "Meeting Friends" phase occurs in feedback loop with algorithmic recommendations, leading to the mediation phenomenon illustrated in Figure 2. We find that mediated edges are common, long-lasting and biased. For instance, in the case of Latent recommendations, the mediated edges are more likely to connect nodes from the same community, resulting in significant effects on homophily.

Prevalence, persistence and bias of mediated edges. Mediated edges play a substantial role in the "Meeting Friends" stage of natural growth dynamics. Figure 4 shows that the proportion of mediated edges increases with the duration of the intervention period and remains constant even after recommendations have stopped for both Latent and FoF recommenders. The bias of mediated edges in terms of bichromaticism is distinct for each recommender, relative to the proportion of bichromatic edges under natural growth. For instance, about a third of unmediated "Meeting Friends" edges are bichromatic for both recommenders. However, the proportion of bichromatic edges among mediated edges is lower for Latent and higher for FoF. Our findings about homophily in Section 4.2 suggest that algorithmic edges are biased towards monochromatic for Latent and bichromatic for FoF. This experiment supports the

hypothesis that recommendations have a compounding effect by inducing biases in the natural network evolution.

Counterfactual measurements of indirect effects. To measure the indirect effects on structural metrics, we apply the counterfactual procedure from Section 3.3 and isolate the direct effects. We find that without mediation, structural metrics trend faster to their natural evolution. Figure 4 shows the comparison of homophily under natural growth, algorithmic growth, and unmediated algorithmic growth. The difference between the latter two quantifies the magnitude of indirect effects. Additionally, indirect effects grow relatively stronger over time, even after recommendations end, highlighting the persistence of indirect impacts.

4.4 Impacts of group structure

Previous studies have highlighted the unequal impact of recommendations on different communities, especially when they are divided into majority and minority groups and display varying levels of homophily [13, 16, 26]. Here, we examine the effects of differential homophily between majority and minority groups as well as within-community heterogeneity.

Community heterogeneity. We examine the impacts of homogeneity and heterogeneity in the latent representation of nodes on the results of link recommendation. We model within-group heterogeneity by varying the variance of the latent embeddings. In the heterogeneous setting, we set the variance of the embedding distribution to $\sigma^2 = 0.1$, and to $\sigma^2 = 0.01$ in the homogeneous setting; in the extreme case when the variance is 0 this recovers the JR variant of [1]. Figure 5 shows that when there is high within-group heterogeneity, Latent recommendations greatly increase the global clustering coefficient, while for high within-group homogeneity, Latent recommenders have the opposite effect, reducing global clustering. This unexpected phenomenon is due to Latent recommendations favoring nodes with high embedding norms. For any existing edge $i - j$, there is a disproportionately high chance that both i and j will receive recommendations concentrated on a small subset of large-normed nodes, creating closed triangles. In the homogeneous setting, such "collisions" are less likely to occur, as the probability of two nodes independently being recommended with the same node is lower. These findings highlight the importance of investigating not only between-group differences but also within-group differences.

Homophilic minority and heterophilic majority. In previous studies, it has been observed that homophilic minority groups receive an excessive amount of exposure from recommendations, leading to increased disparities in homophily between groups. We simulate the network evolution for a majority fraction of 60%, $\mu_1 = [0, 1]$ and $\mu_2 = [1.1, 1.1]$ ⁵ and find that our results support this conclusion for Latent recommendations. In contrast, Figure 6 shows that while FoF increases homophily for the majority, it decreases it significantly for the minority. The FoF recommender amplifies the visibility of the majority more than that of the minority. This phenomenon can be explained by the heterophilic nature of the

⁵by increasing the norm of latent representations in group 2 we increase the within-group connection probabilities.




Figure 4: Indirect effects. Solid black lines correspond to natural growth dynamics. The blue and orange lines represent trajectories for Latent and FoF recommenders, respectively. Left plot illustrates the prevalence of algorithmically mediated edges in the "Meeting Friends" phase of the natural growth dynamics. Lighter and darker lines correspond to various intervention intervals. Center plot illustrates the bias in the mediated edges relative to unmediated ones. Solid lines correspond to fraction of bichromatic edges among unmediated edges created in the "Meeting Friends" phase; the dashed lines record the fraction of bichromatic mediated edges. Right plot shows the trajectories of homophily metrics. Dotted lines correspond to the trajectories of the metrics for the unmediated counterfactual evolution of the network.




Figure 5: Effects of within-group heterogeneity. Black trajectories correspond to the natural evolution of the global clustering metric. Blue and orange lines indicate the trajectories for Latent and FoF recommenders, respectively.

majority group, where most of its direct connections are with minority nodes. For a majority node i , its neighbors' neighbors that are in the minority group are likely to also be direct neighbors of i . Therefore, there is a larger probability that i is recommended with




Figure 6: Homophily impacts. Trajectories of average degree and community homophily for Latent and FoF for the heterophilic majority and homophilic minority.

other majority nodes at distance 2, increasing the homophily for the majority group.

4.5 Evaluation biases

As we do not have access to the full range of counterfactual measurements, we often resort to either longitudinal or A/B evaluations. However, these methods have their limitations and may fail to accurately capture the impacts of recommendations. In this section, we explore the potential biases and limitations of these evaluation procedures.

Longitudinal evaluation. Simulating a longitudinal evaluation mimics an observational study. In this setting, a single trajectory is

observed. The estimated total effect of link recommendations is the temporal difference between the value of the metric before and after the intervention. If the underlying dynamics for a metric are stationary, as is the case for homophily in Figure 3, then a longitudinal evaluation is unbiased. Conversely, when under natural dynamics a metric is non-stationary, a naive comparison between the initial network at time t and the network at time T could yield qualitatively misleading estimates. For instance, in Figure 3, when measuring the impact of FoF on clustering coefficient, a longitudinal measurement would compare the solid orange trajectory between time $t = 50$ and $T = 400$, thus over-estimating about double the size of the true effects. Furthermore, the naive observational measurement for Latent would incorrectly suggest a negative effect on network clustering, when in reality it has a positive impact compared to the natural evolution.

A/B evaluation. A/B evaluations in dynamic networks become complex due to changes of network structure over time caused by natural dynamics and algorithmic interventions. We simulate A/B tests aimed at estimating causal effects of Latent and FoF link recommendations. We perform random treatment ($p = 0.5$) on node assignment for clustering and Gini coefficient and community-based treatment assignment for measuring homophily. The estimates in Figure 7 are adjusted for network interference. Detailed computation of the adjusted estimates as well as the corresponding plots for the naive cases can be found in Appendix A.3. The quality of A/B estimates varies based on the metric and intervention. For example, A/B tests overestimate the impact on homophily for Latent but underestimate it for FoF. The clustering coefficient is accurately estimated for Latent but greatly underestimated for FoF. For the Gini coefficient, both Latent and FoF interventions underestimate the metric for the treatment group while accurately estimating the metric's evolution under natural growth for the control group. Across most settings, the quality of the metric deteriorates over time, further supporting the existence of dynamic effects that compound network interference effects.

5 DISCUSSION AND FUTURE WORK

In this study, we explored the dynamic effects of link recommendations on network evolution through simulations. Our proposed extension of the Jackson-Rogers model provides insight into the impact of link recommendations on network structure. Emphasizing the importance of temporal dynamics and measurement timing, our simulations revealed surprising and persistent effects of link recommendations on network structure.

Using synthetic data and simple network evolution models and recommendation algorithms, we answered "what-if" scenarios in a controlled setting, providing a valuable first step in understanding these effects in real networks. Our results showed that link recommendations can have delayed and indirect impacts on network structure, with long-lasting effects even after recommendations have ceased, which result in significant cascading indirect effects over time. This highlights the need for further research in evaluating link recommendation algorithms in dynamic networks.

Finally, our study sheds light on the potential biases that can arise in evaluating link recommendation algorithms in dynamic networks. We find that evaluating metrics longitudinally or using A/B




Figure 7: A/B evaluations. Solid lines correspond to ground truth counterfactual evaluations of homophily, clustering and Gini coefficient. Dashed lines correspond to trajectories estimated from running an A/B test.

tests can result in biased estimates when the underlying network dynamics are not stationary. This highlights the need for advanced estimation procedures that consider both network interference and underlying dynamic effects to accurately assess the impacts of link recommendation algorithms.

We identify several opportunities for future research. To improve the validity of our conclusions, one avenue is to validate modeling assumptions against real-world networks. Refining the modeling assumptions, such as allowing for non-recommender-driven edge creation between existing nodes, could better reflect reality and potentially result in greater indirect effects. Another important area of interest is to further develop methods for measuring direct and indirect effects in different network formation models.

A second direction of future work is to study downstream impacts of link recommendation. Modeling node embeddings as dynamic properties, influenced by the local neighborhood through biased assimilation [10] or mere exposure effects [8], could provide insight into how opinion formation on networks is influenced by recommendations. Finally, examining recommendation scenarios where network edges in the social network are formed by content recommendations, rather than user recommendations, which is typical of social media sites like Instagram, is a promising area for investigation.

ACKNOWLEDGMENTS

This work is partially supported by the National Science Foundation and the Office of Naval Research.

REFERENCES

- [1] Rediet Abebe, Nicole Immorlica, Jon Kleinberg, Brendan Lucier, and Ali Shirali. 2022. On the Effect of Triadic Closure on Network Segregation. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 249–284. <https://doi.org/10.1145/3490486.3538322> arXiv:2205.13658 [cs, econ].
- [2] Lada A Adamic and Eytan Adar. 2003. Friends and neighbors on the web. *Social networks* 25, 3 (2003), 211–230.
- [3] Luca Maria Aiello and Nicola Barbieri. 2017. Evolution of Ego-networks in Social Media with Link Recommendations. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 111–120. <https://doi.org/10.1145/3018661.3018733> arXiv:1702.01398 [physics].
- [4] Nil-Jana Akpınar, Cyrus DiCiccio, Preetam Nandy, and Kinjal Basu. 2022. Long-term Dynamics of Fairness Intervention in Connection Recommender Systems. *arXiv preprint arXiv:2203.16432* (2022).
- [5] Aili Asikainen, Gerardo Iñiguez, Javier Ureña-Carrión, Kimmo Kaski, and Mikko Kivelä. 2020. Cumulative effects of triadic closure and homophily in social networks. *Science Advances* 6, 19 (2020), eaax7310.
- [6] Vincenzo Auletta, Antonio Coppola, and Diodato Ferraioli. 2022. On the Impact of Social Media Recommendations on Opinion Consensus. <http://arxiv.org/abs/2204.03299> arXiv:2204.03299 [cs].
- [7] Federico Cinus, Marco Minici, Corrado Monti, and Francesco Bonchi. 2022. The effect of people recommenders on echo chambers and polarization. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 90–101.
- [8] Mihaela Curmei, Andreas A Haupt, Benjamin Recht, and Dylan Hadfield-Menell. 2022. Towards Psychologically-Grounded Dynamic Preference Models. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 35–48.
- [9] Elizabeth M Daly, Werner Geyer, and David R Millen. 2010. The network effects of recommending social connections. In *Proceedings of the fourth ACM conference on Recommender systems*. 301–304.
- [10] Pranav Dandekar, Ashish Goel, and David T. Lee. 2013. Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences* 110, 15 (April 2013), 5791–5796. <https://doi.org/10.1073/pnas.1217220110> Publisher: National Academy of Sciences Section: Social Sciences.
- [11] Gianlorenzo D'Angelo, Lorenzo Severini, and Yllka Velaj. 2019. Recommending links through influence maximization. *Theoretical Computer Science* 764 (April 2019), 30–41. <https://doi.org/10.1016/j.tcs.2018.01.017>
- [12] Lisette Espin-Noboa, Claudia Wagner, Markus Strohmaier, and Fariba Karimi. 2022. Inequality and inequity in network-based ranking and recommendation algorithms. *Scientific reports* 12, 1 (2022), 1–14.
- [13] Francesco Fabbri, Francesco Bonchi, Ludovico Boratto, and Carlos Castillo. 2020. The Effect of Homophily on Disparate Visibility of Minorities in People Recommender Systems. *Proceedings of the International AAAI Conference on Web and Social Media* 14 (May 2020), 165–175. <https://ojs.aaai.org/index.php/ICWSM/article/view/7288>
- [14] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2021. Exposure Inequality in People Recommender Systems: The Long-Term Effects. *arXiv:2112.08237 [cs]* (Dec. 2021). <http://arxiv.org/abs/2112.08237> arXiv: 2112.08237.
- [15] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2022. Exposure Inequality in People Recommender Systems: The Long-Term Effects. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 194–204.
- [16] Antonio Ferrara, Lisette Espin-Noboa, Fariba Karimi, and Claudia Wagner. 2022. Link recommendations: Their impact on network structure and minorities. *arXiv preprint arXiv:2205.06048* (2022).
- [17] Huan Gui, Ya Xu, Anmol Bhasin, and Jiawei Han. 2015. Network a/b testing: From sampling to estimation. In *Proceedings of the 24th International Conference on World Wide Web*. 399–409.
- [18] Guido W. Imbens and Donald B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press. Google-Books-ID: Bf1tBwAAQBAJ.
- [19] Matthew O Jackson and Brian W Rogers. 2007. Meeting strangers and friends of friends: How random are social networks? *American Economic Review* 97, 3 (2007), 890–915.
- [20] Fariba Karimi, Mathieu Génois, Claudia Wagner, Philipp Singer, and Markus Strohmaier. 2018. Homophily influences ranking of minorities in social networks. *Scientific reports* 8, 1 (2018), 1–12.
- [21] Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2018. Delayed impact of fair machine learning. In *International Conference on Machine Learning*. PMLR, 3150–3158.
- [22] Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. 2019. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 230–239.
- [23] Karthik Rajkumar, Guillaume Saint-Jacques, Iavor Bojinov, Erik Brynjolfsson, and Sinan Aral. 2022. A causal test of the strength of weak ties. *Science* 377, 6612 (Sept. 2022), 1304–1310. <https://doi.org/10.1126/science.abl4476> Publisher: American Association for the Advancement of Science.
- [24] Diego Ruiz and Jorge Finke. 2017. Stability of the Jackson-Rogers model. 1803–1808. <https://doi.org/10.1109/CDC.2017.8263909>
- [25] Fernando P. Santos, Yphtach Lelkes, and Simon A. Levin. 2021. Link recommendation algorithms and dynamics of polarization in online social networks. *Proceedings of the National Academy of Sciences* 118, 50 (Dec. 2021), e2102141118. <https://doi.org/10.1073/pnas.2102141118> Publisher: Proceedings of the National Academy of Sciences.
- [26] Ana-Andreea Stoica, Christopher Riederer, and Augustin Chaintreau. 2018. Algorithmic Glass Ceiling in Social Networks: The effects of social recommendations on network diversity. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*. ACM Press, Lyon, France, 923–932. <https://doi.org/10.1145/3178876.3186140>
- [27] Jessica Su, Aneesh Sharma, and Sharad Goel. 2016. The effect of recommendations on network structure. In *Proceedings of the 25th international conference on World Wide Web*. 1157–1167.
- [28] LinkedIn Engineering Team. 2022. People You May Know. <https://engineering.linkedin.com/teams/data/artificial-intelligence/people-you-may-know> [Online; accessed 31-Jan-2023].

A ADDITIONAL EXPERIMENTS

A.1 Behavioral assumptions

In this section, we investigate two behavioral assumptions: edge rewiring and choice homophily. Under edge rewiring assumption, upon accepting a recommended edge the node will remove at random one of its existing edges. Under choice homophily, nodes accept recommendations according to embedding similarity rather than with a constant probability, assumed in the main body of the paper. Figure A1 shows the trajectories of homophily, clustering and Gini coefficient under different behavioral assumptions. In the case of edge-rewiring (Fig. A1b, A1d), the delayed impacts on Gini coefficient diminish. The behavioral assumption that nodes accept preferentially nodes with high embedding similarity unsurprisingly leads to further increases in observed homophily for both Latent and FoF recommenders (Fig A1c, A1d).

A.2 Intervention Variants

In Fig. A2a, we consider different choices of temperature parameter β for the Latent recommender. All recorded metrics increase as the recommendations become more deterministic (increasing β). Note that $\beta = 4$ and $\beta = 10$ are nearly indistinguishable in the homophily graph whereas $\beta = 4$ and $\beta = 2$ are close in the graph of the clustering coefficient. This suggests that different metrics have different sensitivities to changes in recommendation intervention.

In Fig. A2b we compare FoF with a neighborhood-based recommendation algorithm that recommends node pairs with higher

Adamic-Adar index. The Adamic-Adar index assigns a higher similarity score to pairs of low-degree nodes that share many neighbors in common. This change leads to an extreme increase in the clustering coefficient.

A.3 Naive and Interference Adjusted estimates for A/B evaluation

First we comment on the naive and adjusted estimation procedure for each metric.

Homophily. Naively, we estimate homophily for an A/B evaluation by assigning one community as treatment and the other as control and measuring the homophily for each community. To adjust the estimate for the control group, we discount the algorithmic edges received by the control group, whereas to adjust for the treatment, we double the count of the algorithmic edges that the treatment group forms with the control group.

Clustering. Naively, we estimate the group-level clustering coefficient for treatment and control by averaging node-level clustering for treatment and control groups respectively. We adjust the estimate for the treatment/control group by calculating its clustering coefficient only among the treatment/control neighbors, i.e., in a sub-network with only the treatment/control nodes.

Gini coefficient. We naively estimate Gini coefficient by measuring the value within the treatment and control groups, respectively. We adjust Gini coefficient by discounting the algorithmic edges for the control group.

Fig. A3 shows the naive estimates derived from the A/B evaluation.




Figure A1: Counterfactual trajectories of homophily, global clustering and Gini coefficient for Latent and FoF recommendations applied over varying intervention intervals. The solid black lines represent the evolution of the metric under natural growth dynamics. The blue lines represent trajectories for the Latent recommender whereas the orange lines correspond to FoF. The shaded area corresponds to the 95% confidence intervals calculated over 5 independent trajectories.





Figure A2: Counterfactual trajectories for variants of Latent and FoF recommenders

Figure A3: A/B evaluations (naively): Solid lines correspond to ground truth counterfactual evaluations of homophily, clustering and Gini coefficient. Dashed lines correspond to trajectories estimated *naively* from running an A/B test.