Inducing Efficient and Equitable Professional Networks Through Link Recommendations

Cynthia Dwork
Department of Computer Science, Harvard University, Cambridge, MA, USA

Chris Hays
Institute for Data, Systems and Society, Massachusetts Institute for Technology, Cambridge, MA, USA

Lunjia Hu
Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA

Nicole Immorlica
Microsoft Research, Cambridge, MA, USA

Juan Perdomo
Massachusetts Institute of Technology, Cambridge, MA, USA
New York University, NY, USA

Abstract

Professional networks are a key determinant of individuals' labor market outcomes. They may also play a role in either exacerbating or ameliorating inequality of opportunity across social groups. We initiate an investigation into the positive role that a professional networking platform can play when network members have different degrees of off-platform privilege. In a theoretical model, we show that the set of link recommendation policies that reduce costs between privileged and unprivileged individuals yield equilibria that are *welfare-improving over all possible equilibria*, compared to those obtained when not recommending links or recommending some smaller fraction of cross-group links.

We next investigate the implications of platforms that do not intervene on the network formation process. We show that, absent intervention, inequality can increase relative to starting privilege levels *even without exogenous in-group preferences*, confirming and complementing existing theoretical literature. Increased inequality emerges from the differential leverage privileged and unprivileged individuals have in forming connections due to their asymmetric *ex ante* prospects. This is a formalization of a source of inequality in the labor market which has not been previously explored.

These two findings reveal a stark reality: professional networking platforms that fail to foster integration in the link formation process risk reducing the platform's utility to its users and exacerbating existing labor market inequality.

2012 ACM Subject Classification Theory of computation → Network formation; Applied computing → Economics

Keywords and phrases Professional networks, Inequality, Link Recommendations

Digital Object Identifier 10.4230/LIPIcs.FORC.2026.13

Related Version *Full Version*: https://arxiv.org/pdf/2503.04542

Funding This work was facilitated by the Hire Aspirations Institute at Harvard, supported in part by Alfred P. Sloan Foundation grant G-2017-9890.
Lunjia Hu: Work done while supported by the Simons Foundation Collaboration on the Theory of Algorithmic Fairness and the Harvard Center for Research on Computation and Society (CRCS).
Juan Perdomo: Work done while supported by the Harvard Center for Research on Computation and Society (CRCS) and Alfred P. Sloan Foundation grant G-2020-13941.

1 Introduction

Individuals' professional networks play a crucial role in their labor market outcomes. They facilitate job placement and lead to higher wages, increased productivity and greater performance [24, 8, 28, 22, 1, 27, 4, 13, 2, 3, 7, 9]. However, they may also exacerbate employment or wage inequality, lead to differential job search costs, yield uneven gains to new technology or otherwise entrench or increase economic or social stratification [15, 21, 30, 5]. Indeed, recent empirical work has demonstrated that minority individuals' connection requests on platforms are accepted at lower rates [18] and that widespread disconnection across social groups is partially attributable to lack of exposure [12].

Digital professional networking platforms crucially influence network formation by providing link recommendations between individuals. Indeed, a large fraction of all connections are formed through recommendations [31]. Thus, the choices that platforms make with respect to their recommendation systems have far-reaching implications for their millions of users.

In our work, we explore the role a platform may play in increasing social welfare and ameliorating inequality through link recommendations. We show link recommendation policies that favor cross-group connections decrease inequality and improve efficiency relative to exogenous levels. The additional welfare comes from a better use of opportunities on the platform, which is arguably good for the platform itself as well, since a primary consideration for professional networking platforms is that users find opportunities through the platform (and thus have a reason to maintain their accounts). Thus, an increase in welfare may be aligned with the incentives of the platform acting in its long term interests.

We also characterize how, if the platform does not intervene on the network formation process, inequality can increase relative to starting levels. Importantly, these results hold even if individuals do not have exogenous preferences for connecting with members of their own group. In other words, inequality can be *self-propagating*. This is because inequality can increase as a result of the differential leverage between individuals who come to the network with a starting level of privilege – and therefore can offer greater benefits to potential connections – and those that do not. This confirms and complements prior work, where in-group preferences are assumed to be exogenous.

At a high level, our model consists of a population of individuals, each belonging to one of two groups. Each individual in a group receives a random number of opportunities. Opportunities may be thought of as jobs, referrals, contracts, projects, or other desirable professional outcomes.[1] They may come from one of two sources: first, exogenously or, second, after exogenous opportunities are drawn, from a contact in a professional network. Exogenous opportunities might come from an individual's educational or socio-economic background, their family, or other pre-determined sources.[2] The arrival of exogenous opportunities is assumed to be out of the individuals' control and unrelated to the individual's ability to make use of an opportunity. If an individual receives more exogenous opportunities than

[1] This model may also be suitable to describing the (costly-per-person) transmission of information within a network. This is not to be confused with "fixed-cost broadcasts", like social media posts that simultaneously go out to all social contacts at once. Our intuition is that costly-per-person transmission is the more relevant kind for professional networking. This is perhaps justified by the fact that fixed-cost broadcasts look similar to network-wide communication, and is therefore less dependent on network structure.

[2] We may even extend the model to consider "off-platform" links as having formed exogenously through pre-determined factors like an indivual's personal history. These off-platform links may be encapsulated into a person's exogenous opportunity distribution and we may assume that all links considered in this work are "on-platform" links. Our results can easily accommodate this framing of exogenous versus endogenous connections. See Appendix A of the full manuscript for further discussion.

they can use, they pass remaining opportunities to their contacts in a professional social network. Sometimes, two individuals may unwittingly pass more opportunities to a single person than that person can use, in which case the extra opportunities are wasted.

Our focus in this work is on *network formation*. Prior to the point at which exogenous opportunities are realized, individuals in the model will form connections with others. Forming each connection bears a small cost for the individuals involved but allows for the possibility that one may pass an opportunity to the other. The cost may be imagined to be a search cost (*e.g.*, the difficulty of finding each other) or an attentional burden (*e.g.*, the time necessary to navigate a user interface or communicate). Each individual will choose how many connections to create, and with whom, according to their own interests: balancing the cost of the connection against the probability that they receive an opportunity as a result of the connection and would not have otherwise received opportunities.

Individuals within each group have the same opportunity distributions, and there will be one group with a higher number of expected exogenous opportunities (the privileged group) and one with a lower number (the unprivileged group). Our focus will be on quantifying *efficiency*, or the number of opportunities that are used, and *inequality*, the ratio of the number of opportunities used within each group.

We analyze the policies of a platform that has the limited capacity to influence outcomes by making link recommendations. Recommendations have the effect of reducing connection costs by reducing search or attentional effort, and we analyze the effect of different policies on efficiency and inequality. We introduce the model in more detail and define notation in Section 2.

1.1 Our contributions

We propose a model of strategic network formation with exogenous social inequality and where a platform can make link recommendations. Our model builds on prior work studying a theoretical model among a homogeneous population and without link recommendations [17].

In our main results, we show that the natural way of alleviating inequality, *i.e.*, encouraging cross-group connections, is also the most effective way of improving total utility, over all possible equilibria, as long as the exogenous level of inequality is sufficiently large. The reason for this is that fewer opportunities are wasted when the network is more integrated: With more cross-group links, it is more likely that, when an opportunity is passed, it is passed to a lower-privileged person who would not otherwise receive an opportunity.

Thus, it can be in the best interest of a purely utilitarian platform to recommend cross-group links and consequently reduce inequality. Somewhat surprisingly, we demonstrate this fairness-utility alignment even without exogenous reasons to favor diversity: all the opportunities in our model are equivalent, so an opportunity obtained from a cross-group link does not bring any extra benefit compared to an opportunity from an in-group link. In reality, opportunities may come in a variety of types and it is often beneficial to have a wide coverage of different types of opportunities from diverse sources [19]. In those cases, the incentive of a utilitarian platform to recommend cross-group links can be even stronger. Our result provides an explanation for the advantages of integration even without *a priori* reasons to favor diversity.

We show how platforms, if they do not intervene on the network formation process, can increase labor market inequality. Specifically, we show that inequality, defined as the ratio of average utilities between the two groups, increases relative to exogenous levels at equilibrium in the model where all connections are organic (*i.e.*, the platform does not make any recommendations). This is because they enable cheap connections of which only the already privileged can take advantage. Thus, a *laissez-faire* platform that does not try to intervene on the link recommendation process increases inequality.

Interestingly, despite the fact that our model does not bake in any *a priori* in-group preferences, homophily may emerge endogenously. This is because, especially when exogenous inequality between groups is a large, individuals who are privileged may only be willing to connect to other privileged individuals; connecting with a less privileged person is not worth the (communication or search) costs associated with connecting, and so *laissez-faire* platforms may be segregated. Thus, our work can serve as a microfoundation for homophily in professional networking, complementing the existing literature where homophily is assumed to be pre-existing.

1.2 Related work

Several previous theory models have explored quantification of or remedies for inequality in networks for hiring and opportunity.

We adapt our model from that of Dwork et al. (2024) [17], which, among other things, quantifies the price of anarchy resulting from strategic network formation. Our model extends theirs by allowing for social groups with distinct levels of exogenous privilege and by allowing the network to influence incentives by making recommendations.

Several prior works have studied theoretical models of inequality in labor markets [6, 29, 10]. In each of these works, the source of inequality is *ex ante* homophily, where social groups have exogenous preferences for members of their own group. In our model, there is no *ex ante* homophily – instead, “neutral” networks increase inequality because privileged individuals have more leverage to form higher value connections. Thus, our analysis describes a different pathway through which inequality can emerge in the labor market. Similar to our work, Bolte et al. (2024) [6] also explore how inequality can be reduced using affirmative action policies: their setting concerns lowering a hiring threshold for the minority group, while ours is about changing the structure of the network.

More broadly, our work fits in a large body of research providing microfoundations for addressing inequality in society. Among them, Heidari and Kleinberg (2021) [20] considers a model of intergenerational opportunities showing that it can be economically efficient to allocate (*e.g.*, educational) opportunities to individuals of lower socioeconomic status rather than higher-performing individuals of higher socioeconomic status. Other work has considered how affirmative action policies may help inform firms about the capabilities of or counteract the biases against minority workers [14, 11, 25]. Our paper complements this literature by providing a microfoundation for cross-group link recommendations in professional networking.

A final related line of related work considers the learning theoretic aspects of determining who to connect in social networks. Formalization of the problem of determining which links to form in a social network dates back to at least Liben-Nowell and Kleinberg (2003) [26]. Dwork et al. (2024) [16] shows how to make fair link predictions based on potentially complex and evolving characteristics. Our work does not consider heterogeneous link formation probabilities between individuals in a population, and extensions to address this aspect of network formation would be a valuable direction for future work.

1.3 Organization of the paper

In Section 2 we formally define the model and notation. In Section 3, we establish several general properties of equilibria in our model. In Section 4, we explore the welfare implications of cross-group link recommendation policies and show that recommending cross-group links is utility improving relative to not making any recommendations or making a smaller fraction of cross-group recommendations, for all equilibria. In Section 5 we explore how *laissez-faire* platforms can increase inequality in equilibrium, relative to exogenous levels. We also analyze how cross-group link recommendations can be used to counteract these effects and reduce inequality relative to exogenous levels. In Section 6, we conclude and discuss avenues for future work. We defer proofs to the full version of the paper.

2 Model

We analyze a population of individuals $i \in [n]$. Individuals belong to one of two groups: a *privileged* green group, denoted G, and an *unprivileged* blue group, denoted B. We assume that $|G|, |B| = \Omega(n)$, but the groups need not be the same size.

Individuals each seek a single opportunity. We assume all opportunities are of equal value to all individuals, normalized to 1. Opportunities arrive either exogenously from nature or endogenously through the job network. Exogenous opportunities are drawn from group-dependent probability distributions. For each $i \in B$ (resp. $i \in G$), the probability of i receiving k exogenous opportunities is denoted b_k (resp. g_k). Exogenous opportunities are drawn IID conditional on group. Privilege relates to the exogenous distribution of opportunities. Namely, we assume that g stochastically dominates b (i.e., for all $m \in \mathbb{Z}_{\geq 0}$ it holds $\sum_{\ell \geq m} g_\ell \geq \sum_{\ell \geq m} b_\ell$ with the inequality strict for some m). We will also impose assumptions, to be formalized later, ensuring that the b_0 is sufficiently large and that g_0 is not too large. These assumptions allow us to ensure that *inequality* is sufficiently large so that our results hold *over all possible equilibria* and thus for any equilibrium selection process. If we do not wish to specify a group membership for an individual i, we will refer to their probability of receiving ℓ exogenous opportunities as $p_{i\ell}$. We also assume that the support of the exogenous opportunity distribution is bounded: there is a constant C so that for all $\ell \geq C$, it holds $p_{i\ell} = 0$ for all $i \in [n]$.

Opportunities also arrive through an endogenously-formed job network, creating a game whose action space for each individual is a selection of potential neighbors. The timing of the game is as follows. A platform first recommends a (potentially empty) subset of links $Q_i \subseteq \{(i, j) : j \in [n]\}$ of size $|Q_i| = k$ for each $i \in [n]$, which have the effect of eliminating the cost of the connection. We assume if $(i, j) \in Q_i$, then $(j, i) \in Q_j$. Define $Q = \cup_{i \in [N]} Q_i$. Individuals – knowing Q, $[n]$, and the exogeneous distribution of opportunities for each individual but *not* their specific realization – then choose neighbors, forming a network $E \subseteq \{(i, j) : i, j \in [n]\}$. Let $N_i(E) = \{j : (i, j) \in E\}$ and $d_i(E) = |N_i(E)|$ (when the edge set E is clear from the context; we drop it from the notation). We will call edges in the network that were not recommended (*i.e.*, $E \setminus Q$) *organic* or *non-recommended* connections. Finally, exogenous opportunities are realized and extra ones are distributed to neighbors in the network uniformly at random. Specifically, if i receives $\ell > 1$ exogenously opportunities, then they select $(\ell - 1)$ of their d_i neighbors uniformly at random and pass each of them a single opportunity (discarding leftover opportunities if $\ell - 1 > d_i$). Let

$$\mu_i(d) = \sum_{\ell \geq 1} \min\{\ell - 1, d\}\, p_{i\ell}$$

be the expected number of opportunities an individual of degree d passes to their neighbors. Let $\mu_B(d)$ (resp. $\mu_G(d)$) be this expectation for a generic (*i.e.*, any) blue (resp. green) group member. We will overload notation and write $\mu_i = \mu_i(\infty)$ when we wish to refer to the expected number of extra opportunities a member receives (sim. μ_G and μ_B). I.e., $\mu_i(\infty)$ is the expected number of opportunities a person with arbitrarily large (hence, ∞) degree will pass to their neighbors. To avoid trivialities where there is no chance a member of a group receives any extra opportunities, we will assume throughout that $\sum_{\ell \geq 1}(\ell - 1) g_\ell$ and $\sum_{\ell \geq 1}(\ell - 1) b_\ell$ are positive.

Given a network and realization of exogenous and endogenous opportunities, an individual i receives a utility of 1 if they receive an (exogenous or endogenous) opportunity and pays a cost of γ for each of their unrecommended neighbors in the network.[3] Thus their *ex ante* expected utility from a network E is:

$$u_i(E) = \left(1 - p_{i0} \prod_{j \in N_i(E)} \left(1 - \frac{\mu_j(d_j)}{d_j}\right)\right) - \gamma \cdot |N_i \setminus Q_i| \tag{1}$$

where the first term is the probability of receiving an opportunity and the second is the cost of their organic connections. Social welfare will be defined, depending on the context, as the sum of individuals' expected utilities or the minimum utility over the population.

We study the equilibria of the induced game. We will use defection-free pairwise Nash (DFPN) as our equilibrium concept, as in Dwork et al. (2024) [17]. We will characterize pure-strategy equilibria in this work (*i.e.*, we will not consider players who select edges non-deterministically). This yields the property that equilibria are uniquely specified by a set of formed edges. DFPN considers defections in which individuals may unilaterally sever links and/or pairs of individuals may add a link if it is mutually beneficial. Notably in contrast to the notion of pairwise stability studied in Jackson and Wolinsky (1996) [23], a single defection may involve both severing and forming links simultaneously. In this sense, DFPN is a refinement of pairwise stability that allows individuals to consider dropping current links when contemplating forming a new one, analogous to stability in stable marriage problems.

Formally, an edge set E is DFPN if

(a) for all i, j such that $(i, j) \notin E$ and for all $S_i \subseteq N_i(E)$, $S_j \subseteq N_j(E)$

$$\begin{aligned} &0 \geq u_i(E \cup \{(i,j)\} \setminus \{(i,\ell) \;:\; \ell \in S_i\}) - u_i(E), \text{ or} \\ &0 \geq u_j(E \cup \{(i,j)\} \setminus \{(j,\ell) \;:\; \ell \in S_j\}) - u_j(E), \end{aligned} \tag{2}$$

(b) and for each i, j such that $(i, j) \in E$, it holds

$$\begin{aligned} &0 \geq u_i(E \setminus \{(i,j)\}) - u_i(E), \text{ and} \\ &0 \geq u_j(E \setminus \{(i,j)\}) - u_j(E). \end{aligned} \tag{3}$$

Intuitively, Equation (2) specifies that no pair of individuals would prefer to form a connection amongst themselves, dropping a (possibly empty) set of their existing connections. Equation (3) specifies that no individual in a connection would rather sever the connection. For fixed parameters, we will write the set of DFPN equilibrium edge sets as $\mathcal{E}$. If we wish to contrast different equilibrium edge sets for different parameters, we will specify a particular choice of parameter (*e.g.*, number of recommendations k) using subscripts (*e.g.*, $\mathcal{E}_k$).

Our main results characterize the inequality and welfare properties of different link recommendation policies on equilibrium outcomes. In particular, we will study policies that are *anonymous* in the sense that they are the same for people who are indistinguishable to the platform. Since in our simple setting, individuals are only distinguished by their group membership, this means that we consider policies where the platform recommends a fixed fraction of cross-group connections to each individual within a group. Since we are not assuming $|B| = |G|$, it may not be possible for these fractions to be the same for members of B

[3] It would be an interesting direction for future work to explore differential connection costs across different pairs (either where one member of a pair must exert more effort to connect or where different pairs pay different costs based on uncertainty about their value to each other). We imagine that most useful connections in the real world impose some costs on their participants, since relationships in which one exerts no maintenance effort do not typically function as conduits of information about job opportunities.

and G. We will denote the fraction of cross-group connections recommended to B as ρ where $\rho \in \{0, 1/k, \ldots, 1\}$ for the number of recommendations k. This implies that members of G receive a $\rho\,|B|\,/\,|G|$ fraction of cross-group recommendations. Throughout, we will assume that ρ is set so as to ensure $\rho\,|B|\,/\,|G| \leq 1$. This inequality is trivially satisfied if $|B| \leq |G|$. Note that Hall's marriage theorem implies that such a set of cross-group recommendations always exists. The remaining same-group recommendations can be constructed if the relevant populations have even size, which we assume for convenience. See Appendix A of the full manuscript for further explanation.

Our model extends the base model in Dwork et al. (2024) [17]. The key differences are that (1) we model multiple social groups, (2) we consider a platform making link recommendations and (3) we allow for the possibility for each individual to receive more than one extra opportunity. We add (3) only for the sake of generality: all of our results hold in the special case that each individual may receive no more than one extra opportunity.

3 Equilibrium properties

The equilibria of our game satisfy several useful properties, also potentially of independent interest. The first is a type of balance condition. It states that almost all individuals who are connected have similar probabilities of passing each other opportunities. In other words, individuals more-or-less get what they give in their relationships.

The formal lemma, stated below and proved in the appendix, extends the results of Dwork et al. (2024) [17] to our more complex setting with recommendations and two groups. Recall that $\mu_i(d)/d$ is the probability that an individual i of degree d passes one of their connections an opportunity. Define constant $C(\gamma, k) = 2(\gamma^{-1} + k)^2(\gamma^{-1} + k + 1)^2$.

▶ **Lemma 1.** *For all equilibrium $E \in \mathcal{E}$, all $\varepsilon > 0$, and the constant $C(\gamma, k)$, there exists a set of individuals S with $|S| \geq n - C(\gamma, k)/\varepsilon$ such that:*

(a) *For all individuals $i \in S$ and $j \in N_i(E)$,*

$$\frac{\mu_i(d_i)}{d_i} \geq \frac{\mu_j(d_j + 1)}{d_j + 1} - \varepsilon, \tag{4}$$

and $j \in S$. Also, if i and j are members of the same group and ε is sufficiently small, then it also holds that $\mu_i(d_i)/d_i \leq \mu_j(d_j - 1)/(d_j - 1) + \varepsilon$.

(b) *For all $i \in S$, there exists some $j \in S$ such that $(i, j) \notin E$, $d_i = d_j$ and i and j are members of the same group (implying that Equation (4) holds).*

Moreover, if ε is small enough (i.e., upper bounded by a constant not depending on n), the inequalities hold exactly without the additive ε.

Intuitively, part (a) of Lemma 1 says that each individual i participating in a connection (i, j) in equilibrium satisfies an inequality where their connection j has probability of passing them an opportunity no less than the probability i would pass j an opportunity if they had degree one larger. It additionally says that, for connected members of the same group, each individual exactly satisfies the condition that they have probability of passing their connection an opportunity no more than if their connection had one fewer connection, respectively.[4] Part (b) further states that each individual has an outside option: they *could* be (but are not) connected to someone else of their same degree with an approximately equal expected number of extra opportunities.

[4] This condition is not trivially satisfied by the fact that $\mu_j(d_j)/d_j \geq \mu_i(d_i + 1)/(d_i + 1) - \varepsilon$ since the d_i, d_j are determined by equilibrium conditions, and may not continue to hold if individuals' degrees change.

The result above is quantified by a slack term ε. It holds for all ε, but there is a trade-off between the tightness of Equation (4) and the size of S. Making ε smaller makes the inequality tighter but decreases the size of the set S. Making ε larger makes the inequality looser but increases the size of the set S. Depending on the situation, it may be desirable to choose different values of ε, but the choice of ε is *not* a parameter influencing the set of possible equilibria. Instead, it is a analyst-defined choice for finding a $n - O(1)$ (as long as ε is not dependent on n) set of individuals in any equilibrium who satisfy the balance conditions in the lemma.

We also observe that by the fact that $j \in S$, Equation (4) holds when we switch the roles of i and j. That is, $\mu_j(d_j)/d_j \geq \mu_i(d_i + 1)/(d_i + 1) - \varepsilon$. Together, Equation (4) and the inequality above ensure that all neighbors of i satisfy a balance condition ensuring that neither individual in a connection will have substantially lower probability of passing the other an opportunity.

It is also not hard to see that the form of the utility function as well as the equilibrium conditions imply the following two facts:

▶ **Fact 2.** *In any equilibrium $E \in \mathcal{E}$, all recommended edges form,* i.e.*, $Q \subseteq E$.*

▶ **Fact 3.** *In any equilibrium $E \in \mathcal{E}$, all individuals have bounded degree. In particular, $d_i \leq 1/\gamma + k$ for all $i \in [n]$.*

We omit the proofs of these elementary facts and remark that, although we assume that recommendations eliminate connection costs in our model, we could relax this modeling choice and assume instead that recommendations only *reduce* (but do not eliminate) the cost of connections. Our results in this paper continue to hold as long as recommended connection costs are sufficiently small (*i.e.*, small enough that members of the privileged group can be incentivized to form connections with members of the unprivileged group in equilibrium). We also note that central planners might extract even greater social benefits in cases where they can subsidize links by more than the connection cost, yielding payments (negative costs) for links formed. We leave exploration of these topics for future work.

Finally, we establish that there always exists an equilibrium, as long as exogenous inequality is large enough.

▶ **Fact 4.** *For all g, b, k, ρ and large enough n such that $\gamma > b_0(1 - b_0 - b_1)$, it holds $\mathcal{E} \neq \varnothing$.*

We can construct equilibria in our model by letting each individual in a group have the same degree as every other member of their group. Individuals in different groups will have different degrees (according to their privilege) distributions. Setting $\gamma > b_0(1 - b_0 - b_1)$ ensures that members of B will only form connections through recommendations, which considerably simplifies the analysis and conforms with the parameter regime considered in this work (where members of B are sufficiently unprivileged).[5] We leave for future work the question of whether there exist equilibria for any possible setting of the parameters.

[5] We remark that our model in this paper concerns *economically meaningful* connections, *i.e.*, connections that carry with them some possibility that the individuals incident to the connection benefit. Our intuition is that the number of economically meaningful connections an individual has may be much lower than their number of "connections" on professional networking platforms might indicate. Thus, depending on the situation, the number of connections in our model may or may not exactly correspond to connections as defined by professional networking platforms. In general, the assumption that individuals must be sufficiently unprivileged so as to form few organic connections is an analytical convenience that is not necessary for the results to hold.

4 Social welfare analysis

We analyze social welfare of the network at equilibrium under different link recommendation policies. We quantify social welfare as either the average utility of all individuals in the population (a utilitarian notion) or the minimum utility of any individual in the population (a Rawlsian notion), and our results hold for both definitions of social welfare. In Proposition 5, we show that, for sufficiently differing privileges between groups, recommending k cross-group link recommendations generates higher – both utilitarian and Rawlsian – social welfare, for all equilibria, than providing no recommendations. In Proposition 6, we show that, keeping k constant, it is always better to recommend a larger fraction of cross-group link recommendations than a smaller fraction for all equilibria, provided there is a sufficient gap between the two fractions.[6]

We first formally state Proposition 5. Intuitively, this proposition says that any equilibrium $E_k \in \mathcal{E}_k$ induced by all cross-group edge recommendations achieves greater social welfare than any equilibrium $E_0 \in \mathcal{E}_0$ induced by a platform that makes no recommendations (where $\mathcal{E}_k$ is the set of equilibria of the game with k cross-group recommendations per person).

▶ **Proposition 5.** *For any $k > 0$, there exist constants $\bar{g}_0, \underline{b}_0 \in (0,1), \underline{n} > 0$ such that for all $n \geq \underline{n}$ and opportunity distributions g, b with $g_0 \leq \bar{g}_0$ and $b_0 \geq \underline{b}_0$, the following holds. There exist constants $\underline{\gamma}, \bar{\gamma}$ with $\underline{\gamma} < \bar{\gamma}$ such that for all edge costs $\gamma \in (\underline{\gamma}, \bar{\gamma})$ and all equilibria $E_k \in \mathcal{E}_k, E_0 \in \mathcal{E}_0$,*

$$\sum_{i \in [n]} u_i(E_k) > \sum_{i \in [n]} u_i(E_0), \text{ and} \tag{5}$$

$$\min_{i \in [n]} u_i(E_k) > \min_{i \in [n]} u_i(E_0). \tag{6}$$

Intuitively, the result holds because an opportunity has a lower chance of being wasted if there are many edges from high privilege to low privilege individuals. The first inequality concerns utilitarian social welfare, and the second concerns Rawlsian social welfare. We will sometimes refer to the utilitarian social welfare as the *average endogenous utility* to contrast it with the *average exogenous utility*, which is simply the sum of utilities of individuals in the absence of the platform: $\sum_{i=1}^{n}(1 - p_{i0})$. Note that because we assume $\rho\,|B|\,/\,|G| \leq 1$, the result requires $|B| \leq |G|$.

We prove the result in the full version of the paper and sketch the intuition here. To prove the statement for utilitarian social welfare, we appeal to Lemma 1, to show that almost all individuals form connections with others who have a similar probability of passing each other opportunities. This allows us to derive upper and lower bounds on the degrees of all but a constant number of individuals for a fixed set of parameters. This, in turn, allows us to derive upper and lower bounds on the *utilities* of all but a constant number of individuals (since their utility is determined by their neighbors' opportunity distributions and degrees which Lemma 1 greatly simplifies). We can then compare the upper bound for all $E_0 \in \mathcal{E}_0$ with a lower bound for all $E_k \in \mathcal{E}_k$ and show that the latter is greater than the former. For the Rawlsian social welfare, Lemma 1 is not useful, since the properties implied by the lemma may not hold for some number of individuals, and one of these individuals could achieve the minimum. Thus, we argue that *all* members of the unprivileged group form connections,

[6] A gap is needed between fractions of cross-group recommendations because of the multiplicity of equilibria: the best equilibrium for a particular setting of ρ may be better than the worst equilibrium for $\rho + \delta$ where δ is a small constant.

that the lowest-utility member of the unprivileged group has lower utility than that of the privileged group, and finally that the lowest-utility member of the unprivileged group under k cross-group recommendations has utility greater than the corresponding individual under no recommendations.

We validate Proposition 5 in Figure 1. In each of the plots in this paper, bounds on utilities are asymptotic (as $n \to \infty$) and set $|G| = |B|$. The non-asymptotic bounds are within additive $O(n^{-1})$ of the asymptotic bounds. For simplicity, we also set $p_{i\ell} = 0$ for all $\ell \notin \{0, 2\}$ and $i \in [n]$ so that exogenous opportunity distributions collapse to a one-dimensional parameter. In Figure 1, we plot upper and lower bounds on average utilities for different numbers of recommendations. Upper bounds are in red and lower bounds are in blue, and different shades indicate different values of g_0. On the vertical axis of Figure 1, we plot the difference between the average endogenous ($\sum_{i=1}^{n} u_i(E)$) and exogenous ($\sum_{i=1}^{n}(1 - p_{i0})$ utilities. We see that as the number of recommendations increases, the lower bound on average utility increases so that it is above the upper bound at $k = 0$. (I.e., for a given g_0, the red line at $k = 0$ is below the blue line at $k > 0$.) However, in general, increasing the number of recommendations – of any fraction cross-group – is not always social welfare increasing, since increased connectivity can lead to more inefficiency. We can see this by the fact that for values of g_0 where the upper and lower bounds are tight towards the right of the plot, the average endogenous minus exogenous utility is decreasing in k past a certain point.

We next state our second result of this section, which concerns the comparison between different fractions of cross-group recommendations for a fixed number of recommendations. (This is in contrast with the result above, which is about comparing k cross-group edges with *no recommendations at all.*) Our result states that a higher fraction of cross-group recommendations always achieves greater utility, for all equilibria, than a lower fraction, as long as the gap between the lower and higher fractions is sufficiently large. Given a positive integer k and $\rho \in [0, 1]$, we use $\mathcal{E}_\rho$ to denote the set of all equilibria induced by any k edge recommendations per person among which ρk recommendations per person are cross-group.

▶ **Proposition 6.** *For any $k > 0$, there exist constants $\bar{g}_0, \underline{b}_0 \in (0, 1), \underline{n} > 0$ such that for all $n \geq \underline{n}$ and opportunity distributions g, b with $g_0 \leq \bar{g}_0$ and $b_0 \geq \underline{b}_0$, the following holds. There exist $\underline{\gamma}, \bar{\gamma}$ where $\underline{\gamma} < \bar{\gamma}$ such that for all edge cost $\gamma \in (\underline{\gamma}, \bar{\gamma})$ and all $\rho \in (0, 1]$, there exists $\delta \in (0, \rho]$ such that for all $\rho' \in \{0, 1/k, \ldots, \lfloor(\rho - \delta)k\rfloor/k\}$ and all equilibria $E_\rho \in \mathcal{E}_\rho, E_{\rho'} \in \mathcal{E}_{\rho'}$,*

$$\sum_{i \in [n]} u_i(E_\rho) > \sum_{i \in [n]} u_i(E_{\rho'}), \text{ and}$$
$$\min_{i \in [n]} u_i(E_\rho) > \min_{i \in [n]} u_i(E_{\rho'}).$$

The proof of Proposition 6 is deferred to the full version of the paper and appeals to the same upper and lower bounds on utilities used in the proof of the previous result. We validate our result in Figure 2: for particular values of k and ρ, we determine the largest ρ' for which the inequalities in the result hold. We also demonstrate how, if inequality is *not* sufficiently large, the inequalities in Proposition 6 may *not* hold. That is, there are some parameter regimes where there exists ρ such that the worst-case social welfare under a ρ-fraction of cross-group recommendations is not greater than the best-case social welfare for any policy with a $\rho' < \rho$ fraction of cross-group recommendations.

In Figure 2, we show upper and lower bounds on the difference between average endogenous utility and average exogenous utility as ρ varies. The setting is the same as in Figure 1, except that we fix $k = 5$ and vary the proportion of cross-group link recommendations ρ.

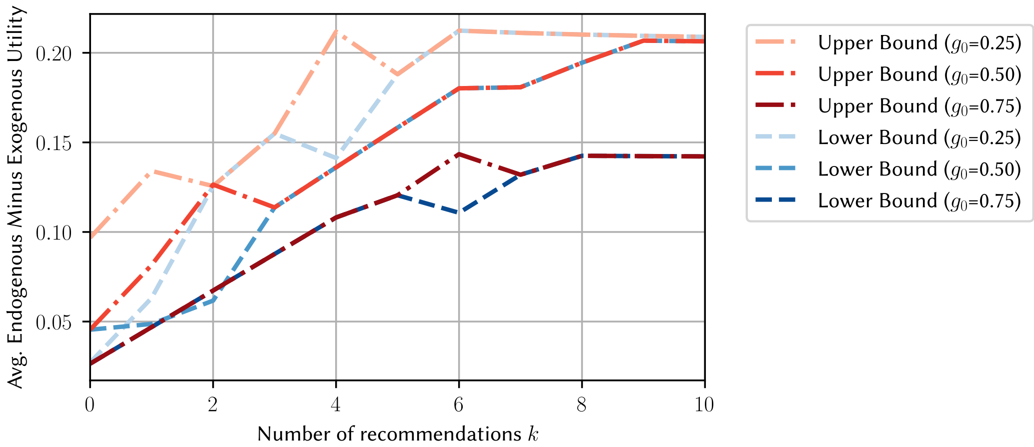

Figure 1 We plot average utility against the number of cross-group recommendations made. It is always better to make $k > 0$ cross-group recommendations than to not make any recommendations. We set b_0 so that the exogenous utility ratio is always 2. We set $\gamma = 0.02$ and $\rho = 1$.

Note that since $k = 5$, the only valid values of ρ are multiples of $1/5$, the ticks on the horizontal axis. The average endogenous minus exogenous utility is increasing, in general, as the fraction of cross-group link recommendations increases, although there are some values of ρ where the change is indeterminate (where the upper and lower bounds are too far apart). These parts of the parameter space, e.g., the bounds when $\rho \leq 0.4$ and $g_0 = 0.5$ demonstrate that inequality was *not* sufficiently large so that for every tick on the horizontal axis greater than 0, the blue line at the tick is above the red line for some tick to its left. (By contrast, if we set the exogenous utility ratio to be sufficiently large, then Proposition 6 tells us that this would occur.) Also, note that when g_0 is higher (at darker shades), the benefits of increasing ρ are more modest. This is because there is a larger chance that members of G receive no opportunities, so integrating the network can only go so far in increasing the prospects of members of B.

5 Inequality analysis

We next explore inequality at equilibrium in our model. In Section 5.1, we show how the existence of a platform can result in increased inequality, even without pre-existing homophily in the network formation process. Then, in Section 5.2, we analyze the ability for platforms to ameliorate the extra inequality they created by selectively subsidizing cross-group links.

5.1 Laissez-faire platforms increase inequality

Here, we analyze the network formation process under a platform does that not subsidize any connections via recommendations, *i.e.*, $k = 0$. We call these "*laissez-faire*" because they do not intervene on behalf of any particular connections. To quantify inequality, we will use the ratio of simple within group averages. Formally, for an edge set E, define the *utility ratio* between greens and blues as

$$\mathsf{UR}(E) = \frac{|G|^{-1} \sum_{i \in G} u_i(E)}{|B|^{-1} \sum_{i \in B} u_i(E)}$$

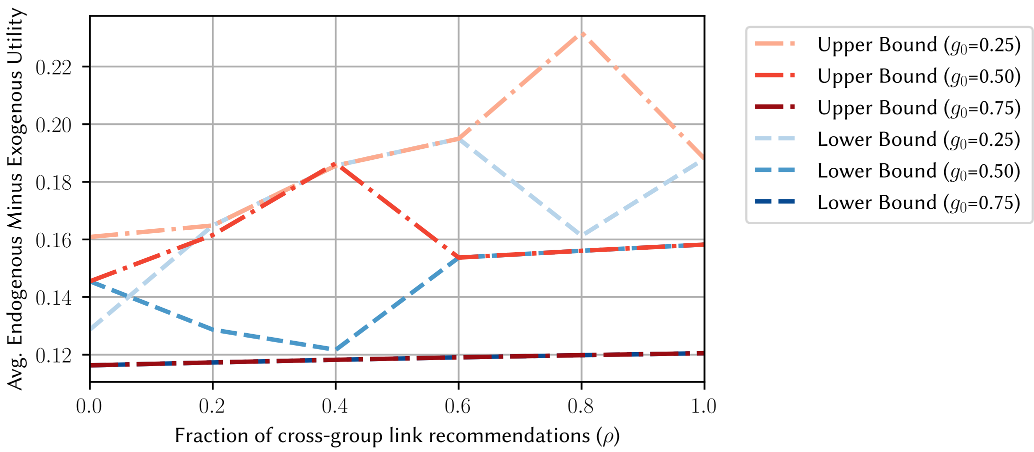

Figure 2 We plot upper (in red) and lower bounds (in blue) on average utility $n^{-1}\sum_{i=1}^{n} u_i(E)$ minus exogenous utility $(\sum_{i=1}^{n}(1-p_{i0}))$ against the fraction of cross-group link recommendations ρ, keeping $k=5$ constant. We set b_0 so that the exogenous utility ratio is always 2. We set $\gamma = 0.02$.

We will want to compare this quantity to the exogenous utility ratio, *i.e.*, the ratio that would be obtained in the absence of the platform and hence any networking:

$$\mathsf{UR}(\varnothing) := \frac{1-g_0}{1-b_0}.$$

We will sometimes refer to $\mathsf{UR}(\varnothing)$ as the *exogenous utility ratio* and $\mathsf{UR}(E)$ for $E \in \mathcal{E}$ as the *endogenous utility ratio*. A high utility ratio implies the privileged group fares much better than the unprivileged one, indicating high inequality. A utility ratio of 1 indicates equality of group welfare.

Our simple first result in this section states that for exogenous inequality sufficiently large, the existence of a *laissez-faire* platform only makes inequality worse and that, moreover, only the privileged group benefits from the platform. This holds for exogenous opportunity distributions where green members have a much larger chance of receiving extra exogenous opportunities, which is natural given their privilege. We also need green members to have some lower-bounded probability of receiving no exogenous opportunities.

▶ **Proposition 7.** *Suppose $k=0$ and there exist at least two green and one blue individual(s). Then for any exogenous opportunity distributions g and b such that $b_0(1-b_0-b_1) < g_0(1-g_0-g_1)$ and edge cost $\gamma \in (b_0(1-b_0-b_1), g_0(1-g_0-g_1))$, for all equilibria $E \in \mathcal{E}$,*

$$\mathsf{UR}(E) > \mathsf{UR}(\varnothing).$$

Moreover, the benefits of the platform accrue exclusively to the privileged group. That is, if for some $i \in [n]$ it holds that $u_i(E) > 1 - p_{i0}$, then $i \in G$.

The conditions in Proposition 7 are sufficient but not necessary for inequality to increase in the model: Under other relations between b and g the platform may still induce more inequality relative to their exogenous levels.

We illustrate Proposition 7 in Figure 3. We vary the exogenous utility ratio $\mathsf{UR}(\varnothing) = (1-g_0)/(1-b_0)$ on the horizontal axis and a lower bound (in blue) and upper bound (in red) on the endogenous utility ratio $\mathsf{UR}(E), E \in \mathcal{E}$ on the vertical axis. The black dashed line indicates where exogenous and endogenous utility ratios are equal. Anything above the line indicates greater endogenous utility than exogenous utility. Each line in the plot corresponds to a different value of g_0. Thus, for each line, moving right along the horizontal axis corresponds to increasing the value of g_0. Discontinuities in the plot occur as a result of changes in our bounds for the set of feasible equilibria as we sweep over different parameter values.[7]

As we can see in Figure 3, for sufficiently large values of $\mathsf{UR}(\varnothing)$, the lower bound on inequality for each value of g_0 becomes larger than the break-even value where $\mathsf{UR}(\varnothing) = \mathsf{UR}(E), E \in \mathcal{E}$. Thus, for all equilibria, the endogenous utility ratio becomes greater than the exogenous utility ratio. Depending on the value of g_0, the point at which $\mathsf{UR}(E) > \mathsf{UR}(\varnothing)$ for all $E \in \mathcal{E}$ occurs at different points. For g_0 small (lighter shades), it occurs at higher values of inequality: in this plot, when $g_0 = 0.25$, this happens close to $\mathsf{UR}(\varnothing) = 20$. For g_0 large (darker shades of red and blue), this happens at low values of inequality: in this plot, when $g_0 = 0.75$, it occurs at $\mathsf{UR}(\varnothing) < 5$. The reason for this is that, when g_0 is low, b_0 is also relatively low, so members of B may form more connections at low values of $\mathsf{UR}(\varnothing)$, resulting in inequality lower bounds below the break-even line. When g_0 is higher, b_0 is also higher, so members of B form few or no connections at low values of $\mathsf{UR}(\varnothing)$, and thus only members of G benefit from the network, leading to increased endogenous utility ratio.

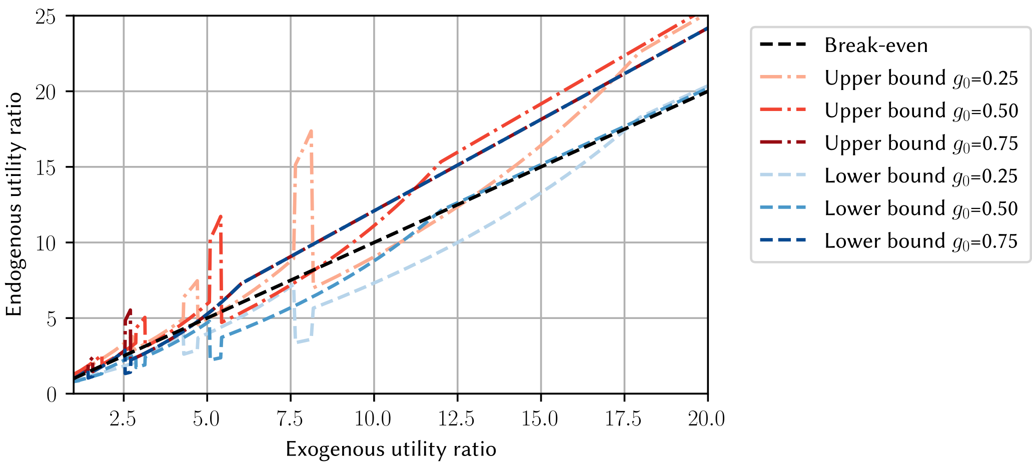

Figure 3 For large enough exogenous utility ratio, a network without link recommendations increases inequality, relative to exogenous levels. In the plot above, for each line, we fixed g and vary b plotting the value of $\mathsf{UR}(\varnothing) = (1 - g_0)/(1 - b_0)$ on the horizontal axis. On the vertical axis, we plot upper bounds (in red) and lower bounds (in blue) for the utility ratio ($\mathsf{UR}(E), E \in \mathcal{E}$) achieved for all equilibria at these parameter values. We set $\gamma = 0.04$.

[7] As the value of b_0 increases sweeping to the right on the horizontal axis, the bounds on the set of degrees for members of both groups change. Since degree bounds are integer values, changes to the degree bounds lead to discontinuous changes in functions of the degree bounds, such as the endogenous utility ratio on the vertical axis. In particular, consider $g_0 = 0.25$, around the $\mathsf{UR}(\varnothing) \in [7.5, 10]$ where the upper bound discontinuously jumps up and then down again. Before the step up, members of B can only have degree 2 in equilibrium. At the step up, they may have degree 2 *or* 1. After the step down, they may only have degree 1. Each of these changes lead to significant changes to the endogenous utility ratio, since the utilities of members of B is already small for these parameter settings and changes to their utility can have large effects through the denominator of the endogenous utility ratio. Hence, we see the sharp jump up and then down again. The changes in degree bounds for the lower bound on the endogenous utility ratio in the same range (and for all of the other "spikes" in Figure 3) lead to the corresponding discontinuous step down and then up.

Before we conclude this subsection, it is worth comparing Proposition 7 to existing results on inequality in the referrals and labor markets literature. An important point of departure between our work and prior work is that there are no exogenous in-group preferences between individuals. Exogenous in-group preferences are key to the results in Bolte et al. (2024) [6] and Okafor (2022) [29], where increased inequality emerges from the information advantage the privileged group has by virtue of receiving more referrals and the draining of opportunities from a minority group to the majority group, respectively. Increased inequality in our model emerges even without exogenous homophily. Instead, increased inequality emerges from the fact that individuals with greater privilege also have more leverage in forming connections. In particular, the equilibrium conditions Equations (2) and (3) imply that, in a sufficiently large population, almost all individuals can connect to others who are at least as valuable to them as they are to others. This, to our knowledge, a novel formalization for how inequality can emerge in a professional network and suggests that the existence of *laissez-faire* platforms may still increase inequality, even if individuals' in-group biases can be eliminated. Of course, additionally introducing exogenous homophily in our model would only induce greater inequality at equilibrium. We leave deeper exploration of the interactions between exogenous homophily and the greater leverage of privileged individuals for future work.

5.2 Cross-group recommendations reduce inequality

We now consider a platform that intervenes in the network formation process by selectively subsidizing connections, *i.e.*, choosing $k > 0$ and making k edge recommendations per person.

In our first result in this section, Proposition 8, we establish that for exogenous utility ratio sufficiently large, if a platform recommends all cross-group link recommendations, then inequality in any equilibrium will be strictly *less* than exogenous levels. Together with Proposition 7, this implies that a policy recommending cross-group links yields lower inequality, for all equilibria, than any equilibrium induced by a platform that does not make link recommendations. We formalize this statement in Corollary 9.

Recall that we denote the fraction of cross-group connections recommended to B as ρ and the corresponding fraction for members of G as $\rho\,|B|\,/\,|G|$, where we assume that ρ is set so as to ensure $\rho\,|B|\,/\,|G| \leq 1$. For our next result, we choose $\rho = 1$, so all edge recommendations are cross-group. We use $\mathcal{E}_k$ to denote the set of all equilibria induced by any k cross-group recommendations per person.

▶ **Proposition 8.** *For any $k > 0$, there exist constants $\underline{b}_0 < 1, \underline{n}$ such that for all $n \geq \underline{n}$, g, and b satisfying $b_0 \geq \underline{b}_0$, there exists constant $\bar{\gamma} > 0$ such that for all edge cost $\gamma \in (0, \bar{\gamma})$ and all equilibria with k cross-group link recommendations per person $E \in \mathcal{E}_k$, it holds*

$$\mathsf{UR}(E) < \mathsf{UR}(\varnothing).$$

Intuitively, the result follows from the fact that unprivileged individuals benefit more from connections to privileged individuals than privileged individuals benefit from connections to other privileged individuals. We visualize how inequality decreases relative to exogenous levels in Figure 4. We plot upper and lower bounds on inequality under different values of the exogenous utility ratio. As in Figure 3, the red lines correspond to upper bounds and the blue lines correspond to lower bounds and different shades of a color indicate different values of g_0, which (due to the fact that $g_\ell = 0$ for all $\ell \neq 0, 2$) uniquely determines the exogenous opportunity distribution of members of G. As before, the discontinuities in the plot indicate points where the set of feasible equilibria change, perhaps by allowing more or fewer connections for members of B.

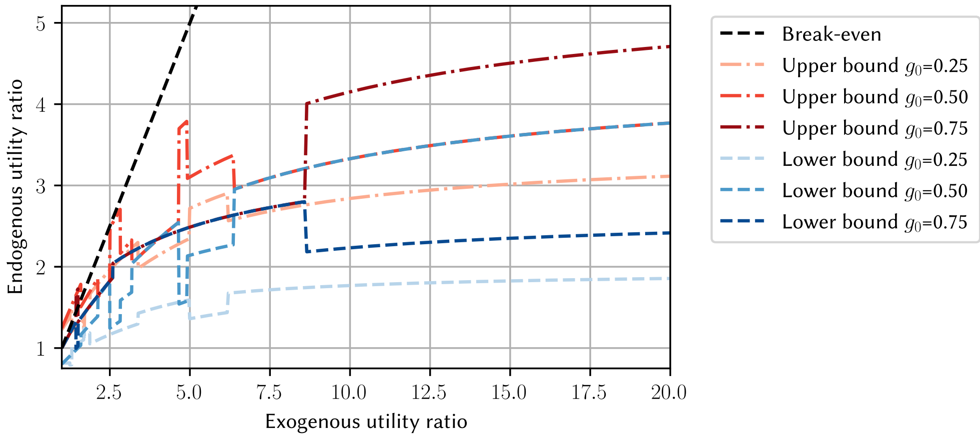

Figure 4 For large enough exogenous utility ratio, cross-group link recommendations reduce inequality, relative to exogenous levels. In the plot above, for each line, we fix g and vary b, plotting the value of $\mathsf{UR}(\varnothing) = (1 - g_0)/(1 - b_0)$ on the horizontal axis. On the vertical axis, we plot upper bounds for the utility ratio ($\mathsf{UR}(E)$, $E \in \mathcal{E}$) achieved for all equilibria at these parameter values. We set $\gamma = 0.04$, $k = 1$ and $\rho = 1$.

We can see that for each of the values of g_0 in the plot, once $\mathsf{UR}(\varnothing)$ is more than about 3, the upper bound on utility is below the break-even line, and therefore the endogenous utility ratio is guaranteed to be less than exogenous utility ratio. The point at which $\mathsf{UR}(E) < \mathsf{UR}(\varnothing)$ happens sooner for lower values of g_0 (displayed by lighter shades of red). This is because members of G are likely to receive an extra opportunity and thus the benefits to members of B are large. Indeed, for lower values of g_0, the endogenous utility ratio remains persistently lower than at higher values of g_0 as $\mathsf{UR}(\varnothing)$ grows large, since the members of B benefit relatively more when members of G are likely to receive an extra exogenous opportunity. Finally, we note that even a single cross-group link recommendation (we set $k = 1$ in Figure 4) results in the endogenous utility ratio growing sublinearly in exogenous utility. This demonstrates the increasing value of cross-group link recommendations for reducing endogenous inequality at greater levels of the exogenous inequality.

We also observe that with even a single link cross-group recommendation, endogenous inequality grows sublinearly in exogenous inequality. Thus, cross-group link recommendations are increasingly valuable at greater levels of exogenous inequality.

We note that the intuitive fact in our model that more cross-group connections leads to less inequality is not necessarily shared by other models of inequality in labor markets. For example, in Okafor (2022) [29], networks that are totally segregated yield equal labor market outcomes, while networks that are partially integrated lead to inequality.

A direct implication of Proposition 8 with Proposition 7 is that cross-group link recommendation policies reduce inequality relative to all equilibria resulting from *laissez-faire* platforms. We state this formally next.

▶ **Corollary 9.** *For parameters g, b, γ, n, k, ρ satisfying the conditions in Proposition 8, let $\mathcal{E}_k$ be the set of associated equilibria. Let $\mathcal{E}_0$ be the set of equilibria with the same parameters except that the platform makes no recommendations. Then for all $E_k \in \mathcal{E}_k$ and $E_0 \in \mathcal{E}_0$, it holds*

$$\mathsf{UR}(E_k) < \mathsf{UR}(E_0).$$

The result follows from the fact that under Proposition 7, it holds $\mathsf{UR}(\varnothing) < \mathsf{UR}(E_0)$ and under Proposition 8, $\mathsf{UR}(E_k) < \mathsf{UR}(\varnothing)$.

6 Discussion and future work

In this work, we explore the effect of platform link recommendation policies on inequality and inefficiency. We show that cross-group link recommendations decrease inequality and increase efficiency in terms of the total fraction of opportunity used. We also show that *laissez-faire* platforms, which do not intervene on the link recommendation process, increase inequality if inequality is large enough, for all equilibria. Our work has a number of natural extensions that would be interesting to pursue in future work.

Search costs. In our model, we assume individuals pay a fixed cost for each connection. Another reasonable model would consider situations in which individuals pay a fixed cost to *search* for a connection. In this version, an individual who pays the search cost might receive a random potential connection and then be able to decide whether to connect or not.

Asymmetric connection costs. In many situations, there may be asymmetric costs for connections, where the individual proposing the connection pays the cost (perhaps viewed as the cost of proposing, going out on a limb, or suffering the possibility of rejection).

More than two groups. Ultimately, real-world platforms consist of more than two levels of exogenous privilege, and it would be interesting to study $m > 2$ groups, each with identical distributions. If we took $k \to \infty$, we could induce a probability distribution over exogenous opportunity distributions, which would allow for more precise quantification of inequality calibrated to real-world populations.

Heterogeneous opportunities. It would also be interesting to model opportunities as (possibly binary) vectors to model complementarities between connections. For example, plumbers and appliance repair technicians may encounter many different types of home repair jobs but each only be able to capitalize on certain types of jobs.

Endogenous opportunity creation. It is also worth considering how connections between individuals may *create* opportunities. Certain situations might be generative, where the intrinsic fact that two people connect creates opportunities, like a pair of researchers with good chemistry who together generate great project ideas.

Noisy or imperfect information of exogenous opportunity distributions. Finally, we note that platforms and individuals may not have full information about individuals' opportunity distributions (or group memberships). Thus, they may need to act on predictions about individuals, which could lead to effects on inequality or efficiency.

We hope our work inspires further exploration of inequality in professional networking, the role of platforms, and how recommendation systems may have important implications in society.

References

1. Attila Ambrus, Markus Mobius, and Adam Szeidl. Consumption risk-sharing in social networks. *American Economic Review*, 104(1):149–182, 2014.
2. Marcelo Arbex, Dennis O'Dea, and David Wiczer. Network Search: Climbing the Job Ladder Faster. *International Economic Review*, 60(2):693–720, 2019. doi:10.1111/iere.12375.

3 Kenneth J Arrow and Ron Borzekowski. Limited Network Connections and The Distribution Of Wages. Technical report, Board of Governors of the Federal Reserve System (U.S.), 2004.

4 Patrick Bayer, Stephen L. Ross, and Giorgio Topa. Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes. *Journal of Political Economy*, 116(6):1150–1196, 2008. doi:10.1086/595975.

5 Lori Beaman, Niall Keleher, and Jeremy Magruder. Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Journal of Labor Economics*, 36(1):121–157, 2018. doi:10.1086/693869.

6 Lukas Bolte, Matthew O. Jackson, and Nicole Immorlica. The role of referrals in immobility, inequality, and inefficiency in labor markets. *Journal of Labor Economics*, 0(ja):null, 2024. doi:10.1086/733048.

7 Meta Brown, Elizabeth Setren, and Giorgio Topa. Do informal referrals lead to better matches? evidence from a firm's employee referral system. *Journal of Labor Economics*, 34(1):161–209, 2016. doi:10.1086/682338.

8 Konrad B. Burchardi and Tarek A. Hassan. The Economic Impact of Social Ties: Evidence from German Reunification. *The Quarterly Journal of Economics*, 128(3):1219–1271, 2013. URL: https://www.jstor.org/stable/26372522.

9 Stephen V. Burks, Bo Cowgill, Mitchell Hoffman, and Michael Housman. The Value of Hiring through Employee Referrals. *The Quarterly Journal of Economics*, 130(2):805–839, 2015. doi:10.1093/qje/qjv010.

10 Antoni Calvó-Armengol and Matthew O. Jackson. The Effects of Social Networks on Employment and Inequality. *American Economic Review*, 94(3):426–454, 2004. doi:10.1257/0002828041464542.

11 L. Elisa Celis, Chris Hays, Anay Mehrotra, and Nisheeth K. Vishnoi. The Effect of the Rooney Rule on Implicit Bias in the Long Term. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 678–689, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3442188.3445930.

12 Raj Chetty, Matthew O Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, et al. Social capital ii: determinants of economic connectedness. *Nature*, 608(7921):122–134, 2022.

13 Federico Cingano and Alfonso Rosolia. People I Know: Job Search and Social Networks. *Journal of Labor Economics*, 30(2):291–332, 2012. doi:10.1086/663357.

14 Stephen Coate and Glenn C. Loury. Will Affirmative-Action Policies Eliminate Negative Stereotypes? *The American Economic Review*, 83(5):1220–1240, 1993. URL: https://www.jstor.org/stable/2117558.

15 Paul DiMaggio and Filiz Garip. Network Effects and Social Inequality. *Annual Review of Sociology*, 38(Volume 38, 2012):93–118, 2012. doi:10.1146/annurev.soc.012809.102545.

16 Cynthia Dwork, Chris Hays, Nicole Immorlica, Juan C. Perdomo, and Pranay Tankala. From Fairness to Infinity: Outcome-Indistinguishable (Omni)Prediction in Evolving Graphs. In Nika Haghtalab and Ankur Moitra, editors, *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 1564–1637. PMLR, 30 June–04 July 2025. URL: https://proceedings.mlr.press/v291/dwork25a.html.

17 Cynthia Dwork, Chris Hays, Jon Kleinberg, and Manish Raghavan. Equilibria, efficiency, and inequality in network formation for hiring and opportunity. *arXiv preprint arXiv:2402.13841*, 2024. doi:10.48550/arXiv.2402.13841.

18 Yulia Evsyukova, Felix Rusche, and Wladislaw Mill. Linkedout? a Field Experiment on Discrimination in Job Network Formation. *The Quarterly Journal of Economics*, 140(1):283–334, 2025.

19 Hoda Heidari, Solon Barocas, Jon Kleinberg, and Karen Levy. Informational Diversity and Affinity Bias in Team Growth Dynamics. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–10, Boston MA USA, October 2023. ACM. doi:10.1145/3617694.3623238.

20 Hoda Heidari and Jon Kleinberg. Allocating Opportunities in a Dynamic Model of Intergenerational Mobility. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 15–25, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3442188.3445867.

21 Yannis M. Ioannides and Linda Datcher Loury. Job Information Networks, Neighborhood Effects, and Inequality. *Journal of Economic Literature*, 42(4):1056–1093, December 2004. doi:10.1257/0022051043004595.

22 Matthew O Jackson, Tomas Rodriguez-Barraquer, and Xu Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, 102(5):1857–1897, 2012.

23 Matthew O. Jackson and Asher Wolinsky. A Strategic Model of Social and Economic Networks. *Journal of economic theory*, 71(1):44–74, 1996.

24 Denneal Jamison-McClung. Mentorship, Sponsorship, and Professional Networking. *Uprooting Bias in the Academy: Lessons from the Field*, pages 175–187, 2022.

25 Jon Kleinberg and Manish Raghavan. Selection Problems in the Presence of Implicit Bias. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, pages 33:1–33:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ITCS.2018.33.

26 David Liben-Nowell and Jon Kleinberg. The Link Prediction Problem for Social Networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, CIKM '03, pages 556–559, New York, NY, USA, November 2003. Association for Computing Machinery. doi:10.1145/956863.956972.

27 Nan Lin. Building a Network Theory of Social Capital. *Social Capital*, pages 3–28, 2017.

28 Nan Lin and Mary Dumin. Access to Occupations through Social Ties. *Social Networks*, 8(4):365–385, 1986. doi:10.1016/0378-8733(86)90003-1.

29 Chika O. Okafor. Seeing Through Colorblindness: Social Networks as a Mechanism for Discrimination. *The Journal of Law and Economics*, 68(3):519–560, 2025. doi:10.1086/735081.

30 David S. Pedulla and Devah Pager. Race and Networks in the Job Search Process. *American Sociological Review*, 84(6):983–1012, 2019. doi:10.1177/0003122419883255.

31 Karthik Rajkumar, Guillaume Saint-Jacques, Iavor Bojinov, Erik Brynjolfsson, and Sinan Aral. A Causal Test of the Strength of Weak Ties. *Science*, 377(6612):1304–1310, 2022.