2023

Multilayer Network Model of Gender Bias and Homophily in Hierarchical Structures

Emerson McMullen

Multilayer Network Model of Gender Bias and Homophily in Hierarchical Structures

Emerson McMullen

Heather Zinn-Brooks, Advisor

Mason A. Porter, Reader

HARVEY MUDD COLLEGE

Department of Mathematics

May, 2023

Abstract

Although women have made progress in entering positions in academia and industry, they are still underrepresented at the highest levels of leadership. Two factors that may contribute to this leaky pipeline are gender bias, the tendency to treat individuals differently based on the person's gender identity, and homophily, the tendency of people to want to be around those who are similar to themselves. Here, we present a multilayer network model of gender representation in professional hierarchies that incorporates these two factors. This model builds on previous work by Clifton et al. (2019), but the multilayer network framework allows us to track individual progression through the hierarchy and relationships at the level of individual agents. We use this model to investigate how the network structure and location of female and male nodes within a given network affect gender representation throughout the hierarchy.

Contents

List of Figures

List of Tables

Acknowledgments

I would like to express my sincerest thanks to my thesis advisor Heather Zinn-Brooks, whose guidance, mentorship, and support have meant so much to me throughout the past few years. She has helped me build my skill set and self-confidence and is a large part of the reason I am taking the next steps that I will be taking.

Numerous mentors at Harvey Mudd College have also greatly shaped my experience here. Jon Jacobsen has mentored and advised me with kindness and respect ever since I stepped foot into my very first class in college, Math 19. Arthur Benjamin helped me discover my love for math, and Andrew Bernoff helped me discover my love for applied math. I am incredibly grateful to all of you.

Thank you to Todd Kresser, who made high school calculus fun, and is probably the reason I majored in math.

To Miles, Jo, and Carmel, thank you for being the best friends I ever could have asked for. I couldn't have done it without you.

To my real family, thank you for twenty-two years of support and love. I couldn't have done it without you either.

There are certainly many more people that helped shape my journey so far, and to all of you, I am incredibly grateful.

And to you, the reader of this thesis, I say thanks for taking the time to pick up this work. I hope you enjoy it!

Chapter 1

Introduction

1.1 Gender Bias and Inequality

Gender bias and gender inequality are prevalent throughout many societies and affect many people in a variety of ways. Gender bias can be formally defined as a person receiving different treatment based on the person's real or perceived gender identity (Reference (2023)). Gender bias affects many people in many ways, but one particularly important inequality is the gender pay gap. According to the United Nations, globally, women earn only 77 cents for every dollar that men earn. Women are concentrated in low-wage, unskilled jobs, and also do two and a half times as much unpaid work as men (United Nations). Although these quantities have improved as compared to historical averages, it will take 257 years to reach equal pay at the current rate.

While there are many factors that contribute to the gender pay gap, in this thesis we will focus on the disparity between the number of men and women in leadership positions. One especially visible example of this is in Chief Executive Officers of Fortune 500 companies, where only 8.8% of them are women (Women Business Collaborative). Another example of a type of hierarchical organization that is especially relevant to this thesis is mathematics departments. According to the Association for Women in Mathematics, in 2017, 38% of mathematics bachelor's degrees were given to women, 29% of PhDs were awarded to women, 21% of new postdocs were women, only 17% of tenured or tenure track faculty were women, and only 12% of full professors at PhD-granting institutions were women. One possible explanation is the so-called leaky pipeline, which describes how women and minorities are steadily lost at each stage in STEM fields; however, more work is needed to determine the cause of this particular example. Because there are very few female full professors who are mentoring students, prospective female graduate students may feel less empowered to pursue graduate study, and even if they do, may not have the same outcomes as men. For example, Schwartz et al. (2021) found that researchers in the life sciences who had been awarded an outstanding distinction trained male graduate students at a higher rate than their other colleagues. They also found that female mentors have less access to resources and that mentors and trainees tend to be the same gender. This implies that mentors of women tend to have fewer resources, contributing to the disproportionately small number of women obtaining independent research positions.

1.2 Existing Mathematical Work on Professional Gender Balance

There are many mathematical tools that researchers have been using to try to better understand gender disparity in professional settings. In this section, we will discuss previous projects which studied networks of collaboration on papers between researchers in academia, academic recommendation letters via natural language processing, and agent-based models of women in corporations. Bellotti et al. (2022) investigated differences between male and female professional networks in academia, specifically whether men and women received the same benefits from having similar collaboration networks. They found that in certain disciplines, despite having similar social networks, women received less funding than men. They suggested that increasing the representation of women in academic fields alone is not enough to reduce inequalities. They concluded that this may be due to the fact that women occupy fewer leadership positions in research, such as principal investigators.

Lerman et al. (2022) analyzed author citation networks of researchers elected to the National Academy of Sciences to better understand gender disparities in senior faculty positions. They were able to identify gender disparities in citation networks, such as women receiving fewer citations over the course of their career than men do and women reciprocating a significantly higher fraction of citations than men do. Furthermore, the authors were able to infer the author's gender based solely on their citation network using a random forest classifier. They also found that women tend to have a higher clustering coefficient than men and that their peers are more productive, which may indicate that women are more often a part of tight-knit research communities. One possible reason that women's advancement in academic careers is limited could be evaluations from those at higher levels in the professional hierarchy. Using natural language processing, Bernstein et al. (2022) found that letters of recommendation tend to describe women more often in terms of their work ethic and drive compared to men. However, both genders are equally likely to be labeled as exceptional or standout students in these letters.

Mathematicians have also built models incorporating individual action and progression to understand how and where gender parity emerges in hierarchical organizations. Du et al. (2022) used an agent-based model to examine how women are affected by bias in the workplace. In this model, a single instance of bias, such as being penalized slightly more for making a mistake, doesn't have a huge impact at that moment, but over time, bias drastically affects women. Du et al. (2022) incorporated six different types of bias into their model, which include that women's successes on projects are valued less than men's, women's failures and errors are penalized more than men's, women receive less credit when working on teams with men, women receive more blame when their mixed-gendered team fails, women are penalized for behaving non-altruistically, and that women receive fewer opportunities for growth. Their results showed that in their model, women must have a higher success rate to be promoted to the highest level of the corporation than men and that individual biases against female success and failure have the greatest effect on women's career paths. They also examined the effect of intervention strategies, such as quotas, on their model and determine that as long as macro-level discrimination still exists, they won't be effective in the long run.

1.3 Previous Model of Gender Balance in Hierarchical Organizations

Clifton et al. (2019) created a model of the fraction of each gender in hierarchical organizations using ordinary differential equations. They model an organization in which people occupy positions at different levels. People are hired into the lowest level of the organization and have the potential to rise through the ranks or quit. To progress through the levels, people on the lower level apply for promotion, and people on the next level decide which applicants to promote. Bias affects the decisions of the higher level, and homophily, the tendency to seek out people who are similar to oneself, affects the decision to seek promotion. In the model, it is assumed that if the next level in an organization has a high fraction of women, then a woman in the lower level is more likely to seek promotion. Clifton et al. (2019) defined bias as all conscious or unconscious decisions made by an employer that are affected by the applicant's gender and assumes that bias is the same in each level of the organization.

In their model, Clifton et al. (2019) made a few key assumptions. They assumed that the probability of promotion is a function of the fraction of people who share the applicant's gender and the fraction of like-gendered individuals in the applicant's current level. The probability of seeking promotion did not change based on gender. Bias was a constant term that exists uniformly throughout the professional hierarchy, and people were promoted, which leaves vacancies that are subsequently filled.

To model the probability of seeking promotion, Clifton et al. (2019) used a sigmoid function that incorporates a gradual switch from being not likely to apply to very likely to apply as homophily increases:

$$P(u,v) = \frac{1}{1+e^{-\lambda(u-v)}},$$

where u is the fraction of like-gendered individuals in the next level up of the hierarchy, v is the fraction of like-gendered individuals in the current level, and λ is the strength of the homophilic tendency. Then, they model the fraction of women who are promoted using the function

$$f(u,v;b) = \frac{bvP(u,v)}{bvP(u,v)+(1-b)(1-v)P(1-u,1-v)},$$

where u is the fraction of women in the higher level of the professional hierarchy, v is the fraction of women in the lower level, and b is the fraction of women promoted if the applicant pool has an equal number of men and women and takes on possible values in the set $[0,1]$. If $b < \frac{1}{2}$, then women are disproportionately not promoted and if $b > \frac{1}{2}$ women are disproportionately promoted (Clifton et al. (2019)). Based on the function $f(u,v;b)$, Clifton et al. (2019) formulated a system of differential equations that incorporate movement between levels.

To understand the effects of bias and homophily, Clifton et al. (2019) considered multiple variations of their model. First, they examined a model with no effects from either bias or homophily. They set the bias constant to

$b = \frac{1}{2}$ and let the probability of seeking promotion be a constant. With these choices, the model reaches a stationary state where women represent half the employees at each level of the professional hierarchy. In the model which only considers the effect of bias, if $b < \frac{1}{2}$, then the steady-state gender fraction is reduced in comparison to when $b = \frac{1}{2}$, with higher levels having fewer women. In the model which only considers the effect of homophily, there are three different types of outcomes. When the strength of the homophilic tendency is low, the organization reaches gender parity on all levels, but when the strength is moderate, they observe oscillations in the fraction of women with respect to time at all levels. When the strength of homophily is high, the system will converge to having either all men or all women on each level. In the model which considers both homophily and bias, the system experiences very similar long-term behavior to the bias-free model.

Clifton et al. (2019) applied their model to various academic fields by estimating bias and homophily values for each field from data. Using these values, they concluded that fields such as law and medicine that have bias values near $\frac{1}{2}$ and weak homophily will reach gender parity, whereas fields with strong homophily, like engineering and nursing, should become either male or female-dominated. Finally, fields with strong gender bias, like math and computer science, will never reach gender parity at the highest levels.

This model is valuable because it incorporates bias and homophily into a model which predicts gender balance over time and can be used to examine different institutions and policies. However, there are a few limitations of this model. They don't consider any individual connections and interactions in the decision to apply for promotion, and they assume that the population of women and men is well-mixed on each level. They also assume that all people on the next level in the professional hierarchy have the same bias, whereas it is likely that women and men have different gender biases in reality. They also assume that homophily is something experienced universally throughout a level, rather than on an individual basis. Finally, while bias and homophily may change over time, they assume that both values are constant.

1.4 Goals of Thesis

In this thesis, we extend the work of Clifton et al. (2019). We model the same hierarchical organizational structure, but instead of using differential equations, we use a probabilistic process on networks. Differential equations assume that populations are well-mixed, but networks allow us to examine relationships and how individual interactions impact gender balance in more depth. We model homophily and bias for each individual in the network in the following way. Homophily depends on the number of individuals of the same gender to which an individual is connected in their current level of the organization. Individuals are promoted to the next level based on the bias of those in the level above and on the homophily of the individual. It is important to include a hiring and retirement process so we can analyze the fraction of women in the organization in a longer time frame. To encode the hierarchical nature of the organization, we use a multilayer network Kivelä et al. (2014). Using this model, we investigate how different models or levels of bias affect the gender balance in the organization, the impact that the location of female and male nodes within the network structure have on the behavior, as well as the effects of the initial network structure.

The rest of the thesis will proceed as follows. In Chapter 2, we will provide mathematical background, and in Chapter 3 the model formulation is discussed. Results from the unbiased model are presented in Chapter 4, and in Chapter 5 we examine the results from the biased model. Finally, the thesis concludes in Chapter 6.

Chapter 2

Mathematical Background

2.1 Networks

2.1.1 Single-Layer Network

A network G is a mathematical structure in which a set of vertices V (also called nodes) are connected by a set of edges E. If there is an edge from vertex u to vertex v, we can write $(u, v) \in E$. An undirected network has undirected edges, indicating that the two nodes are connected in the same way, while directed networks have directed edges that encode one-way relationships. An abstract undirected network is shown in Figure 2.1. Directed networks are useful because they can encode individual relationships and social structures. For example, they can encode the social structures of a second-grade classroom, as seen in Figure 2.2, or could represent members of a family who are in contact with one another.

2.1.2 $G(n, p)$ Random Network Model

One important type of random network model is called the Erdös-Rényi model, also called a $G(n, p)$ model Brooks (2023). When using this model to generate a network, we fix n nodes and let p be the probability of an edge connecting any pair of nodes. Each graph G with n nodes has a probability $P(G) = p^m(1-p)^{\binom{n}{2}-m}$, where m is the number of edges. The $G(n, p)$ model generates a simple, random graph. We will use the $G(n, p)$ generated network as an important building block of our model.

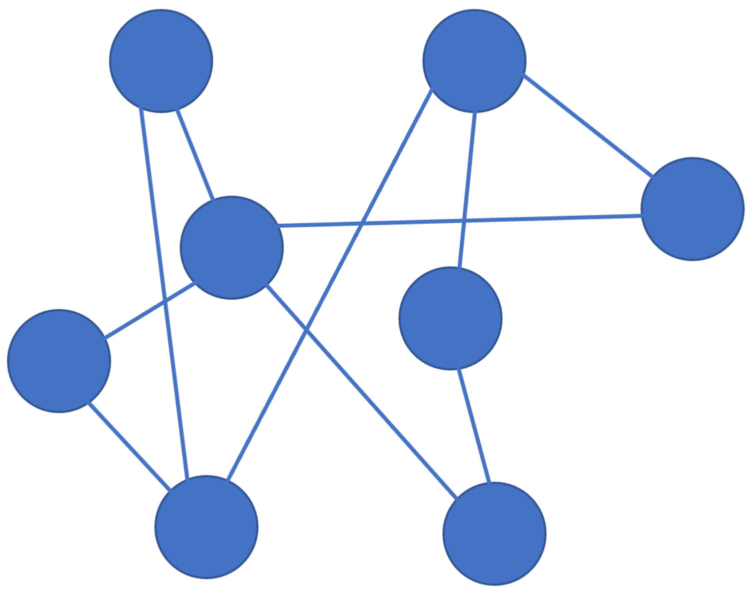

Figure 2.1 An abstract, undirected network. Circles represent nodes, while lines represent edges.

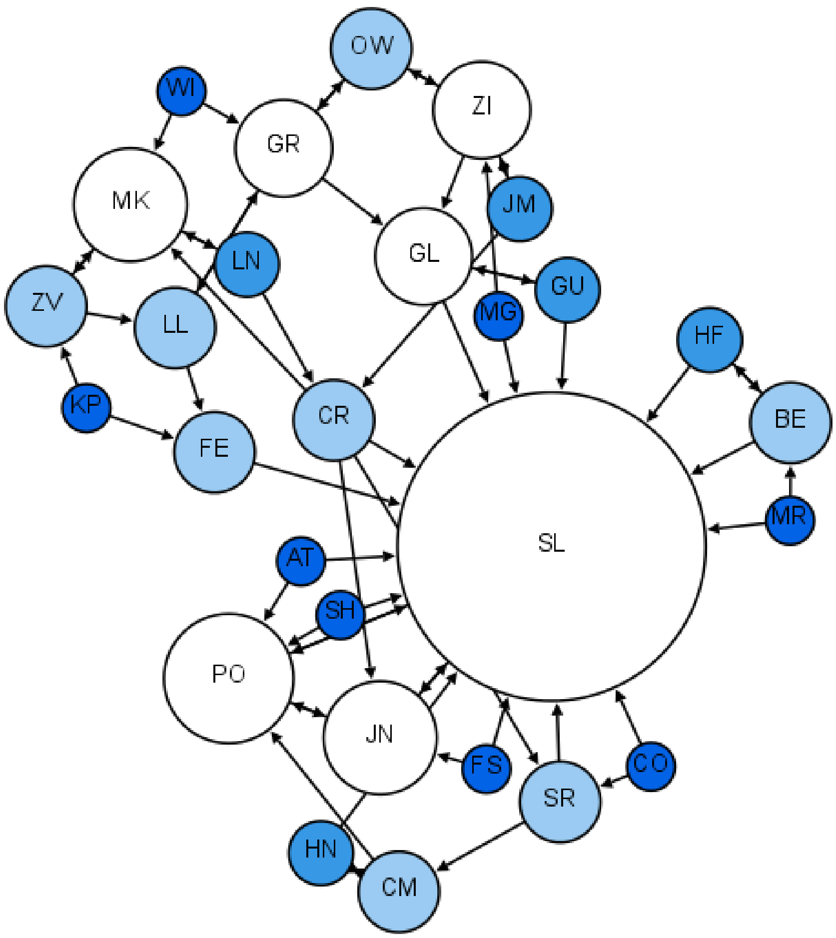

Figure 2.2 Directed network of a second-grade classroom. Nodes are labeled with the initials of each child in the class. Directed edges represent social relationships between the second-graders. Grandjean (2015)

2.1.3 Modularity

Vertices within a network may be assigned additional values or traits. For example, we can assign vertices to be either male or female and can assign values that could represent the amount of money that they make. Categorical traits divide the vertices into groups, and modularity is a number that tells us how connected vertices are to members of their own group. Modularity can be defined for a partition into any number of groups, but because in this thesis we only have two groups, we will focus on that case. When considering two groups, modularity is calculated as

$$M = \frac{1}{4m} \sum_{ij} B_{ij} s_i s_j$$

where m is the number of edges, k_i is the degree of node i, and $s_i = 1$ if it is in the first group and $s_i = -1$ if it is in the second. The expression for B_{ij} is $B_{ij} = A_{ij} - \frac{k_i k_j}{2m}$. The adjacency matrix, A, has entry $i, j = 1$ if there is an edge between nodes i and j. If we sum up $\sum_{ij} A_{ij}\delta_{ij}$ where $\delta_{ij} = 1$ if i and j are in the same group and 0 otherwise, we get two times the total number of edges between members of the same group. If we multiply this sum by $\frac{1}{2m} \sum_{ij} A_{ij}\delta_{ij}$, we get the fraction of edges that belong to members of the same group. Then, $\frac{k_i k_j}{2m}$ is the expected value of the number of same-group edges between nodes with the degree of i and degree of j. The sum $\sum_{ij} \frac{k_i k_j}{2m}\delta_{ij}$ is two times the number of expected edges between nodes of the same group, and we can again divide by $2m$ to get the expected fraction of edges between same-group nodes. Putting these together, $\sum_{ij} B_{ij}\delta_{ij}$ is the actual fraction of edges between same-group nodes minus the expected fraction of edges between same-group nodes. Some algebra leads us to the final equation

$$\frac{1}{4m} \sum_{ij} B_{ij} s_i s_j.$$

When M is positive, members of the same group are more connected with each other, while M being negative means that it is more likely for an edge to exist between nodes of different groups than nodes of the same group. If $M = 0$, then all vertices are equally likely to be connected to nodes in either group. The concept of modularity will help us understand how network structure affects the graph.

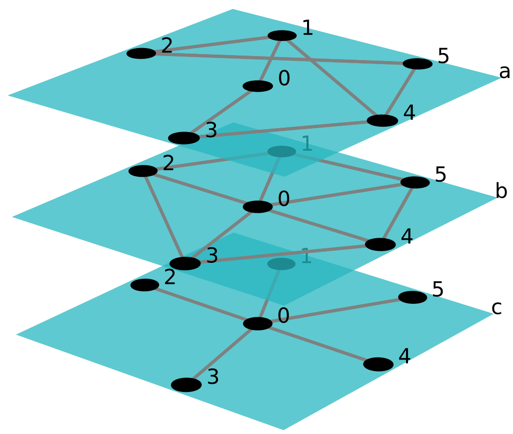

Figure 2.3 An example of a multilayer network with three layers, drawn using Kivelä (2013)

2.2 Multilayer Networks

A multilayer network is a network with multiple layers. Multilayer networks can contain both intra-layer and inter-layer edges. Often, multilayer networks are used to represent different types of relationships between the same nodes. In my model, each level of the multilayer network represents a level in the professional organization. Figure 2.3 shows an example of a multilayer network.

2.3 Binomial Distribution

The binomial distribution is a probability distribution that describes the number of successes in a fixed number of independent trials, each with the same probability of success. It is often used to model situations where there are only two possible outcomes for each trial, such as “heads" or “tails."

The binomial distribution is characterized by the probability of success in a single trial p and the number of trials n. Figure **??** shows two example binomial distributions with different probabilities of success.

In this model, we use a binomial distribution to make random choices between two options with a probability of success p. The binomial distribution models the number of successes given n trials with probability of success p. In the special case where $n = 1$, this is called a Bernoulli distribution.

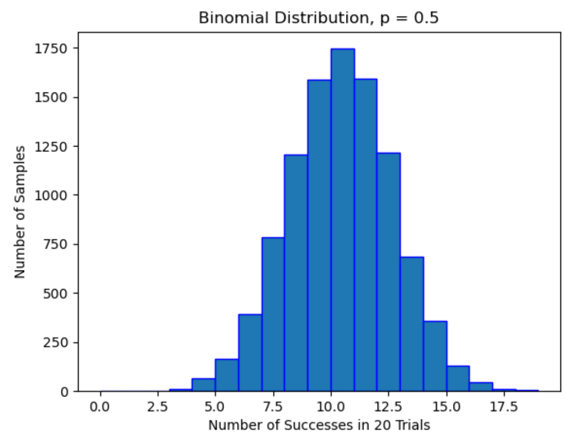

a. $p = 0.5$

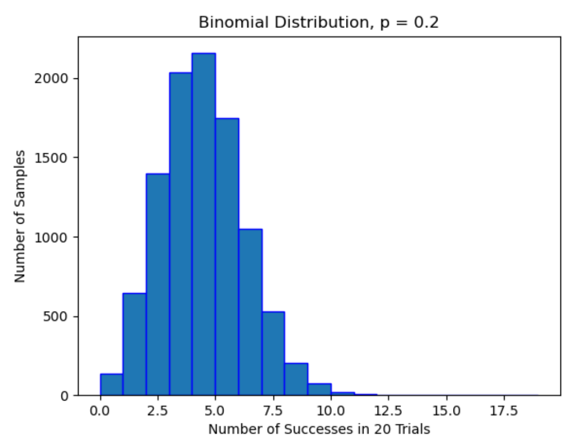

b. $p = 0.2$

Figure 2.4 Two binomial distributions run with $n = 20$ trials and a sample size of 10000. The probability of success differs between the two distributions.

Chapter 3

Model Formulation

3.1 Overview of the problem

In this thesis, we create a network model to study gender balance in hierarchical organizations such as corporations and academic institutions. This work is an extension of the work by Clifton et al. (2019), where they modeled the fraction of women in hierarchical organizations using a system of ordinary differential equations, studying gender balance. In the Clifton et al. (2019) model, the fraction of women at each level of the organization is affected by people hired to the lowest level, people who leave the company, and people who are promoted. To determine the fraction of promoted people who are women from a given level, the authors first determine the fraction of women in an applicant pool using a function of the fraction of like-gendered individuals in the applicant's current level and level above. Then, a bias parameter affects the probability that women will be promoted from the applicant pool.

In my model, we create an agent-based model which incorporates individual interaction. To do this, we use a multilayer network to encode the organizational structure in which individuals are represented by nodes, relationships are represented by edges, and levels of the organization are represented by layers. Nodes are adjacent to other nodes in their current layer, and we do not incorporate edges between layers. Each node is assigned a homophily value, which we take to be the number of like-gendered nodes it is connected to. Individuals decide to apply for promotion as a function of the fraction of the homophily value over the number of neighbors it has. From the applicant pool, nodes are promoted based on a constant probability

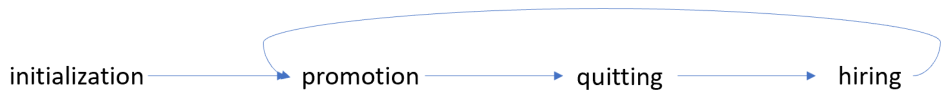

Figure 3.1 This figure depicts the general flow of my model in one level of an organization. It starts by initializing a multilayer network that represents each individual, their gender, and their social connections by their edges. From there, we cycle through promotion, quitting, and hiring for a specified number of time steps.

of promotion, with a bias term reducing the probability of promotion for female nodes. Then, new individuals are hired into the lowest level and some quit from each level depending on some probability. Figure 3.1 shows the progression of the model on a given level of the organization.

My model is structured using a multilayer network. Although each node is only on one layer, each node has a placeholder node on every layer it is not actually on. Every node in the multilayer network is either labeled as an "active" or "shadow" node. Shadow nodes are the placeholder nodes, while active nodes are the nodes that actually are on that layer. Shadow nodes Because we keep track of which nodes are active and which are not, the shadow nodes do not impact the dynamics overall.

In the multilayer model, we calculate homophily as the number of like-gendered neighbors of a node. First, we determine which nodes apply for promotion using a probability based on their homophily values, which are the number of like-gendered neighbors to which they are adjacent. Then, those nodes are promoted with a set probability. Once we know which nodes are promoted, in their current layers, we set their status in their status to "shadow" and remove all of their edges, while in their new layer, we set their status to "active" and add new edges, with the same probability of an edge occurring between any two nodes, p_g, that was used in graph construction. There is a possibility that nodes are promoted into the next layer and leave a node that has no neighbors and no chance of being promoted but is still active. We leave these nodes as they are.

When a node is hired, we add that node to all layers. However, we only set the node to be active in the bottom layer. When a node quits, we set it to be a shadow node in all layers and remove all of its edges, so it is effectively not a part of the network anymore. See Figure 3.2 for an overview of how the model works.

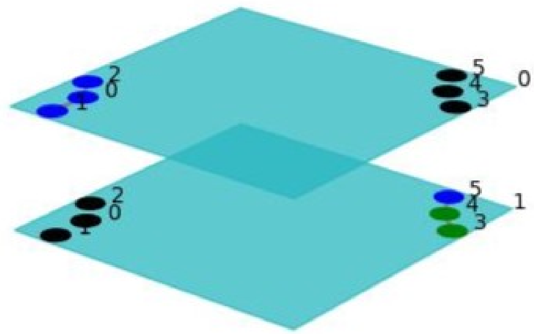

a. The initialization of the multilayer network. Nodes 0, 1, and 2 are active in layer 0, the lowest level of the organization, and nodes 3, 4, and 5 are active in layer 1, the highest. Note how nodes 0, 1, and 2 appear in layer 1 and nodes 3, 4, and 5 appear in layer 0 with no edges as shadow nodes. Shadow nodes are shown in black, where male nodes are blue and female nodes are green.

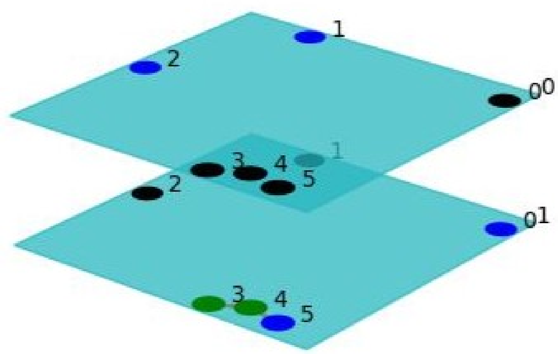

b. The multilayer network after promotion. Node 0 was promoted from layer 0 to layer 1, making no new connections. It was disconnected from its old neighbors and is a shadow node in layer 0.

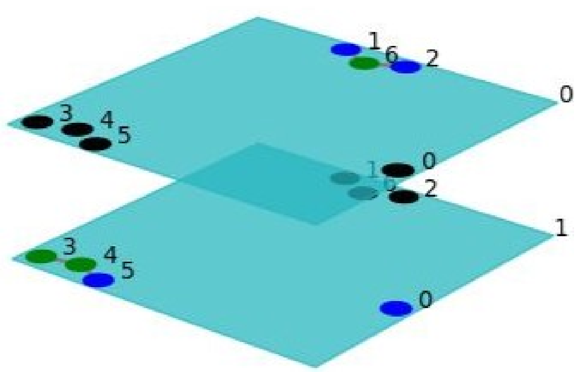

c. The multilayer network after quitting and hiring. None of the nodes quit, while node 6 was hired into layer 0, forming edges with nodes 1 and 2.

Figure 3.2 Example of the multilayer network model

3.2 Initialization

In the initialization step, we need to construct the initial structure of the multilayer network and all of the attributes of the nodes. The details of this algorithm can be found in the appendix in Algorithm 1.

Let p_g be the probability that any given pair of nodes are adjacent, and n_i be the number of nodes on layer i. The first thing we do in this step is to create a $G(n_i, p_g)$ graph for each layer. We use the $G(n, p)$ model because of it's simplicity. When we hire nodes, it is easy to generate new adjacencies as described by the $G(n, p)$ network model. Then, we create a multilayer network by adding each $G(n_i, p_g)$ starting graph to its corresponding layer. For example, if we want to model a structure with one hundred nodes on the first layer and thirty nodes on the second layer, we would generate a $G(100, p_g)$ single-layer network and use that to initialize the bottom layer of the multilayer network and a $G(30, p_g)$ single-layer network which to initialize the top layer.

We assign genders to each node, which we do by drawing from the binomial distribution with mean p_w, the probability of a node being assigned female. Then, each node gets a homophily value, h, which is it's number of like-gendered neighbors. For example, if a female node is adjacent to two nodes, one male and one female, its homophily value will be 1. For each node that we add to the multilayer network, we add a "shadow" version of the node on every other layer of the network. This makes computation easier when we promote nodes from one layer into the next. We give each node the "shadow" or "active" attribute.

3.3 Promotion

In promotion, we first determine which nodes are in the applicant pool, and from that subset we determine which are promoted to the next layer. The implementation can be found in Algorithm 2. We do this process on each layer except for the last one as no one is promoted from the top layer.

The first thing to do is to find which nodes are seeking promotion. To do this, we calculate the probability of being promoted into the pool for each node. This probability is the node's homophily divided by its number of neighbors. Then, we draw from the binomial distribution with that probability, and if that is successful, we add that node to the subset of nodes applying for promotion.

Once we know which nodes are seeking promotion, we go through the process of promoting nodes into the next layer. To determine which nodes are promoted, again we draw from a binomial distribution and a constant probability of success. The probability of promotion is multiplied by b, the bias constant, if we are deciding on promotion for a female node. This serves to incorporate institutional gender bias into the system.

If the draw from the binomial distribution was successful, then the node is promoted. We set that node on its current layer to "shadow," remove all of its edges, and adjust the homophily values of all its former neighbors accordingly. On the next layer up, we change the node to be "active." For each active node on the layer the node was promoted into, we draw from the binomial distribution with probability of success p_g, and add an edge between the promoted node and the other node. We also update the homophily values accordingly.

3.4 Hiring

To perform hiring, we first need to determine how many individuals will be hired. We want to hire some percentage of the number of people on the bottom layer, so we first determine the number of people on the bottom layer. Then, the number of nodes that we hire is $P_{hire} \times activeInLowest$, where P_{hire} is the percentage of the bottom layer we want to hire in. If the number of nodes active in the bottom layer is small enough, we will not hire any nodes. To ensure that we always have at least one node in the bottom layer, we actually hire the maximum of 1 and the calculated number of nodes.

Once we know how many nodes to hire, we add each new node to the network. Each node is only active in layer 0, the lowest layer, has its corresponding "shadow" nodes in the other layers, and is assigned gender in the same way that it is done during initialization. The homophily values are initially set to 0 because they do not have any neighbors yet.

The next step is to add neighbors to the new nodes. We add new nodes using the $G(n, p)$ network model. Each new node has a p_g probability of having an edge with every other active node in the bottom layer. For each possible edge, we add it using a binomial generator with p_g, the probability of having an edge between two random nodes, as the probability. This is consistent with the generation of the initial network structure. Our implementation of this process can be found in Algorithm 3.

3.5 Quitting

In this model, quitting occurs in each time step, and we use quitting as a way to keep the number of people in each layer constant. On a given layer, each node has the same probability of quitting. However, due to the different types of interactions that happen on the layers, each layer can have a different quitting probability.

To find the quitting probability, we first make the key assumption that we want the number of nodes on each layer to be constant. Using the Law of Large Numbers, we know that the mean number of nodes that are promoted is $p_i \times n_i$, where p_i is some probability of promotion, n_i is the number of nodes on layer i, and i is the current layer. Let q_i be the quitting probability from layer i and k be the probability of being hired. We assume that the number of nodes entering the layer equals the number of nodes leaving the layer. Using these facts, we find that in the top layer, where nodes only leave by quitting, the equation for quitting is

$$\begin{aligned} p_{i-1}n_{i-1} &= q_i n_i \\ p_{i-1}\frac{n_{i-1}}{n_i} &= q_i. \end{aligned}$$

In the bottom layer, where hiring, quitting, and promotion from the layer occur, we find that the equation for quitting is

$$\begin{aligned} kn_0 &= q_0 n_0 + p_i n_0 \\ k - p_0 &= q_0. \end{aligned}$$

For the middle layers, where individuals are promoted into and from the layer and individuals quitting, we find that the equation for quitting is

$$\begin{aligned} p_{i-1}n_{i-1} &= q_i n_i + p_i n_i \\ \frac{p_{i-1}n_{i-1} - p_i n_i}{n_i} &= q_i. \end{aligned}$$

The number of nodes is a constant that we choose at the beginning of the simulation. Thus, to calculate the quitting rates, we just need to find the probability of promotion p_i and the hiring probability k.

In promotion, there are two main steps. The first is a node joining the applicant pool and the second is being promoted from the applicant pool. Using the definition of conditional probability, we find that

$$P(\text{Promoted}) = P(\text{Pool}) \times P(\text{Promoted} \mid \text{Pool}).$$

To find the probability of promotion into the applicant pool, we first consider the formula for promoting an individual node. We draw a sample using a probability of $\frac{h_i}{nbr_i}$, where h_i is the homophily of node i and nbr_i is the number of neighbors of i. The number of neighbors on average in a $G(n, p)$ network is $p_g \times n$, where p_g is the probability of two nodes having an edge between them in the graph and n is the number of nodes in the layer. On average, each node will have roughly $p_w \times nbr$ female neighbors and $(1 - p_w) \times nbr$ male neighbors.

On average, female nodes will account for p_w times the total number of nodes, while the other $(1 - p_w)$ will be male nodes. We can calculate a weighted average of the homophily scores using these weights, finding that

$$h_{avg} = p_w^2 \times pg \times n_i + (1 - p_w)^2 \times pg \times n_i.$$

Finally, to find the overall probability of going into the applicant pool, we divide the average homophily by the average number of neighbors:

$$\begin{aligned} P(\text{Pool}) &= \frac{h_{avg}}{nbrs} \\ &= \frac{p_w^2 \times pg \times n_i + (1 - p_w)^2 \times p_g \times n_i}{p_g \times n_i} \\ &= p_w^2 + (1 - p_w)^2 \end{aligned}$$

Next, we need to find the probability of being promoted given being in the pool, $P(\text{Promoted} \mid \text{Pool})$. We promote from the pool to the next layer with a constant probability of promotion, p_{promo}, and multiply that probability by a bias constant b if the node is a female node. Because we expect female nodes to account for p_w of the total node population, we can create a weighted average of the two gender's probability of promotion from the pool:

$$P(\text{Promoted} \mid \text{Pool}) = p_{\text{promo}} \times b \times p_w + p_{\text{promo}} \times (1 - p_w).$$

Finally, to find the overall probability of promotion, we multiply the equations for $P(\text{Promoted} \mid \text{Pool})$ and $P(\text{Pool})$, resulting in the equation

$$P(\text{Promoted}) = (p_{\text{promo}} \times b \times p_w + p_{\text{promo}} \times (1 - p_w))(p_w^2 + (1 - p_w)^2).$$

It is interesting to note that the probability $P(\text{Promoted})$ has no dependence on pg or n_i and only depends on the parameters p_w, p_{promo}, and b. The variables p_w, p_{promo}, and b are all characteristics of the overall model, not the structure of the network, while pg or n_i describe the initial structure of the graph. The formula for $P(\text{Promoted})$ suggests to us that on average, the promotion probability, and long-term behavior of the model, averaged over many trials, will not depend on the different graph structures. However, this result depends on the use of a $G(n, p)$ network as the initial network.

Once we have calculated $P(\text{Promoted})$, we can create probabilities for quitting on each layer. We use these probabilities by giving each node a random chance of quitting. We do this in a layer-by-layer step, going from the bottom up. The first thing we do is calculate the number of active nodes in the layer. This allows us to make sure that we always have at least one node in the layer. Then, for each node, we randomly pick if it will quit by drawing from a binomial distribution with the quitting probability for the layer. If the node is active and there are at least two active nodes in the layer left, then we set that node to be inactive, reset its homophily, and disconnect it from all edges, adjusting its neighbors' homophilies. This algorithm is found in Algorithm 4.

3.6 Shuffling

There are two ways that we reassign gender to nodes. The first is found in Algorithm 5, where we keep the same number of male and female nodes in each layer and permute the locations of the gender labels. We then update the homophily values for each node. Using this method of shuffling keeps the same starting fractions of women on each layer.

The other method of reassigning gender is done by randomly selecting gender again using the probability of being a female node, p_w. This process is found in Algorithm 6. For each node, we find a new gender using a weighted choice where being a female node has probability p_w and being a male node has probability $1 - p_w$. Then, for each node and layer combination, we reassign the homophily value. This method doesn't guarantee that the fractions of women are the same on each layer. Because we use probability to reassign gender, there could be large differences in the number of men and women on each layer between the original network and the shuffled network.

3.7 Modularity

The modularity of a partition of a network is a way to quantify the connectedness of two groups of nodes. In our case, we want to see how connected the male and female nodes are to each other. The formula for modularity is

$$Q = \frac{1}{4m} \sum_{ij} B_{ij} s_i s_j,$$

where $B_{ij} = A_{ij} - \frac{k_i k_j}{2m}$ and shows how connected a node is to members of its own group and k_i is the number of neighbors of node i. The entry of the adjacency matrix for nodes i and j, A_{ij}, is 1 if there is an edge between nodes i and j and 0 otherwise. If node i is in group one, s_i is 1, if i is in group two, $s_i = -1$, and m is the total number of edges in the graph. Algorithm 7 shows the calculation of the modularity of the partition between male and female nodes, and Section 2.1.3 explains what the modularity calculation means in more detail.

3.8 Putting It All Together

When running our model, there are six parameters that we can change: $npl, p_w, p_g, b, years$, and $trials$. We want to test how these different parameters affect the long-term behavior of our model. Table 3.1 shows what these parameters mean and their range of values

The first is the number of nodes on each layer, npl. This is a list of integers. Because we want to directly compare different runs with different parameters, we keep npl the same between all simulations to eliminate one source of variance.

The next variable we can alter is p_w, which is the probability a node is assigned female and can range in value from 0 to 1. We use p_w to assign gender to nodes during initialization, hiring, and reassigning gender to shuffle. We vary p_w greatly through the simulations.

Parameter	Meaning	Range of Values
npl	List of the number of nodes on each layer	list of $x_i \in \mathbb{Z}$
p_w	Probability a node is assigned female	$[0, 1]$
p_g	Probability of any two nodes sharing an edge	$[0, 1]$
b	Bias, 1 is unbiased, 0 is fully biased	$[0, 1]$
$years$	number of years the model should run	$\mathbb{Z}$
$trials$	number of trials the model should run	$\mathbb{Z}$

Table 3.1 Six parameters that influence the model.

The variable p_g is the probability of two nodes being connected in our network. This probability is used during initialization, promotion, and hiring.

The bias constant, b, ranges from values of 0 to 1. When $b = 0$, no women are promoted at all. We call this case full bias. When $b = 1$, women are promoted at the same rate as men, which we call running an unbiased simulation.

Finally, we can adjust the duration of our simulation by varying the number of years and trials.

When we run our simulation, we start by initializing a graph model starting with an initial graph for as many trials as specified. Due to computing constraints, this is typically 10 or 20 trials. We cycle through promotion, quitting, and hiring each year for a specified number of years. We record the fraction of women at every layer for each year and average it with the same data from all other trials, and similarly record data on the number of nodes in each layer for each year.

Chapter 4

Results of Unbiased Model

4.1 Variance

We want to understand the variance between trials and find the number of trials we would need to run to get results with low variance. A low variance value would be when it is less than 5% of the mean value. To test this, we calculate the variance of the fraction of women over time. An example run can be found in Figure 4.1. Note that the variance is rather large compared to the value of each point. This value could be lowered with an increased number of trials. Due to limitations on computing power, running more than twenty trials wasn't feasible for this project. Also, we note that the variance increases as the number of years goes on. This is due to the randomness of the model. As time goes on, more random choices are made in the promotion and quitting steps, which increases the variation in the state of the model. Thus, the general trend of increasing variance is consistent with our intuition of the model.

4.2 Number of Nodes

In Section 3.8, we calculated the probability of quitting to maintain the initial number of nodes in each layer in expectation. Here, we compare the analytical approximation to the direct numerical computations done by our model.

We found that the probability of being promoted is $P_p = (p_{promo} * b * p_w + p_{promo} * (1 - p_w))(p_w^2 + (1 - p_w)^2)$. The quitting rate is defined in terms of P_p and k, which the hiring rate. On the bottom layer, hiring, quitting, and

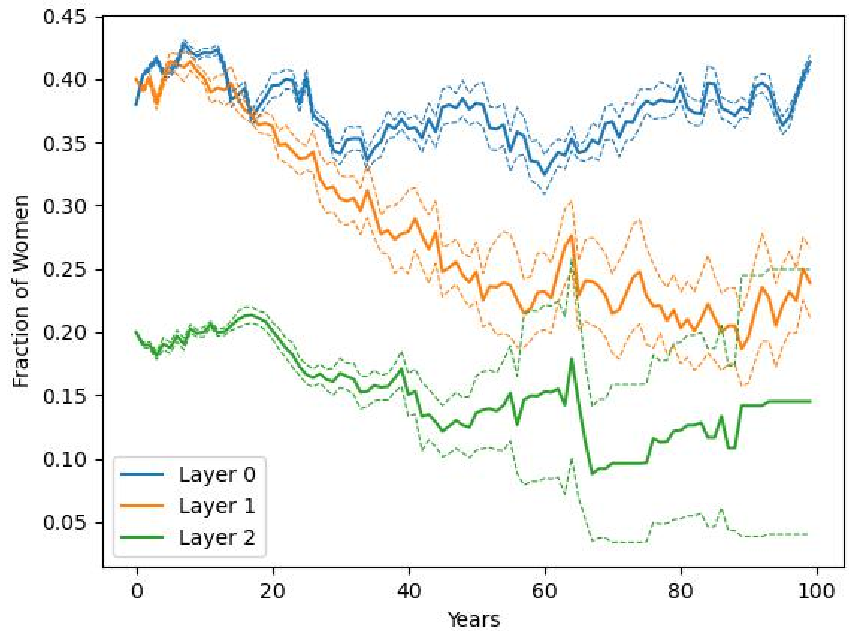

Figure 4.1 A simulation of the model run for 100 years plotting the average fraction of women over time. The dashed lines represent the average value ± the variance over the trials. We ran this simulation for twenty trials with a bias parameter of 1, the probability of being a woman 0.3, and the connection probability of 0.05. This network started with 50, 20, and 10 people in each layer of the graph respectively.

promotion processes are happening, so the quitting rate is

$$q_0 = k - P_p.$$

On the top layer, the only processes are promotion into the layer and quitting, so the quitting rate is calculated to be

$$q_i = P_p \frac{n_{i-1}}{n_i}.$$

On the middle layers, nodes are promoted into and from the current layer and nodes quit, so the quitting rate is

$$q_i = \frac{p_p n_{i-1} - p_p n_i}{n_i}.$$

If we condition $n_{i-1} \geq n_i$ for the number of nodes n_i on layer i, we won't have any issues with a negative probability of quitting for all but the bottom layer. For the bottom layer, we need to set k, the hiring constant, to be larger than P_p to ensure we have a nonnegative quitting constant. To do so, we set k to be

$$k = 0.1 + P_p.$$

Figure 4.2 shows the total number of nodes on average during one simulation of ten trials.

4.3 Sources of Difference

4.3.1 Different locations of genders

The most fundamental source of different behavior between the two simulations is the placement of female and male nodes within the same graph structure given the same number of nodes. For example, if we start with ten nodes on the first layer, five nodes on the second layer, and two nodes on the third layer, we will generate our initial multilayer network. This network has randomly generated connections and genders. As an example, we have three female and seven male nodes in the first layer, which are all randomly connected together. We then randomly shuffle the locations of the genders so that we still have three female and seven male nodes, but they are in different locations within the network. The edges all remain the same, but the homophily values are changed. This process is outlined in Algorithm 5 We do this process for each layer of the multilayer network. Shuffling the

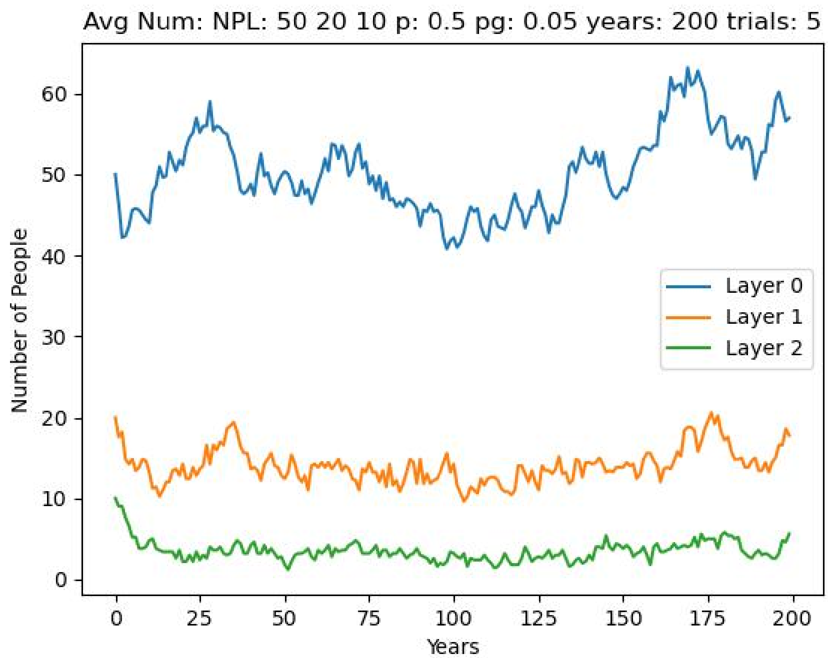

Figure 4.2 Graph of the mean number of nodes with time after the quitting probabilities were calculated. There were initially fifty nodes in layer 0, 20 nodes in layer 1, and 10 nodes in layer 2. We ran this simulation with a probability of being a woman of 0.5.

gender of the nodes is interesting because it reveals whether the placements of female and male nodes within the network play an impactful role.

In Figure 4.3, we run simulations on the same graph that has shuffled genders for probabilities a node is assigned female p_w of 0.5, 0.3, and 0.1. We note that after shuffling the initial locations of the nodes, the model does not produce the same exact graph as the original network. However, the qualitative behavior in each pair of shuffled networks is the same. In Figure 4.3a and Figure 4.3b, when $p = 0.5$, we note that all three levels of the organization fluctuate with respect to time around having half female and half male nodes regardless of the modularity. The simulation in Figure 4.3a has a more negative modularity, but the behavior between the two is qualitatively the same.

The qualitative behaviors of the two runs for the rest of the probabilities p_w are similar, with the steady-state fractions of women decreasing on each level as p_w decreases. Figures 4.3c, 4.3d, 4.3e, and 4.3f show that the level with the highest fraction of women is consistently the lowest level, followed by the middle layer, and then the highest layer has the lowest fraction of women. For any probability of being a female node less than 0.5, female nodes are underrepresented at the highest level of the organization. We also note that as the modularity of the initial networks of each simulation varies, the overall behavior of the model does not change.

These results suggest that the location of women in the model does not play a significant role in the long-term fractions of women in the corporation. However, they do show that the probability of being a female node does have a large impact on the behavior of the model. As the probability of being a female node increases, the fraction of women on each layer increases.

4.3.2 Different random assignment of genders

The next thing we want to test is how the probabilistic reassignment of genders affects the overall trends. To do this, similar to Section 4.3.1, we keep the graph structure the same for both runs of the simulation. However, instead of keeping the number of female and male nodes the same and changing their locations in the graph, we reassign gender to each node using the same probability of being a female node, p_w. This allows us to observe another source of difference between initial graphs. We keep the fraction of female nodes to be roughly the same but have some variance in these numbers. This test shows us how small changes in the number of women on each layer influence the outcome of the model.

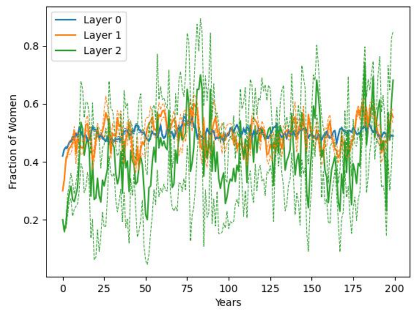

a. p_w = 0.5 and modularity of initial network of -0.12.

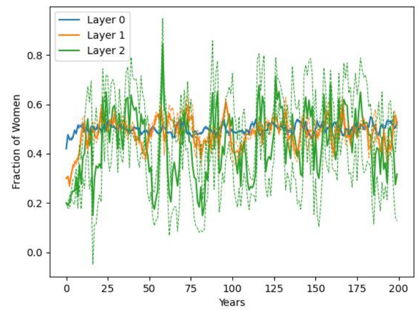

b. p_w = 0.5 and modularity of initial network of -0.11.

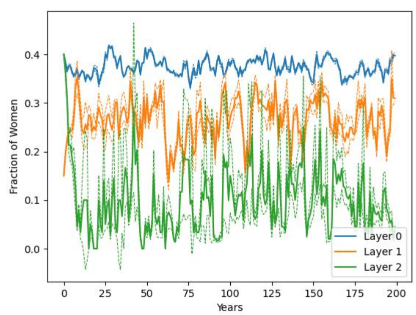

c. p_w = 0.3 and modularity of initial network of -0.13.

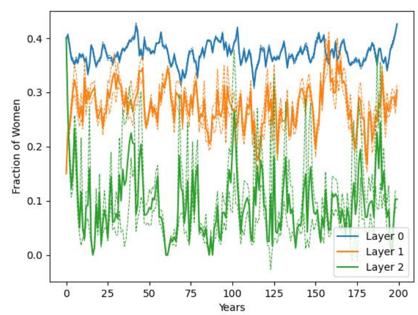

d. p_w = 0.3 and modularity of initial network of 0.01.

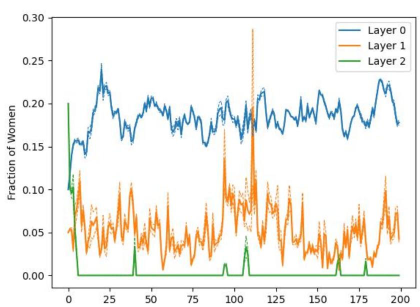

e. p_w = 0.1 and modularity of initial network of 0.2.

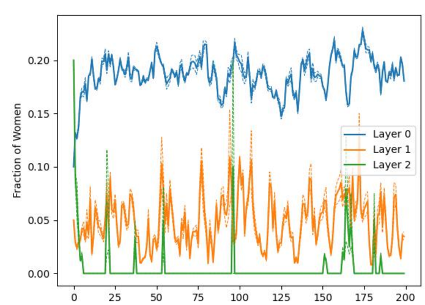

f. p_w = 0.1 and modularity of initial network of 0.28.

Figure 4.3 Three pairs of simulations run with the locations of the female and male nodes shuffled throughout each layer, varying p_w, the probability of being a female node. The dashed lines represent the average value ± the variance over the trials. We ran these with three layers, initially having 50, 20, and 10 people. The simulation was run for 20 trials, 200 years each, and had no bias. The probability of any two nodes in the network being adjacent, p_g, was 0.05.

Figure 4.4 shows the results of this test. We note that the higher the layer, the smaller the number of nodes in that layer. Due to the larger number of nodes, we see an increase in the variance of the fraction of women in higher layers of the network. Similar to the findings when we shuffled the genders, the quantitative behavior of the two simulations is not exactly the same, yet the qualitative behavior is the same. We note that when $p_w = 0.5$, we find that all of the layers quickly tend to have about half female and half male nodes. This is seen in Figures 4.4a and 4.4b. Although the exact fluctuations from steady state in the fractions of women may be different, the overall trends are the same.

In Figures 4.4c and 4.4d, when $p_w = 0.3$, we note the similar quantitative behavior between the two trials. Again note that the top layer has the most women, tending toward a proportion of 0.4 women, which is higher than the probability of women that are being hired. The middle layer hovers around a proportion of 0.3, while the fraction of women in the lowest layer rises above and falls below a value of about 0.1, seeming to stay close to that value.

Finally, we note that when $p_w = 0.1$, we have similar behavior between the two initial graphs with reassigned gender in Figure 4.4e and 4.4f. The bottom layer converges to a proportion of women of 0.2, which is higher than our p_w value of 0.1, while the middle layer oscillates around roughly 0.05. There are hardly any women in the top layer at all, and when one is hired, she quits in about one to five years.

These simulations suggest that the long-term behavior of this model isn't significantly affected by slight changes in the number and location of female and male nodes. We can only draw this conclusion in the context of the $G(n, p)$ network model that we use.

4.3.3 Different Initial Network Structures

The final source of change we want to analyze is the difference between different initial network structures. We keep the number of nodes the same in both of these networks but initialize two different networks which are two different realizations of the $G(n, p)$ network model. When we compare these results with those of the two previous sections, the results of this simulation give us some intuition on how the structure of the network impacts the overall outcome.

Figure 4.5 shows the results of running this test. In each subfigure, a different initial network was used, helping to mitigate the fear that we simply

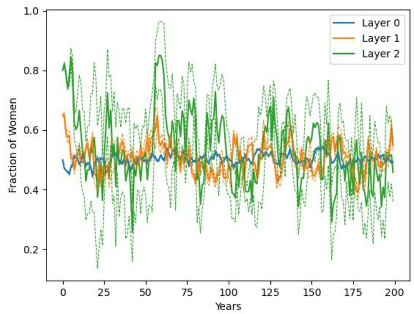

a. p_w = 0.5 and modularity of initial network of 0.03.

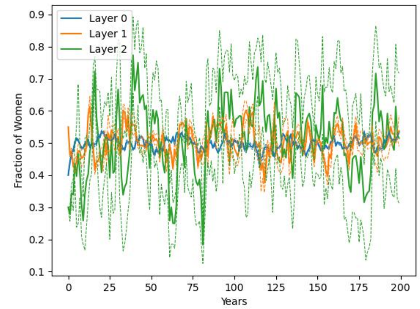

b. p_w = 0.5 and modularity of initial network of 0.03.

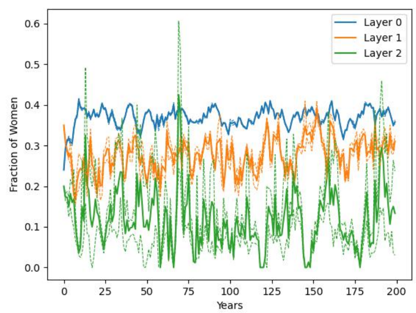

c. p_w = 0.3 and modularity of initial network of 0.04.

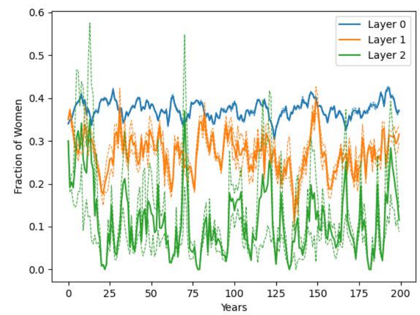

d. p_w = 0.3 and modularity of initial network of 0.09.

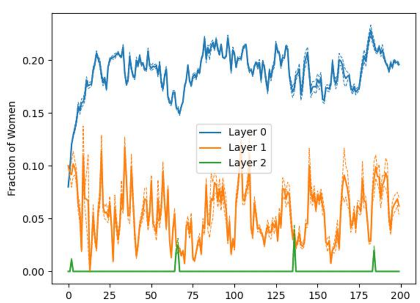

e. p_w = 0.1 and modularity of initial network of 0.24.

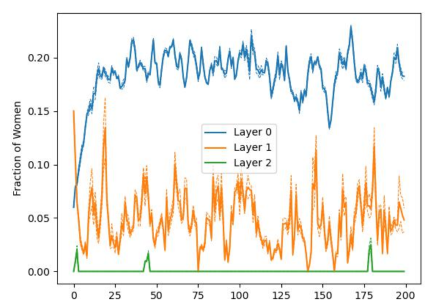

f. p_w = 0.1 and modularity of initial network of 0.16.

Figure 4.4 Three pairs of simulations run with the locations of the female and male nodes reassigned probabilistically throughout each layer, varying p_w, the probability of being a female node. The dashed lines represent the average value $\pm$ the variance over the trials. We ran these with three layers, initially having 50, 20, and 10 people. The simulation was run for 10 trials, 200 years each, and had no bias. The probability of any two nodes in the network being adjacent, p_g, was 0.05. These trials included calculated quitting probabilities.

chose two initial networks that yielded similar results. Regardless of the probability of being a female node, all pairs of graphs have qualitatively similar results. We observe the same trends as seen when we probabilistically reassigned gender in the previous two sections, as seen in Figure 4.3 and Figure 4.4.

These results suggest that the network structure may not have a significant impact on the long-term behavior of the model, and they are consistent with the analytical calculation done in Section 3.5 that found that the network structure should have no impact on average behavior. However, it is important to recognize that the model behavior could be highly sensitive to the type of network model used. Here, we use the $G(n, p)$ network model to create initial graphs and new connections while hiring. If we use a different type of network model, the influence of network structure may change significantly.

When observing the effects of different types of changes in the network structure, we don't see any qualitative differences regardless of the type of change that we make. However, we do see that the probability that a node is assigned female, p_w, has a large effect on the long-term behavior of the model.

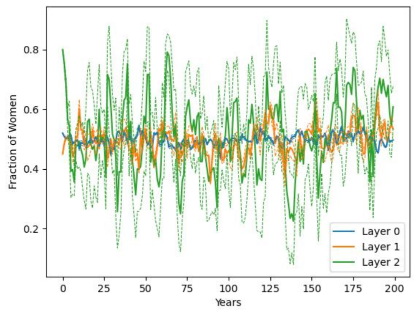

a. p_w = 0.5 and modularity of initial network of 0.0.

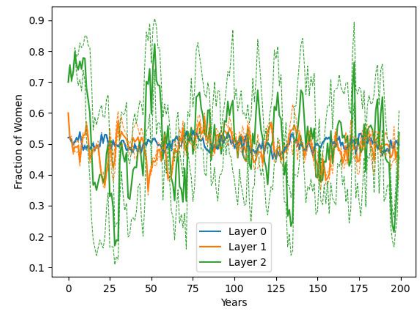

b. p_w = 0.5 and modularity of initial network of 0.15.

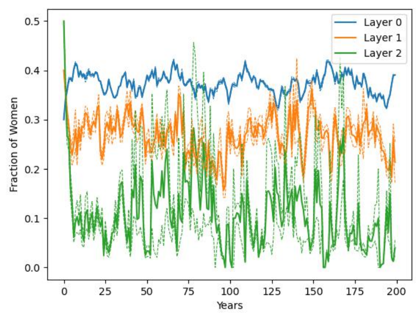

c. p_w = 0.3 and modularity of initial network of -0.01.

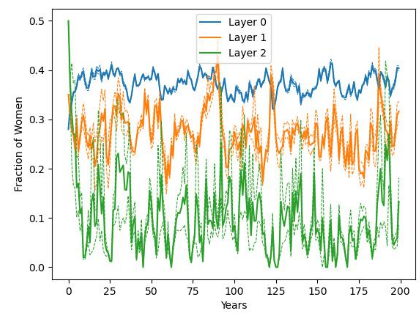

d. p_w = 0.3 and modularity of initial network of -0.13.

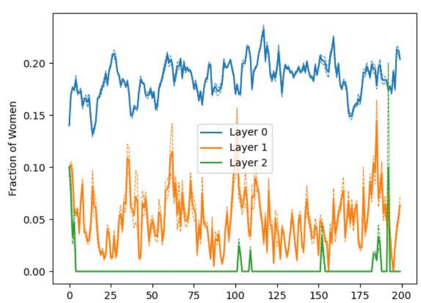

e. p_w = 0.1 and modularity of initial network of 0.17.

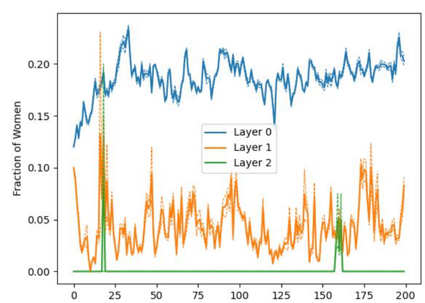

f. p_w = 0.1 and modularity of initial network of 0.15.

Figure 4.5 Three pairs of simulations run with different initial multilayer networks, varying p_w, the probability of being a female node. The dashed lines represent the average value ± the variance over the trials. These were all run with three layers, initially having 50, 20, and 10 people. We ran the simulation for 10 trials, 200 years each, and had no bias. The probability of any two nodes in the network being adjacent, p_g, was 0.05. These trials included calculated quitting probabilities.

Chapter 5

Results of Biased Model

5.1 Incorporating Bias

In this model, we incorporate bias by scaling the probability of promotion from the applicant pool to the next layer, but only for female nodes. To add bias in this way, we set the probability of being promoted is $P(\text{promoted}|\text{pool})$ for male nodes and $P(\text{Promoted} \mid \text{Pool}) * b$ for women. The bias constant ranges from 0 to 1, where $b = 0$ represents complete bias, where no female nodes will be promoted, and $b = 1$ represents no bias, and female and male nodes are equally likely to be promoted. We chose this form taking influence from Clifton et al. (2019), who incorporated bias as a constant term that affected the ratio of women to men promoted. We similarly wanted to affect this ratio. While incorporating it in a way that universally affects all women may not be the most accurate, it is a simple way to include an additional disadvantage that women may face.

We expect that having a bias of $b = 0$ will have no female nodes promoted. To validate that the model is working as intended, we include a graph of this case in Figure 5.1. We note that the fraction of women in layers 1 and 2 quickly drops to zero as women quit the organization and at no point during the entire simulation are women hired to replace them. We also see that in layer 0, the fraction of women is higher than 0.5. Because women are being hired in and quitting at the same rate as men, but men are being promoted, we are left with more women than men in the lowest layer of the organization. This result could show an instance of what is known as the "glass ceiling" effect, which describes an unofficial barrier to entering more advanced professional roles Sharma and Kaur (2019).

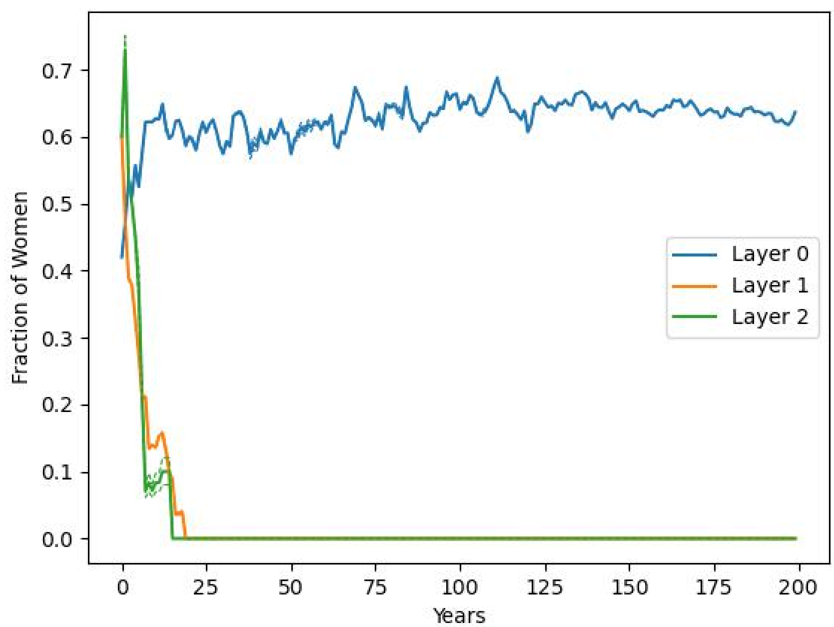

Figure 5.1 A simulation with complete bias against women. When having complete bias, $b = 0$ and no women are hired at all. This simulation was run with p_w = 0.5.

5.2 Simulation Results

To isolate the effect of bias on this model, we run multiple simulations on the same initial graph while varying the bias constant. In Figure 5.2, we see the results of a test where the probability of being a woman is 0.5. We notice that as the bias constant decreases, the fraction of women at the top layer of the organization decreases. This is consistent with our intuition that because women won't be hired as frequently and will quit at the same rate, we won't replace as many women as men and the long-term behavior of the top layer will have a fraction much lower than 0.5.

As the bias constant decreases, the bottom layer of the organization tends to have a larger fraction of women. This could be because women and men are being hired and quitting at the same rate, but men are promoted more often than women. This leaves more women than men in the bottom layer.

Interestingly, in all four of the simulations run with bias, the fraction of women in the middle layer seems to fluctuate around 0.5, seen in Figure 5.2. The steady state fraction of women in this layer may decrease slightly from 0.5 as the bias constant b decreases. This could be due to the increase in women on the lowest layer and the decrease in women on the top layer. Because there are more women in the bottom layer, the lower promotion rate may balance out with the fact that there are fewer men and a higher promotion rate.

Overall, we can conclude that the bias strongly has a strong relationship with the fraction of women on each layer. The fraction of women in the top layer drops by 0.2 as the bias goes from $b = 1$ to $b = 0.5$, while the fraction of women in the bottom layer increases by a smaller amount as the bias constant decreases.

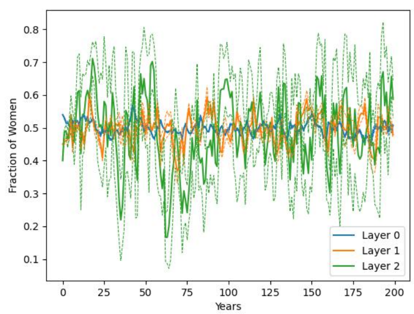

a. $b = 1$ and modularity of initial network of -0.12.

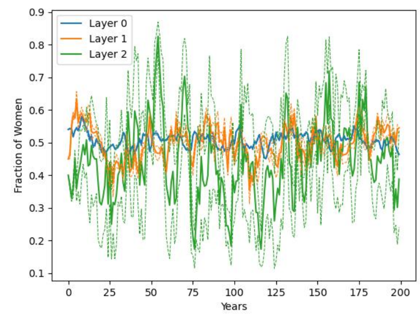

b. $b = 0.9$ and modularity of initial network of 0.06.

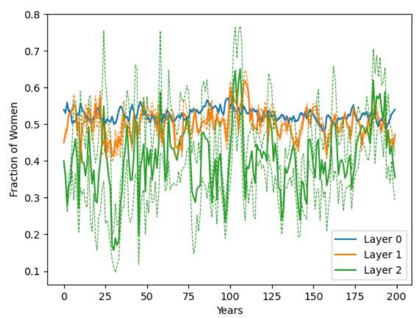

c. $b = 0.7$ and modularity of initial network of 0.06.

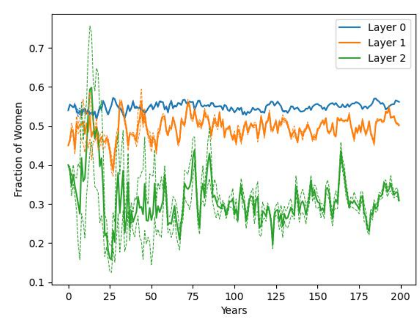

d. $b = 0.5$ and modularity of initial network of 0.06.

Figure 5.2 Four simulations run with probability $p_w = 0.5$ of being a woman, varying b, the bias constant. The simulations were run with three layers, initially having 50, 20, and 10 people. We ran each for 20 trials, 200 years each, and p_g, the probability of any two nodes in the network being adjacent, was 0.05.

Chapter 6

Conclusion

Gender bias and inequality continue to be significant issues facing society today. In this thesis, we had the goal of investigating some of the factors which influence gender balance. We used multilayer networks to create a model of professional structures and examined how our different factors influenced the overall behavior of the model. We found that the probability of a node being assigned female and bias had an impact on overall model behavior, but the initial network structure didn't influence our model at all when we used the $G(n, p)$ network model.

We used networks to encode individual progression through the professional structure. Progression itself was determined probabilistically by a few parameters. These parameters were p_w, the probability of being a female node, which was used to initialize node genders and the gender of hired nodes, and b, the bias constant which negatively impacted all female nodes throughout the organization. As the probability of being assigned to be a female node decreased, the fraction of women at the highest levels of the organization decreased. Middle layers oscillated around a constant fraction of women of roughly p_w, whereas the fraction of women on the bottom layer was slightly higher than p_w. These conclusions can be found in Table 6.1

Bias affects the model in a similar way. As the bias constant decreases, the fraction of women at the highest level of the organization decreases as well. The fraction of women at the middle layer is roughly the same, while the bottom layer has slightly more women than the proportion of women that were coming into the organization.

There are many possible future directions for this project. It would be particularly interesting to investigate how different initial networks affect long-term behavior. The lack of an effect on the dynamics from the network structure could be due to the choice of using a $G(n, p)$ network model. The $G(n, p)$ structure isn't the most realistic social model because all nodes have the same probability of adjacency with other nodes. Using other types of initial network structures that more closely model social relationships could reveal interesting dynamics, such as dependence on the initial network structure.

There are also many more simulations to run in addition to the comparisons done in this work. Simulating different policy changes, such as starting to hire more women into the organization, would yield interesting conclusions. For example, running the model by initially starting with some fraction of female nodes but hiring female nodes at a different rate may suggest that long-term behavior depends on the fraction of hired nodes that are female.

Parameter	Impact
p_w	As p_w decreases, fraction of women decreases
b	As b decreases, fraction of women decreases
$G(n, p)$ network model	No impact

Table 6.1 Parameters investigated and their impact on model dynamics.

Appendix A

Pseudocode

Algorithm 1 initGraph()

```
Inputs: a list of the number of nodes per layer npl, the probability of two
nodes being adjacent p_g, and the probability of being female p
totNodes = ∑ npl
Initialize list graphs
for n ∈ npl do
    G = G(n, p_g)
    Append G to graphs
end for
Initialize multilayer network M
for each layer l do
    Add nodes in graphs[l] to l
    for e ∈ graphs[l].edges do
        Add edge e to l
    end for
end for
Initialize list AD
for n ∈ totNodes do
    Initialize list nodeAtts
    lBound = 0
    uBound = npl[0]
    nGender = weighted.choice((p, 1 − p), (F, M))
    for each layer l do
        Initialize dictionary layerAtts
        nodeLayerAtts[gender] = nGender
        if not in first layer then
            lBound += npl[l − 1]
            uBound += npl[l]
        end if
        if lBound ≤ l < uBound then lnodeLayerAtts[active] = True
        elsenodeLayerAtts[active] = False
        end if
        Append nodeLayerAtts to nodeAtts
    end for
    Append nodeAtts to AD
end for
for n ∈ totNodes do
    for each layer l do
        for nbr ∈ n.neighbors do
            if AD[n][l].gender = AD[nbr][l].gender and l = nbr.layer
then
                AD[n][l].hom += 1
            end if
        end for
    end for
end for
Return AD and M
```

Algorithm 2 promote()

Inputs: multilayer network M, 2D list of attribute dictionaries AD, and probability of edge connection p_g

for each layer l except the last **do**
  Reset the list $seekPromo$ to empty
  for $n \in totNodes$ **do**
    if $AD[n][l].active = True$ and $\frac{AD[n][l].hom}{n.nbrs}$ **then**
      Append n to $seekPromo$
    end if
  end for
  for $k \in seekPromo$ **do**
    $pNum = binomial(1, bias(gen) * p_{promote})$
    if $pNum = 1$ **then**
      $AD[k][l].homo = 0$
      $AD[k][l].active = False$
      $AD[k][l].active = True$
      for $nbr \in k.neighbors$ **do**
        Disconnect n and nbr
      end for
      for $j \in totNodes$ **do**
        if $AD[j][l+1].active$ and $k \neq j$ and $binomial(1, p_g) = 1$ **then**
          Add edge between k and l
          Increment $AD[k][l].hom$ by 1
        end if
      end for
    end if
  end for
end for

Algorithm 3 hiring()

Inputs: multilayer network M, 2D list of attribute dictionaries AD, probability of being a woman p, and probability of edge connection p_g, hiring constant $hire$.

```
for each layer l do
    activeInLowest = 0
    for each node n do
        if AD[n][l].active then
            activeInLowest += 1
        end if
    end for
end for
numHires = max(P_hire * activeInLowest, 1)
for i in range(numHires) do
    Initialize list nodeAtts
    nGender = weighted.choice((p, 1 − p), (F, M)).
    for each layer l do
        Create a new dictionary newDict
        Add node i to multilayer Network M
        Assign gender, bias, and activity entries to newDict
        Initialize homophily as 0 in dictionary
        Append newDict to nodeAtts
    end for
end for
for new in range(nodes, totalNodes) do
    for i in range(new) do
        if AD[i][0] then
            if i ≠ new then Edge (i, 0)(new, 0) = binomial(1, p_g)
                if AD[i][0].gender == AD[i][0].gender then
                    AD[i][0].homo += 1
                    AD[i][0].homo += 1
                end if
            end if
        end if
    end for
end for
```

Algorithm 4 quit()

Inputs: multilayer network M, 2D list of attribute dictionaries AD, list of quitting probabilities $quitProbs$

```
for each layer l do
    numActive = 0
    for each node n do
        if AD[n][l].active then
            numActive += 1
        end if
    end for
    for each node n do
        quits = binomial(1, quitProbs[l])
        if quits == 1 and AD[n][l].active and numActive > 1 then
            numActive −= 1
            AD[n][l].active = False
            AD[n][l].homo = 0
            for nbr ∈ n.neighbors do
                Edge (n, l), (nbr, l) = 0
                if AD[n][l].gen == AD[nbr][l].gen then
                    AD[nbr][l].homo −= 1
                end if
            end for
        end if
    end for
end for
```

Algorithm 5 shuffle()

```
Inputs: multilayer network M, 2D list of attribute dictionaries AD, and a
list of the number of nodes on each layer npl.
Set numFem and numMale to lists of 0s with l entries
for each layer l do
    for each node n do
        if AD[n][l].active then
            if AD[n][l].gender ==FEM then numFem[l] += 1
            elsenumMale[l] += 1
            end if
        end if
    end for
end for
Set list genList to a list of n 0s
indShift = 0
for every level l do
    Create a list nodeLocs of elements 0 through npl(l)
    Randomly permute nodeLocs
    for i in range nodes do
        if i < numFem[l] then
            genArray[indShift + nodeLocs[i]] = FEM
        else
            genArray[indShift + nodeLocs[i]] = MALE
        end if
    end for
    indShift += npl[l]
end for
for every node n do
    for every layer l do
        if genArray[n] == FEM then
            AD[n][l].gender = FEM
        else
            AD[n][l].gender = MALE
        end if
    end for
end for
for every node n do
    for every layer l do
        numSameGen = 0
        for nbr ∈ n.neighbors do
            if AD[n][l].gender == AD[nbr][l].gender then
                numSameGen += 1
            end if
        end for
        AD[n][l].homo = numSameGen
    end for
end for
```

Algorithm 6 resetGenderProb()

Inputs: multilayer network M, 2D list of attribute dictionaries AD, and a list of the number of nodes on each layer npl.

```
for each node n do
    nGender = weighted.choice((p_w, 1 − p_w), (F, M))
    for every layer l do
        AD[n][l].gender = nGender
    end for
end for
for every node n do
    for every layer l do
        numSameGen = 0
        for nbr ∈ n.neighbors do
            if AD[n][l].gender == AD[nbr][l].gender then
                numSameGen += 1
            end if
        end for
        AD[n][l].homo = numSameGen
    end for
end for
```

Algorithm 7 modularity()

```
Inputs: multilayer network M, 2D list of attribute dictionaries AD.
totalDegree = 0
Q = 0
for each node n do
    for every layer l do
        totalDegree += (n, l).degree
    end for
end for
totalEdges = totalDegree/2
for every node m do
    for every node n do
        for every layer l do
            if AD[m][l].active and AD[n][l].active then
                k_m = len(m.neighbors)
                k_n = len(n.neighbors)
                if AD[m][l].gen == FEM then
                    s_m = 1
                else
                    s_m = −1
                end if
                if AD[n][l].gen == FEM then
                    s_n = 1
                else
                    s_n = −1
                end if
                B_{m,n} = edge(m, n) − (k_m k_n)/(2totalEdges)
                Q += B_{mn} s_m s_n
            end if
        end for
    end for
end for
Q = Q/(2totalEdges)
return Q
```

Bibliography

Association for Women in Mathematics. 2019. Data and research. URL https://awm-math.org/resources/academics/data-and-research/.

Bellotti, Elisa, Dominika Czerniawska, Martin G. Everett, and Luigi Guadalupi. 2022. Gender inequalities in research funding: Unequal network configurations, or unequal network returns? *Social Networks* 70:138–151.

Bernstein, Robert H., Michael W. Macy, Wendy M. Williams, Christopher J. Cameron, Sterling Chance Williams-Ceci, and Stephen J. Ceci. 2022. Assessing gender bias in particle physics and social science recommendations for academic jobs. *Social Sciences* 11(2):74.

Brooks, Heather Z. 2023. A tutorial on networks in social systems: A mathematical modeling perspective. 2302.00801.

Clifton, Sara M., Kaitlin Hill, Avinash J. Karamchandani, Eric A. Autry, Patrick McMahon, and Grace Sun. 2019. Mathematical model of gender bias and homophily in professional hierarchies. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 29(2):023,135.

Du, Yuhao, Jessica Nordell, and Kenneth Joseph. 2022. Insidious nonetheless: How small effects and hierarchical norms create and maintain gender disparities in organizations. *Socius* 8:23780231221117,888.

Grandjean, Martin. 2015. Social network analysis and visualization: Moreno's sociograms revisited. URL https://www.martingrandjean.ch/social-network-analysis-visualization-morenos-sociograms-revisited/.

Hagberg, Aric A., Daniel A. Schult, and Pieter J. Swart. 2008. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference*, eds. Gaël Varoquaux, Travis Vaught, and Jarrod Millman, 11–15. Pasadena, CA USA.

Heilman, Madeline E. 2001. Description and prescription: How gender stereotypes prevent women's ascent up the organizational ladder. *Journal of Social Issues* 57(4):657–674.

Jalali, Zeinab S., Krishnaram Kenthapadi, and Sucheta Soundarajan. 2021. On measuring the diversity of organizational networks. In *Complex Networks XII*, eds. Andreia Sofia Teixeira, Diogo Pacheco, Marcos Oliveira, Hugo Barbosa, Bruno Gonçalves, and Ronaldo Menezes, 59–72. Cham: Springer International Publishing.

James, Alex, and Ann Brower. 2022. Levers of change: using mathematical models to compare gender equity interventions in universities. *Royal Society Open Science* 9(9):220,785.

Kivelä, Mikko. 2013. Multilayer networks library. URL http://www.mkivela.com/pymnet/.

Kivelä, Mikko, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. 2014. Multilayer networks. *Journal of Complex Networks* 2(3):203–271.

Lee, Eun, Aaron Clauset, and Daniel B. Larremore. 2021. The dynamics of faculty hiring networks. *EPJ Data Science* 10(1):48.

Lerman, Kristina, Yulin Yu, Fred Morstatter, and Jay Pujara. 2022. Gendered citation patterns among the scientific elite. *Proceedings of the National Academy of Sciences* 119(40):e2206,070.

Nordell, Jessica. 2021. This is how everyday sexism could stop you from getting that promotion. *The New York Times* URL https://www.nytimes.com/interactive/2021/10/14/opinion/gender-bias.html.

Reference, Oxford. 2023. Gender bias. URL https://www.oxfordreference.com/display/10.1093/oi/authority.20110803095846581.

Roy, Marie-Françoise, Colette Guillopé, Mark Cesa, Rachel Ivie, Susan White, Helena Mihaljevic, Lucía Santamaría, Regina Kelly, Merrilyn Goos, Silvina Ponce Dawson, Igle Gledhill, and Mei-Hung Chiu. 2020. *A Global Approach to the Gender Gap in Mathematical, Computing, and Natural Sciences: How to Measure It, How to Reduce It?* Zenodo. doi:10.5281/zenodo.3882609. URL https://doi.org/10.5281/zenodo.3882609.

Schwartz, Leah P., Jean Liénard, and Stephen V. David. 2021. Impact of gender on the formation and outcome of mentoring relationships in academic research. *arXiv preprint arXiv:210407780* .

Sharma, Sakshi, and Rajvir Kaur. 2019. Glass ceiling for women and work engagement: The moderating effect of marital status. *FIIB Business Review* 8(2):132–146.

United Nations. 2022. Closing gender pay gaps is more important than ever. URL https://news.un.org/en/story/2022/09/1126901.

Women Business Collaborative. 2022. Executive summary and forward. URL https://www.wbcollaborative.org/women-ceo-report/the-report/2022-executive-summary/.