# Population-Based Simulation of Gender Inequality Issues

John A. Bullinaria

School of Computer Science, University of Birmingham
Edgbaston, Birmingham, B15 2TT, UK
j.a.bullinaria@cs.bham.ac.uk

**Abstract**

A population-based simulation framework is presented that allows a principled approach for exploring gender inequalities in professional hierarchies such as universities or businesses, and how they might emerge, evolve and be rectified. Results from a representative range of cases involving gender-based discrimination and intrinsic gender-based ability differences are presented to demonstrate the power of the approach. Such artificial life simulations will hopefully inspire and facilitate better approaches for dealing with these issues in real life.

## Introduction

There has been much discussion in recent years about gender imbalance in certain professions, such as university computer science departments (e.g., Camp, 1997; Altonji and Blank, 1999; Handelsman et al., 2005; Moss-Racusin et al., 2012), and how one might go about rectifying such situations, for example by better advertising or positive discrimination. However, it is often difficult to identify the best solutions when it is not clear what the main causes of the imbalance are (Halpern et al., 2007), and applying solutions based on incorrect assumptions could easily make matters worse.

One obvious potential cause of imbalance is simple discrimination against a particular gender (e.g., Davison and Burke, 2000; Moss-Racusin et al., 2012), and if that cannot be prevented, some form of positive discrimination might be an appropriate remedy. Another possible cause is that one gender might have evolved to be intrinsically less able (either on average or in the tails of the distribution) in a particular area (e.g., Geary, 1998; Browne, 2002; Baron-Cohen, 2004; Halpern et al., 2007; Halpern, 2012), and that results in less success in that area, and hence a tendency for that gender to avoid entering related professions in future. It is not obvious what interventions here would be most beneficial, or whether any intervention at all would be a good strategy. Another possibility is that, despite having intrinsically equal ability in the chosen area, one gender is disadvantaged by other factors, such as delays in career progression caused by child rearing and maternity leave (e.g., Ceci and Williams, 2011), and these cases may require different forms of intervention.

The idea of using computer simulations to model such situations and explore the best strategies for intervention in complex processes such as these is not new (e.g., Martell, Lane and Emrich, 1996; Robison-Cox, Martell and Emrich, 2007; Helbing, 2010), but what might not be so widely appreciated is that population-based simulations with ability-based selection of the type commonly used in computational intelligence (e.g., Engelbrecht, 2007) and artificial life (e.g., Bullinaria, 2009, 2010) can be effective for exploring the key causes, effects and solutions here. They can also model the evolution of such factors by natural selection. Moreover, the known methodological pitfalls that commonly arise with agent-based approaches to social and economic simulation (Richiardi, Leombruni, Saam and Sonnessa, 2006) are well understood in the field of artificial life and can thereby more easily be avoided. This paper presents a general framework for performing such simulations, and provides a selection of results that illustrate the power of this approach.

The remainder of this paper is organized as follows: The next section describes the proposed simulation framework and its associated simplifications and assumptions. Then results from some preparatory baseline simulations are presented to establish appropriate values for the various free parameters. The next two sections show how those results differ in the cases of gender-based ability differences and discrimination. Finally, the effect of interventions, and how the evolution of individual preferences affect the results, are explored. The paper ends with some conclusions and discussion.

## Simulation Framework

This study begins by setting a few basic principles, and then explores what is possible within that general framework. The idea is to have an evolving population of individuals, with a range of intrinsic (innate) abilities, who can progress during their lifetimes to improve their position within their chosen professions. To draw reliable conclusions, the simulations need to be kept as clear and unbiased as possible (Bullinaria, 2009, 2010). Therefore, for the purposes of this initial study, a number of simplifying assumptions are made that help avoid any unnecessary confounding factors and also reduce the computational costs of the simulations to feasible levels:

1. There are two distinct genders, which are chosen randomly at birth with equal probability, and overall are equally able.
2. The distributions of innate individual abilities are the result of the evolutionary past, but are fixed for the duration of each simulation.
3. There are two distinct professions, which overall are equally valuable.

4. The initial individual abilities for each profession are determined randomly at birth and follow a normal (Gaussian) distribution. The means and/or the standard deviations of those distributions may depend on gender.
5. If one gender has higher mean ability in one profession, the other gender will have an equally higher mean ability in the other. The effect of the magnitude of such differences is one of the key factors to be explored.
6. Each individual can choose their profession randomly, or according to their abilities, or could have an intrinsic gender-based preference (i.e., probability) for choosing one profession over the other. Such preferences might emerge during the course of the simulations.
7. Individuals grow older, potentially improve their abilities through experience in their profession, and eventually retire and leave the working population.
8. Professional development involves a series of stages, and promotion between them is (by default) determined purely according to the best abilities currently available at each level (Rosenbaum, 1979). Discrimination or intervention in that process are other key factors to be explored.
9. If an individual does not get promoted within a set number of simulated years, they are likely to give up and leave the working population. It is also possible that varying percentages of individuals leave the working population for other reasons. Such details will need to be explored.
10. Individuals leaving the population are replaced by new individuals, and profession preferences may be based on the more successful individuals of previous generations.

There clearly remains much scope for variations within this general framework, and what emerges will depend on the relative magnitudes of the various parameters involved. There is also scope for variations designed to investigate the consequences of these initial simplifications.

The simulations follow common Artificial Life procedures. For each new individual in the population, a record is created and initialized with their innate gender, intrinsic abilities for the two professions, and any preferences for the professions. Thereafter it will be regularly updated with their age, chosen profession, stage in their profession, and number of years since reaching that stage. After updating for a number of simulated years, the population averages will settle down into a steady state, and the relevant results can be computed.

In principle, the above general framework can be used to simulate "professions" in any species. For example, food provision versus offspring protection in wild dogs. However, this paper will concentrate on abstract human professions, and therefore adopt human-like lifetimes and other parameters. It will assume, for simplicity, that all individuals enter their chosen profession at age 20 and retire at age 70, and that each profession has 7 stages, so 6 promotions are required to reach the top stage. Again for simplicity, it will be assumed that there is just one employer for each profession, so there is no need to simulate transfers between employers, or sub-groups of eligible individuals being considered at each promotion stage, as that is already known to bias the results (Lyness and Judiesch, 1999). A population size of 10,000 provides a sufficient number of individuals per profession per stage per gender for a reasonable level of competition at each stage, even when the distributions become skewed. The ability scale is measured in arbitrary units, and that will be set (without loss of generality) by taking the standard deviation of the initial Gaussian distributions to be 1.0, and measuring all other ability differences relative to that.

A workable grain size for the simulations is one round of updates per simulated year, and 10,000 simulated years is plenty for all populations to settle down into a stable final state. Updating the individual ages, applying any ability increments, and replacing retired and removed individuals is straightforward. Dealing with the promotions between the stages of each profession requires further specification. One approach is to maintain pre-chosen numbers at each stage to correspond to typical companies (e.g., Robison-Cox, Martell and Emrich, 2007). An alternative approach, adopted here, is to promote a fixed fraction $x$ of eligible individuals at each stage in each profession each year. Varying the promotion criteria and $x$, and requiring a certain number of years at a given stage before becoming eligible, are factors that will need to be explored empirically. Finally, an important aspect of this study is to incorporate into the standard setup a whole range of parameterized ability differences, discriminations and interventions that might be considered relevant.

The output of each simulation will usually be the final population of individuals, each with a gender, age, preference for profession, profession, ability in their chosen profession, and profession stage. Typically, the main factors of interest will be the various differences in population means between genders, such as how the profession preferences and numbers at each profession stage depend on gender (e.g., Robison-Cox, Martell and Emrich, 2007). Sometimes the evolution of the key parameter values throughout the simulation will also be of interest. To obtain reliable results, means and standard deviations over thirty runs of each simulation are computed, and unpaired *t tests* are used to determine the statistical significances of any differences found.

## Baseline Simulation Results

The approach adopted is to first present the results from the simplest possible simulation set-up, and then systematically investigate how the potential variations affect those baseline results. Such a sequential approach also facilitates the setting of the various parameter values at each stage.

The baseline case simply has the most able individuals promoted at each stage, with an equal promotion fraction $x$ at each stage varied from 0.01 to 0.05. The resulting numbers at each profession stage are shown in Figure 1. For very small $x$ values (0.01), no individuals reach the upper stages. For high values (0.04 and above), more are in the highest stage than in some of the lower stages. Values around 0.02 to 0.03 probably come closest to realistic situations in academia or industry where there is a pyramid structure with fewer individuals as one moves up the hierarchy (Rosenbaum, 1979). Having different promotion fractions $x_s$ for each stage $s$ might be required to model realistic scenarios, and that can easily be done, but, for simplicity, the following will continue with a single value $x$ across all stages.

The first variation requires individuals to wait at each stage for a certain number of years $w$ before they become eligible for promotion. Now, a higher proportion of eligible

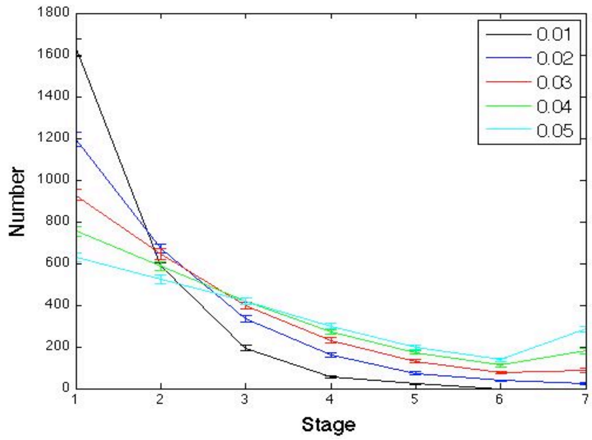


Figure 1: The initial baseline results showing the number of individuals at each profession stage and how that varies across a range of different promotion fractions *x*.

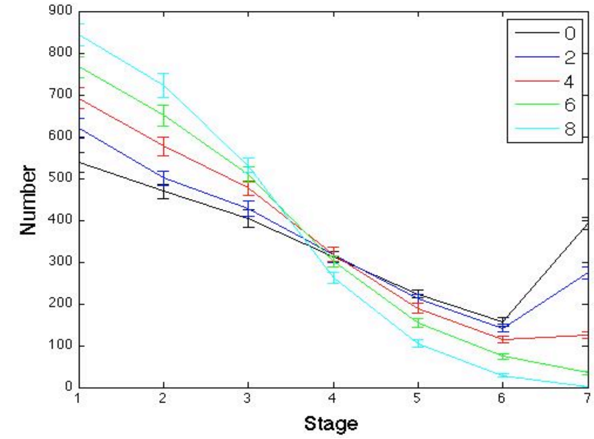


Figure 2: Number of individuals at each stage for promotion fractions $x = 0.06$ for different numbers of years *w* required at each stage before becoming eligible for promotion.

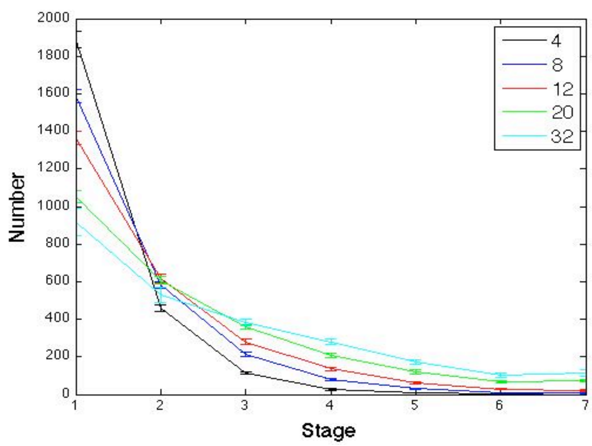


Figure 3: Number of individuals at each stage for promotion fractions $x = 0.06$ and wait $w = 4$ for different numbers of years *g* without a promotion before giving up and leaving.

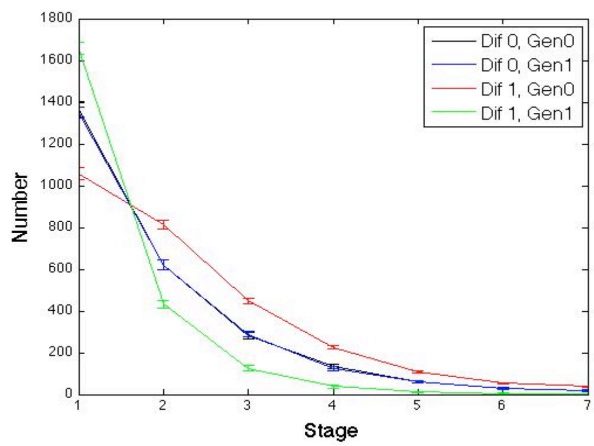


Figure 4: Number of individuals at each stage for $x = 0.06$, $w = 4$ and $g = 12$ for each gender (Gen0, Gen1) with no ability difference (Dif 0) and one std. dev. difference (Dif 1).

individuals need to be promoted each year to fill the higher stages. For a promotion fraction of 0.06, the effect of varying the required number of years *w* from 0 to 8 is shown in Figure 2. For 0 years, the promotion fraction is too high, as seen in Figure 1. Waits *w* of around 6 years lead to a reasonable distribution of individuals across the stages.

The next variation explores the effect of giving up and leaving the profession if promotion is not achieved within a certain number of years *g* after becoming eligible. Now, a slightly shorter wait *w* is needed so sufficient numbers are eligible for promotion at each stage. For a promotion fraction of 0.06 and a 4 years wait for eligibility, the effect of varying the number of years *g* before giving up from 4 to 32 is shown in Figure 3. For $g = 32$ years, there is little difference from never giving up. For fewer years, the numbers at later stages fall more sharply, and since the total number of individuals is fixed, there are more at the initial stage.

For situations that have learning or experience increase the individuals' abilities in line with the number of years in their chosen profession, or at each stage in that profession, it will be interesting to investigate the different age distributions that emerge for each gender at each stage, and to go on to explore the effect of factors such as maternity leave.

## Exploring Gender-based Ability Differences

Having seen how the three promotion parameters (*x*, *w*, *g*) affect the distribution across stages, explorations of the effect of gender differences can begin. The baseline simulations suggest that a promotion fraction of 0.06, a 4 years wait for eligibility, and giving up after 12 years, provides a reasonably realistic basis for the forthcoming simulations. Varying those values by small factors will inevitably change the results, but is unlikely to affect the general emergent patterns.

To begin, separate distributions for the two genders (Gen0 and Gen1) can be plotted for the case when each individual randomly chooses one of the two professions. Figure 4 shows the results for one particular profession when the gender difference is zero as above (Dif 0), and when the mean

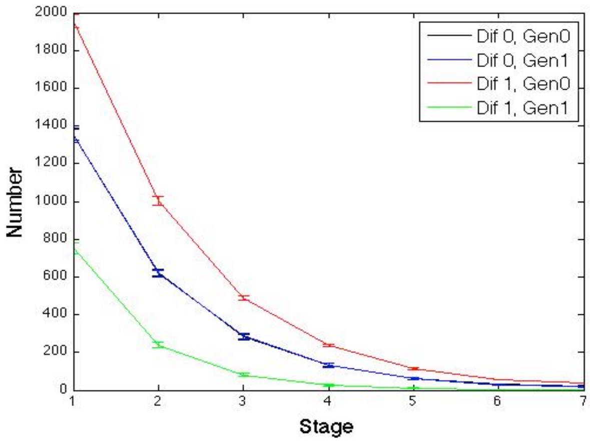


Figure 5: Number of individuals at each stage for same set-up as Figure 4, but with each individual pursuing the profession they are best at, rather than choosing one at random.

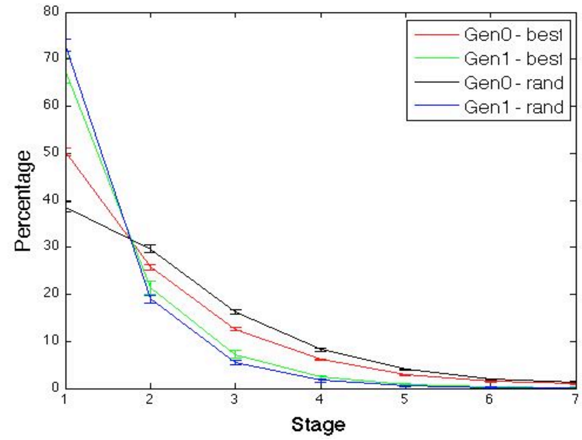


Figure 6: The Dif 1 results of Figures 4 and 5 as percentages of the whole population of each gender, for random choice of profession (rand) and choice of best profession (best).

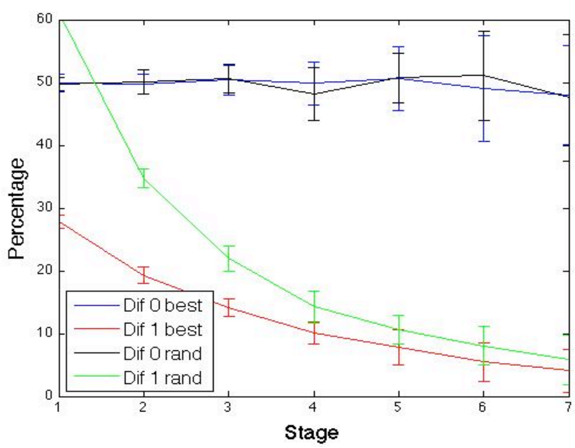


Figure 7: The percentage of Gen1 at each stage, with and without ability differences (Dif 1, Dif 0), for random choice of profession (rand) and choice of best profession (best).

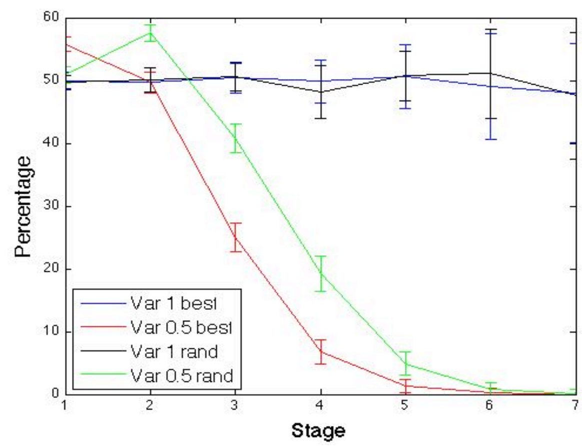


Figure 8: The percentage of Gen1 at each stage, for equal (Var 1) and half (Var 0.5) ability variance, for random choice of profession (rand) and choice of best profession (best).

abilities differ by one standard deviation (Dif 1). Obviously, there is no significant gender difference in the zero difference case. For a unit difference, the more able gender (Gen0) for the given profession has more individuals at the higher stages, and fewer stuck at the initial stage. The effect of dropouts also affects the total number of each gender participating in the profession. For the Dif 1 case there are 2746 (std. dev. 39) of Gen0 and only 2276 (std. dev. 40) of Gen1, which is a significant difference (*t test*, $p < 0.01$). For Dif 0, there are 2511 (std. dev. 54) of Gen0 and 2500 (std. dev. 46) of Gen1, with no significant difference (*t test*, $p > 0.05$).

If, rather than choosing their profession randomly, each individual were to choose the profession for which they have the best ability, the outcome is rather different as shown in Figure 5. Again there is no significant gender difference for the Dif 0 case, but for Dif 1 there is a massive statistically significant (*t test*, $p < 0.01$) reduction in the number of Gen1 individuals choosing the profession, 1114 (std. dev. 44) compared to 3893 (std. dev. 52) for Gen0. Figure 6 shows the Dif 1 results for random choice of profession (rand) and choice of best profession (best), as percentages of the whole population of each gender at each stage in the profession. The ability-based choice of profession brings the gender distributions a little closer together, but Gen1 still has much reduced numbers at the higher stages compared to Gen0.

Another way of representing the data is as the percentages of each gender at each stage. Since these always total 100%, it is sufficient to present the results only for Gen1. These are shown in Figure 7 for best and random profession choice. For no gender difference (Dif 0), the percentage of Gen1 at each stage is not significantly different to 50%. For Dif 1 and random profession choice, the proportion at stage 1 is slightly over 50% (due to weaker individuals waiting for a promotion that never comes) and then falls for later stages. When the best profession is chosen, there is a lower starting point, and a slower fall off at later stages. A gender-based ability difference leads to stage distribution percentage differentials even when self selection leads to reduced participation of the less able gender. This kind of pattern, known as a shrinking pipeline, is found in real populations, (Camp, 1997), though not necessarily for the same reason.

Interestingly, similar shrinking pipelines can arise even when there is no difference between genders in their mean abilities. If the variance in abilities for Gen1 is less than that of Gen0, as apparently happens with some human skills (e.g., Humphreys, 1988), that can give Gen1 a disadvantage at later stages of promotion, even if the means are the same for each gender. Figure 8 shows the effect of a factor of two in ability variances for random and best choices of profession. Similar patterns also arise when there are combinations of mean and variance differences. It is clear that there are many possible types of gender differences in ability that can account for the unequal gender distributions observed in real professions. From the simulation point of view, one can add further realism by replacing the simple Gaussian distributions used here with something more appropriate, but determining what those distributions should be might not be so easy (e.g., Benbow, 1988). Of course, the observed differences may also occur when there are no ability differences at all, and that is what will be investigated next.

## Exploring Gender-based Discrimination

Perhaps the question of most practical importance is: how do the above patterns of gender differences vary when, rather than any intrinsic ability difference, there is discrimination against a particular gender? Given the range of possibilities, it is not feasible to study all types of discrimination here, nor the potential reasons for them. However, to demonstrate the power of the simulation framework, it is sufficient consider a simple abstract case. In particular, suppose an individual of one gender had to be vastly superior to a rival of the other gender before being promoted before them. That might, for example, arise due to different perceived prior probabilities of the abilities for the two genders (that are not necessarily correct) being used in conjunction with the actual evidence submitted with the promotion application. It could also be indirect, rather than direct, discrimination, for example due to one gender being less likely to be awarded prestigious invited talks or prizes (e.g., Gürer and Camp, 2001), or due to the promotion criteria being skewed in favour of one gender (e.g., Schneider, 1998; Ginther and Hayes, 2003; Mixon and Trevino, 2005). To be specific, the above simulations were re-run with a Gen1 individual only being promoted in preference to a Gen0 individual in the profession of interest when their ability is at least one unit higher. The symmetry was maintained by having discrimination in the opposite direction for the other profession. That leads to exactly the same stage distribution as in Figure 4 for the case when the Gen0 ability distribution really was one unit lower.

What is different between the discrimination and ability difference cases is the average abilities at each stage A similar pattern emerges for both random profession choice and ability-based choice, though choosing according to ability not surprisingly leads to better ability levels throughout. In the ability difference case (Dif 1), the ability of the weaker Gen1 is lower than Gen0 at the entry stage 1, but the ability-based promotions lead to much closer ability levels at later stages. In the discrimination case (Disc 1), the Gen0 abilities are reduced to the same degree as the Dif 1 case, because there is less competition for promotion, but the Gen1 abilities

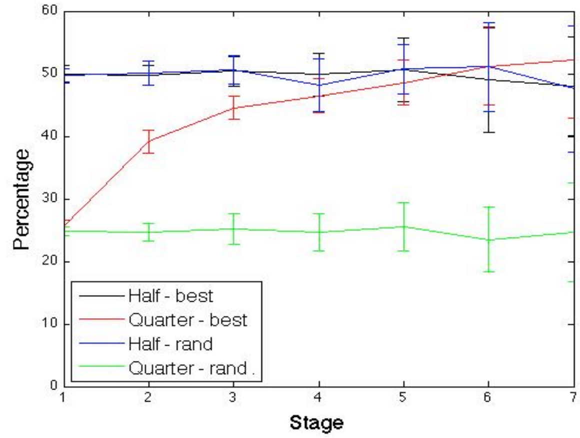


Figure 9: Percentage of Gen1 at each stage, with no ability or discrimination differences, starting at half or quarter total, for profession chosen randomly (rand) and by ability (best).

are much higher, due to the extra ability required to achieve promotion. Similar, though smaller (half a std. dev.) gender differences in ability have been observed in real corporations, suggesting that the presence of gender discrimination there (Lyness and Heilman, 2006). That, of course, does not mean there is necessarily a gender discrimination based glass ceiling in all cases, but there is certainly evidence consistent with that existing elsewhere too (e.g., Sabatier, 2010).

It is often suggested that indirect forms of discrimination are discouraging young women from entering particular professions, such as computer science, in the first place (e.g., Gürer and Camp, 2001). This effect can be modelled too, with the above simulations run in the same way whatever factor is reducing the numbers entering the given profession. The simplest possible case has no other promotion-based discrimination and no intrinsic ability differences, and the results in Figure 9 show for two rather different starting fractions that, as long as the profession is chosen randomly, those fractions persists throughout the stages. However, if the individuals that choose the profession are doing so according to their best abilities, they are likely to be at the higher end of the ability distribution, and fair promotions will allow them to rise quickly through the stages so that the gender proportions become equalized at the highest stages.

Figures 7, 8 and 9 demonstrate how rather different pipeline patterns emerge depending on the situation simulated. These are the "pure" cases. In practice, there is likely to be more than one form of ability or discrimination difference present, and untangling the various factors will be a challenge. This is where the simulation approach proposed here will prove most useful, as it enables all the possible combinations and variations to be simulated relatively easily and reliably, with the inevitable interactions accommodated automatically.

## Exploring Intervention Policies

The shrinking pipelines and gender differences in the numbers entering some professions are often argued to be important issues that need addressing. For example, the lack of women in certain higher stages of academia might discourage women from studying those subjects and that may lead to critical

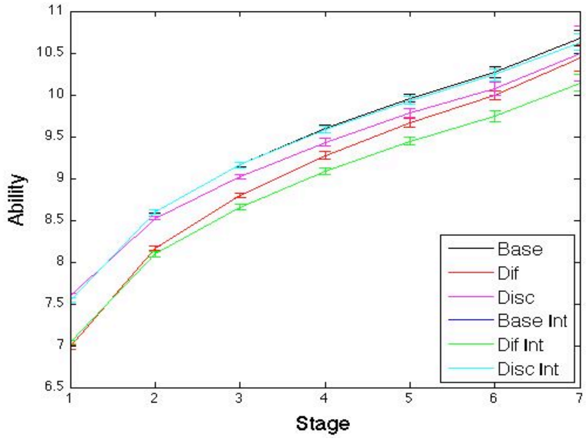


Figure 10: The average overall abilities at each stage for the standard and intervention populations for the three situations. The Base, Base Int and Disc Int results are identical.

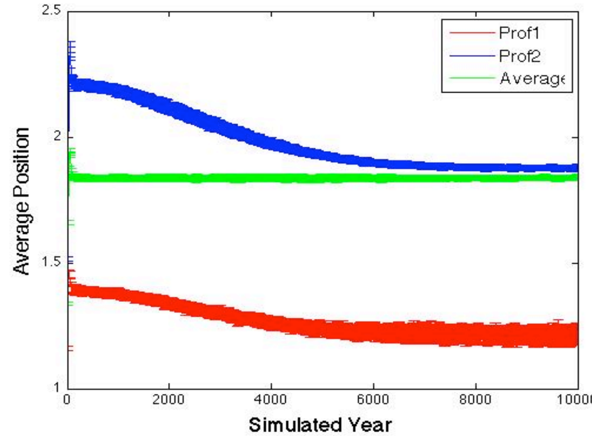


Figure 11: Average positions (profession stages) achieved by one particular gender while strong profession preferences emerge. The speed of change is parameter dependent.

skilled-worker shortages some areas (e.g., Camp, 1997).

The classic intervention would be to simply make sure that the numbers of each gender at each stage of each profession are as equal as possible. That can be implemented easily by ranking the eligible individuals of each gender separately, and promoting equal numbers of each gender to the next stage to give the required total number of promotions overall. The consequence of doing that in the above simulation framework does then lead to no significant number differences between the baseline, gender-based ability difference, and gender-based discrimination cases.

This intervention results in the expected differences in the corresponding abilities. All groups have the same average ability at each stage, except the less able Dif 1 Gen1 case which is one unit below at all stages, because equal numbers of promotions are taking place despite the lower ability levels of that gender. That leads to the important practical question: what is the average ability of the individuals at each stage, irrespective of their gender. That is shown in Figure 10. The baseline (Base) and baseline with intervention (Base Int) results are identical, since there is no gender imbalance for the intervention to correct, and these exhibit the best average abilities overall. The discrimination (Disc) case is slightly worse, since it unfairly allows weaker Gen0 individuals into the upper stages, rather than more able Gen1 individuals. The discrimination with intervention (Disc Int) case is no different to the base case, since the intervention successfully corrects the discrimination-based imbalance, and once again allows the best individuals at each stage to be promoted. The innate ability difference (Dif) case is overall worse than the base and discrimination cases, because that corresponds to Gen1 individuals having lower abilities than the base case, and that inevitably brings the population averages down. Obviously, if the gender difference corresponded to improved abilities for Gen0 over the baseline, rather than reduced abilities for Gen1, that would lead to improved population averages over the base case. The important question is: what will the consequences of intervention be in this case? As Figure 10 clearly shows, this makes the overall population performance (Dif Int) worse, particularly at the higher stages, since it forces the promotion of weaker Gen1 individuals over better Gen0 individuals. This highlights the importance of understanding the problem before trying to correct it.

## Evolving Preferences for Professions

So far, the simulations have been run for many generations to allow enough time for the various population distributions to stabilize, but none of the innate properties or preferences have been allowed to evolve or change from one generation to the next. However, the steady-state evolutionary computation approach underlying the general framework proposed here can automatically allow any inherent parameters (such as gender-based abilities, or preferences for particular professions) to evolve by natural selection if required (Bullinaria, 2009). In some species, such factors may be encoded genetically, but for current human professions it is more likely that such information will be passed on mimetically, in the form of social learning or mimicry (Bullinaria, 2010). For example, low take-ups of particular professions could emerge as a sensible reaction to poor progress in that profession by members of their gender in previous generations. There are clearly many such factors that could usefully be explored in the proposed simulation framework, and many ways they could be implemented, but one simple example should suffice to demonstrate the power of the simulation approach.

Suppose individuals choose their professions stochastically, rather than according to their abilities, but with particular intrinsic probabilities. The initial population would have equal preferences for the two professions, but individuals in later generations will have preferences that vary according to the success of recently replaced individuals. There are many ways that can be implemented, but one approach is enough to illustrate what typically emerges: each new individual copies the preference probabilities of the more successful of the last two replaced individuals (i.e. the one with the highest final position) but with a random "mutation" added from the range [-0.02, 0.02]. Those small mutations are sufficient to allow the preferences to drift away from the symmetric 0.5 values if a final position advantage emerges from doing so.

There are eighteen distinct combinations to consider: three

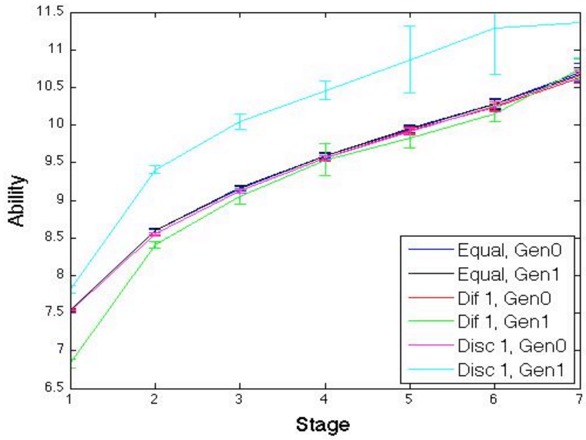


Figure 12: Ability levels for each gender (Gen0, Gen1) in the baseline (Equal), ability difference (Dif 1) and discrimination (Disc 1) cases, for preference-based profession choice.

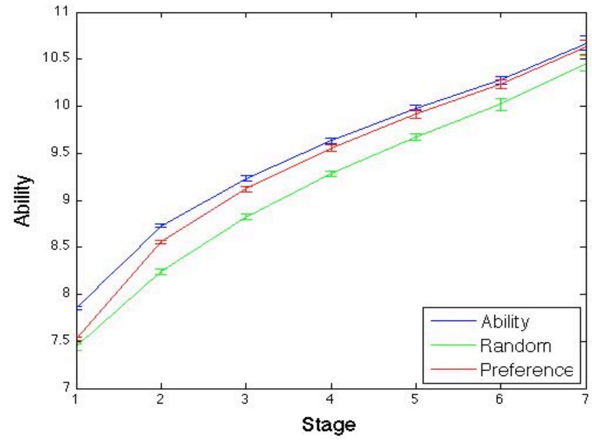


Figure 13: Ability levels for each gender in their appropriate profession for the ability difference (Dif 1) or discrimination (Disc 1) cases, for the three profession choice approaches.

basic conditions (baseline, gender-based ability difference, and gender-based discrimination), three profession choice approaches (by best ability, purely random, and random with intrinsic preferences), and each case can either involve, or not involve, intervention to equalize the numbers of each gender. Clearly, all the baseline cases, all the ability-based and pure random profession choice cases, and all the intervention cases, lead to the average preferences remaining at 0.5 because there is nothing to drive natural selection away from that symmetric case. However, those cases were still run to provide a check that no unexpected biases exist in the simulations. If there is no intervention, both the ability difference and discrimination case preferences shift towards the profession where the greatest success is most likely, while the preferences for the baseline case remain near 0.5 as expected. This is probably the simplest explanation of many of the observed gender differences in the numbers entering certain professions.

Since the professions and genders are set up symmetrically, and there is a fixed promotion rate at each stage, the overall average position at any given time for each gender must be independent of all the other factors, including any changing profession preferences, so it is not obvious what the changes really are optimizing. Figure 11 shows the average positions achieved by one gender as the preferences emerge when either an ability difference or discrimination favors Prof2 advances. The average positions achieved in *both* professions decrease as a result of the preferences changing away from being equal. Both genders gravitate towards the profession they do best at, increasing the competition there, and reducing the average position there for their gender. Those individuals remaining in the other profession face a larger pool of competitors of the better performing gender, and they are worse off on average too. It is the higher numbers in the best profession for each gender that keeps the average position constant throughout.

Often the most important issue for the businesses concerned is the average abilities at each stage in the two professions, irrespective of the genders involved. Figure 12 shows the ability levels at each stage after strong preferences emerge. Each profession employs individuals almost exclusively of just one gender, so the effects of gender-based discrimination or ability differences are very small, and all the ability levels converge, except for the very few individuals who persist in the profession that discriminates against them.

A final question of interest here is how does the emergence of profession preferences affect the society as a whole, given that they have been driven purely by individuals wanting to reach higher positions in their chosen profession. Figure 13 compares the ability levels for each gender in their most appropriate profession. The best abilities overall come from ability-based profession choice, and the worst abilities come from a random profession choice. That is to be expected, given that random choice means many individuals will not end up performing according to their best potential. The not-so-obvious result is that emergent profession preferences are able to bring the ability levels close to the ability-based levels, particularly for the higher stages. This might be important for the population as a whole if the abilities are difficult to assess before the profession choice needs to be made.

## Conclusions and Discussion

A general population-based framework has been proposed that enables the simulations of gender-based differences in various professions that involve ability-determined promotions up some form of hierarchy. The representative results presented were primarily chosen to demonstrate that the models do lead to reliable results in key simplified scenarios, though the approach can also be used to generate novel results for known real-world scenarios. The simulations presented have served to show how the principal factors can be studied effectively within the framework, and illustrated how distinct causes can lead to indistinguishable consequences, how preferences are able to emerge by natural selection, and how inappropriate interventions can make matters worse rather than better.

To simplify the presentation, all the simulations reported in this paper have only involved two professions, and all the gender-based differences have been symmetric across those professions. In reality, of course, there are many more than two professions, and a distinct lack of symmetries, but the proposed simulation framework is general enough to cope easily with such complications. The results presented in this paper will then serve as the baseline against which those more realistic simulations can be compared.

There are clearly many other factors that could be built into the simulations, so hopefully this modeling approach will become more widely used in the future. A key issue is that there are too many potential gender-based effects to simulate all the possible combinations, but there are numerous specific hypotheses that could be tested empirically with the approach. One concerns the effect of gender differences in risk taking leading to differences in the variance of abilities (e.g., Schubert, 2006; Robison-Cox, Martell and Emrich, 2007). Another relates to distinct career paths to the top levels of some professions, with different gender effects on each (e.g., Robison-Cox, Martell and Emrich, 2007). It is also easy to fix the number of individuals at each level, rather than let it emerge by promoting a fixed fraction of eligible individuals each year, and that would lead to simulations more like those of the corporate management study of Robison-Cox, Martell and Emrich (2007) than the merit-based promotions more typical in academia. The consequences of other suggestions could also be explored, such as that women are less aggressive about seeking promotion, or are quicker to give up waiting for promotion, or more likely to leave or take time out for other reasons such as maternity leave, etc. All these ideas could be tested explicitly within the presented framework.

There are a number of further computational complexities that could relatively easily be incorporated into the simulation approach presented in this paper, that have previously been tested in simulations of Life History Evolution (Bullinaria, 2009, 2010), such as allowing abilities and preferences that change with time, or having parameter value distributions rather than parameters fixed at particular values. Hopefully, however, this short paper has been sufficient to demonstrate that the general framework proposed will allow all manner of additional factors to be explored in a more systematic manner than previously, in which the assumptions and simplifications are explicit, and the effects quantifiable. It is inevitable that some readers will disagree with the particular assumptions and simplifications employed in the simulations presented here. Hopefully progress can be made by other researchers using the approach to test the consequences of varying those assumptions and simplifications, and performing simulations more carefully matched to their own data and beliefs.

## References


Altonji, J.G. and Blank, R.M. (1999). Race and gender in labor market. *Handbook of Labor Economics*, **3C**, 3143-3259.

Baron-Cohen, S. (2004). *The Essential Difference.* London, UK: Penguin.

Benbow, C.P. (1988). Sex differences in mathematical reasoning ability in intellectually talented preadolescents: Their nature, effects, and possible causes. *Behavioral And Brain Sciences*, **11**, 169-183.

Browne, K.R. (2002). *Biology at Work: Rethinking Sexual Equality*. New Brunswick, NJ: Rutgers University Press.

Bullinaria, J.A. (2009). Lifetime learning as a factor in Life History Evolution. *Artificial Life*, **15**, 389-409.

Bullinaria, J.A. (2010). Memes in artificial life simulations of Life History Evolution. In: *Proceedings of Artificial Life XII Conference (Alife XII)*, 823-830. Cambridge, MA: MIT Press.

Camp, T. (1997). The incredible shrinking pipeline. *Communications of the ACM*, **40**, 103-110.

Ceci, S.J. and Williams, W.M. (2011). Understanding current causes of women's underrepresentation in science. *Proceedings of the National Academy of Sciences*, **108**, 3157-3162.

Davison, H.K. and Burke, M.J. (2000). Sex discrimination in simulated employment contexts: A meta-analytic investigation. *Journal of Vocational Behavior*, **56**, 225-248.

Engelbrecht, A.P. (2007). *Computational Intelligence: An Introduction*. Sussex, UK: Wiley.

Geary, D.C. (1998). *Male, female: The evolution of human sex differences*. Washington, DC: American Psychological Association.

Ginther, D. and Hayes, K. (2003). Gender differences in salary and promotion for faculty in the humanities 1977–1995. *Journal of Human Resources*, **38**, 34-73.

Gürer, D. and Camp, T. (2001). Investigating the incredible shrinking pipeline for women in computer science. Final Report NSF 9812016.

Halpern, D.F. (2012). *Sex Differences in Cognitive Abilities (4th Edition)*. New York, NY: Psychology Press.

Halpern D.F. et al. (2007). The science of sex differences in science and mathematics. *Psychological Science*, **8**, 1-51.

Handelsman J. et al. (2005). More women in science. *Science*, **309**, 1190-1191.

Helbing, D. (2010). *Quantitative Sociodynamics*. Berlin, Germany: Springer.

Humphreys, L.G. (1988). Sex-differences in variability may be more important than sex differences in means. *Behavioral and Brain Sciences*, **11**, 195-196.

Lyness, K.S. and Heilman, M.E. (2006). When fit is fundamental: Performance evaluations and promotions of upper-level female and male managers. *Journal of Applied Psychology*, **91**, 777-785.

Lyness, K.S. and Judiesch, M.K. (1999). Are women more likely to be hired or promoted into management positions? *Journal of Vocational Behavior*, **54**, 158-173.

Martell, R.F., Lane, D.M. and Emrich C. (1996). Male-female differences: A computer simulation. *American Psychologist*, **51**, 157-158.

Mixon, F. and Trevino, L. (2005). Is there gender discrimination in named professorships? An econometric analysis of economics departments in the US South. *Applied Economics*, **37**, 849-854.

Moss-Racusin, C.A., Dovidio, J.F., Brescoll, V.L., Graham, M.J. and Handelsman, J. (2012). Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, **109,** 16474-16479.

Richiardi, M., Leombruni, R., Saam, N. and Sonnessa, M. (2006). A common protocol for agent-based social simulation. *Journal of Artificial Societies and Social Simulation,* **9**, 1-15.

Robison-Cox, J.F., Martell, R.F. and Emrich, C.G. (2007). Simulating gender stratification. *Journal of Artificial Societies and Social Simulation,* **10**, 1-8.

Rosenbaum, J.E. (1979). Tournament mobility: Career patterns in a corporation. *Administrative Science Quarterly*, **24**, 220-241.

Sabatier, M. (2010). Do female researchers face a glass ceiling in France? A hazard model of promotions. *Applied Economics*, **42**, 2053-2062.

Schneider, A. (1998). Why don't women publish as much as men? *Chronicle of Higher Education*, **45**, 14-16.

Schubert, R. (2006). Analyzing and managing risks – On the importance of gender differences in risk attitudes. *Managerial Finance*, **32**, 706-715.