



EUI Working Papers

ECO 2012/18

DEPARTMENT OF ECONOMICS

REFERRAL NETWORKS AND THE ALLOCATION OF TALENT

David Pothier

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
DEPARTMENT OF ECONOMICS

Referral Networks and the Allocation of Talent

DAVID POTHIER

EUI Working Paper **ECO** 2012/18

This text may be downloaded for personal research purposes only. Any additional reproduction for other purposes, whether in hard copy or electronically, requires the consent of the author(s), editor(s). If cited or quoted, reference should be made to the full name of the author(s), editor(s), the title, the working paper or other series, the year, and the publisher.

ISSN 1725-6704

© 2012 David Pothier

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Referral Networks and the Allocation of Talent

David Pothier*

June 30, 2012

Abstract

We study a model of occupational choice where workers must rely on their social contacts to acquire job vacancy information. Contrary to the existing literature, we allow for worker heterogeneity in terms of their idiosyncratic skill-types. In this case, the allocation of talent (the matching of skills to tasks) becomes a welfare-relevant consideration. A worker's skill-type determines both his relative cost of specialising in different occupations and his productivity on the job. The model shows that relying on word-of-mouth communication for job search generates both positive externalities (due to improved labour market matching) and negative externalities (due to a poor allocation of talent). Which effect dominates depends on the properties of the job search and productivity functions. Taking into account worker heterogeneity shows that the degree of occupational segregation in competitive labour markets is generally not efficient.

1 Introduction

Choosing an occupation or profession is undoubtedly one of the most important economic decisions made by individuals. In many cases, this decision depends not only on an individual's idiosyncratic characteristics (e.g innate ability), but also on the occupations chosen by their family, friends and peers. Occupational segregation refers to the sorting of individuals across occupations based on their social, religious, ethnic and/or gender identity. Several theories have been proposed to explain the presence and persistence of occupational segregation in labour markets. Many of these rely on employers exhibiting some form of discriminatory preferences when recruiting new workers belonging to different social groups. Another line of research first suggested by Arrow [1], which has only recently begun to garner attention in the economics literature, argues that the widespread use of informal referral networks is also an important cause of occupational segregation. Building on this latter branch of the economics literature, this paper studies how reliance on informal referral networks, via its effects on occupational segregation, can lead to an inefficient allocation of talent in a market economy.

*Department of Economics, European University Institute. Contact: dpothier@eui.eu. I would like to thank Fernando Vega Redondo and Piero Gottardi for their comments on earlier drafts of this paper. I would also like to thank participants at the EUI Microeconomics Working Group and the 2011 Workshop on Discrimination and Segregation at Sciences Po Paris for useful discussions. All remaining errors are mine.

Homophilic inbreeding in social networks refers to individuals being disproportionately likely to form social ties with other individuals belonging to the same social group. There exists a wealth of empirical evidence suggesting that such homophilic inbreeding is a widespread social phenomena.¹ One of the key building blocks of this paper is the assumption that workers' reliance on homophilic networks to access job vacancy information introduces a degree of strategic complementarity in the occupational choice decision made before entering the labour market. All else equal, workers who rely on their social contacts to obtain job vacancy information are more likely to choose an occupation that is popular among individuals belonging to their own social group, as this maximises their likelihood of finding a job. In effect, the use of informal referral networks in job search generates a positive externality: by choosing to specialise in a given occupation, an individual increases the probability that other individuals in his social network choosing the same occupation are successfully employed in the labour market.

When workers differ in terms of some idiosyncratic skill characteristic, they face an important trade-off when choosing an occupation. On one hand, as the logic above suggests, workers would like to choose an occupation that is popular among individuals in their social network as this increases the probability that they find a job. On the other hand, workers prefer choosing an occupation in which they are relatively more able, as this minimises the idiosyncratic specialisation cost they must incur before entering the labour market. If in addition to affecting his cost of specialisation, a worker's skill-type also affects his productivity on the job, reliance on informal networks in job search generates a negative externality: when choosing an occupation, a worker does not internalise how his occupational choice decision affects the allocation of talent, and thus aggregate labour productivity.

The model we consider incorporates both *technological* externalities (via the job search technology) and *pecuniary* externalities (via the mismatch of talent and its effects on equilibrium wages). The latter depends critically on two assumptions. The first is that workers' skill-types cannot be observed by firms. If this were not the case, firms could offer a menu of wages which would vary as a function of workers' productivity. As is well known, in such a complete markets setting the pecuniary externalities generated by workers' occupational choice decisions would not imply an efficiency loss, as wages would adjust so that workers receive their marginal product regardless of which occupation they choose.² The second is that workers' occupational choice decisions are strategic complements due to the presence of network externalities in job search. If these were absent, and workers' occupational choice decisions were non-strategic, then the competitive equilibrium would be Pareto efficient: even though workers' skill-types would still be unobservable and markets incomplete, specialisation costs would drive workers to choose the occupation in which they are relatively more able. Hence, both the positive (technological) externalities and negative (pecuniary) externalities fundamentally depend on the degree of homophily.

¹See for example the landmark study by McPherson, et al [14].

²See Greenwald and Stiglitz [11] for a more general discussion of this issue.

The direction and magnitude of these externalities are shown to depend critically on the properties of the job search and productivity functions. An efficient allocation in this economy trades-off an increase in employment from more efficient job search when workers segregate across occupations, with the associated decrease in productivity implied by a misallocation of talent. In a decentralised economy, on the other hand, workers ignore how their occupational choice decision affects both the job finding probability of other workers, and the allocation of talent across occupations. Instead, workers choose an occupation by comparing the expected wage with the associated idiosyncratic cost of specialisation, taking as given the occupational choice decision of other workers in the economy. Thus, whenever the efficiency loss from misallocating talent is sufficiently large, the competitive market can lead to “too much” occupational segregation. Indeed, providing that the costs in terms of loss productivity outweigh the benefits in terms of job matching efficiency, an efficient allocation would have workers choosing the occupation in which they are relatively more able.

The presence of network effects in occupational choice implies that workers’ best response functions are highly non-linear. Because of this, we are unable to provide a full analytical characterisation of the set of equilibria. Nonetheless, we can still answer some interesting questions such as: (1) when can occupational segregation be supported as an equilibrium phenomenon in competitive labour markets, and (2) under what conditions is occupational segregation actually efficient. Heuristically speaking, we identify conditions (most notably restrictions on the variable parameterising the degree of homophilic inbreeding bias) under which occupational segregation can be supported in equilibrium. We find that these conditions are generally not the same as those needed for occupational segregation to be efficient, implying that the degree of occupational segregation observed in competitive labour markets may not be welfare maximising. Moreover, we are able to derive explicit conditions explaining how the two different externalities contribute to the divergence between equilibrium and efficient outcomes. Intuitively, we find that the technological externalities are increasing in the efficiency of the job search technology, while the pecuniary externalities are exacerbated when workers’ productivity on the job is very sensitive to their skill-type.

As a clarifying remark, this paper does not argue that gender-based or race-based discrimination is not a major cause of the observed occupational segregation in labour markets. On the contrary, the sociological evidence in support of discrimination-based theories is overwhelming. Rather, the model developed in this paper serves to highlight a complementary channel through which occupational segregation can arise in competitive labour markets. To this extent, it serves to underline the view that even in societies in which the level of overt discrimination is on the wane (due to either broad societal changes or more specific political causes), there is reason to believe that occupational segregation will persist as a salient feature of labour markets. What is more, the conclusions presented below suggest that in economic environments where social networks play an important role in allocating workers to vacancies, competitive labour markets often fail to obtain efficiency. This suggests that some form of policy intervention in labour markets may be justified, even though a detailed discussion of the specific features of such policy intervention lies beyond the scope of this paper.

The remainder of this paper is organised as follows. Section 2 provides a brief review of the relevant theoretical literature. Section 3 describes the primitives of the model. Section 4 gives a definition of the equilibrium concept and the relevant welfare benchmark. Equilibrium and welfare analysis is presented in Section 5. Section 6 concludes.

2 Literature

The paper most closely related to the model developed below is the recent work of Buhai and van der Leij [5]. They study explicitly how reliance on social networks in job search can lead to occupational segregation in the labour market. More specifically, they develop a network model of occupational segregation between different social groups generated by the presence of a positive inbreeding bias among individuals belonging to the same group. In addition to examining the degree of wage and unemployment inequality between social groups in equilibria characterised by occupational segregation, they also study the welfare implications of such segregation. Interestingly, they find that welfare optimality implies a positive degree of occupational segregation in the labour market. However, contrary to the present paper, Buhai and van der Leij assume all workers to be *ex ante* homogenous. This paper therefore extends their model by introducing a degree of heterogeneity among workers, so that some workers are intrinsically more able in one occupation compared to another. This allows us to study the interaction between occupational segregation and the mismatch of workers to tasks when referrals are important for job search.

Another closely related paper by Bentolila, et al [3] studies the implications of social contacts for occupational mismatch. As in this paper, they find that social contacts imply both benefits (in terms of job finding probability) and costs (in terms of labour force productivity). Their paper, however, makes no mention of occupational segregation. Indeed, they assume an exogenous correlation between workers' skill-type and the skill-type of their social contacts. The distortions that arise due to occupational mismatch depend fundamentally on this exogenous parameter. In this paper, the skill-type of a worker's social contacts is endogenously determined via workers' occupational choice decisions.

Work by Bowles, et al [4] examines the conditions under which inequality can emerge and persist between *ex ante* identical social groups in a competitive market environment. Their results follow from three key factors, namely: (1) the extent of homophily in social networks, (2) the strength of local spillovers in human capital accumulation, and (3) the sensitivity of relative wages to the skill composition in production. They find that inter-group inequality can emerge and persist providing that the degree of homophilic inbreeding bias is sufficiently large. Their model is extended by Kim [12], who also considers how the degree of inequality is affected by network effects in job search. Contrary to these models, this paper focuses on how network externalities in job search affect the allocation of talent in competitive economies, rather than the degree of inter-group wage inequality.

More generally, this paper builds on the theoretical literature studying the role of social interactions in economic environments. The first to tackle this issue formally within the economics literature was Schelling [16] in his landmark paper on residential segregation. Schelling showed that when household preferences exhibit a small degree of discrimination, urban environments can quickly evolve into completely segregated neighborhoods. The labor market implications of Schelling’s “tipping model” were most famously addressed by Benabou [2], who modeled the residential choice and human capital accumulation decision of workers in a competitive labor market. His key finding was that local complementarities in human capital accumulation can lead to occupational segregation in equilibrium, and that this segregated outcome is generally inefficient.

Although the inefficiency results in Benabou may appear very similar to the results found in this paper, it is important to emphasise what features the two models have in common and in what ways they differ. A common element is that the two models incorporate both local and global interactions. In both papers, the global interactions are modeled in terms of neoclassical complementarities in production. The local complementarities in Benabou’s paper take the form of local spillovers to human capital accumulation, while in this paper they take the form of local spillovers to job search. A key difference lies in the fact that Benabou focuses on vertical heterogeneity (skilled versus unskilled) rather than horizontal heterogeneity. Consequently, local spill-overs in human capital investment flow only from the high-skilled to low-skilled workers. Insofar as residential segregation separates highly educated workers from less educated workers, it increases the total cost of education for low-skilled communities and reduces the supply of complementary low-skilled workers, thereby leading to inefficient outcomes. The model developed below differs importantly in this respect, as workers are assumed to be horizontally rather than vertically differentiated. This implies that the local spillovers to job search are symmetric across agents, regardless of their chosen occupation.

Finally, the model of job referral networks we adopt builds largely on the seminal contribution by Montgomery [15], who studied the interaction between an informal referral network and an anonymous competitive market subject to adverse selection frictions. Montgomery argued that the presence of a homophilic inbreeding bias effectively allows firms hiring by referral to costlessly screen job applicants. A key difference between the referral network studied in this paper and Montgomery’s model is that the homophilic bias in Montgomery’s paper exists among workers of similar skill-type, while in this paper it affects individuals belonging to the same social group. Consequently, referrals play no screening role in our paper, but instead simply serve as a mechanism via which workers obtain job vacancy information. Nonetheless, the inbreeding bias remains an important parameter as it determines the degree of strategic complementarity in the occupational choice decision faced by workers, and thus the magnitude of the externalities generated by workers’ reliance on network-mediated job search.

3 The Model

The section describes the primitives of the model. We begin with a detailed discussion of workers' occupational choice decisions. We then explain the structure of the social network via which job vacancy information is disseminated, and briefly describe the production-side of the economy.

3.1 Workers

We consider an economy populated by a continuum of risk-neutral workers. Let N denote the set of workers, with (Lebesgue) measure normalised to two. Workers are *ex ante* heterogenous in terms of their idiosyncratic skill-type. We denote the type space by $\Theta = [0, 1]$. Moreover, workers are equally divided into two social groups: reds (R) and greens (G). For simplicity, we assume skill-types to be uniformly distributed in both groups: i.e $\theta \sim U[0, 1]$ for all $X \in \{R, G\}$.³ Notice that the assumption that skill-types are identically distributed across social groups implies that no one group is *ex ante* predisposed to any particular occupation. Any occupational segregation which arises in equilibrium will therefore be due to strategic considerations among the workers, rather than some presupposed productivity difference between individuals originating from different social groups.⁴ An individual worker's social 'colour' in this context thus only serves as a social marker, and is otherwise completely payoff irrelevant.

There exists a binary set of occupations for workers to choose from, denoted by $\Phi = \{A, B\}$. Before entering the labour market, workers must choose to specialise in one of the two occupations. Importantly, we assume that a worker cannot be hired by a firm unless he has specialised in a particular occupation.⁵ A worker's skill-type determines the idiosyncratic cost he must incur when choosing to specialise in one of the two occupations. We let $c_\phi(\theta) \in \mathbb{R}_+$ denote the cost incurred by a type θ worker choosing to specialise in occupation $\phi \in \{A, B\}$.

Assumption 1: The cost functions $c_\phi(\theta)$ for $\phi \in \{A, B\}$ satisfy the following conditions

1. Symmetry: $c_A(1 - \theta) = c_B(\theta)$
2. Monotonicity: $c'_A(\theta) > 0$ and $c'_B(\theta) < 0$

The monotonicity assumption implies that workers located on the left hand side (near 0) of the unit interval find it relatively easier to specialise in occupation A , while workers located on the right hand side (near 1) of the unit interval find it relatively easier to specialise in occupation B . The symmetry assumption implies that, absent any network effects, workers located at the midpoint of the unit interval will be indifferent between specialising in occupation A or occupation B .

³This uniform assumption is made to simplify the derivations. The qualitative properties of the results would be unchanged had we instead assumed a single-peaked symmetric distribution function.

⁴This is consistent with Loury's [13] axiom of *anti-essentialism*.

⁵The decision should therefore be viewed as an investment in some observable and publicly recognised certificate needed for employment within a particular industry (e.g a law or architecture degree).

A worker’s skill-type also determines his productivity if hired by a firm. Importantly, we assume that workers’ skill-types cannot be observed by firms, implying that the wage a worker receives if employed is independent of his type. Since the labour market is competitive, this means that workers hired in the same occupation will all receive the same wage. In equilibrium, this wage will equal the average productivity of workers specialised in that occupation. We let $z_\phi(\theta) \in \mathbb{R}_+$ denote the productivity of a type θ worker employed in occupation $\phi \in \{A, B\}$.

Assumption 2: The productivity functions $z_\phi(\theta)$ for $\phi \in \{A, B\}$ satisfy the following conditions

1. Symmetry: $z_A(1 - \theta) = z_B(\theta)$
2. Monotonicity: $z'_A(\theta) < 0$ and $z'_B(\theta) > 0$

As with the cost function, the monotonicity assumption implies that workers located on the left hand side of the unit interval are relatively more productive in occupation A , while workers located on the right hand side of the unit interval are relatively more productive in occupation B .

3.2 Social Network

After having chosen an occupation, workers enter a competitive labour market. All matches on the labour market necessarily take place through an informal referral network. We assume network-mediated search to be subject to frictions insofar as workers are more likely to receive a job offer when they have more friends specialised in the same occupation. In fact, we assume social ties to workers specialised in a different occupation provide no job vacancy information whatsoever.⁶

We model workers’ social network as an Erdos-Renyi random graph formed as a result of a binomial link formation process. We assume the stochastic link formation process to be subject to an inbreeding bias: workers are disproportionately likely to form links with other workers belonging to the same social group. Formally, let $\alpha \in (1/2, 1)$ denote the conditional probability that a randomly chosen link from a given worker leads to another worker belonging to the same social group. Importantly, we assume that workers make their specialisation decisions *before* the stochastic network is realised.⁷ This implies that workers’ occupational choice decisions do not depend on the realised structure of the social network. Let η_ϕ^X denote the (expected) measure of a worker’s neighbours specialised in occupation $\phi \in \{A, B\}$ when a worker belongs to group $X \in \{R, G\}$. The probability that a worker receives a referral offer is given by the function $q(\eta_\phi^X) \in [0, 1]$.

Assumption 3: The job search function $q(\eta_\phi^X)$ is monotonically increasing such that $q'(\eta_\phi^X) > 0$.

⁶Intuitively, one should interpret the two occupations as being very different in terms of the skills they require. Building on the example provided in footnote 5, job vacancies for architects are unlikely to be of interest to someone holding a law degree.

⁷This is consistent with the interpretation whereby an individual’s social networks is constituted of “weak” or “instrumental” ties. Granovetter [10], among others, has shown that job vacancy information is more likely to be obtained from social connections made at university or on the labour market (and thus after workers have made some fixed investment in a career path), rather than from family or kinship ties.

This exogenous job search technology should be interpreted as a reduced-form representation of a dynamic job search process. The dynamics of such network-mediated job search have been studied in detail [6], [7], [9], [8]. In such models, the structure of workers' social networks plays a key role in determining the propagation of job vacancy information. As we are only concerned with the overall network externality generated by network-mediated job search, we choose to model the job search process using an exogenous job search function in order to simplify the analysis.

3.3 Firms

Firms employ workers specialised in the two occupations and produce a homogenous consumption good. We assume the two occupations to be essential and complements in production, implying that firms must employ a positive measure of workers specialised in each occupation in order to produce. We normalise the price of the consumption good to one, and let $\tilde{l}_\phi \in \mathbb{R}_+$ denote the *effective* labour supply (in efficiency units) of workers specialised in occupation $\phi \in \{A, B\}$. Note that due to the nature of the job search technology, a positive mass of workers remain unemployed in equilibrium. Hence, the *effective* labour supply will be strictly less than the total measure of workers in the economy. Firms are assumed to have access to a constant returns to scale production technology denoted by $y = f(\tilde{l}_A, \tilde{l}_B) \in \mathbb{R}_+$.

Assumption 4: The production function $f(\tilde{l}_A, \tilde{l}_B)$ satisfies the following conditions

1. Symmetry: $\frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_A} \Big|_{\tilde{l}_A = \tilde{l}_B} = \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_B} \Big|_{\tilde{l}_A = \tilde{l}_B}$
2. Monotonicity: $\frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_\phi} > 0$ for all $\phi \in \{A, B\}$
3. Concavity: $\frac{\partial^2 f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_\phi^2} < 0$ for all $\phi \in \{A, B\}$
4. Essentiality: $\lim_{\tilde{l}_A \rightarrow 0} \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_A} = \lim_{\tilde{l}_B \rightarrow 0} \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_B} = \infty$

Aside from the symmetry assumption, these restrictions consist of the traditional Inada conditions commonly imposed on neoclassical production functions. The symmetry assumption is used to simplify the analysis. It could be relaxed in order to study the effects of occupational segregation on the degree of wage inequality across occupations.

3.4 Timing

The timing of the model can be summarised as follows

- **Stage 1:** Workers choose to specialise in a particular occupation $\phi \in \{A, B\}$ and incur the specialisation cost $c_\phi \in \mathbb{R}_+$.
- **Stage 2:** Workers randomly and non-strategically form a network of social connections.
- **Stage 3:** Workers enter the labour market and are hired with probability $q(\eta_\phi^X) \in [0, 1]$. Conditional on being hired in occupation $\phi \in \{A, B\}$, workers receive the wage $w_\phi \in \mathbb{R}_+$.

4 Equilibrium and Welfare: Definitions

In this section, we provide definitions of the equilibrium concept and the welfare benchmark to which equilibrium outcomes will be compared. Some preliminary results relating to properties of the equilibrium and efficient allocations are also provided in order to facilitate the ensuing analysis.

4.1 Equilibrium

Let $\sigma^X(\theta) \in [0, 1]$ denote the probability with which a type θ worker belonging to group $X \in \{R, G\}$ chooses occupation A , and define $\sigma^X \equiv \int_0^1 \sigma^X(\theta) d\theta$ to be the total measure of workers in group $X \in \{R, G\}$ choosing occupation A . The payoff function of a type θ worker belonging to group $X \in \{R, G\}$ and choosing to specialise in occupation $\phi \in \{A, B\}$ is given by

$$U_\phi^X(\sigma^X, \sigma^{X'}; \theta) = q(\eta_\phi^X(\sigma^X, \sigma^{X'})) w_\phi(\sigma^X, \sigma^{X'}) - c_\phi(\theta) \quad (1)$$

The first term of the right hand side of this equation equals a worker's expected wage when choosing occupation $\phi \in \{A, B\}$, while the second term equals the cost he must incur in order to specialise in this occupation. A worker chooses an occupation in order to maximise this objective function, taking the distribution of workers across occupations and wages as given.

The objective function of the (representative) firm is given by the following profit function

$$\Pi = f(\tilde{l}_A, \tilde{l}_B) - w_A l_A - w_B l_B$$

The firm chooses a labour demand schedule l_ϕ for $\phi \in \{A, B\}$ to maximise its profits, taking the occupational choice decision of workers and wages as given.

Definition 1: An equilibrium is defined as a specialisation strategy $\sigma^X(\theta)$ for all $X \in \{R, G\}$ and $\theta \in [0, 1]$, and a vector of wage rates w_ϕ for $\phi \in \{A, B\}$ such that

1. Workers choose a specialisation strategy $\sigma^X(\theta)$ to maximise their utility, taking wages and the occupational choice decisions of other workers in the economy as given.
2. The representative firm choose a labour demand schedule $(l_\phi)_{\phi \in \{A, B\}}$ to maximise its profits, taking wages and the occupational choice decision of workers as given.
3. The labour market clears.

Definition 2: An equilibrium in *threshold strategies* is an equilibrium such that the specialisation strategies $\sigma^X(\theta)$ of workers satisfy the following condition

$$\exists \hat{\theta}^X \in [0, 1] \quad \forall X \in \{R, G\} : \quad \sigma^X(\theta) = 1 \quad \text{if } \theta < \hat{\theta}^X \quad \text{and} \quad \sigma^X(\theta) = 0 \quad \text{if } \theta > \hat{\theta}^X$$

Due to the assumed properties of the cost functions - in particular, the monotonicity and continuity assumptions - we can show that restricting attention to equilibria in threshold strategies is without loss of generality.

Lemma 1: All equilibria are necessarily in threshold strategies.

Proof: See Appendix A.

Using Lemma 1, we can derive explicit expressions for the average productivity of workers across the two occupations as follows

$$E[z_A] = \begin{cases} 0 & \text{if } \hat{\theta}^X = 0 \quad \forall X \in \{R, G\} \\ \frac{1}{\hat{\theta}^R + \hat{\theta}^G} \sum_{X \in \{R, G\}} \int_0^{\hat{\theta}^X} z_A(\theta) d\theta & \text{otherwise} \end{cases} \quad (2)$$

$$E[z_B] = \begin{cases} 0 & \text{if } \hat{\theta}^X = 1 \quad \forall X \in \{R, G\} \\ \frac{1}{2 - \hat{\theta}^R - \hat{\theta}^G} \sum_{X \in \{R, G\}} \int_{\hat{\theta}^X}^1 z_B(\theta) d\theta & \text{otherwise} \end{cases} \quad (3)$$

We can then define the effective labour supply (in efficiency units) across occupations

$$\tilde{l}_\phi = E[z_\phi] l_\phi, \quad \forall \phi \in \{A, B\} \quad (4)$$

where l_ϕ denotes the measure of workers employed in occupation $\phi \in \{A, B\}$.⁸

Again using Lemma 1, we can derive explicit expressions for η_ϕ^X ; i.e the (expected) measure of a worker's neighbours specialised in each occupation

$$\eta_A^X = \alpha \int_0^{\hat{\theta}^X} d\theta + (1 - \alpha) \int_0^{\hat{\theta}^{X'}} d\theta \quad (5)$$

$$\eta_B^X = \alpha \int_{\hat{\theta}^X}^1 d\theta + (1 - \alpha) \int_{\hat{\theta}^{X'}}^1 d\theta \quad (6)$$

The effective labour supply of workers in each occupation is thus equal to

$$l_A = \sum_{X \in \{R, G\}} q(\eta_A^X) \int_0^{\hat{\theta}^X} d\theta \quad (7)$$

$$l_B = \sum_{X \in \{R, G\}} q(\eta_B^X) \int_{\hat{\theta}^X}^1 d\theta \quad (8)$$

Notice that since $q(\eta_\phi^X) \in (0, 1)$, the measure of employed workers in a given occupation will always be strictly less than the measure of workers specialising in that occupation.

⁸Notice that since types are unobservable and the probability of being employed is independent of types, the average productivity of workers specialised in occupation $\phi \in \{A, B\}$ is the same as the average productivity of workers employed in occupation $\phi \in \{A, B\}$.

4.2 Social Welfare

We now turn to the task of defining the relevant welfare benchmark for the economy described above. The welfare criterion we use is based on the notion of utilitarian efficiency, rather than the more common notion of Pareto efficiency. The utilitarian welfare benchmark is a compelling one, especially if one is interested in comparing different outcomes from an *ex ante* perspective [17]. Given that workers located at the extrema of the type space always have a payoff advantage relative to those located in the middle of the type space, it makes sense to consider *ex ante* rankings of different potential equilibrium outcomes. Aggregate utilitarian welfare in this economy is given by the following function

$$W = \sum_{X \in \{R, G\}} \int_0^1 (\sigma^X(\theta) U_A^X(\theta) + (1 - \sigma^X(\theta)) U_B^X(\theta)) d\theta$$

Rewriting this equation in terms of threshold profiles, we obtain

$$W(\hat{\theta}^R, \hat{\theta}^G) = \sum_{X \in \{R, G\}} \left(\int_0^{\hat{\theta}^X} U_A^X(\theta) d\theta + \int_{\hat{\theta}^X}^1 U_B^X(\theta) d\theta \right)$$

Using the definition of the payoff function as given by condition (1), together with the labour supply conditions (7) and (8) and the assumption that the production function exhibits constant returns-to-scale, we can rewrite this expression as follows

$$W(\hat{\theta}^R, \hat{\theta}^G) = f(\tilde{l}_A, \tilde{l}_B) - C(\hat{\theta}^R, \hat{\theta}^G) \tag{9}$$

where

$$C(\hat{\theta}^R, \hat{\theta}^G) = \sum_{X \in \{R, G\}} \left(\int_0^{\hat{\theta}^X} c_A(\theta) d\theta + \int_{\hat{\theta}^X}^1 c_B(\theta) d\theta \right) \tag{10}$$

The social planner thus seeks to maximise total output net of the aggregate specialisation costs. The salient trade-off which characterises the efficient allocation can be summarised as follows. On one hand, the social planner would like to segregate workers belonging to different groups across occupations, as this increases the efficiency of the job matching technology and thereby maximises aggregate employment in this economy. However, he must balance this against the increase in specialisation costs and the decrease in average labour productivity that arise when workers segregate, as this leads to a misallocation of talent. To gain a better intuition of this underlying trade-off, consider the two benchmark cases without inbreeding bias and with homogenous skill-types, respectively. Absent any inbreeding bias, the social planner would have each worker choose the occupation in which he is most able, as this minimises total costs and allocative inefficiencies while leaving total output unchanged. If workers were homogenous but the referral network exhibited positive inbreeding bias, the social planner would have workers completely segregate across occupations, as this maximises aggregate employment and thus total output.

In general, the symmetry properties imposed on the cost and production functions imply that the efficient allocation must also be symmetric. This leads us to the following result.

Lemma 2: The welfare maximizing threshold profile $(\hat{\theta}^R, \hat{\theta}^G)$ is necessarily symmetric such that $\hat{\theta}^R = 1 - \hat{\theta}^G$.

Proof: See Appendix A.

5 Equilibrium and Welfare: Analysis

The objective of this section is to identify conditions under which occupational segregation can arise in equilibrium. Formally, occupational segregation is defined as any deviation from the mixed threshold profile $(\hat{\theta}^R, \hat{\theta}^G) = (1/2, 1/2)$, whereby each worker chooses the occupation in which he is relatively more able. Complete occupational segregation is defined as one of the two corner solution threshold profiles: i.e $(\hat{\theta}^R, \hat{\theta}^G) \in \{(1, 0), (0, 1)\}$. Any intermediate threshold profile will be referred to as partial occupational segregation. We are particularly interested in characterising the conditions under which complete occupational segregation can be supported as an equilibrium, and compare these to the conditions needed for complete occupational segregation to be efficient. More specifically, we show that there exist cut-off values of the inbreeding bias parameter, $\underline{\alpha}^{EQ} \in (1/2, 1)$ and $\underline{\alpha}^{SW} \in (1/2, 1)$, above which complete occupational segregation can be supported in equilibrium and is efficient, respectively. Interestingly, these two cut-off values generally do not coincide.

5.1 Optimality Conditions

We begin by deriving the general optimality conditions which characterise equilibrium prices and allocations. Profit maximisation implies that firms' labour demand must satisfy the following condition

$$w_\phi = E[z_\phi] \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_\phi}, \quad \forall \phi \in \{A, B\} \quad (11)$$

For notational simplicity, we denote the difference in expected wages across occupations for a worker belonging to group $X \in \{R, G\}$ by

$$\Delta_\phi E[w]^X = q(\eta_A^X)w_A - q(\eta_B^X)w_B \quad (12)$$

where w_ϕ for $\phi \in \{A, B\}$ is given by condition (11) above. We can then write the payoff difference across occupations for workers belonging to group $X \in \{R, G\}$ as follows

$$\Delta_\phi U^X(\theta) = \Delta_\phi E[w]^X - (c_A(\theta) - c_B(\theta))$$

Optimality implies that workers' specialisation decisions must satisfy the following conditions

$$\begin{aligned} \sigma^X(\theta) &= 0 & \text{if } \Delta_\phi U^X(\theta) < 0 \\ \sigma^X(\theta) &\in [0, 1] & \text{if } \Delta_\phi U^X(\theta) = 0 \\ \sigma^X(\theta) &= 1 & \text{if } \Delta_\phi U^X(\theta) > 0 \end{aligned} \tag{13}$$

for all $\theta \in [0, 1]$ and $X \in \{R, G\}$. Any (interior) equilibrium threshold profile $(\hat{\theta}^R, \hat{\theta}^G) \in (0, 1)^2$ must then satisfy the following indifference conditions

$$q(\eta_A^X)w_A(\hat{\theta}^X, \hat{\theta}^{X'}) - q(\eta_B^X)w_B(\hat{\theta}^X, \hat{\theta}^{X'}) = c_A(\hat{\theta}^X) - c_B(\hat{\theta}^X), \quad \forall X \in \{R, G\} \tag{14}$$

For corner solutions, so that $(\hat{\theta}^R, \hat{\theta}^G) \in \{(1, 0), (0, 1)\}$, these conditions may hold as inequalities.

5.2 General Results

Using these general optimality conditions, we can obtain a few results which do not depend on specific functional form assumptions. The first result follows somewhat trivially from the essentiality assumption imposed on the production function.

Proposition 1: No equilibrium exists such that all workers from both social groups specialise in the same occupation.

Proof: See Appendix A.

Proposition 2: The mixed threshold profile $(\hat{\theta}^R, \hat{\theta}^G) = (1/2, 1/2)$ can always be supported as an equilibrium.

Proof: See Appendix A.

This second result is rather intuitive when one interprets workers' occupational choice decision as a classic coordination game. Given that no worker chooses to segregate, the individual gains to choosing an occupation other than the one in which a worker holds a natural skill advantage are zero.

5.3 Constant Productivity

We now turn to a more thorough analysis of the set of equilibria, and analyse their welfare properties using specific functional form assumptions. Although these assumptions are restrictive, they allow us to solve for the allocations in closed-form, and thereby provide a more robust intuition of the

underlying workings of the model. To begin, we assume a Cobb-Douglas production function

$$f(\tilde{l}_A, \tilde{l}_B) = \tilde{l}_A^\beta \tilde{l}_B^{1-\beta}$$

where $\beta = 1/2$ in order to satisfy the symmetry condition of Assumption 4. In this case, profit maximisation implies that wages are equal to

$$w_A = \frac{E[z_A]}{2} \sqrt{\frac{\tilde{l}_B}{\tilde{l}_A}} \quad \text{and} \quad w_B = \frac{E[z_B]}{2} \sqrt{\frac{\tilde{l}_A}{\tilde{l}_B}}$$

The job search function and cost functions are assumed to be linear, so that $q(\eta_\phi^X) = \eta_\phi^X$ and

$$c_A(\theta) = k\theta \quad \text{and} \quad c_B(\theta) = k(1 - \theta)$$

where $k > 0$. In this section, we assume that workers' productivity on the job is independent of their skill-type. This allows us to isolate the technological externalities generated by network-mediated job search, without considering the potential pecuniary externalities engendered by the misallocation of talent. We relax this assumption in the next section, where we consider the case with heterogenous worker productivity. Normalising the productivity of labour to one, we have that $z_A(\theta) = z_B(\theta) = 1 \forall \theta \in \Theta$. Our first task is to explicitly derive the equilibrium conditions for this economy.

Equilibrium Analysis Utility maximisation implies that the threshold skill-type is implicitly determined by the following conditions

$$\hat{\theta}^X = \frac{1 + k^{-1} \Delta_\phi E[w]^X}{2}, \quad \forall X \in \{R, G\} \quad (15)$$

Notice that

$$\lim_{k \rightarrow \infty} \hat{\theta}^X(k) = \frac{1}{2}, \quad \forall X \in \{R, G\}$$

implying that as specialisation costs become large, the unique equilibrium threshold profile is $(\hat{\theta}^R, \hat{\theta}^G) = (1/2, 1/2)$. The difference in expected wages across occupation is given by

$$\Delta_\phi E[w]^X = \frac{E[z_A]}{2} \left(\alpha \hat{\theta}^X + (1 - \alpha) \hat{\theta}^{X'} \right) \sqrt{\frac{\tilde{l}_B}{\tilde{l}_A}} - \frac{E[z_B]}{2} \left(\alpha (1 - \hat{\theta}^X) + (1 - \alpha) (1 - \hat{\theta}^{X'}) \right) \sqrt{\frac{\tilde{l}_A}{\tilde{l}_B}}$$

where, from the common productivity assumption, we have that $E[z_A] = E[z_B] = 1$. The linearity of the job search function, together with the assumption that skill-types are uniformly distributed,

imply that we can write the labour supply functions as follows

$$l_A = \sum_{X \in \{R, G\}} \left(\alpha \hat{\theta}^X + (1 - \alpha) \hat{\theta}^{X'} \right) \hat{\theta}^X \quad (16)$$

$$l_B = \sum_{X \in \{R, G\}} \left(\alpha (1 - \hat{\theta}^X) + (1 - \alpha) (1 - \hat{\theta}^{X'}) \right) (1 - \hat{\theta}^X) \quad (17)$$

These conditions define a threshold equilibrium for this economy.

Result 1: When worker productivity is constant, complete occupational segregation can be supported in equilibrium for sufficiently high values of the inbreeding bias.

Proof: See Appendix A.

Welfare Analysis We now turn to characterising the properties of the efficient allocation for this economy. To make the welfare analysis comparable with the equilibria characterised above, we impose the same functional form assumptions. In this case, the welfare function (9) becomes

$$W = \tilde{l}_A^{1/2} \tilde{l}_B^{1/2} - k \sum_{X \in \{R, G\}} \left((\hat{\theta}^X) (\hat{\theta}^X - 1) + \frac{1}{2} \right)$$

Result 2: When worker productivity is constant, if complete occupational segregation can be supported in equilibrium then it is also efficient.

Proof: See Appendix A.

The divergence between the equilibrium and efficient outcomes stems from the fact that workers do not internalise how their occupational choice decisions affect the job finding probability of other workers in the economy. Indeed, the (positive) externalities that result from network-mediated job search imply that for intermediate levels of the inbreeding bias, workers will choose not to segregate even though it would be efficient for them to do so. It must be emphasised that this result is very specific to the case where workers' productivity is independent of their skill-type, and does not (in general) extend to the case where workers' productivity across occupations varies as a function of their skill-type. We explore the equilibrium and welfare properties of this more interesting case below.

5.4 Heterogenous Productivity

We now do away with the assumption that workers' productivity is constant, and consider the case where workers' productivity on the job varies as a function of their skill-type. In this section, we

consider the case where the productivity functions are linear so that

$$z_A(\theta) = h(1 - \theta) \quad \text{and} \quad z_B(\theta) = h\theta$$

where $h > 0$. In this case, conditions (2) and (3) can be rewritten as follows

$$E[z_A] = \frac{h}{\hat{\theta}^R + \hat{\theta}^G} \sum_{X \in \{R, G\}} \left(\hat{\theta}^X - \frac{\hat{\theta}^{X2}}{2} \right)$$

$$E[z_B] = \frac{h}{2 - \hat{\theta}^R - \hat{\theta}^G} \sum_{X \in \{R, G\}} \left(\frac{1}{2} - \frac{\hat{\theta}^{X2}}{2} \right)$$

Equilibrium Analysis The equilibrium conditions (15)-(17) remain the same as in the case with constant productivity. As before, we are interested in identifying restrictions on the inbreeding bias parametre $\underline{\alpha}^{EQ} \in (1/2, 1)$ above which occupational segregation can be supported in equilibrium, and compare this to the cut-off value $\underline{\alpha}^{SW} \in (1/2, 1)$ above which occupational segregation is efficient.

Result 3: When workers vary in terms of their productivity on the job, complete occupational segregation can be supported in equilibrium for sufficiently high values of the inbreeding bias.

Proof: See Appendix A.

Welfare Analysis The social welfare function is the same as in the constant productivity case. However, the welfare properties of the equilibrium are markedly different when workers' productivity on the job varies as a function of their skill-type.

Result 4: When workers' productivity depends linearly on their skill-type, complete occupational segregation can be supported in equilibrium even though it is inefficient.

Proof: See Appendix A.

Why do the welfare properties of the competitive equilibrium differ so dramatically when workers' productivity varies as a function of their skill-type, compared to the case with constant worker productivity? Heuristically speaking, the difference arises because heterogenous worker productivity, coupled with the fact that individual worker productivity cannot be observed by firms, generates a negative pecuniary externality. In this particular case, this (negative) externality dominates the (positive) externality generated by the job search technology. The following section delves into this issue in more detail.

5.5 Externalities

In this section, we do away with the specific functional forms for the cost and productivity functions. This allows us to gain a more general understanding of the direction and magnitude of the externalities driving the divergence between the equilibrium and efficient outcomes. Moreover, it allows us to identify the conditions under which the degree of occupational segregation in the competitive market exceeds that implied by the welfare maximising allocation.

In what follows, we restrict attention to symmetric allocations so that $\hat{\theta}^R = 1 - \hat{\theta}^G$. This implies that labour supply, average productivity and wage rates will be equal across occupations. Given this, condition (11) implies that wages are given by

$$w = \frac{E[z]}{2} \quad (18)$$

where, using conditions (2) and (3), we have that the average productivity of workers is equal to

$$E[z] = \int_0^{\hat{\theta}} z_A(\theta) d\theta + \int_{\hat{\theta}}^1 z_B(\theta) d\theta \quad (19)$$

Focusing on interior solutions such that $\hat{\theta} \in (0, 1)$, the indifference condition (14) becomes

$$w(q(\eta_A) - q(\eta_B)) - (c_A(\hat{\theta}) - c_B(\hat{\theta})) = 0 \quad (20)$$

Using the symmetry restriction, we can write the social welfare function (9) as follows

$$W = E[z](q(\eta_A)\hat{\theta} + q(\eta_B)(1 - \hat{\theta})) - 2 \left(\int_0^{\hat{\theta}} c_A(\theta) d\theta + \int_{\hat{\theta}}^1 c_B(\theta) d\theta \right) \quad (21)$$

Note that the social welfare function is symmetric about $\hat{\theta} = 1/2$. Consequently, we can restrict attention to allocations belonging to the interval $\hat{\theta} \in (1/2, 1)$ without loss of generality.

The objective of this section is to compare the equilibrium condition (20) to the conditions that maximise the social welfare function (21). As above, we proceed in two stages. First, we consider the case where workers' productivity is constant in order to isolate the effect of the technological externality generated by the job search technology. Second, we allow for heterogeneous worker productivity in order to identify the effect of the pecuniary externality engendered by the misallocation of talent.

Technological Externality We begin by considering the case where workers' productivity is constant. Differentiating the social welfare function (21) with respect to $\hat{\theta}$ and using condition (18)

yields the following first-order condition

$$w(q(\eta_A) - q(\eta_B)) - (c_A(\hat{\theta}) - c_B(\hat{\theta})) + w\left(q'(\cdot)\frac{d\eta_A}{d\hat{\theta}}\hat{\theta} + q'(\cdot)\frac{d\eta_B}{d\hat{\theta}}(1 - \hat{\theta})\right) = 0 \quad (22)$$

where $q'(\eta_A) = q'(\eta_B) > 0$ due to the linearity and monotonicity assumptions imposed on the job search function. Using conditions (5) and (6), together with the symmetry restriction, we have that

$$\frac{d\eta_A}{d\hat{\theta}} = 2\alpha - 1 > 0 \quad \text{and} \quad \frac{d\eta_B}{d\hat{\theta}} = 1 - 2\alpha < 0$$

Substituting these conditions into condition (22) and rearranging yields

$$w(q(\eta_A) - q(\eta_B)) - (c_A(\hat{\theta}) - c_B(\hat{\theta})) + \mathcal{E}^T = 0 \quad (23)$$

where

$$\mathcal{E}^T = w\left(\hat{\theta} - \frac{1}{2}\right)q'(\cdot)\frac{d\eta_A}{d\hat{\theta}} = w\left(\hat{\theta} - \frac{1}{2}\right)q'(\cdot)(2\alpha - 1) > 0, \quad \forall \hat{\theta} \in (1/2, 1) \quad (24)$$

Condition (24) measures the (positive) technological externalities generated by the job search technology. Unsurprisingly, the magnitude of this externality is found to be strictly increasing in the value of the inbreeding bias parameter $\alpha \in (1/2, 1)$ and the efficiency of the job search technology, as measured by the slope of the function $q(\cdot)$. The key term of condition (24) is $q'(\eta_\phi)d\eta_\phi/d\hat{\theta}$, which measures the effect of changes in the threshold $\hat{\theta}$ on the job-finding probability of workers' neighbours specialised in occupation $\phi \in \{A, B\}$. This is the source of the technological externality generated by the job search technology: in the competitive market, a worker do not take into account how his occupational choice decisions affects the job-finding probability of other workers in his social network choosing the same occupation.

Pecuniary Externality We now turn to the case where workers' productivity varies as a function of their skill-type. In this case, differentiating the social welfare function (21) with respect to $\hat{\theta}$ and using condition (18) yields

$$w(q(\eta_A) - q(\eta_B)) - (c_A(\hat{\theta}) - c_B(\hat{\theta})) + \mathcal{E}^T + (q(\eta_A)\hat{\theta} + q(\eta_B)(1 - \hat{\theta}))\frac{dE[z]}{d\hat{\theta}} = 0 \quad (25)$$

where, using conditions (19), we have that

$$\frac{dE[z]}{d\hat{\theta}} = z_A(\hat{\theta}) - z_B(\hat{\theta}) < 0, \quad \forall \hat{\theta} \in (1/2, 1)$$

Substituting this into condition (25) and rearranging yields

$$w(q(\eta_A) - q(\eta_B)) - (c_A(\hat{\theta}) - c_B(\hat{\theta})) + \mathcal{E}^T + \mathcal{E}^P = 0 \quad (26)$$

where

$$\mathcal{E}^P = \frac{f(\tilde{l}_A, \tilde{l}_B)}{E[z]} \frac{dw}{d\hat{\theta}} = \frac{f(\tilde{l}_A, \tilde{l}_B)}{E[z]} \frac{(z_A(\hat{\theta}) - z_B(\hat{\theta}))}{2} < 0, \quad \forall \hat{\theta} \in (1/2, 1) \quad (27)$$

Condition (27) measures the (negative) pecuniary externalities that arise when workers' productivity on the job varies as a function of their skill-type. It is immediate to see that this effect will be greater when workers' productivity is very sensitive to their skill-type, implying that the difference $z_A(\hat{\theta}) - z_B(\hat{\theta})$ is large. The key term of condition (27) is $dw/d\hat{\theta}$, which measures the effect of changes in the threshold $\hat{\theta}$ on wages, and by extension the average productivity of labour. This is the source of the pecuniary externality: workers do not internalise how their occupational choice decisions affect the allocation of talent across occupations, and the consequences for equilibrium wages.

Discussion Comparing conditions (20) and (26), we find that the equilibrium and welfare-optimal allocations differ by two additive terms: \mathcal{E}^T and \mathcal{E}^P . The first measures the magnitude of the technological externalities generated by the job search technology, while the second measures the magnitude of the pecuniary externalities implied by the misallocation of talent. For any allocation $\hat{\theta} \in (1/2, 1)$, if the LHS of condition (20) is less than the LHS of condition (26), this implies that the competitive market supplies “too little” occupational segregation. Similarly, if the LHS of condition (20) exceeds the LHS of condition (26), then the competitive market supplies “too much” occupational segregation. It is immediate to see that the relative magnitudes of the externality terms \mathcal{E}^T and \mathcal{E}^P determine how the degree of occupational segregation supplied by the market compares to the welfare maximising allocation: if the (negative) pecuniary externality is greater/less than the (positive) technological externality, the market will supply too much/too little occupational segregation. Heuristically speaking, if workers' productivity on the job is very sensitive to their skill-type and the job search technology is relatively inefficient (implying that $q'(\cdot)$ is small in magnitude), then the level of occupational segregation supplied by the market will exceed the welfare maximising level.

Notice that in the case with constant productivity, we had that $dw/d\hat{\theta} = 0$. This is because of the symmetry assumption imposed on the production function. An immediate consequence of this symmetry property is that wages across occupations are constant, since aggregate labour supply does not vary as a function of the threshold $\hat{\theta}$. However, even if we were to relax this symmetry assumption, so that wages varied as a function of $\hat{\theta}$ even in the common productivity case, the negative pecuniary externality would remain. Indeed, the change in wages in this (more general) case could be decomposed into two components: the first generated by changes in the relative scarcity of labour across the two occupations, and the second generated by changes in the aggregate productivity of labour across the two occupations. The first effect would not, in itself, generate any divergence between the equilibrium and efficient outcomes. Indeed, even though changing $\hat{\theta}$ would lead to a change in relative wages across occupations, these variations would be perfectly internalised by utility-maximising agents in equilibrium. The second effect, on the other hand, would not be internalised by the agents since firms cannot post type-specific wage contracts. Why this difference? The key lies in the fact that while in equilibrium agents can

perfectly anticipate the marginal effect on wages implied by changes in the *quantity* of labour, they can only anticipate the average effect on wages implied by changes in the *quality* (or productivity) of labour. Hence, the pecuniary externalities discussed above would remain even if the production function failed to satisfy the symmetry assumption.

6 Conclusion

This paper argues that occupational segregation in competitive labour markets is generally not efficient when workers differ in terms of their idiosyncratic skill characteristics. Inefficiencies arise because workers do not internalise how their individual occupational choice decisions affect either: (1) the job finding probability of other individuals belonging to their social network, or (2) the average productivity of labour across different occupations. More specifically, we show that the divergence between the equilibrium and optimal allocations arise due to the presence of a (positive) technological externality and a (negative) pecuniary externality, and that the degree of occupational segregation in competitive labour markets can be either greater or less than the optimal degree of occupational segregation. Which of these two effects dominate depend on the properties of the job search and productivity functions. In particular, whenever the job search technology is relatively inefficient and the productivity of workers on the job is very sensitive to their skill-type, the level of optimal segregation supplied by the competitive market will exceed the welfare maximising level. More generally, the model suggests that the widespread use of social connections in labour market can lead to serious inefficiencies due to the misallocation of talent, even in the absence of discriminatory preferences.

References

- [1] K.J. Arrow. What has economics to say about racial discrimination? *The Journal of Economic Perspectives*, 12(2):91–100, 1998.
- [2] R. Benabou. Workings of a city: location, education, and production. *The Quarterly Journal of Economics*, 108(3):619–652, 1993.
- [3] S. Bentolila, C. Michelacci, and J. Suarez. Social contacts and occupational choice. *Economica*, 77(305):20–45, 2010.
- [4] S. Bowles, G.C. Loury, and R. Sethi. Group inequality. *Unpublished paper, Santa Fe Institute*, 2009.
- [5] I.S. Buhai and M.J. Leij. A social network analysis of occupational segregation. Technical report, Tinbergen Institute, 2006.
- [6] A. Calvó-Armengol. Job contact networks. *Journal of economic Theory*, 115(1):191–206, 2004.
- [7] A. Calvo-Armengol and M.O. Jackson. The effects of social networks on employment and inequality. *The American Economic Review*, 94(3):426–454, 2004.
- [8] A. Calvó-Armengol and M.O. Jackson. Networks in labor markets: Wage and employment dynamics and inequality. *Journal of economic theory*, 132(1):27–46, 2007.
- [9] A. Calvó-Armengol and Y. Zenou. Job matching, social network and word-of-mouth communication. *Journal of Urban Economics*, 57(3):500–522, 2005.
- [10] M.S. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.
- [11] B.C. Greenwald and J.E. Stiglitz. Externalities in economies with imperfect information and incomplete markets. *The quarterly journal of economics*, 101(2):229–264, 1986.
- [12] Y.C. Kim. Lifetime network externality and the dynamics of group inequality. 2009.
- [13] G.C. Loury. *The anatomy of racial inequality*. Harvard Univ Pr, 2002.
- [14] M. McPherson, L. Smith-Lovin, and J.M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pages 415–444, 2001.
- [15] J.D. Montgomery. Social networks and labor-market outcomes: Toward an economic analysis. *The American economic review*, 81(5):1408–1418, 1991.
- [16] T.C. Schelling. Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186, 1971.
- [17] W. Vickrey. Measuring marginal utility by reactions to risk. *Econometrica: Journal of the Econometric Society*, pages 319–333, 1945.

Appendix A: Proofs

Proof of Lemma 1

Since the problem is symmetric for workers of both groups, it suffices to show that the claim holds for workers of one group. Hence, without loss of generality, we restrict attention to workers belonging to group R . Consider first the case of a candidate equilibrium strategy profile $(\sigma^R(\theta), \sigma^G(\theta))_{\theta \in \Theta}$ such that $U_A^R(\theta) > U_B^R(\theta)$ for some $\theta \in \Theta$ and $U_A^R(\theta') < U_B^R(\theta')$ for some $\theta' \in \Theta$. Then, utility maximisation implies that we must have $c_A(\theta) - c_B(\theta) < \Delta E[w]^R < c_A(\theta') - c_B(\theta')$. Since $c_A(\theta)$ is increasing in θ and $c_B(\theta)$ is decreasing in θ , this implies that $\theta' > \theta$. Moreover, by continuity and monotonicity of the functions $c_A(\theta)$ and $c_B(\theta)$ it must be that $\exists! \bar{\theta} \in \Theta$ such that $U_A^R(\bar{\theta}) = U_B^R(\bar{\theta})$. Since $U_A^R(\theta)$ is decreasing in θ while $U_B^R(\theta)$ is increasing in θ , it must be that $U_A^R(\theta) < U_B^R(\theta) \forall \theta > \bar{\theta}$ while $U_A^R(\theta) > U_B^R(\theta) \forall \theta < \bar{\theta}$. This implies that in equilibrium we must have $\sigma^R(\theta) = 1 \forall \theta < \bar{\theta}$ and $\sigma^R(\theta) = 0 \forall \theta > \bar{\theta}$. But this is just the definition of a threshold strategy. The argument easily extends to the case of candidate equilibrium strategy profiles where either $U_A^R(\theta) > U_B^R(\theta) \forall \theta \in \Theta$, so that $\sigma^R(\theta) = 1 \forall \theta \in \Theta$, or $U_A^R(\theta) < U_B^R(\theta) \forall \theta \in \Theta$, so that $\sigma^R(\theta) = 0 \forall \theta \in \Theta$. \square

Proof of Lemma 2

The claim follows from the symmetry assumptions imposed on the cost and production functions. Since the social welfare function is additively separable, we prove the claim by demonstrating that the production function $f(\cdot)$ reaches its maximum value and the cost function $C(\cdot)$ reaches its minimum value on the line defined by $\hat{\theta}^R = 1 - \hat{\theta}^G$.

Without loss of generality, we can write $\hat{\theta}^R = c + \hat{\theta}^G$ where $c \in [-1, 1]$. Substituting this condition into the cost functions yields $C(\hat{\theta}^R, \hat{\theta}^G) = C(\hat{\theta}^G; c)$. Differentiating this function with respect to $\hat{\theta}^G$ and using the symmetry assumption $c_A(x) = c_B(1 - x)$ we obtain

$$\left(c_A(\hat{\theta}^G) - c_A(1 - \hat{\theta}^G) \right) + \left(c_A(c + \hat{\theta}^G) - c_A(1 - c - \hat{\theta}^G) \right) = 0$$

It is easily verified that this condition holds if $\hat{\theta}^G = (1 - c)/2$. To verify that this is indeed a unique minimum, notice that

$$\frac{dC^2(\hat{\theta}^G; c)}{d\hat{\theta}^G} = c'_A(\hat{\theta}^G) + c'_A(1 - \hat{\theta}^G) + c'_A(c + \hat{\theta}^G) + c'_A(1 - c - \hat{\theta}^G) > 0$$

This inequality follows from the assumption that the function $c_A(\cdot)$ is monotonically increasing, implying that the cost function $C(\hat{\theta}^G; c)$ is globally convex.

Turning now to the production function $f(\cdot)$, notice that the maximum must satisfy the following condition

$$E[z_A] \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_A} = E[z_B] \frac{\partial f(\tilde{l}_A, \tilde{l}_B)}{\partial \tilde{l}_B}$$

By Assumption 4 and the definition of \tilde{l}_ϕ as given by condition (4), this implies that

$$\frac{E[z_A]}{E[z_B]} = \frac{\tilde{l}_A}{\tilde{l}_B} \Rightarrow l_A = l_B$$

Using conditions (7) and (8), together with $\hat{\theta}^R = c + \hat{\theta}^G$, this last equality implies

$$\begin{aligned} q(\hat{\theta}^G + \alpha c)(\hat{\theta}^G + c) + q(\hat{\theta}^G + (1 - \alpha)c)\hat{\theta}^G = \\ q(1 - \hat{\theta}^G - \alpha c)(1 - \hat{\theta}^G - c) + q(1 - \hat{\theta}^G - (1 - \alpha)c)(1 - \hat{\theta}^G) \end{aligned}$$

It can be verified that this condition is again satisfied if $\hat{\theta}^G = (1 - c)/2$. Moreover, we know from the concavity and monotonicity conditions of Assumption 4 that $f(\cdot)$ is globally concave. The final step of the proof simply requires to notice that since $c \in [-1, 1]$, the locus of points defined by the condition $\hat{\theta}^G = (1 - c)/2$ is nothing more than the line $\hat{\theta}^R = 1 - \hat{\theta}^G$. \square

Proof of Proposition 1

Formally, the claim implies that there cannot be an equilibrium such that $\hat{\theta}^R = \hat{\theta}^G \in \{0, 1\}$. The claim follows directly from the essentiality condition of Assumption 3. We proceed to prove the claim by contradiction. In either of the two cases, it is clear from conditions (7) and (8) that $\exists \phi \in \{A, B\} : l_\phi = 0$. Given the wage equation specified by condition (11) and the essentiality condition of Assumption 4, this implies that $w_\phi \rightarrow \infty$. But then the optimality conditions of workers as stated in condition (13) must be violated, implying that these threshold values cannot constitute an equilibrium. \square

Proof of Proposition 2

We prove the claim by construction. We begin by showing that this candidate threshold profile implies that labour supply is constant and equal across occupations. Specifically, $l_A(1/2, 1/2) = l_B(1/2, 1/2) = q(1/2)$. Using this, the difference in expected wages $\Delta_\phi E[w]^X$ as given by condition (12) can be shown to be the same across groups and be equal to zero, so that $\Delta_\phi E[w]^X(1/2, 1/2) = \Delta E[w]^{X'}(1/2, 1/2) = 0$. By the symmetry of the cost functions we must have $c_A(1/2) = c_B(1/2)$. As this implies that the indifference condition $\Delta_\phi E[w]^X = c_A - c_B$ is satisfied, it confirms the claim that the threshold profile $\hat{\theta}^X = 1/2$ for all $X \in \{R, G\}$ constitutes an equilibrium. \square

Proof of Result 1

Formally, complete occupational segregation implies a threshold profile such that $(\hat{\theta}^R, \hat{\theta}^G) \in \{(1, 0), (0, 1)\}$. As the problem is symmetric, it suffices to show that one of these threshold profiles can constitute an equilibrium. We focus attention on the case where $(\hat{\theta}^R, \hat{\theta}^G) = (1, 0)$, and proceed to prove the claim by construction. In this case, the labour supply of workers across occupations is equal and given by $l_A = l_B = \alpha$. Wages across occupations are also equal and given by

$w_A = w_B = 1/2$. It follows that the difference in expected wages for workers belonging to group $X \in \{R, G\}$ equals

$$\Delta E[w]^R = \alpha - \frac{1}{2} \quad \text{and} \quad \Delta E[w]^G = \frac{1}{2} - \alpha$$

Plugging these equations into the optimality condition given by condition (15), we find that for the equilibrium to exist the following two inequalities must be satisfied

$$\begin{aligned} \hat{\theta}^R &= \frac{1 + k^{-1}(\alpha - \frac{1}{2})}{2} \geq 1 \\ \hat{\theta}^G &= \frac{1 + k^{-1}(\frac{1}{2} - \alpha)}{2} \leq 0 \end{aligned}$$

Solving this system of inequalities, we find that there exists an equilibrium that supports complete occupational segregation whenever $\alpha \geq \underline{\alpha}^{EQ} = \frac{1}{2} + k$ and $k < 1/2$. \square

Proof of Result 2

From Lemma 2, we know that we can restrict attention to symmetric allocations such that $\hat{\theta}^R = 1 - \hat{\theta}^G$. The social welfare function can then be rewritten as follows

$$W(\hat{\theta}; \alpha, k) = (\alpha\hat{\theta} + (1 - \alpha)(1 - \hat{\theta}))\hat{\theta} + (\alpha(1 - \hat{\theta}) + (1 - \alpha)\hat{\theta})(1 - \hat{\theta}) - 2k \left(\hat{\theta}(\hat{\theta} - 1) + \frac{1}{2} \right)$$

Twice differentiating the social welfare function with respect to $\hat{\theta}$ we obtain

$$\frac{d^2 W(\hat{\theta}; \alpha, k)}{d\hat{\theta}^2} = 8\alpha - 4k - 4 \leq 0$$

It follows that the social welfare function is either globally concave or globally convex depending, on the values of α and k . Global convexity of the social welfare function implies that complete occupational segregation is efficient. Thus, the cut-off value of the inbreeding bias that determines whether or not complete occupational segregation is efficient is equal to $\alpha \geq \alpha^{SW} = \frac{1}{2}(1 + k)$ for $k < 1$. It follows immediately that $\underline{\alpha}^{EQ} > \underline{\alpha}^{SW}$. \square

Proof of Result 3

The proof follows closely the proof of Result 1. Again, without loss of generality we focus on the case where $(\hat{\theta}^R, \hat{\theta}^G) = (1, 0)$. To simplify notation, we set $h = 1$. As in the case with constant productivity, complete occupational segregation implies that the labour supply of workers across occupations is equal so that $l_A = l_B = \alpha$. Moreover, average worker productivity across occupations is also equal and given by $E[z_A] = E[z_B] = 1/2$. This implies that wages across occupations are equal to $w_A = w_B = 1/4$. The difference in expected wages for workers belonging

to group $X \in \{R, G\}$ is thus given by

$$\Delta E[w]^R = \frac{\alpha}{2} - \frac{1}{4} \quad \text{and} \quad \Delta E[w]^G = \frac{1}{4} - \frac{\alpha}{4}$$

Plugging these equations into the optimality condition (15) and solving the system of inequalities, we find that the cut-off value of the inbreeding bias above which complete occupational segregation can be supported in equilibrium to be given by $\alpha \geq \underline{\alpha}^{EQ} = \frac{1}{2} + 2k$ and $k < 1/4$. \square

Proof of Result 4

Again, from Lemma 2 we can restrict attention to symmetric allocations such that $\hat{\theta}^R = 1 - \hat{\theta}^G$. The social welfare function is then given by

$$W(\hat{\theta}; \alpha, k) = \left(\frac{1}{2} + \hat{\theta} - \hat{\theta}^2 \right) ((\alpha\hat{\theta} + (1-\alpha)(1-\hat{\theta}))\hat{\theta} + (\alpha(1-\hat{\theta}) + (1-\alpha)\hat{\theta})(1-\hat{\theta})) - 2k \left(\hat{\theta}(\hat{\theta} - 1) + \frac{1}{2} \right)$$

Unfortunately, unlike the case when workers' productivity was constant, the social welfare function is no longer either globally convex or globally concave. To verify the result, we therefore compare the welfare value of the following two allocations $\hat{\theta} \in \{0, 1\}$ and $\hat{\theta} = 1/2$, and find restrictions on the inbreeding bias parameter such that the segregated allocation yields higher welfare than the mixed allocation. Note that this does not imply that the segregated allocation is welfare optimal, as there may be intermediate values of $\hat{\theta}$ (implying partial segregation) which yield higher welfare. Nonetheless, if welfare under complete mixing exceeds the welfare under complete occupational segregation, it necessarily follows that complete occupational segregation cannot be the efficient outcome. Hence, the cut-off value we derive should be interpreted as a lower bound. Comparing the welfare generated by these two allocations yields

$$W(\hat{\theta} \in \{0, 1\}; \alpha, k) = \frac{\alpha}{2} - k > W\left(\hat{\theta} = \frac{1}{2}; \alpha, k\right) = \frac{3}{8} - \frac{k}{2} \Rightarrow \alpha > \underline{\alpha}^{SW} = \frac{3}{4} + k$$

It is immediate to verify that $\underline{\alpha}^{SW} > \underline{\alpha}^{EQ}$, providing that $k < 1/4$. \square

