Mommy tracks and public policy: On self-fulfilling prophecies and gender gaps in promotion*

Kjell Erik Lommerud and Steinar Vagstad

Department of Economics, University of Bergen,

Hermann Fossgate 6, N-5007 Bergen, Norway

February 8, 2007

Abstract

Consider a model with two types of jobs. The profitability of promoting a worker to a fast-track job depends not only on his or her observable talent, but also on incontractible effort. We investigate whether self-fulfilling expectations may lead to higher promotion standards for women. If employers expect women to do more household work than men, thereby exerting less effort in their paid job, then women must be more talented to make promotion profitable. Moreover, specialization in the family will then result in women's doing most of the household work. Such self-fulfilling prophecies can be defeated: both affirmative action and family policy can make women spend more effort in the market, which can lead the economy to a non-discriminatory equilibrium. However, we find that it is unlikely that temporary policy can move the economy to a symmetric equilibrium: policy must be made permanent. Anti-discrimination policy need not enhance efficiency, and from a distribution viewpoint this is a policy with both winners and losers.

Keywords: self-fulfilling prophecies; gender discrimination; promotion

JEL numbers: D13; J16; J22; J71

*We thank Jim Albrecht, Geir Asheim, Patrick Francois, Robert Gibbons, William R. Johnson, Susan van Velzen and Susan Vroman for discussions. The comments of an anonymous referee have also been most helpful. We have benefitted from presenting this work at the Tinbergen Institute in Amsterdam and at Umeå University. The paper has partly been written during visits to Stanford University and Universitat Pompeu Fabra, and we thank these institutions for their hospitality. Correspondence: University of Bergen, Department of Economics, Fosswinckelsgate 6, N5007 Bergen, Norway. E-mail: kjell.lommerud@econ.uib.no / steinar.vagstad@econ.uib.no

1 Introduction

Choice of education and working career are among the most important economic decisions a person makes. Gender and family status have a major impact on these decisions in practically every society. True, the nature of the sexual division of labor has undergone vast changes over the past decades. The traditional pattern of specialization, with a bread-winning father and a mother working solely at home, is fading in importance. The sexual division of labor in many societies now takes a more subtle form: both men and women work, but women choose working arrangements that are compatible with having primary responsibility for children.[1]

The traditional economic explanation of the division of labor within a family is that specialization occurs according to comparative advantage. However, as regards "natural endowment," comparative advantages of men and women in market work and household production appear to be of limited importance today. The natural advantage of women in taking care of children should apply only to infants, and physical strength has become less and less important in most jobs. So why are there such big differences in men and women's labor market adaptions, given small natural comparative advantages? Of course, with constant rewards in the two lines of work, the least comparative advantage leads to full specialization. But the least reduction in marginal productivity as one spends more time doing a task breaks this result. However, Gary Becker (1991) suggests that comparative advantages might evolve dynamically. For example, small initial comparative advantages might lead to some specialization, and due to learning-by-doing, the comparative advantages grow more important over time.[2]

In this paper we offer a complementary explanation of why a marked sexual division of labor might continue to exist even when comparative advantages are of little importance. Our starting point intuition is the following: suppose there are two types of jobs, "fast track" and "slow track" jobs. A worker is placed on the fast track if a fixed investment, paid by the firm, is undertaken. Future wages are incontractible at the time of investment, and assumed to be set ex post through bargaining. The firm then gets a fraction of the value produced by any given employee. This value in turn depends on the (known) ability of the person, but also on his or her future "effort." "Effort" is to be interpreted broadly —

[1] For some US evidence, see, e.g., Fuchs (1989) and Hersch (1991). Relevant findings from a British setting are found in Joshi (1989).

[2] Parallel thoughts abound in the theory of international trade, see e.g. Krugman (1987).

for instance to include the willingness to work irregular hours and not take long parental leaves. The employers only place those employees on the fast track whose expected output is so large that the investment cost is recouped. Suppose that women have traditionally been the ones to assume primary responsibility for child care and that such responsibilities make it more costly to exert effort in the labor market. If employers expect less effort from women than from men of comparable talent also in the future, then women, more often than men, will be put on the labor market's slow track. The decision concerning who will actually assume primary responsibility for the children and who is allowed to concentrate on the outside labor market is then taken within the family. Most family models predict some degree of specialization according to compara tive advantage, and perhaps the most important source of a comparative advantage in household production is a low outside wage.[3] Can it be that employers' beliefs in this way turn into self-fulfilling prophecies? This would mean that the slow track in the labor market would predominantly be a "mommy track."

The present paper formalizes this story. We associate self-fulfilling prophecies with situations where the equilibrium outcome is asymmetric even when agents are symmetric. Whenever there is an asymmetric equilibrium in which the female assumes primary responsibility for household production, there will be another asymmetric equilibrium with reversed gender roles. Self-fulfilling beliefs can make one of these equilibria focal.

The paper establishes conditions under which discrimination prevails in a imagined world in which men and women are made identical in all economic respects (which in our setup basically means that there are no gender differences in preferences or skills). Two conclusions can be drawn: (i) The symmetric model sometimes have only asymmetric equilibria, meaning that self-fulfilling prophecies discrimination can be an equilibrium phenomenon. But note that this is the case only under given conditions. (ii)We will argue below that it is "unlikely" that a stable nondiscriminatory equilibrium coexists with asymmetric equilibria. Successful anti-discrimination policy will then have to sustain permanently a situation that is not a stable equilibrium – rather than shifting the economy from a discriminatory equilibrium to a nondiscriminatory one. The latter would arguably

[3] Konrad and Lommerud (1995) show that specialization according to comparative advantages can also characterize decisions in a family where family members non-cooperatively and non-altruistically determine their own time use. Related work can be found in Konrad and Lommerud (2000), Konrad, Künemund, Lommerud and Robledo (2002) and Vagstad (2001).

be the easier task.

The general idea that there might be "feedback" from employers' expectations to family members' individual behavior is well understood.[4] In order for this to be an equilibrium phenomenon, however, the "feedback" mechanism must be such that employer expectations turn out to be exactly correct. As mentioned, this can be the case, but there are also parameter sets where the unique equilibrium is a symmetric one. We will also ague below that it is "unlikely" that a stable symmetric equilibrium coexists with equilibrium discrimination. We show that the existence of a stable symmetric equilibrium depends on two things: the effect of own talent on output being "large" relative to the expected effect of whether or not one's spouse is promoted to the fast track, and the distribution of talent in the population.

But what does this imply for policy? Can and should this type of asymmetric equilibrium be countered by policy measures? What are the effects of various policy instruments? Does policy need to be in place only temporarily until the economy is established in a new, symmetric equilibrium in which men and women are promoted according to the same rules? Or will the economy revert to the old, asymmetric equilibrium as soon as the policy measures are lifted? The attempt to answer these questions is the focus of the remainder of the paper.

A first result of our study is that it is quite likely that anti-discrimination policy needs to be permanent. This restates our finding that when the symmetric equilibrium is stable, it is probable that this is a unique equilibrium. A further implication is that if the economy is initially in an asymmetric equilibrium, then there is no stable symmetric equilibrium.[5] Policy can force the economy to an outcome that is not a stable equilibrium, but as soon as the policy in question is lifted, such a situation cannot be upheld.

We argue that many types of policy can, indeed, break a discrimination equilibrium. Affirmative action and family policy are but two examples of policies that might work. Turning to welfare issues, however, it turns out that anti-discrimination policy is a policy with both winners and losers. Perhaps paradoxically, discriminatory promotion standards make the income distribution among families more even. If fair promotion standards are introduced, the likelihood that some families will have two promoted workers and some

[4] See, for example, Blau and Ferber (1992).

[5] This finding is in contrast to Coate and Loury's (1993) work on self-fulfilling prophecies and discrimination. We shall return to this point later in the paper.

families having none will increase. This could suggest that anti-discrimination policies should be complemented by redistributive measures. Even if these distributional effects are neutralized, it is not certain that anti-discrimination measures will increase welfare. On the positive side, fair promotion means that the talent capital in the society is put to better use. However, as someone has to do the housework in a family, it can be wasteful to promote both parties in a couple. Discrimination can be seen as a coordination device that reduces the probability of wasteful promotion of both partners. A further policy problem that we point out is that anti-discrimination policy that is mistakenly applied when the economy is in a symmetric equilibrium can in fact create discriminatory outcomes.

These caveats should not cloud what our analysis reveals: that discriminatory treatment of women can arise even when men and women are completely equal in all economic respects; that policy can effectuate symmetric treatment of the sexes; and that there are cases where anti-discrimination policy is welfare-improving. However, the correct application of policy measures requires a careful analysis of the situation, and may easily go wrong.

Other papers have also examined models in which women face tougher promotion standards than men. A well-known paper with such a message is Lazear and Rosen (1990). Precisely as in our model, these authors assume that a relation-specific investment is needed in order to establish a worker on the fast track, and employers set tougher promotion standards for women because there is a higher probability that a woman will leave a fast-track job in order to perform nonmarket activities.[6] The key factor behind the Lazear-Rosen result, however, is their assumption that although men and women have the same distribution of labor market ability, women have superior ability in nonmarket activities. An important objective of the present paper is to show that the same empirical predictions can arise within a model where men and women are completely equal in all economic respects. The policy implications of the two models are also different. The higher promotion standard for women in the Lazear-Rosen model is socially efficient — in the present paper this may or may not be the case.

More closely related to our paper is a series of papers that discuss self-fulfilling prophecies and promotion standards in various model settings, i.e. Coate and Loury (1993),

[6] Other papers that study discrimination and affirmative action in a setting with promotions and/or investments in education include Renes and Ridder (1995), De Fraja (2002); Booth, Francesconi and Frank (2003) and Austen-Smith and Wallerstein (2003).

Francois (1998), Engineer and Welling (1999) and Moro and Norman (2003, 2004). Coate and Loury's (1993) seminal paper builds on Arrow's (1973) version of the statistical discrimination model. The essence of Arrow's model is that employers' negative stereotypes about a group can in themselves weaken incentives for acquiring skills, and in this way become self-fulfilling prophecies.[7] Coate and Loury (1993) redesign Arrow's model to form a model of job discrimination rather than wage discrimination within one given job category. There are many important differences between the Coate-Loury model and our own. Coate and Loury assume that productivity can only be observed through a noisy signal. The coordination problem in their model is between employers making job assignments and workers investing in accruing productive ability. In our model productivity can be observed, and promotion decisions and human capital investment decisions are made by the same party (the employers). Self-fulfilling prophecies discrimination can still occur in our model, and it is driven solely by the coordination problem that can arise when there are substantial specialization gains within a family and long-term career decisions must be made before the identity of a future spouse is known for certain.

Perhaps the most important claim that Coate and Loury make about policy is that temporary affirmative action may sometimes move the economy from a discriminatory to a non-discriminatory outcome. When black workers know they have a higher promotion probability, they will invest in productivity, and when employers discover that black workers are more productive than they originally thought they were, they will revise their beliefs. When beliefs have changed, there is no longer need for the original policy.

In our setting, it is far more likely that affirmative action must be permanent. This, of course, makes anti-discrimination policy much less attractive. The Coate-Loury model is tailored to a situation with racial discrimination. There are no links between the situations for black and white workers, so in principle black workers could have been studied in isolation. In such a case, there are few constraints on assumptions about beliefs, which makes it rather easy to construct a case with effective temporary anti-discrimination policy. Our case is different, however. Men and women tend to live together in families and have joint children. As we know that someone must take care of the children, any belief we might have about how much effort women will typically exert in the labor market is logically linked to what one must rationally believe about men's behavior. We will show that this

[7] Other contributions to the statistical discrimination literature include Phelps (1972), Aigner and Cain (1977), Lundberg and Startz (1983), Milgrom and Oster (1987), Lang (1990), and Norman (2004).

tends to produce situations in which the economy reverts to an asymmetric equilibrium once policies that encourage women's careers are lifted.[8] Getting rid of discrimination in the Coate-Loury setting is always welfare-improving. In a family context, enforcing gender equality when there are substantial gains from specialization turns out, as already suggested, to be a much more complicated issue from a welfare perspective.[9],[10]

Francois (1998) and Engineer and Welling (1999) present models of self-fulfilling prophecies where family interactions and gains from specialization create coordination problems. In Francois' paper, efficiency wage setting is assumed to be predominant in the "good jobs" labor market, and discrimination takes the form of being shut out from the good jobs sector. Francois assumes that all workers are equally talented, which, in our view, is the main limitation of his model. For example, there is no waste of talent when men rather than women are promoted: this seems to miss an important aspect of the problem. Engineer and Welling study a setting where parents prepare their children either for market work or for household production at a time when the future spouse and his or her talent or training is still unknown. Should parents train children according to gender roles or according to their talents? Unlike Francois, Engineer and Welling allow workers to be heterogenous in talents, but only in a discrete way: children either have a talent exclusively for market work or exclusively for homework. There are many differences between Francois and Engineer-Welling and our own model. One major difference is that we allow

[8]Coate and Loury also stress that affirmative action can be "patronizing". By this, they mean that black workers are discouraged from investing in human capital because of lax promotion standards. This does not occur in the present model. Employers control both investment and promotion decisions, and no one is put on the economy's fast-track without having the proper skills.

[9]As already mentioned, in the Coate-Loury model, employers only receive a noisy signal about workers' productivities. In our model a worker's talent is known at the time of hiring. Instead we focus on the interplay between decision making in the labor market and within the family. The self-fulfilling prophecy in Coate-Loury concerns a belief about workers' current productivity. The self-fulfilling prophecy in our model concerns expectations about whom a given worker will marry and what future family decision making will imply for how much effort this worker will exert. This distinction is important for policy. The Coate-Loury problem can in principle be solved by certification. By this we mean a policy where the government spends the necessary resources to ascertain a worker's qualifications and then announces the result. Certification can play no role in our model: The problem for female workers is not that they cannot credibly convey their present productivity but that they cannot credibly commit their own future choices.

[10]Moro and Norman (2001a, 2001b) contain a general equilibrium generalization of the Coate-Loury model, where the dependence of wage formation on policy is investigated. Welfare results are less clear-cut here than in the original Coate-Loury setting.

workers to have varying talents in a continuous way. Various discreteness assumptions can, in particular, influence the analysis of the type of stable equilibria that exist, which, in turn, is important for policy analysis.[11]

The remainder of the paper is organized as follows. In Section 2 we present our model. Section 3 contains analyses of different means of battling self-fulfilling-prophecy discrimination; welfare issues are addressed in section 4, and some concluding remarks are presented in Section 5.

2 The model

In what follows we develop a model which links household and labor market decisions. A household is a pair of workers of opposite sexes. A worker receives utility from consuming a household good and a market good. The household good is a public good for members of the family. The market good is bought with money earned in the labor market, while the household good is produced in the family.

There are two types of jobs. First there are "ordinary," perhaps piece-rate jobs in which a worker's payment corresponds to the value of his or her marginal product. Second, there are career or "fast-track" jobs. In order to install a worker in a fast-track job, the employer must sink a relation-specific investment. One example of such an investment is on-the-job training in the beginning of the employer-employee relationship. After investing, the employer receives a share of the workers' marginal product, while the worker retains the remaining share. Noteworthly, the main conclusions in this paper would survive if we instead assumed that the employees undertake the investments, as long as firms and employees share the same (lack of) information about employee talent and future spouse characteristics.

The employer's net returns from installing a worker i in a fast-track job will be denoted π, and depend on the worker's ability or talent, b_i, and on his or her on-the-job effort, e_i, according to the following function:

$$\pi = b_i + e_i - I \tag{1}$$

[11] Both Francois and Engineer-Welling lack a systematic discussion of whether policy measures can be transitory or need to be permanent, but their work seems implicitly to suggest temporary policies to a much larger extent than the present analysis.

where I is the investment cost.[12][13] Effort is to be interpreted broadly, and its measurement may include factors such as hours spent on a task, work intensity, and willingness to work extra hours when needed.

The timing of decisions is as follows. First, each worker i's talent is drawn from a twice differenciable common distribution F on an interval $[\underline{b}, \overline{b}]$. The density function is denoted f. Second, employers observe the talent of the workers they are facing and decide simultaneously which workers to put on the fast-track and which not. Third, male and female workers randomly match to form families.[14] Finally, the family members bargain over how each member's total effort is divided between the paid job and household production. To simplify the exposition we assume that effort is either high ($\overline{e}$) or low ($\underline{e} < \overline{e}$), and depending only on whether one's spouse is promoted or not. (A more general formulation is found in Lommerud and Vagstad (2000)). That is, person i's equilibrium job effort, given that he or she is installed in fast-track employment, is given by $e_i = \underline{e}$ if the spouse has a fast-track job and $e_i = \overline{e}$ if not.

Our analysis will rest on the following set of assumptions.

A1. The probability that one employer hires both members of a family is negligible.

A2. It always pays to promote the most talented, and it never pays to promote the least talented: $\underline{b}+\overline{e} < I \leq \overline{b}+\underline{e}$.

A3. Any worker that is offered a fast-track job prefers this job to an ordinary job.

A1 is clearly a context-dependent assumption, which we claim is a reasonable approximation in modern market economies. A2 holds if we employ a talent space that encompasses (possibly very unlikely) persons with extremely low or high talent. Assumption A3 rests on talent being more productive on the fast-track than on the slow-track and that the worker get some of the net gains from promotion.

Consider the investment decision of an employer facing a worker with talent b_i. By eq. (1) the investment decision must be monotonic in a worker's talent: if a worker is

[12] Implicit in equation (1) is an assumption that there are "enough" fast-track jobs in the sense that a given worker's possible promotion only depends on own talent and effort. Real world promotion decisions are often interlinked because there are a given number of promotion slots. Promotion decisions will then be based not only on absolute expected performance (as underlying equation (1)), but also on relative performance. This complication does not affect our key point notably, however: discrimination in promotion will still arise if and only if expected performance of workers differs according to their gender.

[13] Total productivity then is $p = \pi + w(b_i, e_i)$, where $w(b_i, e_i)$ is the fast-track salary. We assume for simplicity that $w(b_i, e_i) = \alpha\,(b_i + e_i)$, where $\alpha \in (0, 1)$ is the workers' share of gross surplus.

[14] Non-random matching – assortative matching – is discussed below.

on the fast-track, so is any worker of the same gender with higher talent. Therefore, any rational investment strategy can be represented by a cutoff talent $\widehat{b}_i$ with the property that the worker is invested in if and only if $b_i \geq \widehat{b}_i$. Since the investment decision is made before learning the spouse's talent and career, the investment decision must be based upon expected effort, expectation taken over the possible talents and careers for the spouse. The employer's belief about the spouse's talent is described by the distribution function F, while the employer's rational belief about the spouse's career can be represented by the cutoff talent $\widehat{b}_j$. Expected returns, denoted Π, from investing in a worker with talent b_i can be written:

$$\Pi(b_i, \widehat{b}_j) \equiv b_i + E[e_i] = b_i + F(\widehat{b}_j)\overline{e} + [1 - F(\widehat{b}_j)]\underline{e} - I. \tag{2}$$

The employer will invest in a worker if and only if $\Pi(b_i, \widehat{b}_j)$ is nonnegative. Therefore, $\widehat{b}_i$ must satisfy the following equation (by continuity and assumption A2 the equation has at least one solution):

$$\Pi(\widehat{b}_i, \widehat{b}_j) \equiv \widehat{b}_i + F(\widehat{b}_j)\overline{e} + [1 - F(\widehat{b}_j)]\underline{e} - I = 0. \tag{3}$$

Equation (3) implicitly defines the best-response cutoff talent for the employer of worker i as a function, denoted B_i, of the equilibrium cutoff talent of the other gender; $\widehat{b}_i = B_i(\widehat{b}_j)$for any $\widehat{b}_j \in [\underline{b}, \overline{b}]$. $B_i()$ can be interpreted as a reaction function in a two-player simultaneous-move game in which the employer of worker i controls $\widehat{b}_i$ and the employer of worker j controls $\widehat{b}_j$. The reaction function for male promotion and that for female promotion are mirror images of each other because of the assumed symmetry between the sexes.

For a pair of strategies $(\widehat{b}_i, \widehat{b}_j)$ to form an equilibrium of this game, the strategies must be a point at which the two reaction curves cross. That is, $\widehat{b}_i = B_i(\widehat{b}_j)$ and $\widehat{b}_j = B_j(\widehat{b}_i)$. Moreover, for an equilibrium to be stable, the absolute value of the product of the derivatives of the two reaction functions must not exceed unity, that is, $\left|B_i'(\widehat{b}_j)B_j'(\widehat{b}_i)\right| \leq 1$. This draws our attention to the slope of the reaction functions. Rewriting eq. (3) yields:

$$\widehat{b}_i = B_i(\widehat{b}_j) \equiv I - \underline{e} - F(\widehat{b}_j)\left(\overline{e} - \underline{e}\right) \tag{4}$$

with slope given by

$$B_i'(\widehat{b}_j) \equiv -f(\widehat{b}_j)\,(\overline{e} - \underline{e}) \tag{5}$$

The negative slope confirms our intuition that if, for some reason, more women are put on fast-track jobs, then, in equilibrium, fewer men will be established in such jobs.

Our main result, which deals with the conditions for the existence of stable non-discriminatory equilibria, is as follows (the proof follows from standard game theory and is omitted here, but can be found in Lommerud and Vagstad (2000)):

Proposition 1 *There is a unique fixed point of the reaction function, denoted b^*. A symmetric equilibrium has to satisfy $\widehat{b}_M = \widehat{b}_F = b^*$ and such a stable equilibrium exists if and only if $|B_i'(b^*)| \leq 1$.*

While there is nothing in our assumptions that rules out multiple equilibria, Proposition 1 says that there cannot possibly be more than one symmetric equilibrium. Moreover, the proposition also establishes a simple check of whether a stable symmetric equilibrium exists at all: all we have to do is to calculate the slope of the reaction function at its fixed point. Below we discuss in more detail the conditions under which the symmetric equilibrium exists. Before we do that, however, we will take a closer look at the shape of the reaction function outside its fixed point, partly to establish conditions for the symmetric equilibrium to be the unique equilibrium whenever it exists, and partly to identify asymmetric equilibria, in which men and women face different promotion standards.

First note that the slope of B_i is proportional to the density $f(\widehat{b}_j)$. The intuition is that changing the cutoff talent of, say, men, will have a larger effect on employers' decisions about female workers if the change is likely to affect a given female's spouse, that is, if the probability density at and around the cutoff level is high. This effect will be dubbed the *spouse density effect*. Second, the slope is also proportional to $(\overline{e} - \underline{e})$: that is, the reduction of a worker's effort when his or her spouse is being invested in. We call this effect the *effort penalty effect*.

Since the effort penalty $(\overline{e} - \underline{e})$ is constant by assumption, variations in the the slope of the reaction function depends only on the density of the spouse's talent. A more general formulation of the model (see Lommerud and Vagstad (2000)) allows for many interpretations of what effort or talent really is, and each of them can be measured along different scales.[15] With our specific assumptions, however, profit is linear in a worker's

[15] Clearly, if the talent we are talking about is measured by some sort of standardized test, it is normally

talent in any equilibrium, and the relevant empirical talent distribution should then mirror the economy-wide distribution of wages.[16] Empirical evidence on wage distributions suggests bell-shapes (with certain spikes that have to do with institutional conditions, not with productivity), see e.g. Donald, Green and Paarsch (2000) and the references therein. Consequently, in what follows we will assume that the distribution of talent is described by a bell-shaped density function.

What remains is straightforward. If the talent distribution is described by a bell-shaped graph, the reaction function itself must be relatively flat at both ends where the density is low, and steeper in the middle where the density is higher. A typical reaction function is shown in Fig. 1 below.

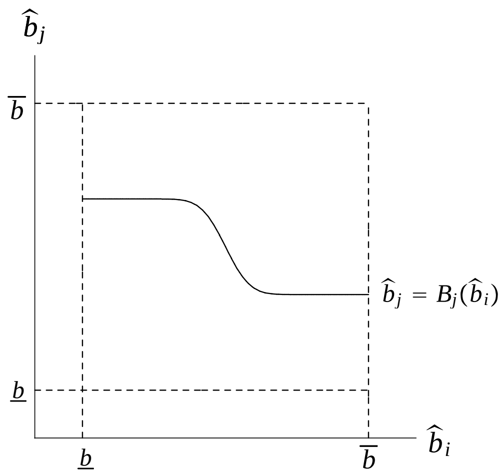

Fig. 1: Typical reaction function

The natural next step is to draw both players' reaction functions in the same diagram, in order to identify the equilibria of the model. Two possible outcomes are shown in Fig. 2 below. By assumption A2 and continuity of the reaction functions, the reaction functions must cross an odd number of times.[17] If they cross once, we get a unique symmetric

distributed by definition. What is important, however, is how talent combine with effort to produce net returns for the employers. For a given measurement of net returns, our choice of talent measure will affect our choice of effort measure or the net returns function, or both.

[16] With a competitive labor market the correspondence will be perfect. But also with labor market imperfections, there will typically be substantial correlation between productivity and wages. For instance, any bargaining models based on rent-sharing will yield perfect correlation.

[17] It is also possible that the two curves coincide for some interval of talents, giving rise to a continuum

equilibrium. If they cross three times, there are two stable equilibria, both asymmetric (cf. Fig. 2). If they cross five times there are three stable equilibria (one symmetric and two asymmetric). With seven crossing points, there are four stable equilibria (all asymmetric), and so on. In short, much can happen.

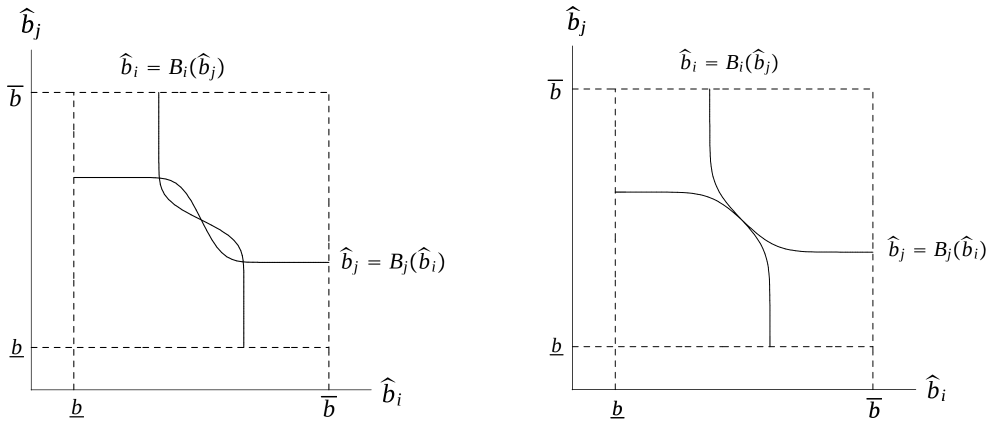

Fig. 2: Asymmetric (left-hand side) and symmetric (right-hand side) equilibria

From a policy viewpoint, one would like to know whether anti-discrimination policy must be used only in a transition period until symmetry is achieved, or whether policy must be made permanent. Presumably, permanent anti-discrimination will be more costly to society. If a stable symmetric equilibrium coexists with asymmetric equilibria, policy is only needed to shift the economy from one equilibrium to another; we dub this an "announcement" policy. For an announcement policy to work, the reaction functions must cross at least five times. With inverse-S shaped (symmetric) reaction functions and reaction functions being mirror images of each other, it is not impossible (though quite difficult) to construct examples of reaction functions crossing each other more than three times. (An example is found in Fig. x. in the appendix.) It turns out that the following assumption is sufficient to preclude the existence of situations with more than 3 crossings:

A4. The talent density is log-concave.

Noteworthly, many commonly used distributions satisfy this property. (See Bagnoli and Bergstrom, 1989, for a survey of the log-concavity properties of many common distributions.) This includes the normal, Chi and Chi-squared distributions.[18] Also note that

of equilibria. This is, however, a knife-edge case without general interest.

[18] Other distributions do not satisfy the requirement for all parameters or for all values. However, the Beta and Weibull distributions have log-concave density for parameters yielding single-peaked density functions, and the Student's t-distribution and the log-normal distribution have log-concave density for variables in

even if the talent distribution should not satisfy the assumption of log-concave density, it would still take a lot to produce situations with more than 3 crossings. (Fig. A.1 in the appendix is drawn using a talent distribution that is a carefully designed linear combination of a uniform distribution and a normal distribution. This produces a density that is sufficiently convex in the relevant region to break the log-concavity condition.)

Given assumption A4, the following result holds (proof in the appendix):

Proposition 2 *The reaction functions cross at most three times.*

With either one symmetric stable equilibrium or two asymmetric ones, policy could still play a role, for example in shifting reaction curves so that we move from a situation with two stable equilibria – both asymmetric – to a situation with only one – a symmetric one. This would call for permanent measures.

The crucial question now is whether the economy is in a symmetric, non-discriminatory equilibrium or in an asymmetric, discriminatory one. The answer is found in Proposition 1. Since, by symmetry, the two reaction functions have to intersect at the 45-degree line, that is, the fixed-point, we need only check the slope at this point. If $|B_i'(b^*)| \leq 1$, there is only one equilibrium: the symmetric one. If not, there are two asymmetric equilibria.

While Proposition 1 is a rather technical characterization of situations leading to or precluding the symmetric equilibrium, its economic content is not fully explained. Suppose that we are in an asymmetric equilibrium in which females face a tougher promotion standard than do their male counterparts: $\widehat{b}_F > \widehat{b}_M$. From Proposition 1 we have that this can be an equilibrium if and only if the reaction function is relatively steep when evaluated at its fixed point. As already established, this requires the density and career effects to be relatively strong. By assumption, the density effect is the determinant of the slope of the reaction function. Therefore, the discriminatory equilibrium requires the fixed point to be somewhere near the peak of the talent distribution.[19]

Moreover, for the career effect to be strong there must be substantial specialization according to career in the family. Promoting a man with as little talent as $\widehat{b}_M$ can be part of an equilibrium if and only if rational beliefs suggest that his effort will be relatively high, despite his lack of talent. With substantial specialization according to career it may

the relevant interval (a sufficiently large interval around the distribution peak, loosely speaking).

[19] Note that assumption A2 and the negatively sloped reaction curves imply that the average slope is less than unity. Consequently, if there is little variation in the slope, there can only be one equilibrium: the symmetric one.

actually be rational to believe that his effort will be high. Since with a relatively high probability (that is, $F(\widehat{b}_F)$) his spouse will have a slow-track job, he can be expected to have a comparative advantage in paid work, with a spouse that has the main responsibility for household production.

Conversely, not promoting a woman with talent almost as high as $\widehat{b}_F$ can only be optimal if her effort is likely to be relatively low. This it may well be, since with a relatively low probability ($F(\widehat{b}_M)$) this woman will have a man in a slow-track job, the woman will therefore not have much comparative advantage in the paid job, there will be little specialization and her employer will correctly anticipate that even if this woman is promoted, she will have to carry out her part of the child care and other housework.

Before turning to policy analysis we will discuss two natural extensions of the model. The first, assortative matching – where high-talent men tend to marry high-talent women – may arise e.g. if people sort into schooling according to talent and prospective spouses meet at school. This will narrow the distribution of spouse's talent conditioned on own talent: On seeing a talented woman, her employer will rationally expect her to have a talented husband, too. This essentially makes selection of the right familiy member to invest in less important and coordination (ensuring that one and only one family member is promoted) more important, thereby making discrimination a more likely equilibrium outcome. To formalize this line of reasoning, suppose (b_i, b_j) are drawn jointly from a bivariate normal distribution with correlation coefficient ρ_t. Then we know that the marginal distributions as well as the conditional distributions are normal (see e.g. Judge et. al., 1988, Section 2.5.6). Moreover, if the marginal distributions have common mean μ and common variance σ^2, the conditional distribution of e.g. b_i given b_j has mean $\rho_t b_j + (1 - \rho_t)\mu$ and variance $\sigma^2(1 - \rho_t^2)$. That is, the conditional distribution exhibit a smaller variance than the marginal distribution, and the mean is a weighted average of the unconditional mean and the value of the conditioning variable.

Consequently, if we wish to find out whether a fixed point at the peak of the talent distribution is an equilibrium, the slope of the relevant reaction function can be written:

$$\frac{d\widehat{b}_i}{d\widehat{b}_j} = -f(\widehat{b}_j | b_i = \widehat{b}_i)\,(\overline{e} - \underline{e}) = -\frac{(\overline{e} - \underline{e})}{\sigma\sqrt{2\pi}\sqrt{1-\rho_t^2}} = -\frac{f(\widehat{b}_j)\,(\overline{e} - \underline{e})}{\sqrt{1-\rho_t^2}} \tag{6}$$

Consequently, assortative mating yields steeper reaction curves and therefore increases the likelihood of asymmetric equilibria. Moreover, correlation yields a "moving peak" of the distribution; we note that the conditional mean of b_i is affected by b_j. In particular, if ρ_t is close to 1, then the conditional mean of b_i approaches b_j, while the slope of the reaction function approaches minus infinity, implying that with perfect correlation there is no symmetric equilibrium. (This should not be surprising, as e.g. Francois' model yields asymmetric equilibria based on an assumption of equal talent, which can be interpreted as perfect assortative mating.)

Second, it can be argued that the correlation between observed talent and actual performance in a given job is less than perfect. Intuitively, this should make employers put less weight on a worker's talent and thereby more weight on the worker's gender. To formalize this argument, let true productivity be denoted by β_i, and assume that b_i and β_i are drawn jointly from a bivariate normal distribution with common mean μ, common variance σ^2 and correlation coefficient ρ. Then the conditional distribution of β_i given b_i has mean $\rho b_i + (1-\rho)\mu$ and variance $\sigma^2(1-\rho^2)$. Consequently, after rewriting eq. (3) the cutoff talent must satisfy

$$\Pi(\widehat{b}_i, \widehat{b}_j) \equiv \left(\rho\widehat{b}_i + (1-\rho)\mu\right) + F(\widehat{b}_j)\overline{e} + [1 - F(\widehat{b}_j)]\underline{e} - I = 0. \tag{7}$$

which yields $\frac{d\widehat{b}_i}{d\widehat{b}_j} = -f(\widehat{b}_j)\left(\overline{e} - \underline{e}\right)/\rho$. Consequently, imperfect correlation between talent and true productivity also contributes to making the reaction curves steeper.

3 Policy

We have presented a model of discrimination that focuses on the interaction between labor market and intra-family decision making. In this section we discuss some policy options available if one wants to ensure that men and women meet the same promotion standards. At this point, we are only concerned with whether or not a policy works. Discussions of efficiency and distribution effects from anti-discrimination policy are left to Section 4. Policy can, in principle, either intervene in the labor market, in an attempt to influence the promotion standards set by employers, or in the family arena, influencing how much effort men and women spend in household production. Also, there is a choice between direct regulation and tax-subsidy schemes. Here we will briefly discuss "affirmative action," which is an example of direct regulation in the labor market, and, more fully, "family policy," which will be taken to mean subsidized provision of a close substitute to the household good (a price-based policy aimed at the family sphere). The most obvious omission as regards possible policy instruments lies in the lack of a discussion of the effect of labor income taxation on employer and family behavior in a setting such as this.

Note further that a high interdependence of a married couple's effort decisions is one factor that makes discriminatory outcomes more likely. This means that the income tax system can be an important factor in creating discrimination, in that tax wedges in many countries make it too expensive to rely on purchased market substitutes to own production in the family. This again should imply that self-fulfilling prophecies about women should be more predominant in countries with little use of outside services than in countries that make more use of such market substitutes, even though time use on own children of course is important everywhere.

A caveat is warranted. The following analysis is based on the assumption that there are no inherent differences between the productivities of the sexes. Given this "feminist" premise, we have wanted to show that self-fulfilling prophecies only exist in some given circumstances, and that even if they do, the application of policy is quite problematic. It would be interesting to study public policy in a setting where there are differences between the sexes and self-fulfilling feedback mechanisms operate, but we leave this type of question open at this stage.

3.1 Affirmative action

The term "affirmative action" is here taken to mean any policy that directly intervenes in the labor market to increase the chance that a woman is promoted to the fast track. Clearly, if talent were observed not only by employers but also by the government, a policy of strict affirmative action could end discrimination: by forcing employers to use the same promotion standard for both genders, the equilibrium is forced to lie on the 45-degree line, and, consequently, the fixed point of the reaction function can be supported as the unique equilibrium point, even if the reaction curves cross from the "wrong" side.[20]

In the real world it is more reasonable to assume that the government has less information about different workers' talents than the workers themselves and their employers. Suppose only "substantially" different promotion standards can be detected and punished. Then the government cannot force firms to choose promotion standards that are located on the 45-degree line, they can only push employers to set promotion standards within a

[20] Clearly, under full information a government can, in principle, enforce any set of promotion standards, regardless of whether these standards constitute an equilibrium of the original game.

more or less narrow band centered around the 45-degree line. Depending on how accurately discrimination can be detected and punished, discrimination may be reduced, but never completely eliminated.[21]

Affirmative action shares many of the same problems as direct regulation in many other contexts. In real life there is not one single promotion standard for men and one for women. Direct regulation requires quite detailed surveillance of millions of decisions made by employers in the private and public sectors. Moreover, speaking as amateur sociologists, we think that a policy that intervenes directly in the promotion processes will often have very identifiable losers – and may therefore meet bitter opposition. Policies that work more indirectly by improving women's qualifications and willingness to exert labor market effort may to a lesser degree upset motivation and workplace social relations.

3.2 Family policy

What drives employers to promote less talented men before more talented women in our context is their expectations of substantial specialization between housework and labor market effort. This suggests that discrimination can be reduced or even completely eliminated if we can reduce the interdependence between someone's equilibrium job effort and the career of his or her spouse. Provision of subsidized daycare is a policy candidate in this respect; with access to cheap, high-quality childcare, a talented woman can be expected to put forth substantial effort even if her spouse has a demanding career job. In some countries, tax breaks for domestic help have been suggested, and the effects would be parallel.

Here, the policy instrument we consider is offering a fixed amount of free, high-quality, reliable daycare for every child below school age. We assume that the cost of this policy is covered by lump-sum taxation of the family itself, so the family policy does not imply that extra resources are transferred to the families in question. If we, for the moment, treat the family members' careers as given, the effects of such a policy will be to increase

[21] One could imagine implementing the non-discrimination equilibrium simply by mandating an equal number of men and women in fast-track jobs. However, unless economic differences between the sexes are eliminated, this would typically entail different promotion standards. Moreover, even without economic gender differences, such an implementation would be problematic: even big firms hire a small number of workers at a time, and if you hire two workers, the two best might easily be of the same sex. Therefore, talent uncertainty makes mandating an equal number of men and women in fast-track jobs imply time-variant as well as firm-specific gender differences in promotion standards.

the family members' job effort. Publicly provided child care is a close substitute to child care provided within the family, and almost any specification of a family model then predicts that the public provision, at least to some extent, crowds out the family's own provision. We assume that both family members reduce their effort in home production and increase at-the-job effort. In reaction function terms, the policy will shift the curve in Fig. 1 downward. Figure 3 below is further drawn under the assumption that the policy affects the promotion of a worker more when this worker is more likely to have a promoted spouse. The leftmost part of the curve will then display the largest shift, since on this part of the curve we find workers that with the highest probability have promoted spouses. Consequently, the policy will not only shift the reaction functions downwards, but also make them in general less steep.

Now we are ready to assess the equilibrium effects of the proposed policy. Depending on the amount of such policies, the equilibria may entail less discrimination or even collapse to a single symmetric stable equilibrium. Both alternatives are illustrated in Fig. 3 below.

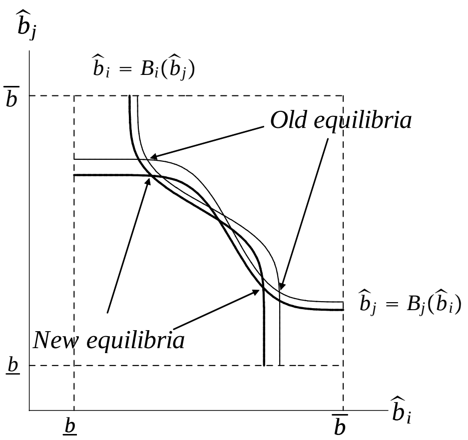

Fig. 3a: Reduced discrimination

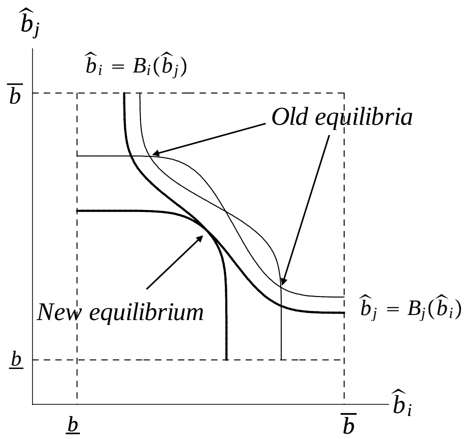

Fig. 3b: Eliminated discrimination

In the new equilibrium women will be promoted more often than before. It is a priori unclear whether the same applies to men. On one hand, it is clear that a large enough downward shift of the reaction curves implies that in the new equilibrium both sexes face less tough promotion standards. On the other hand, the effect of decreasing the slope of the curves will pull the equilibrium point toward the 45-degree line, suggesting that men may face tougher promotion standards in the new equilibrium. Put perhaps more intuitively, there is a tendency that more people are promoted, but also a tendency that promotion standards are more equal between the sexes. Some mediocre men, who before were promoted because they could be relied upon not to do much housework, will now have promoted wives and will have to share the work at home. Consequently, they are not promoted.[22]

It should be emphasized that even if the proposed policy makes the reaction curves flatter in general, family policy does not necessarily imply reduced discrimination. The reason is that the policy also affects the location of the fixed point. If the fixed point moves from a location with low density to a location of higher density of workers, this will, ceteris paribus, increase the slope of the reaction function at its fixed point, and this effect may well dominate the general reduced-slope effect of the policy. As a consequence, we may get the somewhat counterintuitive result that public provision of subsidized daycare creates discrimination. This possibility is illustrated in Fig. 4 below.

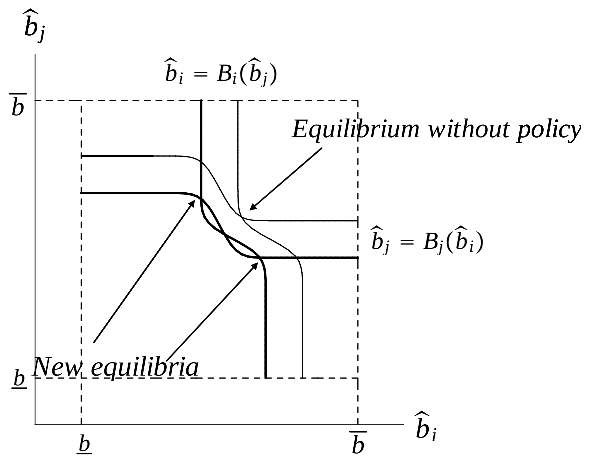

Fig. 4: Policy may create discrimination

[22]Note that the ambiguity of the effect of the policy on men's promotion standards is to some extent caused by our assumption of independent promotion decisions. If promotion decisions were more interdependent, lowering the promotion standard for women will in itself decrease the profitability of promoting men. Clearly, this effect is always dominant in the extreme case in which the number of fast-track positions is fixed.

4 Efficiency and distribution

So far we have concentrated on whether anti-discrimination policy can be expected to work in a framework with possible self-fulfilling prophecies. We now briefly discuss the efficiency and distribution consequences of such a policy.

Concerning efficiency, as a main rule, the total value created in a fast-track job should be large enough to justify the fixed installation cost. This does not help us much when choosing between equilibria, however, since this rule resembles what the investing party has an incentive to do in any situation, whether it is an equilibrium or not. Moreover, externalities may call for deviations from the main rule. First, ex post rent sharing tends to produce underinvestment. (If some of the installation costs are not deductible, certain taxes have the same effect.) Second, if firms make the investment, promoting a man will exert a negative externality on those who employ women, and vice versa. (The same outcome will arise if the workers themselves make the investment, unless the workers internalize the negative effect of their own investment on their future spouses.)

A more worrying problem is uncertainty. Due to uncertainty about the talent and career of one's future spouse, one's future effort is also uncertain. As a consequence, ex post one will observe some poor investment decisions of both types. Effort uncertainty need not be constant across equilibria, however. If effort depends much on the spouse's career and little on her talent per se, much of the effort uncertainty is sometimes resolved in the asymmetric equilibium, and then it may make sense to use discrimination as a coordination device, substantially reducing effort uncertainty and thereby also substantially reducing the problem of ex post inefficient investment decisions. Discrimination entails, to some extent, that talent is wasted, but it gives us families with bread-winning fathers and housekeeping mothers, which makes everyone's effort choices more predictable. It is straightforward to show that discrimination may be good for efficiency: suppose the talents are almost similar, but that effort depends significantly on whether one's spouse is promoted or not (that is, coordination is important). Then discrimination entails good coordination without any significant waste of talent. Francois (1998) takes this point to the extreme, since there are no talent differences in his setup.

Implementing "fair" promotion rules implies that high-talent women are promoted rather than low-talent men. (This effect is easiest to spot when the two are married to each other, cf. what is referred to as area E in Fig. 5 below.) Putting talent to better use in this way is by itself a way to enhance efficiency. While it was easy to find a case in which discrimination is good for efficiency, it is slightly more complicated to find a case in which discrimination is bad for efficiency. The complications arise precisely because discrimination is already limited by efficiency considerations: discrimination can only occur in equilibrium when lack of talent can be compensated for by extra effort. Consequently, it is not enough to identify situations which are opposite to the situation described in the previous paragraph, i.e. little family interaction and/or important talent differences (compared to effort differences), because then discrimination will not occur in equilibrium.

Efficiency gains from non-discrimination therefore require a certain environment. Somewhat loosely, non-discrimination may affect the total number of people invested in, and if the net effect is positive and sufficiently strong, this will in itself imply an efficiency gain. To be more precise, suppose the employers pay the investment and then the gross surplus is shared according to e.g. Nash bargaining (on the extent of rent sharing in labor markets, see e.g. Blanchflower, Oswald and Sanfey 1996). As mentioned above, rent sharing gives underinvestment incentives. The underinvestment problem can be made arbitrarily large by allocating more bargaining power to the workers. If so, then almost any change that increases the number of people on fast-track will be efficiency-enhancing. Increasing the number of promoted people is easiest if the policy in question affects relative prices (e.g. by day care subsidies), but this may also happen with affirmative action.

More generally, it should not be surprising that efficiency effects can go either way in a model that does not specify the exact wage setting regime or how family decisions are made. Efficiency analysis is clearly interesting, but requires a more specialized model and is therefore beyond the scope of the present paper.

From a distribution point of view, anti-discrimination policy is a policy with winners and losers. We will evaluate distributional effects from a family perspective. Any family member is taken to be better off if his spouse is promoted. In a harmony model this is a natural assumption: in a non-cooperative family model (as in Konrad and Lommerud 1995, see also Lundberg and Pollak 1993) there may be a potential conflict between individual and family welfare, in that a husband can lose from his wife's promotion.

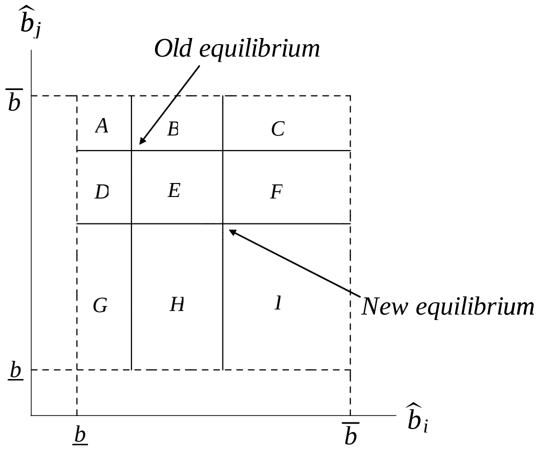

Fig. 5: Effects on workers' utilities

Consider implementing the family policy as studied in the previous section in a situation characterized by women meeting tougher promotion standards than do men, and suppose the family policy implements the symmetric equilibrium as in Fig. 3b. Fig. 5 above replicates the essentials of Fig. 3b, assuming that worker j is the female worker. We have drawn horizontal and vertical lines through the two equilibrium points, in effect dividing the set of possible talents into nine different rectangular areas, marked with capital letters A through I in the figure.

What remains is straightforward. Families consisting of persons drawn from region A are not affected by the policy: the female (worker j) is promoted and her spouse is not in both equilibria. Similarly, families in regions C, G or I are also left unaffected. In regions B, E and H the male who was promoted in the old equilibrium is no longer promoted, while in families from regions D, E and F the female was not promoted in the old equilibrium but will be so in the new. In conclusion, the policy has winners, but also losers. The winners are households with talented women who were previously discriminated against, and the losers are households with less talented men who are now driven out of the fast-track jobs. The effect on families with both winners and losers (region E) is ambiguous, depending on whether the wage premium from being promoted is increasing in talent or not.

Discrimination is to some extent redistributive across families and may from an ex ante viewpoint be seen as a form of insurance: in the discriminatory equilibrium the good jobs are distributed according to gender, and gender, in contrast to talent, is evenly distributed across families (by assumption). However, even though discrimination might lead to a more even income distribution, this is probably not the optimal way to effectuate a redistributive incomes policy. For this reason, this argument instead suggests that anti-discrimination should be supplemented with redistributive policies, not that it should be abandoned. But as long as redistribution is costly, this adds to the cost of pursuing anti-discrimination policy.

We wish to stress that the somewhat counterintuitive welfare effects arise in a model where employers harbor no discriminatory sentiments per se. Therefore, even though discrimination entails wasted talent, this is only to a limited extent: a man can be promoted instead of a more talented woman only as long as he can be expected to compensate for his lack of talent by exerting more effort. Moreover, we have seen that discrimination has some positive side effects. Discrimination serves to insure workers, and it serves to coordinate promotion decisions. We should perhaps also emphasize that even if anti-discrimination policy need not increase welfare, it can very well do so, especially if combined within redistributive measures that sterilize the increased inequality at family level that accompanies fair promotion.

5 Concluding remarks

We have investigated a model of self-fulfilling-prophecies discrimination of women. The fact that women and men live together in families and share responsibility for the home production of family public goods turns out to be crucial. This means that an employer's belief about women's future work effort is by necessity intertwined with his belief about men's future choice between exerting effort in the labor market and at home. Can this give rise to asymmetric equilibria, where men and women are treated differently in spite of their being equal in all relevant economic aspects? One message of the paper is that much can happen. Basing policy recommendation on the belief that self-fulfilling expectations about women are important in the labor market – or that they are not – without further thought or consideration, is therefore a careless strategy.

Self-fulfilling prophecy equilibria can arise, though, and when this is the case, both affirmative action and family policy that encourage female effort supply in the labor market can break discrimination. We base our discussion of policy on an example we find reasonable, with inverse-S shaped reaction curves that cross at most three times. In this example, a temporary “announcement strategy” cannot permanently banish discriminatory beliefs. Permanent use of policy is required. There are many problems with anti-discrimination policy. For example, mistakenly applied policy can create asymmetric equilibria rather than move the economy toward gender-blindness. From an efficiency point of view, eliminating discrimination can ensure better use of the talent pool in society, but an increase in the probability of wasteful fast-track initiation of people that later turn out to exert little effort can be a problem. In sum, it may be the case that the economy is locked in a discriminatory asymmetric equilibrium, that policy can take us to symmetric treatment of the sexes, and that this is efficiency improving. Successful identification of when and how policy measures should be used is quite problematic. By way of conclusion we would like to emphasize, however, that even if fair promotion may be undesireable according to welfarist criteria, we have some sympathy with the view that women simply have a right to be judged on the basis of their talents and not their sex - regardless of how this affects efficiency and distribution in the economy.

Appendix: Proof of Proposition 2.

First we will show that more than 3 crossings are impossible if the fixed point of the reaction function equals the peak of the distribution. Then the reaction function $B_j(\widehat{b}_i)$ is strictly convex for $\widehat{b}_i > b^*$ and strictly concave for $\widehat{b}_i > b^*$, while the opposite applies to its mirror image $B_i^{-1}(\widehat{b}_i)$. If $B_i^{-1}(\widehat{b}_i)$ is the steepest of the reaction functions for $\widehat{b}_i = b^*$ then the same must apply for all other values of $\widehat{b}_i$, hence the two curves cannot possible cross outside the fixed point (as in the right hand side of Fig. 2).

In contrast, if $B_j(\widehat{b}_i)$ is the steepest at $\widehat{b}_i = b^*$ then, as we increase $\widehat{b}_i$, sooner or later$B_i^{-1}(\widehat{b}_i)$ becomes the steepest, and when the two curves eventually cross again this can happen only once to the right of (cf. the left hand side of Fig. 2). By symmetry, the two curves will cross exactly once also for $\widehat{b}_i < b^*$.

What remains is to consider situations in which the fixed point is not at the distribution peak. By assumption A2 and continuity of the reaction functions, the reaction functions have to cross an odd number of times. Suppose the reaction functions cross five times. Then there are three stable equilibria and two unstable ones, as indicated in Fig. A.1 below.

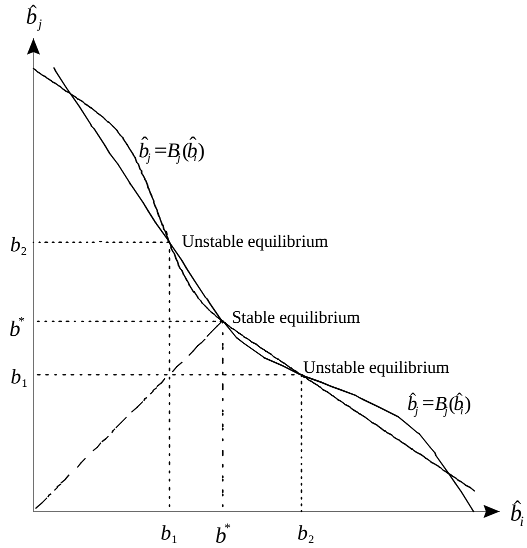

Fig. A.1: Five crossings

Since five crossings is impossible if the fixed point of the reaction function is at the peak μ of the talent distribution, five crossings implies either $b^* > \mu$ or $b^* < \mu$. In what follows we assume the former: $b^* > \mu$. (The proof for the opposite case follows the same steps and is omitted.)

Suppose the two unstable equilibria are at $\widehat{b}_i = b_1 < \mu$ and $\widehat{b}_i = b_2 > \mu$. The reaction function $\widehat{b}_j = B_j(\widehat{b}_i)$ is convex for $\widehat{b}_i > \mu$ and therefore for $\widehat{b}_i > b^*$. Since $\left|B_j'(b^*)\right| < 1$, then $\left|B_j(\widehat{b}_i)\right| < 1$ for all $\widehat{b}_i > b^*$. Symmetry then implies that (cf. Fig. A.1)

$$b_2 - b^* > b^* - b_1 \tag{8}$$

Second, since, by assumption, the crossings at b_1 and b_2 are unstable equilibria while the crossing at b^* is stable, the following must hold:

$$\left|B_j'(b_1)B_i'(b_2)\right| < 1 < \left|B_j'(b^*)\right|^2 \tag{9}$$

Substituting $-f(b_i)(\overline{e} - \underline{e})$ for $B_j'(b_i)$ and rearranging yields

$$f(b_1)f(b_2) < f(b^*)^2 \tag{10}$$

Taking logarithms on both sides preserves the inequality. Hence,

$$\frac{1}{2}\ln f(b_1) + \frac{1}{2}\ln f(b_2) < \ln f(b^*) \tag{11}$$

In Fig. A.2 below we have illustrated this inequality under the assumption that $\ln f(b_i)$ is a concave function and that $b_1 < \mu$ (it is easily verified that the result also holds if $b_1 > \mu$).

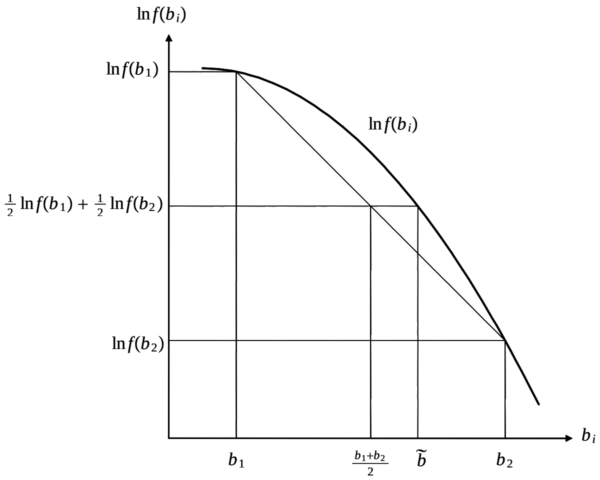

Fig. A.2: The implications of log-concave density

If inequality (A.4) is to be satisfied, it should be clear that $b^* > \widetilde{b}$, hence that $b^* > \frac{1}{2}(b_1 + b_2)$ or, rearranging,

$$b_2 - b^* < b^* - b_1 \tag{12}$$

Comparing inequalities (A.1) and (A.5) reveals a contradiction. Consequently, five crossings is incompatible with having a talent density that is log-concave.

What remains is to show that also seven crossings is incompatible with log-concave talent density. (By induction, if 5 crossings is impossible, so is 9, 13, 17 etc. Similarly, if 7 crossings is impossible, so is 11, 15 etc.) The proof essentially follows the same steps, and we therefore only sketch the proof. With seven crossings, let the numbers b_1 through b_7 denote the values of b_i that corresponds to the crossings, in ascending order. Then we know that the crossings at $b_2, b_4 (= b^*)$ and b_6 are unstable equilibria while the crossings at b_1, b_3, b_5 and b_7 are stable equilibria. Then symmetry implies that (after rearranging)

$$\frac{1}{2}\ln f(b_3) + \frac{1}{2}\ln f(b_5) < \frac{1}{2}\ln f(b_2) + \frac{1}{2}\ln f(b_6) \tag{13}$$

Suppose $b_4 > \mu$. Then it can be shown that convexity of the reaction function $\widehat{b}_j = B_j(\widehat{b}_i)$ for $b_i > b^*$ together with symmetry implies that $b_6 - b_5 > b_3 - b_2$, while with log-concave density it is easily seen that inequality (A.6) implies the opposite, a contradiction. Q.E.D.

References

Aigner, Dennis J. and Glen G. Cain, 1977. Statistical theories of discrimination in the labor market. Industrial and Labor Relations Review, 30, 175-187.

Arrow, Kenneth J., 1973. The theory of discrimination. In Orley Ashenfelter and Albert Rees, eds., Discrimination in Labor Markets. Princeton University Press.

Austen-Smith, David and Michael Wallerstein, 2003. Redistribution in a divided society. Manuscript, Northwestern University.

Bagnoli, Mark and Ted Bergstrom, 1989. Log-concave probability and its applications. Mimeo, University of Michigan. Downloadable at http://www.econ.ucsb.edu/~tedb/Theory/logconc.ps.

Becker, Gary, 1991. A treatise on the family. Enlarged edition. Cambridge, Mass.: Harvard University Press.

Blanchflower, David, Andrew Oswald and Peter Sanfey, 1996. Wages, Profits and Rent-Sharing. Quarterly Journal of Economics, 111, 227-252.

Blau, Francine D. and Marianne A. Ferber, 1992. The economics of women, men and work. Second edition. Englewoods Cliffs, New Jersey: Prentice Hall.

Booth, Alison, Marco Francesconi and Jeff Frank, 2003. A sticky floors model of promotion, pay and gender. European Economic Review, 47, 295-322.

Coate, Stephen and Glenn C. Loury, 1993. Will affirmative-action policies eliminate negative stereotypes? American Economic Review, 83, 1220-1240.

Donald, Stephen G., David A. Green and Harry J. Paarsch, 2000. Differences in wage distributions betwewen Canada and the United States: An application of a flexible estimator of distribution functions in the presence of covariates. Review of Economic Studies, 67, 609-633.

De Fraja, Gianni, 2002. Affirmative action and efficiency in education. CEPR Discussion Papr No. 3357, London.

Engineer, Merwan and Linda Welling, 1999. Human capital, true love and gender roles: is sex destiny? Journal of Economic Behavior and Organization, 40, 155-78.

Francois, Patrick F., 1998. Gender discrimination without gender difference: theory and policy responses. Journal of Public Economics, 68, 1-32.

Fuchs, Victor, 1989. Women's quest for economic equality. Journal of Economic Perspectives, 3, 25-41.

Hersch, Joni, 1991. The impact of nonmarket work on market wages. American Economic Review, 81, Papers and Proceedings, 157-160.

Joshi, Heather, 1988. Changing roles of women in the British labour market and the family. Discussion Paper in Economics 88/13, Birkbeck College.

Judge, George G., R. Carter Hill, William E. Griffiths, Helmut Lütkepohl and Tsoung-Chao Lee, 1988. Introduction to the theory and practice of econometrics, Second Edition, John Wiley & Sons, New York.

Konrad, Kai A. and Kjell Erik Lommerud, 1995. Family policy with non-cooperative families. Scandinavian Journal of Economics, 97, 581-601.

Konrad, Kai A. and Kjell Erik Lommerud, 2000. The bargaining family revisited. Canadian Journal of Economics, 33, 471-487.

Konrad, Kai A., Harard Künemund, Kjell Erik Lommerud and Julio R. Robledo, 2002. Geography of the family. American Economic Review, 92, 981-998.

Krugman, Paul, 1987. The narrow moving band, the Dutch disease, and the competitive consequences of Mrs. Thatcher. Journal of Development Economics, 27, 41-55.

Lang, Kevin, 1986. A language theory of discrimination. Quarterly Journal of Economics, 101, 363-382.

Lazear, Edward and Sherwin Rosen, 1990. Male-female wage differentials in job ladders. Journal of Labor Economics, 8, S106-S123.

Lommerud, Kjell Erik and Steinar Vagstad, 2000. Mommy tracks and public policy: On self-fulfilling prophecies and gender gaps in promotion. CEPR Discussion Paper No. , London.

Lundberg, Shelly and Robert Pollak, 1993. Separate spheres bargaining and the marriage market. Journal of Political Economy, 101, 988-1010.

Lundberg, Shelly and Richard Startz, 1983. Private discrimination and social intervention in competitive labor markets. American Economic Review, 73, 340-347.

Milgrom, Paul and Sharon Oster, 1987. Job discrimination, market forces, and the invisibility hypothesis. Quarterly Journal of Economics, 102, 453-476.

Moro, Andrea and Peter Norman, 2003. Affirmative action in a competitive economy. Journal of Public Economics, 87, 567-594.

Moro, Andrea and Peter Norman, 2004. A general equilibrium model of statistical discrimination. Journal of Economic Theory, forthcoming.

Norman, Peter, 2004. Statistical discrimination and efficiency. Review of Economic Studies, forthcoming.

Phelps, Edmund S., 1973. The statistical theory of racism and sexism. American Economic Review, 62, 659-661.

Renes, Gusta and Geert Ridder, 1995. Are women overqualified? Labour Economics, 2, 3-18.

Vagstad, Steinar, 2001. On private incentives to acquire household production skills. Journal of Population Economics, 14, 301-312.