

Cheap Talk in the Classroom: How Biased Grading at School Explains Gender Differences in Achievements, Career Choices and Wages

LYDIA MECHTENBERG

Technische Universität Berlin

First version received March 2006; final version accepted December 2008 (Eds.)

In this paper, I provide a theoretical explanation for the gender differences in education and on the labour market that are observed empirically in most OECD (Organisation for Economic Cooperation and Development) countries, including the US. Within a cheap talk model of grading, I show that biased grading in schools results in (1) boys outperforming girls in maths and sciences, (2) boys having more top and more bottom achievers in maths and sciences than girls, (3) girls outperforming boys in reading literacy, (4) female graduates enrolling in university studies more often than male graduates, (5) the predominance of female students in arts and humanities at the university, (6) the predominance of male students in maths and sciences at the university and (7) the gender wage gap on the labour market for the highly educated.

1. INTRODUCTION

Persisting gender differences in education and on the labour market constitute a threefold puzzle to economists: achievements in school, enrolments at university and wages show gender-specific patterns that are very similar across different OECD (Organisation for Economic Cooperation and Development) countries. Concerning achievements in school, boys and girls differ both in mathematical and non-mathematical subjects. With regard to mathematical subjects, nearly all existing data on cognitive achievement of school children reveal that boys outperform girls in maths and sciences.¹ Besides, significantly more males than females exceed the magnitude of the highest proficiency level in mathematics.² In about half of the OECD countries, including the US, boys perform also more frequently at the lowest proficiency level in the maths and science tests of TIMSS and PISA. At moderate proficiency levels, females are more strongly

1. See the Third International Mathematics and Science Study (TIMSS) 2000 Report Mullis *et al.*, 2000 and the PISA 2003 Report (Lemke *et al.*, 2004). In almost all OECD countries, average PISA (Programme for International Student Assessment) scale scores in mathematics are higher for males than for females. See, for instance, the PISA 2003 Report, Figure 2.18.

2. In the US, for example, 2.8% of 15-years-old males and only 1.2% of 15-years-old females perform at Level 6 in mathematics, the highest possible proficiency level in PISA 2003.

represented.³ With regard to reading abilities, by contrast, PISA data prove that girls are on average better and have more top achievers and less bottom achievers than boys in almost all OECD countries.⁴ With regard to university enrolments, much more females than males enrol in university studies. But while female students are predominant in arts and humanities, male students occupy the maths and science fields.⁵ Finally, there is a significant wage gap between men and women on the labour market.⁶

While the economics literature has provided some possible theoretical explanations for the gender wage gap, there are so far no attempts to theoretically explain the gender differences in education. More importantly, to my knowledge there is no unified model providing a single explanation for all gender differences mentioned above. The current paper aims at closing this gap. I provide an explanation of gender differences in achievements at TIMSS and PISA tests, gender differences in choices of whether and what to study and gender differences in earnings. The explanation I suggest locates distortions at school.

I analyse a symmetric cheap talk model of teachers and students. In this model, teachers in humanities and maths get signals about the talent of their students. Then, they send messages (grades) to their students, either transmitting their signals or lying. Then, students choose the effort they would invest into the subject at university. I assume that these efforts are measured by the PISA and TIMSS data. Afterwards, the students choose whether or not to go to university. When they decide to go to university, they choose the field—maths or humanities—they want to enrol in. They invest their corresponding effort. At the end of university, their human capital, the product of talent and effort, becomes publicly observable. They enter the labour market and earn a wage that is determined by their human capital.

The perfect Bayesian equilibria of this model are characterized by the equilibrium grading behaviour of teachers and the corresponding equilibrium beliefs that the students have about the meaning of the grades they get. The existence of equilibria without full information transmission is due to the assumption that some teachers are biased. This assumption is justified in Section 2 of this paper in which I discuss empirical studies on grading behaviour of teachers and other related literature. From this literature, one can draw two inferences. First, biased teachers often use grades as messages about their liking or disliking the student's attitude. But second, biased teachers are somewhat scrupulous; they do not want to bias their grades too much. Therefore, I assume that teachers do not want to distort the beliefs that their students have about their own talents too much. They give biased grades if their students do not take these grades too seriously, but if they anticipate a student to internalize a grade fully as message about his (her) talent, they react with using the grade as an unbiased message about talent only.

In the current paper, I concentrate on the asymmetric equilibrium that is characterized by the gender differences observable in schools, at universities and on the labour market. I call this equilibrium the "US equilibrium", because the US is one of the most important OECD

3. For the comparison of achievement distributions for males and females, compare Table 2.5b in the PISA 2000 and 2003 Reports (OECD, 2001; Lemke *et al.*, 2004).

4. In the US, 9.3% of 15-year-old males, but only 3.7% of 15-year-old females, perform below proficiency Level 1 in the reading-literacy test of PISA 2000. But 11% of 15-year-old males and 13.4% of 15-year-old females perform at the highest level, Level 5. For the achievement distributions in reading (by gender), compare Table 5.2a in the PISA 2000 Report. The PISA 2003 Report claims that results did not change much from 2000 to 2003. The gender differences in mean achievement in reading are reported both for 2000 and 2003 in the PISA 2003 Report, Figure 6.6.

5. In the US, for example, 81% of all tertiary degrees rewarded in type-B programmes of humanities, arts and education are allocated to females. But females get only 31% of all such degrees rewarded in mathematics and computer sciences. Compare Table A3.4, OECD, www.oecd.org/edu/eag2006.

6. See Weichselbaumer and Winter-Ebmer (2005) for a meta-study.

countries in which the gender differences described above are observed. The asymmetry of the US equilibrium is due to asymmetric equilibrium beliefs of boys and girls about the meaning of the grades they get. These asymmetric equilibrium beliefs lead to asymmetric equilibrium grading behaviour of teachers towards boys and girls, or vice versa.

More specifically, in humanities, girls do not fully internalize bad grades as bad messages about their talent in humanities. But they believe good grades to be good news about their talent. Boys, by contrast, have the reverse beliefs. Teachers in humanities anticipate that bad grades for a girl are not fully internalized as messages about the girl's talent in humanities. They also anticipate that good grades for a boy are not fully internalized as messages about the boy's talent in humanities. Thus, biased teachers in humanities start to misuse these grades as messages about their liking or disliking the student. This, in turn, justifies the students in their unwillingness to fully internalize these grades. As a result, in humanities, lowly talented girls partly discard their bad grades and become overconfident, exhibiting higher achievement than lowly talented boys. Highly talented boys partly discard their good grades and become underconfident, achieving less than highly talented girls. The resulting gender difference in achievement distributions in humanities is (in qualitative terms) exactly the one that can be observed from the PISA data on reading literacy.

In maths, the equilibrium situation is different because all students internalize bad grades as bad messages about their talent in maths. This is intuitive because in maths, it is difficult for teachers to distort their grades downwards without being detected to do so. Grades in maths can only be biased easily when they include oral contributions; and it is much easier for teachers to bias grades upwards by referring to oral contributions, than downwards. In the US equilibrium, girls do not fully internalize good grades as good messages about their talent in maths. However, boys do. Thus, biased teachers who anticipate these beliefs of their students often use good grades to express their liking of a female student, thereby justifying the reluctance of girls to internalize good grades in maths. Teachers' patronizing of girls in maths has two consequences. On the one hand, more lowly talented girls than lowly talented boys receive praise in maths, and as these girls *partially*, though not fully, internalize the praise, they achieve more on average than the lowly talented boys. But on the other hand, the highly talented girls in maths mistrust the praise they get and achieve therefore less on average than the highly talented boys. In sum, boys outperform girls in maths. Again, the resulting gender difference in achievement distributions in maths is (in qualitative terms) exactly the one that can be observed from the PISA data.

The driving force of these effects is a loss of information due to potential biases in grades. Grades are signals that students receive about their own talent, and potentially biased grades are noisier signals than unbiased grades. In the equilibrium that will be analysed in the current paper, the loss of information is generally larger for girls than for boys.

The above-mentioned self-assessments of boys and girls, which are produced at school, can explain the gender difference in university enrolment decisions and in earnings, too. To grasp the intuition of this result, consider first the group of students with good grades in both humanities and maths. The girls in this group fully trust their praise in humanities, but not their praise in maths. Thus, they decide to enrol in humanities. The boys, however, exhibit the reverse beliefs. They therefore choose to enrol in maths. Next, consider the group of students with bad grades in both subject fields. Because boys believe both the bad grade in humanities and the bad grade in maths to reflect low abilities, they decide not to go to university at all. Girls, however, partly attribute the bad grade in humanities to dislike of the teacher. Thus, they enrol in humanities *in spite of* their mediocre success in this subject at school. As a result, females enrol more often at university than males. Second, they choose humanities instead of maths more often. Third, female university students are less talented on average than male

university students, because too many females who are lowly talented in humanities enrol at university. Because this gender difference in average talents of university students is also reflected in wages, the gender wage gap can be explained without assuming any distortions on the labour market.

At this point, a remark on the rationale of analysing asymmetric equilibria in a symmetric model is in order. In general, an asymmetric equilibrium in a symmetric model should be viewed as an equilibrium that persists even when the (unmodelled) asymmetry that originally has triggered it vanishes.⁷ There are reasons to believe that the persisting gender differences observable in reality should be modelled in this way. Beliefs about talents of males and females and about the right way to educate boys and girls have considerably changed in all industrialized countries. But gender differences in schools, at universities and on the labour market have not vanished yet.

The remaining part of this paper is organized as follows. Section 2 gives a short overview of related literature. Section 3 presents the model. In Sections 4–7, equilibrium results are derived and discussed and in Section 8, conclusions are presented.

2. RELATED LITERATURE

The current paper bridges the gap between two separate branches of the literature, namely, the literature on cheap talk games and the literature on discrimination. The stream of literature on cheap talk games that started with Crawford and Sobel (1982) contains a small but important sidearm of papers on the manipulation of beliefs about one's own self. Benabou and Tirole (2000, 2003) have applied a cheap talk game to the interaction between an agent and a principal who is more informed about the agent's abilities than the agent himself.⁸ The authors consider situations in which the principal has an incentive to manipulate the agent's self-confidence either upwards or downwards. The message space of the principal allows him to evoke, according to his preferences, either underconfidence or overconfidence of the agent. For example, the principal can choose between transmitting information about the agent's abilities or not sending any message to the agent at all. The agent cannot observe whether the principal has received any information. Then, by transmitting only bad news to the agent and keeping silent about good news, the principal can make the agent underconfident.

The core intuition behind the current paper is that mechanisms like those depicted in Benabou and Tirole (2000, 2003) are the reason for gender differences in education and on the labour market. Nevertheless, the model in the current paper diverges in important respects from Benabou and Tirole (2000, 2003). First, I consider not only one cheap talk game but two simultaneous and interrelated cheap talk games. The message that a teacher (principal 1) transmits to a given student eventually also depends on the signal that the teacher of the other subject field (principal 2) receives. This is because a signal received by a teacher does not only, as in Benabou and Tirole (2000, 2003), reflect the agent's (i.e. student's) abilities

7. One way to think about the unmodelled asymmetry that might have triggered the US equilibrium originally is to include parents into the model. Suppose that originally, parents had different priors about the talents of their sons and daughters. These (irrational) priors influenced the way in which they interpreted the grades their children got at school. Suppose that girls and boys originally accepted the way in which their parents interpreted grades and therefore ended up with exactly the asymmetric beliefs that are part of the US equilibrium. Then, the US equilibrium will persist even if parents stop having irrational priors about the talents of their children or if children stop listening to their parents' interpretation of grades.

8. Compare also Fang and Moscarini (2005). In their 2002 paper, Benabou and Tirole analyse the rational management of one's own self-confidence.

but also encodes information about the student's personality that is liked or disliked by the teachers.

Second, the message space of the principal is more standard in my model. The teacher cannot choose to send no message at all; instead, he is able to lie. Third, I consider not only one type, but several types of principals, i.e. teachers. Thus, my paper is related to the literature on cheap talk with two types of senders. As in Benabou and Laroque (1992), Morris (2001) and Sobel (1985), I consider a cheap talk game where the receiver of the message (the student, in my model) does not observe the type of sender. Fourth, I also consider several types of receivers, namely, male and female lowly and highly talented students. This setting allows for underconfidence and overconfidence to occur simultaneously at the two different points of the talent distribution, resulting in achievement distributions as observed from the data.

My paper is also related to the literature on discrimination. The basic ideas of both taste discrimination and statistical discrimination are combined. Taste discrimination has been defined for the labour market by Becker (1957). According to Becker, if "good behaviour" towards a given group of agents costs a principal more than this same good behaviour towards another group of agents, then the principal is likely to discriminate against the first group whenever the two groups are distinguishable. In my model, teachers sometimes have higher utility from dishonest grading than from honest grading, and the amount of this difference in utility depends on the personality of the students receiving the grades. Thus, my model incorporates the basic element of taste discrimination.

Statistical discrimination, by contrast, originates from prior beliefs instead of preferences.⁹ Consider, for example, a principal who believes women to be less likely to be productive than men. Consequently, if the agent's pay-off depends on the principal's belief, a woman has to invest more effort than a man to earn the same pay-off. Put differently, her marginal benefit of effort is often lower, compared to a man. Thus, women end up being less productive than men on average. As it is easy to see, statistical discrimination is a self-fulfilling prophecy.¹⁰

In my model, girls and boys internalize the same kind of feedback differently because they have different expectations about how honestly the teacher behaves to them. As a result, a teacher who wants to use grades as messages about his liking or disliking a student's personality anticipates different reactions from boys and girls. This, in turn, induces the teacher to behave

9. The literature on statistical discrimination started with Arrow (1972), Phelps (1972) and Arrow (1973). See also Coate and Loury (1993).

10. Evidently, the assumption of discrimination, regardless of its being taste discrimination or statistical discrimination, necessitates research on the effects of affirmative action. Coate and Loury (1993) discuss affirmative action with regard to statistical discrimination. They prove that the effect need not be positive. Defining affirmative action as the commitment to equal representation of the advantaged and the disadvantaged group in the labour market, they show the following: although there is a "good" equilibrium in which affirmative action leads to homogeneous beliefs about the two groups, there is also a "bad" equilibrium. In the bad equilibrium, negative beliefs about the disadvantaged group, combined with affirmative action, lead to patronizing of this group. The patronized group faces easier test standards than the other group. Consequently, the disadvantaged group invests less effort than the other group and thereby confirms the negative beliefs held about them. A similar effect is demonstrated in Fryer, Loury and Yuret (2003). In their paper, they show that colour-blind affirmative action (defined as a flattening of the function that relates worker productivity to the probability of being employed) weakens the incentives for all kinds of workers to invest effort. An argument in favour of affirmative action has been put forward by Benoit (1999). Benoit assumes that the disadvantaged group of workers scores lower in assessment tests than the other group, although both groups are equally productive. In such a situation, affirmative action could be an efficient (temporary) solution. In my paper, I abstract from affirmative action because affirmative action has not yet become as prevalent in the education system as in the labour market, especially not with regard to grading. However, I shortly discuss the potential effects of affirmative action with regard to grading in the conclusion of this paper. There, I argue that such a kind of affirmative action would lead to biased grading and would produce exactly the gender differences that it would be aimed against.

in a discriminatory way that fulfils exactly the expectations of boys and girls. This kind of self-fulfilling prophecy is close to the mechanism underlying statistical discrimination.

As the way in which teachers influence the beliefs of their students is at the heart of the model, I relate my assumptions about teacher preferences closely to empirical studies of grading behaviour. Researchers on grading practices broadly agree that a large percentage of teachers consciously use grades not only to inform their students about their talents or achievements but also to signal to them how their “attitude” or “motivation” is evaluated. For example, Bursuck *et al.* (1996, p. 308) report that approximately “50% of all teachers [use] certain specific grading adaptations for their students (...) including basing grades on improvement, giving multiple grades (e.g., grades for tests and effort), and making individual adjustments to grading weights (e.g., counting projects more than tests for some students).” Non-achievement factors included into grades are pure effort, improvement, and compliance or attitude of the students.¹¹ Friedman and Manley (1992) find that according to teachers, 72% of a final grade should reflect the pupil’s achievement. Thus, teachers give biased grades, but do not want to bias their grades too much.

Jussim (1989) and Jussim and Harber (2005) show empirically that self-fulfilling prophecies and biased grading do indeed occur in the classroom. In his 1989 paper, Jussim analyses longitudinal data obtained from 27 teachers and 429 sixth graders in maths classes. He uses a path analysis, controlling for past achievements, grades and motivation of students in order to separate the effects that teachers’ perceptions and grades have on students’ achievements from the reverse effects. With regard to biased grading, he reports the following result: “In comparison to students whom teachers believed to be lazy, those whom teachers believed to try hard received higher grades, but not higher standardized test scores. (...) Perhaps teachers intentionally used grades as a way of rewarding hard-working students or as a way of punishing lazy students” (p. 473). Moreover, Jussim (1989) finds that the correlation between actual effort of the students and the effort that teachers perceive them to invest is very weak and sometimes even negative. Thus, teachers’ perceptions are biased. With regard to self-fulfilling prophecies due to biased grading, Jussim (1989) reports: “Teacher perceptions of performance had no direct effects on achievement, but they did have indirect effects mediated by students’ self-concept of ability and teacher perceptions of effort and talent (...)” (p. 475). Thus, the results “were consistent with the occurrence of modest-sized self-fulfilling prophecies and perceptual biases (...)” (p. 475). Jussim and Harber (2005) give an overview over these and similar results from other studies.¹²

As Howley, Kusimo and Parrot (2000) claim, the grading practices of teachers reported above make a class or race bias of grading more probable. This, of course, is also true with regard to a gender bias.¹³ Yet, this topic is still not widely explored, and there is so far not enough evidence on the relation between students’ grades and gender. What is known, however,

11. Compare Stiggins *et al.* (1989), Hills (1991), Cross and Frary (1996) and Zeidner (1992). Measurement experts agree, however, that grades should be based solely on achievements. See for example Waltman and Frisbie (1994).

12. For another general description of self-fulfilling prophecies in the classroom, see Trouilloud *et al.* (2002, p. 591). Many psychological studies prove that teacher expectations are strongly correlated with the effort choice of students. See, for example, Clifford and Walster (1973), Hoge and Butcher (1984), Clifton *et al.* (1986), Alvidrez and Weinstein (1999), the seminal paper of Rosenthal and Jacobson (1968) and, for an economics paper, Lavy (2004).

13. Most interestingly, Dee (2005) proves that perceptual biases of teachers occur with increasing probability when teacher’s gender and/or race differs from that of the student. For a similar study, compare Ouazad (2008). However, Ouazad does not find effects of gender differences between teachers and students. Thus, I abstract from this point as the overall empirical evidence on effects of gender differences between teachers and their students is still not strong enough.

is that both boys and girls receive confounded grades that reflect both their abilities and their perceived attitude. My model accounts for this.

3. THE MODEL

There are male students of mass one and female students of mass one. Both the group of male students and the group of female students are independently and identically distributed on a unit square. Each student i has talent θ_{ij} in the subject field j , with $j \in \{H, M\}$. The location of student i on the respective unit square is given by (x_i^H, x_i^M) . When $x_i^H \leq \frac{1}{2}$, the student is lowly talented in humanities, $\theta_{iH} = k$, and when $x_i^H > \frac{1}{2}$, the student is highly talented in humanities, $\theta_{iH} = 1$. Correspondingly, when $x_i^M \leq \frac{1}{2}$, the student is lowly talented in maths, $\theta_{iM} = k$, and when $x_i^M > \frac{1}{2}$, the student is highly talented in maths, $\theta_{iM} = 1$. Students do not know their talents; both boys and girls have the priors $\Pr\{\theta_H = 1\} = \Pr\{\theta_M = 1\} = \frac{1}{2}$. Besides, each student i has either a pleasing or displeasing personality, $a_i = 1$ or $a_i = 0$. The two types of personality have equal probability and cannot be observed directly; also students themselves do not know if their own personality is pleasing or not. Hereafter, the two types of personality are called “attitudes”. From the perspective of teachers, a pupil’s good attitude might in general consist of a bundle of traits that teachers happen to believe to be most important for a good “learning climate”, such as manner and cleanliness in one school and curiosity and liveliness in another school. Attitudes and talents are independent.¹⁴

Each student is taught in both humanities and maths. In each field, he or she is matched with probability α with a biased, i.e. potentially dishonest, teacher.¹⁵ (For simplicity, I assume a uniform α for both fields. But this does not affect the main results.) The teacher r gets an imperfect signal $s_{rij} \in \{1, k\}$ on the talent of his student i in field j , with expected signal quality $\sigma \in (\frac{1}{2}, 1)$.

Then, the grades for student i in the fields j and j' are determined in the faculty room. The grading continuation game in the faculty room has the following structure. First, both teachers, r and r' , are forced to reveal publicly the signals s_{rij} and $s_{r'ij'}$ they have got. They are not allowed to lie in the faculty room. Then, one of them is drawn randomly to be the first mover in grading. Without loss of generality, let teacher r (i.e. the teacher of field j) be the first mover. He decides on the grade $m_{rij} \in \{1, k\}$ he wants to give. Afterwards, the other teacher, r' , decides on his grade $m_{r'ij'} \in \{1, k\}$. Both are allowed to opt for messages that are not in line with their signals.

This is the simplest possible structure of the grading continuation game between teachers, which achieves the equilibrium results reported in Section 4. Two assumptions contribute to simplifying the analysis. First, the assumption that teachers cannot lie in the faculty room, although they can do so in the classroom, guarantees that teachers’ updating of their beliefs about their pupils’ attitudes be straightforward. Second, the sequential structure of the game between teachers leads to a unique pure-strategy equilibrium of the continuation game starting

14. Independence is assumed for tractability reasons. Generally, the fundamental assumption is that attitude and talent are not perfectly correlated, i.e. that there is a possibility of biased grading when teachers react to attitudes. However, the assumption of (imperfect) correlation between attitude and talent would complicate the updating of beliefs in the model too much, without contributing much to the general idea.

15. The model would be robust to the assumption that each pupil is taught by one teacher in both fields who is potentially dishonest in field j with probability α , with $j \in \{M, H\}$ and independent probabilities. However, the assumption of two teachers, one per field, is more natural in my view because potential dishonesty seems to be a trait of character rather than a subject-dependent property. I thank an anonymous referee for directing my attention towards this point.

with a given signal-vector $s = (s_j, s_{j'})$. Apart from simplifying the analysis, both assumptions are also intuitive. First, teachers want to learn more about their pupils' attitudes, so they appreciate honesty in the faculty room, and it is easier for a teacher than for a pupil to find out that some fellow teacher has misreported his signal. A teacher, for example, could check whether his colleague employed equal grading standards across all his pupils' written tests while a pupil, by contrast, could not do that. Thus, the social norm of honesty could be enforced in the faculty room, but not in the classroom. Second, the assumption of a sequential structure of the game is also natural because communication is sequential by nature.

The continuation game that starts with the end of the faculty meeting proceeds as follows. Back in the classroom, teachers send their messages $m_{rij}, m_{r'ij'} \in \{1, k\}$ to their student i . Students date up their beliefs about their talents according to Bayes' rule, taking into account the probability that the teacher was not honest about his signal. Then, students decide between a university programme in maths (field M), a university programme in humanities (field H) and the outside option, i.e. entering the labour market directly after school and earning wage w_0 .

A university student i invests effort e_{ij} into the chosen programme, with effort costs e_{ij}^2 . Then, his or her talent in the chosen field j , θ_{ij} , becomes publicly observable. At the end of his or her university career, student i has human capital $\theta_{ij}e_{ij}$. At the labour market, he or she has full bargaining power and earns $w_{ij} = \theta_{ij}e_{ij}$. Then, the game ends.

The equilibrium concept, which applies to the game, is that of perfect Bayesian equilibrium: Given their beliefs, teachers and students must make optimal decisions at all information sets, and they must update their beliefs according to Bayes' rule whenever that is defined.

3.1. Preferences of students

The expected utility of student i at the end of school is¹⁶

$$\max \{U_H(m), U_M(m), w_0\}$$

with

$$m = (m_M, m_H),$$

$$\begin{aligned} U_H(m) &= \pi_H(m)(e_H - e_H^2) + (1 - \pi_H(m))(ke_H - e_H^2) \text{ and} \\ U_M(m) &= \pi_M(m)(e_M - e_M^2) + (1 - \pi_M(m))(ke_M - e_M^2), \end{aligned} \quad (1)$$

where $\pi_H(m)$ (or $\pi_M(m)$) represents the student's posterior subjective probability of being highly talented in humanities (or maths), given the messages m . Accordingly, the student's optimal effort choice with respect to field j is given by

$$e_j^* = e_j(m) = \frac{1}{2} [\pi_j(m)(1 - k) + k]. \quad (2)$$

From equations (1) and (2), one can easily see that the expected utility of choosing field j for university studies increases monotonously in the student's confidence in his abilities in field j :

$$U_j(m, e_j^*) = \frac{1}{4} [\pi_j(m)(1 - k) + k]^2. \quad (3)$$

16. I drop subscripts when they do not contribute to clarity.

I assume that w_0 , the wage of non-academics, lies marginally above the expected utility of studying a field j in which one has received a fully credible bad message about one's talent, so that $\pi_j(m) = \pi_j(k) = (1 - \sigma)$. Thus,

$$w_0 = \frac{1}{4} [1 - \sigma(1 - k)]^2 + \varepsilon, \text{ with } \varepsilon > 0. \quad (4)$$

This assumption guarantees that a fully credible bad grade in field j prevents students from studying this field at university. Nevertheless, a bad grade that is not fully credible (but might be due to the teacher's dislike of the student's attitude) does not necessarily close the door to university studies in the respective field. Thus, students' decisions between studying humanities, studying maths, and entering the labour market are fully determined by their beliefs about their respective abilities, as reported in:

Lemma 1. (1) When $\pi_j(m) = (1 - \sigma) \forall j \in \{H, M\}$, the student enters the labour market directly after school and earns w_0 . (2) Otherwise, the student decides to study field j at the university if $\pi_j(m) > \pi_{j'}(m)$, $j \neq j'$; $j, j' \in \{H, M\}$, and to study field j' if $\pi_j(m) < \pi_{j'}(m)$. (When $\pi_j(m) = \pi_{j'}(m)$, the student is assumed to toss a coin.)

Proof. Lemma 1 follows directly from equations (3) and (4). \parallel

3.2. The signal technology and types of teachers

Consider a signal s_{rij} received by teacher r about the talent θ_{ij} of student i in field j . I assume that the quality of this signal depends on the true attitude a_i of the student. Consider first the case when the student's attitude is such as teachers want it to be, i.e. $a_i = 1$. Then, if the student is highly talented in field j , teacher r will be very likely to detect this:

$$\Pr \{s_{rij} = 1 \mid a_i = 1, \theta_{ij} = 1\} = \sigma_{rH} > \sigma.$$

But, if the student is lowly talented in field j , teacher r will be less likely to find this out:

$$\Pr \{s_{rij} = k \mid a_i = 1, \theta_{ij} = k\} = \sigma_{rL} < \sigma,$$

with

$$\frac{1}{2} (\sigma_{rH} + \sigma_{rL}) = \sigma \text{ and } \sigma_{rL}, \sigma_{rH} \in \{k, 1\}.$$

Thus, if student i has a convenient attitude, the teacher's perception of this student's talent will be biased towards a positive judgement.

Consider now the case when the student's attitude is rather inconvenient, i.e. $a_i = 0$. Then, if the student is highly talented, teacher r will not be very likely to find this out:

$$\Pr \{s_{rij} = 1 \mid a_i = 0, \theta_{ij} = 1\} = \sigma_{rL}.$$

But, if the student is lowly talented, the teacher will be quite likely to detect this:

$$\Pr \{s_{rij} = k \mid a_i = 0, \theta_{ij} = k\} = \sigma_{rH}.$$

Thus, if student i has an inconvenient attitude, the teacher's perception of the student's abilities will be biased towards a negative judgement.

Consequently, the signal s_{rij} received by teacher r is not only a signal about the student's talent θ_{ij} but also a signal about his (or her) attitude a_i . More specifically, a positive signal $s_{rij} = 1$ about the student's talent is also a positive signal about his (or her) attitude and a negative signal $s_{rij} = k$ about the student's talent is also a negative signal about his (or her) attitude.

Teachers' perceptions of their students' talents can be *weakly* or *strongly attitude-sensitive*, $\tau_r \in \{\tau_w, \tau_s\}$. A fraction $\alpha \in (0, 1)$ of teachers is weakly attitude-sensitive ($\tau_r = \tau_w$), i.e. their signal quality depends only weakly on student's attitude:

$$\sigma_{rH} = \sigma_H \text{ and } \sigma_{rL} = \sigma_L, \text{ if and only if } \tau_r = \tau_w.$$

The remaining fraction $(1 - \alpha)$, however, is strongly attitude-sensitive ($\tau_r = \tau_s$), i.e. their signal quality depends strongly on student's attitude:

$$\sigma_{rH} = \sigma_H + \Delta \text{ and } \sigma_{rL} = \sigma_L - \Delta, \text{ with } \Delta > 0, \text{ if and only if } \tau_r = \tau_s.$$

A given teacher himself knows whether he is weakly or strongly attitude-sensitive. But, his type cannot be observed directly by others.

Consider now the first move within the grading continuation game in the faculty room. Both teacher r in field j and teacher r' in field j' of student i disclose their respective signals. Consequently, both teachers have *two* pieces of information about student's attitude a_i , namely, s_{rij} and $s_{r'ij'}$. Consider first the case where r is strongly attitude-sensitive, i.e. $\tau_r = \tau_s$. Applying Bayes' rule, it is easy to see that his posterior belief about a_i is

$$\begin{aligned} \Pr_r \{a_i = 1 \mid s_{rij}, s_{r'ij'}\} &> \frac{1}{2} \text{ iff } s_{rij} = 1 \text{ and} \\ \Pr_r \{a_i = 0 \mid s_{rij}, s_{r'ij'}\} &< \frac{1}{2} \text{ iff } s_{rij} = k, \forall s_{r'ij'} \in \{k, 1\}. \end{aligned}$$

Intuitively, the signal that a strongly attitude-sensitive teacher received is, from his own perspective, more informative about a student's attitude than the signal of his colleague. This is because the strongly attitude-sensitive teacher does not know whether or not his colleague is strongly attitude-sensitive too. Consequently, the posterior belief of a strongly attitude-sensitive teacher r about the student's attitude moves into the same direction as his posterior belief about this student's talent, regardless of the signal that his colleague r' received.

Consider now the case where r is weakly attitude-sensitive. It is straightforward to show that the teacher's posterior belief about a_i is

$$\begin{aligned} \Pr_r \{a_i = 1 \mid s_{rij}, s_{r'ij'}\} &> \frac{1}{2} \text{ iff } s_{r'ij'} = 1 \text{ and} \\ \Pr_r \{a_i = 0 \mid s_{rij}, s_{r'ij'}\} &< \frac{1}{2} \text{ iff } s_{r'ij'} = k, \forall s_{rij} \in \{k, 1\}. \end{aligned}$$

Thus, from the perspective of a weakly attitude-sensitive teacher r , the signal that his colleague r' received is always more informative about student's attitude a_i than the signal that he himself received. Thus, the weakly attitude-sensitive teacher follows his colleague in his belief about the pupil's attitude. This is because a weakly attitude-sensitive teacher must take into account that his colleague might be strongly attitude-sensitive. As a result, the posterior belief of a weakly attitude-sensitive teacher r about a student's attitude runs contrary to his posterior belief about this student's talent if and only if $s_{rij} \neq s_{r'ij'}$. Put differently, his belief about the pupil's attitude diverges from his belief about the pupil's talent if and only if the pupil seems to be talented in one field but not the other. This is intuitive because teachers discuss their pupils in the faculty room, and very likely some influence the others in their opinion.

3.3. Preferences of teachers

Define the perceived attitude \hat{a}_{ri} that teacher r suspects student i to have as follows:

$\hat{a}_{ri} = 1$ iff the posterior belief of teacher r about a_i is

$\Pr_r \{a_i = 1 \mid s_{rij}, s_{r'ij'}\} > \frac{1}{2}$;

$\hat{a}_{ri} = \frac{1}{2}$ iff the posterior belief of teacher r about a_i is

$\Pr_r \{a_i = 1 \mid s_{rij}, s_{r'ij'}\} = \frac{1}{2}$; and

$\hat{a}_{ri} = 0$ otherwise.

Besides, define $d_i(m)$ as the decision of student i , given his or her grades $m = (m_{ij}, m_{ij'})$ in fields j and j' , with $j \neq j'$. Let $d_i(m) = 1$ when student i decides to enrol in subject field j at university, $d_i(m) = 0$ when student i decides not to enrol in field j , and $d_i(m) = \frac{1}{2}$ when student i is indifferent.

Then, the expected utility that any teacher r gets from giving student i the grade m_{rij} in field j amounts to

$$V_r(m_{rij}) = \left(\hat{a}_{ri} - \frac{1}{2}\right)(m_{rij} - s_{rij}) - c[d_i(s_{rij}, m_{r'ij'}) - d_i(m_{rij}, m_{r'ij'})]^2, \quad j \neq j', \quad (5)$$

with

$$c > 2(1 - k). \quad (6)$$

The expected utility of the other teacher r' is defined analogously. It is easy to deduce from equations (5) and (6) that a weakly attitude-sensitive teacher will exhibit the following behaviour. If the student has most probably a likable attitude ($\hat{a}_{ri} = 1$) and the teacher's signal has been positive ($s_{rij} = 1$), then the teacher will report his signal honestly ($m_{rij} = 1$). Similarly, if the student has most probably a less likable attitude ($\hat{a}_{ri} = 0$) and the teacher's signal has been negative ($s_{rij} = k$), the teacher will transmit his information ($m_{rij} = 0$). But, if the teacher's signal runs contrary to the perceived attitude of the student, the teacher has an incentive to lie and use the grade as a positive or negative reaction to the student's attitude alone. He will only do so, however, if distorting the grade in this way does not change the student's decision about his or her future career.

To understand the nature of teachers' preferences and their resulting behaviour more deeply, consider the two expressions in equation (5) separately. The first expression represents the teacher's emotional pay-off from grading. If the teacher suspects the pupil of having an attitude that the teacher does not like, he has a purely emotional incentive to give his pupil a bad grade. However, if the teacher believes his pupil to have most probably a likable attitude, he is emotionally inclined to give a good grade. Thus, the first expression in the teacher's utility function represents the *direct utility* the teacher has from giving a good or bad grade.¹⁷

By contrast, the second expression in equation (5) stands for the teacher's *indirect costs* of biased grading, i.e. costs that he might incur due to his pupil's reaction to the distorted grade. The intuition behind this second term is that teachers feel a (limited) responsibility for their pupils; they do not want to do them too much harm by biased grading. Thus, teachers partly internalize their pupils' welfare. To see this, note that a weakly attitude-sensitive teacher will

17. I assume that a teacher does not react to the intensity of his belief about a pupil's attitude, but only to the direction of this belief, i.e. whether a good or a bad attitude is more probable. In the simple model of this paper with only two possible grades, the assumption that teachers react to the intensity of their belief would not contribute much to realism, although it would of course make sense in a more complicated model with a broader message space. I thank an anonymous referee for pointing that out to me.

have an incentive to distort his grade only if the signals about his pupil's talents are either $s = (1, k)$ or $s = (k, 1)$, i.e. if the pupil seems to be talented in one field but not the other. In this case, biased grading can in general have two consequences for a pupil's future career: either the pupil chooses the field in which he or she is most probably talented, but does not invest the optimal effort or the pupil chooses the wrong career path *and* fails therefore to invest the optimal effort, too. For certain parameters, e.g. for a low k and w_0 , the second consequence (choosing the wrong career path) is worse than the first one (choosing the correct career path but suboptimal effort) in terms of the pupil's welfare. Thus, the assumption that teachers incur costs from making their pupils choose suboptimal career paths but not from making them choose suboptimal efforts can be rationalized by the assumption that teachers partly internalize their pupils' welfare: teachers like to express their emotions by biased grading, but not when the welfare losses that the pupil would incur in consequence would become too pronounced.¹⁸

The following Lemma reports when biased grading actually occurs.

Lemma 2. (1) If $s_{rij} = k$, a weakly attitude-sensitive teacher r who is the first mover in the grading continuation game distorts his grade of student i in field j upwards when $s_{r'ij'} = 1$ and $d_i(1, m_{r'ij'}) = d_i(k, m_{r'ij'})$ for $j \neq j'$. (2) If $s_{rij} = 1$, a weakly attitude-sensitive teacher r who is the first mover in the grading continuation game distorts his grade of student i in field j downwards when $s_{r'ij'} = k$, $j \neq j'$ and $d_i(1, m_{r'ij'}) = d_i(k, m_{r'ij'})$. (3) Strongly attitude-sensitive teachers who are first movers in the grading continuation game always report their signal, $m_{rij} = s_{rij}$. (4) The second mover r' in the grading continuation game behaves as follows. If the first mover r did not distort his grade, r' acts as if he was a first mover. In all other cases, r' reports his signal, $m_{r'ij'} = s_{r'ij'}$.

Proof. See Appendix. ||

To understand this result intuitively, note that only weakly attitude-sensitive teachers have the incentive to distort their grade. The reason is as follows. Strongly attitude-sensitive teachers follow only their own signal when updating their beliefs about their pupil's attitude. Consequently, the belief of a strongly attitude-sensitive teacher about his pupil's attitude is always in line with his belief about his pupil's talent in his field, and he never has an incentive to distort his grade. In contrast, weakly attitude-sensitive teachers follow their colleague's signal when updating their beliefs about their pupil's attitude. The weakly attitude-sensitive teacher has therefore divergent beliefs about his pupil's talent and attitude when his pupil seems to be talented in one field but not the other. In this case, the weakly attitude-sensitive teacher has an incentive to distort his grade. He will do so, however, only if his grade will not distort the

18. Indeed, the model could be extended so as to relate teachers' preferences more directly and explicitly to the welfare of their pupils. First, one could include a direct "consumption" utility from getting good grades (and a disutility from getting bad ones) into pupils' preferences. Then, the first expression in equation (5) would describe how teachers internalize this utility depending on the perceived attitude of the pupil. Second, one could replace the second expression in equation (5) by a weighted function of the pupil's welfare from his/her career decisions. Then, one would have to specify the parameters for which the equilibrium results of the current paper still hold. However, this would make the calculations (and the paper itself) much more complicated and lengthy. Among other things, teachers' updating of their beliefs would become more complicated. Besides, it is an open question whether teachers really fully internalize their pupils' welfare (in the positive or in the negative) or whether they just follow a thumb rule. For a related model in which the teachers internalize their pupils' welfare (but often in a biased way), compare Mechtenberg (2006). I thank two anonymous referees and one editor for directing my attention to the relation between teachers' preferences and pupils' welfare in the current model.

pupil's choice of a future career, given the grade of the other teacher. Thus, a teacher gives a biased grade only if he is not pivotal for the career choice of his pupil. The condition that is necessary and sufficient for non-pivotality of a teacher is given in:

Remark 1. For w_0 low enough, a teacher in field j will not be pivotal for his pupil's career choice if and only if the grade in field j' is credible in field j' but would not be credible in field j .

To see this, consider (w.l.o.g.) the grades (a) $m = (1, m_2)$ and (b) $m = (k, m_2)$. In (a), the teacher in field 2 will not be pivotal if and only if the grade 1 is credible in field 1 but not in field 2, because only in this case, the pupil chooses field 1 independently of m_2 . In (b), the teacher in field 2 will not be pivotal if and only if the grade k is credible in field 1 but not in field 2, because only then, the pupil chooses field 2 independently of m_2 , given w_0 is low enough.

4. THE US EQUILIBRIUM AND GENDER DIFFERENCES IN ENROLMENT DECISIONS

The model has an equilibrium in which all teachers report their signals honestly, so that there are no differences between the decisions of boys and girls. To see this, note that any teacher will be honest if he is pivotal for the pupil's career decision. Suppose now that both boys and girls always believe their teachers to be honest. Then, a pupil with message $m = (1, k)$ or $m = (k, 1)$ will choose the field that he or she is better in for a future career; a pupil with message $m = (1, 1)$ will (supposedly) flip a coin; and a pupil with message $m = (k, k)$ will enter the labour market directly. Consequently, both the maths teacher and the teacher in humanities are always pivotal for their pupils' career decisions and will therefore always be honest, thereby justifying the trusting beliefs of their pupils. This equilibrium is efficient because all the information is transmitted.

However, the empirical literature on grading behaviour discussed in Section 2 implies that *not* all teachers report their signals honestly to all pupils; and there *are* differences between the careers of men and women. Thus, I concentrate on equilibria where at least some pupils suffer from a loss of information due to biased grading.

There is more than one such equilibrium, but they are limited in number. First, there is no babbling equilibrium. To see this, suppose for the moment that teacher r in field j babbles to his pupil i . Then, the pupil does not believe the grade m_{rij} to be in the least informative. Hence, i chooses his or her career path independently of m_{rij} , and teacher r is therefore not pivotal. Consequently, r will want to use m_{rij} as a reaction to his belief about i 's attitude alone. If r is strongly attitude sensitive, this incentive will make him report his signal honestly. But this, in equilibrium, must be anticipated by i , so that i uses m_{rij} to update his or her belief about his or her talent in j . Because this is a contradiction, there is no babbling equilibrium. Second, given that an indifferent pupil flips a coin to decide about his or her future career path, there are no equilibria in which a given pupil can be discriminated (patronized) in both fields, M and H , at once. This follows from the condition for non-pivotality that is stated in the above Remark. Thus, an equilibrium with biased grading in this model will always be asymmetric in one sense or the other.

In the remainder of the paper, I will show that the game described in the sections above has a pure strategy equilibrium that is characterized by all the gender differences observable in half

of the OECD countries, including the US. This is true even though boys and girls are *ex ante* identical. For simplicity, the respective equilibrium shall be referred to as the *US equilibrium*.¹⁹

In order to formally characterize the US equilibrium, define $D_G(m_M^G, m_H^G) \in \{H, M, w_0\}$ as the career decision variable of a female student, given her grade m_M^G in maths and her grade m_H^G in humanities and define $D_B(m_M^B, m_H^B) \in \{H, M, w_0\}$ as the corresponding career decision variable of a male student. Let r_H be a teacher in humanities who is weakly attitude-sensitive and r_M be a teacher in maths who is weakly attitude-sensitive, too. The message that r_H sends to a boy i_B when he would be the first to distort his message to i_B shall be denoted by $m_{riH}^B(s_{riM}^B, s_{riH}^B)$, with s_{riM}^B and s_{riH}^B denoting the two signals on i_B 's abilities in maths and humanities, respectively. The corresponding message that r_H sends to a girl i_G shall be denoted by $m_{riH}^G(s_{riM}^G, s_{riH}^G)$. Similarly, the message that r_M sends to a boy i_B when he would be the first to distort his message to i_B is $m_{riM}^B(s_{riM}^B, s_{riH}^B)$ and the corresponding message that r_M sends to a girl i_G is $m_{riM}^G(s_{riM}^G, s_{riH}^G)$. Note that the equilibrium strategies of the teachers are fully described by giving m_{riH}^B , m_{riH}^G , m_{riM}^B , and m_{riM}^G , because Lemma 2 already implies that all other possible messages, including those of teachers of the other type, must be honest anyway.

Theorem 1. *There exists a pure strategy equilibrium (“US equilibrium”) of the game that is characterized by the following strategies. (1) Career decisions of female students are $D_G(1, 1) = H$, $D_G(1, k) = M$, $D_G(k, 1) = H$ and $D_G(k, k) = H$. (2) Career decisions of male students are $D_B(1, 1) = M$, $D_B(1, k) = M$, $D_B(k, 1) = H$ and $D_B(k, k) = w_0$. (3) r_H grades boys as follows: $m_{riH}^B(1, s_{riH}^B) = 1 \ \forall s_{riH}^B$, and $m_{riH}^B(k, s_{riH}^B) = s_{riH}^B \ \forall s_{riH}^B$. (4) r_H grades girls as follows: $m_{riH}^G(1, s_{riH}^G) = s_{riH}^G \ \forall s_{riH}^G$ and $m_{riH}^G(k, s_{riH}^G) = k \ \forall s_{riH}^G$. (5) r_M grades boys as follows: $m_{riM}^B(s_{riM}^B, s_{riH}^B) = s_{riM}^B \ \forall s_{riM}^B, s_{riH}^B$. (6) r_M grades girls as follows: $m_{riM}^G(s_{riM}^G, 1) = 1 \ \forall s_{riM}^G$ and $m_{riM}^G(s_{riM}^G, k) = s_{riM}^G \ \forall s_{riM}^G$.*

Proof. See Appendix. \parallel

For a better understanding, the directions of the equilibrium grading bias and the equilibrium career decisions of the students in reaction to their grades are depicted in the table below. The meanings of the arrows are as follows. The upward-pointing arrow in the second cell in the first column of the girls' table indicates that by patronizing girls in maths, teachers “shift” some girls from the second to the first cell in the column. The arrow pointing to the right means that by being too strict towards girls in humanities, teachers “shift” some girls from the second cell in the first column to the second cell in the second column. Finally, the arrow pointing to the left in the first cell of the second column in the boys' table indicates that by patronizing boys in humanities, teachers “shift” some boys from the first cell of the second column to the first cell of the first column.

Girls	$m_H = 1$	$m_H = k$	Boys	$m_H = 1$	$m_H = k$
$m_M = 1$	H	M	$m_M = 1$	M	$\leftarrow M$
$m_M = k$	$\uparrow H \Rightarrow$	H	$m_M = k$	H	w_0

Intuitively, the equilibrium can be described like this. Girls trust a good grade in humanities, but they do not trust a good grade in maths. Thus, girls who get a good grade in humanities

19. This does not mean that the phenomena described are more prominent in the US than in other OECD countries that are described by the current model. On the contrary, this conclusion would be wrong. I choose the label “US equilibrium” only because the US is the most important country among those to which the model applies.

always choose humanities for their future career, independently of their grade in maths. A maths teacher is therefore never pivotal for the career choice of a girl who is good in humanities. Consequently, a weakly attitude-sensitive maths teacher indulges in patronizing a girl who seems to be good in humanities. Thereby, he justifies her distrust in good maths grades.

A similar argument holds with regard to girls who get a bad grade in maths. Note that a bad grade in maths must come from an honest teacher. Therefore, girls are justified in taking it seriously. However, they do not believe a bad grade in humanities to be very informative. Therefore, they choose humanities for their future career even if their grade in humanities is bad. Thus, the teacher in humanities is never pivotal for the career choice of girls who get a bad grade in maths. A weakly attitude-sensitive teacher in humanities feels therefore free to give downward-biased grades to girls who seem to be bad in maths. Again, his behaviour justifies girls' beliefs: good grades in humanities are to be trusted, bad ones are not.

Consider now the boys. They always fully believe their maths teacher. Thus, the maths teacher is always pivotal for a boy's career choice and grades him therefore honestly, thereby justifying the boy's beliefs about his grades in maths. However, boys do not trust a good grade in humanities. Consequently, they will never choose humanities for their future career if they get a good grade in maths; they will choose maths instead. Thus, the humanity teacher is never pivotal for the career choice of a boy who seems to be good at maths, and the teacher can therefore feel free to patronize this boy in humanities. Again, the teacher's behaviour justifies the beliefs of the pupil.

In general, both boys and girls suffer from a loss of information at school. However, this loss is greater for the girls. For girls, the bad grades in humanities are noisier than they would be if all teachers were honest, whereas in maths, the *good* grades are too noisy. For boys, only the good grades in humanities are noisier than the other grades.

To analyse the equilibrium more formally, consider first the girls. Define

$$\pi_M^G(m_{iM}, m_{iH}) \equiv \Pr\{\theta_{iM} = 1 \mid m_i = (m_{iM}, m_{iH}), i = i_G\}$$

to be the posterior confidence of a girl in maths after having received the grades $m = (m_M, m_H)$ in maths and humanities, respectively. Accordingly, girls' confidence in humanities and boys' confidence in maths and humanities are denoted as

$$\begin{aligned} \pi_H^G(m_{iM}, m_{iH}) &\equiv \Pr\{\theta_{iH} = 1 \mid m_i = (m_{iM}, m_{iH}), i = i_G\}, \\ \pi_M^B(m_{iM}, m_{iH}) &\equiv \Pr\{\theta_{iM} = 1 \mid m_i = (m_{iM}, m_{iH}), i = i_B\} \text{ and} \\ \pi_H^B(m_{iM}, m_{iH}) &\equiv \Pr\{\theta_{iH} = 1 \mid m_i = (m_{iM}, m_{iH}), i = i_B\}. \end{aligned}$$

Theorem 2 reports the precise distribution of female students over all four possible tuples of grades and their respective posterior beliefs about their talents.

Theorem 2. *In the US equilibrium, there are (1) $\left[\frac{1}{4}(1 + \alpha) - \frac{1}{8}\alpha^2\right]$ girls with $m = (1, 1)$, $\pi_M^G(1, 1) = \frac{\sigma + (1 - \sigma)(\alpha - \frac{1}{2}\alpha^2)}{1 + (\alpha - \frac{1}{2}\alpha^2)} < \sigma$ and $\pi_H^G(1, 1) = \sigma$; (2) $\frac{1}{4}$ girls with $m = (1, k)$, $\pi_M^G(1, k) = \sigma$ and $\pi_H^G(1, k) = (1 - \sigma)$; (3) $\frac{1}{4}(1 - \alpha)^2$ girls with $m = (k, 1)$, $\pi_M^G(k, 1) = (1 - \sigma)$ and $\pi_H^G(k, 1) = \sigma$; (4) $\left[\frac{1}{4}(1 + \alpha) - \frac{1}{8}\alpha^2\right]$ girls with $m = (k, k)$, $\pi_M^G(k, k) = (1 - \sigma)$ and $\pi_H^G(k, k) = \frac{(1 - \sigma) + \sigma(\alpha - \frac{1}{2}\alpha^2)}{1 + (\alpha - \frac{1}{2}\alpha^2)} > (1 - \sigma)$. (5) Girls with grades $m = (1, 1)$ are underconfident in maths and girls with grades $m = (k, k)$ are overconfident in humanities.*

Proof. See Appendix. ||

Consider now the boys. Theorem 3 reports their distribution over tuples of grades and their respective posterior beliefs formally.

Theorem 3. *In the US equilibrium, there are (1) $\frac{1}{4}(1 + \alpha)$ boys with $m = (1, 1)$, $\pi_M^B(1, 1) = \sigma$ and $\pi_H^B(1, 1) = \frac{\sigma + (1 - \sigma)\alpha}{1 + \alpha} < \sigma$; (2) $\frac{1}{4}(1 - \alpha)$ boys with $m = (1, k)$, $\pi_M^B(1, k) = \sigma$ and $\pi_H^B(1, k) = (1 - \sigma)$; (3) $\frac{1}{4}$ boys with $m = (k, 1)$, $\pi_M^B(k, 1) = (1 - \sigma)$ and $\pi_H^B(k, 1) = \sigma$; (4) $\frac{1}{4}$ boys with $m = (k, k)$, $\pi_M^B(k, 1) = (1 - \sigma)$ and $\pi_H^B(k, k) = (1 - \sigma)$. (5) Boys with grades $m = (1, 1)$ are underconfident in humanities.*

Proof. See Appendix. ||

Given Theorem 1, the distributions of males and females over all possible career decisions can be inferred directly from Theorems 2 and 3. As one can easily see, students' career decisions in the US equilibrium reveal significant gender differences: more female than male students go to university after school, but more male students study maths and sciences. The respective formal results are summarized in²⁰

Corollary 1. *(1) In the US equilibrium, all girls but only $\frac{3}{4}$ of the boys go to university. (2) Only $\frac{1}{4}$ of the girls but $\frac{3}{4}$ of the boys study maths and sciences. (3) Boys do not study arts and humanities, but $\frac{3}{4}$ of the girls do.*

In qualitative terms, the career decisions of boys and girls in the US equilibrium, which are reported in Corollary 1, reflect the differences in enrolment decisions between male and female students within the broad majority of the OECD countries. For instance, OECD data reveal that the percentage of first tertiary degrees (degrees in type-B programmes) rewarded to females in maths and sciences lies below 35% in 16 of 23 OECD countries with available data, including the US and the UK, and it lies below 50% in 22 of these 23 OECD countries. By contrast, in humanities, arts and education, the percentage of first tertiary degrees rewarded to females lies above or equals 70% in 14 of 28 OECD countries, including the US with 81%, and it lies above or equals 60% in 26 of 28 OECD countries. Besides, the percentage of first tertiary degrees in all fields that are rewarded to females lies above or equals 60% in 12 of 28 OECD countries and it lies above 50% in 23 of 28 OECD countries.²¹

In my model, the predominance of females at the university in general and in arts and humanities in particular, and the absence of females in maths and sciences originate from biased grading solely. Thus, it is possible to explain these phenomena without referring to any possible differences in innate abilities or *ex ante* preferences of boys and girls.²²

20. Obviously, the quantitative results in Corollary 1 are due to the specific distribution of individuals over the type space. But the qualitative results can be generalized.

21. Compare Table A3.4, OECD, www.oecd.org/edu/eag2006.

22. This, of course, does not imply that no such innate differences exist. It rather means that grading practices should become a more important topic in empirical research.

5. WHY GIRLS OUTPERFORM BOYS IN READING LITERACY

The predominance of female students in arts and humanities at the university suggests that girls perform better than boys when confronted with tasks that are typical for university programmes in arts and humanities. A prevalent task in humanities is to read and interpret complex texts. Therefore, the PISA test on reading literacy necessitates abilities and efforts similar to those needed for a university programme in humanities.²³ Not surprisingly, the high enrolment rates of females in humanities and the respective low rates of males are in line with the empirical fact that girls outperform boys in reading literacy, being better on average and having less bottom and more top achievers.

However, rash conclusions about underlying differences in abilities should be avoided. Obviously, PISA test scores do not measure talent but rather achievement, a combination of talent and effort. Because effort invested by students into the PISA test is supposedly already affected by their self-confidence, biased grading has an influence on the resulting achievement distributions of boys and girls.

How students perceive their (strategic) situation during the PISA test is not clear, but it is plausible to assume that their effort during the test is highly correlated with the effort they would invest if forced to study a subject field comprising similar tasks. Therefore, I assume the hypothetical effort that the boys and girls in my model would invest if forced to study humanities to be equal to the effort they would invest into a PISA test on reading literacy.

These efforts can be defined by substituting into equation (2) the equilibrium beliefs of boys and girls about their respective talents in humanities. The results are

$$\begin{aligned} e_{h1} &\equiv \frac{1}{2} [\sigma (1 - k) + k], \\ e_{h2} &\equiv \frac{1}{2} \left[\frac{\sigma + (1 - \sigma) \alpha}{1 + \alpha} (1 - k) + k \right], \\ e_{l1} &\equiv \frac{1}{2} \left[\frac{(1 - \sigma) + \sigma \left(\alpha - \frac{1}{2} \alpha^2 \right)}{1 + \left(\alpha - \frac{1}{2} \alpha^2 \right)} (1 - k) + k \right] \text{ and} \\ e_{l2} &\equiv \frac{1}{2} [(1 - \sigma) (1 - k) + k], \text{ where} \\ e_{h1} &> e_{h2} > e_{l1} > e_{l2}. \end{aligned}$$

A student's achievement in reading literacy that is measured by the respective PISA scores is assumed to be $\theta_{iH} e_{iH}$, with e_{iH} representing the effort that student i would invest if studying humanities. Thus, bottom achievers in reading literacy will be students lacking talent in humanities ($\theta_{iH} = k$) who choose the lowest possible equilibrium effort, $e_{iH} = e_{l2}$. By contrast, top achievers in reading literacy will be those students who are highly talented in humanities ($\theta_{iH} = 1$) and choose high effort, $e_{iH} = e_{h1}$.

In order to see how biased grading by teachers of humanities affects the achievement distributions of boys and girls in reading literacy, compare first the respective fractions of bottom and top achievers. The fraction n_{l2}^G of female bottom achievers in reading literacy in the US equilibrium is

$$n_{l2}^G = \Pr \{ \theta_{iH} = k \wedge m_i = (1, k) \mid i = i_G \} = \frac{1}{4} \sigma,$$

23. During the PISA tests on reading literacy, 15-year-old pupils are asked to retrieve and interpret verbally coded information. Compare the PISA 2003 Report, p. 268.

while the corresponding fraction of male bottom achievers is

$$\begin{aligned} n_{i2}^B &= \Pr \{ \theta_{iH} = k \wedge (m_i = (1, k) \vee m_i = (k, k)) \mid i = i_B \} \\ &= \frac{1}{4} \sigma (2 - a) > \frac{1}{4} \sigma. \end{aligned}$$

Thus, boys have more bottom achievers in reading literacy than girls. This effect is due to teachers of humanities using their grades partly as a reaction to incongruous attitudes. The corresponding carelessness with which girls consider bad grades in humanities leads to overconfidence among the lowly talented girls: they achieve more than the lowly talented boys who take their bad grades in the non-mathematical classes seriously.

Consider now the top achievers. There again, the difference between boys and girls is in favour of the girls. To see this, consider first the fraction n_{h1}^G of female top achievers in reading literacy:

$$\begin{aligned} n_{h1}^G &= \Pr \{ \theta_{iH} = 1 \wedge (m_i = (1, 1) \vee m_i = (k, 1)) \mid i = i_G \} \\ &= \frac{1}{4} \sigma \left[2 - \left(\alpha - \frac{1}{2} \alpha^2 \right) \right]. \end{aligned}$$

It is straightforward to show that the fraction n_{h1}^B of male top achievers is lower:

$$n_{h1}^B = \Pr \{ \theta_{iH} = 1 \wedge m_i = (k, 1) \mid i = i_B \} = \frac{1}{4} \sigma < \frac{1}{4} \sigma \left[2 - \left(\alpha - \frac{1}{2} \alpha^2 \right) \right].$$

The reason for this gender difference among students gifted with reading literacy is that teachers in humanities often use their grades to “reward” boys for supposedly good attitudes. Boys react with mistrusting praise in humanities and, consequently, highly talented boys achieve less than highly talented girls in tests of reading literacy. These distributional results of biased grading in humanities in the US equilibrium are summarized in:

Proposition 1. *In the US equilibrium, biased grading in school shifts the achievement distribution of girls in the reading literacy test to the right of the corresponding achievement distribution of boys: there are more girls than boys with top achievement e_{h1} and less girls than boys with bottom achievement ke_{l2} .*

This result is already of some interest because the overconfidence of girls with low reading literacy and the underconfidence of boys with high reading literacy already suggest that men will specialize in mathematical subject fields or avoid the university at all, while women will fill the lecture rooms in humanities even if they are not highly talented. Besides, the distributional effect reported in Proposition 1 is exactly what can be observed from the PISA data in all OECD countries.²⁴

Nevertheless, this first observation is not sufficient. In order to fully understand the distributional effects of biased grading in humanities at school, one has to compare the average achievements of boys and girls. Consider first the students with high talents in humanities ($\theta_{iH} = 1$). The average achievement A_{H1}^G in the reading literacy test of girls with $\theta_{iH} = 1$ amounts to

$$\begin{aligned} A_{H1}^G &= \frac{1}{2} (1 - \sigma) e_{l2} + \frac{1}{2} [(1 - \sigma) + \sigma \eta] e_{l1} + \frac{1}{2} [2\sigma - \sigma \eta] e_{h1}, \text{ with} \\ \eta &\equiv \alpha - \frac{1}{2} \alpha^2, \end{aligned} \quad (7)$$

24. Compare the PISA 2003 Report, p. 282 and Table 6.4.

while the corresponding average achievement A_{H1}^B of boys with $\theta_{iH} = 1$ is

$$A_{H1}^B = \frac{1}{2} (1 - \sigma) (2 - \alpha) e_{l2} + \frac{1}{2} [\sigma + (1 - \sigma) \alpha] e_{h2} + \frac{1}{2} \sigma e_{h1}. \quad (8)$$

Let Δ_{H1} represent the difference between the highly talented girls' and the highly talented boys' average achievement in the reading literacy test. Subtracting equation (9) from equation (8), substituting for the efforts and simplifying yield

$$\Delta_{H1} = \frac{1}{4} (1 - k) \frac{\alpha^2 \left[\frac{1}{2} - 2\sigma (1 - \sigma) \right]}{(1 + \alpha) \left[1 + \left(\alpha - \frac{1}{2} \alpha^2 \right) \right]} > 0. \quad (9)$$

Thus, girls who are talented in humanities achieve more on average in the reading literacy test than boys with the same talent. This effect is due to the underconfidence of talented boys who are praised by their teacher of humanities but do not trust this praise.

Consider next the average achievements in the reading literacy test of boys and girls who are lowly talented in humanities ($\theta_{iH} = k$). Lowly talented girls achieve on average

$$A_{Hk}^G = \frac{1}{2} \sigma k e_{l2} + \frac{1}{2} k [(1 - \sigma) \eta + \sigma] e_{l1} + \frac{1}{2} k (2 - \eta) (1 - \sigma) e_{h1}, \text{ with} \quad (10)$$

$$\eta \equiv \alpha - \frac{1}{2} \alpha^2,$$

while the average achievement of lowly talented boys amounts to

$$A_{Hk}^B = \frac{1}{2} \sigma k (2 - \alpha) e_{l2} + \frac{1}{2} k [(1 - \sigma) + \sigma \alpha] e_{h2} + \frac{1}{2} k (1 - \sigma) e_{h1}. \quad (11)$$

After subtracting equation (12) from equation (11), substituting for the efforts and simplifying, the difference Δ_{Hk} between the average achievements of girls and boys who are lowly talented in humanities can be specified as follows:

$$\Delta_{Hk} = \frac{1}{4} k (1 - k) \frac{\alpha^2 \left[\sigma (1 - \sigma) - \frac{1}{2} \sigma^2 - \frac{1}{2} (1 - \sigma)^2 \right]}{(1 + \alpha) \left[1 + \left(\alpha - \frac{1}{2} \alpha^2 \right) \right]} < 0. \quad (12)$$

As one can see from equation (12), the relation between average achievements of girls and boys in the reading literacy test is reversed at the lower end of the achievement distribution, compared to the upper end: lowly talented boys achieve more on average than lowly talented girls. The reason for this effect is that the weakly attitude-sensitive teachers in humanities patronize boys. Thus, although the lowly talented boys mistrust their good grades, they internalize their praise at least partly and invest and achieve more in consequence than the lowly talented girls who got an honest mark.

The question is now how these two effects together, i.e. the higher average achievement of highly talented girls and of lowly talented boys, affect the relation between the overall average achievements of boys and girls in the reading literacy test. Define $\Delta_H \equiv \Delta_{H1} + \Delta_{Hk}$. Adding equations (9) and (12) and simplifying the result yield

$$\Delta_H = \frac{1}{4} (1 - k) \frac{\alpha^2 \left[\frac{1}{2} - 2\sigma (1 - \sigma) \right] + \alpha^2 k \left[\sigma (1 - \sigma) - \frac{1}{2} \sigma^2 - \frac{1}{2} (1 - \sigma)^2 \right]}{(1 + \alpha) \left[1 + \left(\alpha - \frac{1}{2} \alpha^2 \right) \right]}. \quad (13)$$

This specification of the difference between average achievements of girls and boys in the reading literacy test implies:

Proposition 2. *In the US equilibrium, biased grading in school has the effect that girls achieve more on average than boys in the reading literacy test: $\Delta_H > 0 \forall k < 1$.*

Proof. Define

$$N \equiv \frac{1}{4} (1 - k) \left[\alpha^2 \left[\frac{1}{2} - 2\sigma (1 - \sigma) \right] + \alpha^2 k \left[\sigma (1 - \sigma) - \frac{1}{2} \sigma^2 - \frac{1}{2} (1 - \sigma)^2 \right] \right].$$

Because $\text{signum}(\Delta_H) = \text{signum}(N)$, it suffices to analyse N . It is straightforward to show that for $k \in [0, 1]$, both $N = 0$ and $\frac{dN}{dk} = 0$ if and only if $k = 1$, and $\frac{dN}{dk} < 0$ otherwise. Besides, $\frac{d^2N}{dk^2} > 0 \forall k \in [0, 1]$. This implies that N has its minimum value at $k = 1$, this minimum value equals zero, and with decreasing k , N increases monotonously. Thus, $\Delta_H > 0 \forall k \in [0, 1)$. \parallel

Intuitively, the prevalence of girls' advantage in the overall average achievement can be explained as follows. On the one hand, lowly talented boys who are made overconfident by undeserved praise invest more effort than they would have done if they were graded honestly. Besides, they invest also more than the lowly talented girls who are overconfident in spite of their bad grade. The reason is that the overconfidence of the lowly talented boys after undeserved praise is stronger than the overconfidence of the lowly talented girls after a criticism that they suspect to be undeserved. But because these boys are lowly talented, their excess effort does not fully translate into a corresponding excess achievement. On the other hand, however, many highly talented boys are made underconfident because they never fully internalize a good grade, while comparatively few highly talented girls become underconfident. The reduction of boys' effort fully translates into a corresponding reduction of achievement. Therefore, the decrease in achievement of the highly talented boys outweighs the increase in achievement of the lowly talented boys. Thus, the achievement distributions of boys and girls in reading literacy in the US equilibrium differ from each other exactly in the way that the PISA data show for the US and all other OECD countries.

6. WHY BOYS OUTPERFORM GIRLS IN MATHS AND SCIENCES

In maths and sciences, the situation at school is almost, but not exactly, the opposite of the corresponding situation in humanities. According to the PISA data, boys have more top achievers and achieve more on average in the PISA maths test than girls in all OECD countries. But in half of the OECD countries, including the US, boys also have more bottom achievers in the maths test than girls.

Similar as before, the efforts in maths can be defined by substituting into equation (2) the equilibrium beliefs of boys and girls about their respective talents in maths and sciences:

$$\begin{aligned} e_{h1} &\equiv \frac{1}{2} [\sigma (1 - k) + k], \\ e_{h3} &\equiv \frac{1}{2} \left[\frac{\sigma + (1 - \sigma) (\alpha - \frac{1}{2} \alpha^2)}{1 + (\alpha - \frac{1}{2} \alpha^2)} (1 - k) + k \right] \text{ and} \\ e_{l2} &\equiv \frac{1}{2} [(1 - \sigma) (1 - k) + k], \text{ where} \\ e_{h1} &> e_{h2} > e_{h3} > e_{l1} > e_{l2}. \end{aligned}$$

A student's achievement in maths measured by the PISA maths test scores is therefore assumed to be $\theta_{iM}e_{iM}$, with e_{iM} representing the effort that student i would invest if studying maths and sciences. Thus, bottom achievers in maths will be students achieving only ke_{l2} , and top achievers in maths will be students achieving e_{h1} .

The fraction v_{l2}^G of female bottom achievers in maths in the US equilibrium is

$$\begin{aligned} v_{l2}^G &= \Pr \{ \theta_{iM} = k \wedge (m_i = (k, 1) \vee m_i = (k, k)) \mid i = i_G \} \\ &= \frac{1}{4}\sigma \left[2 - \left(\alpha - \frac{1}{2}\alpha^2 \right) \right]. \end{aligned}$$

It is easy to see that the corresponding fraction of male bottom achievers is higher:

$$\begin{aligned} v_{l2}^B &= \Pr \{ \theta_{iM} = k \wedge (m_i = (k, 1) \vee m_i = (k, k)) \mid i = i_B \} \\ &= \frac{1}{2}\sigma > \frac{1}{4}\sigma \left[2 - \left(\alpha - \frac{1}{2}\alpha^2 \right) \right]. \end{aligned}$$

The reason for this difference is that less lowly talented girls than lowly talented boys get a bad grade: girls who are good in humanities are patronized in maths.

Consider now the top achievers. There, the difference between boys and girls is in favour of the boys. To see this, consider first the fraction v_{h1}^G of female top achievers in maths:

$$v_{h1}^G = \Pr \{ \theta_{iM} = 1 \wedge m_i = (1, k) \mid i = i_G \} = \frac{1}{4}\sigma.$$

It is straightforward to show that the fraction v_{h1}^B of male top achievers is higher:

$$\begin{aligned} v_{h1}^B &= \Pr \{ \theta_{iM} = 1 \wedge (m_i = (1, 1) \vee m_i = (1, k)) \mid i = i_B \} \\ &= \frac{1}{2}\sigma > \frac{1}{4}\sigma. \end{aligned}$$

Again, the reason for this gender difference lies in the maths teachers patronizing girls who are good in humanities: although all boys with good marks in maths can trust their praise, girls who are also good in humanities cannot do so. Thus, they are underconfident and achieve less. These distributional results of biased grading in humanities in the US equilibrium are summarized in:

Proposition 3. *In the US equilibrium, biased grading in mathematical classes at school creates mediocrity among girls: there are less girls than boys with top achievement e_{h1} and less girls than boys with bottom achievement ke_{l2} .*

This result does not only mirror the empirical gender differences that the PISA data show for about half of the OECD countries, including the US. But, it also implies that due to having been patronized in maths at school, talented women are on average less confident in their mathematical skills and therefore less productive in this area than men.²⁵ Consider now the average achievements, starting with students with high talents in maths ($\theta_{iM} = 1$). The average achievement A_{M1}^G in the PISA maths test of girls with $\theta_{iM} = 1$ amounts to

$$\begin{aligned} A_{M1}^G &= \frac{1}{2}(1 - \sigma)[2 - \eta]e_{l2} + \frac{1}{2}[\sigma + (1 - \sigma)\eta]e_{h3} + \frac{1}{2}\sigma e_{h1}, \quad \text{with} \quad (14) \\ \eta &\equiv \alpha - \frac{1}{2}\alpha^2, \end{aligned}$$

25. To my knowledge, there is so far no economic paper on favouritism in schools. For a paper on favouritism in organizations in general, see Prendergast and Topel (1996).

while the corresponding average achievement A_{M1}^B of boys with $\theta_{iM} = 1$ is

$$A_{M1}^B = (1 - \sigma) e_{l2} + \sigma e_{h1}. \quad (15)$$

Let Δ_{M1} represent the difference between the highly talented girls' and the highly talented boys' average achievement in the maths test. Subtracting equation (15) from equation (14), substituting for the efforts, and simplifying yield

$$\Delta_{M1} = -\frac{\left(\frac{1}{2} - \sigma\right)^2 \left(\alpha - \frac{1}{2}\alpha^2\right) (1 - k)}{1 + \left(\alpha - \frac{1}{2}\alpha^2\right)} < 0. \quad (16)$$

Thus, girls who are talented in maths achieve less on average than boys with the same talent. This effect is due to the underconfidence of talented girls with good marks in both subject fields: in maths, they do not trust their praise and invest less than boys with the same grades.

Consider next the average achievements of boys and girls who are lowly talented in maths and sciences ($\theta_{iM} = k$). Lowly talented girls achieve on average

$$A_{Mk}^G = k \left[\frac{1}{2} [\sigma (2 - \eta)] e_{l2} + \frac{1}{2} [(1 - \sigma) + \sigma \eta] e_{h3} + \frac{1}{2} (1 - \sigma) e_{h1} \right], \text{ with} \quad (17)$$

$$\eta \equiv \alpha - \frac{1}{2}\alpha^2.$$

The average achievement of lowly talented boys, however, amounts to

$$A_{Mk}^B = k [\sigma e_{l2} + (1 - \sigma) e_{h1}]. \quad (18)$$

Define $\Delta_{Mk} \equiv A_{Mk}^G - A_{Mk}^B$. Then,

$$\Delta_{Mk} = \frac{\left(\frac{1}{2} - \sigma\right)^2 \left(\alpha - \frac{1}{2}\alpha^2\right) k (1 - k)}{1 + \left(\alpha - \frac{1}{2}\alpha^2\right)} > 0. \quad (19)$$

From equation (19), one can see that lowly talented girls achieve *more* on average than lowly talented boys. The reason is that the former receive an undeserved good grade in maths more often than the latter: the weakly attitude-sensitive teachers in maths patronize girls.

In order to see how the patronizing of girls in maths affects the total gender difference in average achievement, define $\Delta_M \equiv \Delta_{M1} + \Delta_{Mk}$. It is easy to see that

$$\Delta_M = -\frac{\left(\frac{1}{2} - \sigma\right)^2 \left(\alpha - \frac{1}{2}\alpha^2\right) (1 - k)^2}{1 + \left(\alpha - \frac{1}{2}\alpha^2\right)} < 0. \quad (20)$$

This implies:

Proposition 4. *In the US equilibrium, biased grading in mathematical classes at school has the effect that girls achieve less on average than boys in the maths test: $\Delta_M < 0 \forall k < 1$.*

Proof. Proposition 4 follows directly from equation (20). \parallel

The reason why being patronized in maths turns out to be, altogether, a disadvantage for the girls is as follows. The excess effort of the lowly talented girls who are made overconfident by undeserved praise does not fully translate into a corresponding excess achievement, because the lack of talent leads to a waste of effort. However, the lack of effort of highly talented girls distrusting their good grades does fully translate into a corresponding lack of achievement. Thus, the negative effect prevails. Accordingly, the gender differences between the achievement distributions of boys and girls in maths in the US equilibrium are exactly as can be observed from the PISA data for half of the OECD countries, including the US

7. WHY THE MAJORITY OF HIGHLY EDUCATED WOMEN EARN LESS THAN HIGHLY EDUCATED MEN

Since the mid-1960s, gender wage gaps have been reported for almost all countries within and outside the OECD.²⁶ Although these gaps are mostly narrowing in the industrialized countries, this process has been slowing down significantly in the US in the 1990s.²⁷ For some high-income jobs, the gender pay gap in the US has even been growing again from the end of the 1990s until 2004.²⁸ Therefore, the question why women earn less than men remains challenging.

In the rich countries, the two main reasons why women earn less than men are job segregation and the so-called “glass ceiling”, an unexplained difficulty for women to become promoted to the top jobs within their firm.²⁹ But even when differences between women and men in job characteristics, endowment, such as education and job experience, and marital status have been controlled for, a significant residual wage gap has been reported for all countries.³⁰ Thus, a woman with the same university degree and the same job experience as her male colleague often earns significantly less within the same occupation. This is also true for high-income jobs.³¹

It is plausible to conjecture that discrimination on the labour market is the unobserved source of the residual gender pay gap; this is indeed often found. But obviously, one could equally well suspect hidden productivity differences to be the cause. Indeed, performance related pay has become more prevalent in the last years, and the increasing gender wage gap in top corporate jobs in the US could also be a result of this and a gap in productivities.³² If this interpretation is true, further institutions against discrimination on the labour market would not be sufficient to eliminate the gender wage gap.

In my model, biased grading at school leads to productivity differences between men and women, which contribute to the residual gender wage gap: a woman with a master’s degree might be less confident and therefore less productive in her job than her male colleague with the same degree and in the same job.

More specifically, compare women and men who study field j at the university in the US equilibrium. During their studies, they invest effort according to equation (2). At the end of their university career, their talents become publicly observable, and they enter the labour

26. Compare Weichselbaumer and Winter-Ebmer (2005) for a meta-analysis of these studies.

27. Compare Blau and Kahn (2006).

28. See Yurtoglu and Zulehner (2006, p. 14).

29. Compare, for example, Yurtoglu and Zulehner (2006).

30. The empirical standard method to estimate this residual gender wage gap is the Blinder–Oaxaca approach. Compare Blinder (1973) and Oaxaca (1973).

31. For example, Yurtoglu and Zulehner (2006) find a significant gender wage gap for top corporate jobs in the US

32. See Yurtoglu and Zulehner (2006) for such an argument on page 16.

market. The wage of such a former student i is his (her) human capital gained at university, i.e. his (her) actual achievement $\theta_{ij}e_{ij}$.

It is easy to see that no gender wage gap exists for graduates in maths. Both male students and the few female students in maths are highly confident because they received honest praise at school for their mathematical skills. They all earn

$$w_M(\theta_{iM}, e_{iM}) = \theta_{iM}e_{h1}, \quad (21)$$

so that their average wage, too, amounts to

$$\overline{w}_M(\theta_M) = \theta_M e_{h1}. \quad (22)$$

But it is only a minority of women—only 25% within the model, according to Theorem 2—who catch up in income with their male colleagues and former fellow students. These women are those who got a bad mark in humanities, but were praised in maths by a trustworthy teacher and studied maths in consequence. The vast majority of women, however, namely 75% in the model, study humanities. In order to see why, on average, they earn less than their male fellow students with the same abilities, average wages of male and female graduates have to be compared.

Consider first the male students. Only those of them who got a bad grade in maths and a good grade in humanities study the latter subject field at the university. They trust their grade in humanities, invest high effort and earn

$$w_H(\theta_{iH}, e_{iH}) = w_H(\theta_{iH}, e_{h1}) = \theta_{iH}e_{h1}. \quad (23)$$

Obviously, the average wage $\overline{w}_H^B(\theta_H)$ of male graduates in humanities who are of type $\theta_H = 1$ amounts to

$$\overline{w}_H^B(1) = e_{h1}, \quad (24)$$

while the average wage of their lowly talented male fellow students is

$$\overline{w}_H^B(k) = ke_{h1}. \quad (25)$$

Consider now the female students of humanities. Define $g^H(\theta_H, 1, 1)$ to be the share of them who are of type θ_H and have received the message $m = (1, 1)$ at school. They all study humanities, and each of them earns $w_H(\theta_{iH}, e_{h1}) = \theta_{iH}e_{h1}$, like the men. Define now $g^H(\theta_H, k, 1)$ to be the share of females of type θ_H who have received the message $m = (k, 1)$ from their teachers. They, too, study humanities, and earn the wage $w_H(\theta_{iH}, e_{h1}) = \theta_{iH}e_{h1}$. But the remaining females of type θ_H who study this field are those who got the bad message $m = (k, k)$ from their teachers. Let $g^H(\theta_H, k, k)$ represent their share. They are not as confident as their fellow students, invest less effort and, consequently, become less productive and earn a lower wage:

$$w_H^G(\theta_{iH}, e_{iH}) = w_H(\theta_{iH}, e_{l1}) = \theta_{iH}e_{l1}.$$

Let $\overline{w}_H^G(\theta_H)$ represent the average wage of females of type θ_H and with a degree in humanities. Then, the average wage $\overline{w}_H^G(1)$ of women who are highly talented in humanities and who studied this field amounts to

$$\begin{aligned} \overline{w}_H^G(1) &= e_{h1} - \gamma(e_{h1} - e_{l1}) < e_{h1}, \text{ with} \\ \gamma &\equiv \frac{g^H(1, k, k)}{g^H(1, 1, 1) + g^H(1, k, 1) + g^H(1, k, k)} \end{aligned} \quad (26)$$

and the average wage $\overline{w}_H^G(k)$ of women who are lowly talented in humanities and who studied it amounts to

$$\overline{w}_H^G(k) = k[e_{h1} - \gamma(e_{h1} - e_{l1})] < ke_{h1}. \quad (27)$$

Thus, there will be a gender wage gap among graduates in humanities, even if one controls for talents. Put differently, this gender wage gap is due to productivity differences and not due to innate differences in abilities. Besides, this result already shows that it is impossible to separate between productivity differences and discrimination: the productivity differences that are reflected in the gender wage gap of former students of humanities are due to discriminatory behaviour (towards both sexes) at school. The gender differences in earnings of graduates are summarized in:

Proposition 5. (1) *Each male graduate i invests maximum effort e_{h1} and achieves $\theta_{ij}e_{h1}$, when he enters the labour market.* (2) *The female maths graduates, i.e. $\frac{1}{4}$ of all female graduates, do so, too.* (3) *But $\frac{3}{4}$ of all female graduates earn less on average than their male colleagues with the same education and abilities. This is because among graduates in humanities, there is a gender wage gap: $\overline{w}_H^G(1) - \overline{w}_H^B(1) = -\gamma(e_{h1} - e_{l1})$ and $\overline{w}_H^G(k) - \overline{w}_H^B(k) = -k\gamma(e_{h1} - e_{l1})$.* (4) *This gender wage gap widens with increasing income: $|\overline{w}_H^G(1) - \overline{w}_H^B(1)| > |\overline{w}_H^G(k) - \overline{w}_H^B(k)|$.*

Proof. Parts (1) and (2) of Proposition 5 follow from equations (22), (24) and (25); part (3) follows from equations (24) and (27), and part (4) is directly implied by part (3). \square

Of course, there are many different reasons for the residual gender wage gap of the highly educated which is observed empirically. This paper offers one possible explanation that differs from the standard explanations. Instead of holding discrimination on the labour market responsible for the gender wage gap, it shows that biased grading in school towards both sexes suffices to produce such a wage gap. In my model, there is no explicit or direct discrimination against women; on the labour market, highly educated women are, on average, simply less productive than highly educated men. Nevertheless, this has nothing to do with differences in innate characteristics between the sexes. The reason for their lack of productivity, compared to male graduates, is twofold. First, too many women study humanities without being sufficiently talented. Second, too many talented women who study humanities are underconfident. Thus, they invest less effort than their male fellow students.

This is the reason why the gender wage gap of the highly educated increases with increasing income, when abilities and effort are complementary: The more able the employee, i.e. the higher his or her income, the more impact does a reduction of effort have. Thus, low confidence of women produces more inequality among the highly talented wage earners than among the lowly talented ones. Actually, there is evidence that the gender wage gap at the higher end of the income distribution does indeed increase with increasing income and abilities: As Yurtoglu and Zulehner (2006, p. 14) note, the gender pay gap in top corporate jobs in the US is larger in higher positions.

8. CONCLUSION AND DISCUSSION

In the current paper, I offer a unified explanation for a large number of persistent gender differences that are observed empirically: the complex gender differences in achievements in

reading literacy and maths, as measured by the PISA test scores, the gender differences in enrolment decisions at the university and finally the gender differences in earnings of the highly educated. I found that biased grading at school, which occurs quite often according to empirical research on grading policies, is sufficient to shift girls' achievement distribution in reading literacy to the right of boys' achievement distribution, to lower girls' average, bottom and top achievement in maths, compared to the corresponding achievements of the boys, to make women shy away from maths at the university, to make them crowd in university programmes in humanities and to produce a gender wage gap among university graduates.

Although biased grading is discrimination in some sense of the word, because students with a supposedly negative attitude are treated differently than those with a supposedly positive one, I did not incorporate any kind of direct discrimination against women (or men) into my model. In particular, teachers in my model do not exhibit any taste for discrimination against girls (or boys). On the contrary, the preferences of teachers are the same towards boys and girls. It is the combination of teachers' preferences—they sometimes like to use grades as expressions of sympathy or dislike—and the different expectations of boys and girls with regard to the meanings of their grades that creates the special situation in which biased grading leads to the gender differences described.

The behaviour of teachers and students in the model could also be derived in a different way. Note that the teachers praise girls too much in maths and too seldom in humanities, while they praise boys too much in humanities. This grading policy could also result from naive counteractive measures against the gender differences described. Some teachers who are do-gooders, working against gender stereotypes, would want girls to become more confident in maths and less obsessed with humanities. In consequence, they might well be tempted to be a bit too strict with girls in humanities and a bit too enthusiastic with them in maths, not taking into account that the girls might start to suspect them of doing so. Analogously, these teachers would want to induce more zeal for humanities in their male students. Thus, they would patronize them in the non-mathematical classes, again not taking into account that the boys would become suspicious. Such naive do-gooders among teachers would produce exactly the gender differences that they would want to abolish, and the results would be the same as in the actual model. Thus, there is room for further research, in particular empirical research, on biased grading and its effects on achievements, career decisions and earnings of males and females.

APPENDIX

Proof of Lemma 1. Consider first a teacher who is the first mover in the grading continuation game. I now prove part (1) of the lemma. Assume that $d_i(1, m_{r'ij'}) = d_i(k, m_{r'ij'})$. An upward bias of a teacher r of field j will occur only if $\hat{a}_{ij} = 1$ but $s_{rij} = k$. This, in turn, happens if and only if $\Pr_r \{a_{ij} = 1 \mid s_{rij} = k, s_{r'ij'}\} > \frac{1}{2}$. Remember that

$$\begin{aligned}\Pr_r \{s_{rij} = 1 \mid a_i = 1, \theta_{ij} = 1\} &= \sigma_{rH} \text{ and} \\ \Pr_{r'} \{s_{r'ij'} = 1 \mid a_i = 1, \theta_{ij} = 1\} &= \sigma_{r'H}\end{aligned}$$

and remember that

$$\begin{aligned}\Pr_r \{s_{rij} = k \mid a_i = 1, \theta_{ij} = k\} &= \sigma_{rL} \text{ and} \\ \Pr_{r'} \{s_{r'ij'} = k \mid a_i = 1, \theta_{ij} = k\} &= \sigma_{r'L}.\end{aligned}$$

Correspondingly, it holds that

$$\begin{aligned}\Pr_r \{s_{rij} = 1 \mid a_i = 0, \theta_{ij} = 1\} &= \sigma_{rL} \text{ and} \\ \Pr_{r'} \{s_{r'ij'} = 1 \mid a_i = 0, \theta_{ij} = 1\} &= \sigma_{r'L}\end{aligned}$$

and

$$\begin{aligned}\Pr_r \{s_{rij} = k \mid a_i = 0, \theta_{ij} = k\} &= \sigma_{rH} \text{ and} \\ \Pr_{r'} \{s_{r'ij'} = k \mid a_i = 0, \theta_{ij} = k\} &= \sigma_{r'H}.\end{aligned}$$

Let λ denote the probability that $\sigma_{rH} = \sigma_{r'H}$ and $\sigma_{rL} = \sigma_{r'L}$ at the information set of teacher r ; i.e. the probability that the teachers in j and j' are either both weakly attitude-sensitive or both strongly attitude-sensitive, given the information of teacher r . Then, the posterior belief of teacher r about the student's attitude, after having observed the other teacher's signal, is either

$$\Pr_r \{a_i = 1 \mid s_{rij} = k, s_{r'ij'} = k\} \text{ or}$$

$$\Pr_r \{a_i = 1 \mid s_{rij} = k, s_{r'ij'} = 1\} = \frac{v}{v + \delta}, \text{ with}$$

$$\begin{aligned}v &= \lambda [(1 - \sigma_{rH}) \sigma_{rH} + (1 - \sigma_{rH}) (1 - \sigma_{rL}) + \sigma_{rL} \sigma_{rH} + \sigma_{rL} (1 - \sigma_{rL})] \\ &\quad + (1 - \lambda) [(1 - \sigma_{rH}) \sigma_{r'H} + (1 - \sigma_{rH}) (1 - \sigma_{r'L}) + \sigma_{rL} \sigma_{r'H} + \sigma_{rL} (1 - \sigma_{r'L})], \\ \delta &= \lambda [(1 - \sigma_{rH}) \sigma_{rH} + (1 - \sigma_{rH}) (1 - \sigma_{rL}) + \sigma_{rL} \sigma_{rH} + \sigma_{rL} (1 - \sigma_{rL})] \\ &\quad + (1 - \lambda) [(1 - \sigma_{r'H}) \sigma_{rH} + (1 - \sigma_{r'H}) (1 - \sigma_{rL}) + \sigma_{r'L} \sigma_{rH} + \sigma_{r'L} (1 - \sigma_{rL})].\end{aligned}$$

It always holds that $\Pr_r \{a_i = 1 \mid s_{rij} = k, s_{r'ij'} = k\} < \frac{1}{2}$, regardless of the types of the teachers. This implies that teacher r wants to bias his grade upwards only if the other teacher's signal has been positive, $s_{r'ij'} = 1$. Then, teacher r wants to send the upward-biased message $m_{rij} = 1$ if and only if $\frac{v}{v + \delta} > \frac{1}{2}$. It holds that if $\sigma_{rH} \neq \sigma_{r'H}$ and $\sigma_{rL} \neq \sigma_{r'L}$, then either $\sigma_{r'L} < \sigma_{rL} \wedge \sigma_{r'H} > \sigma_{rH}$ or $\sigma_{r'L} > \sigma_{rL} \wedge \sigma_{r'H} < \sigma_{rH}$. It follows that $\frac{v}{v + \delta} > \frac{1}{2}$ if and only if $\lambda < 1$ and $\sigma_{r'L} < \sigma_{rL} \wedge \sigma_{r'H} > \sigma_{rH}$. Consequently, teacher r sends the upward-biased message if and only if (i) there is a positive probability $(1 - \lambda)$ that the other teacher is strongly attitude-sensitive and he himself is not, and (ii) he himself is weakly attitude-sensitive. Because teacher r has the first move in the grading continuation game, $(1 - \lambda)$ reduces to $(1 - \alpha) > 0$. Thus, teacher r will always send the upward-biased message if he himself is weakly attitude-sensitive. This proves part (1) of Lemma 2.

The proof of part (2) is skipped because it is completely analogous to the one of part (1). The proof of part (3) follows directly from the proofs of parts (1) and (2). It remains to prove part (4).

The second mover in the grading continuation game, teacher r' , observes what the first mover did and dates up his beliefs about the first mover's type and the attitude of the student. Assume first that $d_i(1, m_{r'ij'}) \neq d_i(k, m_{r'ij'})$. Then, the first mover never distorts his grade, regardless of his type. Accordingly, the second mover does not learn anything new about the first mover's type. He is in the same situation as the first mover has been. Thus, he behaves exactly alike. Consider now the case where $d_i(1, m_{r'ij'}) = d_i(k, m_{r'ij'})$. Assume that, nevertheless, the first mover did not distort his grade. Then, the second mover r' knows that the first mover is strongly attitude-sensitive. Thus, he himself will distort his grade if and only if he is weakly attitude-sensitive, $s_{rij} \neq s_{r'ij'}$, and if $d_i(m_{rij}, 1) = d_i(m_{rij}, k)$. Thus, the second mover r' will again behave as if he was the first mover.

Assume now that the first mover distorts his grade. Then, the second mover knows that the first mover is weakly attitude-sensitive. Thus, $\lambda = 1$, and the second mover r' does not want to distort his grade. This proves part (4) of the lemma. \parallel

Proof of Theorem 1. Suppose the career decisions of the students to be as described in parts (1) and (2) of the theorem. Consider now grading in humanities. Lemma 2 implies that only a teacher who is weakly attitude-sensitive and who would be the first to distort a grade of student i in the grading continuation game ever sends a message to i that is not in line with his signal. Thus, it is sufficient to define the strategy of teacher $r = r_H$ when characterizing equilibrium grading in humanities. Consider first the grading of a boy i_B . When $s_{r'iM}^B = m_{r'iM}^B = 1$, a change in $m_{r'iH}^B$ does not influence the career decision $D_B(1, m_{r'iH}^B)$ of i_B . Hence, $d_i(1, m_{r'ij'}) = d_i(k, m_{r'ij'})$, and biased grading is costless for r_H . Besides, because r_H is weakly attitude-sensitive and the other teacher's signal has been positive, r_H 's belief about the attitude of i_B is $\Pr_r \{a_i = 1 \mid s_{rij}^B, s_{r'ij'}^B = 1\} > \frac{1}{2} \forall s_{rij}^B$. From this and the definition of teacher types and preferences, it follows that $m_{r'iH}^B(1, s_{r'iH}^B) = 1 \forall s_{r'iH}^B$. However, when $s_{r'iM}^B = m_{r'iM}^B = k$, a change in $m_{r'iH}^B$ does influence the career decision $D_B(1, m_{r'iH}^B)$ of i_B . Hence, $d_i(1, m_{r'ij'}) \neq d_i(k, m_{r'ij'})$, and biased grading is costly for r_H . Because $c > 2(1 - k)$, as assumed in inequality (6), these costs of biased grading imply that r_H grades honestly: $m_{r'iH}^B(s_{r'iM}^B, s_{r'iH}^B) = s_{r'iH}^B \forall s_{r'iM}^B, s_{r'iH}^B$. This proves that given the career decisions of boys in reaction to

possible grades in maths and humanities, grading of boys in humanities is as described in part (3) of the Theorem. The proofs that parts (1) and (2) of the theorem imply parts (4), (5), and (6) are analogous and shall therefore be skipped. It remains to prove that grading practices as described in parts (3)–(6) of the theorem imply students' career decisions given in parts (1) and (2). This follows from Lemma 1, Theorem 2, and Theorem 3. Therefore, the proofs of Theorems 2 and 3 below will complete the proof of Theorem 1. ||

Proof of Theorem 2. Define the tuple of signals received by the maths teacher and the humanities teacher on the respective talents of girl i to be $s_i = (s_{iM}, s_{iH})$, and let $m_i(s_{iM}, s_{iH}) = (m_{iM}, m_{iH})$ represent the tuple of messages. It is straightforward to show that s_i takes the values $(1, 1)$, $(1, k)$, $(k, 1)$ or (k, k) with equal probability $\frac{1}{4}$. Thus, there are four grading continuation games for each combination of teachers' types. Each teacher is weakly attitude-sensitive with probability α . Suppose the career decisions of the girls to be as described in part (1) of Theorem 1. Then, Lemma 2 implies that the three grading continuation games following the signals $(1, 1)$, $(1, k)$ and (k, k) have the outcome $m_i(s_{iM}, s_{iH}) = (s_{iM}, s_{iH})$ (both teachers honest), regardless of the types of the teachers. With regard to the grading continuation game subsequent to the signals $(k, 1)$, Lemma 2 implies that the outcome will be $m_i(s_{iM}, s_{iH}) = (s_{iM}, s_{iH})$ (both teachers honest) if and only if both teachers are not weakly attitude-sensitive, i.e. with probability $(1 - \alpha)^2$. If both teachers are weakly attitude-sensitive, the first mover will distort his message and the second mover will not. Thus, the outcome will be $m_i(k, 1) = (1, 1)$ or $m_i(k, 1) = (k, k)$, with probability $\frac{1}{2}\alpha^2$ each. If exactly one of the teachers is weakly attitude-sensitive, he will distort his message, and the other teacher will not. Thus, if only the maths teacher is weakly attitude-sensitive, the outcome of this continuation game will be $m_i(k, 1) = (1, 1)$. This happens with probability $\alpha(1 - \alpha)$. If, however, the humanities teacher is weakly attitude-sensitive, the outcome will be $m_i(k, 1) = (k, k)$. This, too, happens with probability $\alpha(1 - \alpha)$. For the girls receiving their grades, all grading continuation games with the same outcome belong to one information set. Taking this and the prior α about teachers' types into account and calculating girls' posterior beliefs about their talents with the help of Bayes' rule yields the values given in parts (1)–(4) of Theorem 2. Part (5) follows directly from parts (1) and (4). (Besides, given the posterior beliefs of the girls, Lemma 1 implies that their career decisions are as reported in part (1) of Theorem 1.) ||

Proof of Theorem 3. Again, s_i takes the values $(1, 1)$, $(1, k)$, $(k, 1)$ or (k, k) with equal probability $\frac{1}{4}$, leading to four grading continuation games for each combination of teachers' types, and again, each teacher is weakly attitude-sensitive with probability α . Suppose the career decisions of the boys to be as described in part (2) of Theorem 1. Then, Lemma 2 implies that the three grading continuation games following the signals $(1, 1)$, $(k, 1)$ and (k, k) have the outcome $m_i(s_{iM}, s_{iH}) = (s_{iM}, s_{iH})$ (both teachers honest), regardless of the types of the teachers. With regard to the grading continuation game subsequent to the signals $(1, k)$, Lemma 2 implies the following: the outcome will be $m_i(1, k) = (1, k)$ (both teachers honest) if and only if either both teachers are not weakly attitude-sensitive or solely the maths teacher is weakly attitude-sensitive. Thus, the outcome will be honesty of both teachers with probability $[(1 - \alpha)^2 + \alpha(1 - \alpha)] = (1 - \alpha)$. If both teachers are weakly attitude-sensitive, the humanities teacher will distort his message, and the maths teacher will not. This also happens if solely the humanities teacher is weakly attitude-sensitive. Thus, the outcome of the continuation game will be $m_i(1, k) = (1, 1)$ with probability $[\alpha^2 + \alpha(1 - \alpha)] = \alpha$. For the boys receiving their grades, all grading continuation games with the same outcome belong to one information set. Taking this and the prior α about teachers' types into account and calculating boys' posterior beliefs about their talents with the help of Bayes' rule yield the values given in parts (1)–(4) of Theorem 3. Part (5) follows directly from part (1). (Besides, given the posterior beliefs of the boys, Lemma 1 implies that the career decisions of boys are as reported in part (2) of Theorem 1.) ||

Acknowledgements. I thank Helmut Bester, Annette Boom, Paul Heidhues, Kai A. Konrad, Daniel Krämer, Dorothea Kübler, Johannes Münster, Kerstin Puschke, Roland Strausz, Elmar Wolfstetter, the audience at the ES WC 2005, two editors and two anonymous referees for valuable comments on earlier versions of this paper. The usual caveat applies. This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 "Economic Risk". I also thank the Econometric Society for a travel grant.

REFERENCES

- ALVIDREZ, J. and WEINSTEIN, R. S. (1999), "Early Teacher Perceptions and Later Student Academic Achievement", *Journal of Educational Psychology*, **91**, 731–746.
 ARROW, K. J. (1972), "Some Mathematical Models of Race in the Labor Market", in Pascal, A. (ed.) *Racial Discrimination in Economic Life* (Lexington, MA: Lexington Books) 187–204.
 ARROW, K. J. (1973), "The Theory of Discrimination", in Ashenfelter, O. and Rees, A. (eds) *Discrimination in Labor Markets* (Princeton, NJ: Princeton University Press) 3–33.

- BECKER, G. (1957), *The Economics of Discrimination* (Chicago: University of Chicago Press).
- BENABOU, R. and LAROQUE, G. (1992), "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility", *Quarterly Journal of Economics*, **107** (3), 921–958.
- BENABOU, R. and TIROLE, J. (2000), "Self-Confidence and Social Interaction" (NBER Working Paper).
- BENABOU, R. and TIROLE, J. (2002), "Self-Confidence and Personal Motivation", *Quarterly Journal of Economics*, **117** (3), 871–916.
- BENABOU, R. and TIROLE, J. (2003), "Intrinsic and Extrinsic Motivation", *Review of Economic Studies*, **70** (3), 489–520.
- BENOIT, J.-P. (1999), "Color Blind Is Not Color Neutral: Testing Differences and Affirmative Action", *Journal of Law Economics & Organization*, **15** (2), 378–400.
- BLAU, F. D. and KAHN, L. M. (2006), "The US Gender Pay Gap in the 1990s: Slowing Convergence" (Working Paper 508, Princeton University).
- BLINDER, A. S. (1973), "Wage Discrimination: Reduced Form and Structural Estimates", *Journal of Human Resources*, **8** (4), 436–455.
- CLIFFORD, M. M. and WALSTER, E. (1973), "The Effect of Physical Attractiveness on Teacher Expectations", *Sociology of Education*, **46** (2), 248–258.
- CLIFTON, R. A., PERRY R. P., PARSONSON, K. and HRYNIUK, S. (1986), "Effects of Ethnicity and Sex on Teachers' Expectations of Junior High School Students", *Sociology of Education*, **59** (1), 58–67.
- COATE, S. and LOURY, G. C. (1993), "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, **83** (5), 1220–1240.
- CRAWFORD, V. P. and SOBEL, J. (1982), "Strategic Information Transmission", *Econometrica*, **50** (6), 1431–1451.
- DEE, T. S. (2005), "A Teacher Like Me: Does Race, Ethnicity, or Gender Matter?" *American Economic Review*, **95** (2), 158–165.
- FANG, H. and MOSCARINI, G. (2005), "Moral Hazard", *Journal of Monetary Economics*, **52**, 749–777.
- FRYER, R. G., LOURY, G. C. and YURET, T. (2003), "Color-Blind Affirmative Action" (NBER Working Paper No. 10103).
- HOGUE, H. D. and BUTCHER, R. (1984), "Analysis of Teacher Judgments of Pupil Achievement Levels", *Journal of Educational Psychology*, **76**, 777–781.
- JUSSIM, L. (1989), "Teacher Expectations: Self-Fulfilling Prophecies, Perceptual Biases, and Accuracy", *Journal of Personality and Social Psychology*, **57**, 469–480.
- JUSSIM, L. and HARBER, K. D. (2005), "Teacher Expectations and Self-Fulfilling Prophecies: Knowns and Unknowns, Resolved and Unresolved Controversies", *Personality and Social Psychology Review*, **9** (2), 131–155.
- LAVY, V. (2004), "Do Gender Stereotypes Reduce Girls' Human Capital Outcomes? Evidence from a Natural Experiment" (NBER Working Paper No. 10678).
- LEMKE, M., SEN, A., PAHLKE, E., MILLER, D., WILLIAMS, T., KASTBERG, D. and JOCELYN, L. (2004), *International Outcomes of Learning in Mathematics Literacy and Problem Solving: PISA 2003 Results from the US Perspective* (NCES 2005-003) (Washington, DC: US Department of Education, National Center for Education Statistics).
- MORRIS, S. (2001), "Political Correctness", *Journal of Political Economy*, **109** (21), 231–265.
- MULLIS, I. V. S., MARTIN, M. O., FIERROS, E. G., GOLDBERG, A. L. and STEMLER, S. E. (2000), *Gender Differences in Achievement: IEA's Third International Mathematics and Science Study (TIMSS)* (Chestnut Hill, MA: International Association for the Evaluation of Educational Achievement, TIMSS International Study Center, Boston College).
- OAXACA, R. (1973), "Male-Female Wage Differentials in Urban Labor Markets", *International Economic Review*, **14** (3), 693–709.
- OECD (2001), "Knowledge and Skills for Life: First Results from the OECD Programme for International Student Assessment (PISA) 2000" (Paris: OECD).
- OUAZAD, A. (2008), "Assessed by a Teacher Like Me: Race, Gender and Subjective Evaluations" (INSEAD Working Paper No. 2008/57/EPS).
- PHELPS, E. S. (1972), "The Statistical Theory of Racism and Sexism", *American Economic Review*, **62**, 659–61.
- PRENDERGAST, C. and TOPEL, R. H. (1996), "Favoritism in Organizations", *Journal of Political Economy*, **104** (5), 958–978.
- ROSENTHAL, R. and JACOBSON, L. (1968), *Pygmalion in the Classroom: Teacher Expectation and Pupils' Intellectual Development* (New York: Holt, Rinehart and Winston).
- SOBEL, J. (1985), "A Theory of Credibility", *Review of Economic Studies*, **52** (4), 557–573.
- TROUILLOU, D. O., SARRAZIN, P. G., MARTINEK, T. J. and GUILLET, E. (2002), "The Influence of Teacher Expectations on Student Achievement in Physical Education Classes: Pygmalion Revisited", *European Journal of Social Psychology*, **32**, 591–607.
- WEICHSELBAUMER, R. and WINTER-EBMER, D. (2005), "A Meta-Analysis of the International Gender Wage Gap", *Journal of Economic Surveys*, **19** (3), 479–511.
- YURTOGLU, B. B. and ZULEHNER, C. (2006), "The Gender Gap in Top Corporate Jobs is still there" (unpublished manuscript).