Industrial & Labor Relations Review

Volume 63, *Issue* 4 | 2010 | *Article* 7

Understanding the Gender Pay Gap: What's Competition Got to Do with It?

Alan Manning* Farzad Saidi†

*LSE, a.manning@lse.ac.uk
†New York University, saidi@nyu.edu

Understanding the Gender Pay Gap: What's Competition Got to Do with It?

Alan Manning and Farzad Saidi

Abstract

A number of researchers have argued that men and women have different attitudes toward and behavioral responses to competition; that is, women are more likely to opt out of jobs in which performance pay is the norm. Laboratory experiments suggest that these gender differences are rather large. To check these hypotheses and findings against differences in the field, the authors use performance pay as an indicator of competition in the workplace and compare the gender gap not only in incidence of performance pay but also in earnings and work effort under these contracts. They find that although women are less likely than men to work under performance pay contracts, the gender gap is small. Furthermore, the effect of performance pay on earnings is modest and does not differ markedly by gender. Consequently, the authors argue, the ability of these competition hypotheses to explain the gender pay gap seems very limited.

UNDERSTANDING THE GENDER PAY GAP: WHAT'S COMPETITION GOT TO DO WITH IT?

ALAN MANNING AND FARZAD SAIDI*

A number of researchers have argued that men and women have different attitudes toward and behavioral responses to competition; that is, women are more likely to opt out of jobs in which performance pay is the norm and to under-perform in some competitive situations. Laboratory experiments suggest that these gender differences are rather large. To check these hypotheses and findings against differences in the field, the authors use performance pay as an indicator of competition in the workplace and compare the gender gap not only in incidence of performance pay but also in earnings and work effort under these contracts. They find that although women are less likely than men to work under performance pay contracts, the gender gap is small. Furthermore, the effect of performance pay on earnings is modest and does not differ markedly by gender. Consequently, the authors argue, the ability of these competition hypotheses to explain the gender pay gap seems very limited.

Economists have long attempted to offer a complete explanation of the gender pay gap. Although factors such as differences in labor market experience do have considerable explanatory power, a sizeable gap typically remains after the best efforts to explain it (see Altonji and Blank 1999 or Blau and Kahn 2006). Recently, a new set of explanations has been proposed, defined by the proposition that by the time men and women enter the labor market, their psychological attitudes differ with respect to the type of employment contracts they favor, which ultimately affect their on the-job performance.[1] For example, Gneezy et al. (2003) found that women perform less well than men in tournaments, particularly mixed-sex tournaments. In addition, Dohmen and Falk (2010) and Niederle and Vesterlund (2007) found that women avoid variable pay schemes (that tend to raise productivity) and tournaments. All of these studies use evidence from laboratory experiments, but their conclusions are similar to those of other researchers, such as that of Babcock and Laschever (2003), who argued that women's attitudes, not their productivity *per se*, account, at least partially, for their lower earnings. The gender gaps found in these experiments are often very

* Alan Manning is Programme Director of the Labour Markets Programme at the Centre for Economic Performance and Professor of Economics, London School of Economics. Farzad Saidi is a PhD student in the Department of Economics, New York University, and Research Affiliate of IZA.

The data used in this paper is available from the U.K. Data Archive http://www.data-archive.ac.uk/. The files used to analyze the data are available by writing to the authors at A.Manning@lse.ac.uk or saidi@nyu.edu.

[1] There is a debate, which we will not discuss in detail here, about whether these effects are due to nature or nurture. The evidence in Gneezy, Leonard, and List (2008) and Booth and Nolen (2008) suggest some role for nurture, for the first paper finds women are more competitive than men in a matriarchal society, and the second finds no gender differences for those who went to single-sex schools.

Industrial and Labor Relations Review, Vol. 63, No. 4 (July 2010). © by Cornell University.
0019-7939/00/6304 $05.00

large. For example, Dohmen and Falk (2010) showed that, even after the analysis controls for productivity, women are about 15% less likely than men to enter a variable pay scheme. Moreover, Gneezy et al. (2003) reported a 33% gender gap in performance in mixed-sex tournaments.

We evaluate the importance of these ideas in practice for our understanding of the gender pay gap. Our data come from the 1998 and 2004 British Workplace Employment Relations Survey (WERS), which contains information on the nature of the pay schemes under which individuals work; we use the presence of a performance pay contract as an indicator of a more competitive environment in the workplace. It is important to note that because we do not have an experimental design and do not know the exact nature of the incentives that workers are offered in the performance pay contracts we study, one potential criticism of our conclusion is that our variables do not capture the measures that laboratory studies have established to be important. In considering these concerns, we provide various upper bounds on the part of the gender pay gap that can be explained by the competition hypothesis.

Data

The Workplace Employment Relations Survey (WERS) collects detailed information on many aspects of employment relations in Great Britain. In one form or another, it has been conducted five times since 1980, though we only use the two most recent surveys—for 1998 and 2004—since only these contain information on the earnings of individuals. Each of these surveys includes interviews with both managers and workers. The former provide a very wide range of information on human resource practices within the establishment, including the pay system, the present focus of interest. The workers answer questions designed to elicit the usual demographic information about themselves, including information about their earnings.

WERS contains information on the use of a number of different types of variable pay schemes—profit-sharing, employee share-ownership, and performance pay. We focus exclusively in this paper on the impact of performance pay schemes since these seem to have the largest effects on behavior and are the focus of interest in the experimental literature discussed above. In WERS, information on the use of performance pay is collected for nine broad occupational groups (essentially one-digit-level groups) from the management respondent, so we do have some within-plant variation in the use of performance pay because some plants use performance pay for some occupations but not for others.

The category "performance pay" encompasses a wide range of incentive schemes, from piece-rates by which individual pay is related to an objective measure of individual output to merit pay based on subjective assessments of managers. WERS does contain information on the type of performance pay contract in use (whether payment-by-results or merit pay), the level at which performance is assessed (whether individual, group, or workplace), and the method of assessment of performance (for example, piece-rates or subjective assessment). However, this information is not disaggregated by occupation; that is, we only have information on whether these types of performance pay contracts are used for any workers within the plant. This means that the types of performance pay contracts are not mutually exclusive—managers can report the use of multiple types. Because one cannot link the type of performance pay contract to workers in particular occupations, our main analysis uses the presence of performance pay as an indicator of a more competitive environment within the plant. We do, however, experiment at various points with dividing performance pay contracts into different types.

It is important to consider the relationship between the "performance pay" variable that is the focus of interest in this paper and the "competitive" versus "non-competitive" pay schemes that have been the focus of the experimental literature. The latter has typically compared "piece-rates" and "tournaments," with piece-rates being interpreted as non-competitive pay schemes and tournaments as competitive

Table 1. The Incidence of Performance Pay Contracts

	1998	*2004*
Panel A: Basic Workforce Information		
All	16.3	32.1
Men	18.6	36.2
Women	14.2	28.3
Panel B: Sample of Employees		
All	16.1	27.6
Men	17.6	31.6
Women	14.6	24.2

Notes: Reported numbers are percentages. The numbers in Panel A use the reported levels of employment by management in each 1-digit occupation plus the establishment weights so they should be an estimate of the incidence of performance pay contracts in the British economy as a whole. The numbers in Panel B are an unweighted average of the respondents to the employee survey in WERS.

ones (though Dohmen and Falk (2010) also considered fixed-pay schemes). Our classification interprets all performance pay schemes (including piece-rates) as being situated within more competitive environments than those of fixed-pay schemes, though we do make some attempt to classify performance pay schemes into different types, which we describe below.

There are a number of reasons why we think that piece-rates are more competitive pay schemes than have been assumed in the experimental literature. First, in most work environments, people generally do not work in isolation; the differences in pay that result from piece-rates will typically be quite visible, and there will emerge clear "winners" and "losers." Second, workers' earnings under piece-rates are not independent. If some produce a lot under piece-rates, the piece-rate will typically be adjusted downward as employers seek to set total earnings to match what could be obtained in the outside labor market. In the folklore of the working-class, "rate-busters" were disliked because their greater output resulted in lower earnings for others as the piece-rate was adjusted downward.

Nonetheless, it could be argued that some forms of performance pay contracts are more competitive than piece-rates and can come closer to the tournaments considered in the experimental literature. For example, the most common form of merit pay is a system whereby the employer decides on a pot of money to be used for merit pay and then allocates this among a given number of individuals . In this case, there is an effective tournament such that a pound extra for one worker would require a pound less for everyone else. Of course these tournaments are nowhere near as extreme as the "winner-take-all" approach in many of the experimental studies, but it is important to study incentive schemes that are actually used. In what follows, we investigate whether piece-rates and merit pay differ in impact.

Do Women Shy away from Competition?

In this section, we discuss the incidence of performance pay schemes and consider whether women are underrepresented in jobs that have those schemes. Table 1 presents the incidence of performance pay schemes in the WERS data. Panel A presents data from the employee profile questionnaire completed by management together with the weights required to create a profile for the entire British working population; thus, this should be an estimate of the proportion of workers in Britain as a whole on performance pay contracts. In the 1998 sample, 16.3% of workers were in jobs that used performance pay; this rose to 32% in 2004. Although it is not surprising that the use of performance pay rose over this period, the magnitude of the rise does seem large. However, a direct question in the 2004 survey reveals that 12% of establishments reported that they had introduced performance-related pay in the

Table 2. The Nature of Performance Pay Contracts

	All	*Men*	*Women*
Merit Pay	35.3	36.1	40.8
Piece-Rates	35.9	37.1	34.5
Both	25.8	26.8	24.6
Sample Size	6069	3221	2848

Notes: These numbers reflect the percentage of workers in performance pay contracts whose employers report the use of merit pay, piece-rates, or both. Note that the presence of performance pay is defined at the occcupaton–plant level, but the nature of that contract is defined only at the plant level. This data is only available for 2004, so the sample sizes are smaller.

Table 3. The Measures Used to Evaluate Performance

	All	*Men*	*Women*
Individual	57.8	57.4	58.2
Team or Group	34.2	35.7	32.4
Workplace	25.4	26.8	23.8
Organization	32.2	32.2	32.2

Notes: These numbers reflect the percentage of workers in performance pay contracts whose employers report the use of difference measures of performance. Multiple answers are possible, so answers add up to more than 100%. Note that the presence of performance pay is defined at the occcupaton–plant level, but the nature of the measure used to evaluate performance is defined only at the plant level.

last two years, so the increases reported in Table 1 may not be implausible. This table also reports the prevalence of performance pay for men and women, a first indication of whether women are underrepresented in performance pay contracts. In 1998, men were 4 percentage points more likely than women to have performance pay contracts, and this gap rose to 8 percentage points in 2004. Panel B of Table 1 shows the incidence of performance pay among the individuals in the workers sample of WERS—the sample we will use for analysis. The proportions in Panel B are quite similar to those derived using the basic work force information reported in Panel A. These raw gender differences in the incidence of performance pay are much smaller than those typically reported in the experimental literature.

Is there any evidence that women shy away from particular types of performance pay contracts? Table 2 illustrates the proportion of men and women with performance pay who worked in plants with merit pay, payment-by-results, or both. There is no indication here of women avoiding the merit pay schemes that might be considered more competitive; if anything, women are more likely to be included in such schemes, though the gender gap is small. There may also be a difference in the level at which performance is assessed. For example, very competitive individuals might prefer individual-based schemes while others might prefer team or workplace-based schemes. Table 3 shows that the most common form of performance pay contracts are individual-based but that there is little raw difference between men and women in the level of the measures used to evaluate performance.

Tables 1–3 do not control for any other relevant factors, but, when we control for additional characteristics, the gender differential remains small and often insignificant. Table 4 estimates probit models of whether an individual works under a performance pay contract, with gender on the right-hand side and different sets of covariates. The first row includes a dummy for gender, the next row comprises personal demographics, the third row includes establishment characteristics, and

Table 4. Do Women Select out of Performance Pay Contracts?

Dependent Variable	*Female Coefficient [s.e.]*	*Personal Characteristics*	*Job Characteristics*	*Occupation*	*Firm Fixed Effects*	*Sample*
Performance Pay	–0.049 [0.007]	No	No	no	No	All
Performance Pay	–0.053 [0.007]	Yes	No	no	No	All
Performance Pay	–0.004 [0.007]	Yes	Yes	no	No	All
Performance Pay	–0.012 [0.006]	Yes	Yes	Yes	No	All
Performance Pay	–0.005 [0.002]	Yes	Yes	Yes	Yes	All
Merit Pay	0.014 [0.020]	Yes	Yes	Yes	No	Performance Pay
Piece-Rates	–0.005 [0.019]	Yes	Yes	Yes	No	Performance Pay
Individual-based performance pay	0.009 [0.017]	Yes	Yes	Yes	No	Performance Pay

Notes: All rows except the fifth report the marginal effects from a probit model; the fifth row is a linear probability model. Sample sizes are 47367 for the first five rows and 5780 for the last two. All standard errors are roboust and clulstered at the establishment–occupation level. Personal characteristics are education, race, age, job, tenure, marital status, and the presence of dependent children. Job characteristics are industry, log establishment size, public sector dummy, and union recognition.

the fourth indicates occupation dummies. The fifth row estimates a linear-probability model with establishment fixed effects. In all of these estimates, the extent to which women seem to shy away from performance-related pay is much smaller than the gaps reported in the experimental literature, and the inclusion of some regressors can fully explain the gap. In particular, the inclusion of job dummies turns a gap of 4–5 percentage points into something close to one-half a percentage point. We do not see strong support for the findings of Gneezy and Rustichini (2006), that the differences within occupations are smaller than the raw differential; in fact, the inclusion of occupation in the fourth row actually raises the gender gap slightly.

It is not immediately obvious which specification is the best to use as an estimate of the gender difference in performance pay contracts. If it is the case that women avoid competitive situations, this may show up in their choice of occupation and industry, whether they work in the public or the private sector, or whether they are unionized. In this case the best evidence would be the specification that controls only for personal characteristics. There may be other reasons, however, why women and men end up in different sectors, and factors apart from gender differences in the attitude to competition may explain the variation in the use of performance pay across sectors. It would then be wrong to attribute all of the raw gender gap in the incidence of performance pay contracts to different attitudes to competition among men and women. Since we are arguing that the overall explanatory power of the "competition" hypothesis is small, we err on the side of generosity and use estimates that include personal rather than job characteristics as our main specification. However, we do report estimates in some tables with different sets of controls to give readers some idea of the sensitivity of results to specification—these other specifications generally suggest performance pay is even less important in explaining the gender pay gap.

As we discuss above, one might claim that the measure of a performance pay contract used here does not correspond to the idea of "competition" that is used in the experimental literature. We might

Table 5. Do Women Select Out of Performance Pay Contracts? Disaggregation by Occupation

Dependent Variable: Worker on a Performance Pay Contract

Sample	*Coefficient On Female Variable*	*Standard Error On Female Variable*	*Coefficient On Female Variable*	*Standard Error On Female Variable*	*Sample Size*
Managers	–0.0247	[0.0177]	0.011	[0.018]	5244
Professionals	–0.11	[0.0151]	–0.0165	[0.013]	7109
Associate Professionals	–0.0929	[0.0155]	–0.0183	[0.013]	6228
Clerical	–0.0706	[0.0154]	–0.0217	[0.013]	9445
Craft	–0.0654	[0.0326]	–0.0163	[0.036]	3456
Personal Service	0.00145	[0.00837]	0.0112	[0.0048]	3697
Sales	–0.0344	[0.0264]	–0.0247	[0.026]	3416
Operatives	0.0536	[0.0329]	0.0595	[0.031]	3782
Other Occupations	–0.0877	[0.0163]	–0.0292	[0.012]	4923
Controls	Personal Characteristics		Personal + Job Characteristics		

Notes: All rows report the marginal effects from a probit model. All standard errors are robust and are clustered at the establishment–occupation level. The personal and job characteristics are those listed in the notes to Table 4.

investigate this hypothesis in a number of ways. First, we have argued that merit pay schemes are more likely to be "tournaments" than piece-rate schemes. The sixth and seventh rows of Table 4 restrict the sample to those on performance pay contracts and estimate probit models for having merit pay and payments-by-results. Gender gaps are small and insignificantly different from zero but, if anything, women are more likely to be part of merit pay systems. Second, it might be the case that since on average more women than men are averse to competition, they are less likely to end up in jobs that include individual-based performance pay (as opposed to group-based performance pay). The eighth row of Table 4 illustrates, however, that this does not seem to be the case—women and men on performance pay schemes are equally likely to be in schemes using individual measures of performance. The estimates in this table can be thought of as averages across all jobs, but perhaps only women in some jobs tend to shy away from competition. To investigate this, we estimated separate probit models for the nine one-digit occupations reporting estimates both with and without controls for job characteristics (see Table 5). None of the gender gaps is particularly large, but not all of them are negative and there is no discernible pattern to the coefficients.

Our studies suggest that although women are less likely than men to be found in performance pay contracts, the difference is much smaller than the size of effect reported in the experimental literature. However, we need to be able to say something about the effect of performance pay on earnings before we can say anything about the size of the contribution to the gender pay gap, the topic of the next section.

The Effects of Performance Pay on Earnings

We now turn to the effects on pay, beginning with an estimation of simple earnings functions pooling both men and women with a gender dummy. The earnings data in WERS is banded and we experiment with a number of ways of dealing with this—an interval regression model, with and without allowance for heteroskedasticity, as well as a linear regression. The first row of Table 6 includes nothing but a gender dummy showing a log hourly pay differential of 23 log points in line with other estimates (Anderson et al. 2001). The

Table 6. The Gender Pay Gap in WERS

Dependent Variable: Log Hourly Wages

Estimation Method	*Female Coefficient [s.e.]*	*Personal Characteristics*	*Job Characteristics*	*Occupation*	*Fixed Effects*	*Other Comments*
Interval Regression	–0.230 [0.007]	No	No	No	No	
Interval Regression	–0.207 [0.006]	Yes	No	No	No	
Interval Regression	–0.174 [0.006]	Yes	Yes	No	No	
Interval Regression	–0.142 [0.005]	Yes	Yes	Yes	No	
Interval Regression	–0.147 [0.006]	Yes	Yes	Yes	No	Allows for heteroskedasticity
Midpoint Regression	–0.142 [0.005]	Yes	Yes	Yes	No	
Midpoint Regression	–0.105 [0.005]	Yes	Yes	Yes	Firm	
Midpoint Regression	–0.092 [0.005]	Yes	Yes	Yes	Firm* Occupation	
Oaxaca decomposition	–0.208 [0.004]	Yes	No	No	No	Evaluated at average female characteristics
Oaxaca decomposition	–0.217 [0.004]	Yes	No	No	No	Evaluated at average male characteristics

Notes: The sample size is 45527. All standard errors are robust and are clustered at the establishment–occupation level. The personal and job characteristics are those listed in the notes to Table 4. The fifth row allows heteroskedasticity to be related to all variables. The Oaxaca decompositions estimate separate earnings functions for men and women using the midpoint method and then compute earnings gaps for the average woman and man.

next rows illustrate how this changes when one includes personal characteristics, job characteristics (including the gender mix on the job), and occupation. Inclusion of personal characteristics (the second row) reduces the gender pay gap to 21 log points and job characteristics (the third row) to 17 log points. When one includes occupation (the fourth row) the gap is reduced further, to 14 log points. The fifth row then shows that the estimate from the interval regression model is essentially identical if one allows for heteroskedasticity related to all the covariates, and the sixth row indicates that one obtains essentially the same results if one estimates a linear regression model using mid-points of the pay bands to assign weekly earnings (plus 1.4 times the top band). Because recognizing the banded nature of the data makes very little difference to the estimates, we use linear regressions in the remainder of this paper. The seventh row includes establishment fixed effects; the reduction in the estimated gender pay gap to 10 log points reveals that women are concentrated in low-paying occupations. The eighth column includes fixed effects for the plant-occupation combination, reducing the gender pay gap still further, though only modestly. Finally, the last two rows report the results of Oaxaca decompositions of the gender pay gap—the results are consistent with what others have found when studying the U.K. gender pay gap using other data sets (e.g., Anderson et al. 2001).

Table 7 also reports estimates of earnings functions but that now include the presence of performance pay as an additional regressor. The estimate in the first row includes no regressors other than gender and performance pay. The gender coefficient is similar to that reported in the similar specification in Table 6 because, as we have demonstrated in the previous section, women do not necessarily opt out

Table 7. The Effects of Performance Pay on Earnings
Dependent Variable: Log Hourly Wages

Female Coefficient [s.e.]	*Performance Pay Coefficient*	*Personal Characteristics*	*Job Characteristics*	*Occupation*	*Fixed Effects*
–0.228	0.171	No	No	No	No
[0.007]	[0.012]				
–0.205	0.128	Yes	No	No	No
[0.006]	[0.009]				
–0.175	0.086	Yes	Yes	No	No
[0.006]	[0.009]				
–0.142	0.046	Yes	Yes	Yes	No
[0.005]	[0.007]				
–0.105	0.025	Yes	Yes	Yes	Yes
[0.006]	[0.009]				

Notes: The dependent variable is the midpoint of the log hourly wage. Sample sizes are 45527. All standard errors are robust and are clustered at the occupation–establishment level. The personal and job characteristics are those listed in the notes to Table 4. The fourth row allows heteroskedasticity to be related to all variables.

of performance pay schemes. The presence of a performance pay scheme is estimated to raise wages by 17.5 log points. As the other rows of Table 7 illustrate, however, this effect is much reduced if one includes other covariates, reflecting the fact that high-level occupations are more likely to feature performance pay. Including personal controls reduces the effect to 12.8 log points, adding job controls reduces the effect to 8.6 log points, and including occupation reduces the effect still further, to 4.6 log points. If plant fixed effects are included, the estimate is only 2.5 log points, though this remains significantly different from zero.[2] These modest estimates of the effects of performance pay schemes on average earnings correspond to what other studies have found (e.g., Lemieux, MacLeod, and Parent 2009) though they stand in contrast to the very large effects found in the experimental and other studies of specific performance-related pay schemes (Lazear 2000; Bandiera, Barankay, and Rasul 2005). One possible reason for this is that most performance pay schemes are not very high-powered whereas the individual schemes that have been studied are.

The estimates of the gender pay gap in Table 7 are similar to those found when the performance pay variable was excluded, suggesting that performance pay schemes contribute little to our understanding of the gender pay gap. At the same time, it is important to note that these estimates so far have been based on the assumption that male and female earnings functions differ only in the intercept. We investigate whether relaxing this assumption makes any difference by estimating separate male and female wage equations. In Table 8, the first two columns report the coefficient on performance pay in male and female earnings functions. Contrary to what might be expected from the experimental literature, the returns to a performance pay scheme seem very similar for women and men.

The third and fourth columns of Table 8 indicate whether or not a merit-based or piece-rate pay system affects women's and men's earnings. These variables are only available for 2004, so the sample size is smaller, but there is no evidence of any difference in the effects of these pay schemes. The fifth and sixth columns investigate one of the findings of Gneezy et al. (2003), namely that performance in competition is affected by the gender mix in the job. We have information on the gender mix of the job done by the worker, which we divide

[2] We cannot go further and include plant–occupation fixed effects as in Table 8 since performance pay schemes are defined at this level.

Table 8. The Effects of Performance Pay on Earnings by Gender
Dependent Variable: Log Hourly Wages

	Men	*Women*	*Men*	*Women*	*Men*	*Women*
Performance pay	0.124	0.124	0.0922	0.1	0.109	0.11
	[0.0106]	[0.0113]	[0.0392]	[0.0417]	[0.0133]	[0.0139]
Performance pay* merit pay			0.0753	0.0672		
			[0.0323]	[0.0345]		
Performance pay* piece–rates			–0.0153	–0.0367		
			[0.0314]	[0.0344]		
Performance pay * mixed job					0.0358	–0.0003
					[0.0181]	[0.0180]
Performance pay * male job						0.0443
						[0.0365]
Performance pay * female job					–0.0345	
					[0.0342]	
Personal characteristics	Y	Y	Y	Y	Y	Y
Observations	21937	23548	9463	10854	21314	22776

Notes: See notes to Table 7. The personal controls are the same as those listed in Table 4.

into mostly male, mixed, and mostly female. The fifth column estimates a male earnings function in which the omitted category is the most common job mix—male-dominated. Interestingly, the data illustrate that there is a marginally significantly negative effect of being employed in a female-dominated job featuring a performance pay scheme, a finding different from that of Gneezy et al., who found that men perform better when competing against women; in other words, they found that performance in competition is affected by the gender mix. The seventh column estimates a female earnings function in which the omitted category is a female-dominated job. The evidence here does not suggest that the gender mix of the job influences the pay impact of a performance pay contract.

All of the estimates presented so far suggest a modest contribution of competition-based pay schemes to the overall gender pay gap. For example, if we assume a return to a performance pay contract of 12.4% (the first two columns of Table 8) and ask how the gender pay gap would be affected if there were no performance pay schemes, the gender pay gap would be reduced by 0.124*(0.238–0.192) = 0.0057 since 23.8% of men and 19.2% of women currently work under such a scheme. The estimated contribution is approximately half of one percentage point. The reason for this is simple—the returns to performance pay are modest, as is the gender difference in the incidence of performance pay.[3]

We make one final attempt to find a large effect of performance pay on earnings. The estimates presented so far assume performance pay affects only the intercept of the earnings function, but it might affect other coefficients as well (for example, Lemieux, Parent, and MacLeod (2009) found higher returns to education and lower returns to job tenure with performance pay). To investigate this without being too restrictive about the way in which performance pay affects earnings, we use a re-weighting estimate of the

[3] There is perhaps an echo here of the results of Manning and Swaffield (2008), who investigated the impact of psychological variables on earnings. Although they do find these variables affect earnings and that men and women are different, they fail to explain a sizeable effect of the gender pay gap since the differences are not large enough.

impact of performance pay along the lines first used by DiNardo, Fortin, and Lemieux (1996). Under the assumption (which we have been making throughout this paper) that the presence of a performance pay scheme is exogenous conditional on the included covariates, we can estimate a probit model for the presence of performance pay for men and women separately, and then re-weight those observations without performance pay to get an estimate of what the distribution of earnings would have been in the absence of performance pay. One can then take the difference between men and women to get an estimate of what the gender pay gap would have been in the absence of performance pay and compare this to the gender pay gap with performance pay to get an estimate of the contribution of competition to the gender pay gap. The result of this exercise is a gap of less than one percentage point, reinforcing our conclusion that performance pay explains little of the gender pay gap.

Potential Criticisms

Because we have failed to discover a noticeably important role that performance pay plays in explaining the gender pay gap, we believe this suggests there is little evidence that women's and men's differing attitudes toward and responses to competition in the labor market are important. It is, nevertheless, worth reflecting on how our conclusions might be misleading and why our results differ from the experimental evidence. In this section we discuss possible explanations.

First, our approach assumes that, conditional on the covariates, performance pay occurs as an exogenous variable and our research design excludes a random assignment of performance pay. Moreover, existing studies, both experimental (Dohmen and Falk 2010) and from the field (Lazear 2000), tend to find that the more able workers select into performance pay contracts. In this case, small though they are, we have most likely overestimated the effects of performance pay. Second, some might argue that the rewards from competition occur not in current wage schemes but in situations such as promotions, and so our outcome measure is not the most appropriate one. Existing British evidence (notably Booth, Francesconi, and Frank 2003), however, indicates that women actually achieve higher promotion rates than men and that women's lower average position in the occupational hierarchy is the result of career interruptions generally associated with having children that are probably not associated with the outcome of competition. To investigate this theory further, we entertain the idea that, whatever the rewards are, the purpose of incentive schemes is to encourage effort. To this end, we employ a Likert scale to determine the extent to which workers agree or disagree with the statement, "My job requires that I work very hard," coded from 1 (strongly disagree) to 5 (strongly agree). In Table 9 we use this as the dependent variable and estimate equations similar to the earlier earnings equations. The first column demonstrates that women report working significantly harder than men whereas the second column shows that this is robust to the inclusion of other controls. When the presence of performance pay is added, the coefficient on female is unaltered but we do find, as expected, that people work harder under incentive schemes. The fourth column estimates an equation only for men and the fifth, only for women. These estimates do suggest that the gap between effort with and without performance pay is significantly larger for men than for women though they are small in absolute terms. Hard work would have to have a very large material payoff for this differential to explain much of the gender pay gap.[4] The next two columns reveal whether the effort level exerted on the job is a direct effect of the gender mix. Men's effort response to performance pay is largest in female-dominated jobs while women seem to respond less to such an incentive in male-dominated jobs. This is one area where we find an effect that corresponds with

[4] Incidentally, if one includes hard work in an earnings function, it has a coefficient of 0.0038 with a t-statistic of 2. It is likely that this coefficient is biased, however, so not too much weight should be put on that.

Table 9. The Effects of Performance Pay on Work Effort

	1	2	3	4	5	6	7
Sample	All	All	All	Men	Women	Men	Women
Female	0.134 [0.00858]	0.145 [0.00857]	0.147 [0.00858]				
Performance pay			0.0345 [0.0112]	0.0735 [0.0152]	0.00937 [0.0153]		
Performance pay * male job						0.0417 [0.0193]	–0.0944 [0.0498]
Performance pay * mixed job						0.0976 [0.0212]	–0.0107 [0.0202]
Performance pay * female job						0.158 [0.0497]	0.00572 [0.0205]
Personal characteristics	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	46771	46771	46771	22526	24245	21904	23459

Notes: The dependent variable is the response to the question, "My job requires that I work very hard" with a 1 representing "strongly disagree" and 5 "strongly agree." The personal characteristics are the same as those listed in the notes to Table 4.

the experimental evidence; it is unlikely, however, that this can explain much of the gender gap in performance since very few men work in female-dominated jobs and likewise, very few women work in male-dominated jobs. One gets some indication of this in the fourth and fifth columns where the coefficient on performance pay can be considered approximately as a weighted average of the effects of working in jobs comprised of different gender mixes. Here the gender differences, though significant, are small.

So far, we have discussed two possible responses to our approach to performance pay and to the appropriateness of our outcomes measures for competitive rewards. One might also argue that the intensity of competition within a job is not well measured by the presence of a performance pay contract (although the evidence presented does suggest that performance pay raises earnings and increases work effort in line with what theory would predict). It is likely, however, that a degree of misclassification is induced by our equation of performance pay contracts with "competitive" pay schemes. We argue that it is hard for any model of misclassification to reconcile our estimates with a sizeable gender difference in the incidence of competitive contracts, a higher return to "competition" among men (both findings that are emphasized by the experimental literature), and a sizeable true effect of competition on the gender pay gap. We start first with our measurement of the contribution of "competition" to the overall gender pay gap. The measure we use is how much the gap would change if there were no competitive pay schemes. This is given by:

$$\overline{C}^*_m \beta^*_m - \overline{C}^*_f \beta^*_f \quad (1)$$

where $\overline{C}^*_m(\overline{C}^*_f)$ is the true fraction of men (women) in competitive pay schemes and β^*_m (β^*_f) is the true returns for men (women) from being in a competitive pay scheme.

Now let us assume that C^* is measured with error. Assume that the probability of being wrongly classified is r_1 if $C^* = 1$ and r_0 if $C^* = 0$. Denote the observed variable (performance pay in our application) by C. With the assumptions about misclassification we will have

$$\overline{C} = (1 - r_1)\overline{C}^* + r_0(1 - \overline{C}^*) \quad (2)$$

This can be re-arranged to give:

(3) $$(1 - \overline{C}) = r_1 + (1 - r_1 - r_0)(1 - \overline{C}*)$$

This misclassification will tend to produce attenuation bias in the returns to performance pay. Using the formula in Freeman (1984, equation 8) this is given by:

(4) $$p \lim \hat{\beta}_i = \beta *_i (1 - r_1 - r_0) \frac{\overline{C}*_i (1 - \overline{C}*_i)}{\overline{C}_i (1 - \overline{C}_i)}, \quad i = m,f$$

Re-arranging this, we have that:

(5) $$\beta *_i \overline{C}*_i = \hat{\beta}_i \frac{\overline{C}_i (1 - \overline{C}_i)}{(1 - r_1 - r_0)(1 - \overline{C}*_i)}, \quad i = m,f$$

where the left-hand side is the contribution of performance pay to the earnings of men and women as defined in (1). Now, using (3), (5) can be written as:

(6) $$\beta *_i \overline{C}*_i = \hat{\beta}_i C_i \frac{(1 - \overline{C}_i)}{(1 - \overline{C}_i) - r_1}, \quad i = m,f$$

Note that the term $\hat{\beta}_i \overline{C}_i$ (both of whose elements are observable) is the contribution we would use assuming there is no measurement error. The final term on the right-hand side of (6) must be larger than one, indicating that measurement error will attenuate the effects. It is interesting to note that the attenuation bias depends only on r_1, the misclassification rate for those who work in competitive pay schemes. The misclassification rate for those who are not in competitive pay schemes, r_0, plays no role in (6). One implication of this is that if our classification of competitive pay schemes is incorrect in the sense that all the fixed pay schemes we classify as non-competitive are *truly* non-competitive, but some of the performance pay schemes we classify as competitive are, in reality, non-competitive, our estimates of the contribution of performance pay to the gender pay gap will be correct even though there will be attenuation bias in our estimates of the return to "competition." The reason for this is that although we are underestimating the returns to competition, we are overestimating the proportion of workers on competitive pay schemes, and these two effects work in opposite directions and cancel each other out.

If some of the fixed pay schemes are truly competitive, then $r_1 > 0$ and we will have underestimated the importance of competition. However, one has to have very large misclassification rates for the bias to be large. For example, suppose we observe 25% of men and 20% of women in performance pay contracts (an average across the years observed in Table 1). Assume further that the estimated response of hourly wages to working under a performance pay contract is 0.124 (a generous estimate taken from the first two columns of Table 8). Using these numbers, we can compute the "true" contribution of "competition" to the gender pay gap for different values of r_1. The results are shown in Panel A of Table 10. The first row illustrates our measured effect with no misclassification. As we discuss above, this is very small. But what is striking from Panel A is that the misclassification rate must be more substantial to produce any larger effects. A misclassification rate of 60%, for example, is required for the true contribution of "competition" to the gender pay gap to be 5 log points, about one-quarter of the total. Even if the misclassification rate were that high, the "competition" story runs into other problems. First, a misclassification rate of 60% on "competitive pay contracts" puts, from a rearrangement of (2), an upper bound on the misclassification rate for "non-competitive" pay contracts of 20% given that we observe 20% of women in performance pay contracts.

Panel B of Table 10, then, takes a misclassification rate of 60% for "competitive pay contracts" and shows, for different values of r_0, the implied true gender difference in the incidence of "competitive pay contracts" (from (2)) and the implied true rates of return to being in a competitive contract for men and

Table 10. The Effects of Misclassification of Pay Contracts

Panel A: The Effect of Misclassifying "Competitive Contracts" on the Contribution of "Competition" to the Gender Pay Gap

Misclassification rate for "competitive" pay schemes, r_1	*Contribution of "competition" to the gender pay gap*
0	0.006
0.1	0.007
0.2	0.009
0.3	0.012
0.4	0.017
0.5	0.027
0.6	0.056
0.65	0.100
0.7	0.267

Notes: This table uses the formula in (6), an estimated observed return to performance pay of 0.124 (from Table 8), and an observed incidence of performance pay for men of 25% and for women of 20%.

Panel B: The Effect of Misclassifying "Non-Competitive Contracts"

Misclassification rate for "non-competitive" pay schemes, r_0	*True gender difference in incidence of "competitive contracts"*	*Implied true return to performance pay for women*	*Implied true return to performance pay for men*
0	0.13	0.20	0.25
0.025	0.13	0.21	0.26
0.05	0.14	0.23	0.27
0.075	0.15	0.26	0.29
0.1	0.17	0.30	0.31
0.125	0.18	0.36	0.34
0.15	0.20	0.50	0.39
0.175	0.22	0.89	0.47

Notes: This table uses the formulae in (2) for the second column and (4) for the last two, an estimated observed return to performance pay of 0.124 (from Table 8), and an observed incidence of performance pay for men of 25% and women of 20%, and a misclassification rate of 0.6 for "competitive"pay contracts.

women (from (4)). Note that if we choose a value of r_0 that implies a gender gap in the true incidence that is close to 20% as the experimental literature suggests, then this implies that women have a higher return to performance pay than men, the opposite of the experimental findings. Our bottom line is that it is difficult to find any pattern of misclassification errors to reconcile all of our findings with the experimental literature even before we start to debate whether such misclassification rates are plausible.[5] Even these computations are based on what are probably overestimates of the importance of performance pay; other estimates in this paper suggest even more modest effects. We would therefore propose that misclassification of "competitive" pay schemes cannot explain our results.

We address one final possible response to our claim that performance pay may not fully explain the gender pay gap. It might be argued that while differing gender attitudes to competition can explain little of the gender pay gap among the average worker, they are better able to explain this among

[5] It should be noted that we have assumed misclassification errors to be the same for men and women. Introducing gender differences in error rates allows for enough degrees of freedom to solve the problem we identify here but seems a slim thread on which to hang the theory.

Table 11. The Effects of Performance Pay on Earnings in Managerial and Professional Jobs

	1	*2*	*3*	*4*
Sample	All	All	Men	Women
Female	–0.159	–0.146		
	[0.00907]	[0.00890]		
Performance Pay		0.137	0.131	0.144
		[0.0117]	[0.0140]	[0.0172]
Other controls	Yes	Yes	Yes	Yes
Observations	11796	11796	6882	4914
R^2	0.32	0.33	0.33	0.31

Notes: This is the same as Table 7 but with the sample restricted to managers and professionals.

senior managers and professionals. Indeed, both Gneezy et al. (2003) and Niederle and Vesterlund (2007) motivate their studies by citing the under-representation of women in senior jobs. However, the actual participants in their experiments (students) and the tasks they undertake are not particularly demanding (though see Gneezy and Rustichini (2006) for experimental evidence from teachers and executives). It is reasonable to conclude, therefore, that their experimental evidence is of sizeable gender differences in the attitudes to competition among average workers doing an average job. To check these ideas, however, we do run separate regressions for managers and professionals. The first column of Table 11 reveals that the gender pay gap among managers and professionals is lower than the average, not what one might expect if competition is more intense among this group and women fare badly in that. The second column then includes a dummy variable for performance pay; this has a significant positive effect somewhat larger than that estimated for all workers though still fairly modest. Finally, the third and fourth columns estimate separate equations for men and women—if anything, the return to performance pay is higher for women than men. The bottom line is that, even among managers and professionals, there seems to be little evidence of competition being an important aspect of the gender pay gap.

Conclusions

Recent evidence from laboratory experiments suggests that men and women have different attitudes toward and responses to competition in the workplace. These studies indicate, in fact, that a sizeable part of the gender pay gap could be explained in this way. These results may better indicate a direction for future research, however, than provide reliable estimates of gender pay differences attributable to these factors. In this paper, we have attempted to provide an estimate of the portion of the gender pay gap in the United Kingdom that can be attributed to these differing attitudes to competition. We find modest evidence for differential sorting into performance pay schemes by gender and small effects of performance pay on hourly wages. Furthermore, and unlike existing laboratory studies, we find that the gender mix of a given profession produces no significant effect on the responsiveness to performance pay. We have discovered evidence, however, that performance pay affects work effort, corresponding to the experimental evidence discussed above. The bottom line, though, is that a very small part of the gender pay gap can be attributed to these factors. Even an attempt to maximize the possible explanatory power of this hypothesis does not lead to very large contributions. Although we do not have an experimental design, we are reassured by studies that draw from the real world (see, e.g., Bandiera, Barankay, and Rasul 2005; Paarsch and Shearer 2007; Paserman 2007; Lavy 2008) and that also fail to find any large gender gaps.[6] These are studies of particular labor markets, and though it is not clear to what extent they represent the labor market as a whole, our results suggest their conclusions may well be valid more generally.

Why, then, do the experimental findings fail to show up in literature on field studies? We have no definitive answer and can only suggest possibilities. First, the sample sizes in the experimental studies are often not large (sometimes less than 100) so the standard errors will typically be large and one needs to find dramatic effects for them to be significantly different from zero. To obtain interesting results, the experiments may be designed to magnify differences, something the market would not generally do (see Lazear et al. 2004 for a similar point). This may mean that the treatments considered are quite extreme. For example, the tournaments typically employ a "winner-take-all" structure in which those who do not win get nothing (and, in reality, presumably starve to death) and the piece-rates have a one-for-one incentive not seen even in the most high-powered managerial incentive schemes (e.g., Murphy 1999). Moreover, experimental design may affect the results in important ways; indeed, the experiments' conclusions themselves sometimes vary regarding gender differences. For example, gender performance differs significantly in tournaments in Gneezy et al. (2003), but not in Niederle and Vesterlund (2007). Additionally, the results of Antonovics, Arcidiacano, and Walsh (2005, 2008) suggest that the size of the stakes may be important. It is also possible that the laboratory results on performance pay are valid, but other factors are also at work in the field that mitigate or offset these effects. For example, it may be more difficult to discriminate against women in performance pay contracts where pay and productivity end up more closely aligned (see Heywood and O'Halloran 2005; and Fang and Heywood 2006, for similar ideas applied to racial wage differentials).

All of these alternative explanations imply first that it is problematic simply to map the findings of the laboratory studies onto the fraction of the gender pay gap that we observe in actual labor markets. One should perhaps regard the laboratory studies as suggesting directions for further research rather than as good estimates of effects in actual labor markets. The bottom line here seems to be that little of the U.K. gender wage gap can be explained by gender differences in the incidence of and response to performance-related pay.

[6] See Ors, Palomino, and Peyrache (2008), who studied gender differences in exams finding greater dispersion among men than women in competitive exams.

Appendix
Descriptive Statistics and Representative Earnings Functions and Performance Pay Equations

	Descriptive statistics–all	*Descriptive statistics–men*	*Descriptive statistics–women*	*Probit for performance pay*	*Earnings function*
Log hourly wage	2.013	2.131	1.903		
	0.581	0.579	0.560		
Female	0.520	0.000	1.000	-0.012	-0.142
	0.500	0.000	0.000	[0.006]	[0.005]
Performance pay	0.212	0.236	0.190		0.046
	0.409	0.425	0.392		[0.007]
CSE or equivalent	0.100	0.103	0.098	0.000	-0.064
	0.301	0.304	0.297	[0.007]	[0.007]
A-level	0.154	0.148	0.159	-0.004	0.058
	0.361	0.355	0.366	[0.006]	[0.006]
Degree or equivalent	0.196	0.212	0.181	0.014	0.182
	0.397	0.409	0.385	[0.007]	[0.007]
Postgraduate	0.068	0.075	0.061	0.048	0.251
	0.251	0.263	0.240	[0.012]	[0.010]
No qualification	0.194	0.210	0.179	-0.009	-0.137
	0.396	0.407	0.384	[0.007]	[0.007]
Other qualification	0.029	0.029	0.029	-0.022	-0.041
	0.168	0.168	0.168	[0.011]	[0.014]
Age 25-29	0.193	0.181	0.204	0.034	0.236
	0.395	0.385	0.403	[0.014]	[0.015]
Age 30-39	0.265	0.274	0.256	0.042	0.377
	0.441	0.446	0.436	[0.015]	[0.015]
Age40-49	0.264	0.261	0.267	0.024	0.402
	0.441	0.439	0.442	[0.015]	[0.016]
Age 50-59	0.201	0.198	0.203	0.020	0.404
	0.400	0.399	0.402	[0.016]	[0.016]
Age 60+	0.041	0.052	0.031	-0.007	0.329
	0.199	0.222	0.175	[0.018]	[0.019]
White	0.933	0.936	0.931	-0.024	0.002
	0.249	0.245	0.253	[0.011]	[0.010]
Kids	0.402	0.428	0.377	-0.007	0.023
	0.490	0.495	0.485	[0.009]	[0.009]
Married	0.687	0.700	0.676	-0.009	0.042
	0.464	0.458	0.468	[0.009]	[0.008]
Married*Kids	0.340	0.316	0.363	0.007	0.012
	0.474	0.465	0.481	[0.010]	[0.010]
Tenure Less than 1 year	0.160	0.153	0.167	-0.006	-0.018
	0.367	0.360	0.373	[0.007]	[0.008]
Tenure 2 to 3 years	0.247	0.233	0.260	0.002	0.043
	0.431	0.423	0.439	[0.007]	[0.007]
Tenure 5 to 10 yrs	0.205	0.199	0.211	-0.012	0.069
	0.404	0.400	0.408	[0.007]	[0.007]
Tenure 10+ years	0.262	0.295	0.230	-0.012	0.121
	0.439	0.456	0.421	[0.008]	[0.008]
Log plant size	4.791	4.924	4.667	0.031	0.027
	1.395	1.332	1.439	[0.003]	[0.002]

Continued

Appendix
Descriptive Statistics and Representative Earnings Functions and Performance Pay Equations Continued

	Descriptive statistics–all	*Descriptive statistics–men*	*Descriptive statistics–women*	*Probit for performance pay*	*Earnings function*
Manufacturing	0.145	0.220	0.076	–0.042	–0.124
	0.352	0.414	0.265	[0.022]	[0.014]
Construction	0.048	0.078	0.019	–0.033	–0.081
	0.213	0.268	0.137	[0.026]	[0.016]
Wholesale and retail	0.114	0.102	0.125	–0.001	–0.244
	0.318	0.303	0.331	[0.027]	[0.016]
Hotels and restaurants	0.034	0.028	0.040	–0.067	–0.375
	0.182	0.166	0.195	[0.024]	[0.019]
Transport and communication	0.063	0.097	0.031	–0.055	–0.097
	0.242	0.296	0.174	[0.024]	[0.018]
Financial services	0.059	0.049	0.069	0.160	–0.027
	0.236	0.215	0.254	[0.041]	[0.017]
Other business services	0.101	0.105	0.097	–0.002	–0.097
	0.301	0.307	0.296	[0.026]	[0.017]
Public administration	0.093	0.099	0.088	–0.030	–0.102
	0.291	0.299	0.283	[0.029]	[0.018]
Education	0.123	0.068	0.173	–0.173	–0.286
	0.328	0.253	0.378	[0.013]	[0.016]
Health	0.138	0.054	0.214	–0.163	–0.229
	0.345	0.227	0.410	[0.014]	[0.016]
Other community services	0.052	0.052	0.051	–0.055	–0.203
	0.221	0.223	0.220	[0.024]	[0.020]
Year (1=2004)	0.442	0.427	0.454	0.135	0.302
	0.497	0.495	0.498	[0.010]	[0.006]
Union recognition	0.483	0.479	0.487	0.062	0.008
	0.500	0.500	0.500	[0.011]	[0.007]
Public Sector	0.334	0.264	0.399	–0.008	0.046
	0.472	0.441	0.490	[0.017]	[0.009]
Managers	0.110	0.150	0.074	0.156	0.401
	0.313	0.357	0.261	[0.025]	[0.013]
Professionals	0.148	0.161	0.136	0.045	0.364
	0.355	0.367	0.343	[0.022]	[0.013]
Associate Professionals	0.131	0.131	0.131	–0.007	0.210
	0.337	0.337	0.337	[0.020]	[0.013]
Clerical	0.199	0.082	0.306	0.047	0.023
	0.399	0.275	0.461	[0.021]	[0.012]
Personal Services	0.080	0.051	0.107	–0.098	–0.067
	0.271	0.220	0.309	[0.019]	[0.017]
Sales	0.072	0.043	0.099	0.049	–0.140
	0.258	0.202	0.299	[0.028]	[0.016]
Operatives	0.081	0.133	0.033	–0.031	–0.155
	0.273	0.340	0.178	[0.021]	[0.015]
Elementary	0.106	0.112	0.100	–0.050	–0.260
	0.308	0.316	0.300	[0.019]	[0.013]
Job done only by men	0.123	0.253	0.002		
	0.328	0.435	0.049		
Job done mainly by men	0.189	0.344	0.046		
	0.392	0.475	0.209		
Job done by men+women	0.354	0.351	0.357		
	0.478	0.477	0.479		
Job done mainly by women	0.267	0.052	0.467		
	0.442	0.221	0.499		
Job done only by women	0.067	0.001	0.128		
	0.250	0.031	0.334		
Observations				47367.000	45485.000
R^2					0.510

Notes: The first three columns show means and standard deviations of the variables used for both genders together and for men and women. The fourth column show the coefficients and standard errors for a probit regression where the dependent variable is having a performance pay contract: marginal effects are reported. The fifth column shows the coefficients and standard errors of an earnings function where the dependent variable is log hourly wages.

REFERENCES

Altonji, Joseph, G. and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Volume 3, pp. 3143–3259. Amsterdam: Elsevier.

Anderson, Tracy, John Forth, Hilary Metcalf and Simon Kirby. 2001. "The Gender Pay Gap." Report to the Women and Equality Unit, London, U.K.

Antonovics, Kate, Peter Arcidiacono and Randy Walsh. 2005. "Games and Discrimination: Lessons from the Weakest Link." *Journal of Human Resources,* Vol. 40, No. 4, pp. 918–947.

____, Peter Arcidiacono, and Randy Walsh. 2008. "The Effects of Gender Interactions in the Lab and the Field," *Review of Economics and Statistics*, Forthcoming.

Babcock, Linda and Sara Laschever. 2003. *Women Don't Ask: Negotiation and the Gender Divide*, Princeton: Princeton University Press.

Bandiera, Oriana, Iwan Barankay and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data." *Quarterly Journal of Economics*, Vol. 120, No. 3, pp. 917–962.

Blau, Francine D. and Kahn, Lawrence M., 2006 "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence." *Industrial and Labor Relations Review*, Vol. 60, No. 1, pp. 45–66.

Booth, Alison, Marco Francesconi, and Jeff Frank. 2003. "A Sticky Floors Model of Promotion, Pay, and Gender." *European Economic Review*, Vol. 47 No. 2, pp. 295–322.

Borjas, George J. 1981. "Job Mobility and Earnings over the Life Cycle." *Industrial and Labor Relations Review*, Vol. 34, No. 3, pp. 365–376.

Cully, Mark, Stephen Woodland, Andrew O'Reilly, and Gill Dix. 1999. *Britain at Work: The 1998 Workplace Employment Relations Survey*. London: Routledge.

Datta Gupta, Nabanita, Anders Poulsen, and Marie Claire Villeval. 2005. "Male and Female Competitive Behavior: Experimental Evidence." IZA Discussion Paper No. 1833.

DiNardo John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages 1973–1992: A Semiparametric Approach." *Econometrica*, Vol. 64, No.5, pp. 1001–44.

Dohmen, Thomas and Armin Falk. 2010. "Performance Pay and Multi-Dimensional Sorting: Productivity, Preferences and Gender." *American Economic Review*, forthcoming.

Fang, Tony and John S. Heywood. 2006. "Output Pay and Ethnic Wage Differentials: Canadian Evidence." *Industrial Relations*, Vol. 45, No. 2, pp. 173–194.

Freeman, Richard B. 1984. "Longitudinal Analysis of the Effects of Trade Unions" *Journal of Labor Economics*, Vol. 2, No. 1, pp. 1–26.

Gneezy, Uri, Kenneth L. Leonard, and John A. List. 2009. "Gender Differences in Competition: Evidence from a Matrilineal and a Patriarchal Society." *Econometrica,* Vol. 77, No. 5, pp. 1637-1664.

____, Muriel Niederle, and Aldo Rustichini. 2003. "Performance in Competitive Environments: Gender Differences." *Quarterly Journal of Economics*, Vol. 118, No. 3, pp. 1049–1074.

____, and Aldo Rustichini. 2006. "Executives versus Teachers: Gender, Competition and Self-Selection." Working Paper, UC San Diego.

Heywood, John S. and Patrick L. O'Halloran, 2005. "Racial Earnings Differentials and Performance Pay." *Journal of Human Resources*, Vol. 40, No.2, pp. 435–452.

Kersley, Barbara, Carmen Alpin, John Forth, Alex Bryson, Helen Bewley, Gill Dix, and Sarah Oxenbrudge. 2006. *Inside the Workplace: Findings from the 2004 Workplace Employment Relations Survey*. London: Routledge.

Lavy, Victor. 2008. "Gender Differences in Market Competitiveness in a Real Workplace: Evidence from Performance-based Pay Tournaments among Teachers." NBER Working Paper No. 14338, September.

Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review*, Vol. 90. No. 5, pp. 1346–1361.

____, Ulrike Malmendier, and Roberto Weber. 2004. "Sorting in Experiments." Unpublished, Stanford University.

Lemiuex, Thomas, W. Bentley Macleod, and Daniel Parent. 2009. "Performance Pay and Wage Inequality." *Quarterly Journal of Economics*, Vol. 124, Issue 1, pp. 1–49.

Manning, Alan and Joanna Swaffield. 2008. "The Gender Gap in Early-Career Wage Growth." *Economic Journal*, Vol. 118, No. 530, pp. 983–1024.

Murphy, Kevin J. 1999. "Executive Compensation." In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 3. pp. 2485–2563. Amsterdam: Elsevier.

Niederle, Muriel and Lise Vesterlund. 2007. "Do Women Shy Away From Competition?" *Quarterly Journal of Economics*, Vol. 122, No. 3, pp. 1067–1101.

Örs, Evren, Frédéric Palomino and Eloïc Peyrache. 2008. "Performance Gender-Gap: Does Competition Matter?" CEPR Discussion Paper No. 6891, June.

Paarsch, Harry J. and Bruce Shearer. 2007. "Do Women React Differently to Incentives? Evidence from Experimental Data and Payroll Records." *European Economic Review*, Vol. 51, No. 7, pp. 1682–1707.

Paserman, M. Daniele. 2007. "Gender Differences in Performance in Competitive Environments: Field Evidence from Professional Tennis Players." IZA Discussion Paper No. 2834.

Save-Soderbergh, Jenny. 2007 "Are Women Asking for Low Wages? *Gender Differences in Wage Bargaining Strategies and Ensuing Bargaining Success*." Swedish Institute for Social Research, Working Paper No. 7/2007.