Rivista internazionale di scienze sociali, 2020, n. 1, pp. 51-66

SORTING, SOCIAL COMPARISON AND WOMEN'S JOB SATISFACTION

Margherita Agnoletto*, Lorenzo Cappellari**, Astrid Gamba***

ABSTRACT

Using linked employer-employee data for the UK we address competing explanations for gender gaps in job satisfaction. Previous studies have rationalized the puzzling greater satisfaction of women either by pointing out gender differences in competitive attitudes, or through differences in sorting across jobs and industries. Our data allow us to test both explanations within a unified framework. The employer-employee structure of the data enables us to control for workplace unobserved heterogeneity that drives sorting. Moreover, we exploit information on workplace average wages to investigate workers' attitudes through the framework of social comparison within the firm. We show that while social comparison matters empirically, gender differences in social comparison are not enough to account for job satisfaction gaps. Instead, controlling for workplace heterogeneity resolves the puzzle, lending support to the sorting hypothesis.

Keywords: Job satisfaction, Gender, Social comparison, Firms, Linked employer-employee data.
JEL Classification: J16, J28, D91.
ISSN: 0035-676X (print); 1827-7918 (digital)
DOI: 10.26350/000518_000041

1. INTRODUCTION

The greater women's job satisfaction (WJS henceforth) is an empirical regularity emerging from workers' surveys in many different contexts, and a longstanding puzzle for labour market analysis and policy given the well documented wage penalty that women face. Past research has explained the puzzle by focusing on women's attitudes, such as lower competitiveness and risk taking associated with lower expectations towards the job, or on the sorting of women across industries and occupations; see Vladisavljevi and Perugini (2018) for a recent account of the literature. According to the first interpretation, women shy away from competition against men, and are therefore more satisfied than men for any given set of job attributes anything else equal. The sorting hypothesis, instead, maintains that women end up in industries or occupations that are conducive to higher satisfaction. While

* Margherita Agnoletto, Catholic University of Milan. Email: margherita.agnoletto@gmail.com.

** Lorenzo Cappellari, Department of Economics and Finance, Catholic University of Milan. Email: lorenzo.cappellari@unicatt.it.

*** Astrid Gamba, Department of Economics, University of Insubria. Email: astrid.gamba@uninsubria.it.

the two alternative hypotheses have found varying degrees of empirical support, they have never been put to the test in conjunction, raising issues of comparability across studies. In this paper, we contribute to the literature by providing a unified framework of analysis within which we jointly consider the two explanations such as we are able to assess their comparative merits in accounting for the puzzle of higher WJS.

We develop our analysis by making use of the 2011 British Workplace Employment Relations Study (WERS), a survey-based data source that has a linked employer-employee structure which makes it representative of the universe of both firms and workers in the UK. Besides providing a detailed list of job satisfaction indicators, there are two features of WERS that make it ideal for our purposes. First, the linked structure implies that for each firm we observe several employees and can therefore control for unobserved heterogeneity between firms, the driver of workers' sorting into firms. Second, we have information on the workplace average wage, which we use to address workers' competitive attitudes. More specifically, we consider the heterogeneous satisfaction effects of individual deviations from workplace average wages as reflecting differences in competitiveness. These deviations are an example of social comparison, which has a long tradition in economics (Veblen, 1899; Duesenberry, 1949) and has been shown to matter for explaining job satisfaction (Clark and Oswald, 1996; Card et al., 2012). We model the satisfaction effect of social comparison following Prospect Theory (Kahneman and Tversky, 1979) according to which, divergence from the reference point has greater impacts on utility depending on whether it occurs from below rather than above that is, in Kahneman and Tversky's words, *losses loom larger than gains* (Kahneman and Tversky, 1979: 279). This mechanism looks well suited for modelling social comparison and job satisfaction within the workplace, but while it has received considerable attention in laboratory experiments or in field-based case studies (see Gächter and Thöni, 2010; Gamba, Manzoni and Stanca, 2017; Card et al., 2012) we are unaware of its implementation on nationally representative firm-level data, and ours is the first evidence on social comparison within the (nationally representative) firm. By using the wage average from a group of subjects with whom each employee actually interacts (co-workers), our reference point is better suited to study social comparison than the aggregate wage averages that typically feature in job satisfaction analysis to control for comparison income.

We begin our analysis by estimating the standard job satisfaction model in which satisfaction is related to a set of demographics and wages, and obtain the classic puzzling finding of greater WJS.

Next, we set up an estimable version of social comparison and explore gender differences in its impact on satisfaction. WERS reports in the manager's questionnaire the workplace specific wage histogram, from which we can compute the workplace specific average wage, that we use as the reference point for social comparison. We next construct the percentage individual deviation from the reference point, and regress job satisfaction on this deviation and its square, interacted

with an indicator for the sign of the deviation by gender. Our specification is driven by Kahneman and Tversky's (1979) consideration that the intensity of the effect should be stronger for negative deviations, while the concavity of the utility function should change from positive to negative after passing the reference point. We do find that social comparison has stronger impact on job satisfaction below rather than above the reference point, and also that the curvature of the relationship changes from positive to negative after passing the reference point, providing strong support for the existence of workplace-based social comparison. When interacting individual deviations from the reference point with gender, we do find that women care less about social comparison, which is consistent with the idea that they have lower competitiveness. However, the puzzling greater WJS persists also in this specification of the model, suggesting that gender differences in attitudes cannot fully account for it.

We next turn to consider the sorting explanation, and we take advantage of the detailed information on occupation and job characteristics available in WERS to saturate the model, showing that controlling for workplace characteristics almost entirely eliminates the WJS premium. We confirm that sorting into workplace is the driver behind the greater WJS by estimating a job satisfaction equation with workplace fixed effects, thanks to which any source of heterogeneity, both observed and unobserved, is eliminated from the data. We therefore conclude that the sorting hypothesis accounts for the greater WJS.

2. RELATED LITERATURE

Previous empirical studies of job satisfaction have shown that women report equal or greater satisfaction than do men. It is called the gender-job satisfaction paradox (Kaiser, 2007) or the paradox of the contented female worker (Bender, Donohue and Heywood, 2005). Indeed, the women's higher job satisfaction is surprising considering that there is wide evidence that they earn less, and they have a disadvantaged position in terms of hiring/firing and promotion areas.

According to a strand of literature, job satisfaction is associated with workers' expectations. Women may be more satisfied than men because they have lower expectations about their job due to their traditionally poorer position in the labour market (Clark, 1997). Differences in workers' expectations could explain why the job satisfaction gender gap disappears for younger, higher-educated, professionals or women in male-dominated workplaces (*ibidem*), because in those groups, women have higher expectations than usual.

Akin to the expectation hypothesis, the relationship between job satisfaction and gender can be caused by the difference in preferences and weights women and men place on job. Males prefer competitive jobs with high pay, responsibility and leadership, while females are first interested in having a good relationship with their co-workers and receiving positive feedbacks. A rapidly growing literature shows that females (even those with high abilities) shy away from competitive set-

tings and when in competition they tend to perform worse than males (Niederle and Vesterlund, 2007). Niederle and Vesterlund's laboratory experiment finds that in addition to gender difference in preferences for competition, the gender gap in competition is also due to men being more overconfident than women, while differences in risk and feedback aversion seem to play a negligible role.

More recently, De Paola, Ponzo and Scoppa (2017) suggest that the difference in competitiveness could be caused by women being more risk averse and less self-confident. They show that low ability women have a lower probability of applying for competition, while high ability ones act similar to men. Moreover, they find women's willingness to enter competition is also related to past success in promotions. Indeed, they are even more reluctant to compete in fields in which they are expected to face difficulties in pursuing their careers.

Another strand of literature considers differences in satisfaction as related to the sorting of women across firms which entails the overrepresentation of women in certain occupations or industries which provide greater satisfaction. Bender, Donohue and Heywood (2005) suggest that the higher WJS in female dominated jobs is associated to unmeasured characteristics of those jobs that women value and in particular, they find that the satisfaction differential disappears once flexibility between work and home is taken into consideration. Then, Zou (2015) using data from the 2006 Skills Survey shows that the gender satisfaction differential disappears if we control for work orientations. While men value job rewards more than other job characteristics, women are more interested in social relations and flexible work hours. Recently, Redmond and McGuinness (2019) confirm that controlling for job preferences the gender-job satisfaction gap disappears. Specifically, their results reveal two factors which mainly determine the gender gap: the higher importance placed by women on a good work-life balance and on the intrinsic desirability of the work.

Finally, Hauret and Williams (2017) argue that a gender convergence is possible; they perform an Oaxaca decomposition that shows that a small part of the gender differential can be explained by gender differences in weights placed on working conditions, and it is better explained by gender differences in job characteristics and thus it will disappear when sociodemographic, occupational, and workplace characteristics converge.

3. DATA

The dataset used in this paper is drawn from the linked employer-employee British Workplace Employment Relations Study 2011 (WERS), which addresses employees in firms with a minimum size of 5. WERS has the following components: a survey of managers, a survey of worker representatives, a financial performance questionnaire (for trading sector workplaces only) and a survey of employees. The first three sections yield establishment level information, whereas the final section provides employee level information. WERS 2011 is the sixth edition of

the survey, covering 2,680 workplaces and 21,981 employees, with information collected between March 2011 and June 2012.

We use two elements of the survey, the survey of managers composed by the employee profile questionnaire and the management questionnaire, and the survey of employees. The employee profile questionnaire is a self-completion questionnaire that was distributed to the most senior manager responsible for employment relations and personnel issues, before the management questionnaire, a face-to-face interview to collect information on the basic characteristics of the workforce. Then, the surveys of employees, that are self-completion questionnaires, were distributed to a randomly selected representative sample of 25 employees in workplaces where management permitted it. Employees' probability of selection for the survey is the product of the probability that their workplaces is selected and the probability of their own selection. If the workplace had 25 or fewer employees, all were selected to participate.

Thanks to the appropriate weighting, the WERS sample is considered nationally representative of British employees working in workplaces with 5 or more employees in both the private and public sectors and covering all industries except agriculture, forestry and fishing, and mining and quarrying.

For the purposes of our paper we select the estimating sample by excluding cases which do not have complete information on some of the variables involved in the analysis, as well as observations referring to workplaces for which less than 5 employees were sampled in the survey. We end up with an estimating sample with 15,379 employees and 1,824 workplaces.

In our analysis, the dependent variable is an overall job satisfaction indicator. The survey asks each employee to provide a rating from 0 (very dissatisfied) to 5 (very satisfied) answering eight questions on job satisfaction aspects: the sense of achievement they got from their work, the scope for using their own initiative, the amount of influence they have, the training they receive, the opportunity to develop their skills in their job, the amount of pay, their job security and the work itself. We build the indicator for overall job satisfaction using six out of eight aspects, and excluding satisfaction with pay and satisfaction with the work itself. The former is excluded because standard specifications of job satisfaction equations include functions of individual wages on the RHS, making the use of pay satisfaction indicators on the LHS tautological. On the other hand, the indicator of satisfaction with the work itself is a summary of overall job satisfaction, and we prefer to summarise overall satisfaction aggregating reported satisfaction on specific job dimension rather than relying on the individual reports to make sure that the *overall* indicator covers each of them. In order to derive our overall job satisfaction indicator, for each job aspect, we built a dummy variable equal to 1 if the individual was either ''very satisfied'' or ''satisfied'' and 0 otherwise. The overall satisfaction index is the sum of these dummies, thus representing the number of times the respondent rated herself as being either ''very satisfied'' or ''satisfied'', ranging from 0 to 6. This variable is strongly (but not perfectly) associated with the indicator for

satisfaction with the work itself, with a polychoric correlation of 0.65 (s.e.= 0.0069).

Table 1 provides a description of the relevant satisfaction indicators by gender (53% of cases in the estimating sample – and 56% in WERS – mare women). For each satisfaction item, the table reports the proportion of cases classified as ''satisfied'' or ''very satisfied'', while for overall satisfaction the average of the 0-6 indicator is displayed. The table shows that there is a varying degree of satisfaction across items, and satisfaction tends to be lower on items relating to training and skill development. Looking at gender differences, each item and the overall indicator show that reported satisfaction is larger for women than it is for men.

TABLE 1 – *Average job satisfaction by gender*

JOB ASPECTS	*FEMALE*	*MALE*
Sense of achievement	0.77	0.72
Scope for using own initiative	0.77	0.75
Amount of influence	0.62	0.62
Training received	0.56	0.51
Opportunity to develop skills in their job	0.54	0.51
Job security	0.60	0.58
Overall	3.85	3.69

Figure 1 offers the distribution of overall satisfaction by gender. The proportion of women who are completely satisfied (score equal to six) is higher than men (28.3% vs 25.2%). Moreover, the female proportion in the score equal to five out of six is also higher than for men (16.8% vs 16.4%). Conversely, men are more concentrated on the lower tail of the distribution of overall satisfaction.

Key ingredients for our analysis are the individual and average workplace hourly wage that are available in WERS, from the employee and manager questionnaire respectively. As discussed, wage indicators normally feature among the conditioning set of job satisfaction regressions. More specifically, starting from the seminal contribution of Clark and Oswald (1996) not only individual wages should be included, but also an aggregate wage measure meant to capture the tendency to benchmark own wages against a comparison level. Typically, in surveys wage benchmarks are derived as averages across regions or occupational groups. In contrast, in our paper we are able to provide a closer comparison income, that is the firm average wage as reported in the manager's questionnaire. This wage measure is closer to the concept of social comparison because it is an average from a social group to which the individual is (at least partly) effectively exposed, that is co-workers, differently from regional or occupational aggregates that are, in this sense, anonymous.

FIGURE 1 – *Overall job satisfaction*

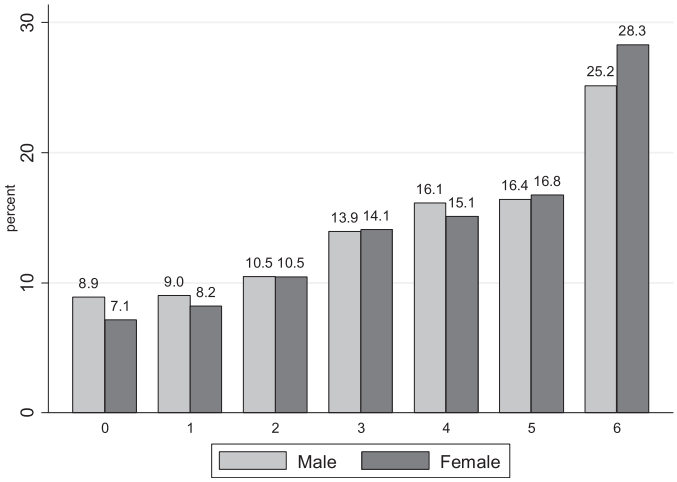

In Figure 2 we provide an illustration of the wage information that we use in the analysis by plotting they estimated kernel densities for three wage concepts of interest: individual wages, workplace average wages, and individual deviations

FIGURE 2 – *Wage densities*

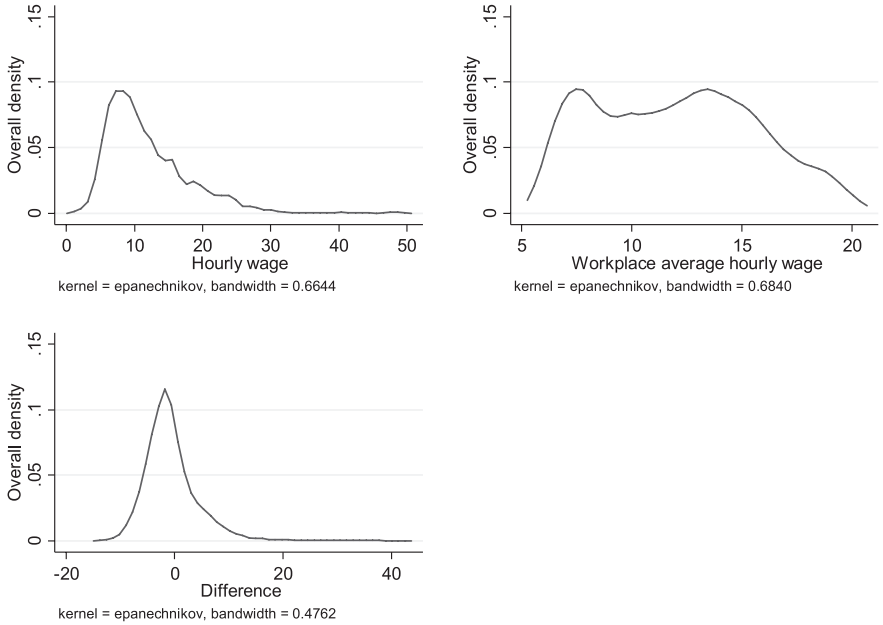

from the workplace average. Individual hourly wages (top left panel) are distributed resembling a log-normal distribution, with the mean located approximately at £12. Workplace average wages (top right panel) depict instead a bimodal distribution that is consistent with the idea of labour market polarization as discussed e.g. by Goos and Manning (2007), according to which the labour market is increasingly divided into good and bad jobs. In our case, it seems that polarization occurs across low and higher paying firms, which is also a concept associated with the idea of workers sorting (Abowd, Kramarz and Margolis, 1999). While the overall average of the workplace average wage distribution is very similar to the average of individual wages (at £12), the two poles of the distribution are located approximately at £7.5 and £14. The bottom panel shows deviations from the workplace mean which are less skewed than the distribution of wage levels, but still reveal that the largest probability mass is below the mean.

4. EMPIRICAL STRATEGY

Our data consist of a cross-section of workers indexed by i, stratified by workplaces; let $j(i)$ be the workplace of worker i. Our baseline empirical model relates the score of overall job satisfaction reported by person i (S_i) to gender (a dummy f_i), hourly wages (w_i) and demographic characteristics (the vector X_i) through a logit specification:

$$S_i^* = X_i'\beta + \phi f_i + \theta_1 w_i + \theta_2 w_i^2 + \varepsilon_i; \quad S_i = \tau\left(S_i^*\right) \tag{1}$$

where S_i^* is the latent level of satisfaction which can be thought of as utility derived from the job, ε_i is an error term assumed standard logistic and the non-linear function τ maps the latent level into the reported satisfaction score. Throughout the analysis, we allow job satisfaction to depend non-linearly on wages because Prospect Theory has predictions for the curvature of the utility function and its changes across the reference point, and omitting non-linearity with respect to own wage may lead to mis-interpret decreasing marginal utility of income as the prediction of the theory. In the basic set-up of Equation (1) the puzzle of greater WJS would result in positive estimate of parameter ϕ.

We first extend the model to incorporate an estimable version of Prospect Theory. As discussed in the data section, we have information on the workplace average wage, that we denote $\bar{w}_{j(i)}$, and we assume that this is the reference point. We define the percentage deviation from the reference point as $\delta_i = \dfrac{w_i - \bar{w}_{j(i)}}{\bar{w}_{j(i)}}$, and in order to address the differential impacts on satisfaction of social comparison from below or above the reference point, we split δ_i into its positive and negative components, that is we define $\delta_i^{(+)} = I(\delta_i > 0)\delta_i$ and $\delta_i^{(-)} = I(\delta_i \leq 0)\delta_i$, I being the indicator function. Using these variables we can extend the baseline model to allow testing Prospect Theory in the following way:

$$S_i^* = X_i'\beta + \phi f_i + \theta_1 w_i + \theta_2 w_i^2 + \gamma_1 \delta_i^{(+)} + \gamma_2 \delta_i^{(+)2} + \lambda_1 \delta_i^{(-)} + \lambda_2 \delta_i^{(-)2} + \varepsilon_i \quad (2)$$

The main prediction of Prospect Theory is that *losses loom larger than gains*, meaning that being below the reference point has a larger utility impact compared with being above; in the framework of Equation (2), the prediction can be tested by testing that the slope of negative deviations is larger than the slope of positive deviations, that is $\lambda_1 + 2\lambda_2\tilde{\delta}_i > \gamma_1 + 2\gamma_2\tilde{\delta}_i$ for any given deviation of size $\tilde{\delta}_i$. The second prediction is that the curvature of the utility function changes from positive to negative when moving from below to above the reference point, which is readily testable as $\lambda_2 > 0 > \gamma_2$.

In order to use this specification to see whether lower competitiveness is the explanation for the puzzling greater WJS, we need to augment Equation (2) with a full set of interactions of all wage-related variables with the gender dummy, such as the model becomes:

$$S_i^* = X_i'\beta + \phi f_i + \theta_1 w_i + \theta_2 w_i^2 + \gamma_1 \delta_i^{(+)} + \gamma_2 \delta_i^{(+)2} + \lambda_1 \delta_i^{(-)} + \lambda_2 \delta_i^{(-)2} + \theta_{1\phi} w_i f_i + \theta_{2\phi} w_i^2 f_i + \gamma_{1\phi} \delta_i^{(+)} f_i + \gamma_{2\phi} \delta_i^{(+)2} f_i + \lambda_{1\phi} \delta_i^{(-)} f_i + \lambda_{2\phi} \delta_i^{(-)2} f_i + \varepsilon_i \quad (3)$$

where the ϕ subscript denotes women-related parameters deviations from baseline effects. If lower competitiveness explains the puzzle of greater WJS, then incorporating deviations from the reference point by gender should show weaker social comparison for women (that is all parameters related to interactions between δ_i and f_i should be negative), and should also eliminate the baseline female effect on job satisfaction (i.e. $\phi = 0$).

Turning to the sorting hypothesis, we address it by first progressively including occupation and industry controls among the vector of regressors X_i', and next by adding workplace fixed effects which are identified thanks to the linked employer-employee structure of the data. Our fixed effect equation for the sorting hypothesis is:

$$S_i^* = X_i'\beta + \phi f_i + \theta_1 w_i + \theta_2 w_i^2 + \eta_{j(i)} + \varepsilon_i;\ S_i = \tau(S_i^*) \quad (4)$$

where $\eta_{j(i)}$ is the fixed effect. If the greater WJS depends on sorting, then estimating equation (4) should deliver $\phi = 0$.

5. RESULTS

In Table 2 we present results (namely estimated ordered logit parameters) obtained from the specifications discussed in the previous section. We use throughout survey weights that account for the fact that the individual probability of being

sampled in WERS depends on the size of the firm, that is once a firm is drawn, workers of smaller workplaces are more likely to be interviewed. Moreover, we account for possibly correlated unobservables within workplaces and cluster standard errors at the workplace level.

TABLE 2 – *Estimates of job satisfaction regressions*

	(1)	(2)	(3)	(4)	(5)
Woman	0.180***	0.351**	0.149***	0.0171	0.0326
	(0.0529)	(0.148)	(0.0503)	(0.0547)	(0.0605)
Hourly wage	0.00563***	0.00628	0.00324***	0.00588***	0.00729***
	(0.00128)	(0.00971)	(0.00117)	(0.00125)	(0.00142)
Hourly wage^2	-7.58e-06***	-2.46e-05	-4.48e-06***	-7.73e-06***	-1.05e-05***
	(1.73e-06)	(2.35e-05)	(1.48e-06)	(1.73e-06)	(1.92e-06)
Hourly wage*		-0.0175			
woman		(0.0110)			
Hourly wage^2*		3.45e-05			
woman		(2.38e-05)			
$\delta^{(+)}$		-0.0577			
		(0.124)			
$\delta^{(+)}$^2		0.00502			
		(0.00403)			
$\delta^{(-)}$		3.054***			
		(0.539)			
$\delta^{(-)}$^2		4.741***			
		(1.045)			
$\delta^{(+)}$*woman		0.253*			
		(0.144)			
$\delta^{(+)}$^2*woman		-0.00841*			
		(0.00430)			
$\delta^{(-)}$*woman		-1.232*			
		(0.668)			
$\delta^{(-)}$^2*woman		-2.638**			
		(1.344)			
Slope $\delta^{(+)}$ *man*		-0.080			
		(0.1241)			
Slope $\delta^{(-)}$ *man*		3.827***			

		(0.719)			
Slope $\delta^{(+)}$ woman		0.191***			
		(0.087)			
Slope $\delta^{(-)}$ woman		2.109***			
		(0.593)			
Occupation	No	No	Yes	No	No
Firm charact.	No	No	No	Yes	-
Demographics	Yes	Yes	Yes	Yes	Yes
Firm FE	No	No	No	No	Yes
Pseudo R-2	0.00377	0.00658	0.00933	0.0139	0.00982

The table reports parameter estimates from ordered logit regressions of the overall job satisfaction index on the relevant set of controls. δ-variables are the percentage deviations of individual wages from workplace average wage. Coefficients in Italic are derived post estimation assuming a 10% deviation (positive or negative) from the reference point. Regressions use survey weight and clustered standard error at the workplace level. *** p<0.01, ** p<0.05, * p<0.1. Number of employees = 15379; number of workplaces = 1824.

Column (1) reports the baseline model of Equation (1) in which job satisfaction is regressed on a gender dummy, a quadratic in hourly wages and a set of demographic controls (age, education, martial and parental status, disability). The regression reveals the classic WJS puzzle with a positive estimated effect for the women's dummy. It is also noteworthy that job satisfaction is concave in hourly wages.

Column (2) considers gender differences in social comparison for explaining greater WJS through the lenses of Prospect Theory, as discussed in Section 4. The regression is augmented with the percentage deviation of individual wages from the workplace reference point, split into positive and negative deviations and interacted with gender. (Consistently, we also interact the baseline wage effect with gender.) To account for the concavity of the utility function, the regression uses a second order polynomial in deviations from the reference points.

Results indicate that Prospect Theory finds varying degrees of support in our data. For both men and women, comparisons from below the reference point have a stronger impact on job satisfaction than comparisons from above. Also, there is clear evidence of a convex relationship below the reference point, that for women becomes concave above the reference point, while for men comparisons from above do not seem to matter. Below the reference point the intensity of the association is stronger for men.

The non-linear behaviour confirms that approaching the reference point from below has an increasingly positive effect of individual welfare, while once the reference point is left behind, further gains are less powerful in increasing satisfaction, consistent with the notion that losses loom larger than gains.

In order to summarize the associations between deviations from the reference

point and satisfaction, in Column (2) we also report the implied slopes, computed at a deviation of 10 percent (positive or negative) from the reference point. These indicate that when social comparison matters the most, that is when individuals are located below the reference point, women care less than men, which indeed supports the view that women have lower competitiveness.

However, these different attitudes are not enough to eliminate the baseline WJS premium, that actually becomes larger in Column (2) compared with the baseline specification. This leads us to conclude that the lower competitiveness hypothesis alone cannot account for the WJS puzzle.

Therefore, in Column (3) and (4) we turn to consider the sorting explanation and add, in turn, occupational dummies and controls for firm characteristics (namely size and industry). Results show that while including occupational dummies has little impact on the estimated WJS premium, the inclusion of firm characteristics induces a sizeable reduction of the WJS premium, which becomes negligible in size (one tenth of the baseline estimate) and statistically insignificant. Amongst firm characteristics (not shown), industry seems to be the most powerful predictor of job satisfaction. Workers in services, utilities, health and education all tend to report higher satisfaction compared with manufacturing, and these are industries where women are more frequently employed. Finally, Column (5) reports estimates from a model with workplace fixed effects, which accounts for both observable and unobservable firm characteristics. All these estimations point towards a disappearance of the WJS puzzle once firm characteristics are accounted for in one way or another.

Finding that the WJS puzzle is resolved once gender-based sorting across firms is taken into account has two alternative interpretations with different policy implications. On one hand, it may be that women are happier than men even before sorting into jobs, that is pre-determined feminine traits such as low competitiveness and risk aversion create a less stressful working environment conducive to higher job satisfaction. Interestingly, if that is the case, also men working with women should be happier than men surrounded by males. Eventually, there may be peer effects in satisfaction (Manski, 1993), whereby exposure to satisfied co-workers may increase own satisfaction. In this case, promoting women's occupation access is expected to increase welfare of the overall workforce. On the other hand, WJS may just be a reflection of differential working conditions between male- and female- dominated industries, such as the higher WJS is not pre-determined but emerges as a consequence of working conditions. In these circumstances, promoting equal occupational opportunities would not affect overall satisfaction; rather a levelling of working conditions between good and bad workplaces would be the effective policy to pursue.

In order to disentangle between the two alternative scenarios behind the sorting explanation, we derive from the manager questionnaire the workplace-specific proportion of women working in each of the occupational categories of WERS. This proportion ranges from 10 percent among machine operatives to 50 percent among sales services and professionals to about 80 percent among secretarial and caring

staff. Next, we assign the occupation-specific proportion of women in the workplace to each employee on the basis of her occupation. We augment the baseline satisfaction regression with this latter variable interacted by gender. Results are in Table 3. Column (1) shows that indeed individuals that are surrounded by women in their occupation report higher levels of job satisfaction, and this is true irrespective of gender. Note also that the inclusion of this variable alone is enough to eliminate any positive effect associated with the baseline gender dummy. *Prima facie*, this evidence would therefore support the view that women, by either affecting the work context or directly though peer effects, exert a positive impact on the satisfaction of their co-workers. However, this interpretation does not resist the inclusion of industry dummies in Column (2), which no-longer shows any significant effect associated with the occupational incidence of women. Overall, the evidence of Table 3 suggests that WJS stems from the particular industries in which women end-up working.

TABLE 3 – *Job satisfaction and the occupational incidence of women*

	(1)	(2)
Woman	0.0792	0.0617
	(0.129)	(0.131)
Hourly wage	0.00608***	0.00604***
	(0.00145)	(0.00147)
Hourly wage^2	-7.96e-06***	-7.79e-06***
	(1.84e-06)	(1.87e-06)
Occupational incidence of women	0.367***	0.182
	(0.128)	(0.127)
Occupational incidence of women*	-0.0319	-0.0907
woman	(0.201)	(0.203)
Industry	No	Yes
Demographics	Yes	Yes
Observations	15379	15379
Pseudo R-2	0.00438	0.0131

See Table 2. Robust standard errors in parentheses*** p<0.01, ** p<0.05, * p<0.1

6. SUMMARY AND CONCLUSIONS

Although women face a more difficult labour market than men, our results (as the ones of many before us) show that women report higher levels of satisfaction with the job. Thanks to the availability of WERS linked employer-employee data,

we have put to the test two competing explanations for this apparent paradox, namely that women have lower competitiveness than men or that women sort into jobs that where satisfaction is intrinsically higher.

While we do find some support for the lower competitiveness hypothesis, accounting for it through a model of social comparison is not enough to explain the paradox. Instead, duly accounting for sorting through industry or firms fixed effects eliminates the puzzling female satisfaction.

Our findings, therefore, show that *coeteris paribus* women are no different from men in terms of reported satisfaction. Given the well documented gender pay gap and the fact that satisfaction increases with wage, our findings actually imply a lower level of female satisfaction within a given industry or firm. Raw gender differences in reported satisfaction are thus of no excuse for delaying the implementation of equal opportunities in the workplace.

REFERENCES

Abowd J., Kramarz F., Margolis D. (1999). High Wage Workers and High Wage Firms. *Econometrica*, **57**, 411-445

Bender A., Donohue S.M., Heywood J.S. (2005). Job Satisfaction and Gender Segregation. *Oxford Economic Papers*, **57**, 479-496

Card D., Mas A., Moretti E., Saez E. (2012). Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review*, **102**, 2981-3003

Clark A.E. (1997). Job Satisfaction and Gender: Why are Women so Happy at Work? *Labour Economics*, **4**, 341-372

Clark A.E., Oswald A.J. (1996). Satisfaction and Comparison Income. *Journal of Public Economics*, **61**, 359-381

De Paola M., Ponzo M., Scoppa V. (2017). Gender Differences in the Propensity to Apply for Promotion: Evidence from the Italian Scientific Qualification. *Oxford Economic Papers*, **69**, 986-1009

Duesenberry J.S. (1949). *Income, Saving and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, Massachusetts

Gächter S., Thöni C. (2010). Social Comparison and Performance: Experimental Evidence on the Fair Wage-Effort Hypothesis. *Journal of Economic Behavior and Organization*, **76**, 531-543

Gamba A., Manzoni E., Stanca L. (2017). Social Comparison and Risk Taking Behavior. *Theory and Decision*, **82**, 221-248

Goos M., Manning A. (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics*, **89**, 118-133

Hauret L., Williams D.R. (2017). Cross-National Analysis of Gender Differences in Job Satisfaction. *Journal of Economy and Society*, **56**, 203-235

Kahneman D., Tversky A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, **47**, 263-291

Kaiser L.C. (2007). Gender-Job Satisfaction Differences across Europe: An Indicator for Labour Market Modernization. *International Journal of Manpower*, **28**, 75-94

Manski C.F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*, **60**, 531-542

Niederle M., Vesterlund L. (2007). Do Women Shy Away From Competition? Do Men Compete Too Much? *The Quarterly Journal of Economics*, **122**, 1067-1101

Redmond P., McGuinness S. (2019). *Explaining the Gender Gap in Job Satisfaction*. IZA Discussion Paper, 12703

Veblen T. (1899). *The Theory of the Leisure Class*. Macmillan, New York

Vladisavljevi M., Perugini C. (2018). Gender Inequality and the Gender-Job Satisfaction Paradox in Europe. *Labour Economics*, **60**, 129-147

Zou M. (2015). Gender, Work Orientations and Job Satisfaction. *Work*, *Employment and Society*, **29**, 3-22