

Does Rosie Like Riveting?

Male and Female Occupational Choices

Grace Lordan and Jörn-Steffen Pischke

LSE

August 2021

Abstract: Occupational segregation and pay gaps by gender remain large while many of the constraints traditionally believed to be responsible for these gaps seem to have weakened over time. We explore the possibility that women and men have different tastes for the content of the work they do. We relate job satisfaction and job mobility to measures that proxy for the content of the work in an occupation, which we label ‘people,’ ‘brains,’ and ‘brawn.’ The results suggest that women value jobs high on ‘people’ content and low on ‘brawn.’ Men care about job content in a similar fashion but have much weaker preferences. High school students show similar preferences in a discrete choice experiment and indicate that they make their choices mainly based on preferences for the work itself. We argue that the more pronounced preferences of women can account for occupational sorting which often leads them into careers with large pay penalties for interruptions due to childbearing.

JEL classifications: J16, J4

Keywords: Occupational choice, job content, gender, preferences

“And finally, in our time a beard is the one thing that a woman cannot do better than a man.” - John Steinbeck, Travels with Charley: In Search of America.

INTRODUCTION

Women’s progress in the labor market has been dramatic since Steinbeck’s travels in the 1960s. The female employment rate has risen, the pay gap with men has declined, and occupational segregation has decreased. Despite all this progress, female convergence has slowed and possibly stopped since about the turn of the millennium, while sizeable gaps still remain in pay and hours. Figure 1 tracks the share of males in the occupations in which women work in the US over time. The share of males in the jobs done by females has been increasing over time, but progression has slowed or stalled in the early 2000’s with substantial differences remaining. One particular concern is that females are still under-represented in many high paying professional and managerial occupations (see Figure 2 and Goldin, 2014). Although a lot of the gender wage gap is within occupations, the lack of women in these high-paying, male-dominated professions contributes to the gap (Bayard et al., 2003; Blau and Kahn, 2016). For example, in 2014, the average hourly wage of individuals in the US who work in majority male occupations (proportion of males ≥ 0.70) was \$23.67, versus \$19.30 for those in minority male occupations (proportion of males ≤ 0.30).¹ We will argue that understanding occupational segregation may help us better understand the pay gap within occupations as well.

[Figure 1 and Figure 2 about here]

As traditional explanations for gender wage gaps, discrimination, labor supply, and human capital investments (Altonji and Blank, 1999), have declined in importance, the literature has turned towards attitudes, personality traits, and gender identity (Croson and Gneezy, 2009; Bertrand, 2010). However, the role of many of the variables suggested as explanations for lower female earnings remains empirically elusive (Manning and Swaffield, 2008; Fortin, 2008). The predominant view among economists seems to be that the main remaining obstacle to more equal labor market outcomes between the sexes is a lack of flexibility to combine a career and family. Goldin (2014) argues this point most forcefully, but it is also shared by Bertrand (2018).² Kleven et al. (2019) and Bütikofer et al. (2018) provide powerful demonstrations of the continuing sharp decline in wages and earnings once a woman has children in Denmark and Norway respectively, countries with long histories of comparatively equal gender attitudes.

The flexibility story raises its own puzzles. In this paper, we explore whether preferences for the content and context of the work done in particular jobs might explain some of the occupational segregation we see in the labor market. We argue that such preferences can help explain some empirical regularities which are at odds with a simple flexibility story. One of Goldin's (2014) metrics for the flexibility of an occupation is the elasticity of individual earnings with respect to hours worked: high elasticities imply a penalty for workers seeking short hours and indicate a lack of flexibility. She demonstrates that less flexible occupations have a larger pay gap. Goldin (2014) classifies occupations into five groups: health, business, tech, science, and other. Women do not necessarily gravitate towards the most flexible groups and sometimes do the exact opposite. Business occupations are the least flexible group with

an average elasticity of 0.93 but women's share in this group is about the same as their overall representation in all occupations, around 40%. On the other hand, women make up only 20% of workers in the much more flexible tech group (with an elasticity of 0.47).³ Across 95 occupations, the share of men (SOM) in an occupation is basically uncorrelated with the earnings-hours elasticity.⁴

Goldin (2014) shows that the lack of flexibility is related to the amount of contact with others and the importance of building relationships in a job: where workers have to communicate with co-workers or clients both parties have to be present at the same time, limiting flexibility. Our conjecture is that women may actually value jobs which incorporate some interpersonal elements over purely abstract tasks (and it seems Claudia Goldin has come to agree with this idea, see EPL Cornell, 2014, 1:21:53-1:23:35).

Jobs differ widely in terms of the tasks performed, and a large literature in economics has classified jobs in terms of task content following the work of Autor et al. (2003). We deviate from this literature by using a statistical classification of the content of work using ONET data on occupations, which we loosely label 'people,' 'brains,' and 'brawn' ex-post. We then relate job satisfaction and exits from an occupation to these measures of job content using panel data on job switchers for three large advanced economies, the US, Britain, and Russia. Both men and women are more satisfied and more likely to stay in 'people' and 'brain' jobs but the pattern is more pronounced for women than for men. An important confounder might be other aspects of the work environment in different occupations. To probe this possibility, we complement the main analysis with cross-sectional regressions from the British Workplace

Employment Relations Study (WERS), which lets us control for firm effects. Overall, we find that firm effects matter strongly themselves, but a similar pattern with respect to the occupation attributes remains as before. We argue that our results point towards an explanation where preferences for the content of the work in a particular job matter for occupational choice.

To substantiate that it is preferences rather than some other job attribute that matters, we conducted a discrete choice experiment with high school students who are mostly university bound. We asked the students to choose between six pairs of occupations. The choices made by the students closely mirror the adult results. To pinpoint what drives differences in choices, we asked the respondents to explain why they made the particular choices they did. The majority of answers indicate that students prefer the activities in one of the jobs, or that their abilities are a better match. Few respondents mention other aspects of the job as important. These results closely mirror survey results by Zafar (2013) on preferences and major choices among Northwestern University students and a choice experiment by Gelblum (2020) on Mechanical Turk.

If women have stronger preferences than men, equilibrium sorting into occupations can explain segregation between men and women. Such an explanation might account for the slowdown in occupational convergence. We view the role of preferences as a natural complement to the flexibility story by offering an explanation as to why women often choose occupations with a large penalty for work interruptions, leading to a within-occupation wage gap. The fact that occupations with a large component of social interaction often have large pay penalties for flexible work is a necessary ingredient to explain other recent findings in the literature as well. Deming (2017)

demonstrates the rising importance of social skills in the labor market and Cortes et al. (2020) show that women have differentially sorted into occupations where interactive tasks have become more important. These occupations are often also cognitive task intensive and well paid, but these trends have not been able to close the gender pay gap. Stronger female preferences for jobs with a social component also implies that this becomes a job amenity for which women are willing to accept lower pay.

Our findings align with a large literature in psychology that has persistently pointed out important sex differences in preferences, particularly along similar lines to our ‘people’ versus ‘brawn’ dimensions (see Su et al., 2009 for an overview). Hakim (2000) and Pinker (2008) have gone further and pushed the idea that these differences in preferences of women and men are a primary driver of the persistent differences in labor market choices. Hakim’s interest is in women’s attitudes towards a role as homemaker, a full-time labor market career, or a combination of family and work. While Hakim offers quantitative evidence using similar variables as we do, occupational choice plays a minor role in her account—it matters primarily to the degree that some occupations are more likely to offer part-time work or accommodate less committed careers. Pinker’s (2008) work is closer to our idea that women may like the nature of male-dominated jobs less, and supports a division along the people-things dimension, but only contains a narrative analysis. Notably, while these literatures have typically stressed gender differences along a people versus things dimension, we also find a strong preference of women according to our ‘brains’ dimension.

A related, concurrent analysis to ours is Gelblum (2020), who carries out a choice experiment on Mechanical Turk, eliciting willingness to pay for jobs which differ only

in terms of the fraction of time spent on tasks typically seen in female and male dominated jobs. She also finds directionally similar preferences by gender but women are willing to pay more for preferred job tasks. Cortés and Pan (2018) discuss a wider range of explanations for occupational segregation of men and women but their empirical analysis considers very similar ONET variables as we do. Fortin (2008) uses a narrower set of survey based variables related to skills and preferences in wage regressions. She shows that they do not explain any of the gender wage gap but does not analyze occupational choice. Also related is Usui (2008), who uses the National Longitudinal Survey of Youth 1979 (NLSY79) from 1979-1982 and shows that women are less satisfied in male-dominated jobs. Hunt (2016) demonstrates that female college graduates in the US are more likely than males to leave engineering jobs but shows that this is mostly due to the fact that women are more likely to leave male-dominated occupations in general.

I. FRAMEWORK AND METHODS

We are interested in an individual's preferences for the content of the work they do in their job, whether these preferences differ in strength between men and women, and whether such differences might explain differences in occupational choices. We would like to know why female academics are more likely to be found in the life sciences than the physical sciences, or why women are more likely to work as financial analysts than electrical drafters. To set the stage for our investigation, suppose utility is given by $U(C, JC)$ where C is consumption, JC is (for simplicity) a unidimensional aspect capturing the content of the work or "job content." A job amenity like JC is typically valued by computing the marginal rate of substitution $(dU/dJC)/(dU/dC)$. Our

conjecture is that $(dU/dJC)/(dU/dC)$ may differ for men and women, and the strength of these differences influences the choices of jobs by gender.

How would we assess this? The economics literature uses three main methods to study preferences: studying choices, asking individuals directly about their preferences, and estimating satisfaction equations. We use all three approaches in this paper.

Studying choices: If women like the attribute JC more than men we might see more women in high JC jobs even if these jobs have lower salary as they are compensated by the utility they get from doing an enjoyable job. We can evaluate this by regressing individual job choices or the share of men (SOM) in an occupation on attributes including JC. There are two obvious complications with this approach. The first is that the list of relevant job attributes may be long, and many of these attributes might be unobservable. If any omitted attributes are correlated with JC, we might get the estimate wrong. The second is that choices are not determined solely by preferences, but by the interaction between preferences and constraints. It may simply be the differences in constraints which give rise to different choices of men and women.

One way to address these issues is not to rely on real choices but rather present individuals with hypothetical choices or vignettes in a survey. The options given to individuals in such a setting can be controlled more tightly in order to minimize the risk of omitted variable bias. This methodology has the advantage that individuals can be confronted with choices from many sets, which produces individual level panel data. Attributes presented can be chosen so as to create a large amount of relevant variation, circumventing many of the problems associated with actual choices. Examples of such

choice experiments are Wiswall and Zafar (2016), who presented university students with hypothetical vignettes, Mas and Pallais (2017), who varied job attributes in a field setting with actual online job applicants, and Gelblum (2020) who varied job tasks in a choice experiment on Mechanical Turk. Drawbacks of hypothetical choice experiments are that choices do not have real consequences, individuals may not be familiar with choice dimensions they have not encountered before, and they may read additional differences into choices which seem artificial to them.

Asking individuals about their preferences: An alternative to studying choices is to simply ask people directly about their preferences. Contingent valuation methods, closely related to choice experiments, have been widely used in settings where valuations are not priced directly by markets, like environmental policy. These methods have been criticized because individuals tend to find it difficult to think about hypothetical choices in areas they are not typically faced with, and as a result give inconsistent responses (see e.g. Diamond and Hausman, 1994). This should be less of an issue in a job choice context. We will ask high school students about their preferences for different occupations. Although this group has no direct experience with these jobs yet, the students are thinking actively about their subject choices which determine their future career options.

Estimating satisfaction equations: An alternative approach is to interpret survey measures of satisfaction (with the job or with life) as measures of $U(.)$, estimate such a satisfaction equation, and treat the estimates as preference parameters. If one of the arguments in the satisfaction equation is income or consumption, the estimates can again be used to calculate a willingness to pay $(dU/dJC)/(dU/dC)$. Frijters and van

Praag (1998) have applied this idea to valuing climate and van Praag and Baarsma (2005) to value airport noise. Finkelstein et al. (2013) use a similar idea to estimate marginal utilities like dU/dJC directly.

Estimating satisfaction equations suffers from the same problem that included job attributes might proxy for omitted ones. One advantage over studying choices is that variation in job attributes which comes about because different individuals face different constraints (or prices), should still lead to valid inferences. As long as variation in constraints move an individual along a single indifference curve, they should report the same satisfaction level.

An important issue in using satisfaction data is that reported job satisfaction may not be the same as choice utility and estimating satisfaction equations may not give the same result as evaluating choices. Kimball and Willis (2006) and Benjamin et al. (2012) consider a utility function of the form $U(C, JC, S(JC))$, where $S(\cdot)$ is the job satisfaction function. JC matters for job satisfaction, and job satisfaction matters for utility relevant for decision making. But JC may also enter the utility function directly, for example, by affecting the happiness of one's family if a person's feelings about work spills over to home. As a result

$$\frac{dU}{dJC} = \frac{\partial U}{\partial JC} + \frac{\partial U}{\partial S} \left(\frac{dS}{dJC} \right). \quad (1)$$

This framework highlights that the strength of preferences of men and women can differ because of differences in either dS/dJC , $\partial U/\partial S$, or $\partial U/\partial JC$. Estimating satisfaction equations at best yields information on the term dS/dJC .

Benjamin et al. (2012) compare vignette-based choices from a variety of diverse scenarios with rankings based on subjective well-being (SWB) measures. Benjamin et al. (2014) make similar comparisons between real choices in the medical Resident Matching Program and SWB measures related to the options. In both studies, there is a fair alignment between choices and SWB ranking, but there are also some systematic deviations. In Benjamin et al. (2012), the differences in rankings are related to other life domains, like control over one's life and a sense of purpose. Various choice scenarios in their paper are work related, and they find a large role for the term $(\partial U/\partial S)(dS/dJC)$ in choices, suggesting that satisfaction equations will contain useful information. Comforting for our purpose, they find no systematic differences in the way choices versus SWB rankings differ for men and women. Any differences we find should therefore reflect real differences in the strength of preferences rather than, for example, different uses of satisfaction scales across sexes.⁵

The previous discussion highlights that none of the methods is likely to give a definite answer to the question whether preferences play a role in the diverging occupational choices of men and women. Therefore, we combine elements of all of these approaches. We start with simple satisfaction and job mobility equations, relating these to a variety of occupational characteristics and find stronger results for females in both. Preferences for the content of a job are one possible explanation for our results but we acknowledge that there could be others, like flexibility or work environment. In order to probe the role of preferences in job choices further, we conducted a choice experiment with high school students. We asked the students to make choices between six paired occupations, distinct in terms of work content. The choice results for the

students are very similar to those for the working adults. The students confirm that interests in the type of work are the primary reason for their choices.

II. ANALYSIS OF LONGITUDINAL DATA

In this section we analyze four datasets: the US National Longitudinal Survey of Youths 1979 (NLSY79), the British Household Panel Study (BHPS), the Russian Longitudinal Monitoring Survey (RLMS) and the British Workplace Employment Relations Study (WERS). We obtain information on job content from the US ONET database.

Measuring Job Content from ONET

To measure job content, we use ONET version 5, which provides a diverse set of information on occupational attributes, requirements, and characteristics of the workers in an occupation; all in all 249 distinct items. Out of these, we use the 79 items describing the work activities and context of a person's occupation. We focus on these 79 items because they capture well what a person does in their job along with the environment that they do their work in, while other items focus on worker attributes like skills requirements (see Appendix B Table B.1 for a list of the items).⁶ We standardize each of these variables to have a mean of 0 and a standard deviation of 1. These variables are later matched to the country specific survey data.

Rather than add the 79 context and activities variables to our regressions directly and risk over-fitting, we follow the psychometric literature and use exploratory factor analysis to reduce the dimensionality first (Gorsuch, 193; Thompson, 2004; see Appendix B for details). This also helps interpretation: a structure of three latent factors

emerges, which we loosely label as ‘people,’ ‘brains,’ and ‘brawn,’ or PBB. These labels appear natural to us based on the ONET items that load on each factor (see Appendix B Tables B.1 and B.2).⁷

US NLSY79

The NLSY79 is a panel of 12,686 individuals who were between 14 and 22 years old when first surveyed in 1979. These individuals were interviewed annually through 1994 and then on a biennial basis. In every wave, respondents were asked about job satisfaction: “How do you feel about the job you have now?” and were given the following response option: ‘I like it very much,’ ‘I like it fairly well,’ ‘I dislike it somewhat,’ ‘I dislike it very much.’ We coded responses so that higher values represent higher satisfaction. Our analysis uses an unbalanced panel of employees who responded to this job satisfaction question.

We create an additional dependent variable that captures movements in the labor market.⁸ This variable is equal to 1 if a person has the same three digit occupation code in year $t+2$ compared to the occupation that they held in t . Conversely, the variable is equal to 0 if an individual has a different occupation code in $t+2$ or has left employment. We call this variable ‘stayers.’ The variable is defined on a biennial basis given the interview schedule of the NLSY79 post-1994.⁹ Our analysis sample spans the years 1982 to 2014. We use sampling weights in the analysis that reflect that the NLSY79 over-sampled blacks, Hispanics, and the economically disadvantaged (see Appendix D for unweighted estimates).

British Household Panel Survey (BHPS)

We use all 18 waves of the original sample of the British Household Panel Survey (BHPS), a longitudinal study of around 5,500 households and over 10,000 individuals in England, Wales and Scotland that began in 1991. This main sample was supplemented with a Welsh extension from 1999 (1,500 households), a Scottish extension from 1999 (1,500 households) and a Northern Ireland extension from 2001 (1,900 households).

We use two questions asking respondents how satisfied or dissatisfied they are with i) their current job overall and ii) the actual work itself. We present additional results on satisfaction with other job domains in Appendix C Table C.3. Answers are on a 7-point scale. We again create an additional binary dependent variable that captures whether a person stayed in the same occupation. We measure mobility in the BHPS between two consecutive years.¹⁰ We present unweighted results from the unbalanced panel of all individuals including the extension samples between 1991 and 2008. We also investigated the sensitivity of our results to i) unweighted regressions of the original BHPS sample only and ii) weighted regressions of the main BHPS sample. See Appendix D for these results.

Russian Longitudinal Monitoring Survey (RLMS)

The Russian Longitudinal Monitoring Survey (RLMS) is a nationally representative annual survey, which started in 1994. However, job satisfaction data is only available from 2002-2012. We restrict our sample to employees who answer the question: ‘How satisfied or unsatisfied are you with your job in general?’. Response options are absolutely satisfied, mostly satisfied, neutral, not very satisfied and absolutely unsatisfied. We code responses so that higher values represent being more satisfied.

We create a binary dependent variable that captures whether a person stayed in the same occupation over two consecutive years. Our RLMS regressions use weights that allow for the complex design of the dataset where many observations are derived from following the housing unit rather than the person, as well as having over-samples from the first wave to allow for attrition. We show unweighted regressions in Appendix D.

British Workplace Employment Relations Study (WERS)

The British Workplace Employment Relations Study (WERS) is a national survey of people at work in Britain, which collects data from employees, employee representatives, and employers in about 2,500 firms. Multiple employees are interviewed from each firm. The WERS is conducted every six to eight years but is not a panel of firms or workers. We use the 2004 and 2011 surveys, which included an individual's three-digit occupation code using the British SOC00 codes (previous versions did not). We utilize the employee responses to the question about satisfaction with the work itself as there is no overall job satisfaction question. Response options are on a 5-point scale.

Matching and Creation of PBB factors

We create and match the three PBB factors to the NLSY, BHPS, RLMS and WERS data in addition to averages of an hourly wage, weekly hours, the proportion of college graduates, and age in each occupation (see Appendix F for further details).

Empirical Model

Our starting point is a fixed effects regression of the form

$$Y_{ijt} = \alpha_i + JC_j\delta' + X_j\beta' + X_{ijt}\gamma' + \mu_t + \omega_a + \varepsilon_{ijt} \quad (2)$$

where Y_{ijt} is either job satisfaction or a binary variable which indicates whether a person stayed in the same occupation in the next period for individual i in occupation j and year t . JC_j refers to the ‘people,’ ‘brains,’ or ‘brawn’ content of the occupation, X_j contains average wages, hours, age, and the proportion of college graduates by occupation, X_{ijt} contains age and age squared of the individual, μ_t are wave effects, and ω_a are region effects.¹¹ α_i is a set of individual fixed effects, so that the effect of job attributes is identified from occupation switchers, while controlling for time invariant individual differences (as a sensitivity analysis we also estimate equation (2) without individual fixed effects see in Appendix C Tables C.1 and C.2. We calculate standard errors using two-way clustering by individual and occupation (see Cameron et al., 2011).

To understand differences by gender, we present estimates separately for males and females. The coefficients of interest in equation (2) are δ . Positive coefficients imply that the job content variables are associated with an increased tendency to stay in an occupation in the stayer regressions and with higher levels of job satisfaction in the satisfaction regressions. To make the interpretation of δ s more intuitive in the job satisfaction regressions (given that the job satisfaction scales differ across country) we follow van Praag and Ferrer-i-Carbonell (2008) and normalize the job satisfaction variables by using the residuals from an ordered probit on the raw sample fractions. Since we also standardize the job content variables, our estimates have the interpretation of effect sizes. Because we want to compare results between men and women, we need to assume that they use the steps in the satisfaction scales in the same

way, but the fixed effects allow the scales to be anchored differently for different individuals.

An important issue in interpreting the results from a regression like (2) is how workers sort into heterogeneous occupations. The standard compensating differentials framework suggests that workers sort into the type of jobs they prefer in equilibrium. Occupation wage differentials reflect the compensating differentials required by marginal workers who are indifferent between two alternative jobs. This framework predicts that men and women may end up working in different jobs in equilibrium if they have different preferences for job attributes or if they face different constraints (say in terms of flexible schedules). In this scenario, it is unlikely that job satisfaction will reflect preferences. In the competitive compensating differentials model everybody works in their most preferred occupation, given equilibrium wages, and hence should report their maximum job satisfaction attainable.

The most natural extension to the simple frictionless, full information framework, which supports job changes, is a job search framework. Such a model with frictions allows for individuals to make choices subject to imperfect information regarding what an occupation's content is in practice and to choose from a limited set of available job offers at any time. Modeling occupational choices and wage differentials in a framework with frictions can lead to very different equilibrium outcomes (see Hwang et al., 1998; Manning, 2003; and Lang and Majumdar, 2004). Importantly, in a setting with frictions, workers may end up in jobs other than their preferred one, but they will switch jobs in future periods in search of better matches. This "frictional disequilibrium" constitutes a natural source for interpreting the results from a job

satisfaction equations like (2). As there are good jobs and bad jobs, as well as high and low-quality job matches for particular individuals in this framework, the coefficients on occupation characteristics have a natural interpretation as individual preferences for these characteristics.

Of course, even in the framework with frictions, individuals are not randomly assigned to occupations. This gives rise to two complications. One is the possibility of reverse causality: the choices women and men make may influence the way they work and how an occupation is structured. For example, Chang (2018) points out that the share of female computer programmers used to be higher in the 1970s than it is now. Programming also used to be organized in a more interactive fashion then. This could be due to the fact that there were enough women in the occupations so that they were able to structure their work environment to suit their own preferences. Once men dominated the profession, work organization changed to a more solitary model with longer working hours in the large firms.

The second complication with the regression strategy we are employing relates to the problem that the ONET variables we are using may proxy for other relevant aspects of the occupations, as discussed above. In order to get at the most important ones, we control for average wages, hours, age, and the proportion of college graduates in an occupation, which are all important factors in the job satisfaction and stayer equations. But we note that the SOM in an occupation is likely to affect variables like wages and hours worked as well, so that these attributes become endogenous. While the controls we use don't vary at the individual level (except for age), the variation in job content we are interested in is an occupation-level variable, and we would expect the bad

controls issue to spill over to the occupation level when the SOM varies across occupations. Like everybody else in the literature on sex differences, we have no solution to offer to this problem.

Another issue in evaluating the valuation of job attributes is that individuals face both a set of jobs with different attributes but also an outside option of not working. We have no information on job satisfaction for the non-employed. We may not see an individual working if a particular job attribute is very important to them (for example, enough flexibility to be able to care for children) and employers may not provide certain amenities because there is no interior market equilibrium where such trade takes place. As a result, those individuals for whom we see job satisfaction may not value an under-provided amenity as much or at all. This selection problem, similar to the problem of estimating wage equations in the presence of employment participation, may distort estimates relating satisfaction to amenities in the sample of working individuals. While we do not address the selection into employment directly, we note that it will likely bias the coefficient estimates on the PBB factors towards zero if the non-employment option offers a better amenity package than the available jobs. The same selection issue also affects the study of observed choices (as we observe no occupation for individuals who do not work) but the student survey we analyze in the next section allows us to elicit responses which are not subject to this problem.

It is typical in the evaluation of job attributes to measure marginal rates of substitution, i.e. $(dU/dJC)/(dU/dC)$. Instead, we simply look at the coefficients of job attributes in the satisfaction equations directly, i.e. dU/dJC . There are a number of reasons for this. First of all, we estimate simple linear satisfaction equations. With a linear income term,

the implied MRS is constant. Of course, we could add non-linear terms of income to the regressions or use a more structural utility framework but we are worried that there is not enough information in the job satisfaction measure, which is measured coarsely in the surveys we use (on a 4 to 7 point scale), and the same is true for our binary mobility equations. We don't believe that these data are particularly well suited to estimate the marginal utility of income well (but see Finkelstein et al., 2013 for an alternative view), and we worry that poor estimates of dU/dC might cloud our results. One cost of this is that our estimates do not have a simple numerical interpretation. We are willing to live with this drawback, as our main interest is the contrast in the strength of preferences between females and males.¹²

A more important reason why we are hesitant to rely on income estimates is the fact that we include various human capital variables like education and age among the occupational averages X_j . These variables capture a lot of permanent income components, and the interpretation of the coefficients on average earnings in the occupation or own earnings of the respondent becomes much more dubious. Average age and education of an occupation are important correlates of job satisfaction, presumably because more educated and experienced workers get paid more but also because they often get to work in more interesting jobs. Finally, even leaving this last issue aside, Benjamin et al. (2012) find that income coefficients are typically underestimated in satisfaction equations compared to the role of income in choice.

Results

We start in Table 1 by presenting a simple linear regression of the SOM on the three latent factors, along with the other occupational averages, time dummies, and area

dummies. We run this at the individual level but note that this is essentially an occupation level regression and the individuals here only serve to give different weights to different occupations. These regressions use data from Census/ACS for the US, QLFS for the UK, and RLMS for Russia.

Table 1 highlights that there is substantial sorting in all three countries along the dimension of ‘people,’ ‘brains,’ and ‘brawn.’ Women are overrepresented in ‘people’ jobs, men in ‘brawn’ jobs, and they share ‘brain’ jobs roughly equally. The pattern is stronger in Russia than in the US and Britain but is important in all three countries. The ‘brawn’ component seems to be the more potent predictor of sorting by gender than the ‘people’ factor. We suspect that this is due to the role of blue-collar jobs in the occupation distribution.

[Table 1 about here]

In Table 2 we turn to individual fixed effects regressions of job satisfaction and occupational mobility on PBB, as in equation (2). In all three countries, both men and women tend to like ‘people’ and ‘brain’ jobs and dislike ‘brawn’ jobs, with the ‘brain’ coefficient for Russia being an exception. Men are more likely to stay in ‘brawn’ jobs, although they are not particularly satisfied. Coefficients for women are generally bigger in absolute value than those for men, suggesting that women have stronger preferences for these job attributes.¹³ In the US, the coefficients of men and women are qualitatively most similar and only magnitudes differ, while in Britain men are indifferent to ‘brain’ jobs. The stayer regressions tend to match these patterns overall, although there are discrepancies for a few coefficients. In general, these results closely

mirror the ones we saw for sorting into occupations in Table 1. We note that these results are from fixed effects regressions and hence are identified from job switchers. In Appendix C, Tables C.1 and C.2 we also report cross-sectional regressions, which show a roughly similar pattern for a more representative population.¹⁴

Recall that the coefficients in the satisfaction regressions reflect effect sizes. As a different way to get a sense of the magnitudes of these effects, consider forming predicted values by multiplying the PBB coefficients from the NLSY job satisfaction equation with the values of the three factors (but ignoring other occupation averages). The female predicted value for heavily female dominated social work (SOM = 0.25) is 0.14, while for male dominated mechanical engineering (SOM = 0.94) it is 0.04. This reflects the fact that mechanical engineering scores much lower on ‘people’ and somewhat higher on ‘brawn’ than social work. Moving between these occupations changes job satisfaction by 0.10 of a standard deviation. For comparison, Stevenson and Wolfers (2008) find that a 33% difference in income is associated with about 0.10 of a standard deviation difference in life satisfaction in within country cross-sectional data.¹⁵ This suggests a potentially sizeable role for job content to us.

For men, the predicted values are 0.06 for social work and 0.04 for mechanical engineering, indicating that men are slightly more satisfied with the social worker bundle of job content as well (since most men dislike the solitary nature of engineering too). The occupations with the most negative predicted values for women are blue-collar jobs with values ranging from 0.0 to -0.2. Men dislike these jobs as well, but less so than women. The fact that men generally care less about the PBB factors is also reflected in the standard deviation of these predicted values across the entire 310

occupations, which is 0.03 for men and 0.09 for women. But for both sexes the influence of the PBB variables on job satisfaction is sizeable.¹⁶

The PBB factors are related to decisions whether to stay in a job or not as well, but the magnitudes are relatively small. The same comparison of the values of PBB implies only a 0.3 percentage points higher probability of a woman quitting her career in mechanical engineering as opposed to one in social work.

[Table 2 about here]

Together, Tables 1 and 2 suggest a role for the PBB variables for satisfaction and job choice. These effects are more important for women than they are for men. Because women strongly shy away from ‘brawn’ jobs, these jobs are left to fill for men who are less averse to them; an implication of the comparative advantage principle.

The results we have presented so far are consistent with the idea that tastes for the content of work differ by gender and influence the occupation choices of women and men. However, the PBB variables are crude measures of work content, and may proxy for environmental or organizational factors, which affect men and women differently.

A lot of aspects related to the work environment might be specific to a workplace and shaped by managers and co-workers. As a result, environment will often be a firm level characteristic rather than a characteristic of the occupation of a particular worker. None of the datasets we have analyzed above allows us to incorporate this in our analysis. We therefore turn to the British Workplace Employment Relations Study (WERS),

which samples multiple employees per firm. The WERS data are cross-sectional but allow us to include firm fixed effects to capture aspects of the environment that may affect females at work. Therefore, we identify the coefficients on PBB from variation caused by having individuals from multiple occupations working in the same firm. Of course, this methodology will not manage to address differences in the work environment within workplaces which are related to different occupations.

The baseline specification for the WERS estimates in Table 3 is a simple cross-sectional regression. The pattern of results is similar to that in Table 2 although coefficients are slightly bigger and the female ‘brawn’ coefficient is small but positive. Including firm fixed effects attenuates the ‘people’ and ‘brawn’ estimates but less so the ‘brains’ coefficient. Notably, the basic conclusion remains intact that female satisfaction is more strongly related to the ‘brains’ and ‘people’ aspects of an occupation compared to males.¹⁷

[Table 3 about here]

III. ANALYSIS OF STUDENT SURVEY DATA

The Survey

Individuals’ satisfaction in a setting may be due to ex-post rationalization; women may have come to like the jobs they chose for some different reason. In order to get at job preferences at an earlier stage in life and to be able to ask individuals directly about the reasons for their choices, we conducted our own survey among students in Year 11 (about age 15–16). We ran the survey in two secondary schools in Greater London,

both of which are high-performing schools with students from relatively advantaged backgrounds (the students go to university at a rate that puts them in the top third in the country). These students are at an age where they are thinking about subject and job choices for the future but will not have engaged in actual work experience. The students completed the surveys in an assembly hall on a day when one of us visited the school. All students who were present on the day participated with nobody choosing to opt out. We received 311 responses and dropped four who provided no gender information. The resulting dataset contains 157 males and 150 females.¹⁸

The survey presented students with a list of 12 occupations and gave them six choices among pairs of occupations. We started by splitting occupations into three classes by earnings, and then each of these into occupations with high or low average hours. These matches, particularly on earnings, are relatively coarse in practice. We picked a pair of occupations for each of these groups. As most of the students in our survey schools will go to university, we started with a list of occupations in which both male and female graduates commonly work. We then picked pairs in order to obtain a large amount of variation in the ‘people’ and ‘brawn’ factors within the pair, as graduate jobs tend to have less variation in the ‘brains’ dimension; see Appendix E for more details.

Why did we choose actual occupations and not vignettes? We are not really interested in varying a discrete and easily described aspect of the job (as in how many hours you work). It is difficult to think of a description of, say, a financial analyst job and an alternative that is similar in all aspects except that it involves more personal interaction. Our respondents are likely to have thought about actual occupations and occupational choice because they are about to make important subject choices in school. But it is

unlikely that they think about these choices in the types of abstract categories like ‘people’, ‘brains’ and ‘brawn’, which we find useful as social scientists. We are also worried that focused descriptions of aspects of an occupation involve priming of the respondents.

In Appendix E Table E.1 we list the six pairs of occupations, together with the average earnings and hours, the PBB scores, and the fraction of males among the students who chose each occupation. The students’ choices closely mimic the sex distribution among actual workers.

Analysis and Results

In order to relate the six occupational choices to the PBB factors, we treat the resulting data as a set of binary choices from a multinomial list of preferences over a large set of occupations. We show in Appendix E that a standard random utility model gives rise to a simple pooled logit regression for these data. Because the choice is one between a pair of occupations, it is only the relative characteristics of the two occupations that matter. Our covariates are therefore the differences in the occupation specific variables between the first and second occupation in the group, and the dependent variable is 1 if the first occupation is chosen.

Table 4 shows odds-ratios from these logit regressions of the occupational choices on the PBB factors. Both genders prefer ‘people’ oriented jobs and are relatively indifferent to the ‘brain’ and ‘brawn’ aspects of the jobs. Despite the qualitative similarities, females gravitate more strongly to people orientated jobs compared to

males. Curiously, in terms of the point estimates, males dislike brawn jobs, while females are indifferent to brawn. However, the male effect is not significant.¹⁹

One worry is that these choices might be spuriously driven by skills the students possess rather than their preferences for the job content. In columns (3) to (6), we therefore control for whether the skills required in the occupation are a particularly good match for the specific talents of the students.²⁰

We define two measures of a skill match for a student-occupation pair, a continuous and a discrete one. Columns (3) and (4) in Table 4 show the results adding the continuous skill match measure, and columns (5) and (6) display estimates with the discrete measure. Skills matter for occupational choices for both females and males. Adding the skill match measures lowers the estimates on the ‘people’ factor a bit, raises estimates on the ‘brains’ factor, and further reduces the ‘brawn’ coefficient for males. But the main message from columns (3) to (6) is that the PBB variables and the skills measures both seem to contribute independently to choices. The fact that in columns (4) and (6) males’ dislike of ‘brawn’ jobs is significant at conventional levels and larger than their preference for people jobs is simply a consequence of our choice of the twelve occupations we analyze (see Appendix E, Table E.2 for more details).

[Table 4 about here]

One advantage of our survey is that we can ask the students directly how they made their choices. In particular, we asked: “For each of the six job choices you made, tell us in a few words why you picked the job you did?” The students gave answers in free

form, without any prompts. There was a fair amount of coherency in the answers, and we coded the answers by hand into seven categories, as shown in Table 5. In most cases, this was straightforward to do. When respondents indicated more than one reason for their choice, we coded the one mentioned first.

More than half of responses indicated that the students found one of the activities more interesting, or that the job related to some desirable goals, like helping people (typical examples of answers are “Interest in helping people” or “I enjoy communicating”). About another 16% of responses indicated that they felt they better qualified for one of the jobs (typical answers are “I am creative” or “I am not good at art”). Another 5% indicated some other clearly articulated reason, either related to the environment of the job or some other reason like higher pay or status and a hodgepodge of other things. Respondents did not mention work hours or flexibility in their answers, although we did set up the comparisons so that pay and hours were similar between the pairs of jobs (but this didn’t stop a few respondents from mentioning pay anyway). There is little difference between males and females in how they report making their choices. Gelblum (2020) asked a very similar question in her experiment and finds very similar responses.

The answers indicate that interest in the activity dominates the thoughts of the students as to their job choices. Of course, this does not rule out that these interests correspond to gender stereotypes or norms, or indeed that these children know little about what it means to juggle work and caring responsibilities. However, English students continue with only three or four subjects after age 16, so the choices they make at that age determine which fields are open to them at university, and which occupations they

might enter later. These results therefore reinforce the idea that differences in the strength of preferences may play an important role in the differences in the jobs in which men and women end up.

[Table 5 about here]

IV. DISCUSSION

Stigler and Becker (1977) have famously cautioned economists against relying on variation in preferences to explain economic outcomes, suggesting that the most worthwhile focus is on the comparative statics induced by variation in constraints. The literature on differences in labor market outcomes and behaviors between men and women has indeed for a long time adopted this approach, and studied the impact of discrimination, human capital investments, and labor supply. Less than two decades ago, Altonji and Blank (1999) devoted two paragraphs of their handbook chapter on race and gender to differences in preferences before moving on to the traditional constraint based explanations.

But stubborn differences in male and female pay and occupational segregation persist while many of the constraints faced by women in the workplace seem to have diminished (which does not mean that these constraints are all gone). At the same time, economists have grown more relaxed in terms of thinking about differences in tastes. The handbook chapter by Bertrand (2010), a mere ten years after Altonji and Blank, focuses almost entirely on explanations based on differences in psychological traits between men and women, as well as gender identity. We have argued that a potent form in which such psychological differences might manifest themselves is in differences in preferences of men and women for the content of the work they do.

Economists should be open-minded that this may help explain occupational sorting, and subject this idea to scrutiny.

Here we have offered an initial attempt at this by analyzing the differences in job satisfaction of women in jobs which we loosely characterize by their ‘people,’ ‘brain,’ and ‘brawn’ content. We find that women care more about these job characteristics than men, however the direction of preference effects are the same for men and women. In addition, the same job content measures predict retention in the occupation more strongly for women than for men. These results are consistent with a role for differences in preferences for the content of the work individuals do in their job and how they feel about their work. Our discrete choice experiment with high school students corroborates the conclusions that males and females differ in the extent that they care about job content, with both genders reporting that affinity to the type of work is most important for their choice.

These results are consistent with a story which runs along the following lines: women care about the content of the work they do more than men, and this influences occupational choices. Most importantly, women stay away from traditional blue-collar jobs, probably because of a combination of tastes and skill based comparative advantage (Weinberg, 2008, and Baker and Cornelsen, 2018). But even within white-collar jobs, women sort systematically into occupations which are high both on ‘people’ and ‘brains’ content. This may explain why women choose occupations in business, law, and the health sector over technical and scientific jobs. Unfortunately, jobs with a lot of human contact are also typically jobs which require coordination and restrictions on work schedules and flexibility (Goldin, 2014). Advancement in these occupations often requires substantial dedication to the job, and career interruptions or

part-time work are heavily penalized (see also Landers et al., 1996). Therefore, the well-educated women in these occupations are exposed to a large pay penalty once they decide to have children. As a result, differences in labor market outcomes between young, childless women and men are slight but large pay gaps that emerge once women have children (Kleven et al., 2019). This story might be stylized, hides a lot of heterogeneity, and leaves out other factors which matter, but we believe it captures the important elements.

ACKNOWLEDGMENTS

We thank Krittika Ray and Liaoliang Zhang for excellent research assistance. We are grateful to the editors, Orianna Bandiera and Noam Yuchtman, two anonymous referees, Alan Manning, Kelly Bedard, Dan Black, Alex Bryson, Deborah Cobb Clark, Stephen Machin, David Neumark, Ron Oaxaca, Barbara Petrongolo, David Ribar, and participants of the CEP Labor Market Workshop at the LSE, 14th IZA SOLE transatlantic workshop, 18th IZA European Summer School in Labor Economics, 2016 Royal Economic Society meeting, and 2016 Society of Labor Economics Meeting for their comments. We are grateful to the ESRC [ES/M010341/1] for support through the Centre for Economic Performance at the LSE.

NOTES

1. Based on the 2014 Current Population Survey (CPS) merged outgoing rotation group data.
2. In the social sciences more broadly, Hochschild (1989) is an early advocate of this view. See also Cortés and Pan (2016).
3. Bütikofer et al. (2018) similarly find a larger childhood penalty for women in law compared to STEM, but more women work in law.
4. These results are from our own calculations based on the data posted with Goldin's (2014) article using file AllOccsWageGaps.xlsx, sheet FullBA, EducTime plus Hours.
5. Bond and Lang (2019) warn that the formal conditions for satisfaction scales to carry the information necessary to draw infallible conclusions are almost certainly not met. We are comforted by the fact that Benjamin et al. (2012) and Benjamin et al. (2014) are a bit more optimistic about the practical validity of satisfaction data.
6. Appendix B tables B.5 and B.6 document estimates that create the 'people' 'brains' and 'brawn' factors based on the full 249 distinct items from ONET version 5, and estimates are robust to this change.
7. See Appendix B Tables B.3 and B.4 for a list of the top and bottom ten occupations for each of the three factors, and also the specific scores for a number of occupations.

8. Give that this outcome relies on comparing occupation codes across periods, this analysis omits the year 2000 given the change in occupation coding.
9. We utilize the 1980 wave of the NLSY to create the stayers variable for 1982, so the stayers sample starts in 1982 comparable to the one for the job satisfaction regressions.
10. This outcome relies on comparing occupation codes across periods, therefore this analysis omits the year 2002 from the analysis given the change in the occupation codes.
11. For the BHPS, this amounts to the inclusion of 19 fixed effects. For Russia, we include eight individual residential site indicators.
12. Marginal rates of substitution would be the same if females also have commensurately higher coefficients on income or consumption. At least in simple regressions including the own wage (shown in Appendix C Tables C4-C6), this is not the case (but these regressions also contain occupational averages).
13. In Appendix C Table C.10 we estimate the same equations with main effects and female interactions. The female differences are significant for two of the ‘people’ coefficients, all the ‘brains’ coefficients except in the Russian satisfaction equation, and all the ‘brawn’ coefficients except for US stayers.
14. In Appendix C Table C.8 we also show estimates for college educated females. While individual coefficients jump around the general pattern of results is very similar to those in Table 2. In Appendix C Table C.9 we also present separate estimates for women with and without children. For about half the coefficients, job satisfaction and retention in the occupations high in the ‘people’ and ‘brains’ factors and low in ‘brawn’ tends to be as strong or stronger for women without children as it is for women with children. In most of the remaining cases, the results for women without children fall in between women with children and men. Only three of the coefficients in the table are virtually the same for women without children as they are for men. While the results are far from clear-cut, they are more aligned with the idea that women differ from men, rather than women differing from each other depending on whether or not they have children.
15. Using their central estimate of 0.3 (Stevenson and Wolfers, 2008, p. 31).
16. We note that personal income is also more significant in explaining job satisfaction and the propensity to stay for males as compared to females (see Appendix C Tables C.4 through C.6). This may suggest that males are more extrinsically motivated than females. Together with the PBB

results, this might explain why females sort more frequently into careers like social work, which are low paid but relatively high on ‘people.’

17. Appendix C Table C.11 shows these estimates with female interactions. The female differences are significant for the ‘people’ and ‘brains’ coefficients but not for the ‘brawn’ coefficients.
18. The questionnaire of the survey is included in Appendix E. Students were advised beforehand they could opt out or choose to passively not answer any or all questions. Ethical approval was received by the authors from their home institution.
19. We note that in a non-linear model like a logit, group comparisons like those between males and females could be done in different ways; e.g. one could compare raw coefficients, odds-ratios, or log odds. We therefore don’t want to over-interpret these results.
20. To proxy students’ skills we asked students which subjects they are taking and which subject is their best one. We combined this information with the fields of study listed by respondents to the American Community Survey from 2009 to 2015 to create measures for the skill match between the best subject of the students and the fields highly represented in the occupation (see Appendix E for more details). These are crude measures of skills and may well capture other factors. As a result, it is far from clear that the regressions with the skill measures are superior.

REFERENCES

- ALTONJI, J. G. AND BLANK, R.M. (1999). Race and gender in the labor market. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Vol. 3C. Amsterdam: Elsevier, 3143-3259.
- AUTOR, D. H., LEVY, F. AND MURNANE, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, **118** (4), 279-1333.
- BAKER, M. AND CORNESLON K. (2016). Gender based occupational segregation and sex differences in sensory, motor and spatial aptitudes. *Demography*, **55** (5), 1749-1775.
- BENJAMIN, D.J., HEFFETZ, O., KIMBALL, M.S. AND REES-JONES, A. (2012). What do you think would make you happier? What do you think you would choose? *American Economic Review*, **102** (5), 2083-2110.
- BENJAMIN, D.J., HEFFETZ, O., KIMBALL, M.S. AND REES-JONES, A. (2014). Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices. *American Economic Review*, **104** (11), 3498-3528.
- BERTRAND, M. (2010). New perspectives on gender. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Vol. 4B. Amsterdam: Elsevier, 1545-1592.
- BERTRAND, M. (2018). Coase lecture – The glass ceiling. *Economica*, **85** (338), 205–231.

- BOND, T. N. AND LANG, K. (2019). The sad truth about happiness scales. *Journal of Political Economy*, **124** (4), 1629-1640.
- BÜTIKOFER, A., JENSEN, S. AND SALVANES, K.G. (2018). The role of parenthood on the gender gap among top earners. *European Economic Review*, **109**, 103-123.
- CAMERON, A. C., GELBACH, J. B., AND MILLER, D. L. (2011). Robust inference with multi-way clustering. *Journal of Business and Economic Statistics*, **29** (2), 238-249.
- CHANG, E. (2018). *Brotopia: Breaking Up the Boys' Club of Silicon Valley*. New York: Portfolio Penguin.
- CORTES, G.M., JAIMOVICH, N. AND SIU, H.E. (2021). The growing importance of social tasks in high-paying occupations: Implications for sorting. Working Paper, New York University.
- CORTÉS, P. AND PAN, J. (2016). Prevalence of long hours and skilled women's occupational choices. IZA Discussion Paper No. 10225.
- CORTÉS, P. AND PAN, J. (2018). Occupation and gender. In S. L. Averett, L. M. Argys and S. D. Hoffman (eds.), *Oxford Handbook on Women and the Economy*. Oxford: Oxford University Press.

- CROSON, R. AND GNEEZY, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, **47** (2), 448-474.
- DEMING, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, **132** (4), 1593-1640.
- DIAMOND, P. A., AND HAUSMAN, J. A. (1994). Contingent valuation: Is some number better than no number? *Journal of Economic Perspectives*, **8** (4), 45-64.
- EPL CORNELL (2014). *Claudia Goldin on gender equality in the labor market*. [online video] Available at: <https://www.youtube.com/watch?v=9kgmmPHxe1E> [Accessed 4 January 2018].
- FINKELSTEIN, A., LUTTMER, E.F.P. AND NOTOWIDIGDO, M.J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, **11** (S1), 221–258.
- FORTIN, N. M. (2008). The gender wage gap among young adults in the United States – The importance of money versus people. *Journal of Human Resources*, **43** (4), 884-918.
- FRIJTERS, P. AND VAN PRAAG, B. M. S. (1998). The effects of climate on welfare and well-being in Russia. *Climatic Change*, **39** (1), 61–81.
- GELBLUM, M. (2020). Preferences for job tasks and gender gaps in the labor market. Working Paper, Harvard University.
- GOLDIN, C. (2014). A grand gender convergence: Its last chapter. *American Economic*

Review, **104** (4), 1091-1119.

GORSUCH, R. L. (1983). *Factor Analysis*. Second edition, Hillsdale: Lawrence Erlbaum Associates.

HAKIM, C. (2000). *Work-Lifestyle Choices in the 21st Century. Preference Theory*. Oxford: Oxford University Press.

HOCHSCHILD, A, AND MACHUNG, A. (1989). *The Second Shift: Working Parents and the Revolution at Home*. New York: Viking.

HUNT, J. (2016). Why do women leave science and engineering? *Industrial and Labor Relations Review*, **69** (1), 199-226.

HWANG, H., MORTENSEN, D.T. AND REED, W.R. (1998). Hedonic wages and labor market search. *Journal of Labor Economics*, **16** (4), 815-847.

KIMBALL, M. AND WILLIS, R. (2006). Utility and Happiness. Mimeographed, University of Michigan.

KLEVEN, H., LANDAIS, C. AND SØGAARD, J.E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, **11** (4), 181-209.


LANDERS, R.M., REBITZER, J.B. AND TAYLOR, L.J. (1996). Rat race redux: Adverse selection

- in the determination of work hours in law firms. *American Economic Review*, **86** (3), 329-348.
- LANG, K. AND MAJUMDAR, S. (2004). The pricing of job characteristics when markets do not clear: theory and policy implications. *International Economic Review*, **45** (4), 1111-1128.
- MANNING, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton: Princeton University Press.
- MANNING, A. AND SWAFFIELD, J. (2008). The gender gap in early-career wage growth. *Economic Journal*, **118** (530), 987-1024.
- MAS, A. AND PALLAIS, A. (2017). Valuing alternative work arrangements. *American Economic Review*, **107** (12), 3722-3759.
- PINKER, S. (2008). *The Sexual Paradox: Troubled Boys, Gifted Girls and the Real Difference Between the Sexes*. New York: Macmillan.
- STEVENSON, B. AND WOLFERS, J. (2008). Economic growth and happiness: Reassessing the easterlin paradox. *Brookings Papers on Economic Activity*, **2008** (1), 1-87.
- STIGLER, G. J. AND BECKER, G. S. (1977). De gustibus non est disputandum. *American Economic Review*, **67** (2), 76-90.

- SU, R., ROUNDS, J. AND ARMSTRONG, P.I. (2009). Men and things, women and people: A meta-analysis of sex differences in interests. *Psychological Bulletin*, **135** (6), 859-884.
- THOMPSON, B. (2004). *Exploratory and Confirmatory Factor Analysis: Understanding Concepts and Applications*. Washington, DC: American Psychological Association.
- USUI, E. (2008). Job satisfaction and the gender composition of jobs. *Economics Letters*, **99** (1), 23-26.
- VAN PRAAG, B. M. S. AND BAARSMA, B. E. (2005). Using happiness surveys to value intangibles: The case of airport noise. *Economic Journal*, **115** (500), 224-246.
- VAN PRAAG, B. M.S. AND FERRER-I-CARBONELL, A. (2008). *Happiness Quantified: A Satisfaction Calculus Approach*. Revised edition, Oxford: Oxford University Press.
- WEINBERG, B. (2000). Computer use and the demand for female workers. *Industrial and Labor Relations Review*, **53** (2), 290-308.
- WISWALL, M. AND ZAFAR, B. (2018). Preference for the workplace, human capital, and gender. *Quarterly Journal of Economics*, **133** (1), 457-507.
- ZAFAR, B. (2013). College major choice and the gender gap. *Journal of Human Resources*. **48** (3), 545-595.

FIGURE 1


THE SHARE OF MALES IN JOBS HELD BY FEMALES



Notes: The lines in this graph show the share of males (SOM) in the occupations in which females work in a particular year in the US. The top line uses Census data and is based on the SOM in each occupation in 1950 using the IPUMS 1950 consistent occupation code. The other lines use annual CPS data. In the second line, SOM in an occupation is calculated based on the 1968 data. The bottom line uses the current occupation codes and fixes the SOM in the year the current code was first introduced. The line is broken whenever a new set of occupation codes comes into use.

FIGURE 2

TRENDS IN THE SHARE OF MALES IN SELECTED WHITE COLLAR JOBS



Notes: This graph shows the share of males in selected white-collar occupations in the US Census.

TABLE 1**THE RELATIONSHIP BETWEEN THE SHARE OF MALES AND PEOPLE, BRAINS, AND BRAWN**

	Samples		
	US – Census	Britain – QLFS	Russia – RLMS
People	-0.031 (0.014)	-0.057 (0.013)	-0.124 (0.029)
Brains	-0.012 (0.017)	-0.029 (0.022)	-0.001 (0.021)
Brawn	0.067 (0.024)	0.102 (0.018)	0.183 (0.025)
Number of Observations	14464167	4266356	328371

Notes: All regressions also include the averages of the log hourly wage, hours, fraction college graduates, and age in the occupation, as well as time and area effects.

Standard errors are clustered by occupation.

TABLE 2

INDIVIDUAL FIXED EFFECTS REGRESSIONS

Dependent Variable	Samples							
	US – NLSY		Britain – BHPS		Britain – BHPS		Russia – RLMS	
	Females	Males	Females	Males	Females	Males	Females	Males
	Overall Job Satisfaction		Overall Job Satisfaction		Satisfaction with Work Itself		Overall Job Satisfaction	
People	0.021 (0.006)	0.011 (0.006)	0.028 (0.010)	0.022 (0.009)	0.063 (0.014)	0.036 (0.010)	0.022 (0.015)	-0.003 (0.017)
Brains	0.072 (0.008)	0.046 (0.008)	0.029 (0.013)	-0.006 (0.011)	0.032 (0.018)	-0.012 (0.012)	-0.009 (0.013)	0.024 (0.014)
Brawn	-0.031 (0.008)	-0.000 (0.006)	-0.046 (0.014)	-0.016 (0.012)	-0.053 (0.017)	-0.010 (0.013)	-0.060 (0.016)	-0.040 (0.015)
Number of Observations	91234	97638	49606	46099	49606	46099	35443	27117
Dependent Variable	Stayers							
People	0.002 (0.003)	0.008 (0.003)	0.033 (0.010)	0.019 (0.009)			0.003 (0.015)	-0.026 (0.015)
Brains	0.033 (0.004)	-0.001 (0.004)	0.022 (0.017)	-0.009 (0.012)			0.030 (0.012)	0.001 (0.012)
Brawn	0.000 (0.004)	0.012 (0.003)	-0.044 (0.013)	0.012 (0.012)			-0.023 (0.015)	0.012 (0.014)
Number of Observations	91234	97638	48116	44862			23449	16792

Notes: All regressions also include age and age squared of the individual, the averages of the log hourly wage, hours, fraction college graduates, and age in the occupation, time and area as well as individual fixed effects. Standard errors are two-way clustered (by individual and their occupation) and shown in parentheses. Models are estimated using `xtivreg2`.

TABLE 3**SATISFACTION WITH WORK ITSELF REGRESSIONS IN THE WERS**

	Samples			
	Females	Males	Females	Males
	Baseline		Firm Fixed Effects	
People	0.106	0.067	0.038	0.006
	(0.010)	(0.009)	(0.011)	(0.012)
Brains	0.052	0.030	0.070	0.020
	(0.010)	(0.009)	(0.013)	(0.013)
Brawn	0.010	0.026	0.000	0.009
	(0.012)	(0.010)	(0.015)	(0.013)
Number of Observations	20964	17231	20964	17231

Notes: All regressions also include age and age squared of the individual, the averages of the log hourly wage, hours, fraction college graduates, and age in the occupation, along with time effects. Standard errors are two-way clustered by firm and worker's occupation and shown in parentheses. Models are estimated using ivreg2.

TABLE 4

LOGIT REGRESSIONS OF OCCUPATIONAL CHOICES ON PEOPLE, BRAINS, AND BRAWN IN THE SCHOOLS SURVEY

	(1)	(2)	(3)	(4)	(5)	(6)
	Females	Males	Females	Males	Females	Males
People	1.63 (0.13)	1.23 (0.09)	1.46 (0.13)	1.19 (0.09)	1.56 (0.13)	1.25 (0.10)
Brains	0.92 (0.16)	0.81 (0.14)	1.13 (0.21)	0.92 (0.16)	1.07 (0.20)	1.07 (0.20)
Brawn	1.02 (0.12)	0.82 (0.09)	0.97 (0.11)	0.76 (0.09)	0.94 (0.11)	0.65 (0.08)
Skill match (continuous)			1.31 (0.07)	1.33 (0.07)		
Skill match (discrete)					1.68 (0.23)	2.28 (0.29)
Equality of male and female PBB coefficients (p-value)	0.000		0.000		0.000	

Notes: Coefficients shown are odds ratios. Regressions have 886 observations on 150 females and 936 observations on 157 males. Robust standard errors in parentheses.

TABLE 5**JUSTIFICATION GIVEN FOR OCCUPATION CHOICE IN THE SCHOOLS SURVEY**

	(1)	(2)	(3)
Reason	All	Females	Males
Like the activity/impact/job interesting	0.562	0.589	0.536
Good at the skills required	0.160	0.158	0.162
Like the environment of the job	0.035	0.034	0.035
Other	0.021	0.019	0.024
Indifferent between the choices	0.013	0.007	0.019
Uninformative/illegible	0.131	0.128	0.134
No answer	0.078	0.066	0.090

Notes: Based on the question “For each of the six job choices you made, tell us in few words why you picked the job you did?” Answers are in free form, without any prompts, and responses are coded into the eight categories above.