Published in final edited form as:
Q J Econ. 2018 February ; 133(1): 457–507. doi:10.1093/qje/qjx035.

PREFERENCE FOR THE WORKPLACE, INVESTMENT IN HUMAN CAPITAL, AND GENDER*

Matthew Wiswall and **Basit Zafar**
UNIVERSITY OF WISCONSIN, MADISON AND NATIONAL BUREAU OF ECONOMIC RESEARCH ARIZONA STATE UNIVERSITY

Abstract

We use a hypothetical choice methodology to estimate preferences for workplace attributes from a sample of high-ability undergraduates attending a highly selective university. We estimate that women on average have a higher willingness to pay (WTP) for jobs with greater work flexibility and job stability, and men have a higher WTP for jobs with higher earnings growth. These job preferences relate to college major choices and to actual job choices reported in a follow-up survey four years after graduation. The gender differences in preferences explain at least a quarter of the early career gender wage gap.

JEL Codes:

J24; J16

I. Introduction

The persistence of gender gaps in labor market earnings and the failure of standard variables to fully explain the gaps has prompted the search for alternative models and evidence. One explanation for gender wage gaps is that these arise in part by women "purchasing" certain positive job attributes by accepting lower wages, and men accepting higher earnings to compensate for negative job attributes. These preferences for job attributes may then affect human capital investments, even prior to job market entry. However, empirically isolating the role of worker-side preferences for job attributes is difficult because the equilibrium matching of jobs to workers reflects not only the workers' preferences but the firms' preferences as well. Various kinds of labor market frictions, which prevent workers from matching with their most preferred job types, also break the direct connection between observed job choices and worker preferences. Even when the labor market is perfectly competitive, jobs likely vary in many unobserved (to the researcher) characteristics, leading

*Ellen Fu and John Conlon provided excellent research assistance. We would like to thank Joe Altonji for feedback on the survey design. We are also thankful to the editor, Larry Katz; the coeditor; five anonymous referees; and participants at various seminars and conferences for valuable comments. This is a revised version of NBER Working Paper 22173 (April 2016).

SUPPLEMENTARY MATERIAL
An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online. Data and code replicating the tables and figures in this article can be found in Wiswall and Zafar (2017), in the Harvard Dataverse, doi:10.7910/DVN/MLOGDLDVN/MLOGDL.

to an omitted variable (selection bias) problem in identifying worker preferences from realized job choices.

To address these empirical challenges, this article estimates individual preferences for workplace attributes using a survey of undergraduates from a selective university, New York University (NYU). We collect data on job attribute preferences by presenting undergraduate students with a series of hypothetical job choice scenarios and eliciting their job choices. The hypothetical job scenarios were constructed to offer students a realistic menu of potential jobs varying in expected earnings and other characteristics such as future earnings growth, dismissal probability, and work hours flexibility. The students' stated preferences for these jobs allows us to construct a "pure" measure of individual preferences—at the time of the survey—for various job characteristics, and estimate, in a simple and robust way, the distribution of their preferences for job attributes. Our data isolate the preference for workplace attributes, free from making explicit assumptions about the equilibrium job allocation mechanism or preferences of employers.

We then use the preference data to examine two channels through which preferences could affect the gender wage gap. First, job preferences could affect college major choice, as students perceive that graduating with certain degrees would result in different types of jobs being offered to them. Second, even for students who choose the same major, preferences for job types could cause men and women to accept different types of jobs and result in different earnings. Our data on job preferences, combined with data on perceptions about the job characteristics given major choices, allow us to quantify these two channels and document how workplace preferences affect the gender gap.

Our hypothetical choice methodology is a kind of "stated choice" analysis, similar to "conjoint analysis" and "contingent valuation" methods, used in fields including marketing, environmental and natural resource economics, and health.[1] Because our data collection in essence conducts a kind of "experiment" at the individual student level, the "panel" data generated by our design allows us to estimate the distribution of preferences, allowing for unrestricted forms of preference heterogeneity. In contrast to our approach, previous work addressing compensating differentials using observed job choices requires generally stronger assumptions about preferences and the firm side of the labor market.[2]

In our sample of recent high-ability undergraduate students from NYU, we find substantial willingness to pay (WTP) for pecuniary and nonpecuniary aspects of jobs and considerable heterogeneity in their preferences for workplace attributes. We find that students have preferences reflecting a distaste for higher job dismissal potential, and a taste for workplace hours flexibility (the possibility of working part-time, rather than full-time, hours). We

1. In the marketing and environmental contexts, these methods are often used to identify preferences for new, as yet unavailable consumer products or for public goods like environmental quality, for which realized choices and markets do not exist. Our primary motivation for collecting hypothetical choice data is not because labor markets and realized choices do not exist, but to resolve problems of endogeneity of realized job choices.

2. Recent work has incorporated nonwage components into rich models of the labor market and education choices, allowing for important features such as search frictions, preferences over unobserved job attributes, and dynamic incentives for occupation and education choices (see for example, Bonhomme and Jolivet 2009; d'Haultfoeuille and Maurel 2013; Bronson 2015; Lim 2015). Motivating our approach, Hwang, Mortensen, and Reed (1998) and Bonhomme and Jolivet (2009) conclude that search frictions can imply small equilibrium wage differentials across jobs when there are in fact substantial preferences for nonwage job amenities.

estimate that on average students are willing to give up 2.8% of annual earnings for a job with a percentage point lower probability of job dismissal and willing to give up 5.1% of their salary to have a job that offers the option of working part-time hours rather than one that does not offer this option. After dividing our sample by gender, we find that women have a much higher average preference for workplace hours flexibility, with an implied WTP of 7.3% compared to 1.1% for men. Women also have a higher average WTP for more secure jobs: they are willing to give up 4% of their salary for a percentage point lower probability of job dismissal (versus a 0.6% WTP for males). On the other hand, men have a higher average WTP for jobs with higher earnings growth: they are willing to give up 3.4% of annual earnings for a job with a percentage point higher earnings growth (the corresponding estimate for women is a statistically insignificant 0.6%).

A natural question is whether preferences recovered from data on hypothetical choices relate to actual occupational outcomes. Using data on reported job characteristics for a subset of our respondents who are employed roughly four years after our original data collection, we find a strong and systematic relationship between estimated preferences and later actual workplace characteristics. Students with strong preferences for flexible hours, distaste for hours, and other nonpecuniary aspects of jobs were later found to be more likely to be working at jobs with those same preferred characteristics. Although these realized job characteristics do not solely reflect preferences (given the issues we raised above), our finding of a correlation between pre–labor market job preferences and later actual job characteristics suggests some added credibility of our research design.

Our finding of large differences in WTP for job amenities between men and women is consistent with prior work noting that women are more likely to be found in jobs offering greater workplace f lexibility (Goldin and Katz 2011; Flabbi and Moro 2012; Goldin 2014; Wasserman 2015; Bronson 2015). However, the observation that women tend to work in certain job types may not reveal women’s preferences alone, but may be affected by firm-side demands for specific workers and discrimination or be driven by some other job attributes that are unobserved in our data sets (Blau and Kahn forthcoming).[3] Our innovation is to quantify the WTP for job attributes using a flexible and robust methodology. In related recent work, Mas and Pallais (forthcoming) conduct a field experiment where call center job applicants are offered various work time schedules and wages. Their finding that women have a higher valuation for worker-friendly alternative work arrangements (and a stronger distaste for employer discretion over their hours) is consistent with our estimates of a higher female valuation of work hours flexibility (availability of part-time work).[4]

We next test whether the job preferences young adults hold in college in fact affect their human capital investments during college. To quantify the importance of job attributes to major choice, we collect additional survey data on students’ beliefs about the characteristics

3. In fact, using a recent nationally representative survey of U.S. workers, Maestas et al. (2016) find that younger college-educated women report less desirable working conditions (including no option to telecommute, higher prevalence of employer setting schedules, and higher incidence of work-related stress).

4. Mas and Pallais also conclude that gender differences in work-time flexibility preferences are not enough to explain any part of the gender gap in earnings, which stands in contrast to our conclusion of a large role (as we discuss later). There could be several reasons for this difference in findings: we measure preferences for several workplace attributes (job stability, earnings growth, hours). In addition, our sample is high skill, and likely to be active in a different segment of the labor market.

of jobs they would be offered if they were to complete different majors. These data are then used to estimate a model of major choice, where students receive utility from major-specific characteristics (such as perceived ability in those majors) and from the job attributes they associate with these majors. We find that job attributes have a sizable impact on major choice. For example, increasing the perceived job firing probability by a standard deviation reduces the probability of pursuing a major, on average, by 5% (4%) for women (men). To put this change in perspective, a standard deviation increase in average earnings leads to a 5% (16%) increase, on average, in the likelihood of majoring in that field for women (men). Thus, for women, this change is equivalent to the effect on major choice of increasing earnings by one standard deviation. We find meaningful effects for other job attributes, such as work hours. In general, we find that women's major choices are more responsive to changes in nonpecuniary job attributes (relative to changes in earnings) than are men's. By linking job preferences directly to human capital investments, we contribute to our limited understanding of how career and workplace preferences shape educational choices.

Prior research on college major choice examines the role of earnings expectations, ability perceptions, college costs, and tastes, but generally does not examine nonpecuniary job attributes.[5] An exception is Zafar (2013), which estimates a model of college major choice that incorporates some nonpecuniary workplace attributes. However, the framework does not allow for unobserved heterogeneity in preferences, and incorporates a smaller set of workplace characteristics. Closely related to our work is Arcidiacono et al. (2015), who study a sample of male undergraduate students and collects expectations about earnings in different major-occupation pairs. They find evidence for complementarities in preferences between different majors and occupations, and conclude that nonmonetary considerations are key determinants of occupational choice (conditional on graduating from a given college major). Our contribution is to directly quantify the role of specific nonmonetary factors in major choice.

Finally, we turn to a key question in the social sciences and ask what our results imply for the gender wage gap. Systematic gender differences in workplace preferences may affect the gender wage gap through two channels: first, it may cause men and women to choose different fields of study, and second, men and women may choose systematically different jobs within the same field. We find that the main channel for preferences to affect the gender gap operates through the second channel, with a smaller effect through major choice. Our analysis reveals that the gender gap in expected earnings early in the career (age 30) would be reduced by at least a quarter if women did not differ from men in the workplace preferences we consider. Remarkably, we find a similar impact on the gender gap in actual earnings for the subset of respondents for whom we have follow-up data.

Our evidence supports the notion that at least part of the early career gender wage gap is the result of women "purchasing" certain positive job attributes by accepting lower wages, and men accepting higher earnings to compensate for negative job attributes. In understanding

5. For examples of recent work, see Arcidiacono 2004; Beffy, Fougere, and Maurel 2012; Arcidiacono, Hotz, and Kang 2012; Stinebrickner and Stinebrickner 2014a; Gemici and Wiswall 2014; Wiswall and Zafar 2015a. Most recently, Bronson (2015) analyzes the importance of work-hours flexibility and changes in divorce law and divorce risk in explaining longer-term gender-specific trends in major choices.

our results, it is important to note that we measure preferences at a particular point in the life cycle of our sample, when our sample was in college. The preferences we measure are not necessarily intrinsic; these preferences were formed by a variety of influences before and during college, and could change substantially after graduation. In addition, it is likely that workplace flexibility issues are much larger determinants of the gender earnings gap for college graduates 10 or more years into their career than for the young college graduates in our study (who are in their mid-20s during our follow-up), as college-graduate women now have children at later ages.[6]

The article is organized as follows. In the next section, we briefly provide some context for our analysis by using nationally representative surveys for the United States on currently employed individuals. Section III describes our data collection; Section IV details the model of job choice and shows how hypothetical data can solve important identification issues with realized choice data. Section V provides the empirical estimates of job preferences. Section VI quantifies the importance of job attributes for college major choice. Section VII investigates the extent to which gender-specific job preferences can explain the gender gap in earnings. Finally, Section VIII concludes. All appendixes are available online.

II. Background: Gender Differences in Job Choices and Human Capital Investments in the United States

To set the stage for the analysis of our hypothetical choice scenario data, we first briefly describe the distribution of college majors, jobs, and associated job characteristics. To do so, we use two large-sample, representative data sets for the United States, the January 2010–December 2012 monthly Current Population Survey (CPS) and the 2013 American Community Survey (ACS).

Table I shows the job attributes across sectors. For this purpose, we use the sample of 25–60-year-old labor market participants with at least a bachelor's degree in the 2010–2012 CPS. The first two columns of Table I show that the gender distribution across work sectors differs (Online Appendix A provides details on how variables in this table were constructed). While nearly half of college-educated women workers are in health or education, less than 20% of college-educated male workers are employed in these sectors. These sectors differ substantially in their labor market returns: column (3) of Table I shows that average annual earnings of full-time workers are the lowest for education and health. These sectors differ along other dimensions as well: more than a quarter of the workers in health and education are employed part-time, possibly suggesting the compatibility of these sectors to work-hours flexibility. Job instability, as measured by the likelihood of being fired, is lowest in the government and education sectors. Of course, jobs in these sectors will also differ in the skills that they demand of their workers. So what explains the propensity of men and women to work in different sectors—is it differences in preferences for workplace attributes, differences in tastes for occupations/industries, or differences in skills? What is the role of the labor market structure, firm labor demand, and discrimination by employers? The

6. Bertrand, Goldin, and Katz (2010) document the rising role of children and hours choices over the first 15 years of the careers of female MBAs from a top U.S. business school.

observed distribution of jobs by gender we see in the data are equilibrium outcomes, and we cannot ascertain from these data alone the extent to which these outcomes are due to worker demand or due to the supply of certain jobs—for example, part-time work may either be a voluntary or involuntary decision.

We next turn to Table II to document the link between field of study and associated job characteristics.[7] The table is based on the 2013 ACS, restricting the sample to 25–40-year-olds with at least a bachelor's degree. The first two columns show that while nearly 55% of women have a bachelor's degree in humanities, less than 40% of men do. While nearly a quarter of men have a bachelor's in engineering, the corresponding proportion for women is only 6%.

Column (3) of Table II shows that these majors differ significantly in their average earnings. Engineering—the field which women are least likely to be present in—has the highest average earnings, while humanities—the most popular bachelor's field for women—has the lowest average earnings. These majors also differ along other dimensions. Columns (4) and (5) show that work-hours flexibility is the highest for jobs associated with humanities: 38% of all humanities graduates are part-time workers, versus 22% of engineering bachelor's graduates. Average hours per week for full-time workers are also the lowest in humanities. The last two columns of the table show that job stability and earnings growth also vary significantly across the fields of study.

So how much do these gender differences in human capital and job characteristics explain the gender gap in earnings? In a recent analysis, Blau and Kahn (forthcoming) find that the gender wage gap is currently larger at the top of the wage distribution (90th percentile), and has decreased more slowly at the top than at other points in the distribution. In addition, they find that traditional human capital variables (experience and degrees earned) explain little of the recent gender gap. They attribute part of the gender gap in high-skilled occupations to a possible compensating differential.

Using the sample of college graduates aged 25–40 from the 2013 ACS (the subsample from Table II that reports nonzero labor income), we find an adjusted gender gap in hourly earnings of about 12 log points (adjusting only for age and full-time status).[8] Demonstrating how important college majors could potentially be in explaining the gender gap among college-educated workers, including four broad college major categories (as defined in our analysis) reduces the gender gap by about 43% (from 12 log points to 6.7 log points). Including indicators for detailed occupation, industry, and race categories as in Blau and Kahn (forthcoming), in addition to indicators for major categories, increases the explained portion of the gender wage gap to 58%. However, this analysis reveals that even conditional on detailed occupation/industry and major controls, a large part of the gender wage gap remains unexplained. The remainder of this article investigates the extent to which

7. Altonji, Kahn, and Speer (2016) provide a more detailed discussion of the relationships between college majors and labor market outcomes.

8. The unadjusted hourly earnings gap is 21.6 log points. For college graduates ages 25–40, the mean earnings for full-time employed men is higher than the mean earnings for full-time employed women by 36%. The median for full-time men is higher than the female median by 28%.

workplace preferences can explain this gender gap, either by influencing human capital choice (major choice) or by influencing job choices conditional on major.

III. Data

This section describes the administration of the data collection, the form of the hypothetical choice scenarios, and the sample we use for the estimation.

III.A. Administration

Our data are from an original survey instrument administered to NYU undergraduate students over a two-week period during May 2012. NYU is a large, selective, private university located in New York City. The students were recruited from the email list used by the Center for Experimental Social Sciences (CESS) at NYU. Students were informed that the study consisted of some simple economic experiments and a survey about educational and career choices. Upon agreeing to participate, students could sign up for a 90-minute session, which was held in the CESS Computer Lab located on the main NYU campus.[9]

The data for this article were collected through a computer-based survey (constructed using the SurveyMonkey software). The survey took approximately 30 minutes to complete and consisted of several parts. Many of the questions had built-in logical checks (e.g., percent chances of an exhaustive set of events such as majors had to sum to 100). Students were compensated $10 as a show-up fee, and $20 for successfully completing the survey.

III.B. Data Collection Instrument

In addition to questions about demographics, family background, and educational experiences, the main survey instrument consisted of two parts. The first part collected data on students' preferences for job attributes using hypothetical job choices, while the second collected data on consequential life activities that would plausibly be key determinants of college major choice, such as attributes of jobs associated with each major and measures of the student's perception of their ability to complete the course-work for each major. We describe the hypothetical job choice data in detail next and leave the description of major-specific data to a later part of the article, where we relate the job attribute preferences to college major choices.

Our hypothetical job choice data were collected by presenting students with a total of 16 job scenarios. Each scenario consisted of three different potential jobs. We exogenously varied different aspects of the job with the intention of creating realistic variation in job attributes. The first eight hypothetical job scenarios were introduced as follows:

> In each of the 8 scenarios below, you will be shown hypothetical jobs offers. Each job offer is characterized by:
>
> Annual earnings when working full-time

[9] During the same session, and immediately prior to completing the survey, students took part in some economic experiments. Students earned additional income through participation in the experiments. See Reuben, Wiswall, and Zafar (forthcoming) for information on this data collection.

> Annual percentage increase in earnings from age 30 onwards until retirement
>
> Full-time work hours per week
>
> Work flexibility (whether part-time work is an option); part-time work is work where you only work at most half as many hours as full-time work and for half of the full-time salary
>
> These jobs are otherwise identical in all other aspects.
>
> Look forward to when you are 30 years old. You have been offered each of these jobs, and now have to decide which one to choose.
>
> In each scenario, you will be asked for the percent chance (or chances out of 100) of choosing each of the alternatives. The chance of each alternative should be a number between 0 and 100 and the chances given to the three alternatives should add up to 100.

Each scenario consisted of three jobs, with each job being characterized by four attributes. The notable point that was highlighted was that these jobs were identical in all other aspects. The jobs did not have any occupation labels on them.[10] The last eight scenarios were introduced in a similar way, except that the job offer was now characterized by a different set of attributes: annual earnings when working full-time, probability of being fired over a one-year period, amount of additional annual bonus pay based on relative performance the respondent may qualify for (in addition to base pay), proportion of men in the firm in similar job positions. All survey respondents received identical scenarios in the same order.

Following the approach of Blass, Lach, and Manski (2010), we asked respondents to provide a choice probability instead of a discrete choice (that is, a 0 or 1). This allows respondents to express uncertainty about their future behavior. It also allows them to rank their choices, providing more information than if we asked only about the most preferred job. As is standard in studies that collect subjective probabilistic data, a short introduction on the use of percentages was provided. In addition, respondents answered some practice questions to become familiar with expressing probabilistic answers.

Besides earnings, the scenarios focus on six different job attributes. We chose not to vary these six dimensions all at once since the cognitive load to process such information could have been overwhelming. We focus on these dimensions based on findings from prior literature, and the fact that there is considerable variation along these dimensions across occupations as well as majors (Tables I and II). Earnings and earnings growth were included since they have been found to be a factor in career/education choice (see Wiswall and Zafar 2015a, and references therein). Work hours and work flexibility are included because they tend to be associated with the remuneration structure in jobs and the associated gender gap in earnings (Flabbi and Moro 2012; Goldin 2014; Cortes and Pan 2016). We recognize that workplace flexibility is a multidimensional concept: for example, the number of hours to be worked matters but perhaps so do the particular hours (Goldin 2014; Mas and Pallais

10. In addition, when presented with each scenario, respondents were told: "Now consider the situation where you are given the jobs offered above when you are aged 30, and you have decided to accept one of these jobs. What is the percent chance (or chances out of 100) that you will choose each of these jobs?" That is, the options were mutually exhaustive, and not working was not an option.

forthcoming). We varied two hours-related attributes: number of hours and the availability of a part-time option, since these are easy to vary in a meaningful fashion. Job stability, as proxied by the likelihood of being fired from the job, is included because of the importance of risk and uncertainty to job choices (Dillon forthcoming) and gender differences in risk preferences (Croson and Gneezy 2009). Finally, relative performance compensation and proportion of men are meant to capture the competitiveness of the job environment, preferences for which have been found to differ by gender (Niederle and Vesterlund 2007; Flory, Leibbrandt and List 2015; Reuben, Wiswall, and Zafar forthcoming).[11]

To keep the scenarios realistic, the job attributes shown to respondents in the scenarios were based on the actual marginal distribution of job characteristics in the CPS (except for the bonus pay variable, since data were not available for that dimension).[12] In addition, no scenario included a job that was clearly dominant or dominated along all dimensions. We also made a conscious effort to keep the variation in job attributes within each scenario relatively "local," so that the claim that the jobs were otherwise identical was credible; for example, two jobs offering $50,000 and $90,000, respectively, with little variation along the specified dimensions are unlikely to be identical. At the same time, we had substantial variation in the job attributes across the scenarios. This ensures that we are not recovering preferences in a local region only. Online Appendix Table A1 shows the range of the attributes across the scenarios.

III.C. Sample Description

A total of 257 students participated in the study. We drop 10 respondents for whom we have missing data for the relevant section of the survey. Sample characteristics are shown in Table III. Thirty-five percent of the sample (86 respondents) is male, 29% is white, and 51% is Asian. The mean age of the respondents is 21.5, with 11% of respondents freshmen, 11% sophomores, 37% juniors, and the remaining seniors or higher. The average grade point average of our sample is 3.5 (on a 4.0 scale), and students have an average Scholastic Aptitude Test (SAT) math score of 696, and a verbal score of 674 (with a maximum score of 800). These correspond to the 93rd percentile of the U.S. national population score distributions. Therefore, as expected, our sample represents a high ability group of college students. Parents' characteristics of the students also suggest that they are overrepresented among high socioeconomic groups. The last panel of the table shows that 48% of the students have a major in the humanities and social sciences category, 31% have a major in business and economics, while the remaining have a major in natural sciences and math (16%), and engineering (5%).

Columns (2) and (3) of Table III report the characteristics by gender. The last column of the table reports the *p*-value of tests of equality of the statistics by gender. We see that male and female respondents are similar in all dimensions, except two. First, male students in our

11. Lordan and Pischke (2016) find a strong relationship between women's job satisfaction and the proportion of men in that occupation.

12. For each job attribute, we constructed a set of hypothetical job scenarios by using uniform random draws from an interval between the 10th and 90th percentile of the observed distribution for each attribute. For each set of job scenarios, we then rejected any set of job scenarios which included jobs which were dominated by another job in all attributes or had earnings differences across jobs which were greater than 30%.

sample have a significantly higher average SAT math score than females, of about 33 points. Second, the two sexes choose very different college majors. Nearly half (49%) of men report majoring in business/economics, with 30% majoring in humanities and social sciences, and 13% in natural sciences/math. On the other hand, 57% of the women report majoring in humanities and social sciences, followed by about 22% majoring in business/economics, and 18% majoring in natural sciences/math. That is, female students are almost twice as likely as men to major in the humanities (the field, as we show below, perceived to have the lowest average earnings among college graduates), and only half as likely as males to major in economics/business (the perceived highest-earnings major category). The gender-specific major distributions are statistically different (p-value $\leqslant .001$, using a chi-square test for equality of distributions). These substantial gender gaps in major choice mirror the national patterns from the ACS data (Table II).

Compared to the NYU population, our sample has a similar proportion female: 63% of students graduating NYU in 2010 are women compared with 65% in our sample (data from the Integrated Post-Secondary Education Data System, IPEDS). For all incoming freshman in 2010, the 25th and 75th quartiles of the SAT math were 630 and 740 and for the SAT verbal were 610 and 710 (IPEDS). The equivalent quartiles in our sample are 650 and 770 for math and 620 and 730 for verbal. Our sample is weighted more toward business/ economics majors than in the actual NYU population graduating in 2010, possibly because the experimental laboratory is located in the building housing the Economics Department. However, the gender differences in major choice are similar.[13]

IV. Model and Identification Analysis

In this section, we present a simple attribute-based job choice model and discuss identification of the model using two types of data: (i) standard realized job choices (as observed after job offers and acceptances are made), and (ii) stated probabilistic job choices (as observed in our job hypotheticals experimental data). We show that under weak conditions the job hypotheticals data identify the distribution of job preferences, while standard realized job choice data do not.

IV.A. A Canonical Random Utility Model of Job Choice

Jobs are indexed by j, and there is a finite set of jobs $j = 1, \ldots, J$. Each job is characterized by a vector of K attributes $X_j = [X_{j1}, \ldots, X_{jK}]$. These job attributes include earnings and various nonpecuniary attributes, such as job dismissal probabilities and work-hours flexibility. Thus, we explicitly allow for the possibility that individuals are not necessarily pure income or consumption maximizers, and may value many other outcomes associated with their job choice.

Let $U_{ij} \in R$ be individual i's utility from job j. The utility from job j is

13. For the NYU population of students who graduated in 2010 (IPEDS), the fraction of students completing degrees in each field are as follows: for women, 14.1% graduated in economics or business, 71.7% in humanities or other social sciences, and 13.7% in natural sciences, math, or engineering. For men, 31.1% graduated in economics or business, 61.2% in humanities or other social sciences, and 7.8% in natural sciences, math, or engineering.

$$U_{ij} = u_i(X_j) + \epsilon_{ij}. \quad (1)$$

$u_i(X) \in R$ is the preferences of individual i over the vector of characteristics X. $\epsilon_{ij} \in R$ is the additional job-specific preference component for job j reflecting all remaining attributes of the job which affect utility, if any. Let ϵ_i be the vector of these components for individual i, $\epsilon_{i1} = \epsilon_{i1}, \ldots, \epsilon_{iJ}$. After observing the attributes $X_1, \ldots, X_J$ for all jobs and ϵ_i, individual i chooses the one job with the highest utility: i chooses job j if $U_{ij} > U_{ij'}$ for all $j' \neq j$.

Population preferences for jobs is the collection of u_i preferences over the job attributes X and the job-specific components ϵ_i. The joint distribution of preferences in the population is given by $F(u_i, \epsilon_i)$. This distribution determines the fraction of individuals choosing each job, $q_j \in [0, 1]$: $q_j = pr(\text{choose job } j)$

$$\begin{aligned} q_j &= pr(\text{choose job } j) \\ &= \int 1\{U_{ij} > U_{ij'} \text{ for all } j' \neq j\} dF(u_i, \epsilon_i). \end{aligned} \quad (2)$$

IV.B. Identification Using Realized Choice Data

Typically empirical research on job choice consists of analyzing data on actual or realized job choices, which provides the one best job chosen by each individual.[14] To analyze the potential advantages of hypothetical data, we first detail the identification using realized choice data.

A common model of realized choice data assumes $\epsilon_{i1}, \ldots, \epsilon_{iJ}$ are i.i.d. Type I extreme value, and independent of preferences represented by u_i. The probability individual i chooses job j, given some characteristics $X_1, \ldots, X_J$ for all jobs, is given by

$$q_{ij} = \frac{\exp(u_i(X_j))}{\sum_{j'=1}^{J} \exp(u_i(X_{j'}))}.$$

The population fraction choosing job j is then

$$q_j = \int \frac{\exp(u_i(X_j))}{\sum_{j'=1}^{J} \exp(u_i(X_{j'}))} dG(u_i), \quad (3)$$

where we have kept the dependence of the job choice on the job characteristics $X_1, \ldots, X_J$ implicit. $G(u_i)$ is the distribution of preferences over attributes u_i in the population. Equation (3) is the mixed multinomial logit model of McFadden and Train (2000). They show that the

14. We confine attention to cross-sectional data. Panel data on repeated job choices over an individual's life cycle may provide more identifying power but at the cost of requiring additional assumptions about the evolution of model features (e.g., preferences) as individuals age.

distributional assumption on the ϵ_i terms that yield the logit form is without any meaningful loss of generality as this model can arbitrarily closely approximate a broad class of random utility models. For ease of exposition, we consider a linear model of utility given by $u_i(X) = X'\beta_i$.

A key concern in using realized job choices is that the data set of job characteristics which the researcher has at hand is not complete in the sense that there are omitted unobserved job characteristics that are potentially correlated with the included observed characteristics. Divide the vector of job characteristics X into observed $X(obsv)$ and unobserved characteristics $X(unob)$, $X = [X(obsv), X(unob)]$. Similarly divide the vector of preference parameters $\beta_i = [\beta_i(obsv), \beta_i(unob)]$. The log odds of job j relative to job j' for individual i is then:

$$\begin{aligned}\ln\left(\frac{q_{ij}}{q_{ij'}}\right) &= (X_j(obsv) - X_{j'}(obsv))\beta_i(obsv) + (X_j(unob) \\ &\quad - X_{j'}(unob))\beta_i(unob) \\ &= (X_j(obsv) - X_{j'}(obsv))\beta_i(obsv) + \eta_{ij},\end{aligned}$$

where q_{ij} and $q_{ij'}$ is the probability of choosing job j and j', respectively, for individual i. $\eta_{ij} = (X_j(unob) - X_{j'}(unob))\beta_i(unob)$ is the omitted variable for individual i.

The omitted variable bias problem is the generic one found in a variety of contexts: the omitted unobserved job characteristics $X_j(unob)$ are correlated with the observed characteristics $X_j(obsv)$. For example, if the researcher's data set includes only current salaries, but not any of the nonpecuniary benefits of the job, we would expect that the estimate of preferences for salaries will be biased. The theory of compensating differentials (Rosen 1987) predicts a close connection among various job characteristics—a trade-off between salary and nonpecuniary benefits—and therefore would suggest important omitted variable bias in estimates of job preferences using realized data.

The omitted variable bias issue could also arise more subtly from the selection/matching mechanism to jobs, reflecting employer preferences over potential job candidates. If the labor market equilibrium is such that employers only offer a limited set of jobs to candidates, then the realized jobs they hold do not reflect their preferences only.[15] Discrimination by employers, by which employers prefer not to hire workers of certain groups (e.g., women, minorities), is one example (Becker 1971). In the presence of important demand-side considerations, one would not want to interpret the equilibrium allocation of jobs as reflecting only worker preferences. As we detail below, our hypothetical data avoid this issue because they experimentally manipulate the characteristics offered to individuals, thereby allowing a "pure" measure of preferences, free from considering the equilibrium job allocation mechanism, preferences of employers, or any omitted unobserved job characteristics.

15. We can represent demand-side restrictions in the omitted variable framework by considering some unobservable job characteristic $X(unob)$, such that $X(unob) \rightarrow -\infty$ if a job is not offered.

Another approach to this issue is to make some assumptions about the structure of the labor market and individual preferences. As in the literature examining identification of these models using observed choices (see Fox et al. 2012 for a recent review), some support condition or restriction on preferences is therefore necessary for identification.

IV.C. Model of Hypothetical Job Choices

We next consider a framework for analyzing hypothetical job choice data, connecting the canonical model of realized job choice specified above in equation (1) with the hypothetical job choice data we collect. Our hypothetical data are asked prior to a job choice (while students are in school). We observe each individual's beliefs about the probability they would take each hypothetical future job offered within the scenario (and not simply the individual's one chosen or realized job). To analyze this type of data, we require a model of hypothetical future jobs. Our model of hypothetical job choices presumes individuals are rational decision makers who anticipate the job choice structure as laid out in the canonical model of job choice, equation (1). To allow for the possibility of uncertainty about future job choices, we assume that the realizations of $\epsilon_{i1}, \ldots, \epsilon_{iJ}$ job-specific utility terms are not known at the time we elicit individual beliefs. Individual i then faces a choice among J hypothetical jobs with characteristics vectors $X_1, \ldots, X_J$. Each individual i expresses their probability of taking a given job j as:

$$p_{ij} = \int 1\{U_{ij} > U_{ij'} \text{ for all } j' \neq j\} dH_i(\epsilon_i), \quad (4)$$

where $H_i(\epsilon_i)$ is individual i's belief about the distribution of $\epsilon_{i1}, \ldots, \epsilon_{iJ}$ elements. As in Blass, Lach, and Manski (2010), ϵ_i has an interpretation of resolvable uncertainty, uncertainty at the time of our data collection but uncertainty that the individual knows will be resolved (i.e., known or realized) prior to making the job choice.[16]

It should be noted that the preferences for workplace attributes elicited in our data collection are potentially specific to the time at which the survey is collected (during the college years in our case). Preferences for job attributes may change as individuals age and may have been different when the students in our sample were younger (say, prior to college) and may be different still when they actually enter the labor market and make job choices. With this caveat in mind, we can still use our research strategy to understand job preferences at a point in time and study how these preferences relate to important human capital investments that are being made contemporaneously.[17]

16. An alternative model is that agents have uncertainty about preferences over attributes, that is the utility function $u_i(\cdot)$ is uncertain. For example, an individual may be uncertain about the number of children she may have at a future date, and the number of young children at home may affect her preference for workplace hours flexibility (an element of the X_j vector). We explore this later by relating preferences for job characteristics as revealed in our hypothetical data with a rich set of beliefs about future outcomes (e.g., individual beliefs about future own fertility and marriage).

17. See Stinebrickner and Stinebrickner (2014a,b) for evidence on the dynamics in beliefs formation among college students.

IV.D. Identification Using Hypothetical Choice Data

We previously analyzed identification of preferences using realized job-choice data and discussed a key shortcoming: realized choice data potentially suffers from omitted variable bias. Hypothetical choice data can overcome this shortcoming and allow a general method to identify heterogeneity in job-choice preferences.

First, because we can experimentally manipulate the hypothetical choice scenarios we provide individuals, we may be able to reduce bias from the correlation of observed and unobserved job characteristics. Rather than use naturally occurring variation in realized job choices—which are in general the result of many unobserved job characteristics and an unknown labor market equilibrium mechanism, as discussed above—we present individuals with an artificial set of job choices. Although the job characteristics we provide are certainly not exhaustive of all possible job characteristics, and are purposely kept limited so as not to "overload" the respondents with too many job features, the key feature of the hypothetical experimental setting is that we instruct respondents that the jobs differ only in the job characteristics we provide, and are otherwise identical. This distinguishes our design from "audit"-based studies in which employers are presented with resumes that are otherwise identical except for the one chosen attribute (say, the gender of applicant). The criticism of audit studies is that even if you make two groups (say, men and women) identical on observables, employers might have very different distributions in mind about unobservables for the two groups, biasing the inference (for an analysis of this issue, see Neumark, Burn, and Button 2015). In our case, students are instructed that the hypothetical jobs are identical in all other ways, instructions that cannot be given to actual employers in audit studies.

The extent of the remaining bias in the preferences we elicit then critically depends on whether respondents fully internalize our instructions that the jobs are otherwise identical. There is reason to suspect this may not strictly be the case. Like audit studies, the participants in our study may still have preconceived notions of what other attributes are related to the attributes we include. For example, they might believe the availability of part-time work (one of the attributes we include) is associated with other aspects of flexibility we do not include, such as time of day one is allowed to work and the ability to take vacations and family care leaves. Dismissal risk (also one of the attributes we include) could be viewed as a proxy for high-stress, high-expectations environments. These types of biases are not different from those present in audit studies where employers have their own prior beliefs about other attributes of workers associated with different observable (on résumé) worker characteristics.

A second advantage of the hypothetical data is that it provides a kind of panel data on preferences which, under fairly weak assumptions, identify the full preference rankings over job attributes. Notice the key distinction between equations (4) and (2). With job hypotheticals data, we observe for each individual i multiple subjective job probabilities p_{i1}, …, p_{iJ}. The job hypotheticals provide a type of panel data allowing less restricted forms of identification, by allowing identification of the $u_i(X)$ preferences without a parametric restriction on the population distribution of preferences. Note that even with a panel of realized choices, it is in general impossible to identify separately preferences for jobs from search frictions or omitted job characteristics. Within our hypothetical setup, these issues are, by design, not a confounding factor.

Our assumption for identification of preferences is that the $\epsilon_{i1}, \ldots, \epsilon_{iJ}$ job-specific terms are i.i.d. and independent of the experimentally manipulated job attributes $X_1, \ldots, X_J$. This is implied by the experimental design: respondents are instructed that the jobs vary only in the listed characteristics and are otherwise identical. Under this assumption, the hypothetical data $p_{i1}, \ldots, p_{iJ}$ identifies the preference ranking for individual i over all jobs J in the choice set: For any two jobs j and j', the characteristics vector X_j is preferred to that of $X_{j'}$ if the probability of choosing that job is higher than that for job j', $p_{ij} > p_{ij'}$.

Our identification concept is that each scenario approximates a multidimensional offer function from which a worker can choose the optimal bundle of job attributes. If this offer function were complete (that is, a continuum of choices rather than three job options in each scenario), the worker would choose the point that is tangent to their indifference curve. Rosen (1987) argues that worker preferences can then be identified if the offer curve shifts, forcing workers to reoptimize in a frictionless labor market, and tracing out the worker's indifference curve. This is effectively what happens when respondents are presented with another job-choice scenario (another set of jobs to choose from) in our survey. The key distinction relative to the Rosen case is that our choice set is discrete, so we can instead think of preferences as being identified by a set of job preference inequalities. This is an important improvement relative to identification using observed job choices because there is information in our data on rejected job opportunities, which are not typically available in real labor-market settings.[18] This rejected-offer information provides both lower and upper bounds on preferences in a discrete-choice setting and can point-identify preferences nonparametrically (up to the distribution of the ϵ_i shocks) with full support of the job offer variation.

In practice, of course we have a only finite number of job scenarios and cannot vary job offers to saturate the full support of the job characteristics. As in the literature examining identification of these models using observed choices (see Fox et al. 2012 for a recent review), some support condition or restriction on preferences is therefore necessary, although more limited than is required using observational data. We assume preferences take a parametric form, $u_i = X_i' \beta_i$, but allow the β_i parameters to be freely varying in the population. This allows for the distribution of preference parameters β_i to be completely unrestricted across individuals; thereby we avoid making assumptions about the population distribution of preferences (such as assuming preferences β_i are normally distributed). In the estimation, we use this identification result constructively and simply estimate preferences for each sample respondent one by one. We then use the sample distribution of preferences as the sample estimator of the population distribution of preferences. Therefore, we allow the distribution of preferences to take any form.[19]

18. In an innovative related approach, Stern (2004) collects data on job offers and accepted jobs from a sample of PhD biologists to estimate the WTP to take a research job over others. However, the limited data on job offers do not allow for identification of heterogeneity in preferences. In addition, this approach only yields unbiased preference estimates in frictionless labor markets.

19. Note that as with any discrete choice setting, the population distribution of preference parameters β_i is identified up to the distribution of the ϵ_i shocks. As we detail below, we assume a logit form for the shocks. For ease of interpretation, we focus on WTP

V. Estimates of Preferences for Job Characteristics

V.A. Variation in Choice Probabilities

Identification relies on variation in probabilities that respondents assign to the various jobs in the hypothetical scenarios. We next present some evidence on this, which should allow the reader to become familiar with the sources of identifying variation. Table IV, Panel A shows two examples from the data sample using the first set of hypothetical scenarios. Recall that each of these eight scenarios included three different job offers, which differed according to the characteristics shown in the table. The last two columns show the mean probability assigned by each gender to the jobs.

Turning to the first example, we see that, for men, Job 3 is the most preferred job in our sample (that is, it received the highest average probability). Job 3 is the job without part-time availability and the highest earnings growth. For women, on the other hand, this job received the lowest average probability. Women assigned the highest probability, on average, to Job 2, the job with a part-time option and an intermediate number of work hours per week and intermediate earnings. In this example, the distribution of choices differs significantly by gender. The gender-specific distributions of average probabilities do not differ in the second example.

Table IV, Panel B shows two examples from the second set of hypothetical scenarios, which vary a different set of attributes. In the first example, the distribution of average probabilities again differs by gender. For women, Job 1 receives the highest probability on average (37%). Job 1 is the job with the lowest probability of being fired and the lowest proportion of men as colleagues. Male respondents, on the other hand, assign the highest average probability to Job 3, the job with the highest earnings and proportion of men but with a high likelihood of being fired.

Another notable aspect of Table IV is the large standard deviation in elicited choice probabilities, reflective of substantial heterogeneity in choices, even within gender. Figure I shows the histogram of elicited percent chance responses for Job 1, pooled across the 16 hypothetical scenarios. Several things are notable. First, responses tend to be multiples of 10 or 5, a common feature of probabilistic belief data (Manski 2004), reflecting a likely rounding bias; this is something we return to below. Second, although there is pooling at multiples of 5, there is little evidence of excessive heaping at the standard focal responses of 0, 50, and 100. The most prevalent response is 20%, but even that receives a response frequency of only 0.11. Third, most respondents (87.5%) report values in the interior (that is, not 0 or 100), reflecting a belief that there is some chance they might choose each of the jobs. This underscores the importance of eliciting probabilistic data, rather than simply the most preferred option, as respondents are able to provide meaningful probabilistic preferences for the full set of choices.

---

implied by the model, where WTP is a function of the ratio of elements of the β_i vector, removing the dependence of WTP on the scale of the shock.

V.B. Empirical Model of Job Preferences

Next, we discuss our empirical model of job preferences, which we estimate using our hypothetical data. Our estimator follows the identification analysis we laid out above. For the job preferences over attributes, we use the form $u_i(X) = X'\beta_i$, where $\beta_i = [\beta_{i1}, \ldots, \beta_{iK}]$ is a K-dimensional vector that reflects individual i's preferences for each of the K job characteristics. The X vector of job characteristics is described below and we consider several different functional forms. We assume beliefs about future job utility $H_i(\cdot)$ in equation (4) are i.i.d. Type I extreme value for all individuals. The probability of choosing each job is then:

$$p_{ij} = \frac{\exp(X_j'\beta_i)}{\sum_{j'=1}^{J} \exp(X_{j'}'\beta_i)}, \quad (5)$$

where it is important to note that the probabilities assigned to each job j are individual i specific.[20] Although we maintain a particular assumption about the distribution of probabilistic beliefs, we place no parametric restrictions on the distribution of preferences, represented by the vector β_i. Our goal is to estimate the population distribution of preferences β_i. We maintain a maximum degree of flexibility by estimating the preference vector β_i separately for each sample member, and do not impose any "global" distributional assumptions about the population distribution of preferences (e.g., that preferences $\beta_i \sim N(\mu, \Sigma)$).

Applying the log-odds transformation to equation (5) yields the linear model:

$$\ln\left(\frac{p_{ij}}{p_{ij'}}\right) = (X_j - X_{j'})'\beta_i .$$

β_i has the interpretation of the marginal change in the log odds for some level difference in the X characteristics of the job. Given the difficulty of interpreting the β_i preference parameters directly, we also present results in which we compute individual-level WTP statistics.

V.C. Measurement Error

One potential issue in using hypothetical data for estimating preferences is that individuals may report their preferences with error. Given that these preferences have no objective counterpart (we cannot ascertain the "accuracy" of a self-reported preference), we cannot point to definitive evidence on the extent of measurement error. The most apparent potential measurement issue is that individuals report rounded versions of their underlying preferences (rounded to units of 5% or 10%). To guard against the potential of rounding bias or other sources of measurement error, we follow Blass, Lach, and Manski (2010) in

[20] Note that utilities across alternatives are correlated through the shared job attributes, therefore the independence of irrelevant alternatives problem does not apply to our model.

introducing measurement error to the model, and in flexibly estimating the model using a least absolute deviations (LAD) estimator.

We assume that the actual reports of job choice probabilities in our data, denoted $\widetilde{p}_{ij}$, measure the "true" probabilities p_{ij} with error. The measurement error takes a linear-in-logs form such that the reported log-odds take the following form:

$$\ln\left(\frac{\widetilde{p}_{ij}}{\widetilde{p}_{ij'}}\right) = (X_j - X_{j'})\beta_i + \omega_{ij}, \quad (6)$$

where ω_{ij} is the measurement error. We assume that the $\omega_{i1}, \ldots, \omega_{iJ}$ have median 0, conditional on the $X_1, \ldots, X_J$ observed job characteristics. Given these measurement error assumptions, we have the following median restriction:

$$M\left[\ln\left(\frac{\widetilde{p}_{ij}}{\widetilde{p}_{ij'}}\right) \mid X_j, X_{j'}\right] = (X_j - X_{j'})\beta_i, \quad (7)$$

where $M[\cdot]$ is the median operator. This median restriction forms the basis for our estimator. Our measurement error assumptions are limited compared to commonly imposed fully parametric models which assume a full distribution for the measurement error process. In contrast, our assumption is that the measurement errors are only median unbiased.[21] Another advantage of the LAD estimator is that it is not sensitive to what the extreme responses (probabilities of 0 and 1) are replaced with.

V.D. Estimation

We estimate the K-dimensional vector β_i by LAD for each student i separately. In our data, each student makes choices across 16 scenarios, assigning probabilities to three possible jobs in each scenario. Equation (7) therefore is estimated for each respondent using 16 × 2 = 32 unique observations. Variation in the job attributes (X_j), which is manipulated exogenously by us, and variation in respondents' choice probabilities allows us to identify the parameter vector β_i. From the full set of estimates of $\beta_1, \ldots, \beta_N$ for our size N sample we estimate population statistics, such as mean preferences, $E(\beta_i)$. We conduct inference on the population statistics using block or cluster bootstrap by resampling (with replacement) the entire set of job-hypothetical probabilities for each student. Online Appendix Section B describes the bootstrapping algorithm. The block bootstrap preserves the dependence structure within each respondent's block of responses, and allows for within-individual correlation across job-choice scenarios.

21. Note we do not impose that ω_{ij} measurement errors are independent across individuals or jobs and do not assume any particular joint distribution for the measurement errors, beyond the conditional median independence with the X variables. For inference, we use a cluster bootstrap method, resampling the entire set of job scenarios for each sample member, to preserve any correlation in residual errors. See Online Appendix B for details.

As discussed in the study design section, we varied four job attributes at a time in each scenario. For estimation, we combine all of these scenarios and assume the dimensions that were not varied in a given scenario were believed by the respondent to be held constant, as we instructed. As mentioned earlier, we instruct respondents that the jobs differ only in the finite number of job characteristics we provide, and are otherwise identical. There is no additional information here that the respondent could use to believe otherwise. The vector of job attributes is as follows: X = {log age-30 earnings; probability of being fired; bonus as a proportion of earnings; proportion of males in similar positions; annual increase in earnings; hours per week of work; availability of part-time}.[22] We also include job-number dummies in equation (7) to allow for the possibility that the ordering of the jobs presented could affect job preferences, although there is no prior reason to suspect this given our experimental design.[23]

V.E. Job Preference Estimates

We first discuss the sign and statistical significance level of the β_i estimates. Because of the difficulty in interpreting the magnitude of these estimates, below we also present results in which we convert the parameter estimates into an individual-level WTP measure. Recall that we can identify the β_i vector without a parametric restriction on the population distribution of preferences. Online Appendix C discusses the estimated heterogeneity in preferences within gender.

The first column of Table V shows the average estimate for each job characteristic (across all individual-level estimates). The standard errors in parentheses are derived from a block bootstrap procedure. We see that the average estimates have the expected signs: estimates for the probability of being fired and work hours per week are negative, while the others are positive. The estimates indicate that individuals, on average, prefer higher salaries and work-time flexibility, and dislike jobs with a high probability of being fired and high numbers of work hours. The only estimate that is not statistically or economically significant is the proportion of males at the job, indicating that we cannot reject that, on average, individuals are indifferent to the gender composition of the workplace. Turning to the average estimates by gender, reported in columns (2) and (3) of Table V, we see similar qualitative patterns. We return to the differences in magnitudes of the preferences by gender below, and also provide a WTP interpretation.

V.F. Willingness to Pay

The parameter estimates in Table V are difficult to interpret given the necessarily nonlinear nature of the model. To ease interpretation, we next present WTP estimates, which translate the differences of utility levels into earnings that would make the student indifferent between giving up earnings and experiencing the outcome considered.

1. Computing WTP.

1. Computing WTP.—WTP to experience job attribute X_k is constructed as follows. Consider a change in the level of attribute X_k from value $X_k = x_k$ to $X_k = x_k + \Delta$, with $\Delta > 0$.

22. We also estimate the model with utility specified as linear in earnings (instead of log earnings). Results are qualitatively similar. Online Appendix D discusses results from several other alternative specifications.

23. This is related to the possibility of "session effects" in laboratory experiments. See Frechette (2012).

Assume X_k is a "bad" attribute. Given our linear utility function, we can write an indifference condition in terms of earnings Y as:

$$x_k\beta_{ik} + \beta_{i1}\ln(Y) = \beta_{ik}(x_k + \Delta) + \beta_{i1}\ln(Y + \mathrm{WTP}_{ik}(\Delta)),$$

where Y is the level of earnings, one of the job attributes included in every job scenario. WTP $_{ik}(\Delta) > 0$ is individual i's willingness to pay to avoid increasing the "bad" attribute k by Δ. Solving, WTP is given by:

$$\mathrm{WTP}_{ik}(\Delta) = \left[\exp\left(\frac{-\beta_{ik}}{\beta_{i1}}\Delta\right) - 1\right] \times Y. \quad (8)$$

WTP for individual i depends on her preference for the attribute β_{ik} versus her preference for earnings β_{i1} (earnings is attribute 1). Given that we allow for a log form to utility in earnings (allowing for diminishing marginal utility in earnings and implicitly consumption), willingness to pay for an individual also depends on the level of earnings at the job.

2. WTP by Gender.—Table VI shows the average and median WTP estimates for changing each of the job characteristics by one unit (for the probabilistic outcomes, this is increasing the likelihood by 1 percentage point; for hours per week, increasing it by an hour; for part-time availability, this is going from a job with no part-time option to one which does).[24] The first three columns of the table present the estimates in dollars, evaluating WTP at the average annual earnings across all scenarios, $75,854 (which is fixed by the experimental setup and does not vary across respondents). The last three columns show the estimates as a proportion of the average earnings. We focus on the latter here.

We estimate, for example, that increasing the likelihood of being fired by 1 percentage point, that is, $X_k = x_k$ 1, would yield an average WTP of 2.8% for the full sample. That is, for students to remain indifferent to moving to a less stable job, they would on average have to be compensated by 2.8% of annual earnings. The gender-specific averages, reported in the last two columns of Table VI, indicate distinct average preferences by gender. Women, on average, have to be compensated by 4% of average earnings for a unit increase in the likelihood of being fired, with the estimate being statistically significant at the 1% level, and statistically different from the much smaller male average of 0.6%. Recall that we fix average earnings at the same level for all respondents, so the gender differences in WTP reflect only differences in preferences, not earnings. The median estimates also differ by gender, with women exhibiting a higher WTP for job stability. The median estimate for women is, however, lower than the average estimate, suggestive of a skewed distribution.

The average and median WTP estimate for the availability of the part-time option is sizable. Individuals, on average, would have to be compensated by 5.1% of their annual salary (that is, they are willing to give up 5.1%) when going from a job with no part-time option to one

[24]. The WTP is computed for each individual, using the individual-specific β_i estimates. The table reports mean WTP across respondents, bootstrap standard errors in parentheses, and median WTP in square brackets.

that does have one. The estimate is driven by the female respondents in the sample, for whom the average WTP is −7.3%, versus −1.0% for males (with the male estimate not being statistically different from 0). The much higher average preference among women for the part-time option is statistically significantly different from 0 and statistically different from the male average, at the 5% level. The median estimate also differs by gender and is larger in magnitude for women.

Examining the WTP for other job characteristics, we see that the average WTP for annual earnings growth is statistically precise for men, who are willing to give up 3.4% of average annual earnings for a 1 percentage point increase in earnings growth; the female average coefficient is indistinguishable from 0 (although not statistically different from the male estimate). The median estimates for the two genders are similar. We see that women have a stronger distaste for the number of hours of work, with the average WTP indicating that they need to be compensated by 1.3% of annual earnings for an increase of one hour in the work week; the male estimate is not precise (but we cannot reject the two gender-specific averages being equal). Both genders are, on average, willing to give up 0.8–1.7% of annual earnings for a percentage point increase in bonus compensation (in addition to base salary).[25] Finally, the average WTP for proportion of men at jobs is economically and statistically insignificant.

Online Appendix C further analyzes the heterogeneity in preferences for workplace characteristics, and investigates how the WTPs are associated with various individual-level characteristics. In addition, note that the utility from jobs, specified in equation (1), is linear and separable in outcomes. Online Appendix D shows that our conclusions are robust to estimating variants of the baseline model.

It is important to note that the timing of our survey is quite important in interpreting the resulting preference estimates. In general, there is no reason to believe that the workplace preferences we elicit are intrinsic, and they may be particular to the age of our survey respondents. Our estimates should not be considered unbiased estimates for intrinsic preferences—preferences are likely not intrinsic at all—but instead unbiased estimates for preferences at the point in each student's life cycle at which we collect our data. Our preference estimates may also reflect past experiences with employment because in some cases, the respondents may have already secured postgraduation employment. Our methodology simply relies on students being able to consider their likelihood of accepting hypothetical job offers, which should be possible even if a student is already employed.

V.G. Estimated Preferences and Actual Workplace Characteristics

Do the pre–labor market preferences we estimate relate to the characteristics of jobs these students actually end up working in?[26] We are able to shed light on this issue through a follow-up survey of a subset of our respondents conducted in 2016, about four years after the original data collection and when respondents were on average aged 25. Of the 247

---

25. That the WTP for a percentage point increase in bonus is greater than 1 in magnitude for women is surprising because it implies that women are on average willing to give up more in base salary to gain a smaller increase in bonus compensation. This is driven by a few outliers. In fact, we cannot reject that the mean WTP for women is different from either −1 (that is, a one-to-one substitution between base pay and bonus pay), or from the mean of −0.8 for male respondents.

respondents who took the survey and answered the hypothetical questions, 112 had also participated in an earlier survey conducted by us in 2010 (data that we have analyzed in Wiswall and Zafar 2015a,b) and given consent for future surveys. In January 2016, we invited these 112 respondents to participate in a 15-minute online survey about their current labor market status. 71 of the eligible 112 respondents (~63%) completed the follow-up survey.[27]

The follow-up survey collected information about respondents' workplace characteristics (for those currently working). Of the 71 respondents, 59 were working (either full-time, part-time, or self-employed) at the time of the follow-up survey, with the remainder enrolled in school. Online Appendix Table A7 shows the earnings and various other workplace characteristics for the overall sample, as well as for male and female workers, separately. Earnings, conditional on working full-time, are higher for men (by nearly $70,000). Bonus, hours of work, likelihood of being fired, fraction of male employees, and typical annual growth in earnings are all higher for our male respondents (though not all of the differences are statistically significant). The last row of the table shows that women's workplaces are more likely to have a part-time or flexible work option.[28]

Are these systematic gender differences in actual workplace characteristics consistent with our estimates of job preferences elicited several years prior, before labor market entry? To investigate this, we regress characteristics of each respondent's current job onto our individual-specific estimate of their past WTP for that attribute. WTP is defined as the amount the individual needs to be compensated by for a unit change in a given characteristic, with a higher WTP reflecting a lower taste (or greater distaste) for that outcome. Therefore, we expect a negative relationship between WTP and the current job characteristic. Estimates are presented in Table VII. Directionally, all six estimates are negative, with three significant at the 5% level or better. A joint test that all coefficients are 0 can be rejected (the *p*-value of this joint test is .012).

To interpret the magnitude of the estimated coefficients in Table VII, we also report "effect sizes" in the table. The effect size provides the estimated change in the dependent variable (that is, the actual workplace attribute) for a one standard deviation change in the WTP for that workplace characteristic. For example, we see that a one standard deviation increase in

---

26. Although being able to document a systematic relationship can provide some credibility to our methodology, on the other hand, a failure to find a systematic relationship between the two would not necessarily invalidate our method because students' preferences for jobs may change over time, or labor market frictions may prevent workers from matching with jobs that they prefer. Answering this question most directly would require both revealed-choice data that are free of any confounds and stated-choice data, which are usually not available. However, the little evidence that exists shows a close correspondence between preferences recovered from the two approaches (see Hainmueller, Hangartner, and Yamamoto 2015).

27. Respondents were initially contacted through email addresses provided in our earlier data collections. Those with inactive email addresses were then approached through LinkedIn. Respondents received a link to the survey that was programmed in SurveyMonkey and were compensated for completing the survey. As shown in Online Appendix Table A6, there is little evidence of selection on observables (reported in 2012) in terms of who participates in the follow-up survey. Based on a joint *F*-test, we cannot reject that the covariates are jointly zero (*p*-value = .360) in predicting survey response. Note that for students to have taken both the 2010 and 2012 surveys, the sample from which we have consent would have had to be in the junior year or higher in 2012.

28. Online Appendix Table A8 shows that the sample for which we have consent, the sample that takes the follow-up survey, and the sample that was working when the follow-up survey was conducted are all very similar to the full sample. The only dimensions along which they differ are school year and age (which, as explained above, is by construction) and race. Importantly, there are no statistical differences along the dimensions of gender, major, ability, or socioeconomic background. Also note that the follow-up samples, in columns (3) and (4), are not statistically different from the consent sample along any dimension.

the WTP (that is, higher distaste) for work hours translates into an estimated decrease of 4.1 hours worked. Given that the standard deviation of hours worked is 14.8 in the sample, this is a sizable impact. Likewise, a one standard deviation increase in the WTP (that is, lower taste) for availability of flexible work options is associated with a 15 percentage point decline in the actual availability of these options in the respondent's workplace (on a base of 61). The effect sizes for bonus percentage and proportion of male are also economically meaningful.

While we have shown that estimated preferences for attributes are jointly systematically related to actual future workplace characteristics in the cross section, a natural question to ask is whether the relationship also holds within the individual, that is, whether a higher WTP for a given attribute translates into more of that attribute for an individual. For each attribute, we rank the 59 individuals in terms of both the estimated WTP and the actual value at the job. This gives us a six-dimensional vector of ranked WTPs and a six-dimensional vector of ranked attribute values for each individual. We then compute the individual-level correlation between the two vectors. We expect a negative correlation: higher WTP (that is, a lower taste or a greater distaste) for an attribute causes an individual to be working in a job with lower values of that attribute. That is exactly what we find: the mean correlation coefficient across the individuals is −0.158 (significant with a *p*-value = .017) and the median correlation coefficient is −0.250 (*p*-value = .36), indicative of a systematic relationship between estimated WTPs and actual attributes even within individuals.

Overall, these results strongly indicate that our estimated preferences capture true underlying heterogeneity that is also reflected in actual job outcomes several years later. We view these results as a joint validation of our methodology, data quality, and empirical specification. Our finding that estimated WTPs predict respondents' actual workplace choices is all the more remarkable given that the hypothetical scenarios were fielded to respondents when they were still in college (though some of the respondents may have already secured postgraduation employment at the time of the survey). In the next section, we investigate whether these workplace preferences impact major choice.

VI. Job Preferences and Major Choice

The preceding sections used a robust hypothetical choice methodology to estimate individual-level preferences for various job attributes. This section relates these preferences to human capital investments, quantifying the importance of job characteristics to college major choices.

First, to set the stage for this analysis, we describe the anticipated major choices reported by our sample. Given that our sample consisted of currently enrolled students, we asked the students to provide their beliefs they would complete a degree in one of the five major categories: "What do you believe is the percent chance (or chances out of 100) that you would either graduate from NYU with a PRIMARY major in the following major categories or that you would never graduate/dropout (i.e., you will never receive a bachelor's degree from NYU or any other university)?" The first column of Online Appendix Table A9 shows the response to the question: the most likely major for males is economics/business (43%), followed by humanities/social sciences (29%). For women, on the other hand, the most likely major is humanities/social sciences (53%), followed by economics (23%). The probability of not graduating is less than 3% for both sexes. The average probabilities assigned to the majors differ significantly by gender for all majors except engineering and natural sciences. Our model of major choice allows for some uncertainty in major choice: at least part of the sample is not 100% certain of their final major at graduation and the data reflect that (a majority of students, 53, do not assign a 100% probability to their most likely major). Our model of probabilistic major choices nests the standard model of deterministic major choice.

We next decompose the anticipated major choices into various factors, including potential job characteristics associated with each major. To gauge the importance of job attributes to major choice, we estimate a model of major choice incorporating our flexible estimates of preferences for job attributes and separate data we collected on students' beliefs about the likelihood they would be offered jobs with these characteristics, conditional on major choice (that is, estimates of students' perceptions of the firm or demand side of the labor market). We then use this estimated model to quantify the importance of each job attribute to major choice. Given that prior literature on educational choice finds that the residual unobserved "taste" component is the dominant factor in major choice (Arcidiacono 2004; Beffy, Fougere, and Maurel 2012; Gemici and Wiswall 2014; Wiswall and Zafar 2015a), our approach can be viewed as trying to get into the black box of tastes by directly incorporating certain nonpecuniary dimensions into these choice models. The estimation details for the major-choice model are provided in Online Appendix E. Here, for the sake of brevity, we comment on only its main features.

We start with a simple framework in which we suppose that utility for student i from major m is given by:

$$V_{im} = X'_{im}\alpha_i + Z'_{im}\gamma + \kappa_m + \eta_{im}, \quad (9)$$

where X_{im} is i's perceived job attributes in major m. With hopefully minimal confusion, we use the same notation X to refer to job attributes as in our hypothetical job choice analysis and to refer to perceptions about job attributes associated with each major, a separate set of variables collected in our survey. Note that here the X vector is indexed by i as these attributes are each student's perception of the job attributes (that are allowed to depend on the major m) rather than the exogenously determined attributes in the hypotheticals we created. Z_{im} is a vector of other major-specific characteristics perceived by student i (including major-specific perceptions of ability and perceived hours of study needed to obtain a GPA of 4.0 in that major). κ_m is a major-specific constant, capturing overall tastes for the major, and η_{im} captures the remaining unobservable attributes of each major.

To estimate the model, we use data on students' perceptions of the likelihood of being offered jobs with various characteristics conditional on each major, as well as their beliefs regarding major-specific ability. Our survey collected data from respondents on their perceptions of characteristics of the jobs that would likely be offered to them if they were to complete each type of major. An important characteristic of our data set is that we gather students' beliefs about workplace characteristics (such as likelihood of being fired and earnings) for a set of different majors, not just for the one major they intend to complete. These data are described at length in Online Appendix E.[29]

In equation (9), the student-specific preference for each job attribute is given by the vector $\alpha_i = [\alpha_{i1}, \dots, \alpha_{iK}]$. α_i, the preference or job characteristics as it relates to the utility from each major, is potentially distinct from the preferences for job characteristics in the job-choice problem, given by β_i (in equation (5)). Job characteristics, such as earnings at the job, may be quite important when choosing among different job offers but might have a more limited value to choosing majors, relative to other major characteristics given by Z_{im}, κ_m, and η_{im}. To allow for this possibility, for each job characteristic k, we specify that each α_{ik} is proportional to the β_{ik} up to some free (to be estimated) parameter δ: $\alpha_{ik} = \beta_{ik}\delta$. δ indicates the importance of job attributes to major choice, relative to other determinants of college major as given by Z_{im}, κ_m, and η_{im}. δ could also reflect standard discounting given that the utility from working at jobs occurs later in life than utility derived from taking courses while in school.

Table VIII presents the LAD estimates of equation (9) using the hypothetical data to estimate the job-preference vector β_i for each student, and a robust cluster bootstrap over all estimation steps for inference (see Online Appendix E for estimation details). The estimate of δ is positive and precise, indicating that the preferences of students over job attributes and the major-specific beliefs about the distribution of job attributes have a statistically significant relationship with major choices. Estimates on the major-specific ability measures are negative, as one would expect (note that higher "ability rank" denotes lower ability in our data). The major-specific dummy terms are all negative, indicative of negative median tastes for the nonhumanities majors (the omitted category): all else equal, students prefer to major in humanities.

Given the nonlinear nature of the model, it is difficult to assess the importance of job attributes in major choice from the estimated coefficients alone. To quantify the effects, we use standard methods to evaluate "marginal effects" in nonlinear models (see Online Appendix E for details). The marginal effect of a job attribute in major choice is computed for a standard deviation change in the value of that specific job attribute, while keeping the other (job- and major-specific) attributes and preferences fixed at their sample average values.

Table IX presents the marginal effects for specific changes in job attributes, averaged across the majors, and separately by sex (in the two panels of the table). The table also shows the start and end value for the attribute at which the marginal effect is computed. The start value is the sex-specific belief for that attribute (averaged across majors and respondents), and the

[29]. Because the vast majority of our sample is either in their junior or senior year, and some have already chosen a major, one concern is that the students' preferences and beliefs, as elicited in our survey data, may be different from the preferences and beliefs they held in the past as they were deciding on a college major. Although we can of course still estimate the relationship between major choice and the data we collect, the interpretation of our estimates in these cases is less clear. One solution is to collect longitudinal data on preferences and beliefs to directly examine the extent to which they change over the life cycle and how this influences college major choices. See Stinebrickner and Stinebrickner (2014b) for an important example.

end value is the start value shifted by one sex-specific standard deviation (again, across majors and respondents) in the beliefs for that attribute. Column (1), for example, shows that increasing the perceived probability of being fired from jobs by one standard deviation decreases the likelihood of majoring in that major, on average, by 4% for men and by 5% for women. A standard deviation increase in part-time availability increases the probability of completing a major by 0.2%. Column (3) shows that a standard deviation increase in weekly work hours reduces the likelihood of majoring in a major by 2.5% for men and 1.4% for women. Bonus pay and earnings growth both also have sizable average marginal effects. The last column of Table IX shows the percent change in the major probability for a standard deviation increase in log age-30 earnings.

A comparison of the effects in the first three columns with those in the last column for earnings (also shown in the last row of each panel) gives a sense of the relative importance of these other job attributes in major choice. We see that, for women, the average effect for the probability of being fired is as large as that for earnings, and for hours is nearly a third of the effect of earnings. For men, the relative impacts are smaller (though still sizable). Overall, this indicates that job attributes matter for major choice and that they are particularly relevant for women's choices.

VII. Job Preferences and the Gender Gap in Earnings

In the previous sections, we have shown systematic gender differences in workplace preferences and quantified the importance of these preferences to major choices. In this section, we explore the extent to which gendered job preferences explain the "gender gap" in earnings. Differences in job preferences can give rise to differences in earnings through two channels. First, as explored above, job preferences can affect college major choices and, given the wide dispersion in earnings across fields, affect the overall distribution of earnings for men and women. Second, even conditional on major choice, gender differences in workplace preferences can affect the distribution of earnings. The gender gap in earnings we observe could be at least partially the result of women "purchasing" certain positive job attributes by accepting lower wages, or conversely, men accepting higher earnings to compensate for negative job attributes.

To quantify the first channel (job preferences affecting earnings through major choice), we conduct the following exercise. Using the estimated major-choice model in Section VI, we predict the likelihood of women choosing different majors if their workplace preferences were shifted by the average male minus female mean (that is, we preserve the heterogeneity in women's preferences but shift them by the average gender difference in the preferences).[30] We then predict the likelihood of each female respondent choosing the different majors and use these to weight the individual's major-specific expected earnings. This provides the impact on the gender wage gap if women's preference distribution was shifted to have the male average, but only through the major choice channel. Note that, for this exercise, we keep women's earnings expectations fixed in that major (which could also be affected by

30. Because the estimated preference parameters are not scale-free, this exercise of shifting the preference parameter by some amount implicitly assumes that the variance of the unobserved factors is the same for all individuals (Train 2003).

workplace preferences). In this exercise, we find that the change in women's major choices lowers the expected gender gap in age-30 earnings by 2.6%.[31] Given our highly aggregated major categories, this is likely a lower bound on the importance of preferences to the gender earnings gap through major choices, and human capital more generally. Previous work has emphasized that important job segregation by gender occurs through choices of subfields (see for example, Goldin and Katz 2016, on choice of medical specialties).

Turning to the second channel, we consider the following simple exercise. We ask how the gender gap in expected earnings changes once we "control" for individual-specific workplace preferences (the estimated preference parameters in Section V). If the gender gap in earnings is solely because women are accepting lower wages for desirable jobs, and/or men are compensated with higher wages for undesirable jobs, then men and women with identical workplace preferences would have equal earnings. If, on the other hand, a gender gap remains, even after conditioning on preferences, then we can conclude that demand-side factors, such as employment discrimination, still play a role in the gender gap.

We implement this exercise using a simple set of regressions in Table X, Panel A. Column (1) of the table reports a regression of an individual's log expected earnings for the major they are most likely to graduate with onto a female dummy. We see a gender gap of about 35 log points in age-30 expected earnings, a gap similar to that in realized earnings data.[32] The second column shows that the gender gap declines to about 20 log points once the individual's major is controlled for, reflecting the fact that women are less likely to graduate in higher-earnings majors. Columns (3) and (4) show how the gender gap changes once we control for the estimated vector of workplace preferences. Importantly, a comparison of column (4) with column (2) shows that, even conditional on major choice, workplace preferences reduce the expected earnings gender gap by about a quarter, from about 20% to 15%. Note that workplace preferences are also likely to impact major choice, which is held fixed here.

Table X, Panel B repeats the exercise using actual earnings reported by the follow-up respondents. The sample here is smaller, but the qualitative results are strikingly similar to those that we observe for expected earnings: conditional on major, the gender gap in realized earnings declines from 45 log points to 32 log points (that is, by nearly 30%) once we control for respondents' workplace preferences.

We conclude from this analysis that gender differences in workplace preferences can explain a sizable part of the gender gap in expected earnings early in the life cycle. And, albeit with a smaller sample, our evidence points to similar conclusions for realized earnings as well. We also find that the main channel by which workplace preferences affect the gender earnings gap is through job choices, not through major choices, at least at the aggregated major level we have available in this data set.

31. More specifically, the gender gap declines by about 0.9 percentage points from a baseline predicted gender gap of 35.1%. This is primarily a result of women's predicted probability of majoring in humanities declining from 55.0% to 53.8%, and their predicted probability of majoring in economics increasing from 18.3% to 19.5%.

32. As described in Section II, in the sample of all college graduates ages 25–40 in the ACS, the mean earnings for full-time employed men is 36% higher than the mean earnings of full-time employed women.

VIII. Conclusion

Economists have long recognized that job and occupational choices are not solely determined by expected earnings.[33] Although simple models based on earnings maximization abound (see, for example, the classic Roy 1951 model) and are quite useful in some applications, it is also clear that individuals have a rich set of preferences for various aspects of jobs beyond expected earnings, including earnings and dismissal risk, and various non-pecuniary aspects such as work hours flexibility. Human capital investments too could be affected by these workplace preferences as individuals alter their human capital investment in anticipation of particular future job choices. Key features of the distribution of labor earnings in the economy, such as the gap in earnings between men and women, need careful consideration, as differences in earnings may reflect, at least in part, heterogeneity in preferences and compensating differentials for various nonpecuniary attributes of employment.

Using a novel hypothetical job-choice framework that experimentally varies different dimensions of the workplace, this article robustly estimates individual preferences for workplace attributes. For a sample of high-ability undergraduate students enrolled at a selective private U.S. university, we document substantial heterogeneity in willingness to pay for job amenities, with large differences in the distribution of preferences between men and women. For a subset of the sample for whom we collect data on actual workplace characteristics (nearly four years after the original survey), we find a robust systematic relationship between estimated preferences and the characteristics of their current jobs. The predictive power of the estimated preferences at the individual level strengthens the credibility of our approach, and makes a case for employing this methodology in other settings to understand decision making.

Combining these workplace preferences with unique data on the students' perceptions of jobs which would be offered to them given their major choice, we quantify the role of anticipated future job characteristics—particularly the nonpecuniary aspects of these jobs—in choice of major, a key human capital investment decision. Women, in particular, are found to be more sensitive to nonpecuniary job aspects in major choice than men. Our analysis indicates that at least a quarter of the gender gap in early career earnings—expected as well as actual—can be explained by the systematic gender differences in workplace characteristics. Our analysis indicates that a substantial part of the early gender gap in earnings we observe is a compensating differential in which women are willing to give up higher earnings to obtain other job attributes.

There are several potential areas for future research. Although we find substantial variation in workplace preferences for our sample of high-ability students at a selective university, it is not clear how these preferences compare to that of the broader population. It would clearly be useful to follow our design and collect similar data in other settings. In particular, preference data collected at older ages would be useful in studying how preferences for

33. See the famous quote by Adam Smith who lists a number of nonpecuniary job attributes which "make up for a small pecuniary gain in some employments, and counterbalance a great one in others" (*Wealth of Nations*, 1776, Book 1, Chapter 10).

nonpecuniary dimensions of the workplace, especially those related to accommodations for raising children, evolve over the life cycle (Bertrand, Goldin, and Katz 2010). Our work also does not directly indicate the sources of the systematic gender differences in workplace preferences that we document. For example, they may be a consequence of social factors including anticipated discrimination (Altonji and Blank 1999). We cannot therefore claim that these preferences are intrinsic and immutable in the sense that they may be due, at least in part, to environmental influences particular to this cohort of students. Research that sheds light on the underlying channels would be immensely valuable.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

REFERENCES

Altonji Joseph, and Rebecca Blank, “Race and Gender in the Labor Market,” in Handbook of Labor Economics, Vol. 3c, Orley Ashenfelter and David Card, eds. (Amsterdam: Elsevier Science, 1999), 3144–3259.

Altonji Joseph, Lisa Kahn, and Jamin Speer, “Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success,” Journal of Labor Economics, 34 (2016), 361–401.

Arcidiacono Peter, “Ability Sorting and the Returns to College Major,” Journal of Econometrics, 121 (2004), 343–375.

Arcidiacono Peter, Joseph Hotz, and Songman Kang, “Modeling College Major Choices using Elicited Measures of Expectations and Counterfactuals,” Journal of Econometrics, 166 (2012), 3–16.

Arcidiacono Peter, Joseph Hotz, Arnaud Maurel, and Teresa Romano, “Recovering Ex Ante Returns and Preferences for Occupations Using Subjective Expectations Data,” NBER Working Paper no. 20626, 2015.

Becker Gary, The Economics of Discrimination (Chicago: University of Chicago Press, 1971).

Beffy Magali, Denis Fougere, and Arnaud Maurel, “Choosing the Field of Study in Post-Secondary Education: Do Expected Earnings Matter?,” Review of Economics and Statistics, 94 (2012), 334–347.

Bertrand Marianne, Claudia Goldin, and Lawrence Katz, “Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors,” American Economic Journal: Applied Economics, 2 (2010), 228–255.

Blass Asher, Saul Lach, and Charles Manski, “Using Elicited Choice Probabilities to Estimate Random Utility Models: Preferences for Electricity Reliability,” International Economic Review, 51 (2010), 421–440.

Blau Francine, and Lawrence Kahn, “The Gender-Wage Gap: Extent, Trends, and Explanations,” Journal of Economic Literature, forthcoming.

Bonhomme Stephen, and Grégory Jolivet, “The Pervasive Absence of Compensating Differentials,” Journal of Applied Econometrics, 24(5) (2009), 763–795.

Bronson Mary Ann, “Degrees Are Forever:Marriage, Educational Investment, and Lifecycle Labor Decisions of Men and Women,” Working paper, Georgetown University, Department of Economics, 2015.

Cortes Patricia, and Jessica Pan, “Prevalence of Long Hours and Women’s Job Choices: Evidence across Countries and within the U.S.,” Working paper, 2016.

Croson Rachel, and Uri Gneezy, “Gender Differences in Preferences,” Journal of Economic Literature, 47 (2009), 448–74.

d’Haultfoeuille Xavier, and Arnaud Maurel, “Inference on an Extended Roy Model, with an Application to Schooling Decisions in France,” Journal of Econometrics, 174 (2013), 95–106.

Dillon Eleanor, “Risk and Return Tradeoffs in Lifetime Earnings,” Journal of Labor Economics, forthcoming.

Flabbi Luca, and Andrea Moro, “The Effect of Job Flexibility on Female Labor Market Outcomes: Estimates from a Search and Bargaining Model,” Journal of Econometrics, 168 (2012), 81–95.

Flory Jeffrey, Andreas Leibbrandt, and John A. List, “Do CompetitiveWorkplaces Deter FemaleWorkers? A Large-Scale Natural Field Experiment on Job Entry Decisions,” Review of Economic Studies, 82 (2015), 122–155.

Fox Jeremy, Kyoo Kim, Stephen Ryan, and Patrick Bajari, “The Random Coefficients Logit Model Is Identified,” Journal of Econometrics, 166 (2012), 204–212.

Fréchette Guillaume, “Session-Effects in the Laboratory,” Experimental Economics, 15 (2012), 485–498.

Gemici Ahu, and Matthew Wiswall, “Evolution of Gender Differences in Post-Secondary Human Capital Investments: College Majors at the Intensive Margin,” International Economic Review, 55 (2014), 23–56.

Goldin Claudia, “A Grand Gender Convergence: It’s Last Chapter,” American Economic Review, 104 (2014), 1091–1119.

Goldin Claudia, and Lawrence Katz, “The Cost of Workplace Flexibility for High-Powered Professionals,” Annals of the American Academy of Political and Social Science, 638 (2011), 45–67.

——, “A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation,” Journal of Labor Economics, 34 (2016), 705–746.

Hainmueller Jens, Dominik Hangartner, and Teppei Yamamoto, “Validating Vignette and Conjoint Survey Experiments Against Real-World Behavior,” Proceedings of the National Academy of Sciences, 112 (2015), 2395–2400.

Hwang Hae-shin, Dale Mortensen, and Reed W. Robert, “Hedonic Wages and Labor Market Search,” Journal of Labor Economics, 16 (1998), 815–847.

Lim Katherine, “Self-Employment, Workplace Flexibility, and Maternal Labor Supply: A Life-Cycle Model,” Working paper, Office of Tax Analysis, US Department of the Treasury, 2015.

Lordan Grace, and Pischke Jorn-Steffen, “Does Rosie Like Riveting? Male and Female Occupational Choices,” NBER Working paper 22495, 2016.

Maestas Nicole, Mullen Kathleen J., David Powell, Jeffrey Wenger, and von Wachter Till, “2015 American Working Conditions Survey: First Findings,” University of Michigan Retirement Research Center (MRRC) Working paper 2016–342, 2016.

Manski Charles, “Measuring Expectations,” Econometrica, 72 (2004), 1329–1376.

Mas Alexandre, and Amanda Pallais, “Valuing Alternative Work Arrangements,” American Economic Review, forthcoming.

McFadden Daniel, and Kenneth Train, “Mixed MNL Models for Discrete Response,” Journal of Applied Econometrics, 15 (2000), 447–470.

Neumark David, Ian Burn, and Patrick Button, “Is It Harder for Older Workers to Find Jobs? New and Improved Evidence from a Field Experiment,” NBER Working paper 21669, 2015.

Niederle Muriel, and Lise Vesterlund, “Do Women Shy away from Competition? Do Men Compete too Much?,” Quarterly Journal of Economics, 122 (2007), 1067–1101.

Reuben Ernesto, Matthew Wiswall, and Basit Zafar, “Preferences and Biases in Educational Choices and Labor Market Expectations: Shrinking the Black Box of Gender,” Economic Journal, forthcoming.

Rosen Sherwin, “The Theory of Equalizing Differences,” in Handbook of Labor Economics, vol. 1, Ashenfelter O and Layard R, eds. (Elsevier, 1987), 641–692.

Roy AD, “Some Thoughts on the Distribution of Earnings,” Oxford Economic Papers, New Series, 3 (1951) 135–146.

Stern Scott, “Do Scientists Pay to Be Scientists?,” Management Science, 50 (2004), 835–853.

Stinebrickner Ralph, and Todd Stinebrickner, “A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout,” Review of Economic Studies, 81 (2014a), 426–472.

Stinebrickner Todd R., and Ralph Stinebrickner, “Academic Performance and College Dropout: Using Longitudinal Expectations Data to Estimate a Learning Model,” Journal of Labor Economics, 32 (2014b), 601–644.

Train Kenneth, Discrete Choice Methods with Simulation (Cambridge: Cambridge University Press, 2003).

Wasserman Melanie, “Hours Constraints, Occupational Choice and Fertility: Evidence from Medical Residents,” Working paper, UCLA Anderson School of Management, 2015.

Wiswall Matthew, and Basit Zafar, “Determinants of College Major Choice: Identification Using an Information Experiment,” Review of Economic Studies, 82 (2015a), 791–824.

——, “How Do College Students Respond to Public Information about Earnings?,” Journal of Human Capital, 9 (2015b), 117–169.

——, “Replication Data for: ‘Preference for the Workplace, Investment in Human Capital, and Gender’,” Harvard Dataverse (2017), doi:10.7910/DVN/MLOGDL

Zafar Basit, “College Major Choice and the Gender Gap,” Journal of Human Resources, 48 (2013), 545–595.

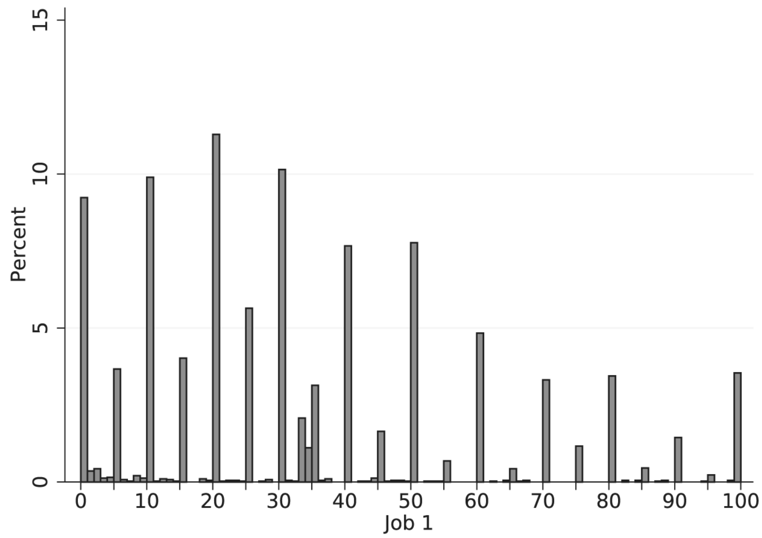

Figure I.
Choice Probabilities for Job 1 (Pooled across Hypothetical Scenarios)

TABLE I

Job Attributes by Broad Sector for College Graduates

	% of males working in [a] (1)	% of females working in (2)	Annual earnings for full-time (3)	Hrs/wk for full-time (4)	Prop. of part-time workers (5)	Yearly firing rate[b] (6)	Prop. male workers[c] (7)	Annual % raise in earnings[d] (8)
Sectors								
Science	9.6	4.0	82,739 (35,989)	44.2 (7.2)	15.5 (3.0)	3.5 (1.6)	67.7 (1.1)	3.8 (19.9)
Health	8.6	22.5	65,427 (35,246)	43.6 (7.8)	28.6 (1.1)	4.0 (0.7)	20.6 (0.5)	4.2 (22.9)
Business	14.2	11.5	77,079 (39,023)	45.0 (7.8)	19.9 (1.8)	4.0 (1.3)	44.6 (0.8)	4.9 (21.4)
Government	6.8	5.8	67,603 (32,322)	43.3 (6.9)	16.2 (5.0)	1.4 (0.9)	52.9 (0.6)	5.8 (22.4)
Education	11.2	25.5	60,588 (29,159)	44.0 (7.5)	30.0 (2.9)	1.8 (1.2)	30.9 (0.4)	4.2 (21.1)
Manufacturing & agriculture	22.4	7.99	77,354 (37,257)	45.4 (7.93)	17.6 (1.6)	6.2 (0.8)	78.0 (0.4)	5.7 (22.4)
Services & trade	27.2	22.8	65,734 (37,883)	45.3 (8.3)	34.4 (1.0)	6.6 (0.9)	51.7 (0.3)	4.9 (22.2)
p-value[e]	.000	.000	.000	.000	.000	.000	.000	0.170

Notes. Table reports means, with standard deviations in parentheses. Statistics are based on the 2010-2012 CPS monthly data. Sample restricted to those with at least a bachelor's degree, between ages 25 and 60. See Online Appendix A for details on construction of variables and definition of the broad sectors. Variables in columns (3), (4), and (8) are based on full-time workers, and are based on individual-level data. Columns (5)-(7) show the average statistics by sector, with the sector-level standard deviation across the months in parentheses.

[a] Proportion of all male workers who are employed in each sector (column sums to 100).

[b] Derived from the monthly firing rate, which is the ratio of workers who are laid off in a given month and have been unemployed for less than one month divided by all employed workers at the beginning of the previous month.

[c] Males as a proportion of all workers in that sector.

[d] Constructed by using the outgoing rotation groups, from the reported earnings in the respondent's fourth and eighth interview (which are separated by 12 months).

[e] *F*-test of equality of means/proportions across the industry categories.

Author Manuscript Author Manuscript Author Manuscript

TABLE II

Job Attributes by College Major for Young College Graduates, 25 to 40 Years Old

	Shares		Annual earnings ($) for full-time	Hrs/wk for fall-time	% Part-time workers[b]	UE rate[c]	Ann % salary raise[d]
	Males[a] (1)	Females (2)	(3)	(4)	(5)	(6)	(7)
Bachelor's (or more) in: Business	24.6	18.8	77,002 (68,110)	45.3 (8.1)	26.8	3.3	4.4
Engineering	23.3	6.1	86,679 (60,494)	44.8 (8.1)	22.2	2.5	4.8
Humanities	38.2	55.8	59,328 (49,697)	44.4 (7.9)	37.8	3.5	4.9
Natural science	13.9	19.3	75,992 (65,921)	44.9 (9.6)	35.1	2.5	5.9
F-test[e]	.000	.000	.000	.000	.000	.000	.000

Notes. Table shows statistics from the 2013 American Community Survey (ACS), restricting the sample to 25-40-year-olds with at least a bachelor's degree. Sample size is 204,190 respondents. 173 majors are grouped into four broad categories. Means (std. dev.) shown for annual earnings and hrs/week for full-time workers.

[a]Proportion of all 25-40-year-old college-educated males with the specified broad major (column sums to 100).

[b]Proportion of part-time workers from pool of those currently employed.

[c]Unemployment rate is number of individuals not employed and currently looking for a job, divided by sum of unemployed and employed respondents.

[d]Calculated by linearly regressing log earnings for a given major group on age (coefficient on age reported).

[e]*p*-value of *F*-test of equality of means across majors (rows).

TABLE III

Sample Statistics

	All (1)	Males (2)	Females (3)	*p*-value (4)
Number of respondents	247	86	161	
School year:				
Freshmen	10.9%	9.3%	11.8%	.549
Sophomore	10.9%	11.6%	10.6%	.798
Junior	36.4%	32.6%	38.5%	.355
Senior or more	41.7%	46.5%	39.1%	.262
Age	21.49 (1.5)	21.69 (1.8)	21.37 (1.2)	.103
Race:				
White	29.2%	33.7%	26.7%	.248
Asian	50.6%	51.1%	50.3%	.898
Non-Asian minority	17.8%	14.0%	19.9%	.247
Parent's characteristics:				
Parents' income ($1,000s)	137 (121)	141 (126)	135 (118)	.731
Mother B.A. or more	67.6%	74.4%	64.0%	.095
Father B.A. or more	69.6%	72.1%	68.3%	.539
Ability measures:				
SAT math score	696.0 (88)	717.7 (72)	684.3 (94)	.006
SAT verbal score	674.0 (84)	677.0 (78)	672.5 (88)	.704
GPA	3.5 (0.32)	3.5 (0.33)	3.5 (0.32)	.938
Intended/current major				
Economics/business	31.2%	48.8%	21.7%	.000
Engineering	4.9%	8.1%	3.1%	.080
Humanities and soc sciences	47.8%	30.2%	57.1%	.000
Natural sciences/math	16.2%	12.8%	18.0%	.289

Notes. For the continuous outcomes, means are reported in the first cell, and standard deviations are reported in parentheses. *p*-value reported for a pairwise test of equality of means (proportions) between males and females, based on a *t*-test (chi-square test).

TABLE IV

Example Choice Scenarios

Panel A	Earnings per year at age 30 if working full time	Annual percentage increase in earnings from age 30 on	Average work hours per week for full-time	Work flexibility: part-time work available?	Probability assigned by: Males	Probability assigned by: Females
Example 1						
Job 1	$96,000	3	52	Yes	31.93 [30] (22.48)	31.46 [30] (21.36)
Job 2	$95,000	2	45	Yes	31.16 [30] (23.71)	39.34^{***} [40] (22.71)
Job 3	$89,000	4	42	No	36.91 [30] (24.71)	29.20^{**} [25] (22.57)
Example 2						
Job 1	$76,000	4	50	Yes	19.38 [20] (19.34)	20.65 [20] (15.23)
Job 2	$81,000	3	44	Yes	49.47 [50] (26.63)	49.45 [50] (22.08)
Job 3	$88,000	2	49	No	31.15 [25] (25.36)	29.91 [25] (21.98)
Example 1						
Job 1	$87,000	1	$4,350 (5)	49	30.34 [30] (22.48)	36.68^{*} [30] (24.33)
Job 2	$84,000	6	$10,920 (13)	67	26.86 [30] (23.71)	30.27 [30] (21.36)
Job 3	$95,000	5	$4,750 (5)	69	42.80 [31.5] (24.71)	33.05^{***} [30] (20.83)
Example 2						
Job 1	$61,000	1	$6,710 (11)	41	25.48 [20] (26.57)	26.80 [20] (23.20)
Job 2	$65,000	5	$7,800 (12)	71	12.14 [9.5] (12.98)	15.53^{**} [10] (11.81)
Job 3	$67,000	2	$10,050 (15)	60	62.38 [60] (31.55)	57.67 [60] (27.19)

Notes. Means [median] (std. dev.) reported in the last two columns. Pairwise *t*-tests conducted for equality of means by gender. Significance denoted on the female column by asterisks:

$^{*}p<.10$,

** $p < .05$,

*** $p < .01$.

TABLE V

Estimates of Job Choice Model

	Overall (1)	Males (2)	Females (3)
Age-30 log earnings	15.40*** (1.65)	22.86*** (3.88)	11.42*** (1.43)
Probability of being fired	−0.38*** (0.04)	−0.39*** (0.10)	−0.37*** (0.04)
Bonus, as a prop. of earnings	0.28*** (0.03)	0.38*** (0.05)	0.22*** (0.03)
Prop. of males in similar positions	0.00 (0.00)	−0.01 (0.01)	0.005 (0.01)
% increase in annual earnings	0.55*** (0.10)	1.09*** (0.22)	0.27** (0.10)
Hours per week of work	−0.15*** (0.02)	−0.21*** (0.05)	−0.12*** (0.02)
Part-time option available	0.79*** (0.11)	0.86*** (0.22)	0.76*** (0.12)
observations	247	86	161

Notes. Table reports the average of the parameter estimates across the relevant sample. Asterisks denote estimates are statistically different from zero based on bootstrap standard errors. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Author Manuscript

Author Manuscript

Author Manuscript

TABLE VI

Willingness-to-Pay (WTP) Estimates

	WTP ($)			WTP (as % of average earnings)		
	Overall (1)	**Male (2)**	**Female (3)**	**Overall (4)**	**Male (5)**	**Female (6)**
Percent chance of being fired	2,147.40*** (525.46) [1,125.94]***	467.79 (670.29) [504.69]**	3,044.58***+++ (715.28) [1,841.27]***+++	2.83%*** (0.69%) [1.48%]***	0.62% (0.88%) [0.67%]**	4.01%***+++ (0.94%) [2.43%]***+++
Bonus as % of earnings	−1,069.78*** (258.47) [−975.36]***	−645.92* (368.77) [−761.57]***	−1,296.19*** (345.77) [−1,139.44]***++	−1.41%*** (0.34%) [−1.29%]***	−0.85%* (0.49%) [−1.00%]***	−1.71%*** (0.46%) [−1.50%]***++
Percent of men at jobs	43.20 (38.31) [59.15]***	63.74 (46.81) [71.82]**	32.24 (53.93) [43.17]	0.06% (0.05%) [0.08%]***	0.08% (0.06%) [0.09%]**	0.04% (0.07%) [0.06%]
Annual % raise in earnings	−1,186.28 (773.44) [−2,596.68]***	−2,564.93** (1,226.42) [−2,934.51]***	−449.86 (957.85) [−2,514.55]***	−1.56% (1.02%) [−3.42%]***	−3.38%** (1.62%) [−3.87%]***	−0.59% (1.26%) [−3.31%]***
Hours per week of work	854.65*** (235.25) [626.52]***	594.70 (416.17) [600.89]***	993.50*** (267.19) [634.22]***	1.13%*** (0.31%) [0.83%]***	0.78% (0.55%) [0.79%]***	1.31%*** (0.35%) [0.84%]***
Part-time option available[a]	−3,892.01*** (1,024.91) [−2,709.40]***	−829.94 (1,822.82) [−1,866.56]***	−5,527.65***++ (1,221.72) [−3,237.95]***+	−5.13%*** (1.35%) [−3.57%]***	−1.09% (2.40%) [−2.46%]***	−7.29%***++ (1.61%) [−4.27%]***+

Notes. Table reports mean (bootstrap standard errors) [median] WTP (amount of earnings an individual needs to be compensated for a unit change in the job attribute).

[a]WTP for moving from a job without a part-time option to one that has it.

***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Tests conducted for differences in means and medians by gender. +++, ++, + denote estimates differ at 1%, 5%, and 10% levels, respectively.

TABLE VII

Actual Job Characteristics and Estimated WTP

	Prob. of fired	Bonus percentage	Prop. of males	Earnings growth	Hours worked	Flex work option
Willingness to pay[a]	−0.07 (0.20)	−1.00 (1.23)	−7.32** (2.82)	−0.02 (0.08)	−1.70** (0.64)	−0.94** (0.29)
Constant	10.70*** (1.90)	3.64 (2.67)	52.60*** (2.82)	7.32*** (1.71)	46.37*** (2.00)	55.61*** (6.37)
Effect size[b]	−0.658	−4.35	−6.89	−0.319	−4.09	−14.75
p-value[c]			0.012			
Mean of dep. var.	10.4	5.8	50.9	7.3	44.6	61.0
Std. dev. of dep. var.	(14.72)	(12.79)	(22.79)	(13.34)	(14.76)	(49.19)
R-squared	0.002	0.16	0.092	0.0001	0.077	0.090
observations	59	59	59	59	59	59

Notes. The table investigates the relationship between the estimated WTP for a given job attribute for a respondent (derived from the 2012 survey) and the value of that job attribute in the respondent's actual workplace (reported in the 2016 follow-up survey). Each column is a separate OLS regression, with the dependent variable (column title) being the value of the job characteristic in the respondent's actual job (reported in the 2016 survey). Bootstrap standard errors in parentheses. ***, **, * denote significance at 1%, 5%, and 10% levels, respectively.

[a]The estimated WTP of the respondent based on the hypothetical job choice scenarios.

[b]The predicted change in the dependent variable for a one std. dev. change in the WTP.

[c]*p*-value of a test that the six estimates on the WTP (in the first row) are jointly zero.

TABLE VIII

LAD Estimates of Major Choice

LAD estimates	
Job attributes (δ)	0.018** (0.007)
Ability rank	−0.064*** (0.006)
Study time	−0.009 (0.025)
Economics dummy	−0.590 (0.444)
Engineering dummy	−1.16** (0.37)
Natural sci dummy	−0.822* (0.375)
Total observations	741
Number of individuals	247

Notes. Bootstrap standard errors in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE IX

Marginal Contribution of Job Attributes in Major Choice

	Fired prob. (1)	Part-time available (2)	Hours (3)	Bonus (4)	Earnings growth (5)	Prop. males (6)	Log earnings (7)
Panel A: Males							
start value[a]	8.82	0.29	48.88	18.37	1.42	54.86	11.38
End value[b]	19.38	0.53	60.40	41.37	5.20	73.21	12.03
Avg. change[c]	−4.10%	0.23%	−2.48%	9.20%	4.30%	−0.20%	15.90%
Relative change[d]	−0.26	0.01	−0.16	0.58	0.27	−0.01	–
Panel B: Females							
start value	16.51	0.35	47.36	19.12	1.35	56.13	11.22
End value	33.14	0.58	61.15	43.11	4.44	75.22	11.70
Avg. change	−5.12%	0.15%	−1.40%	4.58%	0.70%	0.10%	4.78%
Relative change	−1.07	0.03	−0.29	0.96	0.15	0.02	–

Notes. Table shows the average percent change in the probability of majoring in a given major ("marginal effect") for a standard deviation change in the job attribute (column variable). See Online Appendix E for details.

[a(b)]The initial (final) value of the attribute at which the major probability is computed, with all other attributes fixed at the sample mean.

[c]The average change (across majors) in the probability of majoring in a given major for a standard deviation change in the attribute.

[d]The average change in the probability of majoring in a given major for a std. dev. change in an attribute, relative to a corresponding change in earnings.

TABLE X

Workplace Preferences and Gender Gap in Age-30 Expected and Actual Earnings

	(1)	(2)	(3)	(4)
Panel A: Dependent variable: log(age-30 expected earnings)				
Female	−0.346*** (0.060)	0.195*** (0.055)	−0.289*** (0.065)	−0.150*** (0.057)
Constant	11.483*** (0.048)	11.69*** (0.051)	11.36*** (0.065)	11.55*** (0.068)
Major controls[a]	N	Y	N	Y
Workplace preferences controls[b]	N	N	Y	Y
Mean of dep. var	11.26	11.26	11.26	11.26
R-squared	0.1209	0.3386	0.2013	0.3967
Number of observations	247	247	247	247
Panel B: Dependent variable: log(actual 2016 earnings)				
Female	−0.612*** (0.169)	−0.451*** (0.167)	−0.442** (0.191)	−0.318 (0.230)
Constant	12.12*** (0.145)	12.31*** (0.147)	11.91*** (0.190)	12.12*** (0.188)
Part-time work dummy	Y	Y	Y	Y
Major controls[c]	N	Y	N	Y
Workplace preferences controls	N	N	Y	Y
Mean of dep. var	11.65	11.65	11.65	11.65
R-squared	0.226	0.384	0.395	0.495
Number of observations	56	56	56	56

Notes. OLS estimates presented. Block bootstrap standard errors in parentheses. ***, **, * denote significance at 1%, 5%, and 10% levels, respectively. Dependent variable in Panel A is the log of age-30 expected earnings for the individual's reported major. Dependent variable in Panel B is the log of actual earnings for the subset of individuals who took the follow-up survey and were working in 2016.

[a]Dummy for the major the respondent is majoring in (the major with the modal probability).

[b]Controls for the estimated workplace preferences (from the job-choice model).

[c]Dummy for the major the respondent graduated with.