

Beauty, Job Tasks, and Wages: A New Conclusion about Employer Taste-Based Discrimination*

Todd Stinebrickner[†]
University of Western Ontario
and NBER

Ralph Stinebrickner[‡]
Berea College

Paul Sullivan[§]
American University

January 2018

Abstract

We use novel data from the Berea Panel Study to reexamine the labor market mechanisms generating the beauty wage premium. We find that the beauty premium varies widely across jobs with different task requirements. Specifically, in jobs where existing research such as Hamermesh and Biddle (1994) has posited that attractiveness is plausibly a productivity enhancing attribute—those that require substantial amounts of interpersonal interaction—a large beauty premium exists. In contrast, in jobs where attractiveness seems unlikely to truly enhance productivity—jobs that require working with information and data—there is no beauty premium. This stark variation in the beauty premium across jobs is inconsistent with the employer-based discrimination explanation for the beauty premium, because this theory predicts that all jobs will favor attractive workers. Our approach is made possible by unique longitudinal task data, which was collected to address the concern that measurement error in variables describing the importance of interpersonal tasks would tend to bias results towards finding a primary role for employer taste-based discrimination. As such, it is perhaps not surprising that our conclusions about the importance of employer taste-based discrimination are in stark contrast to all previous research that has utilized a similar conceptual approach.

*We would like to thank Jonathan Williams, seminar participants at Arizona State, Washington University in St. Louis, and the University of Virginia for helpful comments, and Lori Scafidi for her work as data manager. The project was made possible by funding from The National Science Foundation, The Mellon Foundation, The Spencer Foundation, the SSHRC, and the Center for Human Capital and Productivity.

[†]Email: trstineb@uwo.ca

[‡]Email: ralph_stinebrickner@berea.edu

[§]Email: pauljsullivan@gmail.com

1 Introduction

It is now well-established that physically attractive workers tend to earn higher wages than less attractive workers (Hamermesh and Biddle, 1994; Biddle and Hamermesh, 1998; Harper, 2000; Robins, Homer and French, 2011; Scholz and Sicinski, 2015). From the standpoint of considering the merits of potential policy related to this finding, a crucial first step involves understanding the extent to which the attractiveness premium arises because of employer taste-based discrimination of the type described by Becker (1957). If such discrimination is present, then legislating away pay differences based on attractiveness has a fundamental anti-discriminatory appeal.¹ However, whether the net benefits of such legislation are positive is less obvious if the premium arises for other reasons. For example, if attractive workers are simply more productive in certain jobs, then removing the incentive to sort into jobs on the basis of this particular characteristic might negatively affect overall worker-employer match quality. Moreover, an unintended inequity could arise because some individuals (e.g., those with high cognitive ability) would be allowed to take full advantage of their most productive attributes while others (e.g., attractive workers) would not.

In this paper, we use novel data about college graduates from the Berea Panel Study to provide new evidence about the labor market mechanisms generating the beauty premium. These data contain measures of physical attractiveness along with unique information about the type of job tasks that individuals perform. We provide new evidence that the beauty premium varies widely across jobs with different task requirements. Specifically, in jobs where existing research such as Hamermesh and Biddle (1994) has posited that attractiveness is plausibly a productivity enhancing attribute—those that require substantial amounts of interpersonal interaction—a large beauty premium exists. In contrast, in jobs where attractiveness seems unlikely to truly enhance productivity—jobs that require working with information and data—there is no beauty premium. This stark variation in the beauty premium across jobs is inconsistent with the employer-based discrimination explanation for the beauty premium, because this theory predicts that all jobs will favor attractive workers. Moreover, we provide strong evidence that physically attractive workers sort into jobs where beauty is valued, just as theory would predict for any other productive

¹Hamermesh (2011) describes localities, such as Washington D.C., where some form of this type of legislation exists.

attribute. Taken together, the results suggest that beauty should perhaps be simply viewed as a productive attribute that is highly valued in certain jobs.

While our results are consistent with the notion that a profit motive may serve as a “natural antidote” to employer taste-based discrimination based on beauty (Mankiw, 2015), they are in stark contrast to past empirical work involving longitudinal surveys, which has consistently concluded that this type of discrimination is of central importance. For example, in their seminal work, Hamermesh and Biddle (1994) find that “The strongest support is for pure Becker-type discrimination based on beauty and stemming from employer/employee tastes.” Similar conclusions are reached more recently by Harper (2000), who finds that “The bulk of the pay differential for appearance arises from employer discrimination...,” by Scholz and Sicinski (2015), who find that the substantial overall beauty premium “is not solely driven by greater productivity, leaving a role for taste-based discrimination...,” and by Robins, Homer and French (2011), who find that the beauty premium appears to span all occupations.

In considering why our findings deviate substantially from previous work, it is worth noting that one gets a sense of some reluctance in this previous work to view currently available evidence in favor of employer discrimination as the final word on the subject. As is the case in this paper, empirical tests in past work are based on the insight that employer taste-based discrimination would produce beauty premiums across all types of jobs, while a worker productivity-based explanation would produce beauty premiums in only a subset of jobs: those jobs where the presence of substantial interpersonal interactions provides an opportunity for attractiveness to potentially enhance worker productivity. Unfortunately, accurately characterizing the amount of interpersonal interaction on one’s job is a major challenge using standard survey data. This is critical because measurement error introduced during this characterization likely creates an attenuation bias that works in favor of an employer discrimination conclusion, by making it difficult/impossible to detect the scenario where a beauty premium exists only in a predetermined subset of jobs where attractiveness potentially enhances worker productivity. Perhaps this empirical difficulty served as motivation for Biddle and Hamermesh (1998) to reexamine their earlier conclusions obtained from general longitudinal surveys, finding some evidence inconsistent with employer discrimination in a sample of attorneys.²

²Biddle and Hamermesh (1998) find that the beauty premium is no smaller for self-employed attorneys than it is for other types of attorneys.

While our access to college ID pictures allows us to construct an attractiveness variable that possesses measurement benefits described in Scholz and Sicinski (2015) and while our data contain necessary, detailed wage information, our primary motivation for taking a fresh look at the employer discrimination issue comes from our access to unique data describing the tasks performed on jobs.^{3,4} From the standpoint of identifying the set of jobs where attractiveness may be productivity enhancing, the task data represent a critical advance for two related reasons: 1) they represent the only longitudinal data in which job tasks are measured directly for a worker's actual job and 2) they are unique in containing explicit time allocation information, which produce quantitative task measures that are easily interpretable and conceptually appealing. The importance of 1) stems from the fact that, in order to characterize the work that a person performs on her actual job, general longitudinal surveys typically provide, at most, the occupation to which the job belongs. This implies that the tasks associated with the person's actual job must be imputed using information about what is "typical" in the occupation from external (occupation-level) data sources such as the Dictionary of Occupational Titles (DOT). However, recent evidence that variation in tasks tends to be large within occupations implies that this imputation can create substantial measurement error (Autor and Handel, 2013; Robinson, 2011). In our substantive context, the concern is that many jobs requiring substantial interpersonal interactions may be present even in occupations where the median, modal, or average amount of interpersonal interaction is not particularly large.

The importance of directly measuring time spent on job tasks arises because the nature of the information provided in external sources such as the DOT can make it challenging to partition the set of occupations into those where attractiveness potentially enhances worker productivity and those where attractiveness seems unlikely to enhance worker productivity. As an example, consider the fifth digit of the DOT code that is used in Hamermesh and Biddle (1994). This digit

³Investigations into the role of beauty in the labor market have been hampered by a dearth of datasets that contain measures of beauty. Hamermesh and Biddle (1994) finds that only three North American datasets include both measures of beauty and labor market information: the 1977 Quality of Employment Survey, the 1971 Quality of American Life Survey, and The 1981 Canadian Quality of Life Study. Scholz and Sicinski (2015) uses the Wisconsin Longitudinal Study to study students from the high school class of 1957, and Harper (2000) uses the National Child Development Study to study British workers born in 1958. To the best of our knowledge, the Add Health data used in Mocan and Tekin (2010) is the only relatively recent North American longitudinal survey that includes attractiveness measures.

⁴A growing literature has explored the value of task-based approaches to classifying jobs, although not in the context of the attractiveness wage premium. See, for example, Poletaev and Robinson (2008), Gathmann and Schönberg (2010), Sanders and Taber (2012), Yamaguchi (2012, 2018), Sanders (2017), Lise and Postel-Vinay (2015), Speer (2017), Stinebrickner, Stinebrickner and Sullivan (2018), and Roys and Taber (2017).

identifies the highest level of People tasks that is typically required in an occupation, with all occupations being assigned a level from a list consisting of mentoring, negotiating, instructing, supervising, diverting, persuading, speaking, serving, and following instructions. The empirical challenge arises because, while the list is ordered from the highest skilled task to the lowest skilled task, theory does not provide guidance about whether this is the correct order when the objective is to predetermine the set of jobs where beauty is potentially productive. For example, if beauty affects productivity by facilitating communication, symmetric effects of beauty might be expected for those giving instructions (third highest skill) and those following instructions (lowest skill), and it also seems quite possible that a server (second lowest skill) will benefit as much from beauty as a manager who spends time mentoring or negotiating (the first and second highest skills). Further ambiguity about the correct task order in this context arises because there is no theoretical reason why the time spent on interpersonal interactions must increase as the level of People tasks increases across the skills in the list above, and the DOT provides no information about time allocation that could be used to examine whether this is the case.⁵

Our primary contribution is showing that the collection and use of conceptually appealing task measures leads to much different findings about the importance of employer discrimination on the basis of attractiveness than have been reached in the past. Thus, while it is necessary to be cautious when thinking about exactly how our results would generalize to graduates from other educational institutions or to workers with less than a college degree, our paper makes the general point that conclusions about the importance of employer discrimination on the basis of attractiveness may depend critically on a researcher's ability to characterize the tasks performed on a job. We stress that our approach is relevant for studying the issue of attractiveness, but provides no evidence about the importance of discrimination on the basis of race or other characteristics (see, e.g., Black, 1995; Neal and Johnson, 1996; Bertrand and Mullainathan 2004; Charles and Guryan, 2008, 2009). The key difference is that, in the context of testing for discrimination based on attractiveness, it seems reasonable to posit a certain set of jobs where attractiveness might be productive.

⁵Hamermesh and Biddle (1994) form two groups: 1) those whose highest level is receiving instructions and 2) everyone else. Highlighting the difficulty of the measurement challenge, they also try two other approaches for partitioning occupations: 1) soliciting the opinions of eight adults about the importance of attractiveness in occupations and 2) taking advantage of a survey in which employers were asked whether an applicant's appearance was an important consideration in filling the most recent job vacancy in an occupation.

As with previous research in this area, it is not our objective to make specific policy prescriptions. One reason for this is that, while we do focus on understanding the importance of employer preference-based discrimination, it would perhaps never be possible to conclusively rule out all other, sometimes subtle, potential sources of discrimination. For example, for convenience we refer to the primary alternative to an employer preference-based discrimination explanation as a worker productivity explanation. However, although our unique data allow an opportunity to identify assumptions under which we can differentiate between this explanation and a customer preference-based discrimination explanation, in general making this distinction is difficult for both conceptual and empirical reasons. It is also possible that our results are consistent with a scenario in which employers unintentionally discriminate because they have incorrect views about the productivity of attractiveness in jobs that require a large amount of interpersonal interaction (see, e.g., the experimental work of Mobius and Rosenblat, 2006, and Deryugina and Shurchkov, 2015).^{6,7}

As has been frequently recognized, attractiveness may be correlated with other cognitive or non-cognitive traits that influence wages (Mobius and Rosenblat, 2006; Deryugina and Shurchkov, 2015; Scholz and Sicinski, 2015).⁸ Our access to detailed administrative data from the college period allows us to control for the most widely recognized proxy for cognitive human capital at the time of workforce entrance, college GPA, and we also take advantage of access to self-reported measures of non-cognitive skills such as communication ability. While the exact interpretations of estimated attractiveness effects will depend on whether these controls effectively eliminate correlations between attractiveness and the unobservable in the wage equation, we stress that our primary result - that the data are inconsistent with employer taste-based discrimination - does not rely on whether this is the case.⁹

Of relevance for thinking about whether attractiveness is best viewed as just another productive characteristic, we highlight a strong parallel between the manner in which attractiveness

⁶The insight that longitudinal data can potentially help understand the importance of issues such as employer learning (see, e.g., Farber and Gibbons, 1996, and Altonji and Pierret, 2001) helps motivate the examination of the long-run effects of attractiveness in Scholz and Sicinski (2015).

⁷Evaluating the net effects of a specific policy is further complicated if, as has been argued, the presence of an attractiveness premium can lead to negative psychological and physical health effects.

⁸In closely related research, Persico, Postlewaite and Silverman (2004) find that the positive relationship between height and wages is likely due to unobserved attributes that are cultivated by individuals who are tall when young.

⁹In contrast, whether controls effectively eliminate correlations between attractiveness and the unobservable is of first order importance in research motivated by the fact that concerns about employer discrimination are alleviated if the beauty premium is reduced/removed when controls are added.

influences wages and the manner in which GPA influences wages. The estimated predicted change in wages from a one standard deviation change in attractiveness is larger than the estimated predicted change in wages from a one standard deviation change in college GPA. Further, we find that, like our attractiveness measure, the GPA premium is generated entirely by jobs in which one would expect GPA to be productivity enhancing, that is jobs that specialize in high skilled tasks. In sharp contrast, there exists no evidence of a GPA premium in jobs that specialize in low skilled tasks. Further evidence of GPA being a productivity enhancing attribute comes from finding that high-GPA individuals are more likely than other workers to sort into jobs in which they spend the most time on high skilled tasks.

2 Data

This section provides general information about the Berea Panel Study (Section 2.1), describes how attractiveness is measured in the BPS (Section 2.2), and explains how the BPS job task information is used to classify jobs by their skill level and degree of interaction with people and information (Section 2.3).

2.1 The Berea Panel Study

Designed and administered by Todd Stinebrickner and Ralph Stinebrickner, the Berea Panel Study (BPS) is a longitudinal survey that was initiated to provide detailed information about the college and early post-college periods.¹⁰ The project involves surveying students who entered Berea College in the fall of 2000 and the fall of 2001 approximately sixty times from the time of college entrance through 2014. In this paper, we take advantage of post-college surveys that were collected annually after students left school. More than ninety percent of all graduates in the two BPS entering cohorts completed one or more of these annual surveys, and the response rate on these surveys leveled off at slightly more than eighty percent. Our analysis uses all yearly employment observations from the time of graduation until an individual first fails to complete a post-college survey. Since our analysis relies on wages and job tasks, observations where an

¹⁰For previously published work that has used the BPS to study issues in education, see Stinebrickner and Stinebrickner (2003*b*; 2003*a*; 2004; 2006; 2008*b*; 2008*a*; 2010; 2012; 2014*b*; 2014*a*). Stinebrickner, Stinebrickner and Sullivan (2018) describes the job task data in the BPS in detail, and Stinebrickner, Stinebrickner and Sullivan (2017) takes advantage of the job task data to examine the gender wage gap. The job task data also play a secondary role in the analysis of mismatch in Agosowicz et al. (2017).

individual reports not holding a job are dropped from the analysis. For the 506 individuals in the sample, the maximum number of yearly observations is 10 and the average number of yearly observations is 6.2.

Berea College is unique in certain respects that have been discussed in previous work. For example, the school focuses on providing educational opportunities to students from relatively low income backgrounds. As noted throughout this paper and in our previous work that used the BPS to study other issues, it is necessary to be appropriately cautious about the exact extent to which the results from our case study would generalize to graduates from other specific institutions. However, important for the notion that the basic lessons from our work are pertinent for thinking about what takes place elsewhere, Berea operates under a standard liberal arts curriculum, and the students at Berea are similar in academic quality, for example, to students at the University of Kentucky (Stinebrickner and Stinebrickner, 2008). In addition, academic decisions and outcomes at Berea are similar to those found elsewhere (Stinebrickner and Stinebrickner, 2003*a*, 2014*b*). For example, dropout rates at Berea are similar to dropout rates at other schools (for students from similar backgrounds) and patterns of major choice and major-switching at Berea are similar to those found in the NLSY by Arcidiacono (2004). Further, even putting aside the obvious issue of data collection feasibility, there are benefits of studying one school. In particular, the ability to hold school quality constant is beneficial for a variety of reasons, including that it makes academic measures such as college GPA directly comparable across individuals. An additional virtue of our data collection is that it allows us to examine individuals starting at the very beginning of their careers.¹¹

This paper is made possible by merging the unique survey data from the early portion of an individual's working life with detailed administrative data from the college period, which provides basic demographic information as well as academic information characterizing one's human capital at the time of workforce entrance. Table 1 shows descriptive statistics separately for men and women in the BPS. Throughout the paper, we follow previous research on the labor market returns to attractiveness by conducting our analyses separately by gender. However, because Table 1 shows that almost two-thirds of our sample is female, the subsequent empirical work focuses primarily on females. Nonetheless, while our analyses of men are hindered to

¹¹In many respects, our general motivation for studying one school is similar in spirit, for example, to that in Bertrand, Goldin and Katz (2010) who examine the earnings of a sample of MBA graduates from one top business school.

Table 1: Descriptive Statistics by Gender

	(1) Women	(2) Men
Log-wage	2.562 (0.613)	2.628 (0.659)
Attractiveness	2.590 (0.779)	2.576 (0.624)
College GPA	3.204 (0.443)	3.082 (0.473)
High school GPA	3.538 (0.406)	3.352 (0.466)
Family income	27.283 (19.553)	25.231 (18.940)
<u>Task Specialization</u>		
High skilled interpersonal	0.390 (0.488)	0.390 (0.488)
Low skilled interpersonal	0.413 (0.493)	0.392 (0.488)
High skilled information	0.183 (0.387)	0.287 (0.453)
Low skilled information	0.159 (0.366)	0.182 (0.386)
Observations	2,038	1,129
Individuals	324	182

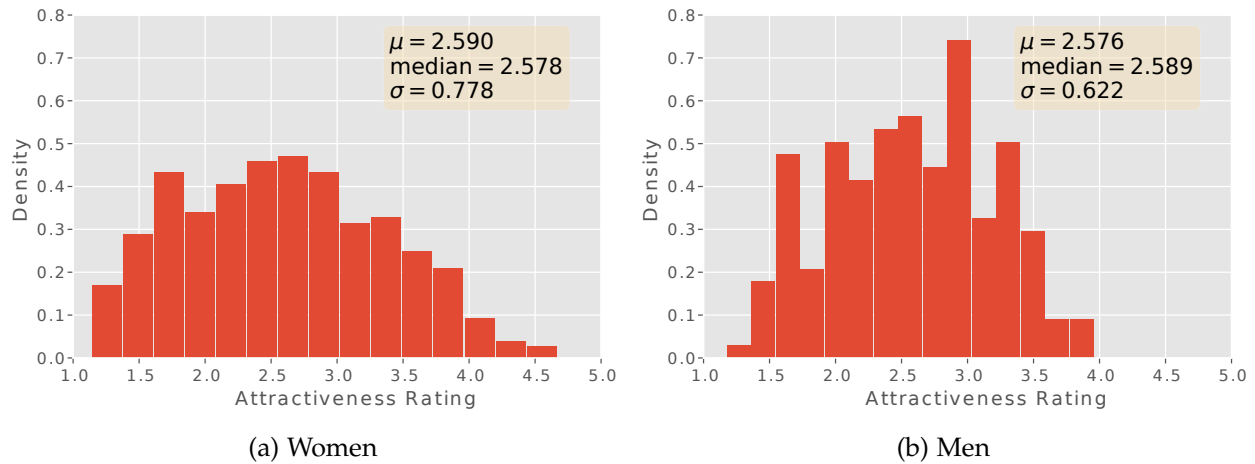
Notes: Entries are means, with standard deviations in parentheses.
Family income measured in \$10,000.

some extent by the smaller sample, we show that our main conclusions for females also apply to men.¹²

Hourly wages are constructed from Survey Question D2 (Appendix A), which gave respondents flexibility over whether earnings were reported over an hourly, weekly, monthly, or yearly period, and Survey Question D1, which elicited a worker's hours in a typical week. For females, Column 1 shows that the mean log hourly wage in 2005 dollars is 2.56 and that the average (std. deviation) of college GPA is 3.20 (0.44). For males, Column 2 shows that the mean log hourly wage in 2005 dollars is 2.63 and that the average (std. deviation) of college GPA is 3.08 (0.47). Table 1 also shows descriptive statistics for the attractiveness measure, which is discussed in detail in Section 2.2, and the job task classification, which is defined in Section 2.3.

¹²As discussed in Stinebrickner and Stinebrickner (2003*b*, 2006, 2008*b*, 2012, 2014*b*), females make up a higher percentage of matriculants at Berea and a lower proportion of the students that drop out. It is worth noting that in this sample of recent college graduates, the labor market experiences of males and females are similar across many dimensions. For example, both men and women have high employment rates (91% men, 85% women).

Figure 1: Distribution of Attractiveness by Gender



2.2 Measuring Attractiveness

Each student's attractiveness was assessed by 50 evaluators on the basis of a color student ID picture, using a five point scale where: 5 = very attractive, 4 = above average, 3 = average, 2 = below average, and 1 = significantly below average. Both oral and written instructions asked each evaluator to first look through all of the 100 pictures that he/she was given and then to assign a number taking into account that, "The number you write should represent how attractive you think the person in the picture is relative to people in the other pictures you received." Our attractiveness measure is the average of the fifty assessments. The histograms in Figure 1 show that the averages and medians of this measure for women and men are virtually identical, but that attractiveness varies somewhat more between women than between men.

Our student-ID approach shares the virtues of the yearbook picture approach described in Scholz and Sicinski (2015). In particular, they find that, as would be expected when variation in assessments exists across evaluators, estimated effects of attractiveness tend to be larger when one has the luxury of reducing measurement error by averaging assessments across multiple reviewers. However, it is worth noting that the facial attractiveness measures in this paper and in Scholz and Sicinski (2015) may have a slightly different interpretation than those in, for example, Hamermesh and Biddle (1994), which are perhaps best viewed as measures of overall beauty because they come from a single in-person assessment of a survey collector.

2.3 Classifying Jobs Based on Tasks

As discussed in the introduction, identifying the subset of jobs where attractiveness is likely to enhance worker productivity is a natural way to distinguish between the employer taste-based discrimination explanation and the productivity-based explanation for the attractiveness premium (Hamermesh and Biddle, 1994). If a wage premium for attractiveness exists in all jobs—including those where it seems implausible that physical attractiveness is truly a productive attribute—this supports the employer discrimination hypothesis.

It seems reasonable to believe that attractiveness may enhance a worker’s productivity in performing interpersonal tasks such as serving customers and negotiating (Hamermesh and Biddle, 1994). Conversely, it seems unlikely that attractiveness directly enhances productivity performing information tasks such as data entry or analyzing data, which are also very common among college graduates. Unfortunately, past research has been forced to infer the degree of interpersonal interaction in a job on the basis of only occupation codes. As detailed in the introduction, this may represent a critical limitation; measurement error arising from fundamental difficulties in characterizing the importance of interpersonal tasks on a job may stack the deck in favor of the employer taste-based discrimination explanation that has been found to be of primary importance in past research.

We address this fundamental difficulty by taking advantage of access to direct, job-level measures of the amount of time spent performing interpersonal tasks and the amount of time spent performing information tasks. Survey Question C.3 (Appendix A) elicits the percentage of total work time that an individual spends on tasks in the general People category and the percentage of total work time that an individual spends on tasks in the general Information category (as well as the percentage spent on tasks in an Objects category). As shown in the Appendix, Questions C.1 and C.2, respectively, elicit the time allocation measures that indicate the percentages of time that are spent on each specific sub-task within the People and Information categories. Classifying the first two sub-tasks (1 and 2) within each of the People (C.1) and Information (C.2) task categories as low skilled and the last two sub-tasks (3 and 4) as high skilled, questions C.1, C.2, and C.3 allow us to compute the percentage of total work time in a year that is spent on each of the following four tasks: high skilled People, low skilled People, high skilled Information, and

low skilled Information.¹³

One could imagine a variety of ways that this task information could be used for our purposes. Our primary results involve taking an approach that is appealing in its simplicity. Specifically, we differentiate between jobs on the basis of the primary task on the job, which is defined to be the task (from the list: high skilled People, low skilled People, high skilled Information, and low skilled Information) in which a worker spends the most time. Table 1 shows that, pooling jobs in all years for females, 39% of jobs have a primary task of high skilled People, 41% of jobs have a primary task of low skilled People, 18% of jobs have a primary task of high skilled Information, and 16% of jobs have a primary task of low skilled Information.¹⁴ The sum of the proportions is greater than one in these descriptive statistics because, for the 14.5 percent of jobs where ties occur, a job is assigned multiple primary tasks. For example, a woman who reports spending 50% of her time on high skilled People tasks and 50% of her time on low skilled People tasks on a particular job would be classified as specializing in both of these tasks.¹⁵

However, we also examine the robustness of our results to alternative ways of using the task data to differentiate between jobs that require large amounts of interpersonal interactions and jobs that do not (Appendix B). Specifically, we classify a job as requiring a substantial amount of a particular task (high skilled People, low skilled People, high skilled Information, or low skilled Information) if the time spent on the task is above a particular percentile (e.g., the median and 75th percentile) of the empirical distribution for that task. The descriptive statistics associated with these robustness classifications appear in Appendix B.

¹³For the purposes of this paper, it would also likely be informative to divide the overall People category into jobs that involve face-to-face interpersonal interactions and jobs that do not require face-to-face interactions. Unfortunately, this type of disaggregation is difficult using Survey Question C. Stinebrickner, Stinebrickner and Sullivan (2018) contains a detailed discussion of issues related to survey design and provides evidence of the usefulness of the survey questions.

¹⁴Given that C.3 elicits information about Objects and given that our survey contains a question about Objects that is analogous to C.2 and C.3 (not shown), we are able to compute the percentage of time spent in high skilled and low skilled Objects tasks. However, we find that, in our sample of female college graduates, only about 7 percent of all jobs have a sole primary task of either high skilled Objects or low skilled Objects. Thus, for all jobs, the job's primary task is based on the maximum of high skilled People, low skilled People, high skilled Information and low skilled Information.

¹⁵Motivated by work such as Sanders and Taber (2012), who note that the primary usefulness of categorizing jobs by occupation is that occupations serve as observable proxies for the true task requirements of jobs, the BPS did not have a focus on creating standard occupational codes. As a result, the empirical analysis in the remainder of the paper focuses solely on this direct, task based method of classifying jobs.

3 Empirical Results- Females

For our primary sample of women, we begin in Section 3.1 by presenting evidence of large wage returns to attractiveness. Inconsistent with an important role for employer taste-based discrimination, Section 3.2 shows that attractive workers earn higher wages only in jobs where attractiveness is plausibly a productive attribute - those that involve a large amount of interpersonal tasks. As would be expected if attractive workers do indeed have a comparative advantage in jobs that involve a large amount of interpersonal tasks, Section 3.3 shows that attractive workers tend to sort into these types of jobs. Section 3.4 discusses the potential alternatives to the employer taste-based discrimination explanation for the attractiveness premium.

3.1 Wages and Attractiveness

Pooling all female observations across all sample years, Column 1 of Table 2 shows the results obtained from regressing log wages on age and our attractiveness measure, which has been normalized to have a standard deviation of one.¹⁶ Consistent with previous research, we find a large, statistically significant coefficient on the attractiveness measure. Specifically, increasing attractiveness by one sample standard deviation (0.78 on the five-point scale) is associated with a 7.8% increase in wages, with the associated t-statistic having a value of 3.55.

A question that has received attention in the literature is whether an observed relationship between attractiveness and wages can be interpreted as a “causal” effect of attractiveness *per se*. This question represents a good starting point given our objectives because concerns about employer taste-based discrimination would be largely irrelevant if attractiveness only appears to affect wages due to a correlation between attractiveness and omitted productivity-enhancing variables. One prominent concern is that a worker’s attractiveness may be related in some way to her cognitive human capital (Scholz and Sicinski, 2015). College GPA has been widely recognized as a leading proxy for one’s human capital at the time of entrance to the labor market, and, because our study involves one school, also has the virtue in our context of being directly

¹⁶Neal and Johnson (1996) and Persico, Postlewaite and Silverman (2004), along with Heckman (1998), argue against controlling for potentially endogenous variables, such as work experience, when estimating the effects of discrimination. However, in part because the large majority of women are working in each year, the results in this section change very little if the model is estimated using actual experience rather than age. Neal and Johnson (1996) study the black-white wage gap using a parsimonious wage regression that includes a measure of skills (AFQT). Persico, Postlewaite and Silverman (2004) investigate whether wage differentials between tall and short adults can be attributed to discrimination, or a number of alternative explanations.

Table 2: Regression of Log-Wage on Attractiveness for Women

Variable	Specifications							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Attractiveness (standardized)	0.078 (0.022)	0.071 (0.022)		0.070 (0.023)	0.055 (0.024)	0.054 (0.024)	0.073 (0.022)	
College GPA (standardized)		0.053 (0.022)	0.055 (0.022)	0.048 (0.027)	0.054 (0.025)	0.057 (0.025)	0.049 (0.022)	0.050 (0.021)
High School GPA				0.011 (0.069)				
Family income				0.010 (0.012)				
<u>Attractiveness (categorical)</u>								
Second quartile			0.063 (0.059)					0.060 (0.057)
Third quartile			0.102 (0.062)					0.101 (0.061)
Fourth quartile			0.181 (0.067)					0.183 (0.066)
<u>Self-Rated Skills</u>								
Communication top 25%					0.089 (0.038)			
Personality top 25%					-0.033 (0.045)			
Relateability top 25%					0.003 (0.040)			
<u>Task Specialization</u>								
High skilled people							0.053 (0.042)	0.055 (0.043)
Low skilled people							-0.042 (0.041)	-0.040 (0.041)
High skilled information							0.171 (0.044)	0.171 (0.045)
R ²	0.045	0.053	0.051	0.054	0.034	0.029	0.068	0.066
Observations	2,037	2,037	2,037	2,028	1,517	1,517	2,037	2,037

Notes: Regressions also include a constant and age. Attractiveness and College GPA are standardized to have a mean equal to zero and a standard deviation equal to one. Log-wages are in 2005 dollars. Family income is measured in \$10,000. Relateability is "ability to relate to others." Standard errors are clustered by individual.

comparable across the students. Column 2 of Table 2 shows that adding College GPA to the wage regression has little effect on the estimated effect of attractiveness; the estimate falls from 0.078 to 0.071, reflecting the existence of a relatively small, positive correlation between college GPA and attractiveness. Access to the GPA measure is also beneficial because it provides a benchmark against which the importance of attractiveness can be compared. The (normalized) GPA estimate indicates that increasing GPA by one sample standard deviation (0.44 GPA points) is associated with a wage increase of 5.3 percent (t-statistic of 2.41), somewhat smaller than the estimated wage increase of 7.1 percent associated with a one standard deviation increase in attractiveness.

While the detailed administrative data contain a variety of other measures that may be related to human capital, the conceptual motivation for controlling for these measures is not particularly strong if, for example, the wage influence of a variable such as high school GPA comes through its effect on college GPA. Nonetheless, as a robustness check, in Column 4 we add high school GPA and family income to the specification in Column 2. As expected given that the specification also contains college GPA, we find that neither of these variables are statistically significant at conventional levels and that the effect of attractiveness remains substantively unchanged (0.070 from 0.071). Similar results are found if the ACT college entrance exam score is included in place of high school GPA or if both of these variables are included together.¹⁷

In terms of other potentially important sources of omitted variables bias, it has been recognized that attractiveness or other physical features may be related to non-cognitive skills such as one's communication ability (Feingold, 1992).¹⁸ Self-reported non-cognitive measures are available for 78% of our person-year observations. Turning to Column 5, we find some evidence that non-cognitive skills influence wages. In particular, students who report having communication skills in the top 25% of Berea graduates have wages that are 8.9 percent higher than other students, with the associated t-statistic having a value of slightly more than two. However, comparing the attractiveness coefficient in Column 5 to the attractiveness coefficient in Column 6,

¹⁷Given the strong correlation between high school GPA and ACT, interpretation of coefficients is perhaps difficult when both variables are included. Regardless, the coefficient on attractiveness is 0.067 when ACT is added by itself to the specification in Column 4 and is 0.068 when both variables are added to the specification in Column 4. The estimated effects of high school GPA and ACT are not significant in either of these specifications.

¹⁸Lewis (2011) raises the possibility that attractiveness ratings may potentially be related to race. About 15 percent of the students in our sample are black. Adding a variable indicating whether a student is black is found to not have an effect on our estimates related to attractiveness. For example, adding race to the specification in Column 2 of Table 2 leads to an attractiveness estimate (std. error) of 0.072 (0.022). The estimate (std. error) associated with the coefficient on the race variable is 0.053 (0.060).

which is obtained by estimating the specification in Column 2 for the smaller sample available in Column 5, reveals that the addition of non-cognitive skills produces a change of only 0.001 in the estimated effect of attractiveness. Even though self-reported communication skills are positively related to wages, their inclusion in the wage regression does not lead to any appreciable reduction of the estimated effect of attractiveness.

Thus, Columns 4 and 5 indicate that adding other measures of cognitive human capital or adding available measures of non-cognitive skills has very little effect on the estimated effect of attractiveness. Given our desire to keep our sample as large as possible and given the conceptual/interpretation issues that arise when measures such as high school GPA and ACT are added to a specification that also includes college GPA, in the remainder of the paper we focus on the parsimonious specification shown in Column 2. We stress that our approach for examining the importance of employer taste-based discrimination in Section 3.2 - which involves concluding that there does not exist evidence of employer taste-based discrimination if attractiveness leads to higher wages only in jobs where the presence of substantial interpersonal interactions provides an opportunity for attractiveness to potentially enhance productivity - does not rely on establishing that attractiveness is uncorrelated with all other unobserved wage-influencing characteristics. Indeed, if uncertainty is present about whether this type of correlation exists, it would only make it difficult to know whether, for example, higher wages for attractive workers in jobs with substantial amounts of interpersonal interaction are due to attractiveness per se or are due to other unobserved characteristics that are correlated with attractiveness.

To conclude our current discussion of Table 2, the third column reports a specification that relaxes the linearity assumption for attractiveness by including indicator variables for attractiveness quartiles. There is relatively little evidence of strong non-linear effects. The wages of individuals in the second, third, and fourth quartiles are 6.3%, 10%, and 18.1%, respectively, higher than individuals in the lowest attractiveness quartile (t-statistics of 1.07, 1.65, and 2.70, respectively).

3.2 Wages and Attractiveness across Job Task Specializations: Evidence about Employer Based Discrimination

In this section, we provide evidence about the presence of employer taste-based discrimination. Our test is based on the insight that employer taste-based discrimination would tend to produce

Table 3: Log-Wage Regressions by Primary Job Task and Skill Level for Women: Linear Attractiveness Specification

Primary Job Task Skill Level	(A) People		(B) Information	
	High (1)	Low (2)	High (3)	Low (4)
Attractiveness (standardized)	0.097 (0.025)	0.093 (0.029)	0.011 (0.045)	-0.039 (0.047)
College GPA (standardized)	0.090 (0.031)	0.034 (0.033)	0.064 (0.035)	0.023 (0.048)
R ²	0.064	0.064	0.038	0.043
Observations	804	832	375	322

Notes: Attractiveness and College GPA are standardized to have a mean of zero and standard deviation of one. Regressions include a constant and age. Standard errors are clustered by individual.

attractiveness premiums across all types of jobs, while a worker productivity-based explanation would produce attractiveness premiums in only a subset of jobs: those jobs where the presence of substantial interpersonal interactions provides an opportunity for attractiveness to potentially enhance worker productivity.

We operationalize our test by estimating wage regressions after stratifying our sample into four groups on the basis of jobs' primary tasks: two groups where the primary task involves interpersonal interactions (high skilled People and low skilled People) and two groups where the primary task does not involve interpersonal interactions (high skilled Information and low skilled Information). Table 3 shows wage equation estimates for our linear attractiveness specification from column 2 of Table 2. Providing very strong evidence that the attractiveness premium should not be attributed to an employer taste-based explanation, the results show that attractiveness has a strong effect on wages in jobs that specialize in People tasks, but not in jobs that specialize in Information tasks. Specifically, a one standard deviation increase in attractiveness leads to a 9.7% wage increase in high skilled People jobs (column 1) and a 9.3% wage increase in low skilled people jobs (column 2). The associated t-statistics are each greater than three. In sharp contrast to the large wage premia in People jobs, columns 3 and 4 of Table 3 show no evidence of an attractiveness premium in Information jobs. Specifically, the coefficients on attractiveness for high and low skilled Information jobs are 0.011 and -0.039, and neither parameter is statistically significant (t-statistics of 0.24 and -0.83).

Appendix B shows that the general conclusion from Table 3 is robust to alternative ways of using the task data to differentiate between jobs that require large amounts of interpersonal interactions and jobs that do not. Specifically, Appendix B presents regressions in which a job is classified into the four task specializations shown in Table 3 based on whether or not the time spent on each task is above a particular percentile (e.g., the 50th or 75th percentiles) of the empirical distribution for that task. For example, using the 75th percentile, Table 9 (Appendix B) shows that the attractiveness coefficient (std. error) is 0.101 (0.030) for high skilled People jobs, is 0.098 (0.035) for low skilled People jobs, is 0.016 (0.033) for high skilled Information jobs, and is -0.009 (0.031) for low skilled Information jobs.

To tie our empirical approach directly to the type of regression specification that has often been estimated in previous research (Hamermesh and Biddle, 1994), we turn to a single regression model that uses an interaction term to allow the attractiveness wage premium to vary with the task specialization of a job. Defining the binary variable *People* to have a value of one if a worker has a primary task of either high skilled People tasks or low skilled People tasks, we estimate the log-wage regression in equation (1) using OLS,

$$\log(w) = \beta_1 \text{People} + \beta_2 \text{Attr} + \beta_3 (\text{Attr} \times \text{People}) + \beta_4 \text{GPA} + u \quad (1)$$

$$= \underset{(0.038)}{-0.053} \text{People} - \underset{(0.039)}{0.006} \text{Attr} + \underset{(0.037)}{0.103} (\text{Attr} \times \text{People}) + \underset{(0.022)}{0.053} \text{GPA}, \quad (2)$$

$$N = 2,037; R^2 = 0.060$$

where *Attr* is the attractiveness rating, *GPA* represents college GPA, and *u* represents the unobservable determinants of log wages. The point estimates and standard errors for the coefficients in equation (1) are shown in equation (2).¹⁹ This model provides evidence against the employer discrimination explanation for the beauty premium if $\beta_2 \approx 0$ and $\beta_3 > 0$.²⁰ As expected, given the results in Table 3, this specification provides strong evidence that attractiveness has a payoff only in jobs that specialize in People tasks. Specifically, the estimated coefficient on attractiveness in non-People jobs (β_2) is small in magnitude, at -0.006 , and with a *t*-statistic of only -0.15 , is

¹⁹Following all other models in the paper, this regression also includes a constant and age, which are omitted above for the sake of brevity.

²⁰The results in Table 3 are highly informative about what we would expect from this specification. Minor differences may arise because, for example, the regressions in Table 3 do not constrain the effect of College GPA or the variance of the error term in the regression to be the same across columns.

Table 4: Log-Wage Regressions by Primary Job Task and Skill Level for Women: Categorical Attractiveness Specification

Primary Job Task Skill Level	(A) People		(B) Information	
	High (1)	Low (2)	High (3)	Low (4)
<u>Attractiveness</u>				
Second quartile	-0.014 (0.084)	0.080 (0.088)	0.009 (0.090)	0.180 (0.091)
Third quartile	0.142 (0.072)	0.120 (0.092)	0.024 (0.111)	-0.092 (0.135)
Fourth quartile	0.235 (0.075)	0.213 (0.093)	0.058 (0.136)	-0.050 (0.137)
College GPA (standardized)	0.089 (0.031)	0.038 (0.033)	0.062 (0.034)	0.016 (0.047)
R ²	0.067	0.057	0.039	0.070
Observations	804	832	375	322

Notes: College GPA is standardized to have a mean of zero and standard deviation of one. Regressions include a constant and age. Standard errors are clustered by individual.

not statistically different from zero at conventional levels. In stark contrast, the estimate of the additional attractiveness premium in People jobs (β_3) is large in magnitude at 0.103, and has a t-statistic of 2.74. Thus, this regression demonstrates in a parsimonious way that the attractiveness wage premium exists only in jobs where workers specialize in interpersonal tasks (People = 1).

A very similar story appears in Table 4, which shows the stratified results using the non-linear specification from Column 3 of Table 2. For example, for workers with a primary task of high (low) skilled People, Column 1 (Column 2) shows that being in the top attractiveness quartile leads to a 23.5 percent (21.3 percent) increase in wages relative to the bottom attractiveness quartile, with the associated t-statistic having a value of 3.13 (2.29). However, for workers with a primary task of high (low) skilled Information, Column 3 (Column 4) shows that being in the top attractiveness quartile leads to a 5.8 percent (-5.0 percent) increase in wages relative to the bottom attractiveness quartile, with the associated t-statistic having a value of only 0.42 (-0.36).

Thus, we find that the attractiveness premium is only present in jobs that specialize in interpersonal tasks. These results provide some of the strongest existing evidence contradicting the hypothesis that employer discrimination is the primary source of the attractiveness premium.

Our findings suggest that perhaps attractiveness should be viewed in the same way as other

attributes that are widely accepted as having an influence on a worker's productivity. This notion is further bolstered by evidence of strong parallels in Tables 3 and Table 4 between one such attribute, college GPA, and attractiveness. The results show that GPA has a significant wage effect in jobs with substantial amounts of high skilled tasks (where GPA is likely to be most productive), but not in jobs with substantial amounts of low skilled tasks (where GPA is likely to be less productive). Specifically, for workers with a primary task of high skilled People (Information), Column 1 (Column 3) of Table 3 shows that a one standard deviation increase in GPA leads to a 9.0% (6.4%) increase in wages, with the associated t-statistic having a value of 2.90 (1.82). However, for workers with a primary task of low skilled People (Information), Column 2 (Column 4) shows that a one standard deviation increase in GPA leads to a 3.4% (2.3%) increase in wages, with the associated t-statistic having a value of 1.03 (0.47). If anything, evidence of wage differences across task categories for which productivity differences were expected a priori is even stronger for attractiveness than for GPA.²¹

3.3 Sorting into Job Tasks on the Basis of Attractiveness

The results in Section 3.2 suggest that attractive workers have a comparative advantage in jobs that involve large amounts of interpersonal tasks. If this is the case, we would also expect to see that attractive individuals tend to sort into jobs that require more People tasks. We find that this type of sorting is present. Table 5 shows estimates of two probit specifications where the binary dependent variable is defined to have a value of one if a worker's primary job task is either high skilled People or low skilled People (so that it has a value of zero if the primary task is either high skilled Information or low skilled Information).²² The estimates of the marginal effects in column 1 show that increasing attractiveness by one standard deviation increases the probability of having a primary task of People by 0.053, with the associated t-statistic having a value of approximately five.²³ Column 2, which allows attractiveness to enter non-linearly, shows that much of the attractiveness effect in Column 1 arises due to the strong sorting of

²¹Of course, this might be caused, in part, by GPA having some productivity benefits in low skilled tasks, while attractiveness is unlikely to have any productivity benefits in tasks that do not involve interpersonal interactions.

²²All of the discrete choice models estimated in this section exclude observations where there are ties in the task time-use amounts that determine the primary task.

²³Sorting based on attractiveness is present at all stages of the career. When the probit model from Table 5 is estimated using the first three years of data, the estimated marginal effect of attractiveness is 0.050, with a standard error of 0.018. Based on the final three years of data, the corresponding estimate and standard error are 0.047 and 0.022.

Table 5: Probit Models of Sorting into Interpersonal Job Tasks for Women

Variable (Marginal Effects)	Specifications	
	(1)	(2)
Attractiveness (standardized)	0.053 (0.010)	
College GPA (standardized)	-0.043 (0.011)	-0.042 (0.011)
<u>Attractiveness (categorical)</u>		
Second quartile		-0.015 (0.028)
Third quartile		0.044 (0.027)
Fourth quartile		0.122 (0.025)
Pseudo R ²	0.021	0.023
Observations	1,698	1,698

Notes: Binary dependent variable = 1 if the job specializes in People tasks. Coefficients are marginal effects evaluated at the mean values of the explanatory variables. The sample consists of women. Models also include a constant and age.

highly attractive people into jobs that specialize in interpersonal tasks. Specifically, being in the top attractiveness quartile increases the probability of having a primary task of People by 0.122 relative to individuals in the bottom attractiveness quartile, with the associated t-statistic again having a value of close to five.

The strong influence of attractiveness on whether or not a person holds a job with a substantial amount of interpersonal tasks is also apparent in Table 6, which shows results from a multinomial probit model that distinguishes between the four primary-task outcomes: high skilled People, low skilled People, high skilled Information, and low skilled Information. Specification (A), shown in columns 1-3, allows attractiveness to enter the model linearly. The estimated coefficients on attractiveness (t-statistics) for the high and low skilled People tasks, respectively, are 0.23 (4.24) and 0.17 (3.05). With low skilled Information specified as the base case, these estimates indicate that attractive individuals are more likely to choose to work in both high skilled People jobs and in low skilled People jobs than in low skilled Information jobs. Examining whether attractive individuals are more likely to choose to work in high skilled People jobs and/or in low skilled People jobs than in high skilled Information jobs requires differencing the coefficients in the first two columns from the coefficient in the third column. The differences of 0.144 and 0.083

Table 6: Multinomial Probit Models of Jointly Sorting into Job Tasks and Skill Levels for Women

Task	Specifications					
	(A) Linear Attractiveness			(B) Categorical Attractiveness		
	People		Information	People		Information
Skill Level	High (1)	Low (2)	High (3)	High (4)	Low (5)	High (6)
Attractiveness (standardized)	0.229 (0.054)	0.168 (0.055)	0.085 (0.061)			
College GPA (standardized)	0.045 (0.053)	-0.172 (0.053)	0.042 (0.058)	0.045 (0.053)	-0.172 (0.053)	0.045 (0.058)
<u>Attractiveness</u>						
Second quartile				-0.125 (0.144)	0.018 (0.142)	0.026 (0.159)
Third quartile				0.180 (0.148)	0.093 (0.148)	0.200 (0.163)
Fourth quartile				0.587 (0.158)	0.459 (0.159)	0.289 (0.177)
Observations						1,799

Notes: The omitted category is low skilled Information. The sample consists of women. Models also include a constant and age.

have associated p-values of 0.005 and 0.105, respectively. Our results based on direct measures of the task content of jobs seemingly represent the strongest and most consistent evidence found to date that individuals systematically sort into jobs based on attractiveness.²⁴

Specification (B) in Table 6, which allows attractiveness to enter the model through quartile dummy variables, shows that sorting into People tasks is driven, to a large extent, by very attractive individuals. With the low skilled Information category again specified as the base case, the estimates show that being in the top quartile of attractiveness significantly increases the “value” of both the high and low skilled People options relative to the low skilled Information option (t-stats of 3.71 and 2.89, respectively). Examining whether very attractive individuals are more likely to choose to work in high skilled People jobs and/or in low skilled People jobs than in high skilled Information jobs requires differencing the fourth quartile coefficients in the fourth and fifth columns from the fourth quartile coefficient in the sixth column. The differences of 0.298 and 0.170 have associated p-values of 0.045 and 0.257, respectively.

As a bit of an aside, we note that, paralleling the sorting by attractiveness, there is some

²⁴Hamermesh and Biddle (1994) find some evidence consistent with the notion that workers sort across occupations based on beauty. In their study of lawyers, Biddle and Hamermesh (1998) find evidence that attractive lawyers are more likely to work in private sector jobs.

evidence that high GPA individuals sort into jobs that involve high skilled tasks. For example, differencing the college GPA coefficients in the first column and the third column, respectively, from the college GPA coefficient in the second column indicates that high-GPA workers are substantially more likely to work in both high skilled People and high skilled Information jobs than in low skilled People jobs. The p-values for the tests that the difference equals zero are both less than 0.01.

The results from the two multinomial probit specifications confirm our prior finding from the binary probit models that attractive individuals tend to sort into jobs that involve substantial amounts of interpersonal interaction. The multinomial aspect of the model is important because it helps bolster the conclusion that the observed sorting arises because of issues related to People tasks per se. Specifically, knowing that the sorting into People jobs on the basis of attractiveness remains even when we explicitly differentiate between high and low skilled tasks rules out the possibility that the attractiveness results in Table 5 is simply a byproduct of attractive workers sorting into high-skilled jobs. Sufficient for generating this conclusion, Table 6 finds no evidence that attractive workers tend to sort into high skilled jobs; statistical tests reveal no evidence of a difference between the attractiveness coefficients for high skilled People (column 1) and low skilled People (column 2) and also reveal no evidence of a difference between the attractiveness coefficients for high skilled Information (column 3) and low skilled Information (which is normalized to zero).

It is an empirical question whether the overall attractiveness premium discussed in Section 3.1 is primarily due to the higher pay for attractive workers in particular types of jobs seen in Section 3.2 or due to the sorting of attractive workers into particular types of jobs seen in Section 3.3. We find evidence in support of the former explanation. A comparison of the attractiveness coefficient in Column 7 of Table 2 to the attractiveness coefficient in Column 2 of Table 2 reveals that adding primary tasks in Column 7 does not have a substantial effect on the overall attractiveness coefficient. A similar result is seen by comparing Column 3 and Column 8, which adds primary tasks to the specification in Column 3. These results arise to a large extent because, as seen in Column 7 and Column 8, the People tasks into which attractive people sort are not necessarily the highest paying tasks; the wage order of the primary tasks (from highest wage to lowest wage)

is: high skilled Information, high skilled People, low skilled Information, low skilled People.²⁵

3.4 Exploring Alternatives to Employer Discrimination

The results in Section 3.2 and Section 3.3 are inconsistent with an employer discrimination explanation for the attractiveness wage premium. While we tend to refer to the alternative to this explanation as a productivity-based explanation, it is worth stressing that it is difficult, even from a conceptual standpoint, to distinguish between a productivity-based explanation and a customer taste-based discrimination explanation. To help fix ideas, consider a standard textbook-type example in which a customer is willing to pay more to interact with an attractive waiter/waitress. This preference might be viewed as productivity-based if the attractiveness leads to more efficient employee/customer interactions that help a customer arrive at the best possible food order. On the other hand, this preference might be viewed as customer taste-based discrimination if attractiveness does not influence the customer's order, but the customer simply enjoys looking at a more attractive employee.

The existence of only fairly nuanced differences between the two scenarios in the example highlight why it will always be difficult to conclusively distinguish between the customer discrimination and productivity-based explanations (Hamermesh and Parker, 2005), and we stress that making such a distinction is not the primary objective of this paper. Nonetheless, given the nature of our data, it seems worthwhile to provide some further discussion about how one might attempt to probe a little into this issue. An empirical attempt to differentiate between the two explanations might posit that the latter (productivity-based) channel, if important, would affect wages across a wide range of interpersonal job tasks, while the former (customer discrimination) channel would only affect wages in jobs that involve directly interacting with consumers. In our data, low skilled People jobs are largely defined to have tasks that involve interacting with consumers (e.g., serving customers or sales), while high skilled People jobs are largely defined to have tasks that involve interacting with other employees within the firm (supervising or exchanging ideas/negotiating). Then, under the posited assumption, our finding that the attractiveness premium is virtually identical for high skilled People jobs as it is for low skilled People jobs might be viewed as suggestive evidence that, at the very least, not all of the attractiveness

²⁵This is consistent with the findings in Stinebrickner, Stinebrickner and Sullivan (2018).

premium is arising due to customer taste-based discrimination.²⁶

4 The Relationship between Attractiveness and Wages for Men

The primary focus on women in the paper is motivated by the fact that our data contain only half as many person-year observations for men. Nonetheless, in this section we show that the main result from Section 3 - that, inconsistent with an employer discrimination explanation, attractiveness is positively related to wages only in jobs where substantial interpersonal interactions take place - is also true for men. While the larger sample of women gave us the luxury of differentiating between high and low skill levels within the People and Information task categories, this is not a crucial conceptual distinction for our primary objectives. Moreover, Section 3 found strong similarities in the effects of attractiveness by skill level (high or low) for women. For our examination of the role of attractiveness for men in this section, these considerations motivate us to aggregate: 1) the high skilled People primary task category and the low skilled People primary task category into a single overall People category and 2) the high skilled Information primary task category and the low skilled Information primary task category into a single overall Information category.

The first two columns of Table 7 show the attractiveness results for men using the linear attractiveness specification. The results show the same role for attractiveness as was found earlier for women. For workers with a primary task of (either high or low skilled) People, Column 1 shows that a one standard deviation increase in attractiveness leads to a 6.8% increase in wages, with the associated t-statistic having a value of 2.42. However, for workers with a primary task of (either high or low skilled) Information, Column 2 shows that the estimated effect has a slightly negative value of -0.024 (t-statistic -0.65).

Analogously, to examine the results related to GPA, we aggregate 1) the high skilled People primary task category and the high skilled Information primary task category into a single overall High-Skilled category and 2) the low skilled Information primary task category and the low

²⁶This finding that a robust wage premium exists across different interpersonal tasks that might be expected to have substantially different amounts of direct exposure to consumers is seen even when we further disaggregate People tasks. Specifically, when we estimate our linear specification from column 2 of Table 2 for the four skill sub-tasks in Question C.1, the estimated effect (standard error) of attractiveness is 0.082 (0.029) for subtask 1, 0.085 (0.058) for subtask 2, 0.092 (0.026) for subtask 3, and 0.053 (0.043) for subtask 4. Combining males and females to increase sample sizes, the estimated effects (standard error) of attractiveness is 0.072 (0.022) for subtask 1, 0.069 (0.043) for subtask 2, 0.067 (0.023) for subtask 3, and 0.063 (0.031) for subtask 4.

Table 7: Log-Wage Regressions by Primary Job Task for Men and Women

Gender	(A) Men				(B) Women			
	People	Information	Both		People	Information	Both	
Primary Job Task	Both	Both	High	Low	Both	Both	High	Low
Skill Level	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Attractiveness (std.)	0.068 (0.028)	-0.024 (0.037)	0.021 (0.029)	0.035 (0.032)	0.097 (0.022)	-0.008 (0.037)	0.070 (0.025)	0.065 (0.028)
College GPA (std.)	0.050 (0.035)	0.063 (0.038)	0.093 (0.035)	0.011 (0.035)	0.056 (0.024)	0.047 (0.031)	0.077 (0.026)	0.028 (0.031)
R ²	0.061	0.122	0.068	0.062	0.067	0.040	0.050	0.050
Observations	767	422	715	606	1,501	623	1,131	1,106

Notes: Attractiveness and College GPA are standardized to have a mean of zero and standard deviation of one within each gender. Models also include a constant and age. Standard errors are clustered by individual.

skilled People primary task category into a single overall Low-Skilled category. The results in Column 3 and Column 4, which use these aggregated categories, are consistent with the earlier GPA results for women. For workers with a High-Skilled primary task (where GPA is likely to be most productive), Column 3 shows that a one standard deviation increase in GPA leads to a 9.3% increase in wages, with the associated t-statistic having a value of 2.66. However, for workers with a Low-Skilled primary task (where GPA is likely to be less productive), Column 4 shows that the estimated effect is only 0.011 (t-statistic 0.31).

For comparison purposes, the final four columns of Table 7 present the aggregated-task results for women. When comparing the female and male results, it is useful to remember that attractiveness is standardized separately by gender so that the standardized measure is mean zero with a standard deviation of one for each gender. Standardizing both the male results and the female results instead by the overall (pooled) mean and standard deviation makes it easier to directly compare the magnitude of the estimates between genders. When attractiveness is standardized in this manner, the attractiveness coefficient estimates (t-statistics) for People jobs are 0.079 (2.38) and 0.090 (4.41) for men and women, respectively.²⁷

²⁷As would be expected, standardizing the variables in this slightly different way does not alter the finding that the attractiveness coefficient estimates in Information jobs are small in magnitude and statistically insignificant at standard levels. Specifically, the estimates (with t-statistics in parentheses) for Information jobs are -0.028 (-0.66) and -0.007 (-0.22) for men and women.

5 Conclusion

Our approach to examining the role that employer taste-based discrimination plays in the determination of the attractiveness wage premium is based on the insight that such discrimination would tend to produce premiums across all types of jobs, while a worker productivity-based explanation would produce premiums in only a subset of jobs: those where the presence of substantial interpersonal interactions provides an opportunity for attractiveness to potentially enhance worker productivity. Unfortunately, past efforts to take full advantage of this conceptually appealing approach may have been undermined to some extent by unavoidable empirical difficulties. Perhaps most prominently, it may be challenging using traditional data sources to accurately characterize the importance of interpersonal interactions on jobs. Notably, a burgeoning recent literature interested in using explicit measures of job tasks to provide new insight into a variety of labor market topics has raised awareness of the substantial measurement error that may be present when job tasks are characterized using information from (occupation-level) sources such as the Dictionary of Occupational Titles.

Measurement error in tasks is particularly worrisome in the attractiveness context; the nature of the approach described above implies that its presence would tend to stack the deck in favor of finding evidence of employer taste-based discrimination. Our reexamination of the role of employer taste-based discrimination in this paper is motivated by the raised awareness of issues related to task measurement, along with our collection of longitudinal task data that directly address the limitations of traditional data. In particular, our data represent the first longitudinal data that characterize tasks at the level of individual jobs (rather than at the level of occupations) and also contain explicit time-allocation measures. Together, these features allow us to observe exactly how much time individuals spend on high skilled People tasks and low skilled People tasks in each year of work.

Our conclusion that the data are inconsistent with a primary role for employer taste-based discrimination stands in direct contrast to all previous work that takes advantage of similar conceptual approaches. For a variety of reasons, including the reality that our study involves one school and the reality that our sample includes only college graduates, it seems most warranted to draw attention to the general message of our work - that conclusions about the importance of employer taste-based discrimination in this context may depend critically on a researcher's

ability to characterize the tasks performed on a job. However, new initiatives to collect task data that share some virtues of our unique task data may make it possible in the future to compare our estimates to those obtained using samples from wider populations.

Appendix A: Survey Questions

Question C: How does your JOB1 require you to relate to PEOPLE and INFORMATION?

- Question C1: Below are 4 ways that you may interact with PEOPLE on a job.
 1. Following instructions from others such as supervisors or directly serving the needs of customers or animals.
 2. Persuading others about a company product/service or point of view (e.g. sales) or entertaining others.
 3. Supervising others or instructing/teaching others.
 4. Exchanging ideas/information/opinions or negotiating with others to make decisions or formulate policies.
 - Think about the TOTAL time that you spend **interacting with PEOPLE** as part of your JOB1. Indicate what percentage of the total time is spent interacting in each of the four ways. **Note:** Each percentage should be between 0 (the item plays no role) and 100 (all interactions are from the one item) **and the four items should sum to 100.**

- Question C2: Below are 4 ways that you may interact with INFORMATION on a job.
 1. Entering data; typing documents written by others; posting information etc.
 2. Gathering or classifying information/data and performing simple calculations using data.
 3. Analyzing data/information
 4. Using data analysis done by yourself/others to develop knowledge/solutions and make important decisions.
 - Think about the TOTAL time that you spend **interacting with INFORMATION** as part of your JOB1. Indicate what percentage of the total time is spent interacting in each of the four ways. **Note:** Each percentage should be between 0 (the item plays no role) and 100 (all interactions are from the one item) **and the four items should sum to 100.**

- Question C3: Now think about your TOTAL job responsibilities on your JOB1. Indicate the percentage of your responsibilities that involve interacting with PEOPLE, INFORMATION, and OBJECTS, respectively. Each percentage should be between 0 and 100 and the three percentages should sum to one.

Question D: Hours and Earnings for JOB1

- Question D1: How many hours do you typically work each week in your JOB1?
- Question D2: Approximately how much do you earn in your JOB1? NOTE: Please indicate both a dollar amount and whether this amount is your pay per hour, per day, per week, per month, per year etc. For example, if you earn \$8.50 an hour, please write \$8.50 per hour. If you earn \$30,000 per year, please write \$30,000 per year.

Appendix B: Wage Regression when Jobs are Classified Based on Task Percentiles

The wage regressions below are robustness checks for the estimates shown in Table 3. In these regressions, jobs are classified into each task category (high skilled People, low skilled People, high skilled Information, low skilled Information) based on whether or not the fraction of

time spent on a particular task is above the median (Table 8), or the 75th percentile (Table 9). These results provide additional evidence in favor of the hypothesis that attractiveness is a valuable attribute in jobs that require interpersonal interactions. This is particularly evident in Table 9, where in order to be classified as a "People" job, interpersonal task usage must be at or above the 75th percentile.

Table 8: Log-Wage Regressions for Women: Jobs Above Median Tasks Usage in Each Category

Primary Job Task Skill Level	(A) People		(B) Information	
	High (1)	Low (2)	High (3)	Low (4)
Attractiveness (standardized)	0.074 (0.024)	0.088 (0.027)	0.040 (0.026)	0.015 (0.026)
College GPA (standardized)	0.082 (0.028)	0.042 (0.028)	0.062 (0.026)	0.035 (0.026)
R ²	0.053	0.052	0.049	0.043
Observations	1,022	1,028	1,055	1,063

Notes: Observations are assigned to each task category if task usage is above the median in each category. The 50th percentiles for each task are as follows: High Skilled People: 0.21, Low Skilled People: 0.14, High Skilled Information: 0.15, Low Skilled Information: 0.15. Attractiveness and College GPA are standardized to have a mean of zero and standard deviation of one. Standard errors are clustered by individual. All regressions include a constant and age.

Table 9: Log-Wage Regressions for Women: Jobs Above 75th Percentile Tasks Usage in Each Category

Primary Job Task	(A) People		(B) Information	
	High (1)	Low (2)	High (3)	Low (4)
Attractiveness (standardized)	0.101 (0.030)	0.098 (0.035)	0.016 (0.033)	-0.009 (0.031)
College GPA (standardized)	0.060 (0.026)	0.009 (0.029)	0.069 (0.024)	0.021 (0.030)
R ²	0.072	0.058	0.079	0.023
Observations	525	523	550	507

Notes: Observations are assigned to each task category if task usage is above the 75th percentile in each category. The 75th percentiles for each task are as follows: High Skilled People: 0.38, Low Skilled People: 0.36, High Skilled Information: 0.24, Low Skilled Information: 0.23. Attractiveness and College GPA are standardized to have a mean of zero and standard deviation of one. Standard errors are clustered by individual. All regressions include a constant and age.

References

- Agopsowicz, Andrew, Chris Robinson, Ralph Stinebrickner, and Todd Stinebrickner.** 2017. "Careers and Mismatch for College Graduates: College and Non-college Jobs." *Working Paper*.
- Altonji, Joseph G, and Charles R Pierret.** 2001. "Employer learning and statistical discrimination." *The Quarterly Journal of Economics*, 116(1): 313–350.
- Arcidiacono, Peter.** 2004. "Ability sorting and the returns to college major." *Journal of Econometrics*, 121(1-2): 343–375.
- Autor, David H, and Michael J Handel.** 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics*, 31(2): S59–S95.
- Becker, Gary S.** 1957. *The economics of discrimination*. University of Chicago Press.
- Bertrand, Marianne, and Sendhil Mullainathan.** 2004. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *The American Economic Review*, 94(4): 991–1013.
- Bertrand, Marianne, Claudia Goldin, and Lawrence F Katz.** 2010. "Dynamics of the gender gap for young professionals in the financial and corporate sectors." *American Economic Journal: Applied Economics*, 2(3): 228–255.
- Biddle, Jeff E., and Daniel S. Hamermesh.** 1998. "Beauty, Productivity, and Discrimination: Lawyers' Looks and Lucre." *Journal of Labor Economics*, 16(1): 172–201.
- Black, Dan A.** 1995. "Discrimination in an equilibrium search model." *Journal of Labor Economics*, 13(2): 309–334.
- Camargo, Braz, Ralph Stinebrickner, and Todd Stinebrickner.** 2010. "Interracial friendships in college." *Journal of Labor Economics*, 28(4): 861–892.

- Charles, Kerwin Kofi, and Jonathan Guryan.** 2008. "Prejudice and wages: an empirical assessment of Becker's The Economics of Discrimination." *Journal of Political Economy*, 116(5): 773–809.
- Charles, K Kerwin, and Jonathan Guryan.** 2009. "Taste-based discrimination." *The New Palgrave Dictionary of Economics*.
- Deryugina, Tatyana, and Olga Shurchkov.** 2015. "Now you see it, now you don't: The vanishing beauty premium." *Journal of Economic Behavior & Organization*, 116: 331–345.
- Farber, Henry S, and Robert Gibbons.** 1996. "Learning and Wage Dynamics." *The Quarterly Journal of Economics*, 111(4): 1007–1047.
- Feingold, Alan.** 1992. "Good-looking people are not what we think." *Psychological bulletin*, 111(2): 304.
- Gathmann, Christina, and Uta Schönberg.** 2010. "How General Is Human Capital? A Task-Based Approach." *Journal of Labor Economics*, 28(1): 1–49.
- Hamermesh, Daniel S.** 2011. *Beauty pays: Why attractive people are more successful*. Princeton University Press.
- Hamermesh, Daniel S, and Amy Parker.** 2005. "Beauty in the classroom: Instructors pulchritude and putative pedagogical productivity." *Economics of Education Review*, 24(4): 369–376.
- Hamermesh, Daniel S., and Jeff E. Biddle.** 1994. "Beauty and the labor market." *The American Economic Review*, 84(5): 1174–1194.
- Harper, Barry.** 2000. "Beauty, stature and the labour market: a British cohort study." *Oxford Bulletin of Economics and Statistics*, 62(s1): 771–800.
- Heckman, James J.** 1998. "Detecting Discrimination." *Journal of Economic Perspectives*, 12(2): 101–116.
- Lise, Jeremy, and Fabien Postel-Vinay.** 2015. "Multidimensional Skills, Sorting, and Human Capital Accumulation." *Working Paper*.
- Mankiw, Gregory.** 2015. *Principles of Microeconomics*. Vol. 10, Cengage Learning.
- Mobius, Markus M, and Tanya S Rosenblat.** 2006. "Why beauty matters." *The American Economic Review*, 96(1): 222–235.
- Mocan, Naci, and Erdal Tekin.** 2010. "Ugly Criminals." *The Review of Economics and Statistics*, 92(1): 15–30.
- Neal, Derek A., and William R. Johnson.** 1996. "The Role of Premarket Factors in Black-White Wage Differences." *Journal of Political Economy*, 104(5): 869.
- Persico, Nicola, Andrew Postlewaite, and Dan Silverman.** 2004. "The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height." *Journal of Political Economy*, 112(5): 1019–1053.
- Poletaev, Maxim, and Chris Robinson.** 2008. "Human capital specificity: evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000." *Journal of Labor Economics*, 26(3): 387–420.

- Robinson, Chris.** 2011. "Occupational mobility, occupation distance and specific human capital." *CIBC Working Paper*.
- Robins, Philip K, Jenny F Homer, and Michael T French.** 2011. "Beauty and the labor market: accounting for the additional effects of personality and grooming." *Labour*, 25(2): 228–251.
- Roys, Nicolas, and Christopher Taber.** 2017. "Skills Prices, Occupations and Changes in the Wage Structure." *Working Paper*.
- Sanders, Carl.** 2017. "Skill Accumulation, Skill Uncertainty, and Occupational Choice." *Working Paper*, 1–44.
- Sanders, Carl, and Christopher Taber.** 2012. "Life-Cycle Wage Growth and Heterogeneous Human Capital." *Annual Review of Economics*, 4(1): 399–425.
- Scholz, John, and Kamil Sicinski.** 2015. "Facial Attractiveness and Lifetime Earnings: Evidence from a Cohort Study." *Review of Economics and Statistics*, 97(1): 14–28.
- Speer, Jamin D.** 2017. "Pre-Market Skills, Occupational Choice, and Career Progression." *Journal of Human Resources*, 52(1): 187–246.
- Stinebrickner, Ralph, and Todd R Stinebrickner.** 2003a. "Understanding educational outcomes of students from low-income families evidence from a liberal arts college with a full tuition subsidy program." *Journal of Human Resources*, 38(3): 591–617.
- Stinebrickner, Ralph, and Todd R Stinebrickner.** 2003b. "Working during school and academic performance." *Journal of Labor Economics*, 21(2): 473–491.
- Stinebrickner, Ralph, and Todd R Stinebrickner.** 2004. "Time-use and college outcomes." *Journal of Econometrics*, 121(1): 243–269.
- Stinebrickner, Ralph, and Todd R Stinebrickner.** 2008a. "The causal effect of studying on academic performance." *The BE Journal of Economic Analysis & Policy (Frontiers)*, 8(1): 1–42.
- Stinebrickner, Ralph, and Todd Stinebrickner.** 2006. "What can be learned about peer effects using college roommates? Evidence from new survey data and students from disadvantaged backgrounds." *Journal of public Economics*, 90(8): 1435–1454.
- Stinebrickner, Ralph, and Todd Stinebrickner.** 2008b. "The Effect of Credit Constraints on the College Drop-Out Decision: A Direct Approach Using a New Panel Study." *American Economic Review*, 98(5): 2163–84.
- Stinebrickner, Ralph, and Todd Stinebrickner.** 2014a. "Academic performance and college dropout: Using longitudinal expectations data to estimate a learning model." *Journal of Labor Economics*, 32(3): 601–644.
- Stinebrickner, Ralph, and Todd Stinebrickner.** 2014b. "A major in science? Initial beliefs and final outcomes for college major and dropout." *The Review of Economic Studies*, 81(1): 426–472.
- Stinebrickner, Todd, and Ralph Stinebrickner.** 2012. "Learning about Academic Ability and the College Dropout Decision." *Journal of Labor Economics*, 30(4): 707–748.
- Stinebrickner, Todd, Ralph Stinebrickner, and Paul Sullivan.** 2017. "Dynamics of the Gender Wage Gap Among College Graduates." *Working Paper*.

Stinebrickner, Todd, Ralph Stinebrickner, and Paul Sullivan. 2018. "Job Tasks, Time Allocation, and Wages." *Journal of Labor Economics, Forthcoming.*

Yamaguchi, Shintaro. 2012. "Tasks and Heterogeneous Human Capital." *Journal of Labor Economics*, 30(1): 1–53.

Yamaguchi, Shintaro. 2018. "Changes in returns to task-specific skills and gender wage gap." *Journal of Human Resources, Forthcoming.*