Working Draft

Assessing the Causal Effect of Gendered Market Structures on Wage Attainments and in Evaluating the Gender Wage Gap: An Intention-to-Treat Analysis*

Scott R. Eliason**
University of Minnesota

Melissa Bonstead-Bruns***
University of Wisconsin – Eau Claire

Jennifer Lee**
University of Minnesota

March 1, 2006

*Please do not quote without permission. Prepared for the 2006 Population Association of American meetings. Please direct all correspondence to Scott R. Eliason, eliason@atlas.socsci.umn.edu, Department of Sociology, University of Minnesota, 267 19th Avenue South, Minneapolis, Minnesota 55455. The authors have benefited greatly from discussions with Robin Stryker and Cecilia Ridgeway on various aspects of the ideas presented in this paper.

** Department of Sociology, University of Minnesota.

*** Department of Sociology, University of Wisconsin – Eau Claire.

Assessing the Causal Effect of Gendered Market Structures on Wage Attainments and in Evaluating the Gender Wage Gap: An Intention-to-Treat Analysis

Abstract

Previous research has documented a persistent gender wage gap. On average, women earn less than men. Empirical evidence also suggests that gendered market structures play an important role in creating this gap. Recent research, however, has cast doubt on this conclusion, suggesting that the effects of gender structures on wages dissolve with properly specified models. In this paper we seek to inform this debate by using Imbens and Rubin's (1997) and Hirano, Imbens, Rubin, and Zhou's (2000) formulation of the intention-to-treat design to assess causal effects. By considering gender-of-job as treatment and gender-of-person as the mechanism encouraging respondents to obtain gendered jobs, we obtain the compliers average causal effect of gender-of-job on wages, and isolate that portion of the gender wage gap due to gender-of-job. Using data from 1988 and 2000 Current Population Surveys, we show that a substantial amount of the observed gender wage gap can be causally attributed to gendered market structures.

Assessing the Causal Effect of Gendered Market Structures on Wage Attainments and in Evaluating the Gender Wage Gap: An Intention-to-Treat Analysis

Previous research on gender wage inequality has clearly documented the persistence of a gender wage gap. On average, women earn less than men. Moreover, empirical evidence amassed in the last two decades give substantial support to the idea that gendered market structures play an important, and perhaps causal, role in creating and maintaining this gap (e.g., England 1984, 1992; Reskin and Ross 1992; Kilbourne, England, Farkas, Beron, and Weir 1994; England, Herbert, Kilbourne, Reid and Megdal 1994). Subsequent research, however, has cast doubt on this conclusion, suggesting that the effects of these gender structures on wages dissolve once researchers properly account for vocational preparation. (See Tam 1997, and the interchange between England, Hermsen, and Cotter 2000 and Tam 2000.)

We seek to inform this debate making use of Imbens and Rubin's (1997) and Hirano, Imbens, Rubin, and Zhou's (2000) formulation of the intention-to-treat design to assess causal effects from non-randomized data. In the general intention-to-treat design, causal effects due to some treatment are obtained in part through a mechanism that randomly distributes an encouragement for subjects to obtain treatment. While this encouragement is randomly distributed, the distribution of treatment is subject to the degree and character of compliance to that encouragement, and is thus typically not randomly distributed.

In assessing the effects of gendered market structures on wages through an intention-to-treat lens, we show that by considering (a) gender-of-job as the treatment and (b) gender-of-person as the mechanism encouraging respondents to obtain gendered jobs, we obtain the compliers average causal effect of gender-of-job on market wages and thus isolate that portion of the observed gender wage gap causally due to gendered market structures. Using data from 1988 and 2000 Current Population Surveys, we show that a substantial amount of the observed gender wage gap can be causally attributed to gendered market structures. We further show how these results are dependent on only a handful of readily defensible assumptions about the mechanisms linking gendered persons to gendered jobs and are necessarily robust to the specification issues that have plagued this research question from the beginning.

Previous Research on Matching Individuals to Gendered Jobs

Empirical research on the market value of sex-typed skills has consistently shown that qualities commonly associated with women, such as nurturance, function to reduce market wages for both men and women in jobs that require those skills (e.g., Kilbourne, England, Farkas, Beron, and Weir 1994; England, Herbert, Kilbourne, Reid, and Megdal 1994; England 1992; Steinberg 1990).

Research has also consistently shown that hiring decisions are often based on stereotypes about which gender is best suited for which job (e.g., Glick, Zion, and Nelson 1988; Ridgeway 1997). In general, labor markets operate in such a way that an individual's gender becomes salient in matching workers to jobs. This operates through the relationship between gender role expectations and task performance expectations on the part of employers and employees (Ridgeway 1997). More specifically, institutionalized expectations about men's and women's abilities and talents, coupled with the perceived gendered demands of tasks, create a match between gender of the person, on the one hand, and gender of job tasks, on the other.

Importantly, this line of work has demonstrated that it is neither necessary, nor in all situations likely, that men and women will actually possess the skills that are commonly attributed to them. Rather, it is the *expectation* or *belief* that men and women possess certain skills which is enough to shape behavior and guide hiring decisions. These expectations or beliefs are very robust; an individual will likely disregard these expectations or beliefs only in the face of manifest countering evidence. See Ridgeway (1997) and Ridgeway and Smith-Lovin (1999) for relevant discussion and a more detailed theoretical treatment regarding gender stereotypes and market inequalities.

Thus, an individual's gender becomes a rather strong and robust signal directing both employers' and employees' to be predisposed to match women to jobs that emphasize traditionally female tasks (such as nurturance and an orientation to people) and men to jobs that emphasize traditionally male tasks (such as a drive for hierarchy and an orientation to things and data). Important to our current research, gender is randomly distributed across individuals and exogenous to the labor market. That is, whether someone is male or female is a random event and gender per se is not created by the labor market. These properties of an individual's gender – that it constitutes a strong signal matching individuals to gendered jobs, and that it is randomly distributed and exogenous to the labor market – make it ideal for the "intention" component of the intention-to-treat analysis described below.

Previous Attempts to Assess the Effect of Gendered Market Structures on the Gender Wage Gap

(Please note that this section is incomplete. We are aware of considerably more research that has been done in this area, which will be included in the final draft.)

A substantial portion of the history of modeling the gender wage gap, and the effects of gendered job structures on that gap, can be viewed as the pursuit for a proper model specification. The culmination of much of this pursuit can be found in an exchange between Tam (1997, 2000) and England, Hermsen, and Cotter (2000). But long before these more recent issues played out, researchers were already divided on how to best model this relationship.

One of the first attempts to model the effect of market structures on the gender wage gap stems back to almost forty years ago with research by Sanborn (1966) and Fuchs (1971), suggesting that the earnings differential between men and women is due to factors other than discrimination. Both researchers argued that evidence of wage differences between men and women within occupations is flawed because occupational categories are too broad. With narrower occupational categories, the female to male earnings ratios within each occupation would be closer to one. In addition, Sanborn (1966) found that differences in productivity and work experience also contributed to wage differentials between men and women.

Ferber and Lowry (1976) criticized this research, arguing that both Sanborn (1966) and Fuchs (1971) use the average earnings of *all* workers to estimate the potential wage gain if a woman were to move into a male occupation. Ferber and Lowry (1976) instead estimated separate equations for men and women, in addition to specifying models with interactions between sex, median education, and the proportion of males in the occupation. They found then that there are different returns to education for men and women, and that men benefited more from greater years of schooling. This could have been due to women being in occupations that did not require high levels of education. Using early understandings of regression standardization techniques, Ferber and Lowry (1976) examined this and found that if women were given the same occupational and educational distributions as men, part of the wage differential disappeared. In addition, they found that men lose some of their education advantage in occupations with a high proportion of women.

Snyder and Hudis (1979) responded to this work, arguing that Ferber and Lowry's (1976) models suffered from mutlicollinearity with the inclusion of non-significant interactions. (See also Snyder and Hudis [1976] for related work.) They also argued that the inclusion of only education, sex and gender composition ,and their interactions, resulted in omitted variable bias. In their reanalysis, they conclude that gender composition of occupations was indeed an important determinant of women's lower earnings, but less so than what Ferber and Lowry (1976) concluded.

Moving away from an aggregate analysis to an individual-level analysis, England et al (1988) extended our understanding of the role played by occupational sex segregation in the gender wage gap. Using fixed effects estimators with longitudinal data, and corrections for sample selectivity, England et al. (1988) found that, net of human capital, skill demands, and working conditions, individuals who work in female occupations earned less.

Suggesting that mechanisms in the market other than just occupational sex segregation play a role in the gender wage gap, Reskin (1988) showed that despite a decline in the index of occupational sex segregation of 10 percent between 1970 and 1980, the gap in wages between men and women would only decline by less than 2 percent. She argued that the basic cause of the gender wage gap was, and still is, men's propensities to maintain their privileges. Thus, even if occupational sex segregation were to be eliminated, other mechanisms would surely arise to maintain a wage gap.

In the 1990s, scholars such as England et al (1994), Steinberg (1990), and Ridgeway (1997), pointed to the gendered nature of tasks, and specifically the devaluation of those tasks expected to be performed by women, that in part contributes to the maintenance of lower wages in female occupations. These scholars provide compelling empirical and theoretical evidence that nurturant skills, like providing face-to-face service to others, are rewarded less than other social skills (e.g., authority). Tying the sex composition of occupations to gendered tasks, England et al (1994) show that, after controlling for a host of other occupational characteristics that would affect earnings – such as racial composition, cognitive skill, physical skill, non-monetary rewards, and industrial and organizational characteristics – sex composition still maintained a significant effect on wages.

All of this work was questioned when, in the late 1990s, Tam (1997) showed that yet another model specification yielded a null finding. Specifically, using 1988 Current Population Survey data, Tam (1997) showed that, after controlling for specialized training and industry, on average there is no wage discrimination against female occupations. His research instead showed that men are more likely than women to invest in firm specific training, and that this specialized training is positively compensated. England and her colleagues (England, Hermsen, and Cotter 2000) shot back, taking issue with Tam's model specification and arguing that his models suffered from omitted variable bias. By merely including one additional control, general educational development (GED), England et al (2000) showed how Tam's results would in fact reconfirm the conclusion that female occupations paid less than male occupations. Not surprisingly, Tam (2000) replied that England et al.'s (2000) analyses was incorrectly specified, biased due to measurement error. Once adjusting for measurement error and including GED in the analysis, Tam (2000) shows that the analysis in fact reinforce his earlier findings.

Intention-to-Treat Analysis of the Gender Wage Gap

The story of this past research – at least the lion's share of the empirical component of that story – is in effect a battle over model specification. Within the context of these regression-style models and analysis, this debate is doomed to be one without resolution. That is, there is currently no agreed upon model specification that satisfies both sides of this debate, and there is no indication that there will ever likely be one that both sides would agree upon.

Therefore, to inform this debate and to hopefully move it in a more fruitful direction, we use an intention-to-treat analysis of the gender wage gap, with data from the 1988 and 2000 February Current Population Surveys (CPS). We use the 1988 CPS to facilitate comparison with Tam's (1997) results on these data, and with the discussion and results found in Tam (2000) and England, Hermsen, and Cotter (2000). We use the 2000 February CPS for comparison with a more recent labor market context.

In general, the intention-to-treat analysis developed by Imbens and Rubin (1997) and further by Hirano, Imbens, Rubin, and Zhou's (2000) is designed to assess causal effects from non-randomized data. In the general intention-to-treat design, causal effects due to some treatment are obtained in part through a mechanism that randomly distributes an encouragement for subjects to obtain treatment. While this encouragement is randomly distributed, the distribution of treatment is subject to the degree and character of compliance to that encouragement, and is thus typically not randomly distributed.

In Imbens and Rubin's (1997) study, the authors examine the causal effect of vitamin supplements (the treatments) on children's survival rates in various Indonesian communities in the 1990s. In this case, the encouragement to treat was given by the imperfect assignment of children to receive the supplements, while the treatment itself was given by the actual receipt of the supplements. In the Hirano et al (2000) study, the effectiveness of a flu vaccine is studied, where there is imperfect compliance with the encouragement to be vaccinated

Here, we follow Hirano et al's (2000) specific formulation of the intention-to-treat analysis. However, rather than using the Bayes estimation strategy suggested by Hirano et al (2000), we instead base our analysis on the data likelihood to assess the causal effect of gendered job structures on wages. We prefer the data likelihood, rather than the Bayes analysis performed by Hirano et al, in our case because past research in assessing the gender gap has typically not used a Bayes analysis. Using a Bayes analysis at this point would introduce, through specification of some prior distribution, yet another source of difference between ours and past research on this problem. That is, we wish to keep the results free of any assumptions that a Bayes analysis would necessitate when imposing some prior distribution on the parameters in the data likelihood. Our analysis, therefore, lays the groundwork from which any subsequent Bayes analysis – with varying degrees of informative priors which, in turn, necessarily require more assumptions overlaid on the analysis than does one based on the data likelihood alone – may be compared.

As discussed above, we posit that gender-of-the-person operates as a randomly distributed societal-wide encouragement for individuals to be matched to gendered jobs. Gendered jobs are therefore considered the so-called treatment in our analysis. For our analysis, we make only the following assumptions.

1. Gender of the person (or person-gender) is randomly distributed and exogenous to the labor market.[1]

[1] Assumption 1 is equivalent to what Imbens and Rubin (1997) or Hirano et al (2000) call the ignorability-of-treatment-assignment assumption.

2. Mechanisms exist, though need not be observed (such as those discussed above), such that person-gender acts as an encouragement to be matched to gendered jobs.

3. There are tendencies for men and women to comply with these mechanisms and tendencies for men and women to not comply with these mechanisms.

4. The tendency to not comply with these mechanisms is considered a mixture of
    a. the tendency to always take a predominantly female-tasked job regardless of one's gender,
    b. the tendency to never take a predominantly female-tasked job regardless of one's gender, and
    c. the tendency to take no job.

No further assumptions – save for the parametric form of the log-likelihood (to be discussed below) – are necessary to generate what Hirano et al (2000) call the Intention-to-Treat (ITT) causal effect for compliers, which is equivalent in this case to what Imbens and Rubin (1997) term the Compliers Average Causal Effect (CACE) for the more general case.[2] Below we adopt Imbens and Rubin's (1997) terminology.

Importantly, to estimate the CACE we needn't impose on the analysis the usual exclusion restriction assumptions which are often a source of considerable concern in identifying certain causal effects from nonexperimental data. This is critical, as it allows the analysis, and thus us, to remain agnostic with respect to the possible (and, some would argue, likely) direct effects of a person's gender on wages for those whose tendencies to comply outweigh their tendencies to not comply with the encouragement mechanism described above. We discuss this in more detail below in describing the partitioning of the person-gender and job-gender effects on log-wages.

Additionally, to estimate the CACE we needn't impose on the analysis any further assumptions about model specifications for any regression functions on, in our case, log-wages (variables to be discussed below). This remarkable property of the CACE is critically important to the debate on the gender wage gap. As described above, much, if not the overwhelming majority, of the methodological content of that debate has centered around proper model

[2] Note that assumptions 4a and 4b encapsulate in our case what is more generally called the monotonicity assumption. See Imbens and Rubin (1997) or Hirano et al (2000) for details.

specifications to accurately estimate the effect of gendered structures and/or an individual's gender on wage attainments. This property of the CACE effectively leaves that portion of the debate without meaningful content in this context.

In general, the CACE identifies the expected counterfactual causal effect on some outcome were a complier to change treatment statuses.[3] In our case, the CACE identifies, for the population of compliers as described in assumptions 2 and 3 above, the counterfactual causal effect on wages due to a change in gendered job types as a function of an individual's gender. More precisely, the CACE from our analysis identifies the average wage penalty suffered by women precisely because (a) they are not male and thus (b) encouraged to attain female-tasked jobs. Another way to think about the CACE in this context is from the counterfactual standpoint of the expected wage gain a woman would have enjoyed had she been male and, thus, encouraged to attain a non-female-tasked job.

Maximum Likelihood Estimation of the CACE

To formalize ideas, let π_c be the probability of compliance with the person-gender mechanism in being matched to gendered jobs. That is, π_c gives the conditional probability that, because of their gender, women will be matched to predominantly female-tasked jobs and that men will be matched to predominantly non-female-tasked jobs. Let π_a and π_n be the probabilities of noncompliance with these mechanisms as given in 4a and 4b respectively. That is, π_a gives the probability of always taking a predominantly female-tasked job regardless of one's gender and π_n gives the probability of never taking a predominantly female-tasked job regardless of one's gender. We refer to these (unobserved) groups as always-takers and never-takers respectively. Finally, let π_u be the probability of noncompliance as given in 4c, the tendency to take no job.

[3] This is the effect that would have been observed in the population of compliers had the treatment been randomly assigned, as would be the case in a typical experimental design framework.

Next, define variables for the market outcomes, in this case the natural log of per-hour wages, such that Y_{1i} is the log-wages for individual *i* had s/he been employed in a female-tasked job and Y_{2i} is *i*'s log-wages had s/he been employed in a non-female-tasked job. Finally, define D as a trichotomous variable indexing gendered job types – 1. predominantly female-tasked jobs, 2. predominantly non-female-tasked jobs, and 3. no job – and Z as a dichotomous variable indexing gender of the person – 1. female and 2. male.

From this, the log-likelihood function described in Imbens and Rubin (1997) and Hirano et al (2000) used to estimate the CACE is given, in our case, by

$$\begin{aligned}\mathcal{L}_{\text{CACE}} = &\sum_{i\in(D=1,Z=1)} \ln\left\{\pi_{ci} f_{c1}\left\{Y_{1i}\right\} + \pi_{ai} f_{a1}\left\{Y_{1i}\right\}\right\} + \sum_{i\in(D=2,Z=2)} \ln\left\{\pi_{ci} f_{c2}\left\{Y_{2i}\right\} + \pi_{ni} f_{n2}\left\{Y_{2i}\right\}\right\} + \\ &\sum_{i\in(D=1,Z=2)} \ln\left\{\pi_{ai} f_{a2}\left\{Y_{1i}\right\}\right\} + \sum_{i\in(D=2,Z=1)} \ln\left\{\pi_{ni} f_{n1}\left\{Y_{2i}\right\}\right\} + \sum_{i\in(D=3)} \ln\left\{\pi_{ui}\right\}\end{aligned} \tag{1}$$

Here, $f_{jk}\left\{.\right\}$, with $j = c,a,n$ and $k = 1,2$, refers to the density functions for log-wages for compliers, always-takers, and never-takers for outcomes Y_{1i} and Y_{2i}. Note, further, that the two observed components of the log-likelihood consistent with compliance – women in female jobs $i \in (D=1, Z=1)$ and men in male jobs $i \in (D=2, Z=2)$ – are in fact mixtures of compliers and always-takers for $i \in (D=1, Z=1)$ and compliers and never-takers for $i \in (D=2, Z=2)$. Men in female jobs, $i \in (D=1, Z=2)$, derive from the population of always-takers; women in non-female jobs, $i \in (D=2, Z=1)$, derive from the population of never-takers. Finally, in addition to the typical likelihood function described by Imbens and Rubin (1997) and Hirano et al (2000) for this type of analysis, we include those not employed as censored. In our case it is important to include this information in the log-likelihood given that estimation of π_u tells us the probability of this non-compliance status, and can be compared to that for other groups.

From this log-likelihood function, the CACE can be defined as

$$\text{CACE} = \text{E}\{Y_{1i} - Y_{2i} \mid \text{ Compliance}\} = \int y f_{c1}\{Y_1\}\,dy - \int y f_{c2}\{Y_2\}\,dy \tag{2}$$

(See Imbens and Rubin's (1997) or Hirano et al (2000) for details.) As described above, this gives the average wage penalty, on the log scale, suffered by women precisely because they are not male and, thus, encouraged to attain female-tasked jobs. To retrieve the average wage penalty on the original dollar scale, calculate e^{CACE}.

We identify two additional quantities of interest for non-compliers. These are (1) the average person-gender effect for those who would always take a female job regardless of one's gender and (2) the average person-gender effect for those who would never take a female job regardless of one's gender. The first is given by

$$\text{APGE}_a \equiv \text{E}\left\{Y_{1i(a1)} - Y_{1i(a2)} \mid \text{ Always-Taker}\right\} = \int y f_{a1}\{Y_1\}\,dy - \int y f_{a2}\{Y_1\}\,dy \tag{3}$$

where $Y_{1i(a1)}$ and $Y_{1i(a2)}$ represents the log-wage outcome Y_{1i} for women and men always-takers respectively. The second quantity is given by

$$\text{APGE}_n \equiv \text{E}\left\{Y_{2i(n1)} - Y_{2i(n2)} \mid \text{ Never-Taker}\right\} = \int y f_{n1}\{Y_2\}\,dy - \int y f_{n2}\{Y_2\}\,dy \tag{4}$$

These two quantities give the effect of a person's gender on log-wages for those embedded in predominantly female-tasked jobs – as given by the APGE_a estimator – and for those embedded in predominantly non-female-tasked jobs – as given by the APGE_n estimator. As with the CACE, we obtain the effect on wages by exponentiating the APGE_a and APGE_n estimators. A comparison of all three quantities, therefore, reveals the relative influence of gendered market structures – as given by the CACE – and a worker's gender embedded in those market structures – as given by the APGE_a and APGE_n – on log-wages, with the exponentiated versions revealing the relative influences on wages.

Estimating the CACE and APGE from the 1988 and 2000 CPS

We use a bootstrapped estimator of the CACE, APGE_a and APGE_n, derived from the above log-likelihood on 100 replications, to examine these effects for the 1988 and 2000 Current Population Surveys, as described above.[4] Using bootstrapped estimators alleviates the need to impose any distributional assumptions on the CACE, APGE_a and APGE_n (e.g., normality). We thus present the empirical distribution function of the bootstrapped estimator for each of these quantities, delineating the inter 95% percentile band for each empirical distribution. The only distributional assumptions imposed on our estimation strategy is that log-wages are assumed derived from a normal distribution. That is, we impose normality on the $f_{jk}\{\cdot\}$ distributions in the log-likelihood function. Other distributions can easily be imposed, and the sensitivity of the CACE, APGE_a and APGE_n estimators can be examined under different distributional forms of the $f_{jk}\{\cdot\}$. We leave that to future work.

To be consistent with much of the past work on assessing the effects of gendered market structures on wages, we use here percentage female in an occupation as our measure of predominantly female-tasked jobs. Specifically, those occupations with 75% or more women, as indicated by distributions from the 1988 and 2000 CPS data, are considered to contain predominantly female-tasked jobs. While we recognize the limitations in this measure and that better measures are available that directly reflect the gendered nature of tasks necessary to a specific job, we chose the percent female measure for the current analysis to maintain closer

[4] Given the rather large sample sizes of the 1988 and 2000 CPS, 100 replications for the bootstrap estimator proved sufficient to obtain relatively tight 5th to 95th percentile bands in the empirical distribution.

comparability with the research presented by Tam (1997) and debated by England, Hermsen, and Cotter (2000) and Tam (2000). In subsequent research (Bonstead-Bruns, Eliason and Lee 2006) we examine the CACE, APGE_a and APGE_n under different measurement strategies.

Our analysis includes all of those aged 25-65 who are not enrolled in school full-time and not out of the labor force. Thus, the censored proportion π_u reflects the traditional notion of unemployment, that is, those not employed who are looking for work. The total sample size for the 1988 CPS is 17,426, with 4,329 and 413 women and men respectively employed in predominantly female-tasked jobs, 3,433 and 7,738 women and men respectively employed in predominantly non-female-tasked jobs, and 809 and 704 women and men not employed. The total sample size for the 2000 CPS is 8,005, with 1,851 and 216 women and men respectively employed in predominantly female-tasked jobs, 1,455 and 2,765 women and men respectively employed in predominantly non-female-tasked jobs, and 753 and 965 women and men not employed. Finally, sampling weights found in the CPS are used in the analysis to ensure inferences properly apply to the sampled population.[5]

1988 CPS Results

Figure 1 compares bootstrapped distributions of the CACE, APGE_a and APGE_n, exponentiated to reflect per-hour dollar amounts, for the 1988 CPS sample. The horizontal axis gives dollar amounts reflecting, respectively in the three panels in Figure 1, the exponentiated CACE, APGE_a and APGE_n.[6] The vertical axis gives the frequency of occurrence and the height of the bars in the graph represent the relative density of the exponentiated effects distributed across the full range of those effects. The smoothed curve superimposed on the distribution represents a reasonably well-fitting function to the histogram, and is included primarily as a visual aid to highlight the continuous character of these distribution.

[5] See documentation for the February 1988 and 2000 Current Population Surveys for details on sampling weights in the CPS.
[6] The dollar amounts reported on the axis represent equidistant markers for that continuous distribution.

As indicated by the 5th to 95th percentile range for the bootstrap estimated CACE (Figure 1, top panel), the compliance mechanism drawing women into predominantly female-tasked jobs resulted in a wage penalty for women of between -\$6.61 and -\$6.11 per hour in 1988. In other words, because of social and market mechanisms operating to match women to predominantly female-tasked jobs, women workers influenced by that constraint – that is, female compliers – lost on average between \$6.61 and \$6.11 dollars per hour in the market precisely because they were not male and, thus, encouraged to attain female-tasked jobs.

This estimate of the CACE is relevant to the degree of compliance with the social and market pressures matching women to predominantly female-tasked jobs and men to predominantly non-female-tasked jobs. An estimate of the probability of compliance with these mechanisms, $\hat{\pi}_c$, gives us some indication of the degree of compliance. For the 1988 CPS sample, $\hat{\pi}_c = 0.46$. There are two ways to interpret this estimate. The first is as the proportion of compliers in the population. That is, a $\hat{\pi}_c = 0.46$ would indicate 46% of the population are compliers with this mechanism. This appears to be the interpretation favored by Imbens and Rubin (1997).

The second is as a compliance rate, or the likelihood that, at any given moment, an individual would feel compelled to comply with this mechanism. This interpretation is consistent with the idea that all members of the population are susceptible to compliance at any given moment, and that no individual is inherently always a complier or not a complier. With this interpretation, the $\hat{\pi}_c = 0.46$ would indicate a 46% compliance rate with this mechanism over the entire population. This interpretation appears to resonate more with the gender mechanisms elaborated by Ridgeway (1997) and Ridgeway and Smith-Lovin (2000).

The bottom two panels in Figure 1 also shows, for the 1988 CPS data, the person-gender wage differences *not attributable* to these person-gender/job-gender matching mechanisms. For those with tendencies to always take a predominantly female-tasked job regardless of their gender, the distribution of the APGE_a bootstrap estimate provides evidence that women suffer a wage loss on average of between \$6.64 and \$5.16 dollars per hour. The rate of this tendency is given by $\hat{\pi}_a = 0.05$, which is considerably lower than the rate of compliance given above and indicates that this tendency is a rare event in the population.

Nevertheless, this result for the 1988 CPS reveals a compounded hardship for women in predominantly female-tasked jobs. That is, this evidence points to dual mechanisms in the 1988 labor market – one through the matching of women to predominantly female-tasked jobs and one through an individual's gender independent of the matching mechanism – that operate concurrently to significantly suppress women's wages in those jobs. Even if we take the most conservative lower bound on that estimate, and presume that only one mechanism can operate at any given time (which appears to us unlikely), for a full-time (35 hour per week) female worker this results in an average loss over the course of a 50 week work year of \$5.16 x 35 x 50 = \$9,030 per year in 1988 dollars.

The final panel in Figure 1 gives the distribution of the APGE_n estimate, the person-gender effect for those with tendencies to never take a predominantly female-tasked job regardless of their gender. The rate for this tendency is estimated at $\hat{\pi}_n = 0.40$, which is the second highest next to the tendency for compliance.[7] On average, those women who do not comply with the pressures given by the person-gender/job-gender matching mechanisms in society and instead always take a non-female-tasked job, enjoy a modest wage gain of about \$1.00 relative to their male counterparts. Importantly, this cannot be due to the nature of these non-female-tasked jobs themselves, as the APGE_n estimate compares across men and women who are found in non-female-tasked jobs only (as measured by less than 75% females in a given occupation).

[7] The three probabilities do not add to one because there is a fourth, the probability of being censored. For the 1988 CPS sample, the censored probability is estimated at $\hat{\pi}_u = 0.09$.

Further analysis of the characteristics of men and women who tend toward never taking a predominantly female-tasked job and the characteristics of jobs obtained by these men and women would prove valuable in informing this result. We leave this exploration for future work, modeling those factors that influence the probability of tendencies men and women.

2000 CPS Results

Figure 2 gives bootstrapped results for the exponentiated CACE, APGE_a and APGE_n estimates for the 2000 CPS sample. As indicated by the 5th to 95th percentile range for the exponentiated CACE (top panel), the compliance mechanism drawing women into predominantly female-tasked jobs in 2000 resulted in a lower wage penalty when compared to that for 1988. In 2000, these women on average suffered between -\$5.32 and -\$4.23 per hour. While this constitutes over a \$1.00 per hour gain when compared to 1988 women compliers, it is still a substantial wage penalty. Taking the median estimate of \$4.82, a full-time (35 hour per week) female complier working over the course of a 50 week work year can be expected to net a loss of \$4.82 x 35 x 50 = \$8,435 per year in 2000 dollars. That is, because of the social and market gender matching mechanisms in 2000, women workers lost on average \$8,435 precisely because they were not male and thus encouraged to attain female-tasked jobs. The estimated tendency for men and women to comply with these matching mechanisms in 2000 is $\hat{\pi}_c = 0.41$, down slightly from the estimate in 1988. Nevertheless, this still constitutes a large rate of compliance for women and men with these mechanisms.

While the CACE results are similar for 1988 and 2000, the APGE_a and APGE_n results tell a very different story for 2000 compared with 1988. Recall that these estimates reveal person-gender wage differences that cannot be attributed to the person-gender/job-gender matching mechanisms. For those tending toward always taking a predominantly female-tasked job regardless of gender, the distribution of the APGE_a estimate shows that these women in 2000 in fact enjoyed a significant wage gain relative to their male counterparts. This wage gain appears substantial, between \$9.68 and \$11.54 dollars per hour, for female always-takers. However, similar to the 1988 results, this is a rather rare event in the population, with an estimated rate of $\hat{\pi}_a = 0.06$.

Nevertheless, this result for the 2000 CPS is contrary to the 1988 results, revealing competing mechanisms for women in predominantly female-tasked jobs. On the one hand, those women with tendencies toward yielding to social and market pressures in taking a predominantly female-tasked job precisely because they are female suffer a wage penalty. On the other hand, those women with tendencies toward taking a predominantly female-tasked job regardless of those person-gender/job-gender matching pressures in the market enjoyed wage gains. Further investigation into those factors that influence women to give in to those pressures would reveal the alternate routes to either wage gains or wage penalties for women in predominantly female-tasked jobs in the 2000 job market. As indicated above, we explore this very question in subsequent work.

The final panel in Figure 2 gives the distribution of the APGE_n estimate, the person-gender effect for those with tendencies to never take a predominantly female-tasked job regardless of gender. This, too, reveals a difference between the 1988 and 2000 labor markets where gender differences are concerned, though more modest than those already discussed. On average, the 2000 CPS results reveal that those women who tend to never take a female-tasked job regardless of gender suffer a modest wage loss of about -\$1.87 relative to their male counterparts. The tendency for women and men to never take a female-tasked job is estimated at $\hat{\pi}_n = 0.52$, and is thus the modal tendency in the 2000 job market and up from the 40% estimate in 1988. As mentioned above, we leave the investigation into those factors influencing all of these different tendencies for future work.

Conclusions

In this paper we have introduced an intention-to-treat analysis into the debate on the causal effects of gendered market structures on the gender wage gap. This analysis, along with the CACE estimator, has numerous advantages over standard regression analyses typically used, and long debated, in assessing those effects. First, this analysis requires no assumptions about any regression specifications on the mean structures of, in our case, log-wages (or any market outcomes for that matter) in assessing the causal effects. This property of the CACE estimator is all the more important in this case given that much of the methodological content over the history of this substantive debate has focused on regression model specification and misspecification. As mentioned above, this property of the CACE leaves that portion of the debate without meaningful content.

Second, this analysis nicely partitions the gender wage gap into components due to gendered market structures and the gender of individuals embedded in those structures. The effects of gendered market structures are estimated by the CACE in Eq. 2., and are relevant to the degree that individuals tend toward compliance with those social and market pressures matching men and women to gendered jobs, as estimated by the $\hat{\pi}_c$. The person-gender effects can be found in the APGE_a and APGE_n, and these are relevant to the degree that men and women tend toward always or never taking predominantly female-tasked jobs, as estimated by the $\hat{\pi}_a$ and $\hat{\pi}_n$ respectively. In our CPS samples, we estimated compliance rates at 46% in 1988 and 41% in 2000, with median wage penalties for women of -\$6.36 and -\$4.82 respectively. This shows both the strength of these person-gender/job-gender matching mechanism in society, as well as the wage consequences for women.

Finally, our future work will seek to assess the impact that various factors – such as, for example, socioeconomic background, aspirations and expectations, human capital, and family configurations – have in influencing women and men to comply with these person-gender/job-gender matching mechanisms. Future work will also address the issue of the measurement of gendered tasks more directly. In this analysis we used a rather blunt instrument to measure predominantly female-tasked jobs. It would be useful to move beyond this measure to more precise measures in assessing, over different time periods and different labor market segments, both the compliance rates and the wage penalties (or gains) suffered by (or enjoyed by) women in the context of various female-typed job tasks.

References

(Incomplete in the current version.)

England, Paula, Joan M. Hermsen, and David A. Cotter. 2000. "The Devaluation of Women's Work: A Comment on Tam." *American Journal of Sociology* 105:1741-175.

England, Paula, George Farkas, Barbara Stanek Kilbourne and Thomas Dou. 1988. "Explaining Occupational Sex Segregation and Wages: Findings from a model with Fixed Effects," *American Sociological Review*, 53: 544-558.

England, Paula, Melissa S. Herbert, Barbara Stanek Kilbourne, Lori L. Ried, and Lori McCreary Megdal. 1994. "The Gendered Valuation of Occupations and Skill: Earnings in 1980 Census Occupations," *Social Forces*, 73: 65-100.

Ferber, Marianne A. and Helen M. Lowry. 1976. "The Sex Differential in Earnings: A Reappraisal" *Industrial and Labor Relations Review*, 29: 377-387.

Fuchs, Victor R. "Differences in Hourly Earnings Between Men and Women," *Monthly Labor Review*, 94: 9-15.

Hirano, Keisuke, Guido W. Imbens, Donald B. Rubin, and Xiao-Hua Zhou. 2000 "Assessing The Effect Of An Influenza Vaccine In An Encouragement Design." *Biostatistics* 1,1:69–88.

Hodson, Randy and Paula England. 1986. "Industrial Sector and Sex Differences in Earnings," *Industrial Relations*, 25: 16-32.

Imbens, Guido W. and Donald B. Rubin. 1997. "Bayesian Inference for Causal Effects in Randomized Experiments with Noncompliance." *The Annals of Statistics* 25:305-327.

Reskin, Barbara F. "Bringing the Men Back In: Sex Differentiation and the Devaluation of Women's Work," *Gender and Society*, 2: 58-81.

Ridgeway CL. 1997. "Interaction and the Conservation of Gender Inequality: Considering Employment." *American Sociological Review* 62:218-35.

Ridgeway, Cecilia L. and Lynn Smith-Lovin. 1999. "The Gender System and Interaction." *Annual Review of Sociology* 25:191-216.

Sanborn, Henry. 1966. "Pay Differences between Men and Women," Industrial and *Labor Relations Review*, 17: 534-550.

Snyder, David and Paula M. Hudis. 1976. "Occupational Income and The Effects of Minority Competition and Segregation: A Reanalysis and Some New Evidence" *American Sociological Review* 41:209-234.

Snyder, David and Paula M. Hudis. 1979. "The Sex Differential in Earnings: A Further Reappraisal: Comment," *Industrial and Labor Relations Review*, 32: 378-384.

Steinberg, R. 1990. “Social Construction of Skill: Gender, Power, and Comparable Worth.” *Work and Occupations* 17.

Tam, Tony. 2000. “Occupational Wage Inequality and Devaluation: A Cautionary Tale of Measurement Error” *The American Journal of Sociology*, 105: 1741-1760.

Tam, Tony. 1997. “Sex Segregation and Occupational Gender Inequality in the United States: Devaluation or Specialized Training?” *The American Journal of Sociology*, 102: 1652-1692.

Figures

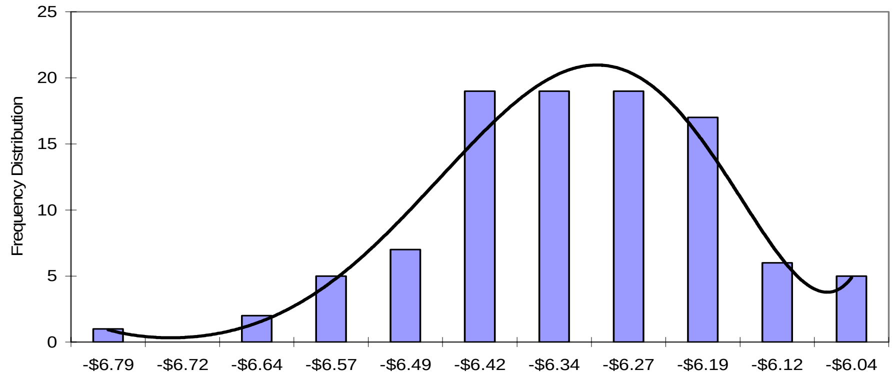

Bootstrapped Distribution of 1988 CPS Exponentiated CACE
Median = -$6.36; [5th, 95th] Percentiles = [-$6.61, -$6.11]

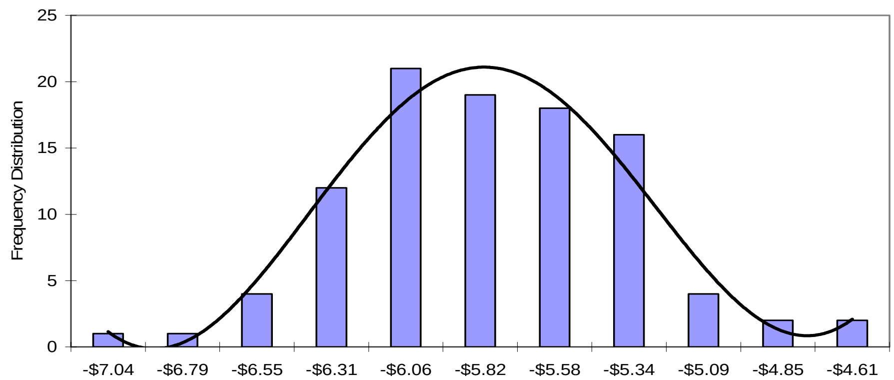

Bootstrapped Distribution of 1988 CPS Exponentiated $APGE_a$
Median = -$5.93; [5th, 95th] Percentiles = [-$6.64, -$5.16]

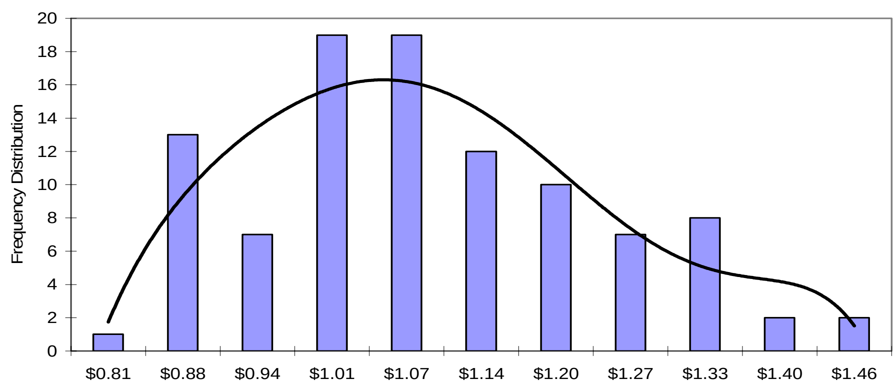

Bootstrapped Distribution of 1988 CPS Exponentiated $APGE_n$
Median = $1.04; [5th, 95th] Percentiles = [$0.85, $1.32]

Figure 1. Bootstrapped Distributions of CACE, $APGE_a$ and $APGE_n$ for the 1988 CPS sample.

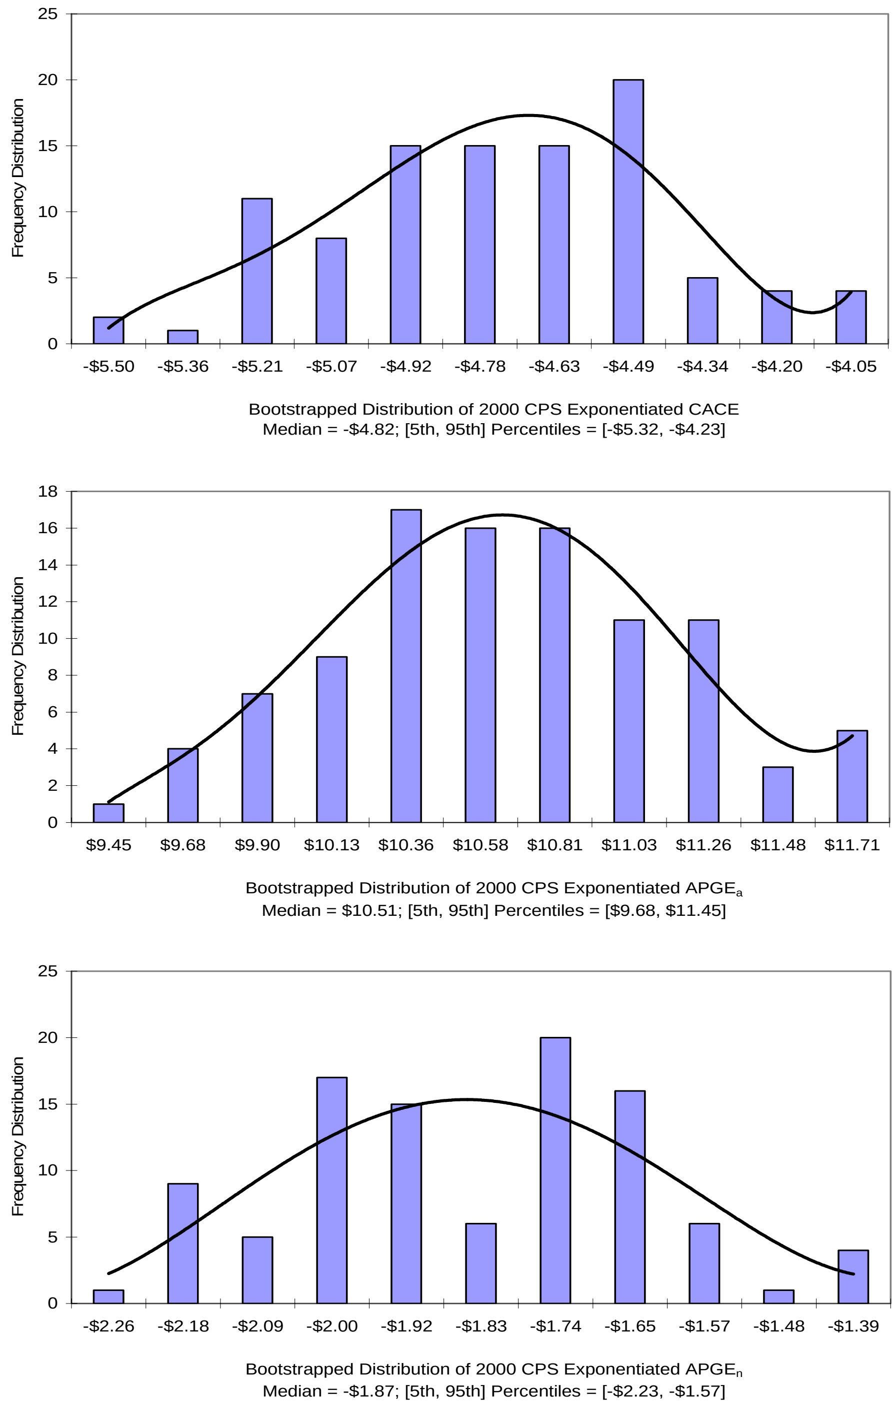

Figure 2. Bootstrapped Distributions of CACE, $APGE_a$ and $APGE_n$ for the 2000 CPS sample.