

WORKING PAPER · NO. 2018-38

The Glass Ceiling

Marianne Bertrand

June 2018

Marianne Bertrand

Booth School of Business, University of Chicago

October 27, 2017

Abstract

Despite decades of progress, women remain underrepresented in the upper part of the earnings distribution, a phenomenon often referred to as the “glass ceiling.” We review the recent research trying to explain this phenomenon. After briefly revisiting gender differences in education, we turn our attention to a body of work that has argued that gender differences in psychological attributes are holding back women’s earnings; we pay particular attention to the research that has aimed to test the relevance of these gender differences in psychological attributes in the field. We then review another active area of research that has returned to a more classical explanation focused on the challenges women may face when trying to juggle competing demands on their time in the workplace and in the home, particularly when the home includes children. We discuss recent work documenting women’s greater demand for flexibility in the workplace, as well work measuring the labor market penalties associated with such demand for flexibility, particularly in the higher paying occupations in the economy. We highlight possible countervailing forces (both at work and at home) that may explain why these work-family considerations may remain highly relevant to today’s glass ceiling despite reduced time spent in non-market work and a trend toward a more equal division of non-market work between the genders. Finally, we discuss the role that public policy and human resource practices may play in adding more cracks to the glass ceiling.

I. Introduction

Despite having made significant progress in the labor market, women remain as of today underrepresented in the upper part of the earnings distribution. This continued underrepresentation has been widely documented in some sectors of the economy, such as business and finance. For example, in a 2017 census of Fortune 500 companies in the U.S., Catalyst found that women hold only 19.9 percent of corporate board seats and comprise only 5.8 percent of CEO positions in those companies. In Europe, only 23.3 percent of board members of the largest publicly listed companies are women and only 5.1 percent are CEOs (European Union 2016), despite women accounting for 45 percent of the labor force. This phenomenon, which extends much beyond the corporate sector, is often referred to as the glass ceiling.

The Glass Ceiling In Numbers

Table 1 illustrates the extent of this underrepresentation in the US context, on which we will draw heavily throughout this paper, and how it has evolved over the last 40 years. The table reports women's labor force participation rates over time, as well as the share of working women with earnings that put them above various thresholds of the men's earnings distribution.¹ We define earnings as the sum of annual earnings from wage, business and farming. We restrict the analysis to those between 25 and 64 years of age. In particular, in each census year, we measure the 50th, 80th, and 90th percentiles of the distribution of earnings among men working full-time, full-year by education level (less than high school, high school degree, some college, college degree, post-graduate degree). We then compute, in each census year, the share of women whose earnings are above these percentiles. We do this computation both for all working women and for the subset of women who report working full-time, full-year.

¹Throughout this paper, we use 1970, 1980, 1990 and 2000 data from the US Census and 2008-2011 data from the American Community Survey (ACS) single-year files. For convenience, we refer to the 2008-2011 ACS as "2010" data.

Panel A of Table 1 considers all women in the relevant age group. As has been widely documented, women's participation in the labor force and likelihood to be employed full-time, full-year have been increasing over time, but also appear to have reached plateaus since the late 1990s: rapid improvements between 1970 and 1990 have turned into much slower progress since. Only 48 percent of women were in the labor force in 1970 compared to 69 percent in 1990, 70 percent in 2000 and 72 percent in 2010. Similarly, only 26 percent of women were working full-time, full-year in 1970, compared to 45 percent in 1990, 50 percent in 2000 and the same 50 percent in 2010.

A similar dynamic of rapid improvement to more muted progress can be seen when analyzing the share of women who have broken into the upper deciles of the similarly educated men's earnings distribution. In this case, the most rapid improvements occurred between 1980 and 2000, with the slowdown occurring over the last 10 years. Only 5 percent of working women in 1970 had earnings that put them above the median of the similarly educated men's earnings distribution in 1970, compared to 7 percent in 1980, 13 percent in 1990, 18 percent in 2000 and 19 percent in 2010. Only 1.2 percent of working women had earnings that put them above the 80th percentile of similarly educated men's earnings distribution in 1970, compared to 1.4 percent in 1980, 2.8 percent in 1990, 4.4 percent in 2000, and 5 percent in 2010. Similar dynamics can be observed in the subset of women working full-time, full-year.

Panel B focuses on women 25 to 64 years of age who have earned at least a college degree. This is the sample that is of most interest in any discussion of the glass ceiling, given that most top earners in the economy will have achieved such higher education credentials. The dynamics in this subsample are very comparable to those in Panel A. There was a rapid increase in labor force participation and likelihood of working full-time, full-year in this group until 1990, but there has been much slower progress since. The gains in the share of working women who have broken to the upper part of the men's earnings distribution were most pronounced between 1980 and 2000, but this trend appears to have somewhat weakened since. The glass ceiling in today's US context can be summarized as follows: only 25 percent of college-educated women working full-time, full-year have earnings above the median of similarly educated men working full-time, full-year; only 6 percent have earnings that put them in the top 20 percent of the men's distribution; and only 2.7 percent have earnings that put them in the top 10 percent of the men's distribution.

In a recent paper that relies on administrative tax data and is not subject to the limitations induced by the top coding of income in the Census, Piketty et al. (2016) track the shares of women in the top 10, top 1 and top 0.1 percent of the labor income distribution from 1962 to 2014. While all of these shares have been growing over time, women's underrepresentation at the very top of the income distribution remains extreme. While women accounted for about half of the employed population in 2014, they accounted for only 25 percent of the top 10 percent of earners in that year, only 15 percent of the top 1 percent of earners, and only 10 percent of the top 0.1 percent of earners. Moreover, Piketty et al.'s analysis of the tax data confirms that the slowdown in the most recent decade we observe in the Census data extends to the rate at which women are breaking into the group of very high earners. The steady gains recorded for women between the late 1970s to the late 1990s appear to have been followed by much more modest improvements since.

Finally, as most recently discussed by Goldin (2014), it is not the case that women's earnings deficit compared to men's is solely a reflection of differential occupational sorting. Women earn less than men within occupation as well. We illustrate this in the bottom rows of Table 1. In particular, we report in those rows the share of women working full-time, full-year, whose earnings are at or above the 50th and 80th percentiles of the earnings distribution of similarly educated men working full-time full-year in the *same occupation*.² Occupations are defined using the 1990 3-digit occupation code. One can see that, across all years, women are also underrepresented in the distribution of occupation-specific earnings. Moreover, the changes over time mimic those observed across occupation, with most rapid progress between 1980 and 2000 and a slowdown since then.

Why Should We Care?

Much of the discussion of the glass ceiling in the popular media and by advocacy groups is typically framed as an issue of rights and fairness. There should be "equal pay for equal work," advocates argue, and the observed gender gaps in earnings are enough of a proof to them that this is not happening. More often than not, this discussion assumes that gender discrimination in the labor market is the driving force behind the glass ceiling. There is no doubt such discrimination still exists today, with news of gender discrimination lawsuits and reports of experienced sexism

² Because of smaller sample sizes within occupation, we do not explore the 90th percentile threshold.

at the workplace regularly hitting the front page of the newspapers, and that is both unfair and illegal.

Much of the research we discuss below, however, which reflects current academic thinking on the glass ceiling, has been offering alternative explanations for the glass ceiling – explanations that do not rely on employers treating women unfairly, explanations that suggest that there might be *observed* gender gaps in earnings even when employers practice “equal pay for equal work.”

Even if there is disagreement about whether the glass ceiling raises concerns about employer discrimination and sexism in the workplace, all should agree that an economy that is tapping into a limited pool (men) to find its leaders must be operating inside the efficiency frontier. Starting from the position that innate talent is equally distributed between men and women, it must be the case that superior economic outcomes would be achieved if women had the same odds as men to make it to the top of the earnings distribution. Hsieh et al. (2017) performed an exercise to quantify how much economic growth is being “left on the table” because of the under-representation of women and minorities in some occupations. They estimate that roughly one-quarter of growth in US GDP per person between 1960 and 2010 can be explained by declining barriers to the entry of white women, black men and black women in occupations where they were previously heavily underrepresented.

Finally, another efficiency-based argument as to why one should care about the glass ceiling is related to how diversity in leadership roles might be productivity-enhancing. This argument has penetrated many organizations that are making the business case for “diversity and inclusion.” This argument was also one of the main motivations behind the introductions of quotas in the corporate sector in various European countries over the last decade. While economic research has fallen short so far of providing robust empirical demonstrations of the economic benefits of diversity and has been more careful to stress the theoretical ambiguities at the core of the argument, both the human resources and management literatures have strongly adopted the view that more diverse organizations will achieve superior outcomes.

Outline of Paper

Having argued that the glass ceiling is a real and concerning phenomenon, we devote the rest of this paper to discussing how the recent literature attempts to explain this phenomenon. We briefly revisit gender differences in education as a potential explanation, arguing that we might have been too quick in dismissing education as a factor because of the reversal of the gender gap in number of years of completed schooling. Second, we turn our attention to a large body of work that has argued that gender differences in some psychological attributes hold the keys to understanding what is holding back women's earnings, paying particular attention to the more recent work that has taken these hypotheses out of the laboratory and aimed to test their relevance in real-world contexts. Third, we review another active area of research that has returned to a more standard explanation centered on the particular challenges women may face in the labor market because of the relatively higher demand on their time outside the labor market (such as in childcare and other forms of non-market work) and has shown this more standard explanation might be particularly relevant for women who are vying for careers in the highest-paid occupations. We highlight a set of countervailing forces that may explain why these work-family tradeoffs may remain highly relevant to today's glass ceiling despite reduced time spent in non-market work and a trend toward a more equal division of non-market work between the genders. Finally, and in light of this acquired understanding of what is still holding back women's earnings, we discuss whether and how public policy and human resource practices can be leveraged to add more cracks to the glass ceiling.

II. Gender Gap in Education

As has been widely documented, women have overtaken men when it comes to completed years of schooling. This phenomenon has been documented throughout most of the developed world. Again, we use the US context for quantification. For example, Goldin et al. (2006) reported on college graduation rates for men and women born in the US between 1876 and 1975. While women lagged behind men until the mid-1950s birth cohorts, there has since been a reversal of the education gap. This reversal has been particularly salient for the cohorts born after the late 1960s, with men's college graduation rates having reached an apparent plateau at about 30 percent while women's college graduation rates continued on a steep upward trajectory until the 1975 birth cohort.

Figure 1 updates this analysis to include more recent birth cohorts. In particular, we use data from the American Community Survey Data from 2012 to 2015 (collated one-year files) and focus on individuals who were born between 1950 and 1985. For each birth cohort, we report college graduation rates (by 30 years of age). It is apparent from Figure 1 that the trend documented by Goldin et al. (2006) has persisted over the more recent birth cohorts for which we have complete schooling information. In particular, men's college graduation rates have remained pretty much unchanged since the 1970 birth cohort (about 30 percent) while women's college graduation rates have kept going up, reaching about 40 percent in the most recent cohort (1985).

There might however be a tendency, based on the figure above, to overstate how large women's educational advantage is in today's labor market without paying closer attention to the particular degrees men and women pursue in college and after college. While there has been convergence on this dimension as well, with women entering "men's fields" at a greater rate, it is not the case that women have overtaken men in their odds of completing the type of educational degrees that provide the highest labor returns.

Figure 2 illustrates this point. We start from the same research sample as in Figure 1 but restrict it to those that have completed at least a four-year college degree by age 30. We proxy for the earnings potential of a given individual based on his or her educational attainment as follows. For each combination of highest degree type (bachelor's degree, master's degree, professional degree, doctorate) and field of study (economics, English literature, etc.), we compute mean earnings and 80th and 90th percentile earnings among men working full time who have completed that degree-field of study combination. We then report, by birth cohort, the gender gap (men-women) in such education-based earnings potential.

The figure shows that, in the oldest birth cohort of college educated men and women (the 1950 birth cohort), women chose degrees and fields of studies with mean earnings that were about 14 percent below those in the degrees and fields chosen by men. Similarly, the figure shows that, among this oldest birth cohort of college-educated men and women, women chose degrees and fields of concentration where the 90th percentile of earnings was about 22 percent below those in the degrees and fields chosen by men. While there is clear convergence until the birth cohorts born in the late 1960s in educational choice, these gaps have remained roughly unchanged since. In

particular, women born in 1985 chose educational degrees that mapped into degrees and fields with about 6 percent lower average earnings than men and 10 percent lower 90th percentile earnings.

Thus, while women's gains over men in terms of number of years of completing schooling have been growing bigger and bigger over time, this masks what has been a somewhat more disappointing story when it comes to women's rate of entry into those educational tracks that are associated with the best labor market prospects (as measured by earnings). Hence, while much of the discussion about what is still holding women's earnings back in the labor market often assumes that education is no longer a relevant factor, this would be the wrong conclusion. More needs to be understood about what still stops women from entering the fields of study that would allow them to compete for top incomes. As we will argue below, a woman's decision to stay away from particular educational tracks might be in part an informed decision in light of the constraints and challenges women expect in the jobs that are associated with those educational tracks. But additional research should be devoted to understanding the full set of considerations that go into the educational choices today's young women make.

III. Gender Differences in Psychological Attributes

One of the most active areas of research over the last 15 years or so when it comes to explaining remaining gender gaps in labor market outcomes, and especially women's underrepresentation in the top echelons of the earnings distribution, has highlighted the role of gender differences in psychological attributes. In particular, there is a large amount of laboratory-based research documenting gender differences in attitudes toward risk and competition as well as attitudes toward negotiation. Much of this literature is discussed in Bertrand (2011) as well as Croson and Gneezy (2009).

In particular, in reviews of earlier work, both Croson and Gneezy (2009) and Eckel and Grossman (2008) come to the conclusion that the published experimental findings are broadly consistent with the view that women are more risk-averse than men. For example, in one of the first papers on this topic, Eckel and Grossman (2002) asked subjects to choose from five alternative gambles that differed in expected return and variance and paid subjects according to the outcome of the gamble that they

choose. They found that men on average choose riskier gambles with higher expected payoffs. Also, Dohmen et al. (2011) show evidence of higher risk aversion among women in the general population. Their study relies on both a large representative survey of the German population and a complementary experiment (choices in real-stake lotteries) carried on a representative subsample that validates the subjective self-assessment question in the survey data. In their large representative sample of the German population, Dohmen et al. (2011) find that gender has a quantitatively significant effect on one's self-assessed willingness to take risk: the gender effect corresponds to about a quarter of a standard deviation reduction in the willingness to take risk.

Such gender differences in risk aversion have the potential to explain why women are underrepresented in top earnings occupations because, as shown among others by Dohmen et al. (2011), there is a positive relationship between average occupational earnings and occupational earnings volatility. In other words, maybe women stay away from occupations or jobs that pay more on average (and stay away from educational tracks that lead into occupations or jobs that pay more on average) because they dislike more than men the associated greater earnings risks.

Many high-profile, high-earning occupations often take place in highly competitive settings where winners and losers are singled out and winners are disproportionately rewarded. A few recent experimental papers have proposed a new explanation for why women may be relatively underrepresented in these "winner-take-all" or "winner-take-most" occupations. These papers suggest that women may systematically underperform relative to men in competitive environments and that many women, even among the most able, may simply prefer to stay away from such environments.

In a very influential study, Gneezy et al. (2003) bring students to the lab in groups of six. Each student is asked to solve mazes for a period of 15 minutes under one of two possible compensation schemes: a piece rate scheme, or a tournament scheme. Under the piece rate scheme, students are paid a fixed prize for each maze that they solve; under the tournament scheme, only the student in the group that solves the highest number of mazes receives compensation. While there was no gender difference in performance under the piece rate scheme, the men strongly increased their performance in the tournament setting while the women did not. The gender gap in performance in the tournament setting was large, with men solving about 40 percent more mazes than women.

Gneezy et al. (2003) further show that women do as well as men in the tournament setting if the groups are single-sex; hence, the authors attribute the gender gap in the tournament setting to women's relative failure to perform at a high level when competing against men, but not when competing in general.

Niederle and Vesterlund (2007) push this research agenda further by studying the compensation choices men and women make in a mixed-sex environment (groups of 2 men and 2 women). As in the previously discussed paper, the compensation schemes under consideration are a piece rate scheme and a tournament-like winner-take-all scheme. The task in this case consisted of solving a series of addition problems. The experiment takes place in three rounds. The first two rounds are used to assess gender differences in performance in this task under either the piece rate setting or the tournament scheme. In neither case (and in contrast with Gneezy et al. (2003)) did the authors observe gender differences in performance in the experimental task. At the end of each of these two rounds, the participants are informed about their own performance but provided no information about their relative performance. In the third round, participants get to choose which compensation scheme they would prefer for their performance in that round. Despite the lack of gender differences in performance in the first two rounds, Niederle and Vesterlund (2007) find that close to three quarters of the men, but only one third of the women, choose the tournament scheme. Most strikingly, even the women who perform in the top performance quartile in the first rounds of the experiment are less likely to choose tournament compensation than the men who performed in the lowest quartile. From a payoff maximization perspective, there are too few (high ability) women and too many (low ability) men entering the tournament. Niederle and Vesterlund (2007) show that a residual gender gap in "willingness to compete" remains after accounting for gender differences in risk aversion, overconfidence and negative feedback aversion. The residual gender gap, Niederle and Vesterlund (2007) concluded, is best interpreted as women having less of a taste for competition.

While there has been much follow-up work on these two seminal studies, sometimes showing their sensitivity (such as to the stakes or to the experimental task) and sometimes questioning their interpretation (such as revisiting the role of risk aversion as the driver), the influence of this work has been large and, inspired by these provocative findings in the laboratory, a growing body of work has aimed to test the relevance of these factors in real world settings. In particular, a few

recent papers have studied how the gender gap in educational choices, job entry decisions and labor market earnings might be affected by gender differences in attitudes toward risk and competition.

Buser et al. (2014) study Dutch secondary school students who are enrolled in the pre-university track and the decision these students make, at the end of 9th grade, to choose between one of four study tracks. In the Dutch context, there is a clear ranking between the tracks in how prestigious and challenging they are, with the science-heavy Nature & Technology NT track being perceived as the most challenging and the humanities-oriented Culture & Society (CS) track as the least challenging. Buser et al. (2014) measured competitiveness at the student level using the lab-based tournament entry task from Niederle and Vesterlund (2007). They then correlated study track choice to this competitiveness measure. They find that, despite the fact that the girls are as good at math as boys and have higher GPAs, 17 percent of the girls, compared to 40 percent of the boys, choose the most prestigious NT track while 15 percent of the girls, compared to only 8 percent of the boys, choose the least prestigious CS track. Most interestingly, Buser et al. (2014) show that the differences in lab-based measures of competitiveness explain up to 23 percent of the gender gap in this educational track choice.

Moving to the job entry context, Flory et al. (2015) conduct a field experiment where they randomized job seekers who had expressed interest in a position into one of six different compensation regime treatments. The first two treatments involved a fixed wage per hour and differ only in whether they advertise that the job is to be done in a team or solo. The third and fourth treatments described a job where the worker will be matched with another person also hired for the position and where pay will consist of both a fixed hourly wage as well as a bonus if the applicant outperforms his or her match. The fifth treatment described a job that involved a team-based tournament: the work will take place in a team of 2 and the employee will be paid a fixed hourly wage as well as a bonus if his or her team outperforms another team he or she is matched to. Finally, the sixth treatment described a job where earnings are more uncertain. Importantly, the authors conduct this evaluation for two different versions of the job: one that emphasizes “maleness” (a sports-related job) and the other does not (a general job). The authors then study, across the six treatment arms and two job types, the probability of receiving an application by gender. Consistent with the lab-based studies, the authors find that female job seekers appear

especially turned off (e.g. less likely to apply) when the job advertisement mentions that they will compete with another worker in order to receive their bonus (third and fourth treatments). There is also some evidence that female job seekers are more turned off than male job seekers by a compensation structure that leaves the condition under which a bonus will be paid more uncertain (sixth treatment), which is consistent with the view that higher levels of risk aversion among women may get them to shy away from jobs with greater earnings volatility. Most interesting, though, is what the authors find when they compare the impact of the different treatments on the number of received job applications across the two job types. In a nutshell, the authors find that women's lower likelihood to apply when the compensation structure involves competing with a co-worker or has more uncertain earnings is very much concentrated in the sports-related job posting, e.g. when the job may be construed as more "male." The gender differences in application rates in the more tournament-like compensation regime do not extend to the more general job type. Hence, while attitudes toward competition and risk may affect patterns of job entry, the paper suggests that these gender differences in psychological attributes might not be fixed traits, but rather are more likely to manifest themselves in tasks or activities that are perceived as more "male."

Finally, Reuben et al. (2015) assess the relevance of various psychological attributes (including willingness to compete) for the labor market outcomes of recent MBA graduates. This study's population is of course most relevant to the "glass ceiling" discussion, in that MBA graduates are overly represented in the group of top earners in the economy. In particular, the research sample is a group of about 300 male and 125 female students from the University of Chicago Booth School of Business. The authors found that male students earn more in their first job after graduation than female students do. They also show that there is a correlation between a student's willingness to compete, as measured in a lab-based experiment à la Niederle and Vesterlund (2007), and this student's earnings in his or her job after graduation. There is also, even in this highly self-selected group of women who have decided to enter business school and hence have expressed interest in obtaining a job in the competitive business sector, evidence that female students are on average less willing to compete than male students. The difference is however quantitatively not very large so that, further controlling for willingness to compete in a regression of first job earnings on a female dummy does not much reduce (in absolute value) the estimated negative and significant

coefficient on the female dummy. Only about 10 percent of the gender gap in compensation in the first job after graduation can be accounted for by differences in willingness to compete.

Blau and Kahn (2017) perform a more systematic review of the studies, aiming to assess the role of psychological traits in accounting for the gender gap in pay. While the Reuben et al. (2015) study discussed above is most relevant to our topic given its focus on students entering the financial and corporate sectors of the economy and on psychological differences in attitudes toward competition, other studies have considered the gender pay gap in broader sections of the economy and have assessed how competitiveness and also other psychological traits such as self-esteem and self-confidence may account for the gender pay gap. One caveat is that many of these studies rely on already pre-existing datasets that were not specifically designed for this purpose and hence have clear limitations in the quality and comprehensiveness of the measurement of psychological traits. Overall, this work does tend to find that accounting for psychological traits does reduce the gender pay gap, but the magnitudes are not large. Across the reviewed studies, psychological traits do not appear to account for much more than 10 percent of the gender gap in earnings.

In summary, there is growing evidence from the field that the gender gaps in some key psychological traits diagnosed in the laboratory are relevant to the well-documented gender gaps in educational choices, job choices, and earnings. However, the quantitative importance of these psychological attributes in explaining the gender pay gap and the underrepresentation of women in top earnings positions remains debatable in light of the existing research. The work by Flory et al. (2015) also very much “highlights” how fragile the role of these attributes might be to the specific domain (e.g. “male job” or “general job”) and suggests that, though we should continue to research the role of these psychological factors in explaining women’s relative underperformance in the labor market, it would be a mistake to view these traits as hardwired. Nurture rather than nature may be responsible for women’s lower willingness to compete as well as lower willingness to take risk. If nurture is indeed the dominant force, this further suggests that “soft” policies that would reframe or recast certain educational and occupational choices to make them less threatening to women, very much along the line of a large literature in psychology on stereotype threat and the negation of these stereotypes, may help undo whatever role these traits have on holding back women.

IV. Gender Differences in Demand for Flexibility

Many of the higher-paying jobs in the economy involve long hours and inflexible schedules. Also, those financially more rewarding careers require continuous labor force attachment in order to stay on the “fast track,” which makes it difficult to combine those careers with job interruptions. Because women remain the dominant providers of child care as well as of other forms of non-market work, these various job features might be particularly detrimental to them.

In a study of the earnings trajectories of male and female graduates of the University of Chicago Booth School of Business, Bertrand et al. (2010) find that, 10 years post-graduation, employed female graduates earn about 50 percent less than their male counterparts. The authors also document that most of this gender gap in earnings 10 years out can be accounted for by differences in labor supply between men and women. Female graduates work shorter hours; they also have fewer years of actual labor market experience, as they are more likely to have taken some time out of the workforce since graduation from business school. Remarkably, the gender gaps in labor supply are not quantitatively large. The differences in hours worked between men and women 10 years out is 8 hours, with women working on average 49 hours a week compared to 57 hours for men. Yet flexibly accounting for hours work dramatically reduces the gender gap in labor market earnings. In other words, the labor market returns to working long hours are particularly large in the modal profession entered by business school graduates and women’s earnings are low relative to men because they are much less likely than men to put in such long hours.

In a recent and very influential contribution, Goldin (2014) more generally demonstrates that much can be understood about the gender pay gap within occupations by accounting for the elasticity of earnings in that occupation with respect to hours worked. Goldin first shows that there is a systematic relationship between the gender pay gap within an occupation and mean full-time, full-year earnings (wage and business income) among men in that occupation: in higher paying occupations, women’s earnings constitute a lower percentage of men’s earnings. Goldin further shows that those high earnings occupations where women experience a particular large deficit compared to men are also occupations where the elasticity of annual income with respect to weekly hours worked is particularly large. In other words, women are particularly unable to match men’s labor market achievement in those occupations where the rewards for working long hours are

particularly large. Business-type occupations show particularly large elasticities of income with respect to weekly hours worked. This is in contrast to technology-based occupations, where this elasticity is much lower. Trying to unpack what job features are associated with relatively higher returns to working long hours, Goldin (2004) shows that business-type occupations, in contrast to technology and science-based occupations, tend to involve work that is more structured, work that is done under more time pressure, and work that requires establishing and maintaining interpersonal relationships. Hence, women appear to be held back compared to men from reaching the top echelons of the highest paying occupations in the economy because of the greater inflexibility of work in these occupations, an inflexibility that might represent a greater disamenity to them than it does to men.

Indeed, there is growing robust empirical evidence supporting one's intuition that women value flexibility at work more than men do. A couple of recent papers provide clear demonstrations that women have a greater willingness to pay for job features that are consistent with more flexibility or shorter hours worked. For example, Mas and Pallais (forthcoming) use a field experiment to study how applicants to call center jobs value alternative work arrangements. Job applicants were randomly offered choices between typical 9-5, Monday to Friday positions and alternative work arrangements, such as more flexible work scheduling, the ability to work from home, or on-demand scheduling where the employer has discretion over the worker's schedule. By randomly varying the wage difference between the traditional option and the alternatives, the authors can estimate the distribution of willingness to pay for the alternatives. In their field context they find, maybe surprisingly, that even though there is a tail of workers who have a very high willingness to pay to choose days and time of work or number of hours of work, a majority of workers are not willing to pay for this flexibility. They document that workers particularly dislike on-call scheduling (where the employer can decide the worker schedule and only announces it a week in advance), with the average worker willing to give up 20 percent of wages to avoid such employer discretion. More relevant to us, they also show that women, and particularly those with young children, have a higher willingness to pay to be able to work from home and to avoid employers' scheduling discretion. They further show that women are indeed more represented in jobs that have such amenities but conclude that this difference in gender representation across jobs with these amenities are not large enough to explain the gender wage gap.

In a related study conducted in a research sample that is more immediately relevant to the glass ceiling discussion, Wiswall and Zafar (forthcoming) present New York University undergraduate students with hypothetical choices for jobs with different menus of job attributes. Again, by randomly varying the wage difference between these different menus, the authors can quantify willingness-to-pay for the attributes. They show that female undergraduate students have a higher willingness than male students to pay for jobs that offer greater work flexibility; female students are also willing to accept lower earnings in exchange for more job stability and have a greater willingness to pay for jobs that offer a part time option. On the other hand, they show that men have a higher willingness to pay than women for jobs that offer stronger earnings growth over time. They conclude that gender differences in preferences for job attributes explain at least a quarter of the early-career gender wage gap.

V. Gender Differences in the Impact of Childcare and Non-Market Work on Earnings

One of the core reasons as to why women may be at a particular high risk of underperforming in occupations that provide little flexibility, and demand more flexibility at work than men, is that they face greater additional pressures on their time as they try to balance market and non-market work commitments. A particularly important component of this non-market work involves taking care of children. Bertrand et al. (2010) found that one key factor explains why women with MBAs work shorter hours than men with MBAs as well as why they have fewer years of accumulated labor market experience: children. In particular, the MBA sample reveals that women without kids do not differ much from men (whether or not they themselves have kids) in terms of their labor supply. The group that accounts for the lower labor supply (and hence lower earnings) is women with children.

A recent wave of papers has leveraged rich administrative panel datasets to quantify child penalties in the broader population. Particularly relevant are recent papers by Angelov et al. (2016) and Kleven and Landais (2017), who compute child penalties in Sweden and Denmark, respectively. The empirical approach both papers follow consists in measuring within-couple changes in earnings following parenthood, with the identifying assumption being that the decision of when to enter parenthood is not induced by unobservable (to the econometrician) information about a coming change in the earnings trajectory of one of the spouses. Angelov et al. (2016) find that,

while earnings of husbands and wives move in parallel fashion in the years that precede the birth of the first child, wives' earnings start diverging from their husbands' immediately post-birth. They find that, 15 years after entering parenthood, the male-female gender gap in income has increased by 28 percentage points over its pre-child level. In further analyses where they estimate these earnings dynamics separately for groups with varying within-couple relative education, they show that the effect of parenthood on the male-female gap increases with the predicted income and wage gap the couple would have experienced in the absence of parenthood. In other words, matching of couples is crucial for the magnitude of the effect of parenthood on the gender gap. Women married to men of higher earnings potential experience especially large income losses following parenthood.

Kleven and Landais (2017) report on very similar dynamics in the Danish context, with wives experiencing sharp drops in labor force participation rate, earnings, hours worked, and wage rates compared to their spouses immediately after giving birth to their first child. In a back-of-the-envelope exercise aimed at determining how much this child penalty accounts for the gender wage gap over time, Kleven and Landais (2017) conclude that nearly all of the remaining gender wage gap in Denmark in the early 2010s can be accounted for by this child penalty. This is in contrast with 30 years ago where the gender wage gap can be decomposed into 3 about equal-magnitude components: child-related, education-related, and an unexplained residual. Moreover, the child-related contribution to gender inequality in earnings appears as large in the early 2010s as it was in the early 1980s. One must note how remarkable these results are given that, of all the developed nations, Scandinavian countries have smaller gender imbalances in the contribution to non-market work.

A couple of additional papers are worth mentioning here in that they have taken a different approach to also document how household work, of which childcare is one element, contributes to holding women's earnings back in the labor market. Cortes and Tessada (2011) and Cortes and Pan (forthcoming) exploit regional variation throughout the US in low-skill immigration. Such immigration, as shown in Cortes (2008), makes the outsourcing of various home production activities (nannies, gardeners, etc.) more affordable, which can be hypothesized to lift some of the time pressures women, especially those who have higher earnings potential in the marketplace, disproportionately face as they try to balance home and work activities. And indeed, Cortes and

Tessada (2011) find positive effects of low-skilled immigration on the hours worked per week of women in the top quartile of the female wage distribution (with no effects for working women in the lower part of the wage distribution). In a complementary paper, Cortes and Pan (forthcoming) show that low-skill immigration, and hence access to cheaper outsourcing options for home production, has helped high-skilled women working in occupations where the premium for working long hours is particularly large and where women face particularly large earnings deficits compared to men (as studied in Goldin, 2014). In particular, Cortes and Pan (forthcoming) show that women working in those inflexible occupations increase their working hours and experience earnings gains compared to men when exposed to a larger pool of low-skilled immigrants in their locality. The effects appear particularly large among women married to husbands who also work long hours, consistent with the evidence above of the crucial role played by the matching of men and women into couples. Moreover, the authors report an increased entry of young women in the occupations that reward long hours when they have access to a greater supply of low-skilled immigrants.

VI. Why Do Women Continue to Bear Most of the Penalty for Non-Market Work?

The discussion of the literature above strongly suggests that one of the most prominent factors preventing women back from operating on an even playing field with men in the labor market is that they continue to bear the brunt of the labor market cost of non-market work, including childcare. This is particularly relevant to any discussion about the glass ceiling in that this non-market demand on women's time is particularly disruptive when they are trying to compete in those higher earnings occupations that reward more intensive time commitments and penalize flexibility.

However, at first glance, it seems that several forces should make non-market work, or double-shift considerations, less of a factor for women today than in the past. First, there has been a sharp decline in the overall amount of time spent on household tasks (Aguiar and Hurst, 2007). Technological change embodied in dishwashers, microwaves and other timesaving technologies have made it easier to maintain a household without a large time investment. Similarly, as discussed above, low-skilled immigration has made it easier to relegate whatever is left of household work to third parties. These changes have been particularly relevant for the more

educated and those with the highest labor market earnings potentials as they can more easily afford these new technologies and the outside help. The labor-savings technologies in the home in fact have been shown to play a non-trivial role in explaining the historical rise in women's labor force participation.

Yet, a possible countervailing force, at least in the US context, has been the rise in the amount of time spent on parenting activities among the more highly educated (Guryan et al., 2008). It remains unclear what has driven this increase. One factor that has been discussed has been a growing race between better-off families to best position their kids for a limited number of slots in the most prestigious universities (Ramey and Ramey, 2010). Whatever the explanation, it appears that the willingness to outsource child-rearing activities to an outside party might not be a winning proposition for many families today, especially when the quality of outsourced child-rearing might fall well below what could be achieved by a well-educated parent. These trends in parenting time are an understudied factor in the quest for more equal gender representation at the top of the labor market and are deserving of more research attention.

Second, even if there remains some amount of household and childcare work that cannot be bought in the market, the gender norms that would imply that women and mothers would be the ones expected to be the main providers of this non-market work have weakened over time. In a panel of OECD countries, Fortin (2005) shows that the share of men (and women) holding conservative gender norm attitudes has been going down over time: among those born after 1965, only about 15 percent of women (and 21 percent of men) agree that "when jobs are scarce, these jobs should go to men first"; this is in contrast with the 36 percent of women (and 38 percent of men) among those born prior to 1935 who agreed with this statement. Similarly, while about 57 percent of women born and 63 percent of men born after 1965 agree with the statement that "being a housewife is fulfilling," these shares are 69 and 72 percent, respectively, among those born prior to 1935.

However, as Fortin (2005) also shows, there is not much difference in gender attitudes between the more recent birth cohorts and the post-1940 birth cohorts (especially when it comes to the views regarding whether there might be a conflict between working mothers and the well-being of their children), suggesting some possible stagnation on the way toward fully neutral gender norms.

Moreover, some “dormant” gender identity norms may only start to become empirically relevant when women’s position in the labor market improves. Consider in particular the gender norms related to wives’ earnings. Women’s improved potential in the labor market has increased the likelihood that a wife would out-earn her husband. Hence, even if attitudes on “whether a man should earn more than his wife” have become more liberal over time, this particular norm may matter more today than it did in the past because it has become binding. Bertrand et al. (2015) explore the possible manifestations of this gender norm in patterns of relative income within households, marriage formation, wives’ labor force participation, marital satisfaction, and the division of home production. Using US administrative data, they show that the distribution of the share of the household income earned by the wife exhibits a sharp drop to the right of 0.5, i.e. when the wife starts to earn more than the husband. This pattern is consistent with the existence of gender identity norms that induce an aversion to a situation where the wife earns more than her husband. Bertrand et al. (2015) also show that this specific gender norm appears to distort labor market outcomes within couples. In particular, they show that when the probability that the wife’s income would exceed her husband’s actual income is higher, the wife is less likely to participate in the labor force. Moreover, if she does work, the gap between her realized and her potential income is greater (in part due to lower hours of work). Furthermore, when gender identity norms are violated in this way, marriages appear to suffer: couples where the wife earns more than the husband are less happy, report greater strife in their marriage and are ultimately more likely to end in divorce. Finally, using time-use data, they show that, controlling for both individuals’ incomes, when the wife starts to earn more than the husband, she starts taking on more of the household chores. In other words, these “threatening” wives may take on a greater share of the non-market work in order to assuage their husbands’ unease with the situation.

In a recent paper, Burstyn et al. (forthcoming) show that conservative gender norms are still very much relevant today even among MBA students. In particular, the authors hypothesize that women who are or appear professionally too ambitious may be less attractive in the marriage market, or at least believe that they will be perceived as less attractive. Using both administrative and survey data from MBA students at UCLA Anderson School, Burstyn et al. (forthcoming) document a set of descriptive facts that appear consistent with this hypothesis. For example, the authors show that single female students have lower class participation grades than married female students, despite

not having systematically worse exam grades: this is as if they are trying hard not to appear too smart or too ambitious in front of the class (where their future spouse might be). More interesting and convincing than the descriptive facts is a simple controlled experiment the authors ran with the MBA students. In particular, the authors leveraged a questionnaire on job preferences and personality traits that must be filled by newly admitted students. Answers provided by students to this questionnaire have “high stakes,” in that they will be used by career services for internship placement. The students were then randomized into being told that their answers to the questionnaire would be shared for discussion with classmates (the public condition) or that they would be anonymized prior to being shared for discussion (the private condition). Overall, the results show that single female students report less ambitious career goals in the public vs. the private condition, while the public vs. private condition does not greatly affect the answers provided by non-single female students. In particular, single women appear reluctant to publicly signal that they desire high level of compensation, or that they are willing to travel many days per month.

But let’s accept as a fact that the trends in overall amount of non-market work and the strengths of gender norms regarding who is responsible for this non-market work have all moved toward a situation where women should have an easier time than in the past balancing work in the labor market and work at home. As conditions are thus improving on the home front, are there countervailing forces that may have contributed to these improvements not translating into more labor gains and more cracks in the glass ceiling? We consider two such possible countervailing forces below, the first based in the labor market and the second based in the home.

First, even as women find it easier than in the past to navigate market and non-market commitments, it is possible that the rewards for inflexible work have gone up in the labor market. In other words, while women can now devote more time to work than in the past, the structure of jobs in the economy, and especially the higher paying jobs, may have changed in a way that increasingly penalizes their smaller (but remaining) higher demand for flexibility.

We take a stab at looking into this possibility in Table 2. In particular, we follow Goldin (2014) and compute for each occupation and census year a measure of the elasticity of occupational earnings to hours worked. This measure is undeniably a very imperfect proxy for how much an

occupation rewards a wholesome commitment to work. It is also subject to a lot of measurement error. Nevertheless, it is a proxy that can be easily computed across all occupations and over time.

We proceed as follows to compute this occupation-census year metric. In each census year, we focus on the set of college-educated (or more) men and women who are between 25 and 64 years of age and report working full-time, full-year. We then regress log earnings on a quadratic in age, race-fixed effects, state-fixed effects, highest degree completed fixed effects, weeks of work fixed effects, occupation fixed effects interacted with gender, and occupation fixed effects interacted with log(weekly hours worked). We then retrieve the coefficients estimated on log (weekly hours worked) for each occupation. These coefficients represent the estimated elasticity of earnings in an occupation to weekly hours worked. Occupations are defined using the 3-digit 1990 occupation code. Finally, we perform this exercise only starting in 1980 given the coarseness of the measurement of weekly hours worked in the 1970 Census.

Table 2 presents means and medians of these occupation-specific elasticities for 1980, 1990, 2000, 2010 and 2015 (2015 corresponds to the ACS 2012-2015 files; we add 2015 to minimize concerns about how the financial crisis may have impacted the 2010 estimates). When computing means and medians, we weigh each occupation by the number of individuals (between 25 and 64 years of age) working full-time, full-year in that occupation. Finally, we present means and medians across all occupations (Panel A), as well the subset of the 100 highest (Panel B), 50 highest (Panel C) and 10 highest (Panel D) paying occupations (based on full-time, full-year male earnings). These higher earnings occupations are identified separately across years.

This table reveals a few important patterns. First, as shown in Goldin (2014), it is generally true that the elasticity of earnings to hours worked is larger in the higher paying occupations. This pattern appears particularly clearly in the comparison of medians across the panels. Second, and most relevant to our argument, are the changes over time in the estimated elasticities. The overall trend is clearly up, and this is true across all occupations, as well as in the subset of higher-paying occupations. The biggest increase is measured between 1980 and 2000, but there is no overall pattern of a decline between 2000 and 2015. Hence, Table 2 provides one clue to a possible countervailing force: as women manage to free up more of their time for work, work might ask for even more time.

A second countervailing force, we argue, might be in the home. While women today may have an easier time than in the past balancing the demand of labor market work and home-based work, home-based work is a factor for a growing share of college educated women in high-paying occupations because a growing share of them *have* a family. We illustrate this point in Table 3. In particular, we compute for each decade between 1980 and 2010 the percent of employed college-educated women between 25 and 64 who are currently married, were never married, and have a child living in the household. We do this separately for women employed in occupations below the top 100 highest paying, in the top 100, in the top 50 and in the top 10. As in the prior table, the determination of which occupation is highest paying is done separately in each year. Also, the means reported in the table are weighed by the number of college-educated women employed in each occupation.

Consider first 1980 (Panel A). There is a clear relationship between the likelihood of being married, having never been married or having a child at home, and occupational rank. Women working in higher-paying occupations are less likely to be married, more likely to have never been married and less likely to have a child at home. Only 54 percent of women working in one of the 10 highest-paying occupations in 1980 report being married, compared to 68 percent of those working in occupations below the 100 rank. Only 34 percent of women working in one of the top highest-paying occupations have a child at home, compared to 51 percent of those working in occupations below the 100 rank.

Over time (Panels B to C), this relationship weakens, and by 2010 the relationship is essentially gone. Women working in 2010 in the top 10 highest-paying occupations in the economy are in fact now about 2 percent more likely to be married than those working in occupations below the 100 rank (68 percent versus 66 percent); they are as likely to report having a child at home (49 percent).

It is beyond the scope of this paper to determine why these changes happened. In a related project, Bertrand et al. (2017a) discuss how improvements in labor market outcomes for college-educated women, while first making them less attractive in the marriage market because of gender norms against having a working wife who cannot provide as much of the household good, may eventually start helping these women in the marriage market as men see enough compensation (via their

wives' growing earnings) for the loss of the household good. It is also possible that women today more than in the past believe (rationally or not) that they can "have it both" and are acting accordingly. But whatever the explanation, Table 3 suggests indeed another countervailing force: while women in high-paying occupations should have an easier time than in the past navigating the demands of work and family, more of these women today than in the past have a family and are confronting its demands.

Furthermore, as suggested by the last four columns of Table 3, the families these women are entering into are families where one should expect there will be competing demands on husbands' and wives' time. We illustrate this point by computing, for each group of employed college-educated women, the average elasticity of earnings to hours worked in their *husbands'* occupation. We do this for all married, employed, college-educated women (by occupation groups), as well as for the subset of employed, college-educated women who are married to a college-educated man. As seen in column 4, the likelihood of getting married to a college-educated man is higher among women employed in higher-paying occupations.

First, we note, not surprisingly given the patterns in Table 2, that the elasticity of earnings to hours worked in husband's occupation has grown up over time for all groups of college-educated women. It is also the case, even though the gradient is not steep, that women employed in higher-paying occupations are married to men working in occupations where the elasticity of earnings to hours worked is larger. Finally, as a proxy for the relative labor market demand on wives' and husbands' time, we compute for each occupation group the fraction of couples where the elasticity of earnings to hours worked is higher in the husband's occupation than it is in the wife's occupation. While we see the expected gradient (e.g. husbands of women in higher-paying occupations are relatively more flexible at work) across all years, it is not the case that women employed in higher earnings occupation today are more likely than in the past to have husbands who are more flexible at work than they are. In fact, except for a decline between 1980 and 1990, the share of couples with wives working in one of the top 100, top 50 or top 10 occupations where the husband's elasticity of earnings to hours work is higher than the wife's has been going up.

In summary, we argue that both changes in the structure of work, with greater rewards over time for fuller commitment work in the higher-paying occupations, as well as changes in the home, with

a growing share of women in higher-paying occupations having families with men who are also looking at greater labor market rewards for fuller commitment to work, might have been important countervailing forces to the steadier progress one might have expected for women given the weakening of gender norms and decline in the length of the “second shift.”

VII. Policy Responses to the Glass Ceiling

Family-Friendly Policies

Many firm-level human resource policies as well as public policies have as a goal to address the greater demand for flexibility women are requesting from the workplace, and, more broadly, want to make the workplace more family-friendly. Such policies include, among others: longer and paid maternity leave, options to work part-time, options to work shorter hours, options to work remotely and more employee control on the structure of the work day and work week. While offering such alternative work arrangements may achieve the objective of attracting and retaining women in the workforce, it will not reduce the gender gap as long as this flexibility is negatively priced in the market. In other words, while these policies offer more flexibility at work, they do not reduce the labor market penalties associated with taking on that flexibility. Of course, one cannot rule out some more general equilibrium medium-to-long-run effect where, as a growing share of employees opt into alternative work arrangements, employers may face the need to reorganize the structure of work in a way that would no longer penalize as much the exercise of this flexibility. In the short-term, though, it is reasonable to worry that many of these policies may backfire and only reinforce the glass ceiling if women take them up at a higher rate than men. For example, longer maternity leave raises the costs for employers of hiring women of child-bearing age. This may lead employers to not assign women to the most important jobs or clients as they expect these women not to be able to generate as much surplus for the organization because of their inability to constantly maintain those important interpersonal relationships with the clients. These policies may also keep women out of the workforce for “too long” to ensure a re-entry on the fast track. Overall, there very well might be, in the short-term at least, a tradeoff in that these family-friendly policies might succeed on the one hand in further reducing (and maybe even flipping) the gender gap in labor force participation but might, on the other hand, increase the gender gap in earnings, especially at the top of the earnings distribution.

Blau and Kahn (2013) report on some empirical evidence consistent with such a trade-off. In a panel of OECD countries, Blau and Kahn (2013) provide suggestive evidence that the plateauing in female labor force participation in the US compared to other countries can be accounted for by more aggressive work-family balance policies in these non-US OECD countries. But, on the other hand, Blau and Kahn (2013) also show that the US has been relatively more successful than the other OECD countries in reaching a higher representation of women in the high-paying managerial and professional occupations in the economy.

Gender-Neutralizing Childcare

There has been effort in some countries to encourage more fathers to take up parental leave. For example, Sweden, Norway and Quebec have introduced dedicated paternity leave into their parental leave policy. This is most often achieved by reserving some months of the parental leave for fathers, with these months being lost if not taken up by the father (e.g. via “daddy quotas” or “daddy months”). Such policies are appealing in that they attempt to go to the core of what seems to hold women’s earnings back in the labor market by trying to speed up the shift in social norms that still too often associate the mother as the dedicated provider of childcare and non-market work.

Correlational studies, while obviously subject to endogeneity and selection bias, have raised the possibility that these policies might be effective policy tools. Such studies have shown that fathers that take more leave around the time of the birth of their child remain more involved in child care throughout the life of the child (Haas, 1990; Nepomnyaschy and Waldfogel, 2007; Tanaka and Waldfogel, 2007). Furthermore, there appears to be some correlation between paternity leave and fathers’ involvement in housework (Brandth and Kvande, 1998; Hook, 2006). Finally, there is also some correlational evidence that paternity leave is associated with improved labor market outcomes for mothers (Pylkkanen and Smith, 2003).

An emerging and still growing body of work has been trying to develop more causal research design to get around the obvious selection and endogeneity concerns in the descriptive work above by exploiting changes in parents’ behaviors and outcomes around the time of the introduction of the new parental leave policies. The results of this work are quite mixed so far, but it is fair to say that the overall optimistic findings of the cross-sectional research have been somewhat tempered

by the more robust research designs. On the positive side, this new research confirms that daddy quotas are effective at increasing paternity leave. For example, Dahl et al. (2014) report that the introduction of a daddy quota in Norway increased fathers' take-up of paternity leave by about 30 percentage points. Ekberg et al. (2013) report on similar positive effects on take-up for Sweden, while Patnaik (2016) report on similar large positive effects on take-up for Quebec. On the other hand, the evidence in these papers also suggests that the policies' impact is often limited to the period of parental leave time that is reserved to the father; in other words, fathers take up their allocated daddy quota but rarely extend the duration of their leave beyond that. Also, this new research offers mixed findings on the impact of the daddy quotas on the division of childcare and housework between spouses, as well as mothers' labor market outcomes. Ekberg et al. (2013) does not find evidence that the daddy quotas in Sweden increased fathers' share of child care duties. While Johansson (2010) and Cools et al. (2015) find that paternity leave leads to higher maternal earnings, Rege and Solli (2013) find no causal effect on mothers' earnings and Cools et al. (2015) find negative effects on maternal earnings. Patnaik (2016) finds the most consistently positive picture of how paternal leave might help reshape the division of responsibilities within the family. In particular, he finds much more robust evidence of an impact of the daddy quota policy in Quebec on the division of household work, with exposure to the policy increasing the amount of time fathers spend on housework activities and decreasing the amount of time mothers spend on these activities. Moreover, Patnaik (2016) finds that the daddy quotas increase the amount of time mothers spend in market work as well as the share of mothers working full-time.

Recent work focusing on the academic labor market also highlights the possibility that daddy quotas may have perverse effects, at least when there is no guarantee that that the "daddy months" will be as fully devoted to child care as the "mommy months." In particular, Antecol et al. (2016) build a dataset on the universe of assistant professor hires at top 50 economics departments from 1985-2004 and study the impact of gender-neutral tenure clock stopping policies on the tenure rates for male and female faculty members. They find that the adoption of such gender-neutral tenure clock stopping policies substantially reduced female tenure rates while substantially increasing male tenure rates: after the implementation of a gender-neutral clock stopping policy, the probability that a female assistant professor gets tenure at that university decreases by 22 percent while male tenure rates rise by 19 percent. When probing for the underlying mechanisms,

the authors show that men publish more in the “top five” economics journals after the policy is implemented, while there is no such effect for women. In other words, it appears that fathers use the extra time on their tenure clock to write papers, while mothers devote more of that extra time to being mothers.

Despite the mixed results so far, and even possible perverse effects in some institutional contexts, the daddy quotas remain one of the most promising concrete policy proposals in that, by nudging families toward more neutral childcare arrangements, they aim to redress a key differential barrier women face in the workforce.

Affirmative Action: Quotas and Gender Diversity in Leadership

Affirmative action-type policies have also been gaining traction as a response to the glass ceiling. While many countries around the world have gender quotas in the political system, gender quotas in other sectors of work have been rarer. Yet in recent years, many European countries have introduced gender quotas in the corporate sector. The first quotas for women in business were introduced in Norway in 2003, with a law requiring 40 percent representation of each gender on the board of directors of public limited liability companies. Following Norway’s lead, Belgium, France, Germany, Iceland, India, Israel, Italy, and Spain have all passed similar reforms. In 2014, the new German coalition government passed legislation requiring that corporate boards be comprised of at least 30 percent women by 2016 (or else the seat would be left vacant). In 2013, the European parliament voted in favor of a proposed draft law that would require 40 percent female board members in about 5,000 listed companies in the European Union by 2020.

While gender quotas represent a very visible response by the political class to the dearth of women in the top corporate echelons, it is unclear whether such quotas can be an effective tool to break the glass ceiling beyond the mechanical impact they will have on the very limited number of women who will directly benefit by becoming board members. Indeed, the most logical theories of change when assessing how quotas may improve labor market conditions for women do not seem directly related to the considerations regarding job design and unequal demand on women’s time outside of work, which, as we argued above, might be the two key factors holding women’s earnings back. In particular, quotas might be most effective if the main constraint faced by women

is one of path dependence: qualified women are harmed by an absence of networks to help them climb the corporate ladder, and quotas can provide the push that women need to break this cycle and change the equilibrium. Also, quotas might also be a powerful policy if discrimination is the key factor for the underrepresentation of women: women in position of leadership may favor promoting other women (or at least not discriminate as much as men against promoting women) and exposure to women in positions of leadership might reduce men's prejudice toward having more women in leadership as they start to see less of a conflict between the traits associated with leadership (strength, assertiveness) and the traits that are typically associated with women (being nice, accommodating). Any effects quotas may have on job design or views about how non-market work should be allocated between the genders, while not out of the question, seem at the very least much more indirect.

Given this context, it is maybe not surprising that the evidence that exists so far on whether gender quotas can help break the glass ceiling is somewhat disappointing. In particular, Bertrand et al. (2017b) study the effects of the 2003 Norwegian law. On the positive side, they document that the women appointed to these boards post-reform were observably more qualified than their female predecessors along many dimensions, and that the gender gap in earnings within boards fell substantially. This is an important and positive outcome for the policy in that the boardrooms did not become simply mechanically more gender balanced in terms of body counts: they also became more gender balanced in terms of the qualifications of the individuals sitting around the table. On the other hand, the authors find no robust evidence that the gender quota reform benefited the larger set of women employed in the companies subject to the quota. Moreover, the reform had no clear impact on career women whose qualifications mirror those of board members but who were not appointed to boards. Overall, they conclude that the policy had very little discernible impact on women in business as of the mid-2010s beyond its direct effect on the few women who made it into boardrooms.

Corporate policies may also incorporate some elements of affirmative action, more or less formally. Anecdotally, there is a sense that at least a subset of companies are trying to improve outcomes for their female employees, for example by taking proactive steps toward achieving some female representation on all key committees within the organization. The logic behind these policies, and how they may help the broader set of women within the organization, appears quite

similar to that of the explicit government-mandated quotas: women in leadership will be watching out for other women, they will be better placed to “network in” other women into important decisions the committee may take, and they may help reshape the views of male committee members by weakening in their mind inconsistencies between the female gender stereotype and qualities associated with leadership, so that prejudice toward female leaders will diminish.

We are not aware of much publicly available data demonstrating how common these practices are among organizations. One exception is Misra et al. (2012) who survey hundreds of faculty members at the University of Massachusetts Amherst in 2008–09. The survey results indicate that three-quarters of female associate professors, compared with only half of their male counterparts, had played major service roles at the university; a third of the women had served as undergraduate directors, compared with only 17 percent of men. Overall, compared to men, women faculty spent 7.5 fewer hours per week on research and 4.6 more hours per week on university service committees. Also, in a sample of 1,400 political science faculty in the US, Mitchell and Heslie (2013) find that women participated in more department and college-level committees.

Babcock et al. (2017) provide evidence suggesting that these gender imbalances in the amount of committee work might be in part a supply response: women are less likely than men to decline requests to participate in what Babcock et al. (2017) call “non-promotable” tasks, e.g. tasks that might be good for the organization (and other women in the organization) but of limited direct benefits to them. A more charitable explanation, as stressed above, is that employers are actively trying to achieve greater gender diversity in decision-making within the organization, with the hope of positive ripple-down effects on other women.

Lack of data on the prevalence of these practices also means there is no research we are aware of on the impact of these practices. A priori, the main reason to be skeptical about any large positive impact on other women in the organization is that, just like quotas, these practices only seem to very indirectly deal with the key barriers for women coming from job design and allocation of non-market work within the household. Moreover, in this case as well, there is a clear possibility that these practices may have some perverse effects. As companies have to draw from a limited pool of women to operationalize more diversity in these “non-promotable tasks,” this may result in a non-trivial additional tax on the time of these few women. This additional tax, as discussed at

length above, may be particularly costly to women in those higher-paying occupations that offer large rewards for time spent focused on “promotable” work.

VIII. Concluding Remarks

In this paper, we have reviewed the lines of enquiries that are currently dominating the academic discussion of the glass ceiling, and have argued about the relative merits of various policy responses in light of what this literature has diagnosed as the most relevant factors explaining why women still lag behind men in the labor market.

A less academic treatment of the glass ceiling would certainly have devoted more space to the topic of pure (e.g. taste-based) gender discrimination and sexism in the workplace. And there is certainly plenty of anecdotal evidence that such sexism exists across many occupations and that it has prevented many talented women from achieving their full potential work. We have all read these stories in newspapers and most of us know someone (or is the “someone”) whose career has been derailed by such workplace practices. Moreover, because we argue above that gender norms within the home may prevent many women from as fully engaging in the labor market as men do, it would be inconsistent to then claim that such gender norms are not relevant to what is happening in the workplace. What we however strongly believe is that the other factors we have prioritized in this paper are not only relevant but also quantitatively important and that it would be wrong to have a glass ceiling discussion that is solely about labor market discrimination and sexism. But to the extent that there is a residual when all the factors we have discussed are accounted for, sexism in the workplace should be high on the list to name that residual. Directly measuring gender discrimination in the workplace, in particular in type of jobs that are relevant to the glass ceiling discussion, is difficult. Audit or resume studies are much easier to implement in lower-skilled entry-level jobs (Bertrand and Duflo, 2017). Moreover, our discussion above suggests that statistical discrimination should be an important factor and audit and resume studies are not well equipped to separate taste-based and statistical discrimination. But a direct measurement of the role of sexism in the glass ceiling is a worthwhile challenge for future work, as this obviously would improve on “naming the residual.”

Finally, we believe that one of the biggest unknowns when trying to predict how the glass ceiling will involve in the future is the role of technology. There is no doubt that many trends are moving in the “right direction” for women, some very quickly (such as the large and rising reverse gender gap in completed schooling) and some more slowly (such as the declining conservativeness of gender norms). It is possible though, as we hinted at above, that changes in the structure of work and job design over the last 40 years may not have not been as beneficial to women. How the next wave of technological change in the workplace (i.e. artificial intelligence) will change the structure of work is anyone’s guess.

References

Aguiar, Mark, and Erik Hurst. "Measuring Trends in Leisure: The Allocation of Time Over Five Decades." *The Quarterly Journal of Economics* 122, no. 3 (August 2007): 969-1006.

Angelov, Nikolay, Per Johansson, and Erica Lindahl. "Parenthood and the Gender Gap in Pay." *Journal of Labor Economics* 34, no. 3 (July 2016): 545-579.

Antecol, Heather, Kelly Bedard, and Jenna Stearns. "Equal but Inequitable: Who Benefits from Gender-Neutral Tenure Clock Stopping Policies?" Institute for the Study of Labor (IZA) Discussion Paper No. 9904, (April 2016).

Babcock, Linda, Maria P. Recalde, Lise Vesterlund, and Laurie Weingart. "Gender Differences in Accepting and Receiving Requests for Tasks with Low Promotability." *American Economic Review* 107, no. 3 (March 2017): 714-747.

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics* 2, no. 3 (July 2010): 228-255.

Bertrand, Marianne. "New Perspectives on Gender." In *Handbook of Labor Economics* 4B, edited by Orley Ashenfelter and David Card, 1543-1590. North Holland: 2011.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. "Gender Identity and Relative Income within Households." *The Quarterly Journal of Economics* 130, no. 2 (2015): 571-614.

Bertrand, Marianne and Esther Duflo. "Field Experiments on Discrimination." In *Handbook of Field Experiments* 1, edited by Abhijit V. Banerjee and Esther Duflo, 309-393. New York: Elsevier, 2017.

Bertrand, Marianne, Patricia Cortés, Claudia Olivetti, and Jessica Pan. "Social Norms, Labor Market Opportunities, and the Marriage Gap for Skilled Women." University of Chicago Working Paper, (2017a).

Bertrand, Marianne, Sandra Black, Adriana Lleras-Muney and Sissel Jensen. "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway." University of Chicago Working Paper, (2017b).

Blau, Francine D., and Lawrence M. Kahn. "Female Labor Supply: Why Is the United States Falling Behind?" *The American Economic Review* 103, no. 3 (May 2013): 251-256.

Blau, Francine D., and Lawrence M. Kahn. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55, no. 3 (2017): 789-865.

Brandth, Berit, and Elin Kvande. "Masculinity and Child Care: The Reconstruction of Fathering." *The Sociological Review* 46, no. 2 (May 1998): 293-313.

- Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais. “‘Acting Wife’: Marriage Market Incentives and Labor Market Investments.” *The American Economic Review*, forthcoming.
- Buser, Thomas, Muriel Niederle, and Hessel Oosterbeek. “Gender, Competitiveness, and Career Choices.” *The Quarterly Journal of Economics* 129, no. 3 (August 2014): 1409-1447.
- Catalyst. Pyramid: Women in S&P 500 Companies. New York: Catalyst, March 1, 2017.
- Cools, Sara, Jon H. Fiva, and Lars J. Kirkebøen. “Causal Effects of Paternity Leave on Children and Parents.” *The Scandinavian Journal of Economics* 117, no. 3 (July 2015): 801-828.
- Cortés, Patricia. “The Effect of Low-skilled Immigration on US Prices: Evidence from CPI Data.” *Journal of Political Economy* 116, no. 3 (June 2008): 381-42.
- Cortés, Patricia and Jose Tessada. “Low-skilled Immigration and the Labor Supply of Highly Skilled Women.” *American Economic Journal: Applied Economics* 3, no. 3 (July 2011): 88-123.
- Cortés, Patricia and Jessica Pan. “When Time Binds: Substitutes to Household Production, Returns to Working Long Hours and the Gender Wage Gap among the Highly Skilled.” *Journal of Labor Economics*, forthcoming.
- Croson, Rachel, and Uri Gneezy. “Gender Differences in Preferences.” *Journal of Economic Literature* 47, no. 2 (June 2009): 448-74.
- Dahl, Gordon B., Katrine V. Løken, and Magne Mogstad. “Peer Effects in Program Participation.” *American Economic Review* 104, no. 7 (July 2014): 2049-2074.
- Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner. “Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences.” *Journal of the European Economic Association* 9, no. 3 (June 2011): 522-550.
- Ekberg, John, Richard Eriksson, and Guido Friebel. “Parental leave – A policy evaluation of the Swedish ‘Daddy-Month’ reform.” *Journal of Public Economics* 97, (January 2013): 131-143.
- Eckel, Catherine C., and Philip J. Grossman. “Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk.” *Evolution and Human Behavior* 23, no. 4 (July 2002): 281-295.
- Eckel, Catherine C., and Philip J. Grossman. “Differences in the Economic Decisions of Men and Women: Experimental Evidence.” In *Handbook of Experimental Economics Results* 1, edited by Charles R. Plott and Vernon L. Smith, 509-519. New York: Elsevier, 2008.
- European Union. “Gender balance on corporate boards. Europe is cracking the glass ceiling.” Directorate-General for Justice and Consumers, Fact sheet, July 2016.

Flory, Jeffrey A., Andreas Leibbrandt, and John A. List. "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions." *The Review of Economic Studies* 82, no. 1 (January 2015): 122–155.

Fortin, Nicole M. "Gender Role Attitudes and the Labour-market Outcomes of Women across OECD Countries." *Oxford Review of Economic Policy* 21, no. 3 (October 2005): 416-438.

Gneezy, Uri, Muriel Niederle, and Aldo Rustichini. "Performance in Competitive Environments: Gender Differences." *The Quarterly Journal of Economics* 118, no. 3 (2003): 1049-1074.

Goldin, Claudia, Lawrence F. Katz, and Ilyana Kuziemko. "The Homecoming of American College Women: The Reversal of the College Gender Gap." *The Journal of Economic Perspectives* 20, no. 4 (Fall 2006): 133-156.

Goldin, Claudia. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104, no. 4 (April 2014): 1091-1119.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney. "Parental Education and Parental Time with Children." *Journal of Economic Perspectives* 22, no. 3 (Summer 2008): 23-46.

Haas, Linda. "Gender Equality and Social Policy: Implications of a Study of Parental Leave in Sweden." *Journal of Family Issues* 11, no. 4 (December 1990): 401-423.

Haas, Linda, and C. Philip Hwang. "The Impact of Taking Parental Leave on Fathers' Participation In Childcare And Relationships With Children: Lessons from Sweden." *Community, Work & Family* 11, no. 1 (February 2008): 85-104.

Hook, Jennifer L. "Care in context: Men's unpaid work in 20 countries, 1965-2003." *American Sociological Review* 71, no. 4 (August 2006): 639-660.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones and Pete Klenow. "The Allocation of Talent and U.S. Economic Growth." University of Chicago Working Paper, (August 2016).

Johansson, Elly-Ann. "The effect of own and spousal parental leave on earnings." Institute for Labour Market Policy Evaluation Working Paper, (March 2010).

Kleven, Henrik Jacobsen, and Camille Landais. "Children and Gender Inequality: Evidence from Denmark." LSE Working Paper, (February 2017).

Mas, Alexandre, and Amanda Pallais. "Valuing Alternative Work Arrangements." *American Economic Review*, forthcoming.

Misra, Joya, Jennifer Hicke Lundquist, and Abby Templer. "Gender, Work Time, and Care Responsibilities Among Faculty." *Sociological Forum* 27, no. 2 (June 2012): 300-323.

Mitchell, Sara McLaughlin, and Vicki L. Hesli. "Women Don't Ask? Women Don't Say No? Bargaining and Service in the Political Science Profession." *PS: Political Science & Politics* 46, no. 2 (April 2013): 355-369.

Nepomnyaschy, Lenna, and Jane Waldfogel. "Paternity Leave and Fathers' Involvement with Their Young Children: Evidence from the American Ecls-B." *Community, Work & Family* 10, no. 4 (November 2007): 427-453.

Niederle, Muriel, and Lise Vesterlund. "Do Women Shy Away from Competition? Do Men Compete Too Much?" *The Quarterly Journal of Economics* 122, no. 3 (August 2007): 1067-1101.

Patnaik, Ankita. "Reserving Time for Daddy, The Short and Long-Run Consequences of Fathers' Quotas." Working Paper, (January 2016).

Pylkkänen, Elina, and Nina Smith. "Career Interruptions due to Parental Leave – A Comparative Study of Denmark and Sweden." *OECD Social, Employment and Migration Working Papers No. 1*, (March 2003).

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. "Distributional National Accounts: Methods and Estimates for the United States." NBER Working Paper No. 22945, (December 2016).

Ramey, Garey, and Valerie A. Ramey. "The Rug Rat Race." *Brookings Papers on Economic Activity*, (Spring 2010): 129-176.

Rege, Mari, and Ingeborg F. Solli. "The Impact of Paternity Leave on Fathers' Future Earnings." *Demography* 50, no. 6 (December 2013): 2255-2277.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. "Taste for Competition and the Gender Gap Among Young Business Professionals." NBER Working Paper No. 21695, (November 2015).

Tanaka, Sakiko, and Jane Waldfogel. "Effects of Parental Leave and Work Hours on Fathers' Involvement With Their Babies: Evidence From the Millennium Cohort Study." *Community, Work & Family* 10, no. 4 (November 2007): 409-426.

Wiswall, Matthew, and Basit Zafar. "Preference for the Workplace, Investment in Human Capital, and Gender." *The Quarterly Journal of Economics*, forthcoming.

Figure 1




Figure 2




Table 1: Women's Labor Force Participation & Representation in the Upper Part of the Earnings Distribution

<i>Panel A</i>		<i>Sample: All Women</i>				
		Year is:				
		1970	1980	1990	2000	2010
In workforce		0.48	0.585	0.689	0.696	0.725
Working full time-full year		0.256	0.363	0.452	0.499	0.499
Share of working women with earnings at or above the xth percentile of the distribution of earnings among men working full time-full year, where xth percentile is:						
	50th percentile	0.051	0.071	0.127	0.18	0.194
	80th percentile	0.012	0.014	0.028	0.044	0.05
	90th percentile	0.006	0.007	0.0119	0.019	0.022
Share of women working full time-full year with earnings at or above the xth percentile of the distribution of earnings among men working full time-full year, where xth percentile is:						
	50th percentile	0.073	0.102	0.175	0.232	0.256
	80th percentile	0.017	0.019	0.037	0.057	0.067
	90th percentile	0.008	0.009	0.016	0.025	0.028
<i>Panel B</i>		<i>Sample: Women with a College Degree or More</i>				
		Year is:				
		1970	1980	1990	2000	2010
In workforce		0.611	0.734	0.817	0.807	0.811
Working full time-full year		0.282	0.437	0.548	0.588	0.602
Share of working women with earnings at or above the xth percentile of the distribution of earnings among men working full time-full year, where xth percentile is:						
	50th percentile	0.058	0.067	0.131	0.195	0.198
	80th percentile	0.01	0.012	0.026	0.043	0.049
	90th percentile	0.005	0.004	0.01	0.017	0.021
Share of women working full time-full year with earnings at or above the xth percentile of the distribution of earnings among men working full time-full year, where xth percentile is:						
	50th percentile	0.086	0.098	0.177	0.247	0.247
	80th percentile	0.016	0.017	0.035	0.054	0.062
	90th percentile	0.008	0.006	0.013	0.021	0.027
Share of women working full time-full year with earnings at or above the xth percentile of the distribution of earnings among men working full time-full year <i>in the same occupation</i> , where xth percentile is:						
	50th percentile	0.235	0.214	0.258	0.327	0.343
	80th percentile	0.058	0.047	0.060	0.090	0.102

Note: Data is 1970 to 2000 Census and 2008-2011 ACS. Sample in Panel A includes all women between 25 and 64 years of age; sample in Panel B includes all college-educated (or more) women between 25 and 64 years of age.

Table 2: Occupation-Specific Elasticity of Earnings to Hours Worked

<i>Panel A</i>			
<i>Sample of occupations: All</i>			
Year	Mean	Median	
1980	0.26	0.25	
1990	0.45	0.44	
2000	0.56	0.54	
2010	0.57	0.63	
2015	0.59	0.63	

<i>Panel B</i>			
<i>Sample of Occupations: 100 Highest Paying Occupations</i>			
Year	Mean	Median	
1980	0.32	0.28	
1990	0.54	0.52	
2000	0.67	0.69	
2010	0.64	0.64	
2015	0.66	0.64	

<i>Panel C</i>			
<i>Sample of Occupations: 50 Highest Paying Occupations</i>			
Year	Mean	Median	
1980	0.31	0.27	
1990	0.61	0.66	
2000	0.72	0.71	
2010	0.67	0.64	
2015	0.69	0.72	

<i>Panel D</i>			
<i>Sample of Occupations: 10 Highest Paying Occupations</i>			
Year	Mean	Median	
1980	0.24	0.41	
1990	0.57	0.66	
2000	0.62	0.54	
2010	0.52	0.78	
2015	0.55	0.78	

Note: Data is 1980 to 2000 Census, 2008-2011 ACS and 2012-2015 ACS. See text for details on construction of occupation-specific elasticity of earnings to hours worked. Occupations are ranked based on mean earnings of men 25-64 years of age working full time-full year in that occupation. Reported means and medians are weighted by the number of individuals working full time-full year in each

Table 3: Families of Employed College Educated Women

Family Charistic:	Married	Never Married	With Child	Married to College Educated Man	Elasticity of earnings to hours worked in husband's occupation (All husbands)	Husband's occupation has higher elasticity of earnings to hours worked than wife's (College-educated husbands)	Elasticity of earnings to hours worked in husband's occupation	Husband's occupation has higher elasticity of earnings to hours worked than wife's	
<i>Panel A</i>					<i>Year = 1980</i>				
Women's occupational earnings rank:									
Below top 100	0.68	0.19	0.51	0.45	0.25	0.72	0.24	0.72	
Top 100	0.61	0.23	0.42	0.45	0.27	0.46	0.27	0.46	
Top 50	0.58	0.24	0.39	0.44	0.28	0.51	0.28	0.51	
Top 10	0.54	0.30	0.34	0.47	0.24	0.64	0.24	0.66	
<i>Panel B</i>					<i>Year = 1990</i>				
Women's occupational earnings rank:									
Below top 100	0.69	0.17	0.54	0.44	0.42	0.73	0.43	0.74	
Top 100	0.64	0.21	0.46	0.44	0.46	0.45	0.48	0.48	
Top 50	0.63	0.22	0.43	0.45	0.48	0.36	0.50	0.40	
Top 10	0.62	0.23	0.43	0.46	0.47	0.47	0.48	0.52	
<i>Panel C</i>					<i>Year = 2000</i>				
Women's occupational earnings rank:									
Below top 100	0.67	0.18	0.50	0.40	0.54	0.73	0.55	0.74	
Top 100	0.65	0.20	0.47	0.42	0.59	0.47	0.60	0.50	
Top 50	0.64	0.21	0.46	0.44	0.61	0.41	0.63	0.45	
Top 10	0.64	0.21	0.46	0.50	0.60	0.55	0.60	0.59	
<i>Panel D</i>					<i>Year = 2010</i>				
Women's occupational earnings rank:									
Below top 100	0.66	0.19	0.49	0.40	0.58	0.67	0.56	0.67	
Top 100	0.65	0.19	0.47	0.43	0.60	0.51	0.59	0.52	
Top 50	0.65	0.20	0.47	0.43	0.61	0.43	0.61	0.46	
Top 10	0.68	0.19	0.49	0.51	0.57	0.58	0.55	0.62	

Note: Data is 1980 to 2000 Census, and 2008-2011 ACS. Sample includes all employed college educated (or more) women between 25 and 64 years of age. See text for details on construction of occupation-specific elasticity of earnings to hours worked. Occupations are ranked based on mean earnings of men 25-64 years of age working full time-full year in that occupation. Reported in each cell is the mean of the family charistic in that column for the group of women identified in that row.