

Discussion of “The Glass Ceiling and the Paper Floor:  
Gender Differences among Top Earners, 1981-2012” by  
Güvönen, Kaplan, and Song

Raquel Fernández  
NYU, NBER, CEPR, ESOP

June 2020

This is a nice paper that uses data from the US Social Security Administration (SSA) to provide empirical evidence on gender differences in earnings. It focuses on differences in gender representation at the top 0.1 percent and the next 0.9 percent of the earnings' distribution. It examines the persistence of an individual's presence in these top percentiles, how age and industry composition matter, and gives some feel for life-cycle dynamics, all the while contrasting the presence of women versus men.

The authors have access to a 10% representative sample of individual earnings histories from the SSA (constructed by selecting all individuals with the same last digit of a transformation of the social security number). This is a panel data set spanning 32 years: 1981-2012. There is basic demographic information available: age, sex, race, type of work (farm/nonfarm, employment/self-employment), and earnings. The latter consists of wages and salaries, bonuses, and exercised stock options as reported in Box 1 on a W-2 form. For most of their analysis, they select from the 10% sample all individuals who in that year are between 25 and 60 years old and whose annual earnings exceed a minimum threshold (equivalent to 13 weeks, full time, at 1/2 minimum wage).

Among the advantages of a panel set are that one can, for example, study earnings over a number of years to smooth out temporary fluctuations, ask questions about persistence, and examine measures of lifetime income. What this data set sacrifices, however, is any rich information about other characteristics of these individuals such as their education, marital status and children, spousal attributes, and occupation other than that captured by broad industry categories. This makes it virtually impossible to attempt to study what mechanisms are contributing to changes over time.

The paper's focus is on two findings: the increased female representation at the very top of the earnings distribution – the cracking of the “glass ceiling” – and the higher persistence of this status over time – the mending of the “paper floor.”

Turning first to the growth in the female share of top earners. Whereas women constituted 1.9% of the top 0.1 pct in 1981-85, by 2008-12, this had grown to 10.5%. Their representation in the next 0.9 pct over the same time period also quintupled: from 3.3% to 17.0%. A closer look at their figure, reproduced below (see Figure 1) for convenience, shows that the dynamics of this increase have been different. The share in the top 0.1 pct has barely budged since the early 2000s whereas the share of women in the next 0.9 pct has kept steadily increasing. An extension of the data to more recent years would have been useful here to see whether both trends have continued. It is interesting to note that for other percentiles of the earnings distribution in the top 20 percent, the share of women has been increasing throughout although women still make up less than 35 percent of these percentiles (see Figure 4A of the paper). More generally, it is useful to note that from 2004 onwards, the median earnings of full-time working women has been more or less constant at around 81% of their male counterpart's (see Figure 2 below).

Can the increase in the female share of top earners be attributed to the growth in female labor force participation over this period? At least mechanically, the answer is no in the sense that the share of women at the very top has grown substantially more than the share of women in the labor force.

Turning next to the paper floor, the authors emphasize the greater propensity in recent times for women to stay in the top of the earnings distribution. Whereas in 1981, the probability that a man transited out of the top 0.1 pct to the bottom 99 pct was 24%, the comparable statistic for a woman was 67%. By 2011, the gender gap is almost closed and the persistence has increased for both: the probability of a woman transiting became 8.1% versus 6.6% for a man. The authors discuss another feature which is worthy of note. The five-year transit probabilities depicted in the bottom panel of their figure 7 show that gender parity in permanence was reached quite early in the sample period. By 1994, both genders have essentially the same probability of remaining in the top 0.1 pct. A striking feature over this time period, however, is the initially large difference in the rate in which these top earning females leave the sample as compared to males. In the mid 1980s, in fact over 20% of these women are exiting the sample. Is this the result of aging out (per the rules followed by the author in sample construction) or because women are behaving differently, e.g., taking earlier retirement or leaving to take care of a family member? The authors' data should allow them to answer the first part of the question but not, unfortunately, the mechanisms underlying differential behavior. It would also be of interest to know whether these women reenter the labor force, and, if so, at what percentile of the earnings distribution. This question should be answerable with the authors' data.

Perhaps the most interesting part of the paper is the cohort analysis. The cohort picture below (Figure 3) graphs the share of women at the top of *their own cohort's* labor earnings' distribution. It paints a mixed picture regarding the march towards gender equality at the top. On the one hand, each successive cohort is entering with a larger share of women at the top. On the other hand, the youngest 3-4 cohorts - those born in each 5 year period starting from 1934 to 1974 - show a rapidly declining representation at the top as they age. This is true for both the top 0.1 pct and the second 0.9 pct, though less markedly in the case of the latter. Why is the female share of younger cohorts decreasing? A few possible explanations are (i) men are promoted more rapidly; (ii) Men change jobs in an advantageous fashion more frequently (either because they obtain better offers or are more able to take advantage of the ones they receive); (iii) Women are exiting or reducing their hours as they age and take on a greater share of family responsibilities. This seems like a fruitful question to explore in greater depth in future work. In fact, it would have been interesting to have conducted the entire analysis from a cohort perspective.

Although the authors cannot observe occupations, they are able to use the SIC code to construct 13 industry groups. One unsurprising finding is that the representation of finance in the top .1% has increased markedly. Overall, there appears to be little variation across industries in the gender composition of the top earners during the latest set of years 2008-2012. This points at a common issue affecting the lack of women at the top rather than their under-representation in particularly high-earning industries. This makes it even more important to understand the occupational structure by gender as what may explain the under-representation of women at the top is a disproportionately higher employment in lower-earning occupations (e.g., if they are less likely to work, say, as a top executive in an industry).

It is instructive to end with some reflection on what factors might be responsible for the low representation of women in the top percentiles of the earnings distribution. As noted previously, although the authors focus on the very top, the under representation of women is true for the top half of the distribution. Therefore there may be common factors that affect women generally rather than simply impact women at the very top of the labor earnings' distribution.

A life-cycle perspective may be useful in terms of thinking about possible contributors. Although we have seen a significant reversal of the gender gap in education, the share of women at the top has been fairly stagnant over the 2000s. One contributor to this is likely to be their field of study as this helps determine their occupation. Across OECD countries, on average boys score slightly better than girls in the PISA (Program for International Student Assessment) math test whereas girls on average score *much* better than boys in the reading test. [Breda and Napp \(2019\)](#) use 2012 PISA scores of 15-year-old students in OECD countries (a sample of around 300,000 students) as well as answers to questions regarding their intention to study math-intensive fields in the future. They show that both girls and boys seem to respond to their *relative* advantage in performance across field, instead of their absolute advantage and the difference in relative rewards. This means that a girl who is good at math but even better in reading may favor the humanities because she perceives herself as a “verbal” person, leading girls to disproportionately opt out of mathematical fields. This response in high school contributes to the large STEM gender gap which is then amplified in college and resulting choice of occupation.

Turning to the work stage of women's lives, studies by [Gayle, Golan and Miller \(2012\)](#) on top executives in publicly traded firms and [Bertrand, Goldin and Katz \(2010\)](#) on MBAs from Chicago Business School show that for these select high-earning groups there are large differences in exit rates from the labor force between women and men. Bertrand et al. use survey data to study the dynamics of wages of University of Chicago MBAs from the classes of 1990-2006. They find that a man in the ninetieth percentile of the male distribution earns over \$1 million at 10-16 years out as compared to \$438 thousand in the ninetieth percentile of the female distribution. Even after accounting for loss in work experience and weekly hours worked, taking any time out of work results in a 23 log point reduction in annual earnings. Furthermore, it is not the case that women suffer a greater penalty for taking time away from work. In fact, the authors find that the reduction in earnings is even larger for men. The fraction of women, however, who take time out of work is much larger: 27% of women as opposed to 11% of men in their sample have a work interruption. Relatedly, recent work by [Kleven et al. \(2019\)](#) shows that, in general, women and men's earnings diverge when they have children. Using an event study methodology, they show that whereas men's earnings barely budges upon the birth of a child, that of a woman's falls (presumably a combination of factors ranging from lower labor force participation, hours worked, and career consequences/choices) resulting in a child “penalty” of around 31%. These papers all point to the greater need to understand the interaction between family and work in order to identify the mechanisms responsible for the gender gap in labor earnings.

I conclude with two main suggestions. First, although the authors' SSA data does not allow them to

study education, marital status, or fertility, it does have information on race and ethnicity. The paper would be greatly enriched if the study were broadened to include an analysis of minority groups at the top. How has the share of Blacks, or Hispanics, or Asians evolved over time. Within these groups, how do the results differ by gender or by area of the country? These seem like the sort of questions that the data is well-suited to address. A second suggestion is to exploit more fully the panel nature of the data which is where its strength lies. In particular, it would be great to look beyond one-period transitions to paint a fuller picture of the labor earning dynamics of these top earners, especially by cohort. I look forward to seeing this work in the future.

Figure 1: Share of Women among Top Earners

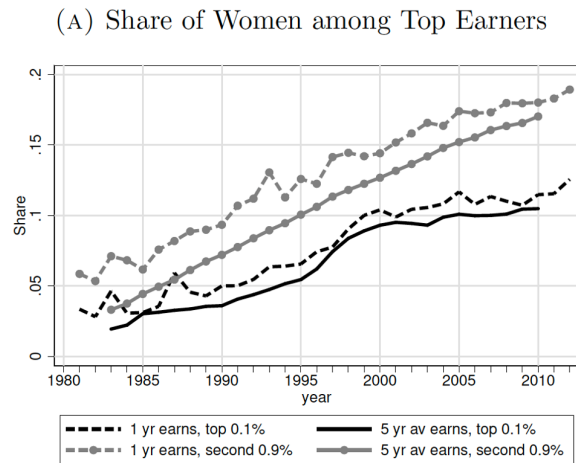
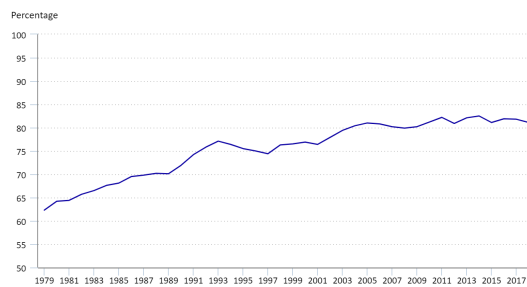


Figure 2: Women’s Earnings as a Share of Male Earnings

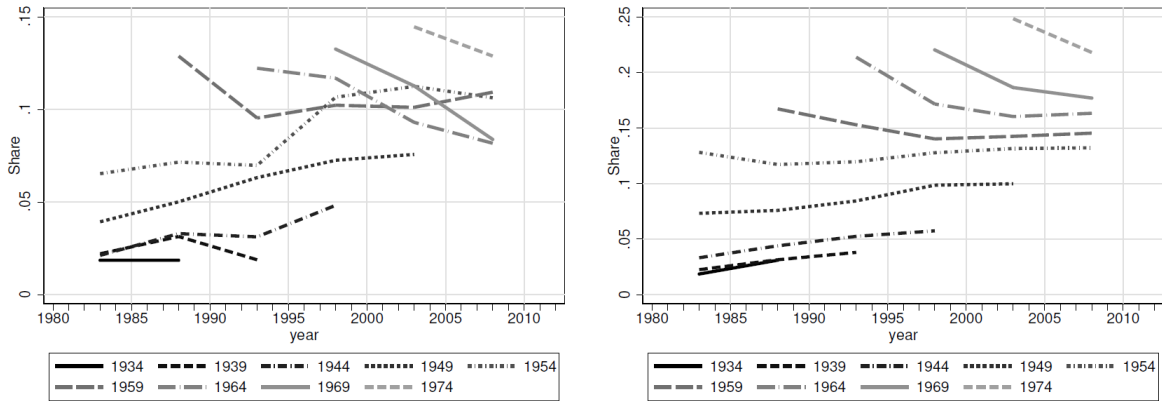
Chart 1. Women’s earnings as a percentage of men’s, for full-time wage and salary workers, 1979–2018 annual averages



Hover over chart to view data.  
 Note: Percentages are calculated from annual averages of median usual weekly earnings for full-time wage and salary workers.  
 Source: U.S. Bureau of Labor Statistics, Current Population Survey.

Figure 3: Share of Women in Top pct of Cohort Distribution

(c) Share of women among top 0.1pct by cohort (D) Share of women among second 0.9pct by cohort



## References

- Bertrand, Marianne, Claudia Goldin and Lawrence F Katz. 2010. "Dynamics of the gender gap for young professionals in the financial and corporate sectors." *American economic journal: applied economics* 2(3):228–55.
- Breda, Thomas and Clotilde Napp. 2019. "Girls' comparative advantage in reading can largely explain the gender gap in math-related fields." *Proceedings of the National Academy of Sciences* 116(31):15435–15440.
- Gayle, George-Levi, Limor Golan and Robert A Miller. 2012. "Gender differences in executive compensation and job mobility." *Journal of Labor Economics* 30(4):829–872.
- Kleven, Henrik, Camille Landais, Johanna Posch, Andreas Steinhauer and Josef Zweimuller. 2019. Child penalties across countries: Evidence and explanations. In *AEA Papers and Proceedings*. Vol. 109 pp. 122–26.