Can Women "Have It All"? The Trade-Off Between Social Tasks and Workplace Flexibility

Savannah Noray*
Harvard University

January 31, 2025

PLEASE CLICK **HERE** FOR MOST CURRENT VERSION

Abstract

Do the jobs best suited to women's talents offer the flexibility that they need? I document that temporal flexibility and social skill intensity are inversely correlated across occupations. Guided by an occupational choice model, I show evidence that this correlation generates a trade-off between women's *preference for flexible work hours* and their *comparative advantage in social tasks.* As a consequence, women's relative labor market returns to both flexibility and social skills are attenuated, which in turn widens the gender wage gap. Event-studies around first births – a shock to women's relative demand for flexibility – help confirm that gender gaps in inflexible jobs are driven by women's time constraints rather than other factors.

*savannahnoray@g.harvard.edu. I thank Michela Carlana, Claudia Goldin, and Gordon Hanson for their guidance and encouragement. I thank Patricia Cortes, David Deming, Will Dobbie, Ed Glaeser, Jessica Goldberg, Chinhui Juhn, Larry Katz, Claudia Olivetti, Ryan Oprea, Heather Sarsons, Mark Shepard, Tom Vogl, and participants in the 2023 NBER Summer Institute Gender in the Economy, Harvard Graduate Student Labor Workshop, and Harvard Kennedy School Economics and Social Policy Workshop for helpful comments and discussions. Special thanks to Kadeem Noray and Rafael Proença for valuable extended discussions about this paper. This work was supported by the Bill and Melinda Gates Foundation through the NBER Dissertation Fellowship Program on Gender in the Economy. All errors are my own.

1 Introduction

When workers face barriers that prevent them from fully utilizing their unique skills, they produce and earn less than they otherwise would, which may lead to misallocated talent and lower economic output overall (Guvenen et al., 2020; Hsieh et al., 2019). Evidence suggests that women excel in social skills, which are becoming increasingly valued in the labor market (Deming, 2017; Cortes et al., 2023; Cortés et al., 2024). Yet, a well-documented barrier to women's advancement in certain careers is the requirement to work long, irregular, or particular hours, which conflicts with their preference for flexible work schedules (e.g., Goldin (2014); Cortés and Pan (2023)). Do these inflexible (or "greedy") work environments affect women's ability to enter and excel in roles that utilize their talent? In this paper, I show that *exactly* the jobs that play to women's skills – social task intensive occupations – are also less flexible. Women therefore face a trade-off: in order to be in a job that rewards their skill set, they will have to forgo flexibility, or vice versa. As a result, inflexible work environments exacerbate the gender wage gap more than previously understood.

I draw on work in psychology to show that women are more skilled in "social perceptiveness," or the ability to accurately infer others' mental states. I then document that jobs requiring these skills also tend to be less flexible. To explore the implications of this correlation between job attributes, I develop a model of occupational choice that yields four predictions. First, the inverse relationship between flexibility and social intensity across occupations obscures their positive effects on women's relative wages and employment. Second, relative to men, women's social skills will be less predictive of their employment in less flexible jobs. Third, flexibility will have a more pronounced positive impact on gender gaps among parents. Fourth, the marginal productivity returns to adopting technologies or practice that enhance flexibility will be highest in social occupations. I test these predictions using the American Community Survey (ACS), National Longitudinal Survey of Youth 1979 (NLSY79), Current Population Survey (CPS), and American Time Use Survey (ATUS).

As *prima facie* evidence that women face this trade-off, consider a simple but surprising set of facts. Among full-time, college-educated workers in high-income jobs, the gender wage gap is larger in occupations where long, irregular, or particular hours are incentivized (Goldin, 2014). Figure 1a replicates this finding by plotting women's relative earnings against an index that proxies for an occupation's tendency to be structured in this way. In less flexible occupations, the price women pay (in terms of lower wages) to customize their schedules is higher, thus generating a larger gender wage gap. Given that their preferences come at a high cost in these occupations, we would expect women to be disproportionately discouraged from entering these types of jobs. However, according to this measure of flexibility, women's relative employment is

higher in *less* flexible occupations (Figure 1b). In tech and science occupations such as computer programmers or engineers, workers can often choose when and how much they work (or even where they work). Yet, women's relative employment is higher in occupations such as financial managers, lawyers, or physicians, where workers are disproportionately rewarded for working longer hours or being available continuously for their co-workers, clients, or patients.[1] This puzzle suggests that less flexible occupations tend to have *some other attribute* which draws women in, despite the larger pay-cut they receive to work flexible hours. As I will demonstrate, this other attribute is an occupation's social task intensity, which acts as a countervailing force on women's wages in these less flexible jobs.

Gender differences in skills and preferences for social tasks suggest that some occupations are better-suited for women. In a meta-review of gender differences along various dimensions, men's interests in things and women's interests in people ranks among one of the largest and persistent gender differences (Hyde, 2014). Moreover, psychologists have documented that on average, relative to men, women excel in accurately reading nonverbal elements of communication, such as body language or tone. This leads them to excel in "social perceptiveness," or the skill of being aware of others' reactions and understanding why they react the way they do (Hall, 1978, 1990; Greenberg et al., 2023). These skills are useful in jobs requiring interactions with people, particularly jobs requiring relatively complex social interactions, such as psychologists, lawyers, or managers. In order to be successful in these jobs, workers must infer accurately how someone is feeling or what someone is thinking. Doing so is often a complicated, hard-to-specify function of facial expression, tone of voice, body language, and the social context in which these things are done. For example, a good psychiatrist will take into account that social stigma may prevent clients from admitting to certain thought patterns. Or, a good manager will consider how certain workers get along with one another before assigning them to the same team.

Yet, as I document, jobs that reward workers' social perceptiveness also tend to be less flexible. While this correlation may arise for a variety of reasons, one explanation is that jobs involving social tasks require workers to build and maintain "relationship-specific capital" with clients or co-workers. As a result, workers are rewarded for being "on call" for others' needs and therefore take a pay cut to work flexible hours. In contrast, jobs like computer programming or accounting are centered around producing deliverables – *when* the work is done is less directly tied to the quality of the output. Workers can run code or crunch numbers on their own schedule and do not necessarily need to interface with clients or team members at specific times. Therefore, insofar as women value both flexibility and social tasks, they will trade-off between these two

[1] Relative to women, men are over three times as likely to become computer scientists or programmers and approximately 1.5 times as likely to become financial managers, lawyers, or physicians.

attributes when selecting a job.

To build intuition and explore the implications of this trade-off, I develop a simple model of occupational choice. In the model, less flexible tasks are ones for which labor supply cannot easily be substituted across time periods. Workers with "particular" preferences for *when* they work are attracted to jobs in which the temporal elasticity of substitution across time periods is higher. Yet, workers must also consider the fact that jobs high in social intensity are also less flexible. The fact that these job attributes are bundled together disproportionately constrains female workers, whose wages are more positively impacted by both temporal substitutability and social tasks. The model delivers four predictions. The first is that social task intensity will have a positive impact on women's relative wages and employment, but this impact is offset by the countervailing *negative* effect of lower flexibility that is bundled with social tasks (and vice versa). The second is that relative to men, women's social skills will be less predictive of their employment in less flexible jobs, despite the fact that such jobs reward social skills. Third, because childbirth shocks women's relative demand for flexibility, the positive impact of flexibility on women's relative wages and employment is more pronounced among parents. Fourth, the marginal productivity gains from enhancing flexibility – and thus increasing women's relative employment – will be highest in social occupations.

Consistent with the model's first prediction, using ACS data, I find that women's employment and wages in highly social (flexible) occupations is muted by the lack of flexibility (social intensity) in these occupations.[2] In other words, women's employment and wages are positively related to both flexibility and social task intensity, but this relationship is amplified – or in some cases only emerges – when controlling for both attributes simultaneously. These results resolve the initial puzzle: while women's relative employment is higher in less flexible occupations, this relationship is completely explained by the social perceptiveness intensity of these jobs. The results are generally robust to limiting the sample to the college educated in high-income occupations (as in Figure 1), using alternative measures of flexibility, and controlling for other job attributes. Moreover, supporting the notion that women's time constraints are a key mechanism, the negative correlation between women's relative wages and employment and an occupation's inflexibility is more pronounced among parents.

I leverage data on individuals' skills in the National Longitudinal Survey of Youth 1979 (NLSY79) to test the second prediction. Consistent with less flexible jobs demand-

[2]Women may be drawn to social tasks either because of their relative ability or their stronger preference for socializing. Insofar as the latter mechanism is operative, firms may pay women *less* than men in social jobs due to compensating differentials. Therefore, the positive correlation between social intensity and women's relative *wages* suggests that women's comparative advantage in social tasks, rather than preferences for socializing, is the dominant mechanism influencing their entry into social task intensive jobs.

ing more social skills, I find that high social skill workers are more likely to be employed in jobs with low flexibility. However, this correlation is muted for women. I find a similar pattern for cognitive skills, which are also rewarded more in less flexible occupations. In other words, despite being at the same social and cognitive skill level, women are less likely to sort into inflexible jobs despite the fact that such jobs reward these skills. And, again consistent with women's time constraints driving these results, this gender gap in skill-based sorting is partially amplified among mothers.[3]

I then test the third prediction of the model, which states that the positive impact of flexibility on women's relative wages and employment is amplified among parents. Less flexible jobs may have unobserved attributes (e.g. masculine culture, gender discrimination) that deter women or lower their wages. Therefore, complementing the tests for heterogeneity by parental status in the previous results, I use childbirth to proxy for a shock to women's relative demand for flexibility to pinpoint the mechanism. Using event-studies in the NLSY79 and pseudo difference-in-differences in the CPS, I test how this shock impacts their employment and wages in less flexible versus more flexible jobs. Consistent with the proposed mechanism, I find that the "child penalty" on women's relative employment and wages is larger in less flexible occupations (vis-à-vis flexible ones).[4]

Finally, I present evidence consistent with the notion that the marginal returns to enhancing flexibility are highest in social occupations (the fourth prediction). There are several well-documented case studies of innovation leading to more flexible work structures (Goldin and Katz, 2011, 2016; Goldin, 2021).[5] The trade-off suggests that these innovations will be most beneficial in social occupations, as they help firms attract, retain, and better utilize female talent otherwise constrained by inflexible work. I use two different proxies for an occupation's growth in flexibility: the change in the child penalty estimated in the CPS and, to complement this "data-driven" approach, the growth in work-from-home (WFH) arrangements based on ATUS data.[6] In both cases, I find that among inflexible occupations, highly social occupations have experienced the most growth in flexibility. Although this evidence is only suggestive, it aligns with the idea that firms in social occupations have the greatest incentive to en-

[3] I only find heterogeneity by parental status along the cognitive skill dimension.

[4] Following past literature, I use the "child penalty" terminology, which is shorthand for the disproportionate effect of children on women's labor market outcomes (e.g. Kleven et al. (2019)). The terminology is not intended to reflect a normative judgment. Child penalties could be driven by a variety of mechanisms, such as traditional gender norms or women's preferences for spending time with children.

[5] For example, Goldin and Katz (2016) document how pharmaceutical work became more "family friendly." This is attributed to a confluence of factors, such as greater standardization in the pharmaceutical industry and improved technology for storing information on customers. These factors allowed workers to become better substitutes for one another, therefore reducing the wage penalty for working shorter hours.

[6] Harrington and Kahn (2023) argue that location flexibility due to WFH arrangements likely enhances temporal flexibility.

hance flexibility, thereby capitalizing on female talent.

Building on the seminal work of Goldin (2014), a large body of literature examines how the lack of workplace flexibility – broadly defined – affects gender gaps in occupational sorting and wages. My paper contributes to this literature in a few ways.

First, I provide additional evidence in favor of the existing hypothesis that women's higher willingness to pay for flexibility drives gender inequality in the labor market. Experimental and quasi-experimental evidence shows that women are more willing to accept lower wages in exchange for certain flexible work arrangements, such as the option to telecommute or avoid irregular schedules (Mas and Pallais, 2017, 2020; Wiswall and Zafar, 2018; Maestas et al., 2023; Bolotnyy and Emanuel, 2022). In line with the compensating differentials framework put forward by Goldin (2014), descriptive evidence suggests women's (especially mothers') greater demand for flexibility leads to a larger wage gap within inflexible occupations (Cubas et al., 2023; Goldin et al., 2024) and, along the extensive margin, may prevent entry into these occupations (Cortés and Pan, 2016; Juhn and Rubinstein, 2022; Erosa et al., 2022; Denning et al., 2022). This evidence suggests that the price one pays to customize work hours varies systematically by work structure and therefore impedes women's advancement in careers with this structure. Quasi-experimental literature focused on specific subset of professions or dimensions of flexibility provides converging evidence to support this notion. (Wasserman, 2023; Ciasullo and Uccioli, 2024; Cortés and Pan, 2019; Kleven et al., 2019; Pertold-Gebicka et al., 2016; Bütikofer et al., 2018).[7] I extend this quasi-experimental evidence by explicitly testing whether child penalties in employment and wages are higher in occupations that, according to a measure of the occupation's typical work environment, are pre-disposed to be less flexible along a variety of dimensions.[8]

Second, this paper helps establish what (and how much) society stands to gain from resolving the tension between greedy work environments and women's demands for flexibility. Specifically, I identify an additional cost of inflexible work: the loss of women's returns to their comparative advantage in social tasks. Without accounting for this cost, we risk understating the positive impacts of a range of policies on women's relative wages. For example, neglecting the trade-off understates the benefits of policies that reduce women's relative time constraints (e.g., subsidized childcare as studied in Humphries et al. (2024)), reduce the price of flexibility in social task in-

[7] Long hours requirements by certain medical specialties discourage women from entering those fields (Wasserman, 2023), irregular schedules characterizing some occupations reduce mothers' labor supply (Ciasullo and Uccioli, 2024), and convex returns to hours worked among highly skilled occupations contribute to larger gender gaps that can be reversed with more access to cheaper household services (Cortés and Pan, 2019). Child penalties are lower in the public sector, STEM, or medicine when compared to management, business, and law (Bütikofer et al., 2018; Kleven et al., 2019; Pertold-Gebicka et al., 2016).

[8] This part of my analysis is empirically adjacent to Bang (2021), who studies how occupational flexibility impacts couples' specialization in domestic versus market work after childbirth. I discuss how my analysis differs in Section 7.

tensive jobs (e.g., work-from-home technologies as studied by Harrington and Kahn (2023)), or reduce gendered specialization in social versus analytical skills (e.g., educational interventions as in Alesina et al. (2024)). Yet, the trade-off not only affects the relationship between inflexible work environments and *gender inequality* – a major focus of the literature – but also helps us understand how inflexibility influences talent allocation and, consequently, *aggregate productivity*. A growing literature suggests there may be aggregate productivity gains to individuals sorting into jobs according to their comparative advantage (Hsieh et al., 2019). Building on this intuition, Erosa et al. (2022) estimate substantial productivity gains from reallocating husbands and wives to jobs based on their relative abilities rather than limiting women to less time-intensive roles. My findings suggest that the average woman is more skilled in the time-intensive job than the average man, therefore magnifying these productivity gains.

In this literature, the work that is most similar to mine, Juhn and Rubinstein (2022), explores an analogous trade-off. They provide evidence that women's comparative advantage in abstract tasks (as opposed to routine or manual tasks) conflicts with their comparative disadvantage in working long or unusual hours, which abstract tasks often require. Consequently, women experience more occupation-skill "mis-match," which has impeded the convergence of the gender wage gap. I complement their work in several ways. First, drawing from work in psychology, I focus on social tasks, a specific type of abstract task. Second, I analyze the within-occupation wage gap, which allows me to disentangle whether women's attraction to social tasks are primarily due to their preferences or their skills. Finally, I leverage event studies around births to verify the mechanisms at work, and I provide evidence that firms adjust in the long-run to this trade-off.

Third, my paper provides a way of understanding why certain occupations may eventually become more flexible. Several well-documented case studies document that policies or innovations have enhanced flexibility in various sectors (Goldin and Katz, 2011, 2016; Goldin, 2021; Harrington and Kahn, 2023; Wasserman, 2023). Yet these case studies do not necessarily provide a systematic framework for understanding whether we should expect firms to adjust and, if so, *which* firms we should expect to adjust. My findings suggest that firms in social intensive sectors may have the most to gain from enhancing flexibility, and that they seem to do so in the long-run.

Fourth, my findings help explain why women do not always opt for more flexible jobs despite valuing these amenities. Mas and Pallais (2020) highlight for many of the various dimensions of flexibility (e.g., work-from-home), women are often employed in jobs which do *not* offer these amenities despite the fact that they value them.[9] They resolve this by suggesting that different aspects of flexibility may be valued differently,

[9]Echoing this finding, Barrero et al. (2023) show that women are only slightly more likely to work from home than men.

may be inversely correlated with wage premiums, or may be inversely correlated with one another. I show that this puzzle persists even when leveraging a multidimensional measure of flexibility. I then show how an occupation's task content, which both influences job choice and is correlated with flexibility, may help resolve the puzzle.

Finally, this paper also contributes to the growing literature exploring the rising returns to social and non-cognitive skills in the labor market (e.g. Atalay et al. 2020; Deming 2017; Edin et al. 2022). Other work documents that women benefited from this rise in social tasks or service sector jobs (Cortes et al., 2023; Cortés et al., 2024; Ngai and Petrongolo, 2017). The results in this paper suggest these gains could be limited to the degree that the prevalence of greedy work has grown alongside these trends.

I begin by describing the data used throughout the paper in Section 2. Then I document two key stylized facts (Section 3): women are more socially skilled and less flexible jobs tend to be social task intensive. I then present the conceptual framework in Section 4. The subsequent sections test the four predictions of the model (5-8).

2 Data

Below I describe the data sources used throughout this paper. Detailed descriptions are provided in the Data Appendix (Appendix A).

2.1 Occupation-level attributes

Throughout the analysis, I crosswalk individuals' occupations to a set of 320 3-digit occupation codes which form a balanced panel of occupations over time (Autor and Dorn, 2013; Deming, 2017; Hanson, 2021). These time-consistent codes allow me to hold constant how I am measuring job attributes across datasets and time periods. The key attributes I measure are an occupations' flexibility and skill requirements, for which I primarily rely on the U.S. Department of Labor's Occupational Network (O*NET) database. The O*NET is a survey administered to a sample of workers (or occupational experts) in each occupation and is commonly used in research on tasks and skills in labor markets (e.g. Acemoglu and Autor (2011)). The survey covers a wide-range of questions about the worker's occupation, such as the typical tasks required for a certain job, how important it is to have certain skills, or features of the typical work environment.

2.1.1 Flexibility

Flexibility is a complex and multi-dimensional concept. The ideal measure of flexibility would capture the *price* that workers pay (in terms of lower wages) to customize their work hours *however they desire*. Workers may desire the flexibility to customize

their hours along a variety of dimensions, and the relevant dimensions may vary by individual (or even within-individual over time). For example, workers may desire to work fewer hours, particular hours, or predictable hours.

To proxy for this ideal measure, I follow Goldin (2014) and use an index intended to capture *underlying features of production* that may generate a higher price of flexibility. The idea, captured by the index, is that firms will find it more costly to offer workers flexibility along any dimension when tasks are time-sensitive and workers are not easily substitutable for one another. These features of a work environment will vary systematically by occupation. Accordingly, the index is the standardized average of five dimensions (each of which are themselves standardized) from the O*NET: time pressure, contact with others, establishing and maintaining relationships, unstructured work, and the freedom to make decisions. The first dimension captures time sensitivity in task completion. The second two dimensions capture the necessity of being in contact with people at certain times. The final two dimensions capture the level of worker discretion: the more freedom over one's task set, the less likely others can easily substitute for you in your absence. Occupations that are ranked highly in this "inflexibility" index include lawyers, CEO's, and managers. Workers in occupations with less flexibility (e.g. lawyers) pay a larger amount to be in specific positions or jobs which grant them more flexible hours (e.g. in-house counsel versus partner at a large law firm). In contrast, workers in more flexible occupations can move into such positions with little impact on their wages.

Other literature has incorporated more "data-driven" measures, such as the share of men in an occupation working over 50 hours per week (e.g. Cortés and Pan (2016)). Relative to the O*NET index, it is possible that these approaches reduce some measurement error by ensuring some real-world manifestation of flexibility (or lack thereof). The O*NET index is my preferred measure of flexibility for two main reasons. First, workers' *desired* labor supply is imperfectly reflected in their labor supply, which may be constrained or a function of labor demand. Therefore, these data-driven measures may inadvertently reflect general equilibrium features of the labor market unrelated to flexibility. For example, workers may work longer hours because the wages in an occupation are very low (or very high). Second, at the end of the paper, I explore how occupations evolve over time, with a focus on occupations which are both high in social intensity and historically predisposed to be less flexible. For that analysis, the flexibility index is more suitable as it reflects an occupation's inherent tendency toward flexibility based on its fundamental production features.

Nevertheless, I show that the results are generally robust to two alternative measures incorporating these other approaches. First, I follow Goldin et al. (2024) in creating a binary variable indicating if an occupation is in the top tercile of both the O*NET index and the long hours measure. I refer to this measure as "top $Inflexibility_o$ and

Overwork." Second, I construct a standardized index of three alternative measures based on past literature that does not rely on the O*NET: a measure of coordinated work hours (Cubas et al., 2023), the estimated elasticity of earnings to hours worked (Cortés and Pan, 2019), and the prevalence of long hours (Cortés and Pan, 2016).[10] I refer to this measure as $AltFlexPremIndex_o$. Reassuringly, both alternative measures are highly correlated with the O*NET index (ρ of 0.60 and 0.49, respectively).[11]

2.1.2 Skill requirements

To measure an occupation's social skill requirements, I use the O*NET's measure of "social perceptiveness," which captures the degree to which a worker must be "aware of others' reactions and [understand] why they react as they do." I focus on social perceptiveness (rather than a broader measure of social skills as in Deming (2017)) because it most closely reflects the type of social skills in which women seem to excel, which I will discuss in Section 3.1.[12] According to O*NET documentation, a job where workers are required to "notice that customers are angry because they have been waiting too long" likely requires low levels of social perceptiveness, while a job where workers must "counsel depressive patients during a crisis period" requires a high level of social perceptiveness. Occupations that rank highly in social perceptiveness include counselors, physicians, social workers, and managers.

I follow Deming (2017) to measure occupations' analytical skill requirements and routine task intensity, which are primarily used as controls throughout the analysis. Analytical intensity is measured via an index of components reflecting the importance of mathematical knowledge and reasoning in an occupation. These skills are most likely to be required in science and technology jobs, but also are relevant for jobs in the business (e.g. financial manager) and health sectors (e.g. physicians). Routine task intensity is measured via an index capturing the propensity for an occupation's tasks to be subject to repetitive, precise processes, such as typing or working in an assembly line. Occupations scoring highly in the routine measure include plant operators, telephone operators, and bank tellers.

[10] Cubas et al. (2023) uses the American Time Use Survey to measure the concentration of work hours within the 8am to 5pm window at the occupation-level. Cortés and Pan (2019) estimate the elasticity of earnings with respect to hours by regressing log earnings on log hours. The Cortés and Pan (2016) measure is simply the share of men in a given occupation working over 50 hours per week. See Appendix A for more details.

[11] It is worth noting that different measures may more heavily weight certain dimensions of flexibility (e.g. long hours versus particular hours), which could, in turn, be correlated with different types of tasks and therefore imply different trade-offs worth exploring. For example, Juhn and Rubinstein (2022) focus on the correlation between long or unusual hours (measured in the ATUS) and abstract tasks.

[12] Social perceptiveness is highly correlated with the Deming (2017) composite measure ($\rho = 0.93$). I test the robustness of the results when using the composite measure in Section 5.2.

2.2 Other data sources

2.2.1 American Community Survey, 2015-2019

To study the implications of the model on gender gaps in employment and wages across occupations, I rely on the 2019 5-year American Community Survey (ACS) 5% sample where I observe individuals' wages and occupation choices. Employing similar sample restrictions as Goldin (2014), I focus on native-born civilian individuals aged 25-64 who are employed full-time, full-year (at least 40 weeks, at least 35 hours per week) and report positive, nonmissing annual earnings (wage income plus business and farm income).

2.2.2 National Longitudinal Survey of Youth: 1979 Cohort, 1979-2019

I use data from the Weekly Work History files in the National Longitudinal Survey of Youth for the 1979 Cohort (NLSY79) to build a person-year panel of labor market activity (e.g., employment, occupation, and wages) covering years 1979-2019. These data are useful for two separate analyses. First, I study how individuals' social and cognitive skills influences occupation choices differently by gender. Second, I study how childbirth impacts gender gaps in employment and wages for workers in flexible versus less flexible jobs.

Following Deming (2017), I proxy for cognitive skills with the respondent's score on the Armed Forces Qualifying Test (AFQT) residualized by the age at which the respondent took the test. And, to proxy for individuals' social perceptiveness skills, I take the average of three standardized variables: (1) young adulthood self-reported sociability, (2) self-reported retrospective sociability at age 6, and (3) the number of extracurricular activities in which the respondent participated in high school (e.g. yearbook, student government, newspaper). Given that gender equal access to athletic activities in high school was not widespread for this cohort, I diverge from Deming (2017) and do not include a fourth component indicating whether the respondent participated in an high school athletics.[13]

Similar to Guvenen et al. (2020), I focus on individuals in the core sample starting when they are least 22 years old, have signaled strong labor force (e.g. two consecutive years of full-time full-year), and are likely done acquiring their education. I drop observations lacking complete information on labor market activity or individual skills. To ensure some balance in the event studies around births, among individuals who have a child, I only include those who are observed in the period just prior to birth and

[13] Athletic activities were not widely accessible to girls in the early 1980's, when most of the NLSY79 cohort was in high school. Title IX, intended to prohibit sex-discrimination in educational institutions, was first passed in 1972. The U.S. Department of Education Office of Civil Rights later began initiatives specifically promoting equality athletics in 1979. Deming (2017)'s social skill index is highly correlated with my modified version ($\rho = .94$).

at least one period after. My final sample is comprised of 3,396 individuals and 103,264 person-year observations.

2.2.3 Current Population Survey, 1979-2019

I use the Current Population Survey (CPS) pooled across years 1979-2019 to build a synthetic panel of individuals before and after the birth of their first child. I use these data for two separate analyses. First, I estimate child penalties in each occupation based on a pseudo difference-in-differences approach (as employed by Kleven (2023)). Second, I repeat this analysis before and after 1995 in order to construct a proxy for the growth of motherhood friendliness within an occupation over time.

I follow Kleven (2023) to build my sample. I focus on individuals aged 20-55 in the CPS basic monthly files and Annual Social and Economic Supplements (ASEC) pooled across years 1972-2019. I identify individual's approximate age at first birth based on their own age and the age of the eldest child living with them. I then match these individuals to five pre-birth surrogate observations. To do this, I retrieve outcomes of childless individuals in the same birth cohort who share the same demographic characteristics: gender, age, year, education category, marital status, race, and state. For example, consider a woman observed in year 2000 at age 35 whose eldest child is age 5. To retrieve her surrogate observation in the year prior to birth ($k = -1$), I locate a childless, 30-year old woman in the year 1995 who shares the same demographic characteristics. When there are multiple people with this set of demographic characteristics, I use the average of the outcome for everyone in the cell to form the surrogate observation.

2.2.4 American Time Use Survey, 2003-2019

Finally, I use the American Time Use Survey (ATUS) to measure which occupations have adopted work from home (WFH) arrangements, an additional way of measuring how firms may be adjusting in the long-run. The ATUS is a nationally representative survey conducted annually since 2003. It asks respondents to report the location and duration of each of their activities over a 24 hour period. Following Harrington and Kahn (2023), I focus on hours during the 8am - 5pm window and I limit the sample to respondents aged between 22 and 55, who are employed, whose reported time diary covers a non-holiday weekday, and who contribute at least 5 hours of work that day.

3 Two stylized facts

3.1 Women are more socially skilled than men

Gender differences in social skills – especially "social perceptiveness" – have been well-documented. While men and women tend to exhibit average differences in a variety of dimensions, many differences are small or do not persist across studies (Hyde, 2005). Yet women's greater expressed interest in people as opposed to things is among one of the most robust gender differences (Hyde, 2014; Bertrand, 2020). Therefore, it is unsurprising that women tend to excel in *social perceptiveness* – the ability to accurately infer how someone is feeling and react appropriately. In particular, experimental research in psychology finds that relative to men, women can more accurately *decode nonverbal cues*, such as body language, tone of voice, and facial expression (Hall, 1978, 1990). Developmental psychology rationalizes these differences by positing that men tend to understand the world through explicit, well-defined cause-and-effect systems, while women more willingly infer others' mental states based on subtle cues which enables them to interact more effectively with others (Baron-Cohen, 2002; Greenberg et al., 2018).

In line with this evidence, Figure 2a depicts gender differences in the "Reading the Mind in the Eyes" test (RMET) scores reported by Greenberg et al. (2023), which capture one's ability to read facial cues. In the test, respondents review a series of pictures of peoples' eyes and are asked to guess how the person is feeling. This measure of social perceptiveness has been psychometrically validated and, more recently, has been shown to predict one's ability to do social tasks such as work in a team (Baron-Cohen et al., 2001; Weidmann and Deming, 2020). These gender differences persist across a number of countries and across the age profile (Greenberg et al., 2023).

The remaining panels in Figure 2 show that women often excel in social skills when defined more broadly. Figure 2b reports gender differences in the number of friends one has during adolescence. This is based on data from the AddHealth Longitudinal Study, which asks students to list their top ten friends. Following Lleras-Muney et al. (2020), I report the "in-degree," or the number of other peers who list the subject as their friend. Figure 2c depicts averages in self-reported extraversion and sociability from the NLSY79 and NLSY97. In contrast to the RMET, such self-reported measures may suffer from biases in one's self-perception (Joseph and Newman, 2010). Moreover, both of these measures may partially conflate one's *preferences* for socializing rather than one's skill. And, even if they do capture one's *skills*, they may reflect types of social skills other than social perceptiveness (e.g., popularity, friendliness). Nevertheless, in most cases, women tend to exhibit significantly higher levels of social skills.[14]

[14]The NLSY79 measure exhibits the smallest female advantage in social skills, while the NLSY97 mea-

Throughout this paper, I will take these gender differences as given. Average differences across men and women in personality or other traits may arise for a variety of reasons – "nature" (e.g. natural selection pressures), "nurture" (e.g. gender socialization), or some combination of both these factors (Muthukrishna et al., 2021). Although an interesting area of research, it is beyond the scope of this paper to identify why these patterns emerge.

3.2 Flexibility & social task intensity are inversely correlated across occupations

Figure 3a reveals that occupations structured to be less flexible tend to demand more social perceptiveness. This fact, together with Section 3.1, provides *prima facie* evidence that the average woman faces a trade-off between her comparative advantage in the labor market and her desire for workplace flexibility.

This relationship is driven primarily by occupations in business, health, and legal sectors, which are relatively more intensive in social perceptiveness and also less flexible. In contrast, science, technology, retail, and administrative work tends to populate the middle of the distribution. Transportation and craft occupations tend to score low on both measures. As an additional way of visualizing which occupations are driving this relationship, Figure D.1 limits the sample to the largest occupations and labels a handful of illustrative occupations in each quadrant. Few occupations populate the northwest or southeast quadrants. For example, in the northwest quadrant, child care workers and nursing aides are structured to be more flexible and demand a relatively high amount of social perceptiveness. Yet, an occupation at a similar level of social perceptiveness (e.g. managerial roles) is more likely to be inflexible, leaving women with only a handful of jobs that both match their social skills and preferences for flexibility.

As presented in Table C.1, this correlation is robust to limiting the sample to high-income occupations (as in Figure 1),[15] weighting by only college-educated workers, dropping individual sub-components of the flexibility index, and using the alternative measures of flexibility discussed in Section 2.1.1.

To build intuition for why this correlation might emerge, Figure D.2 helps unpack which specific dimensions of flexibility may be lacking in social occupations. Work-

sure exhibits the largest. Perhaps consistent with this variation, research on the "Big Five" personality traits tends to find that gender differences in extraversion vary depend on which aspect of sociability is being captured. For example, men tend to be more assertive, yet women tend to report higher levels of warmth and gregariousness (e.g. Hyde (2014); Costa Jr et al. (2001); Weisberg et al. (2011)).

[15]Note that Figure 1 replicates Goldin (2014) and therefore focuses on college-educated individuals in occupations within the top quartile of average male earnings. Because I am interested in explaining occupational sorting patterns, not just the within-occupation gender wage gap, my analysis here and throughout the paper focuses on the full set of occupations and does not exclude noncollege individuals. As I discuss in Section 5.2, the main results are robust to limiting the sample to college educated individuals in high-income occupations.

ers in occupations that are high in social perceptiveness also tend to work long hours, work at the same time of day (possibly signifying a need for coordination), and report being contacted outside of work hours at a higher rate. These patterns are consistent with the notion that social tasks may involve more "relationship specific" capital, which naturally lends itself to demands on workers' time. That is, occupations that require interactions with people are more likely to demand "people-skills," such as social perceptiveness. And, as a result, such occupations will involve working with clients or teammates who may have unpredictable worker-specific demands or require certain workers to be present at particular times. To the degree that these people-oriented tasks are complemented by *relationship-specific* knowledge or capital that is not easily codified, workers will not be able to find someone to cover for them in their absence. For example, after many meetings with a particularly picky client, the creative director at an advertising agency has built up specific, tacit knowledge about what may satisfy the client's demands. This knowledge could inform how an ad is pitched in ways that would be difficult to communicate to others in the ad agency. And, even if it could be communicated, clients may have built up trust specifically with the director. Or, similarly, a team manager may need to coordinate a complex set of tasks across subordinates. Managing these moving pieces efficiently may require re-optimizing task allocation on-the-fly (and with speed) as workers run into unexpected roadblocks. Firms employing workers in such people-intensive roles may, in turn, reward those who make themselves more available. This could manifest in a variety of ways, such as end-of-year bonuses or promotion to a higher paying position.

It is worth noting that although this discussion suggests that Figure 3a reflects a *causal* relationship, this need not be the case in order to generate a trade-off. All that is required is jobs that happen to be less flexible also reward social skills. This bundling – regardless of the underlying reason – results in a trade-off. A related point is that this relationship is not necessarily *tautological*. In fact, empirically, there must be enough variation in order to estimate how sorting and wage gaps vary independently with flexibility and social perceptiveness. The analysis later in this paper confirms this is the case. Conceptually, it is useful to think of flexibility as characterizing the *timing of certain tasks* ("when") while skill demands reflect the *content of these tasks* ("what"). The conceptual framework in the next section clarifies this point further.

In addition to demanding social skills, Figure 3b highlights that less flexible occupations are *generally* skill-intensive: analytical skill intensity is also inversely correlated with flexibility. That analytical skills are rewarded in less flexible jobs suggests that women's higher willingness to pay for flexibility may also negatively impact women who excel in these skills. But, it is less clear how to think about gender differences in analytical skills.[16] This paper therefore focuses on how flexibility constrains women's

[16]On the one hand, women attend college at higher rates than men (Nietzel, 2024), and men are likely

ability to capitalize on their comparative advantage in social tasks, but this does not preclude the possibility that particular subgroups of women (e.g. highly analytically and socially skilled) face the *largest* trade-off.[17]

4 Conceptual Framework

Workers choose an occupation that maximizes their utility. To do so, they first determine how they would optimally allocate their labor supply across two production time periods within each occupation ("labor supply sub-problem"). Workers trade-off between their desire to work in particular time periods and the fact that output from one period may not be perfectly substitutable with output produced in another period. Having solved this sub-problem for each occupation, workers then choose the occupation that yields the highest utility, taking this labor supply choice as given ("occupational choice problem").[18] For this decision, workers trade-off between occupation-level attributes: their skills in an occupation and the substitutability of output across time periods. Below I describe each of these maximization problems, then I present comparative statics.

4.1 Labor supply sub-problem

A worker's utility is equivalent to their earnings. As such, they seek to maximize their total output Q_o, which is a constant elasticity of substitution (CES) function of their labor supplied across two time periods:

$$Q_{og} = \theta_{og}^{\mu} \left(0.5 t_1^{\rho_o} + 0.5 t_2^{\rho_o}\right)^{\frac{1}{\rho_o}} \tag{1}$$

where g indexes the worker's gender and o indexes the occupation. t_1 and t_2 represent labor supplied in each different time periods, such as two different weeks or two different parts of the day. Each time period is weighted equally (by 0.5) in total output.

A key feature of the worker's production function is the elasticity of substitution, $\sigma_o \equiv \frac{1}{1-\rho_o}$, which indexes how easily tasks can be substituted across the two time periods. This substitutability determines the flexibility of an occupation. Units of labor supplied in each period are assumed to be imperfect substitutes, and therefore σ_o is

to have a comparative advantage in manual labor intensive tasks. On the other hand, men may exhibit higher variance in analytical skills and therefore stack up in both the low and high end of the distribution (Hyde, 2014).

[17] This would be consistent with the results in Juhn and Rubinstein (2022), who find that women cannot fully leverage their comparative advantage in abstract tasks (which demand both cognitive and social skills) due to fact that abstract task intensive jobs typically require long or unusual hours.

[18] This approach of nesting a labor supply sub-problem within an occupational choice problem is similar to the approach taken by Erosa et al. (2022).

always greater than 1.[19] An occupation with high σ_o is one in which workers have ample flexibility to supply their labor in one period or the other with little impact on output. Occupations with low σ_o are ones where labor supplied in one period cannot easily "make up for" lack of labor supply in another period.

Another key feature is worker's skills, θ_{og}, which vary by the worker's gender and the occupation. $\mu \in (0, 1)$ captures diminishing returns to these skills. θ_{og} does not impact the labor supply decision but plays a pivotal role in the occupational choice problem.

Within each occupation, workers choose t_1 and t_2 to maximize 1 subject to their time constraint:

$$T = \alpha_g t_1 + t_2 \tag{2}$$

where T represents a fixed unit of labor. $\alpha_g \geq 1$ represents the relative "price" (to the worker) of supplying labor in period 1. α_g therefore captures the propensity for workers to be relatively "selective" about when they work.

Solving this problem yields labor supply functions $t_1^* (\alpha_g, T, \sigma_o)$ and $t_2^* (\alpha_g, T, \sigma_o)$, which are described in Appendix B. The solutions are analogous to a standard constrained utility maximization problem with CES preferences.[20] As α_g increases, workers shift their labor supply from period 1 to period 2. At larger values of σ_o, production is more substitutable across time periods, and this shift has a smaller negative impact on output. In the extreme, as $\sigma_o \to \infty$, workers allocate all of their labor to the cheaper time period (period 2). For this reason, workers with higher α_g will find higher σ_o more valuable.

4.2 Occupational choice problem

Having determined worker's allocation of labor supply conditional on occupation, we can use t_1^* and t_2^* to determine which occupation yields the highest utility. Substituting t_1^* and t_2^* into the production function (Equation 1), we get worker's indirect utility in a given occupation:

$$\overline{Q}_{og} \equiv Q_{og}^* (\theta_{og}, \sigma_o, \alpha_g, T) = \theta_{og}^{\mu} T 0.5^{\frac{\sigma_o}{\sigma_o - 1}} \left[\alpha_g^{1-\sigma_o} + 1\right]^{\frac{1}{\sigma_o - 1}} \tag{3}$$

Workers then choose a bundle of θ_{og} and σ_o (i.e., an occupation) to maximize 3 subject to the feasible set of (θ_{og}, σ_o) pairs (i.e. possible occupations). This constraint deter-

[19] That is, $\sigma_o \in (1, \infty)$. As $\sigma_o \to \infty$, t_1 and t_2 become perfect substitutes; as $\sigma_o \to 1$, production approaches the Cobb-Douglas case.

[20] This type of problem generally involves maximizing utility $U = \left(bx^{\frac{\sigma-1}{\sigma}} + (1-b)y^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$ subject to a budget constraint. My set up is the same, except I have set $b = 0.5$. This ensures that only the time constraint leads workers to favor working in one period or the other, which in turn simplifies the comparative statics.

mines how these two attributes tend to be bundled across occupations and, therefore, whether (and how much) workers are trading off between returns to their skills and their desire for flexibility.

4.3 Specifying gender differences and social task intensity

Assume female workers are more selective about when they want to supply their labor: $\alpha_f > \alpha_m$. This could be driven by gender differences in preferences for work-life balance, societal expectations about childcare responsibilities, or some combination of these factors.

Second, let a worker's skills in an occupation be fully determined by their social skills ($s_g \geq 0$) and the social task intensity of the occupation ($s_o \geq 0$). Specifically, let $\theta_{og} = s_g s_o + \kappa_0$ where κ_0 is a positive constant which guarantees workers have non-zero productivity in all occupations. Moreover, motivated by the previous discussion in Section 3.1, assume that women are more socially skilled ($s_f > s_m$).

Third, assume that social tasks are often harder to substitute across time periods. Section 3.2 discussed why social tasks may be less flexible. For example, if a client has urgent demands, it may be costly for workers to push off the meeting to a later date. Therefore, let $s_o = -\kappa_1 \sigma_o$ where $\kappa_1 > 0$ captures the relationship between social intensity and flexibility across occupations. We can then write the constraint in the occupational choice problem as

$$\theta_{og} = -\kappa_1 s_g \sigma_o + \kappa_0 \tag{4}$$

κ_1 and s_g therefore jointly mediate how much σ_o a worker must give up to get more θ_{og}.

To build intuition, Figure 4 plots sets of loci of $\overline{Q}_{og}$ for a representative female and male worker. The female worker faces a higher α_g and thus her loci have a steeper slope: she more willingly gives up θ_{og} in favor of more σ_o. Yet, at the same time, occupations in which she excels – social intensive occupations – tend to have lower σ_o. Therefore her feasible set (represented by the solid line) has a steeper slope. Compared to her male counterpart, she must give up relatively more θ_{og} in order to achieve a similar level of σ_o.

4.4 Empirical predictions

Below I summarize four key empirical predictions of the model. Throughout, let women's relative earnings be denoted as $\overline{Q}_o = ln\left(\frac{\overline{Q}_{of}}{\overline{Q}_{om}}\right)$. See Appendix B.2 for all derivations.

1. **The positive impacts of flexibility and social task intensity on women's relative wages and employment will be concealed by the fact that these two job attributes are inversely correlated.**
   To see this, consider three facts. First, because $s_f > s_m$, women's earnings will be boosted more by an occupation's social intensity relative to men ($\frac{\partial \overline{Q}_o}{\partial s_o} > 0$). Consequently, women will be more attracted to social occupations, holding all else constant. Second, as discussed in the labor supply sub-problem, α_g has a negative impact on worker's utility in an occupation, but this negative impact is offset to the degree that time periods are more substitutable for one another (higher σ_o). Therefore, because $\alpha_f > \alpha_m$, women's earnings will be boosted more by an occupation's flexibility relative to men ($\frac{\partial \overline{Q}_o}{\partial \sigma_o} > 0$). Consequently, women will be more attracted to flexible occupations, holding all else constant. Third, in equilibrium, worker's demand for social intensity and flexibility must satisfy the constraint: $s_o^* = -\kappa_1 \sigma_o^*$. As a consequence, women's relative employment and earnings in highly social or highly flexible occupations is *lower* than it otherwise would be in the absence of the constraint. Empirically, this implies that the correlation between an occupation's social perceptiveness and women's relative earnings and employment will become *more positive* once holding the occupation's flexibility constant. And, conversely, the correlation between an occupation's flexibility and women's relative earnings and employment will become *more positive* once holding the occupation's social intensity constant.

2. **The correlation between women's social skills and their propensity to choose less flexible jobs (that reward these skills) will be weaker than men's.**
   In other words, for men and women at the same skill level, women will be less likely to select into less flexible (more social) jobs. To see this, consider a man and woman at the *same* level of social skill, $\tilde{s}$. For both workers, if $\tilde{s}$ were to increase, flexibility would become relatively "expensive" vis-à-vis θ_{og}. The optimal σ_o^* would decrease as a result. But, this decreasing relationship will be weaker for the female worker, who is more reluctant to give up σ_o and therefore exhibits a higher marginal rate of substitution ($\frac{MU^f_{\sigma_o}}{MU^f_{\theta_o}} > \frac{MU^m_{\sigma_o}}{MU^m_{\theta_o}}$ due to $\alpha_f > \alpha_m$). Therefore, $0 > \frac{\partial \sigma_o^{*,f}}{\partial s_g} > \frac{\partial \sigma_o^{*,m}}{\partial s_g}$.

3. **The positive impact of flexibility on women's relative employment and wages will be amplified among parents.**
   Assume that after having a child, women become even more selective about when they supply their labor (α_f increases).[21] Occupations with greater flexi-

[21] This assumption aligns with past literature which shows childbirth has no impact on men's labor market behavior (e.g., Kleven et al. (2019)). If anything, there may be a positive impact (e.g., Goldin et al. (2024)).

bility soften the negative impact of higher α_g on one's wages (that is, $\frac{\partial^2 \overline{Q}_{og}}{\partial \sigma_o \partial \alpha_g} > 0$). Therefore, because childbirth disproportionately shocks women's demand for flexibility, the positive impact of flexibility on women's relative employment and wages will be larger among parents. Testing this prediction helps establish that time demands of less flexible jobs (rather than other unobserved factors) are constraining women's success in these jobs.

4. **The marginal productivity returns to enhancing flexibility will increasing in social intensity of the occupation.**
   In the model, σ_o is a worker's choice variable and is exogenously bundled with social tasks s_o. Yet, as I will discuss in Section 8, case studies have documented cases where firms have chosen to enhance flexibility. The model suggests that the marginal return to adopting such practices or technologies will be highest in social occupations. Women's relative productivity is increasing in social tasks, thus implying that relative labor demand for women is increasing in s_o. As technologies become available to increase flexibility, firms in social occupations will therefore be the most eager to adopt higher σ_o, given that it will attract relatively more women.

I test Prediction 1 in the ACS data, where I observe gender gaps in employment and wages across occupations. I test Prediction 2 in the NLSY data, where I observe individual's skills and their occupation choices. I primarily test Prediction 3 by leveraging event studies around first births, but I additionally assess heterogeneity by parental status in the analyses for Prediction 1 and 2 as further evidence of the mechanism. I indirectly test Prediction 4 by examining whether an occupation's social intensity predicts its likelihood of adopting flexibility-enhancing technologies, using two proxies from the CPS and ATUS.

5 How the trade-off influences gender gaps in employment and earnings

In this section, I test whether the positive impact of social tasks and flexibility on women's relative employment and wages is muted by the inverse correlation of these attributes across jobs. This is prediction 1 of the model. I use the ACS 2015-2019 data to regress a group's employment or wages within occupations on occupation-level characteristics and a set of fixed effects:

$$y_{oegp} = \beta_0 + \beta_1 Inflexibility_o * Female_g + \beta_2 Social_o * Female_g + X'_{og}\vec{\alpha} + \gamma_{oep} + \gamma_{egp} + \varepsilon_{oegp} \tag{5}$$

where o indexes the occupation, e indexes education (college vs noncollege), p indexes parental status (any kids vs no kids), and g indexes gender (male vs female). The outcome, y_{oegp}, represents either a demographic group's propensity to sort into an occupation or their average wages within an occupation. I construct a group's propensity to sort as $ln\left(Emp_{oegp}/Emp_{egp}\right)$ where Emp represents the number of workers. For wages, $\overline{ln\left(wage\right)}_{oegp}$, I take the average of log earnings residualized with respect to demographic characteristics and hours worked.[22]

The main variables of interest are the interactions between a female dummy ($Female_g$) and occupation-level characteristics measured in the O*NET. β_1 and β_2 represent how sorting and wage gaps change with a standard deviation increase in the inflexibility index ($Inflexibility_o$) and social perceptiveness ($Social_o$), respectively.

I additionally control for analytical intensity and routine task intensity interacted with female (vector X_{og}). These form the main set of control tasks used in Deming (2017). The robustness section explores the sensitivity to the results when including additional task measures.

I include fixed effects for demographic group γ_{egp} to account for factors that affect groups' sorting or wages that are constant across occupation (e.g. some groups may systematically experience lower pay across all occupations). I also include fixed effects γ_{oep} for all combinations of occupation, education, and parental status. This ensures my estimates are "within" education and parental status. Note that these fixed effects absorb occupation fixed effects (γ_o) which account for factors affecting wages and sorting that are constant across group within occupation (e.g. some occupations are higher-paying in general). Standard errors are robust and clustered on occupation, and observations are weighted by the total employment in the o-e-g-p cell.

5.1 Results

The results are presented in Table 1. Column 1 of Panel A regresses a group's propensity to sort on $Inflexibility_o$. As in Figure 1, women's relative employment is increasing in $Inflexibility_o$, confirming that the same puzzling fact emerges in the full sample of occupations. Taken at face value, this would (incorrectly) suggest that women dislike flexibility or that flexibility negatively impacts women's employment

[22] Specifically, I regress log earnings on age dummies, race dummies, year dummies, education year dummies, log weeks worked, and log weekly hours worked. Regressions are weighted by the ACS provided sample weights. To form $\overline{ln\left(wage\right)}_{oegp}$, I take the average of the residuals within each occupation and demographic group.

in an occupation. Yet, as discussed in Section 2.1, less flexible jobs are intensive in social skills, in which women have a comparative advantage. Consistent with this, Column 2 then shows women's relative employment is increasing in social perceptiveness ($Social_o$). And, Column 3 shows that this completely explains women's attraction to less flexible jobs. Once controlling for social intensity, the coefficient on $Inflexibility_o$ is negative (but indistinguishable from zero), and the coefficient on $Social_o$ becomes more positive.

Yet, as Section 3.2 alluded to, less flexible occupations are not only comprised of social tasks; they also often involve analytical tasks (e.g. a financial consultant may perform both social and analytical tasks). To test whether other attributes contribute significant explanatory power, I repeat the results in Columns 4-6 while controlling for analytical and routine task intensity. The pattern remains, suggesting that although flexibility is correlated with analytical skills, this does not explain why women enter these jobs (or why their wages are lower in these jobs).[23]

Columns 1-3 of Panel B in Table 1 reports the analogous wage results. A similar pattern emerges: comparing Columns 1 and 2 to Column 3, women's relative wages in less flexible occupations is biased upwards by $Social_o$. And, conversely, women's returns to social tasks increases nearly threefold once controlling for the fact that these occupations tend to lack flexibility. Columns 4-6, which additionally control for analytical and routine task intensity, depict a similar pattern. Together with the sorting results, these patterns are consistent with the inverse correlation between flexibility and social tasks muting the positive impact of both of these features on women's relative employment and wages (Prediction 1).

Interestingly, unlike Figure 1, there is no detectable correlation between women's relative wages and $Inflexibility_o$ when I do not control for social perceptiveness (Column 1, Panel B). There are several possible explanations. First, it is worth noting that relative to sorting gaps, wage gaps will generally be a "noisier" way of identifying the occupations with the most barriers or constraints for a group. This is due to Roy selection or composition effects, similar to the ones highlighted by Hsieh et al. (2019), which will bias gaps towards zero in occupations where a disadvantaged group faces some sort of barrier.[24] This suggests it is worth interpreting the wage and sorting results in conjunction with one another. Second, wage distributions in noncollege occupations

[23] In fact, women's relative employment is *decreasing* in analytical skills. This is consistent with women disliking these tasks, having a comparative disadvantage in these tasks, or facing other sorts of barriers (e.g., masculine culture) in jobs demanding these tasks.

[24] For example, perhaps only very high ability women will find it worthwhile to be in the least flexible occupations, which would raise the average ability of women in these occupations relative to men. This would close the gender wage gap in these occupations to some degree. In contrast, the sorting gap would not exhibit these types of competing effects: in such occupations we would see a large sorting gap, even if the gender wage gap appeared small. Hsieh et al. (2019) find evidence of such composition effects.

may be relatively more compressed. Consistent with this, Table C.2 reports the results when limiting the sample to the college-educated, and the coefficients in Panel B are generally larger (and more statistically significant).[25]

Could these results be explained by women's "preferences" for interacting with people? An alternative explanation for the sorting results (Panel A) is that women are attracted to social tasks because they *prefer* working with people, not that they are necessarily good at it. Indeed, as discussed earlier in Section 3.1, women tend to exhibit a greater interest in "people" over "things." Especially insofar as this emerges in childhood, it is difficult to see how this interest would not (at least inadvertently) lead one to cultivate more social skills. Nevertheless, while I cannot rule out the preference channel, the wage results in Panel B help establish evidence in favor of the comparative advantage story. To the degree that women excel in social tasks, firms should be willing to pay them more than men in order to attract them to the firm (i.e., social tasks positively impact women's relative wages). On the other hand, if women prefer social tasks, firms will not need to pay them as much to do the same tasks as men (i.e., social tasks negatively impact women's relative wages). To the degree both channels exist, the positive correlation between $Social_o$ and women's relative wages are consistent with the comparative advantage channel being the dominant force.[26]

5.2 Robustness

Alternative measures of flexibility As discussed in Section 2.1.1, the index I construct follows Goldin (2014) and is based on features of the work environment which may lead firms to reward workers for working long, irregular, or particular hours. And, while alternative approaches have the upside of ensuring some real-world manifestation of flexibility, they have the downside of potentially conflating general equilibrium features of labor markets. I report the results when using these alternative approaches in Tables C.3 and C.4. The same patterns emerge and, in some ways, are stronger. The main difference is that the initial "puzzle" is no longer as apparent. That is, women are not necessarily attracted to less flexible jobs on net. This does not preclude the fact that, nevertheless, Prediction 1 is borne out in these results. That is, women are *more* deterred from (attracted to) less flexible (more social) jobs upon holding social tasks (flexibility) constant.

Alternative measure of social intensity Motivated by the discussion in Section 3.1, I focus on social perceptiveness in order to more precisely capture the skills in which

[25] Relatedly, it is possible that the "greedy work" story is *most* relevant to college educated workers in the high-income careers studied by Goldin (2014). I explore this possibility in the robustness checks.

[26] And, in line with this, Cortes et al. (2023) exploit within-occupation changes to social tasks to show occupations which grow in social intensity attracted more women and exhibited higher wages, the latter of which is more consistent with a comparative advantage channel.

women excel. This is one of four components employed by the social skill composite used in Deming (2017), which additionally includes coordination ("adjusting actions in relation to others' actions"), persuasion ("persuading others to approach things differently"), and negotiation ("bringing others together and trying to reconcile differences"). Social perceptiveness is highly correlated with the Deming (2017) composite measure ($\rho = 0.93$). And, the social composite is even more highly correlated with the inflexibility index (ρ of 0.79 vs 0.74).

Table C.5 reports the results when using this social composite. Consistent with the notion that this less precisely captures women's comparative advantage, this version of social tasks is generally less predictive of women's relative employment and wages. Interestingly, despite being positively correlated with both the inflexibility index and social perceptiveness, controlling for this measure of social skills no longer leads the coefficient on $Inflexibility_o$ to flip to a negative sign. One explanation for this is that there is not enough variation to separately identify these factors. To explore this, I test whether the pattern emerges when using the alternative flexibility measures, which are less correlated with social perceptiveness and therefore may allow me to estimate things more precisely (Table C.6 and C.7). Consistent with this explanation, we see the main pattern emerge, especially when using the index which does not rely on the O*NET at all.

Non-routine manual tasks and physical strength My controls for analytical skill demands and routine task intensity follow Deming (2017). Comparing this control set to the taxonomy of task intensities put forward in Acemoglu and Autor (2011), these controls may fail to capture *non-routine* manual tasks. Non-routine manual tasks require "situational adaptability, visual and language recognition, and in-person interactions" and characterize jobs like pilots, crane and tower operators, and mechanics. Especially insofar as these tasks require physical strength, these are likely to be jobs in which men have a comparative advantage.

An index capturing non-routine manual tasks[27] is inversely correlated with social perceptiveness ($\rho = 0.623$). Are women sorting into high-social jobs not because of comparative advantage in social tasks, but rather because of their *comparative disadvantage in non-routine manual tasks*? If true, this would be a slightly altered narrative, although it would still be one about women's relative comparative advantage. To explore this, Table C.8 reports the results when controlling for non-routine manual tasks. The wage patterns are strikingly similar. The sorting results are qualitatively similar but weaker – women's relative propensity to sort into social occupation is positive but

[27] I follow Acemoglu and Autor (2011) and take the average of the following dimensions from the O*NET: the importance of operating vehicles, mechanized devices, or equipment; spending time using hands to handle, control or feel objects, tools or controls; manual dexterity; spatial orientation.

no longer statistically significant (Column 6, Panel A).

As an additional test, I also construct an index for "physical strength" skill demands to more closely parallel the analytical and social skill demand measures. My index is a standardized average of four components capturing the level of each of the following that the job requires: static strength, explosive strength, dynamic strength, and trunk strength. The results are presented in Table C.9. In this case, the sorting results are robust (Panel A), but the wage results exhibit some sensitivity (Panel B).

Taken together, these results suggests that while this explanation has some validity, it does not fully account for the narrative.

High-income, college-educated sample Goldin (2014) focuses primarily on college educated workers in high-income occupations, similar to the sample reported in Figure 1. In these occupations, there is are clear, well-known vignettes describing the propensity for firms to reward long or irregular hours (e.g. lawyers). Because I am focused on also explaining sorting patterns *across* occupations, it is more natural to expand my focus to the entire set of occupations and workers. Yet, as a proof-of-concept, I also report these results for the smaller sample.

Even in this smaller sample, there is evidence supporting Prediction 1. Table C.10 shows how the puzzle presented in Figure 1 is resolved when controlling for social perceptiveness. In comparison to the main results, the coefficients are generally larger but less precisely estimated. To provide continuity with Figure 1, I also present the results at the occupation level in Table C.11. Column 1 in Panels A and B replicates the correlations depicted in Figure 1. Consistent with the main results, Column 3 shows how controlling for social perceptiveness makes the coefficient on $Inflexibility_o$ more negative. For example, when controlling for social intensity, the correlation between $Inflexibility_o$ and the wage gap doubles. The pattern is unaffected when controlling for routine and analytical task intensity (Columns 4-6).

5.3 Heterogeneity by parental status

Before moving to the next section, I present evidence that women's time constraints, rather than other factors, limit their success in less flexible jobs. Section 7 offers a more comprehensive examination of this mechanism by using event studies around first-births. In this section, I provide corroborating descriptive evidence by exploring heterogeneity in the ACS results based on parental status. Specifically, I modify Equation 5 to incorporate additional interaction terms with a dummy variable indicating parental status ($Parent_p$):

$$
\begin{aligned}
y_{oegp} =& \beta_0 + \beta_1 Inflexibility_o * Female_g + \beta_2 Social_o * Female_g \\
&+ \beta_3 Inflexibility_o * Female_g * Parent_p + \beta_4 Social_o * Female_g * Parent_p \\
&+ X'_{og}\vec{\alpha} + X'_{ogp}\vec{\mu} + \gamma_{oep} + \gamma_{egp} + \varepsilon_{oegp}
\end{aligned} \tag{6}
$$

where estimates of the coefficients on the interaction terms, β_3 and β_4, reflect how gender gaps in sorting and wages change by parental status. For example, β_3 can be interpreted as representing the expected change in the gender gap for each standard deviation increase in the inflexibility index among parents relative to non-parents. Note the fixed effects γ_{oep} absorb the interaction between parental status and occupation characteristics. I also interact parental status with the control tasks (X_{ogp}).

Table 2 report the results. In line with the proposed mechanism, the relationship between the inflexibility index and women's relative employment is more negative among parents (Columns 1-2). A similar pattern emerges for wage gaps, yet this is only true when controlling for simultaneous interactions with other task intensities (Column 4). And, interestingly, the women's relative wages are more positively related to social tasks among parents. This could indicate that women's relative social skill advantage is larger among parents. This could happen if the women who remain in high-social occupations after becoming parents are positively selected. In Section 7, I address concerns like this by leveraging longitudinal data to retrieve *within-person* estimates of how gender gaps change around births.

6 Skill-based sorting gaps in inflexible occupations

The results in Section 5 suggest that women's returns to social tasks are lowered by the lack of flexibility bundled with these tasks (Prediction 1). In this section, I test the second implication of the model: women's social skills are less predictive of their entry into less flexible jobs. The model made the simplifying assumption that occupations are 1-dimensional in their skill demands. Of course, occupations may reward multiple skills; consistent with this, recall that less flexible jobs also tend to reward analytical skills (see Figure 3 and discussion in 3.2). To accommodate the multi-dimensional nature of skill demands, I test a slightly alternative implication: women's social *or* cognitive skills are less predictive of their entry into less flexible jobs. Regardless of whether the results are "driven by" social or cognitive skills, the point is to test whether given the same level of skill, women are less likely to sort into less flexible jobs. If so, this suggests women face obstacles which keeps them from fully realizing their skills in these jobs.

To do this, I use the NLSY79 to run the following regression:

$$y_{it} = \beta_0 + \beta_1 Female_i + \mu' \vec{Skill}_i + \alpha' \vec{Skill}_i * Female_i + \theta' \vec{X}_{it} + \delta' \vec{X}_{it} * Female_i + \varepsilon_{it} \tag{7}$$

where i indexes individuals and t indexes the year. $\vec{Skill}_i$ is a vector of individual i's social skills and cognitive skills as defined in Section 2. $\vec{X}_{it}$ is a vector of age fixed effects, year fixed effects, noncognitive skills, a dummy indicating the noncognitive skills were missing (and filled in with zeros), and a dummy for college completion.[28] I additionally interact these controls with a female dummy. The college and noncognitive skills controls are potentially other measures of ability through which it is possible a "sorting penalty" operates. Therefore, I present the results with and without the interaction of these measures with the female dummy. Standard errors are robust and clustered on individual.

Mirroring the analysis in Section 7.1, I construct two outcomes. The first is a dummy indicating employment in a low flexibility occupation (meaning, above median in $Inflexibility_o$). The second is a dummy indicating employment in a high flexibility occupation. Again, both dummies are equal to zero when the person is not employed; and, the latter acts as a sort of "placebo" test which helps me assess whether the patterns are unique to less flexible jobs.

Results are presented in Table 3. I begin by including social skills and cognitive skills separately (Columns 1-2) then simultaneously (Column 3). Focusing on Columns 1-3 of Panel A, we can see that men's social and cognitive skills both positively predict employment into less flexible occupations. Conversely, turning to Panel B, these skills are *negatively* correlated with employment in highly flexible jobs. This pattern is consistent with Figure 3, which used occupation-level information (from the O*NET) to show that analytical and social skills are more valuable in less flexible occupations.

Consistent with Prediction 2 of the model, this pattern is reversed to some degree among women. For example, Column 3 of Panel A in Table 3 suggests that a one standard deviation increase in social skills is associated with a 0.03 increase in probability that men are employed in a less flexible job. For cognitive skills, this value is 0.10. Yet, this is less true for women, particularly for social skills: the same increase in social skills is only associated with a 0.007 increase in the probability they are employed in less flexible jobs. At the bottom of Table 3 in Column 3, I report a test of joint significance on the female dummy interacted with social and cognitive skills ($\alpha_{social} + \alpha_{cognitive} = 0$). While the coefficient on cognitive skills and the female dummy is not statistically distinguishable from zero, social and cognitive skills are jointly less predictive of women's

[28] Three individuals have missing noncognitive skills. My measure of noncognitive skills follows Deming (2017) and represents the normalized average of the Rotter Locus of Control and the Rosenberg Self-Esteem Scale.

entry into less flexible jobs (p-value of 0.005).

Columns 4-6 present the results while additionally interacting the ability controls (college education and noncognitive skills) with the indicator for female. The overall pattern remains although it is generally weaker, indicating that some of the skill-based sorting penalty operates through these other measures of ability. This is primarily driven by the interaction between cognitive and female, which becomes smaller upon including these controls. The joint significance test yields a larger p-value (0.047).

Finally, I examine heterogeneity by parental status in Column 7. To do this, I include additional interactions with a dummy indicating the individual has had a child. Interestingly, there is no heterogeneity by social skills. But, comparing men and women at the same level of cognitive skills, upon becoming mothers, women are less likely to sort into less flexible occupations.[29]

In line with Prediction 2, these results suggest that women's skills are less predictive of their propensity to sort into less flexible jobs. And, consistent with the notion that women's time constraints are driving these results, there is some evidence that this is driven by mothers.

7 Mechanism: women's time constraints

In this section, I test more explicitly for the proposed mechanism (Prediction 3). Attributes of inflexible occupations that are unrelated to time-sensitive tasks, long hours, or other demands on workers' time may lead women to either avoid these occupations or earn less within them. For example, less flexible occupations may be more likely to have "masculine" firm culture, which could impede women's success in these occupations. In this section, I leverage the fact that the birth of a child will disproportionately shock women's time constraints. Through the lens of the model, this can be considered a positive shock to α_f. If gender gaps widen more in inflexible jobs after this event, then this helps establish that gender gaps in these jobs are driven by women's relative time constraints.

I perform this analysis in two ways. First, I exploit event-studies around first births in the National Longitudinal Survey of Youth 1979 (NLSY79) to estimate the so-called "child penalty" on women's employment and wages in less flexible jobs. Second, to complement these event-studies while leveraging more statistical power, I exploit a pseudo difference-in-difference strategy around births using synthetic cohorts in the Current Population Survey (Kleven, 2023).

[29] Juhn and Rubinstein (2022) perform a similar analysis, but their focus is abstract tasks and cognitive skills. They find that highly cognitively skilled women are under-represented in abstract task intensive jobs, and this is explained by women with children. My findings here could partially reflect the tension they explore.

Two comments are in order before proceeding. First, typical models of labor market decisions bias me "away" from observing heterogeneity by parental status. Insofar as individuals are forward-looking, we would expect women's expectations about their *future* demand for flexibility to impact their labor market decisions even before having children.[30] Therefore, the true treatment effect of children (and a shock to women's time constraints) may be underestimated by the results presented below. Second, it is possible that childbirth not only shocks women's relative time constraints, but in doing so, it also triggers more skill depreciation. If women supply less hours or experience nonemployment spells after childbirth, their skills (e.g. manual dexterity to perform a surgery) may deteriorate. This is a slightly separate mechanism that may partially be driving the results below, insofar as less flexible jobs also reward skills that are easily deteriorated. This would, however, still align with the idea that women's unique labor supply patterns associated with child-rearing lead them to face a trade-off.

7.1 Event-studies around first births

In this section, I leverage event studies around first births in the NLSY79. To build intuition for this test, consider the "ideal experiment" that would capture the causal effect of workplace flexibility on women's labor market outcomes. One approach would be to randomly increase flexibility in one occupation and not the other, then test whether gender gaps in employment and wages close more in the first than the second ("supply-side" experiment). A second approach would be to randomly shock some worker's demand for flexibility and compare how the workers who were in less vs more flexible jobs fare ("demand-side" experiment). Below, following the intuition of the second experiment, I leverage childbirth as a shock to women's relative demand for flexibility and ask how this impacts their success in less flexible jobs.

Empirical strategy I build on the growing "child penalty" literature, which exploits variation in the timing of first births to estimate the impact of children on labor market outcomes (e.g. Kleven et al. (2019, 2021); Kleven (2023)). To ensure the estimates compare treated individuals to untreated (rather than already-treated) individuals, I apply recently developed methods and employ a "stacked" event study design similar to those deployed by Cengiz et al. (2019) and Deshpande and Li (2019).[31] To do

[30] For example, a medical school graduate may choose a specialty with lower hours requirements during residency in anticipation of a desire to start a family; or, a young woman with a B.S. degree may choose nursing over becoming a medical doctor. In line with this, Wiswall and Zafar (2021) document that women's fertility expectations impact their choice of college major.

[31] See Dube et al. (2024) for an overview of the issues that arise in two-way fixed effects designs with staggered adoption. While event study specifications can accommodate dynamic treatment effects (e.g., the emergence of a larger treatment effect over time), the leads and lags may be contaminated by effects in other periods insofar as treatment effects are heterogeneous (e.g., larger treatment effects for women who give birth at an earlier vs later age).

this, for each "treatment cohort" (meaning, age at which the individual has their first child), I create a dataset that contains their observations and *all pre-birth observations* of a set of control individuals. Because I am interesting in outcomes that may take time to materialize, my main specification uses individuals who never have kids as a control group.[32] I then stack these individual "experiments" and estimate the event study within the stacked dataset by controlling for fixed effects which index these experiments. Specifically, I run the following regression separately for men ($g = m$) and women ($g = f$):

$$y_{itc}^{g} = \beta_0^g + \underbrace{\sum_{\forall\, k \neq -1} \beta_k^g \mathbf{1}\left\{t \in k\right\}}_{\text{event time dummies}} + \gamma_{ac}^g + \gamma_{rc}^g + \gamma_{ic}^g + \varepsilon_{itc}^g \tag{8}$$

where i indexes individual, t indexes event time (years since the individual's first-birth), and c indexes the observations associated with each treatment cohort. $\mathbf{1}\left\{t \in k\right\}$ are a dummy variables indicating event time.[33] $\mathbf{1}\left\{t \in k\right\} = 0\ \forall\, k$ for individuals in the control group, thus indicating they are "never treated" over the window of time relevant to the treatment cohort. The period just prior to first-birth, -1, is omitted from the regression. I additionally include age (a) fixed effects, year (r) fixed effects, and individual (i) fixed effects, each interacted with fixed effects for experiment c. These fixed effects are represented by γ_{ac}^g, γ_{rc}^g, and γ_{ic}^g, respectively. Standard errors are robust and clustered on individual.[34] To aid interpretation, estimates of β_k^g are converted to percentage changes by dividing the estimates (and the confidence interval) by the average of the outcome for the pre-birth observations of the treated g individuals for $k \in [-6, -1]$. The β_k^g's therefore represent the average within-person percentage change in y_{itc} in period $t = k$ relative to $t = -1$ net of the same change calculated for individuals in the control group who are the same age and observed in the same year.

I first examine how births impact women's relative propensity to sort into less flexible jobs. My main outcome of interest is a dummy indicating employment in a *low* flexibility (above-median in the inflexibility index) occupation. Estimates of β_k^g therefore reflect the impact of first births on the propensity to sort into less flexible jobs vis-à-vis either flexible jobs or nonemployment. To the degree the results are driven by women differentially sorting into nonemployment, I need to assess whether I see the same patterns for flexible jobs. Therefore, I compare these results to the impact on a dummy variable indicating employment in a *highly* flexibility occupation (those below

[32] Below I show that the results are robust to instead forming control groups of individuals who have children 3 or 5 years later (relative to the treatment cohort).

[33] Note that in practice k varies from -27 to 39 but my plots focus on $k \in \{-6, 9\}$ for visibility.

[34] Given my stacked event study approach, individuals may appear multiple times. Clustering standard errors on the individual helps account for this redundancy.

the median in the inflexibility index).[35]

I then examine whether births disproportionately lower women's wages in less flexible occupations. To do this, I run the same specification as Equation 8 and additionally include interaction terms with a variable equal to the average inflexibility index of the individuals' occupation prior to birth ($\overline{Inflexibility_i}$):[36]

$$y_{itc}^g = \beta_0^g + \sum_{\forall\, k \neq -1} \beta_k^g \mathbf{1}\left\{t \in k\right\} + \sum_{\forall\, k \neq -1} \alpha_k^g \mathbf{1}\left\{t \in k\right\} * \overline{Inflexibility_i} \tag{9}$$
$$+ \gamma_a^g * \overline{Inflexibility_i} + \gamma_r^g * \overline{Inflexibility_i} + \gamma_{ac}^g + \gamma_{rc}^g + \gamma_{ic}^g + \varepsilon_{itc}^g$$

where the coefficients of interest, α_k^g, represent how the impact on wages varies for a one standard deviation increase in the inflexibility of the individual's prebirth occupations (on average). Note that to avoid conditioning on an outcome, I let wages equal zero if the person is not employed.

Results Figure 5 reports the occupational sorting results. In Panel A, the outcome is a dummy variable indicating employment in a job with low flexibility. After childbirth, compared to their male counterparts, women become disproportionately less likely to be in jobs with low flexibility. In contrast, Panel B depicts a similarly sized (though noisier) treatment effect on women's relative propensity to sort into flexible jobs initially, but this begins to fade four years after the birth.[37] The fact that women become relatively less represented in occupations with a low flexibility after childbirth could reflect either that men *outpace* women in entering these jobs or that women *disproportionately drop out* of these careers. The former may be more relevant for jobs like CEOs or managers, in which workers typically began their career in a lower ranked position within the firm. The latter may be a more accurate description of the impacts of childbirth on occupations like lawyers or physicians where one's occupation is the same along the career path. Either way, this corroborates the heterogeneity results in the previous section and suggests that women's relative time constraints bind more in

[35] Adding these two outcomes together yields a dummy variable equal to 1 if the person is employed. A simpler approach would be to use a variable equal to the inflexibility index of the person's job in period t. However, it is not clear how to deal with periods of nonemployment (which is common among mothers), given that the index is defined only for occupations. I instead binarize the outcome in this way so that I can accommodate nonemployment spells. The analysis in the next section does not require me to binarize the flexibility index.

[36] For the never-treated individuals, this is equivalent to the average inflexibility of their occupations across all observations.

[37] As discussed above, these estimates are in terms of percentage change. Much of the initial large decrease observed in Panel B is due to this transformation (because women are less likely to be in these jobs relative occupations with low flexibility). In the absence of this transformation, the drop observed in Panel A is larger than that observed in Panel B in all years.

less flexible jobs.[38]

Figure 6 reports the wage results. Specifically, the figure plots the estimates of α_k^f and α_k^m, the coefficients on the interaction between the event dummies and the flexibility of one's occupation prior to birth. Starting in the second year after childbirth, women in less flexible jobs prior to childbirth experience a larger decrease in their wages. The opposite is true for men initially in less flexible jobs. After childbirth, men's wages are (if anything) larger in these occupations. This disproportionately *positive* impact is consistent with the notion of a "fatherhood premium," in which childbirth has a positive treatment effect on labor market outcomes of men.[39] Overall, these results support the idea that women's time constraints more negatively impact their returns to being in jobs with low flexibility.

Robustness: alternative control group As described above, I use the never-treated (childless) individuals as the control group for each treatment cohort. When choosing control individuals, there is a trade-off between assessing long-run effects and comparing groups of people who are likely to be similar along unobservables (due to their current or eventual parental status). If I were to limit the control group for each treatment cohort to individuals who eventually have children x years later, then I could only assess outcomes up to the point at which the control has children (i.e., x years after the treatment cohort gives birth), at which point they become treated. The advantage of this approach is that for small x's (e.g. 3 or 5), it is easier to argue that the treatment and control groups are similar in unobservables. However, treatment effects in my setting may take a while to materialize, which this approach would miss.[40] Reassuringly, the lack of pre-trends suggest my approach may be suitable. Nevertheless, Figure D.3-D.4 report the results when setting $x = 3$ and $x = 5$. Results are similar when using the longer time horizon ($x = 5$) but are less apparent when comparing to those who give birth 3 years later.

[38] Bang (2021) uses event studies around first births in the NLSY79 to study how occupational flexibility impacts couples' division of labor after childbirth. Bang (2021) finds that the child penalty is larger for women whose *husbands* are in less flexible occupations prior to the birth. Most closely related to my results, Bang (2021) finds that the impacts of births on women's overall employment does not vary by an occupation's pre-birth flexibility. In contrast, I focus on the impact of childbirth on employment in flexible jobs, rather than examining heterogeneity by pre-birth occupational flexibility (which may fail to capture differences in *entry into* less flexible jobs). Additionally, I use a stacked event study approach and employ individual fixed effects.

[39] This could be driven by couples specializing in market versus home production after having a child. After childbirth, men's wives may take on a larger share of home production, leading men to devote more time and effort to their careers. Additionally, childbirth may cause men to become more motivated to succeed in work in order to provide for the family. See Goldin et al. (2024) for a more thorough discussion. The fatherhood premium is generally not well-understood and is a place for future research.

[40] To assess outcomes in the long-run using the eventually-treated individuals, I could set x to a large value, such as 10. But, in this case the advantage of the approach is less clear. Women who give birth at age 25 versus 35 are not likely to be very similar to one another.

Robustness: spouse's income I interpret the wage results as providing evidence that women's demand for flexibility is more costly in occupations with a higher inflexibility index (which translates into a larger negative impact on their wages). This is, implicitly, a "supply side" interpretation in which women in my sample all experience the same shift in their demand for flexibility, but the "price" they pay (in terms of lower wages) is higher in these jobs. One concern is that these results are driven by heterogeneity in the women's *shift in demand* for flexibility. For example, women who are initially in less flexible occupations may have richer husbands, which could imply that childbirth increases their demands for flexibility more than other women. To explore this, I construct a proxy for spousal income prior to childbirth,[41] and I rerun the wage results while controlling for this interacted with the event dummies, age dummies, and year dummies. Results are in Figure D.5. Consistent with this demand-side interpretation having some merit, I find that this explains some, but not all, of the heterogeneity. Note that this interpretation cannot explain the sorting results, in which individual fixed effects hold constant pre-birth spousal income. Of course, this does not preclude the possibility that there is *heterogeneity* in the sorting results by spousal income.

7.2 Pseudo difference-in-difference around births

As an additional test, I rely on a pseudo difference-in-difference around births (similar to that employed by Kleven (2023)) to test whether the sorting gap widens after childbirth in less flexible jobs. Specifically, I use the CPS to estimate child penalties in *each* occupation, then test whether these estimates are related to the inflexibility index. This is a higher-powered version of the NLSY results evaluating how childbirth impacts women's relative propensity to sort into less flexible jobs. In contrast to the NLSY, I observe a sufficient number of people in the CPS to retrieve estimates of how childbirth impacts women's relative propensity to sort in each individual occupation.

Using the sample of synthetic individuals in the CPS between 1972-2019 (described in Section 2), I then run the following regression separately by gender g and occupation o:

$$y_{it}^{o,g} = \alpha^{o,g} + \beta^{o,g} Post_{it} + \gamma_{ard}^{o,g} + \varepsilon_{iard}^{o,g} \tag{10}$$

where i indexes synthetic individuals, t indexes event time, a indexes age, r indexes year, and d indexes the demographic group defined by all other matching characteristics. $Post_{it}$ is a dummy indicating the individual has given birth ($k \geq 0$). $\gamma_{ard}^{o,g}$ are fixed effects at the level of the cell defined by all matching characteristics.

[41] Spousal income is only observed in survey years, so this is missing for many observations. To circumvent this, I take the average conditional on being observed. For individuals who not married, spouse income is set to zero.

My outcome is a dummy variable indicating employment in occupation o. It equals zero when the person is not employed or employed in different occupation. To aid interpretation, I divide the estimates by the average in the outcome among the childless sample. $\hat{\beta}^{o,g}$ can thus be interpreted as the percentage change in the likelihood that g sort into occupation o after they experience their first birth. I then take these estimates and calculate $\hat{\beta}^{o,f} - \hat{\beta}^{o,m} = \hat{\beta}^{o}$ for each occupation.

Figure 7 correlates $\hat{\beta}^{o}$ (standardized for interpretability) with the inflexibility index. Consistent with the notion that less flexible jobs may be more difficult to balance with time constraints associated with child-rearing, the estimated "child penalty" is increasing in $Inflexibility_o$.

To build intuition and to check pre-trends, I present the event study version of these estimates (where event time lag and lead dummies are included) in Figure D.6 for a few illustrative occupations. To further help visualize the gradient of the impact, Figure D.7 plots the female-male difference in the coefficients. We can see that after childbirth, women experience a disproportionately larger negative impact on their propensity to sort into legal work or marketing management, but this is less true for both waiters and computer scientists. And, reassuringly, there are no pre-trends leading up to the event.

8 Adoption of flexibility-enhancing technology across occupations

In this section, I provide evidence supporting the idea that the productivity returns to adopting flexibility-enhancing technology – and thereby retaining female talent – are highest in social occupations (Prediction 4). Specifically, I examine whether social occupations are more likely to *grow* in flexibility, where I assume that this growth reflects firms' profitable decisions to adopt such technology.

A range of technologies, practices, and policies can reduce the wage penalty workers experience to customize their schedule, thereby enhancing flexibility. The most effective approach likely depends on the specific job and context. My model offers a unifying principle for identifying effective strategies: those that make time periods in a worker's production process more substitutable. For example, law firms may offer part-time contracts to associates who aspire to make partner but cannot meet billable hour requirements during their prime childbearing years. Or, when possible, firms can schedule essential team meetings for times that are not likely to conflict with family responsibilities (e.g. hold meetings prior to school dismissal at 3 p.m.). Case studies highlight additional solutions. For example, information technology can enable workers to efficiently store and share clients' information, which can then allow workers

to operate as a "team of substitutes" and reduce the need for any *particular* worker to be present at certain times. Evidence suggests that this technology has facilitated more predictable schedules for OB-GYNs and increased part-time opportunities for pharmaceutical workers (Goldin and Katz, 2016; Goldin, 2021). These impacts can be further amplified by the standardization of products or services, which simplifies the storing and sharing of client information (e.g., the corporatization of pharmacies). Work can also be reorganized to separate on-call tasks from predictable tasks, as seen in the case of veterinarians, who eventually created distinct roles for managing day-to-day care and emergency care (Goldin and Katz, 2011). Occupation-level associations may also consider imposing norms or rules, such as the cap on hours worked in medical residency (Wasserman, 2023).

The breadth of approaches to making work more flexible poses a significant measurement challenge. I develop two proxies for changes to an occupation's flexibility that attempt to capture this breadth. The first is a "data driven" measure of changes in flexibility within occupations over time based on child penalties estimated separately for each occupation. This measure has the upside of capturing potentially a wide-range of ways that firms may adjust, but in doing so, it may also capture other factors (e.g., shifting ability composition of women) while also lacking interpretability. Therefore, to complement this measure, I also study the growth in the prevalence of *specific* practice which applies fairly widely across occupations, work-from-home arrangements.

Before discussing the measures, it is worth noting that firms' decisions to enhance flexibility need not stem from *deliberate* efforts by firms to attract female talent. Technology shocks unrelated to such efforts may persist in occupations that benefit from female talent due to market-driven selective pressures rather than firms' *intentional* pursuit of strategies that retain women. In fact, this "spontaneous order" model may better explain the pharmacy case study in Goldin and Katz (2016), where seemingly unrelated macro trends contributed to the restructuring of pharmaceutical work.[42] However, as institutions become more aware of the unique challenges women face in the workplace, firms may increasingly pursue such restructuring efforts with the explicit goal of attracting female talent.

Measure 1: motherhood-friendliness, 1972-2019 The first measure relies on the pseudo DiD strategy presented in Section 7.2. Specifically, I re-run equation 10 for two separate periods, 1972-1995 and 1996-2019.[43] I then store estimates $\hat{\beta}^{j,o,g}$ for each period j, occupation o, and gender g. As in Section 7.2, I transform these estimates to percentage

[42] Goldin and Katz (2016) document how a confluence of three factors made the pharmaceutical industry more egalitarian: the expansion of pharmacy chains, the adoption of information technology for storing client data, and the standardization of pharmaceutical products.

[43] The regression equation is

$$y_{it}^{o,g} = \alpha^{o,g} + \beta^{o,g} Post_{it} + \gamma_{ard}^{o,g} + \varepsilon_{iard}^{o,g} \tag{11}$$

changes by dividing the estimates by the share of childless g individuals in that occupation in the relevant period. $\hat{\beta}^{j,o,g}$ therefore represents the percentage change in the propensity for g to sort into occupation o after childbirth in period j. I then calculate

$$\Delta Mother\hat{F}riendly_o = [\hat{\beta}_{o,f}^{\text{1996-2019}} - \hat{\beta}_{o,m}^{\text{1996-2019}}] - [\hat{\beta}_{o,f}^{\text{1972-1995}} - \hat{\beta}_{o,m}^{\text{1972-1995}}]$$

which represents the change over time (in percentage points) in women's relative tendency to sort into an occupation after the birth of their first child. A positive $\Delta Mother\hat{F}riendly_o$ indicates a reduction in the sorting penalty experienced by mothers in that occupation.

The main advantage of this measure is that it uses sorting behavior as an indicator of individuals' preferences to identify which occupations have become more "motherhood-friendly." This revealed preference approach circumvents the need to identify *specific* innovations or technologies, the relevance of which for an occupation's flexibility may vary by sector.

Measure 2: supply of work-from-home arrangements, 2003-2019 My second measure focuses on work-from-home (WFH) arrangements. Harrington and Kahn (2023) provide evidence that the rise in WFH technologies since the 2000's has reduced motherhood penalties by a substantial margin, especially in traditionally motherhood unfriendly careers such as marketing and finance. They argue that WFH not only enables locational flexibility, but, in doing so, also enables greater *temporal* flexibility. For example, WFH technologies may enable workers to break up their work day to accommodate childcare responsibilities. In the absence of WFH, workers may be required to stay late at the office to ensure a deadline is met. This may not be possible for those responsible for making dinner or reading the bedtime story. With the possibility of WFH, however, such workers have the option of making up the hours later in the evening after the kids are asleep. Given this, while my paper is focused on temporal flexibility, my second measure measures the prevalence of WFH arrangements across occupations using the ATUS.

I use ATUS's time diary information to calculate the share of work hours in which a worker reports doing their work from home ($ShareWFH_i$). In an attempt to isolate changes in the *supply* of WFH technologies (rather than shifts in *demand*, e.g. due to the influx of women into an occupation), I residualize $ShareWFH_i$ with demographic characteristics to form $Share\tilde{W}FH_i$.[44] I then take the average of $Share\tilde{W}FH_i$ in each

where i indexes synthetic individuals, t indexes event time, a indexes age, r indexes year, and d indexes the demographic group defined by all other matching characteristics. $Post_{it}$ is a dummy indicating the individual has given birth ($k \geq 0$). $\gamma_{ard}^{o,g}$ are fixed effects at the level of the cell defined by all matching characteristics.

[44] I residualize with respect to female, number of kids, female by number of kids, age (5 year bins) dummies, age dummies by female, married dummy, race dummies, education dummies (below high

occupation to get $Avg\tilde{ShareWFH}_o$. I repeat this process for two different periods, 2003-2010 and 2011-2019, then I take the difference:

$$\Delta Avg\tilde{ShareWFH}_o = Avg\tilde{ShareWFH}_o^{\text{2011-19}} - Avg\tilde{ShareWFH}_o^{\text{2003-10}}$$

The measure therefore proxies for the change in the prevalence of WFH technologies in occupation o. Due to the small sample size in the ATUS, only 313 occupations are represented.

8.1 Estimating equation

I now examine whether social occupations are the most likely to become more flexible according to these measures. The model assumes, for simplicity, that a job's social intensity is perfectly correlated with its inflexibility. In reality, this correlation is less than one. Therefore, I specifically test whether occupations that are *both* socially intensive *and* structured to be inflexible are the strongest predictors of adopting these technologies. Put differently, I assess whether, among inflexible occupations, the most social ones have the highest incentive to adopt such changes. Section 2.1 discussed the characteristics and work contexts that may make firms less flexible, which motivated the use of the O*NET-based inflexibility index, $Inflexibility_o$. I use this index again in this analysis. My empirical specification is:

$$\Delta y_o = \beta_0 + \beta_1 Inflexibility_o + \beta_2 Social_o + \beta_3 Inflexibility_o * Social_o + X_o'\alpha + \varepsilon_o \quad (12)$$

where y_o is either $Mother\hat{F}riendly_o$ or $\widetilde{ShareWFH}_o$. $Inflexibility_o$ and $Social_o$ are the O*NET characteristics used in previous sections of the paper. X_o is a vector of base period controls and their interactions with $Social_o$ and $Inflexibility_o$. This includes the share of those employed in the occupation who are college educated and the share of those employed who are female, both calculated in the CPS for the relevant set of base years. Because I am interested in the coefficient on the interaction term (β_3), I interact these base period controls with $Inflexibility_o$ and $Social_o$. These base period controls are based on the CPS data and calculated in the period one years relevant to the outcome in question (e.g. either 1972-1995 or 2003-2010). Standard errors are robust and the regressions are weighted by total employment in the base period.

8.2 Results

Table 4 reports the results with the motherhood friendliness measure in Panel A and the WFH measure in Panel B. Both outcomes yield remarkably similar patterns. On its

school, high school, some college, college, advanced degree), foreign born dummy, year dummies, and class of worker dummies. Regressions are weighted with ATUS sampling weights.

own, social-intensive occupations do not predict more growth in flexibility (Column 1). On the other hand, an occupation's predisposition to lower flexibility is positively correlated with the propensity for that occupation to become *more* flexible (Column 2). Yet, importantly, Column 3 reveals that this is driven by occupations which are *both* historically less flexible *and* social intensive. This is consistent with the idea that these firms have the most to gain from enhancing flexibility and therefore invest in flexibility (e.g. work-from-home technologies) in the long-run.

To test whether this pattern is unique to social intensive occupations, I investigate whether we see the same pattern for *analytical* intensity in Column 4.[45] In both cases, the interaction term between $Inflexibility_o$ and $Analytical_o$ is close to zero and statistically insignificant. This placebo test provides additional evidence in favor of the notion that women have a comparative advantage in *social* tasks which leads firms in those occupations to adopt more flexible work structures.

As alluded to above, neither measure perfectly captures firms' investments in flexibility. How does this influence how we interpret these results? For example, $\Delta \hat{MotherFriendly}_o$ may partially reflect the propensity for some occupations to be increasingly comprised of high ability women whose reservation wages are higher. For these composition shifts to explain these patterns, these shifts would need to be systematically more true in occupations high in both $Inflexibility_o$ and $Social_o$. Given the plausibility of such labor supply-side trends, I leave it to future research to explore this question in a more targeted manner. The suggestive evidence I present here indicates that this could be a promising area for future research.

9 Conclusion

In this paper, I have presented evidence that the lack of workplace flexibility constrains women's ability to leverage their comparative advantage in the labor market. While women tend to excel in social skills, jobs involving people-oriented tasks also tend to require long, irregular, or specific hours. This may arise, for example, because clients, team members, or supervisees have time-sensitive needs that can only be met by a particular worker with whom they have built relationship-specific capital. As a result, in order to be in careers which reward their skills, women often face a trade-off between their desire for flexible hours and earning returns to their unique set of skills.

To understand what empirical patterns would be consistent with this trade-off, I presented a simple conceptual framework. Jobs vary in whether they reward social skills and whether the tasks involved are time-sensitive. The latter is represented by an occupation's elasticity of intertemporal substitution for production in two separate periods. Women, who are more "selective" over the time periods in which they supply

[45] In this column, I additionally interact the controls with $Analytical_o$.

their labor, must trade-off between flexibility and being in a job which rewards their social skills.

I show evidence consistent with the model's four predictions. First, in line with the first prediction of the model, I used the ACS to show that gender gaps in employment and wages are understated in less flexible jobs due to the high social intensity inherent to these occupations. In other words, women are drawn into less flexible jobs to the degree these jobs also reward social skills, which biases their relative wages upwards. Second, I showed that while one's social (and cognitive) skills positively predict entry into less flexible jobs, this is less true for women. Third, I used childbirth as a shock to women's relative demand for flexibility to confirm that time constraints, rather than other unobserved job characteristics, limit women's success in less flexible occupations. Fourth, among traditionally inflexible occupations, the most social occupations are the most likely to adopt flexibility-enhancing technologies, in line with the notion that doing so facilitates greater access to talented women.

This trade-off amplifies the potential benefits of a range of policies and interventions. On the labor supply side, public policies such as funding for childcare or pre-kindergarten could enable women to excel in time-intensive careers (e.g. Humphries et al. (2024)). Childhood interventions could promote analytical skills among girls or social skills among boys, thus enabling women and men's skills to better match with jobs that match their desired flexibility.[46] On the labor demand side, professional associations may consider regulating work hours, similar to the cap on hours worked among medical residents implemented by the American Medical Association (Wasserman, 2023). Or, more broadly, institutions and firms may consider investing in technologies that reduce the cost of flexibility, such as information technologies that enable workers to seamlessly hand-off clients (Goldin and Katz, 2016).

The fact that there are potential *talent losses* stemming from the lack of workplace flexibility implies that *firms* will reap some benefit from such policies or interventions. Consistent with this, my final analysis – my test of the fourth prediction – suggests that firms, driven by the need to attract talented female workers, have begun introducing more flexible work structures in historically less flexible, highly social-intensive occupations. A promising area for future research is to explore whether this process is slowed by frictions (e.g., adjustment costs or organizational inertia) and whether policies can be developed to address these barriers.

[46]Making teachers aware of their biases can reduce the extent to which their stereotyping negatively impact students' performance (Alesina et al., 2024), which evidence suggests partially drives the gender gap in math performance in school (Carlana, 2019). Social interactions during elementary school can impact childrens' acquisition of cognitive empathy (Alan and Turkum, 2024).

References

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier.

Alan, S. and B. Turkum (2024). Abstract reasoning, theory of mind and character development in the school. *Journal of Economic Behavior & Organization 221*, 307–326.

Alesina, A., M. Carlana, E. La Ferrara, and P. Pinotti (2024). Revealing stereotypes: Evidence from immigrants in schools. *American Economic Review 114*(7), 1916–1948.

Atalay, E., P. Phongthiengtham, S. Sotelo, and D. Tannenbaum (2020, April). The evolution of work in the united states. *American Economic Journal: Applied Economics 12*(2), 1–34.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review 103*(5), 1553–1597.

Bang, M. (2021). Job flexibility and household labor supply: Understanding gender gaps and the child wage penalty. Technical report, Working Paper.

Baron-Cohen, S. (2002). The extreme male brain theory of autism. *Trends in cognitive sciences 6*(6), 248–254.

Baron-Cohen, S., S. Wheelwright, J. Hill, Y. Raste, and I. Plumb (2001). The "reading the mind in the eyes" test revised version: a study with normal adults, and adults with asperger syndrome or high-functioning autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines 42*(2), 241–251.

Barrero, J. M., N. Bloom, and S. J. Davis (2023, November). The evolution of work from home. *Journal of Economic Perspectives 37*(4), 23–50.

Bertrand, M. (2020). Gender in the twenty-first century. In *AEA Papers and proceedings*, Volume 110, pp. 1–24. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Bolotnyy, V. and N. Emanuel (2022). Why do women earn less than men? evidence from bus and train operators. *Journal of Labor Economics 40*(2), 283–323.

Bütikofer, A., S. Jensen, and K. G. Salvanes (2018). The role of parenthood on the gender gap among top earners. *European Economic Review 109*, 103–123.

Carlana, M. (2019). Implicit stereotypes: Evidence from teachers' gender bias. *The Quarterly Journal of Economics 134*(3), 1163–1224.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Ciasullo, L. and M. Uccioli (2024). What works for working couples? work arrangements, maternal labor supply, and the division of home production. *Working Paper* , .

Cortes, G. M., N. Jaimovich, and H. E. Siu (2023). The growing importance of social tasks in high-paying occupations: Implications for sorting. *Journal of Human Resources 58*(5), 1429–1451.

Cortés, P., Y. Feng, N. Guida-Johnson, and J. Pan (2024). Automation and gender: Implications for occupational segregation and the gender skill gap. Technical report, National Bureau of Economic Research.

Cortés, P. and J. Pan (2016). Prevalence of long hours and skilled women's occupational choices. Technical report, IZA Discussion Paper.

Cortés, P. and J. Pan (2019). When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics 37*(2), 351–398.

Cortés, P. and J. Pan (2023). Children and the remaining gender gaps in the labor market. *Journal of Economic Literature 61*(4), 1359–1409.

Costa Jr, P. T., A. Terracciano, and R. R. McCrae (2001). Gender differences in personality traits across cultures: robust and surprising findings. *Journal of personality and social psychology 81*(2), 322.

Cubas, G., C. Juhn, and P. Silos (2023). Coordinated work schedules and the gender wage gap. *The Economic Journal 133*(651), 1036–1066.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics 132*(4), 1593–1640.

Denning, J. T., B. A. Jacob, L. J. Lefgren, and C. Vom Lehn (2022). The return to hours worked within and across occupations: implications for the gender wage gap. *ILR Review 75*(5), 1321–1347.

Deshpande, M. and Y. Li (2019). Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy 11*(4), 213–248.

Dube, A., D. Girardi, O. Jorda, and A. M. Taylor (2024). A local projections approach to difference-in-differencess. Technical report, National Bureau of Economic Research.

Edin, P.-A., P. Fredriksson, M. Nybom, and B. Öckert (2022). The rising return to noncognitive skill. *American Economic Journal: Applied Economics 14*(2), 78–100.

Erosa, A., L. Fuster, G. Kambourov, and R. Rogerson (2022). Hours, occupations, and gender differences in labor market outcomes. *American Economic Journal: Macroeconomics 14*(3), 543–590.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review 104*(4), 1091–1119.

Goldin, C. (2021). *Career and Family: Women's Century-Long Journey toward Equity*. Princeton University Press.

Goldin, C. and L. F. Katz (2011). The cost of workplace flexibility for high-powered professionals. *The Annals of the American Academy of Political and Social Science 638*(1), 45–67.

Goldin, C. and L. F. Katz (2016). A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics 34*(3), 705–746.

Goldin, C., S. P. Kerr, and C. Olivetti (2024). The parental pay gap over the life cycle: Children, jobs, and labor supply. *Journal of Economic Dynamics and Control* , .

Greenberg, D. M., V. Warrier, A. Abu-Akel, C. Allison, K. Z. Gajos, K. Reinecke, P. J. Rentfrow, M. A. Radecki, and S. Baron-Cohen (2023). Sex and age differences in "theory of mind" across 57 countries using the english version of the "reading the mind in the eyes" test. *Proceedings of the National Academy of Sciences 120*(1), e2022385119.

Greenberg, D. M., V. Warrier, C. Allison, and S. Baron-Cohen (2018). Testing the empathizing–systemizing theory of sex differences and the extreme male brain theory of autism in half a million people. *Proceedings of the National Academy of Sciences 115*(48), 12152–12157.

Guvenen, F., B. Kuruscu, S. Tanaka, and D. Wiczer (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics 12*(1), 210–244.

Hall, J. A. (1978). Gender effects in decoding nonverbal cues. *Psychological bulletin 85*(4), 845.

Hall, J. A. (1990). *Nonverbal sex differences: Accuracy of communication and expressive style.* Johns Hopkins University Press.

Hanson, G. (2021). Immigration and regional specialization in ai. In *Robots and AI*, pp. 180–231. Routledge.

Harrington, E. and M. E. Kahn (2023). Has the rise of work-from-home reduced the motherhood penalty in the labor market. *Working paper* , .

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and us economic growth. *Econometrica 87*(5), 1439–1474.

Humphries, J. E., C. Neilson, X. Ye, and S. D. Zimmerman (2024). Parents' earnings and the returns to universal pre-kindergarten. Technical report, National Bureau of Economic Research.

Hyde, J. S. (2005). The gender similarities hypothesis. *American psychologist 60*(6), 581.

Hyde, J. S. (2014). Gender similarities and differences. *Annual review of psychology 65*, 373–398.

Joseph, D. L. and D. A. Newman (2010). Emotional intelligence: an integrative meta-analysis and cascading model. *Journal of applied psychology 95*(1), 54.

Juhn, C. and Y. Rubinstein (2022). Task requirements, time requirements, and the gender gap in jobs and pay. *Preliminary working paper* , .

Kleven, H. (2023). The geography of child penalties and gender norms: A pseudo-event study approach. Technical report, National Bureau of Economic Research.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics 11*(4), 181–209.

Kleven, H., C. Landais, and J. E. Søgaard (2021). Does biology drive child penalties? evidence from biological and adoptive families. *American Economic Review: Insights 3*(2), 183–198.

Lleras-Muney, A., M. Miller, S. Sheng, and V. T. Sovero (2020). Party on: The labor market returns to social networks and socializing. Technical report, National Bureau of Economic Research.

Maestas, N., K. J. Mullen, D. Powell, T. Von Wachter, and J. B. Wenger (2023). The value of working conditions in the united states and implications for the structure of wages. *American Economic Review 113*(7), 2007–2047.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–59.

Mas, A. and A. Pallais (2020). Alternative work arrangements. *Annual Review of Economics 12*, 631–658.

Muthukrishna, M., J. Henrich, and E. Slingerland (2021). Psychology as a historical science. *Annual Review of Psychology 72*(1), 717–749.

Ngai, L. R. and B. Petrongolo (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics 9*(4), 1–44.

Nietzel, M. T. (2024). Women continue to outpace men in college enrollment and graduation. *Forbes* , . Accessed: 2024-10-16.

Pertold-Gebicka, B., F. Pertold, and N. D. Gupta (2016). Employment adjustments around childbirth. *Working paper IZA DP No. 9685*, .

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations research 26*(2), 305–321.

Wasserman, M. (2023). Hours constraints, occupational choice, and gender: Evidence from medical residents. *The Review of Economic Studies 90*(3), 1535–1568.

Weidmann, B. and D. J. Deming (2020). Team players: How social skills improve group performance. Technical report, National Bureau of Economic Research.

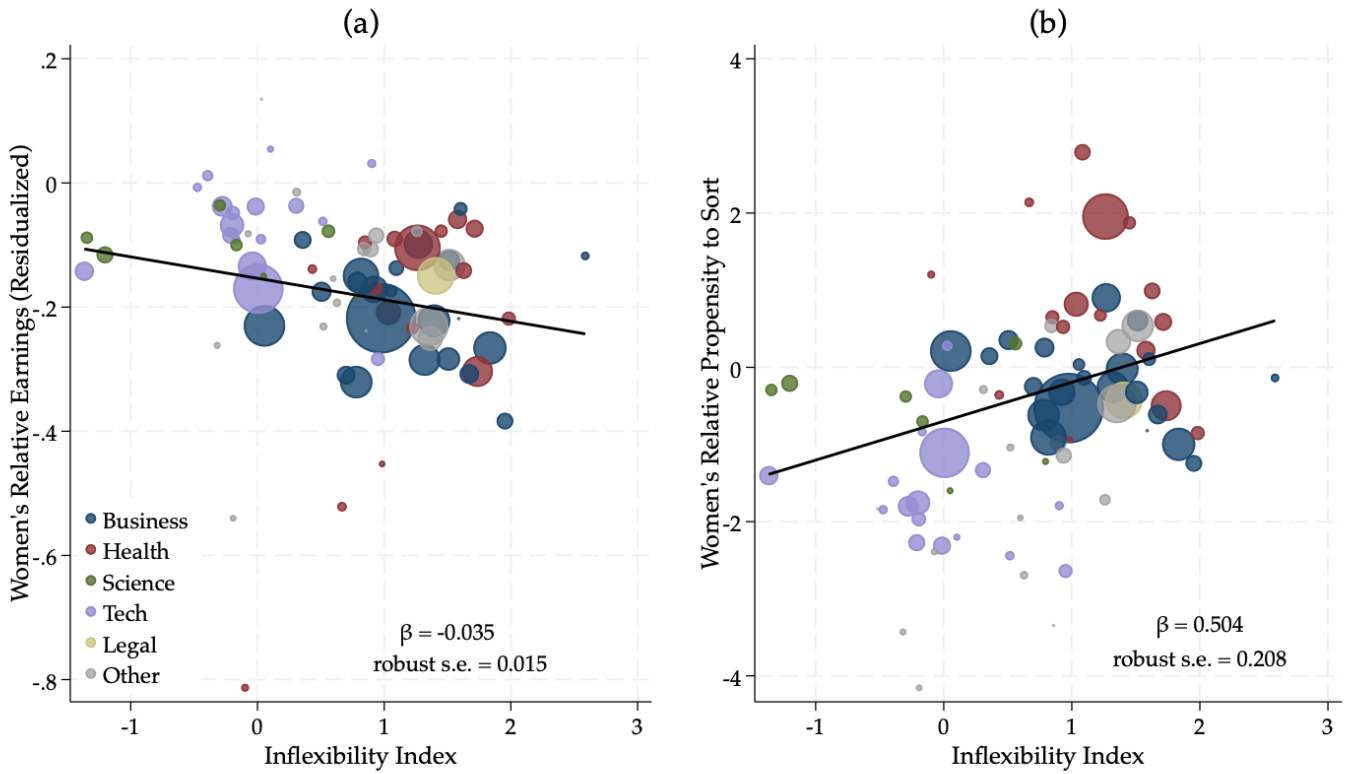
Weisberg, Y. J., C. G. DeYoung, and J. B. Hirsh (2011). Gender differences in personality across the ten aspects of the big five. *Frontiers in psychology 2*, 11757.

Wiswall, M. and B. Zafar (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics 133*(1), 457–507.

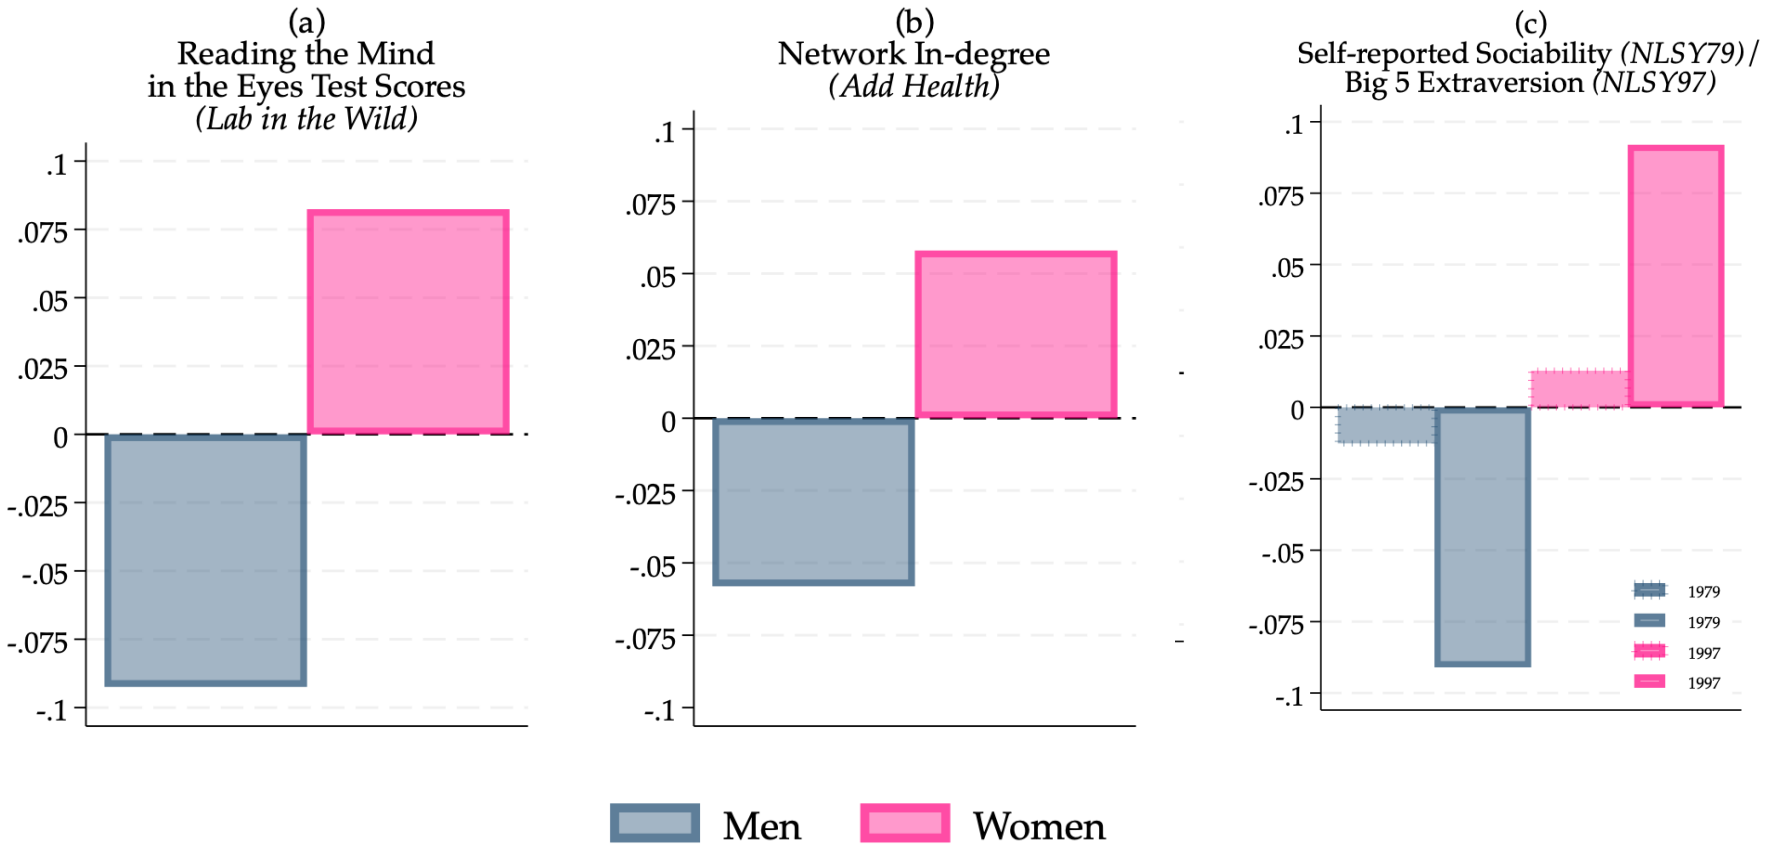
Wiswall, M. and B. Zafar (2021). Human capital investments and expectations about career and family. *Journal of Political Economy 129*(5), 1361–1424.

10 Main Figures and Tables

FIGURE 1: RELATIONSHIP BETWEEN AN OCCUPATION'S FLEXIBILITY AND GENDER GAPS IN WAGES AND SORTING

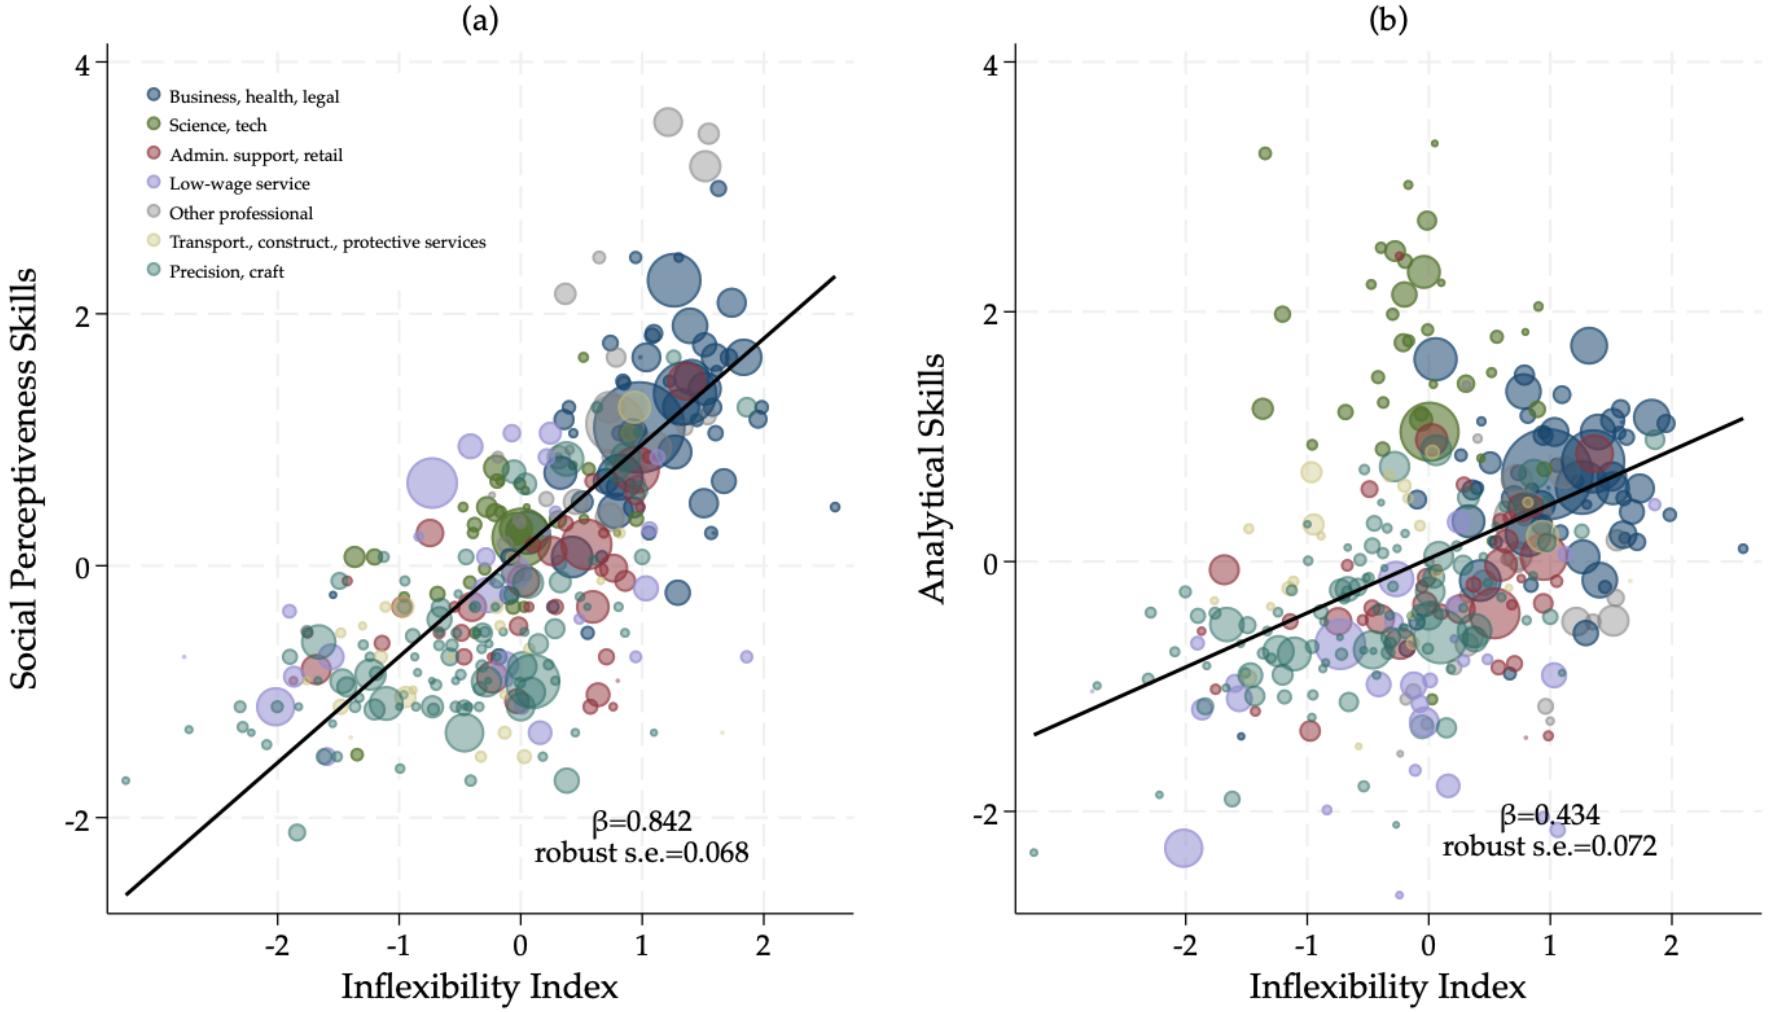


Notes: data are from ACS 2015-2019 and O*NET 2013. Each observation is a 3-digit occupation. The sample includes only occupations where the average male income is in the top quartile. The sample of workers includes native-born, college-educated, full-time-full-year workers aged 25-64. The left panel replicates Figure 5 from Goldin (2014). It depicts the relationship between the log gender wage gap within occupations and an index which proxies for the price of flexibility in that occupation (i.e., the amount of wages a worker gives up in order to customize their schedule). As discussed in Section 2.1.1, the index is based on data from the O*NET and captures how substitutable workers are for one another. The right panel plots the log of the relative propensity for women to sort into an occupation (the fraction of women sorting into an occupation divided by the fraction of men sorting into the occupation) against the same index. Observations are weighted by number employed within the sample. The unweighted relationship is very similar.



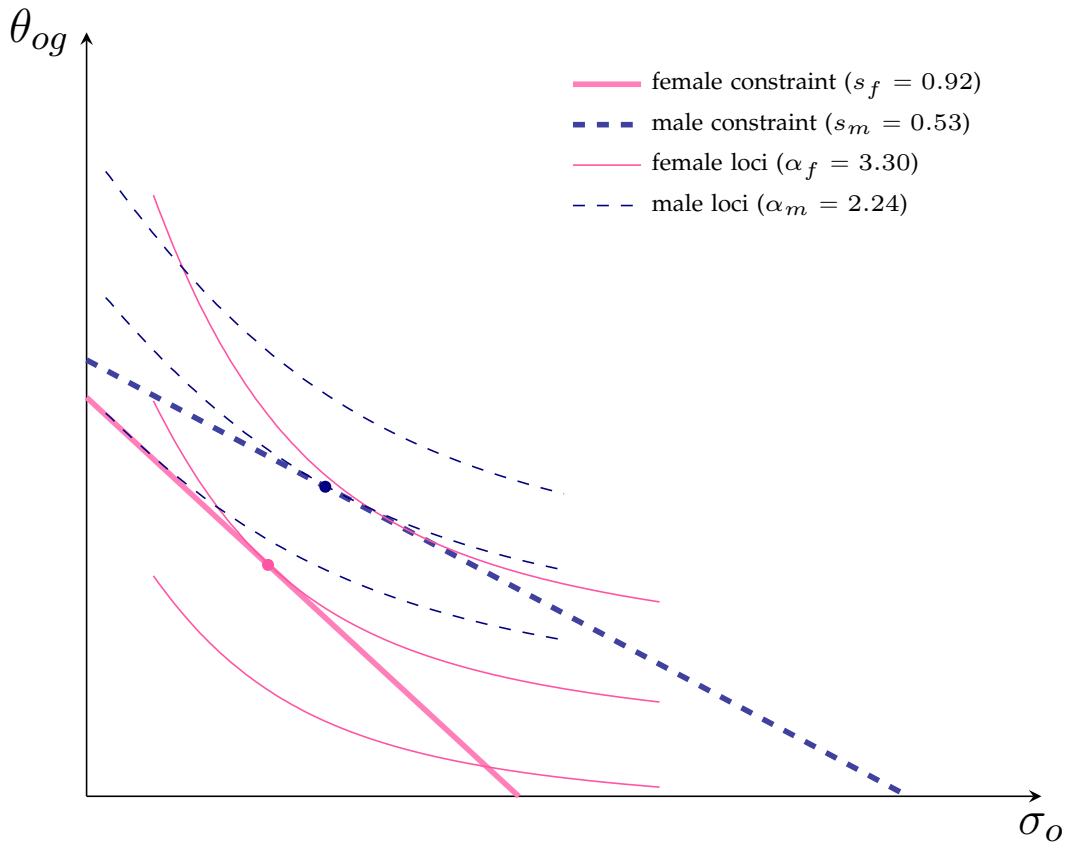
Notes: bar graphs depict average sociability by gender using various measures of social skills. The leftmost figure is based on standardized test scores in the Reading the Mind in the Eyes test from a Lab in the Wild dataset, as used in Greenberg et al. (2023) (reported in Table 1 and Figure 3 of their paper). The middle figure is based on the number of other students who nominated you as their friend in middle school or high school (standardized), as used in Lleras-Muney et al. (2020). The rightmost figure is based on NLSY79 and NLSY97 data, similar to the measures used by Deming (2017) (see Section 2 for details).

FIGURE 3: SKILLS DEMANDED BY LESS FLEXIBLE OCCUPATIONS



Notes: data are from ACS 2015-2019 and O*NET 2013. Observations are 320 3-digit occupations and are weighted by number of workers employed. The scatterplots depict the relationship between an occupation's skill demands and the O*NET index proxying for the price of flexibility. Social perceptiveness represents how important it is that workers are "aware of others' reactions and understanding why they react as they do." Analytical skills capture the importance of skills in mathematics and logical reasoning.

FIGURE 4: OCCUPATIONAL CHOICE: TRADE-OFF BETWEEN COMPARATIVE ADVANTAGE AND FLEXIBILITY



Notes: This figure plots the loci representing worker's earnings/utility at a few different levels for both a representative female worker (solid pink) and a representative male worker (dashed blue). The x-axis represents the occupation's elasticity of temporal substitution (or its "flexibility"), σ_o, and the y-axis represents a worker's skill level in an occupation, θ_{og}. θ_{og} is increasing in the worker's gender's social skill level (s_g) and the occupation's social skill demands (s_o). The solid pink and dashed blue lines represent the feasible set of occupations for the female and male worker, respectively. Workers choose a bundle of θ_{og} and σ_o to maximize their earnings subject to the feasible set of occupations (i.e., bundles of these job attributes).

TABLE 1: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times Inflexibility_o$	0.199*		-0.102	0.300***		-0.084
	(0.106)		(0.165)	(0.104)		(0.154)
$Female_g \times Social_o$		0.321***	0.393**		0.481***	0.538***
		(0.105)	(0.163)		(0.105)	(0.155)
$Female_g \times Analytical_o$				-0.263**	-0.347***	-0.335***
				(0.106)	(0.116)	(0.116)
$Female_g \times Routine_o$				0.246**	0.338***	0.343***
				(0.112)	(0.117)	(0.117)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.792	0.801	0.801	0.802	0.818	0.818
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.002		-0.019**	0.001		-0.016*
	(0.007)		(0.009)	(0.007)		(0.009)
$Female_g \times Social_o$		0.008	0.022**		0.013*	0.024***
		(0.007)	(0.009)		(0.007)	(0.008)
$Female_g \times Analytical_o$				-0.008	-0.014	-0.012
				(0.008)	(0.009)	(0.009)
$Female_g \times Routine_o$				-0.003	-0.000	0.001
				(0.008)	(0.008)	(0.008)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.967	0.968	0.968	0.968	0.968	0.968
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, Female_g. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social perceptiveness intensity ($Social_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE 2: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - HETEROGENEITY BY PARENTAL STATUS

	(1)	(2)	(3)	(4)
	Sorting, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$		Earnings, $\overline{ln(earnings_{oegp})}$	
$Female_g \times Inflexibility_o$	-0.055	-0.044	-0.017*	-0.011
	(0.150)	(0.144)	(0.008)	(0.008)
$Female_g \times Social_o$	0.538***	0.514***	0.024***	0.016**
	(0.155)	(0.145)	(0.008)	(0.008)
$Female_g \times Analytical_o$	-0.335***	-0.316***	-0.012	-0.011
	(0.116)	(0.110)	(0.009)	(0.008)
$Female_g \times Routine_o$	0.343***	0.346***	0.001	0.006
	(0.117)	(0.105)	(0.008)	(0.007)
$Female_g \times Parent_p \times Inflexibility_o$	-0.068**	-0.095**	0.001	-0.013**
	(0.028)	(0.042)	(0.005)	(0.007)
$Female_g \times Parent_p \times Social_o$		0.055		0.018***
		(0.041)		(0.006)
$Female_g \times Parent_p \times Analytical_o$		-0.042		-0.000
		(0.031)		(0.006)
$Female_g \times Parent_p \times Routine_o$		-0.008		-0.013**
		(0.047)		(0.006)
Obs.	2440	2440	2440	2440
Adj. R-sq.	0.818	0.817	0.968	0.969
γ_{egp}	X	X	X	X
γ_{oep}	X	X	X	X

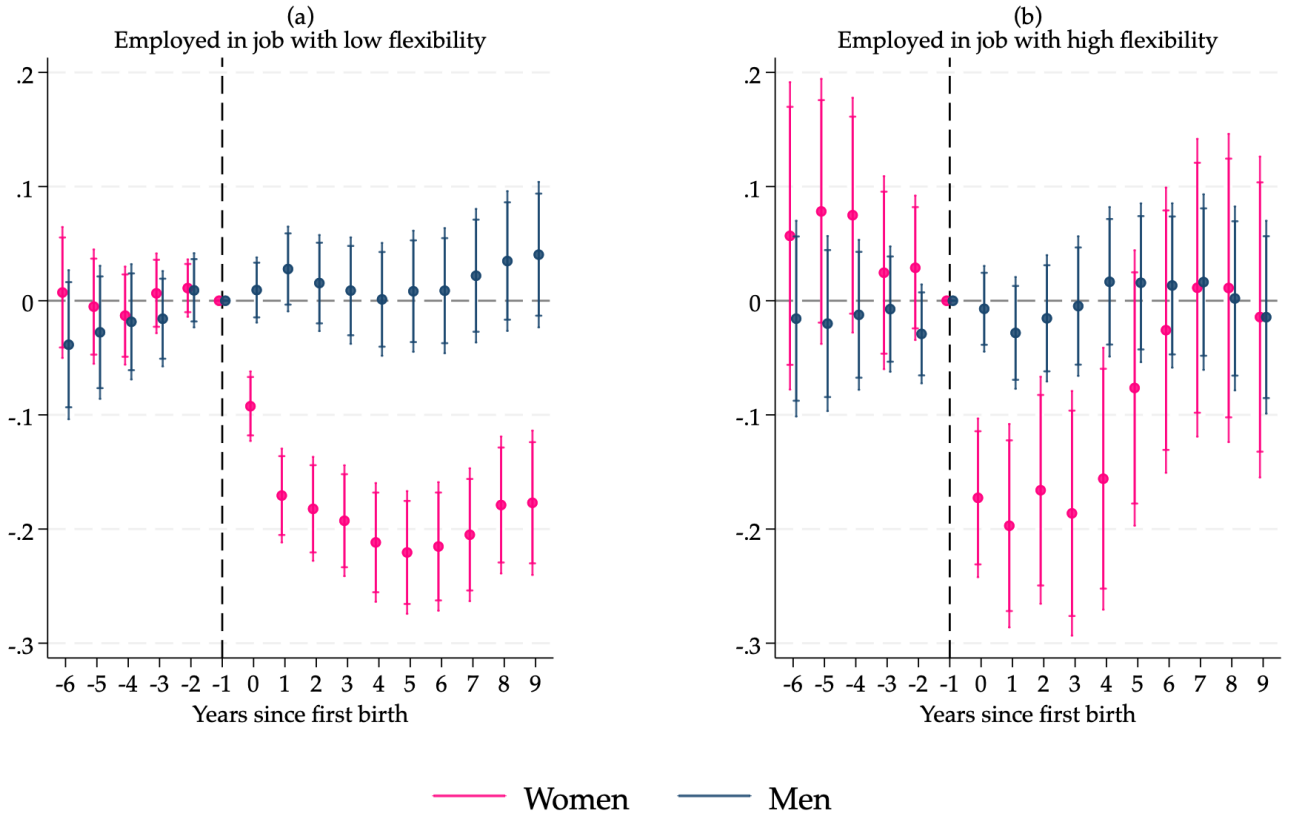
Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female ($Female_g$) and female times parental status ($Female_g \times Parent_p$). Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in columns 1-2 is log employment share for a demographic group in that occupation. The outcome in columns 3-4 is the average of log earnings (residualized with respect to hours worked and demographics). All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE 3: CORRELATION BETWEEN SKILLS AND SORTING INTO JOBS WITH HIGH/LOW FLEXIBILITY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A	Dependent variable: Employed in job with low flexibility						
Social Skills	0.047***		0.030***	0.044***		0.030***	0.029***
	(0.009)		(0.009)	(0.009)		(0.009)	(0.011)
Cognitive Skills		0.108***	0.103***		0.104***	0.100***	0.083***
		(0.008)	(0.008)		(0.009)	(0.009)	(0.010)
Female × Social Skills	-0.035***		-0.023**	-0.028**		-0.023*	-0.022
	(0.012)		(0.012)	(0.012)		(0.012)	(0.015)
Female × Cognitive Skills		-0.022*	-0.017		-0.015	-0.012	0.025
		(0.012)	(0.012)		(0.014)	(0.015)	(0.018)
Female × Has Child × Social Skills							0.001
							(0.020)
Female × Has Child × Cognitive Skills							-0.064***
							(0.022)
Person-Year Obs.	103264	103264	103264	103264	103264	103264	103264
Person Obs.	3396	3396	3396	3396	3396	3396	3396
R-sq.	0.097	0.122	0.124	0.099	0.123	0.125	0.128
F-test: Female × Skills			0.005			0.047	0.881
F-test: Female × Has Child × Skills							0.013
Panel B	Dependent variable: Employed in job with high flexibility						
Social Skills	-0.044***		-0.032***	-0.041***		-0.032***	-0.031***
	(0.008)		(0.008)	(0.008)		(0.008)	(0.010)
Cognitive Skills		-0.076***	-0.071***		-0.071***	-0.067***	-0.051***
		(0.008)	(0.008)		(0.009)	(0.009)	(0.010)
Female × Social Skills	0.024**		0.016	0.018		0.015	0.015
	(0.010)		(0.011)	(0.011)		(0.011)	(0.014)
Female × Cognitive Skills		0.013	0.011		0.003	0.002	-0.023
		(0.011)	(0.012)		(0.014)	(0.014)	(0.018)
Female × Has Child × Social Skills							-0.001
							(0.018)
Female × Has Child × Cognitive Skills							0.050**
							(0.021)
Person-Year Obs.	103264	103264	103264	103264	103264	103264	103264
Person Obs.	3396	3396	3396	3396	3396	3396	3396
R-sq.	0.106	0.118	0.120	0.107	0.118	0.121	0.122
F-test: Female × Skills			0.051			0.313	0.707
F-test: Female × Has Child × Skills							0.040
Controls	x	x	x	x	x	x	x
Ability Controls × Female				x	x	x	x

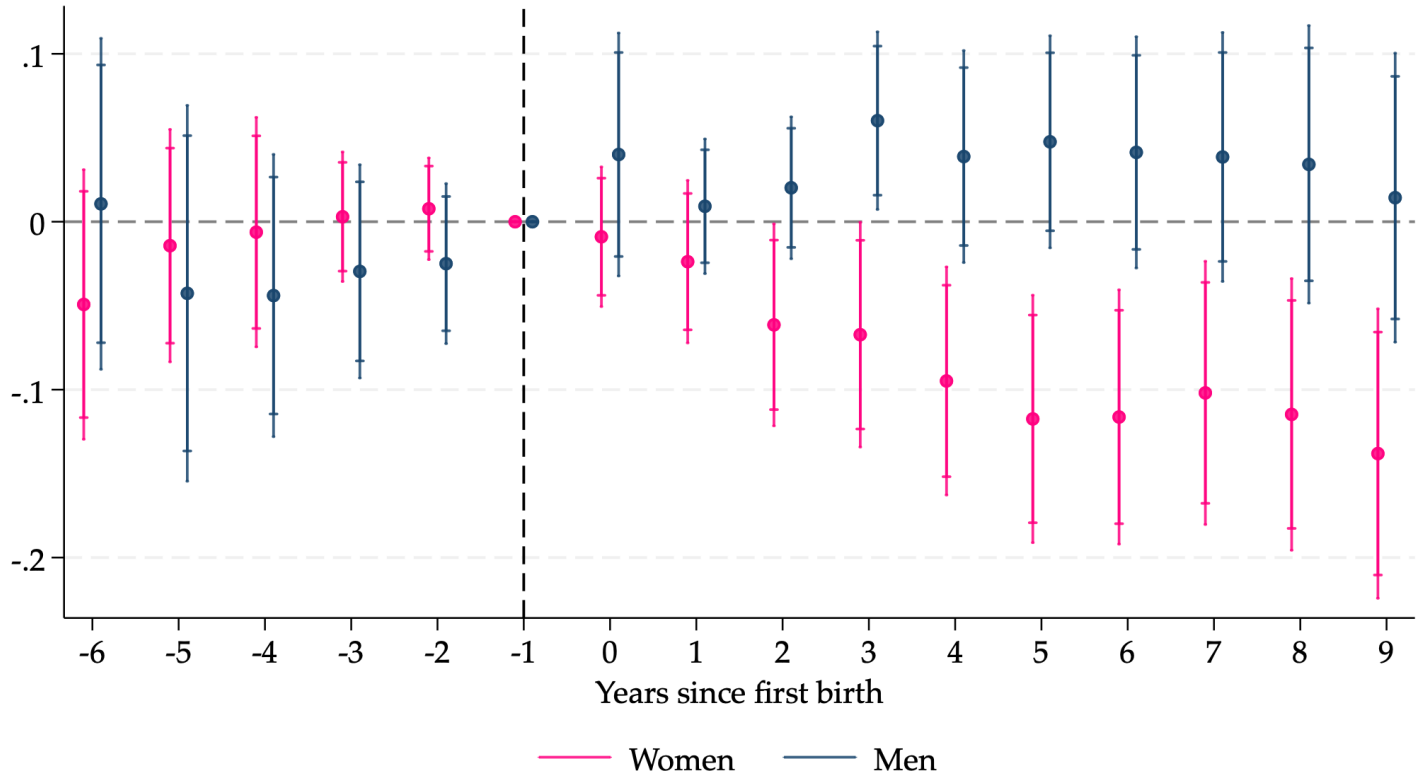
Notes: data are from the NLSY79. The outcome in panel A is a dummy indicating the person is employed in an occupation with low flexibility (i.e., above median in $Inflexibility_o$). The oucome in panel B is a dummy indicating the person is employed in an occupation with high flexibility (i.e., below median in $Inflexibility_o$). Each column presents results from separate regressions of the outcomes on a female dummy, a child dummy, and interactions of these variables with social and cognitive skills. All regressions include a femaly dummy, age fixed effects, year fixed effects, age by female fixed effects, year by female fixed effects, noncognitive skills, a dummy indicating noncognitive skills were missing (and filled with zero), and a dummy for college completion. In columns 4-7, interactions of the female dummy with noncognitive skills and the college dummy are also included. Note that in the final column, all possible interactions of the female dummy, parental status dummy, and skill variables are included although only a subset are displayed. At the bottom of the table, I report two types of F-tests to test for joint significance of the interaction terms with cognitive and social skills. The first F-test tests for significance of Female × Cognitive + Female × Social = 0. The second tests for the significance of Female × Has Child × Cognitive + Female × Has Child × Social = 0. In both cases, I report the corresponding p-value. Standard errors are robust and clustered on individual.

FIGURE 5: CHANGE IN THE PROPENSITY TO BE IN JOBS WITH LOW/HIGH FLEXIBILITY AFTER CHILDBIRTH



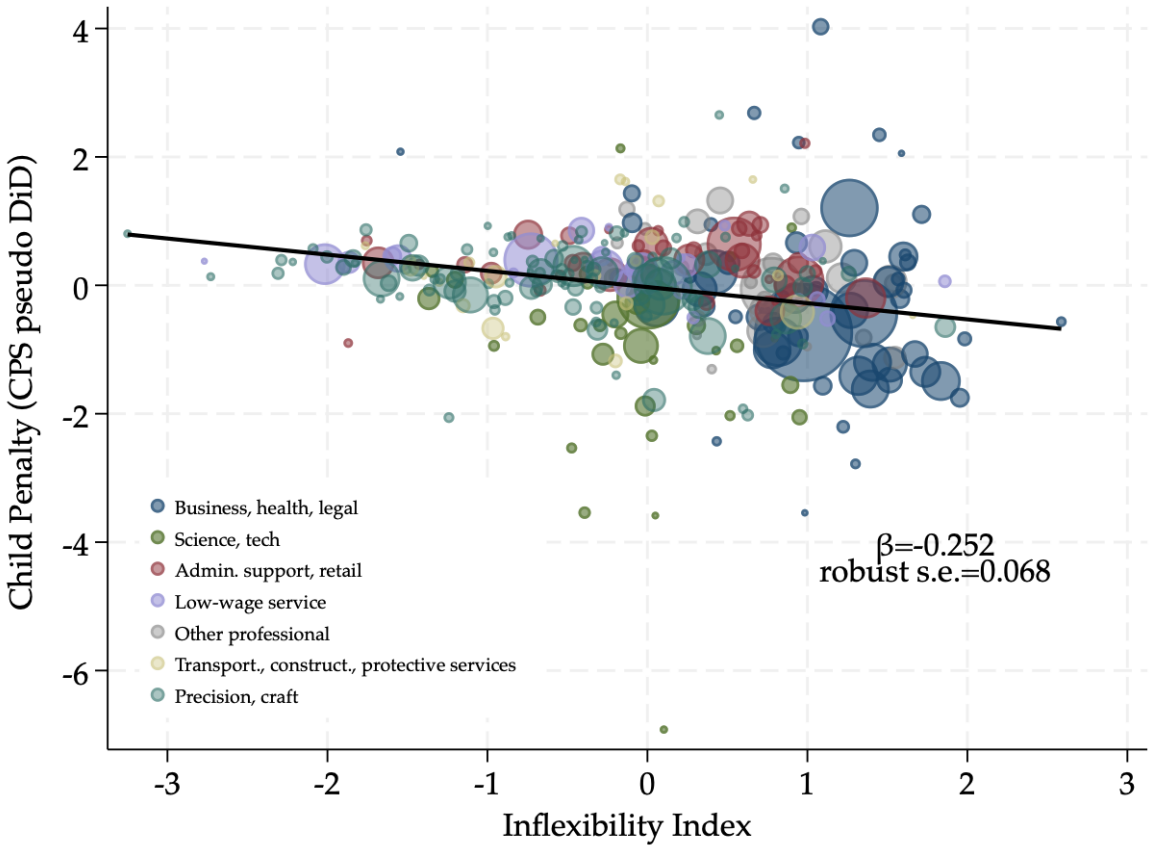
Notes: data are from the NLSY79. This figure presents event studies of first child birth separately by gender. In Panel A, the outcome is whether the person is employed in a job with low flexibility (above median in $Inflexibility_o$). In Panel B, the outcome is whether the person is employed in a job with high flexibility (below median in $Inflexibility_o$). The plotted points represent coefficients on dummies indicating the number of years since an individual's first child birth (k). These coefficients are estimated from $k \in \{-27, 39\}$ but my plots focus on $k \in \{-6, 9\}$ for visibility. Estimates have been divided by pre-birth averages in the outcomes and therefore can be interpreted as percentage changes. 90% and 95% confidence intervals displayed. Standard errors robust and clustered on individual. As discussed in the text, the "control group" are individuals who never give birth over the sample period.

FIGURE 6: HETEROGENEITY IN WAGE CHANGES AFTER CHILDBIRTH BY FLEXIBILITY OF PRE-BIRTH OCCUPATION



Notes: this figure presents heterogeneity in how wages respond to child birth by the flexibility of the individual's occupation prior to birth. The outcome is an individual's wage (coded as zero if not working). The plotted points represent coefficients on event time dummies interacted with a variable representing the lack of flexibility in the individual's occupation. Specifically, the variable is the average of the index proxying for inflexibility ($Inflexibility_o$) over all pre-birth observations for the individual. These coefficients are estimated from $k \in \{-27, 39\}$ but my plots focus on $k \in \{-6, 9\}$ for visibility. Estimates are divided by pre-birth averages in the outcome and therefore represent percentage changes. The estimates in pink represent the results for women, while the estimates in blue represent the results for men. 90% and 95% confidence intervals displayed. Standard errors robust and clustered on individual. Data are from the NLSY79. As discussed in the text, the "control group" are individuals who never give birth over the sample period.

FIGURE 7: CORRELATION BETWEEN THE Δ SORTING GAP AFTER CHILDBIRTH & THE OCCUPATION'S FLEXIBILITY



Notes: in this figure, each observation is an occupation. The index proxying for inflexibility of an occupation ($Inflexibility_o$) is plotted on the x-axis. The y-axis plots the estimated change in women's relative propensity to sort into occupation after childbirth. To estimate these changes, I build synthetic individuals (matched on gender, age, year, marital status, education, race, state) in the CPS pooled from 1972-2019. Then, separately by gender, I regress a dummy indicating that the person is employed in that occupation on a dummy indicating that the (synthetic) person has given birth (including FE for the match variables). To transform the estimates into percentage changes, I divide the coefficient on the post-birth indicator by the average in the outcome among the childless sample. The result represents the difference in this estimate between women and men for a given occupation. For interpretability, I standardize this to have a mean of 0 and standard deviation of 1.

TABLE 4: ENDOGENEITY OF FLEXIBILITY OVER TIME

	(1)	(2)	(3)	(4)
Panel A	Dep. Var.: $\Delta \widehat{MotherFriendly}$ (CPS)			
$Social_o$	0.019		0.020	0.028
	(0.052)		(0.096)	(0.106)
$Inflexibility_o$		0.090***	0.249**	0.258**
		(0.034)	(0.101)	(0.121)
$Inflexibility_o \times Social_o$			0.128**	0.124*
			(0.055)	(0.065)
$Analytical_o$				0.002
				(0.115)
$Inflexibility_o \times Analytical_o$				0.005
				(0.056)
$Social_o \times Analytical_o$				0.026
				(0.074)
Obs.	320	320	320	320
R-squared	0.002	0.016	0.065	0.071
Panel B	Dep. Var.: Δ Avg. $\widetilde{ShareWFH}$ (ATUS)			
$Social_o$	0.006		0.001	0.001
	(0.005)		(0.008)	(0.009)
$Inflexibility_o$		0.008**	0.013	0.006
		(0.004)	(0.010)	(0.011)
$Inflexibility_o \times Social_o$			0.013***	0.011*
			(0.005)	(0.006)
$Analytical_o$				0.012
				(0.011)
$Inflexibility_o \times Analytical_o$				-0.004
				(0.007)
$Social_o \times Analytical_o$				0.010
				(0.006)
Obs.	313	313	313	313
R-squared	0.022	0.033	0.063	0.071
Base Period Controls	x	x	x	x
Base Period Controls $\times$ All Tasks			x	x

Notes: The outcome in Panel A is the change between periods 1972-1995 and 1996-2019 in the female-male employment penalty in each occupation. As discussed in the text, these are estimated for each occupation via pseudo difference in differences in the CPS data. The outcome in Panel B is the change in the average share of hours workers work-from-home in each occupation between periods 2003-2010 and 2011-2019. As discussed in the text, the share of a hours a worker works from home has been residualized by demographic characteristics. Both panels report the results from regressing these outcomes on various occupation-level characteristics from the O*NET. In all columns, I have included controls for the fraction of workers with a college education in the occupation and the fraction of workers who are female, both calculated in the CPS for the relevant base period. Standard errors are robust and observations are weighted by the number employed in the occupation in the relevant base period.

Appendix

Appendix A Data Appendix

This data appendix describes the occupation codes, data sources, and various indices I use in the paper in more detail.

Appendix A.1 Occupation codes

Across the various datasets and time periods I analyze in this paper, I rely on a consistent and balanced set of occupation codes, first created by Autor and Dorn (2013), then subsequently modified by Deming (2017) and, later, Hanson (2021) to accommodate more recent changes to the occupational classification system. The codes represent 320 3-digit occupation codes. For each data source below, I crosswalk occupations into these more aggregate codes, and I describe this process as needed. The number of occupations represented in my analysis fluctuates slightly between analyses because certain occupations are not represented in the dataset or the occupation is too small.

Appendix A.2 Occupational Network (O*NET)

I use the 2013 O*NET data to measure various attributes of occupations. More recent versions of the O*NET are available, but I use this version in order to most closely mimic Goldin (2014). When the measure is a composite of individual components, each component is first standardized, then the measure is re-standardized after taking the average.

The raw O*NET data are at the 8-digit 2010 SOC level. Before constructing these measures, I aggregate the O*NET variables up to 3-digit occupation codes. I begin by taking the simple average across occupations to aggregate to the 6-digit 2010 SOC level. Next I crosswalk 2010 SOC codes to the more aggregate Deming (2017) occupation codes. To build this crosswalk, I use crosswalks developed by the Census Bureau and Deming (2017) to aggregate 2010 SOC codes to 2010 Census Occupation Code and then subsequently to Deming (2017) codes. Finally, I aggregate these codes to the Hanson (2021) codes, taking the 2010 ACS employment weighted average.

Inflexibility Index My flexibility measure follows Goldin (2014) and represents the average of five standardized components (which is itself re-standardized). Below I list the survey questions and response scales for each component. For each component, the larger the value, the less flexible. Each of these dimensions come from the Work Context module, except the Establishing and Maintaining Interpersonal Relationships, which comes from Work Activities.

1. ***Time pressure***: How often does this job require the worker to meet strict deadlines? Never/Once a year or more but not every month/Once a month or more but not every week/Once a week or more but not every day/Every day.

2. ***Contact with others***: How much does this job require the worker to be in contact with others (face-to-face, by telephone, or otherwise) in order to perform it? No

contact with others/Occasional contact with others/Contact with others about half the time/Contact with others most of the time/Constant contact with others.

3. ***Establishing and Maintaining Interpersonal Relationships***: How important is establishing and maintaining interpersonal relationships to the performance of this job?[47] Not important/Somewhat important/Important/Very important/Extremely important.

4. ***Unstructured Work***: How much freedom do you have to determine the tasks, priorities, or goals of your current job? No freedom/Very little freedom/Limited freedom/Some freedom/A lot of freedom.
   Note that I retrieved this survey question from the questionnaire itself. There is a discrepancy in the O*NET lookup file "Content Model Reference", which indicate the survey question as "To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?" This is the survey question wording that is printed in Goldin (2014). At the same time, "Survey Booklet Location" lookup file refers to O*NET Work Context questionnaire question #52, which states "How much freedom do you have to determine the tasks, priorities, or goals of your current job?" (it is unclear which archived questionnaire maps to this version of the O*NET, but this is true of questionnaires from 2008, 2012, and 2015.) The survey question "To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?" comes from the 1998 O*NET data dictionary. In either case (the O*NET questionnaires from 2008, 2012, and 2015 and the 1998 O*NET data dictionary), the response scale is such that a higher value is more freedom or less structure. Therefore, the exact wording of the survey question is not likely to be consequential from a coding perspective. To verify my work, I cross-checked that my standardized value for this dimension for Lawyers and Judges (1.04) is similar to the value printed in Table 2 of Goldin (2014), which is 1.22. (They will not line up exactly because I am standardizing within a set of occupations that may be aggregated slightly differently.)

5. ***Freedom to Make Decisions***: How much decision making freedom, without supervision, does the job offer? No freedom/Very little freedom/Limited freedom/Some freedom/A lot of freedom.

These components are intended to triangulate on time-intensity and the substitutability of workers, which jointly make it more valuable for workers to work particular, long, or continuous hours. One way of thinking about this is that these are two complementary features in inflexibility. Each feature is necessary but not sufficient. Consider one extreme: if workers are not substitutable for another, but there is no time-sensitivity in task completion, then workers can do their tasks whenever they want. Consider another extreme: if there is time-sensitivity, but workers are perfectly substitutable for one another, then it does not matter *who* the firm calls on to complete a task. The first dimension captures some time sensitivity in task completion. The second two dimensions capture the necessity of working alongside people (who may need to be served at certain times), with the latter specifically capturing the continuity of relationships. The final two dimensions capture the level of worker discretion; the more

[47] The questionnaire defines establishing and maintaining interpersonal relationships as "developing constructive and cooperative working relationships with others, and maintaining them over time."

freedom over one's task set, the less likely others can easily substitute for you in your absence. Occupations that are considered highly inflexible according to this measure include lawyers, CEO's, and managers.

Social perceptiveness skill intensity: I use the O*NET component for social perceptiveness from the skills module which is defined as "being aware of others' reactions and understanding why they react as they do."[48] The response scale ranges from 1-7, with examples given for a rating of 2 as "notice that customers are angry because they have been waiting too long" and for a rating of 6 as "counsel depressive patients during a crisis period."

Routine task intensity: Routine task intensity is the average of two components from the Work Context module. The first component is the "Degree of Automation," which asks "How automated is the job? Not at all automated/Slightly automated/Moderately automated/Highly automated/Completely automated." The second component is the Importance of Repeating Same Tasks, which asks "How important is repeating the same physical activities (e.g., key entry) or mental activities (e.g., checking entries in a ledger) over and over, without stopping, to performing this job? Not important at all/Fairly important/Important/Very important/Extremely important."

Analytical (or cognitive) skill intensity: is the average of three components. The first, "Mathematical Reasoning" from the Abilities module, captures the level of "the ability to choose the right mathematical methods or formulas to solve a problem" needed in a job. The second is from the knowledge module and captures the level of "knowledge of arithmetic, algebra, geometry, calculus, statistics, and their applications" required in the job. The third is from the skills module and captures the level of "using mathematics to solve problems" required in the job. Each has a response scale varying from 1-7.

Appendix A.3 Alternative measures of flexibility

As discussed in Section 2.1.1, I use two alternative proxies of inflexibility in robustness checks.

Inflexibility and Overwork This measure follows (Goldin et al., 2024) and is a binary variable equal to one if the occupation is in the top tercile of both the O*NET index ($Inflexibility_o$) and a measure of overwork (i.e. long hours). This measure equals the share of men in an occupation who work over 50 hours per week. Consistent with the other ACS data I use throughout the paper, I base my measure on of overwork on the 2015-2019 5% sample (described below). (Goldin et al., 2024)'s measure uses a 45 hour cut-off, but I follow (Cortés and Pan, 2016) and use a 50 hour cut-off.

Index of Overwork, Coordinated Hours, and Convexity of Earnings This measure is the average of three standardized components, each from a separate paper in the literature. The first is the overwork measure referenced above.

[48] Specifically, the O*NET asks "What level of social perceptiveness is needed to perform your current job?" and defines social perceptiveness as "being aware of others' reactions and understanding why they react as they do."

The second is a measure of coordinated work hours from Cubas et al. (2023). It represents the ratio of hours worked in the 8am to 5pm interval based on time diary data from the American Time Use Survey (ATUS), pooled between 2003-2019. To construct this, following Cubas et al. (2023), I limit the sample to 18-65 year old full time workers who worked full-time (minimum of 35 hours) in their main job. Both week days and weekends are included. Then, for each occupation I construct $ratio8to5_o = \frac{B_o}{A_o+B_o+C_o}$ where A, B, and C are the total minutes worked by individuals in the occupation (weighted by their sample weight) for time intervals 12-8am, 8am-5pm, and 5pm-12am, respectively. This ratio is higher if individuals' work hours are more concentrated in the 8am-5pm window. Cubas et al. (2023) show that the gender wage gap, especially among parents, is increasing in this measure.

The third is a measure of the convexity of earnings with respect to hours worked similar to that used by Cortés and Pan (2019). I use the ACS 2015-2019 sample (described below) to regress log annual income on log annual hours, occupation fixed effects, and log annual hours interacted with occupation fixed effects. I also control for a female dummy, parental status dummy, age dummies, race dummies, year dummies, and dummies for years of education. The regression is weighted by the ACS sampling weights. I store the coefficient on the interaction between log annual hours and the occupation fixed effects. For interpretability, I then add these coefficients to the coefficient on log annual hours. The final measure thus varies at the occupation level and represents the elasticity of earnings with respect to hours worked. In contrast to Cortés and Pan (2019), I use annual hours worked (rather than weekly), I include the full sample of occupations rather than limit to skilled occupations, I include both men and women, and I do not limit to college educated individuals (although I do control for education).

Appendix A.4 American Community Survey (2015-2019) 5% Sample

I rely on the 2019 5-year American Community Survey (ACS) 5% sample where I observe individuals' wages and occupation choices. My sample restrictions primarily follow Goldin (2014). I restrict the sample to native-born, civilian individuals aged 25-64 who are employed full-time, full-year (at least 40 weeks, at least 35 hours per week) and report positive, non-missing annual earnings (wage income plus business and farm income). I omit occupation-demographic group cells for which I do not observe at least 25 men and 25 women so that I can reliably observe the gender wage gap within occupations. Note that weeks worked is only provided in binned values for years 2015-2018. Therefore, I impute the average number of weeks worked within each bin using the individuals in the 2019 sample.

Appendix A.5 National Longitudinal Survey of Youth - 1979 Cohort

I use the Weekly Work History files in the National Longitudinal Survey of Youth for the 1979 Cohort to build a person-year panel of labor market activity (e.g., employment, occupation, and wages) covering years 1979-2019. The NLSY79 Weekly Work History files compile information reported by respondents about their job held at the time of the survey as well as retrospective information on any jobs since the last interview. This allows me to form a relatively complete annual panel.

Most of my sample restrictions follow Guvenen et al. (2020) and are intended to

focus on individuals who signal some initial attachment to the labor market and observations for which all I observe sufficient information. Table 5 summarizes how each restriction affects the number of observations. The raw data are at the person-job (or nonemployment) spell level. In each year, I focus on the job to which the respondent contributed the most hours. I begin by dropping observations in which the individual is not yet 22 or who are likely still in school (based on their age relative to their current years of education). I then drop observations prior to the person reporting 2 consecutive years of full-time full-year (1400 annual hours) work. Next, I drop individuals with two or more years of military experience as well as any observations constituting military spells. I also drop brief labor market spells which occur prior to finishing one's education. Specifically, I drop observations where the person had 7 or less years of experience before pausing to acquire another degree. In order to clean the observations pertaining to one's job, I temporarily drop spells where the individual is not working (but I append these back to the panel later). I then clean duplicated information about a job within a year, given that I am aggregating the panel to the person-year level.[49] Next, I drop observations with missing occupation or wage information. I also drop observations with wages either below the 0.5th or above the 99.5th percentile of wage distribution in a given interview year. I drop individuals without information on social skills, cognitive skills, or their age at first birth. At this point, I merge back in the non-working spells and drop years for which there is no information on labor market status in any week. Finally, among individuals who have a child, I drop individuals who are not observed in the period just prior to birth and at least one period after treatment. While sample size precludes me from forcing a perfectly balanced sample, this helps me ensure a reasonable amount of balance in the sample. The final sample is 3,396 individuals and 103,264 observations.

TABLE 5: SUMMARY OF SAMPLE RESTRICTIONS

	Person-Year Obs.	**Individuals**	**% Obs. Remaining**
Full Sample	360221	6111	100.00
Reason observations dropped			
Below age 22 or in school	311403	6111	86.45
Prior to working FTFY for 2 years	281811	5755	78.23
Military individuals/spells	280674	5732	77.92
Prior to finishing education	280213	5732	77.79
Temporarily dropping non-working spells	232247	5732	64.47
Clean duplicates	154243	5732	42.82
Missing labor market info	146357	5716	40.63
Missing demographic info	180329	5229	50.06
Missing labor market status	154585	5229	42.91
Not observed pre/post birth	103264	3396	28.67

[49] This occurs because people report their labor market information for each job between the current and last interview. In this process, I favor observations with more information, e.g. sometimes the person reports the wage a job in one survey round but not the other.

Appendix A.6 Current Population Survey, 1972-2019

I follow Kleven (2023) to build "synthetic" individuals in pooled CPS data in order to estimate difference-in-differences in Sections 7.2 and 8.

I focus on individuals aged 20-55 in the CPS ASEC and basic monthly files pooled for years 1972-2019. I drop active duty military and those without an identifiable race or occupation (conditional on being employed). I identify individual's approximate age at first birth based on their own age and the age of the eldest child living with them. I then match these individuals to 5 pre-birth surrogate observations. To do this, I retrieve outcomes of childless individuals in the same birth cohort who share the same demographic characteristics. For example, consider a woman observed in year 2000 at age 35 whose eldest child is age 5. To retrieve her surrogate observation in the year prior to birth ($k = -1$), I locate a childless, 30-year old woman in the year 1995 who shares the same demographic characteristics. The same process is performed for surrogate observations $k \in [-5, -4]$.

The demographic characteristics that I match on include: gender (male/female), age, year, education category (below high school, high school, some college, BA or more), marital status (married with spouse present, separated or spouse absent, divorced, widowed, never married), race (white non-hispanic, black non-hispanic, hispanic, other), and state. When there are multiple people with this set of demographic characteristics, I use the average of the outcome for everyone in the cell to form the surrogate observation. Among the treated individuals, I focus on those whose approximate first birth occurred between ages 25-45 and observations up to 10 years after the event. My final sample is comprised of synthetic individuals, each with 6 observations (1 post birth and 5 pre-birth).

Appendix B Conceptual Framework

B.1 Labor supply sub-problem solutions

The labor supply sub-problem is given by

$$\underset{t_1,t_2}{max}\ \theta_g^{\mu}\left(0.5t_1^{\rho}+0.5t_2^{\rho}\right)^{\frac{1}{\rho}}+\lambda\left(T-\alpha_g t_1-t_2\right)$$

where I have dropped the occupation (o) subscripts for brevity. The First Order Conditions are:

$$\theta_g^{\mu}\left(0.5t_1^{\rho}+0.5t_2^{\rho}\right)^{\frac{1}{\rho}-1}0.5t_1^{\rho-1}-\alpha_g\lambda=0$$

$$\theta_g^{\mu}\left(0.5t_1^{\rho}+0.5t_2^{\rho}\right)^{\frac{1}{\rho}-1}0.5t_2^{\rho-1}-\lambda=0$$

$$T-\alpha_g t_1-t_2=0$$

Letting $\rho=\frac{\sigma-1}{\sigma}$, the solutions are therefore:

$$t_1^*\left(\alpha_g,T,\sigma\right)=\left(\frac{0.5}{\alpha_g}\right)^{\sigma}\frac{T}{\alpha_g^{1-\sigma}0.5^{\sigma}+0.5^{\sigma}}$$

$$t_2^*(\alpha_g, T, \sigma) = 0.5^\sigma \frac{T}{\alpha_g^{1-\sigma} 0.5^\sigma + 0.5^\sigma}$$

$$\lambda^*(\theta_g, \alpha_g, \sigma) = \theta_g \left[0.5^\sigma \alpha_g^{1-\sigma} + 0.5^\sigma\right]^{\frac{1}{\sigma-1}}$$

B.2 Derivations of empirical predictions

Proof of Prediction 1 Prediction 1 states that the positive impact of social tasks or flexibility on women's relative employment and wages will be muted to the degree that occupations high in social task intensity also lack flexibility. Below I derive this result for women's relative wages. Given that individuals choose the occupation which maximizes their wages, these results can also be applied to women's relative employment (e.g., if women's relative wages are increasing in social task intensity, so too is their propensity to choose that occupation over other ones).

To fix ideas, I begin by stating Prediction 1 more formally. Recall the worker's occupational choice problem, which we can restate as follows. Workers' log earnings are given by

$$ln\overline{Q}_{og} = ln\left[\left(\underbrace{s_g s_o + \kappa_0}_{\theta_{og}}\right)^\mu T 0.5^{\frac{\sigma_o}{\sigma_o - 1}} \left[\alpha_g^{1-\sigma_o} + 1\right]^{\frac{1}{\sigma_o - 1}}\right] \tag{13}$$

They choose s_o and σ_o to maximize 13 subject to the constraint $s_o = -\kappa_1 \sigma_o$. In equilibrium, the two following equations must hold:

$$ln\left(\frac{\overline{Q}_{of}^*}{\overline{Q}_{om}^*}\right) = \mu ln\left[s_f s_o^* + \kappa_0\right] + \frac{1}{\sigma_o^* - 1} ln\left[\alpha_f^{1-\sigma_o^*} + 1\right] - \mu ln\left[(s_m s_o^* + \kappa_0)\right] - \frac{1}{\sigma_o^* - 1} ln\left[\alpha_m^{1-\sigma_o^*} + 1\right] \tag{14}$$

$$s_o^* = -\kappa_1 \sigma_o^* \tag{15}$$

where the first equation represents the wage gap between a representative female and male worker in occupation o, and where s_o^* and σ_o^* represent individuals' optimal choices of social task intensity and flexibility, respectively. Let $ln\left(\overline{Q}_{of}^*/\overline{Q}_{om}^*\right) = \bar{Q}_o^*$ and define

$$\bar{Q}_o^* = f\left(s_o^*\left(\kappa_1, \sigma_o^*\right), \sigma_o^*\right)$$
$$\bar{Q}_o^* = g\left(s_o^*, \sigma_o^*\left(\frac{1}{\kappa_1}, s_o^*\right)\right)$$

and therefore

$$\frac{\partial f}{\partial \sigma_o^*} = \frac{\partial \bar{Q}_o^*}{\partial s_o^*}\frac{\partial s_o^*}{\partial \sigma_o^*} + \frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} \tag{16}$$

$$\frac{\partial g}{\partial s_o^*} = \frac{\partial \bar{Q}_o^*}{\partial s_o^*} + \frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*}\frac{\partial \sigma_o^*}{\partial s_o^*} \tag{17}$$

Prediction 1 implies that while s_o and σ_o increase women's relative wages, the inverse correlation between these two attributes leads the positive impact of one to be offset

by the negative impact of a reduction in the other attribute:

$$\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} > 0$$
$$\frac{\partial \bar{Q}_o^*}{\partial s_o^*} > 0$$
$$\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} > \frac{\partial f}{\partial \sigma_o^*}$$
$$\frac{\partial \bar{Q}_o^*}{\partial s_o^*} > \frac{\partial g}{\partial s_o^*}$$

Now I show that $\frac{\partial \bar{Q}_o^*}{\partial s_o^*} > 0$. Taking the derivative of 14 with respect to s_o^*, we get

$$\frac{\partial \bar{Q}_o^*}{\partial s_o^*} = \mu \left[\frac{1}{s_o^* + \frac{1}{s_f}\kappa_0} - \frac{1}{s_o^* + \frac{1}{s_m}\kappa_0} \right] > 0 \tag{18}$$

because $s_f > s_m$. This implies that social task intensity has a positive impact on women's relative wages, all else constant. Next, I show that $\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} > 0$. Notice that

$$\frac{\partial ln \overline{Q}_{og}^*}{\partial \alpha_g} = \frac{\partial}{\partial \alpha_g} \left[\frac{1}{\sigma_o^* - 1} ln \left[\alpha_g^{1-\sigma_o^*} + 1 \right] \right] = -\frac{1}{\frac{1}{\alpha_g^{\sigma_o^* - 1}} + 1} \frac{1}{\alpha_g^{\sigma_o^*}} < 0 \tag{19}$$

where I have applied the envelope theorem. And,

$$\frac{\partial^2 ln \overline{Q}_{og}^*}{\partial \alpha_g \partial \sigma_o^*} = \frac{ln\left(\alpha_g\right)}{\left(\alpha_g^{1-\sigma_o^*} + 1\right)^2 \alpha_g^{\sigma_o^*}} \geq 0$$

because $\alpha_g \geq 1$. Therefore,

$$\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} = \frac{\partial}{\partial \sigma_o^*} \left(\frac{1}{\sigma_o^* - 1} ln \left[\alpha_f^{1-\sigma_o^*} + 1 \right] - \frac{1}{\sigma_o^* - 1} ln \left[\alpha_m^{1-\sigma_o^*} + 1 \right] \right) \geq 0 \tag{20}$$

thus implying that flexibility has a positive impact on women's relative wages, all else constant. Finally, it is easy to see that $\frac{\partial s_o^*}{\partial \sigma_o^*} = -\kappa_1 < 0$ and $\frac{\partial \sigma_o^*}{\partial s_o^*} = -\frac{1}{\kappa_1} < 0$. Therefore

$$\frac{\partial f}{\partial \sigma_o^*} = \underbrace{\frac{\partial \bar{Q}_o^*}{\partial s_o^*}}_{+} \underbrace{\frac{\partial s_o^*}{\partial \sigma_o^*}}_{-} + \underbrace{\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*}}_{+} \rightarrow \frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} > \frac{\partial f}{\partial \sigma_o^*}$$
$$\frac{\partial g}{\partial s_o^*} = \underbrace{\frac{\partial \bar{Q}_o^*}{\partial s_o^*}}_{+} + \underbrace{\frac{\partial \bar{Q}_o^*}{\partial \sigma_o^*} \frac{\partial \sigma_o^*}{\partial s_o^*}}_{+ \quad -} \rightarrow \frac{\partial \bar{Q}_o^*}{\partial s_o^*} > \frac{\partial g}{\partial s_o^*}$$

While moving to an occupation with higher σ_o results in higher wages for women, this transition is bundled with a *lower* s_o which acts as a countervailing force on their wages. And, equivalently, while higher s_o closes the wage gap, an occupation with

higher s_o also results in a lower σ_o which serves to widen the gap.■

To further build intuition for this result, I briefly outline how I test Prediction 1 in the data. Let y_o represent either women's relative average wages ($ln\left(\bar{wage}_o^f\right) - ln\left(\bar{wage}_o^m\right)$) or women's relative propensity to sort into an occupation ($ln\left(\frac{employed_o^f}{employed^f}\right) - ln\left(\frac{employed_o^m}{employed^m}\right)$), which I measure in the ACS. Let s_o and σ_o represent social task intensity and flexibility as measured in the O*NET. Consider the regression models:

$$y_o = \beta_0 + \beta_1\sigma_o + \beta_2 s_o + \varepsilon_o \tag{21}$$
$$y_o = \rho_0 + \rho_1\sigma_o + \epsilon_o \tag{22}$$
$$y_o = \gamma_0 + \gamma_1 s_o + \upsilon_o \tag{23}$$

Prediciton 1 implies that

$$\beta_1 > 0$$
$$\beta_2 > 0$$
$$\beta_1 > \rho_1$$
$$\beta_2 > \gamma_1$$

To see why the final two inequalities hold, notice that $\rho_1 = \beta_1 - \kappa_1\beta_2 \implies \beta_1 > \rho_1$ and $\gamma_1 = \beta_2 - \frac{1}{\kappa_1}\beta_1 \implies \beta_2 > \gamma_1$ where κ_1 represents the relationship between flexibility and social task intensity from the regression model $s_o = -\kappa_1\sigma_o + e_o$. In other words, the positive relationship between flexibility and social task intensity and women's relative wages and employment is biased towards $-\infty$ in models 22 and 23.

Proof of Prediction 2 Prediction 2 states that the correlation between women's social skills and their propensity to sort into less flexible jobs will be weaker than men's. That is:

$$\frac{\partial\sigma_o^{*,m}}{\partial s_g} < \frac{\partial\sigma_o^{*,f}}{\partial s_g} < 0$$

where $\sigma_o^{*,g}$ represents the optimal choice of flexibility for a representative worker of gender g. First, I show that the optimal flexibility is decreasing in a group's social skills: $\frac{\partial\sigma_o^*}{\partial s_g} < 0$. To see this, reformulate the worker's occupational choice problem as an unconstrained problem where they choose σ_o to maximize the objective function

$$\mu ln\left(-\kappa_1 s_g\sigma + \kappa_0\right) + ln\left(T\right) + \sigma\left(\sigma-1\right)^{-1} ln\left(0.5\right) + \left(\sigma-1\right)^{-1} ln\left(\alpha_g^{1-\sigma}+1\right)$$

where I have plugged the constraint $s_o = -\kappa_1\sigma_o$ into expression 14. This objective function has increasing differences in $-s_g$ and σ_o:

$$\frac{\partial^2}{\partial\left(-s_g\right)\sigma_o}\left[\mu ln\left(-\kappa_1 s_g\sigma_o + \kappa_0\right) + ln\left(T\right) + \sigma_o\left(\sigma_o-1\right)^{-1} ln\left(0.5\right) + \left(\sigma_o-1\right)^{-1} ln\left(\alpha_g^{1-\sigma_o}+1\right)\right]$$

$$= \frac{\partial}{\partial \sigma_o}\left[\frac{\mu\kappa_1\sigma_o}{-\kappa_1 s_g\sigma_o + \kappa_0}\right]$$
$$= \frac{\partial}{\partial \sigma_o}\left[\frac{\mu\kappa_1}{-\kappa_1 s_g + \kappa_0\frac{1}{\sigma_o}}\right] > 0$$

Therefore, by Topkis (1978)'s monotonicity theorem,

$$\frac{\partial \sigma_o^*}{\partial\left(-s_g\right)} > 0 \rightarrow \frac{\partial \sigma_o^*}{\partial s_g} < 0$$

Next, I show that because women value σ_o more, they *less willingly forgo* σ_o in favor of more s_o. Consider the formulation of the worker's occupational choice problem in the main text (see Equation 3). The First Order Conditions of the worker's occupational choice problem are

$$\frac{\partial \mathcal{L}}{\partial \theta_o} = \underbrace{\frac{\partial ln\bar{Q}_{og}}{\partial \theta_o}}_{MU_{\theta_o}} - \lambda = 0$$

$$\frac{\partial \mathcal{L}}{\partial \sigma_o} = \underbrace{\frac{\partial ln\bar{Q}_{og}}{\partial \sigma_o}}_{MU_{\sigma_o}} - \lambda\kappa_1 s_g = 0$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \kappa_0 - \theta_o - \kappa_1 s_g \sigma_o = 0$$

This implies that in equilibrium,

$$\underbrace{\frac{MU^g_{\sigma_o}}{MU^g_{\theta_o}}}_{MRS^g} = \kappa_1 s_g \tag{24}$$

Notice that marginal utility with respect to θ_o does not vary by gender:

$$MU^g_{\theta_o} = \frac{\mu}{\theta_o}$$

But,

$$MU^g_{\sigma_o} = \frac{\partial}{\partial \sigma_o}\left[\sigma_o\left(\sigma_o - 1\right)^{-1} ln\left(0.5\right) + \left(\sigma_o - 1\right)^{-1} ln\left(\alpha_g^{1-\sigma_o} + 1\right)\right]$$

In the derivation of Prediction 1, I showed that

$$\frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_f^{1-\sigma_o} + 1\right]\right) - \frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_m^{1-\sigma_o} + 1\right]\right) > 0$$

Therefore,

$$MU^f_{\sigma_o} > MU^m_{\sigma_o}$$

which further implies

$$MRS^f > MRS^m$$

In other words, women more willingly trade θ_o for σ_o and, conversely, less willingly give up σ_o in favor of more θ_o. Given equation 24, as s_g increases, both men and women

will choose less σ_o^*, but this will be less true for women and thus:

$$\frac{\partial \sigma_o^{*,m}}{\partial s_g} < \frac{\partial \sigma_o^{*,f}}{\partial s_g} < 0$$

■

Proof of Prediction 3 Assume that after having a child, women become even more selective about when they supply their labor, while there is no effect on men's selectivity:

$$\alpha_{f,p_1} > \alpha_{f,p_0}$$
$$\alpha_{m,p_1} = \alpha_{m,p_0}$$

where p_1 indicates the worker is a parent and p_0 indicates the worker is not. Prediction 3 states that the positive impact of flexibility on women's relative employment and wages will be amplified among parents. Formally:

$$\frac{\partial \bar{Q}_o^{p_1}}{\partial \sigma_o} > \frac{\partial \bar{Q}_o^{p_0}}{\partial \sigma_o}$$

Notice that

$$\frac{\partial \bar{Q}_o^{p_1}}{\partial \sigma_o} - \frac{\partial \bar{Q}_o^{p_0}}{\partial \sigma_o}$$

$$= \frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_{f,p_1}^{1-\sigma_o} + 1\right] - \frac{1}{\sigma_o - 1} ln\left[\alpha_{m,p_1}^{1-\sigma_o} + 1\right]\right) - \frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_{f,p_0}^{1-\sigma_o} + 1\right] - \frac{1}{\sigma_o - 1} ln\left[\alpha_{m,p_0}^{1-\sigma_o} + 1\right]\right)$$

$$= \frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_{f,p_1}^{1-\sigma_o} + 1\right]\right) - \frac{\partial}{\partial \sigma_o}\left(\frac{1}{\sigma_o - 1} ln\left[\alpha_{f,p_0}^{1-\sigma_o} + 1\right]\right)$$

In the proof for Prediction 1, I showed that $\frac{\partial^2 ln \overline{Q}_{og}}{\partial \alpha_g \partial \sigma_o} \geq 0$. Therefore, because $\alpha_{f,p_1} > \alpha_{f,p_0}$,

$$\frac{\partial \bar{Q}_o^{p_1}}{\partial \sigma_o} > \frac{\partial \bar{Q}_o^{p_0}}{\partial \sigma_o}$$

■

Appendix C Additional Tables

TABLE C.1: CORRELATION BETWEEN SOCIAL PERCEPTIVENESS AND VARIOUS MEASURES OF INFLEXIBILITY

	Coefficient	Robust S.E.
$Inflexibility_o$ (benchmark)	0.842***	(0.068)
$Inflex_o$ weighted by college; top income occs	0.755***	(0.090)
$Inflexibility_o$ weighted by college	0.868***	(0.081)
$Inflex_o$ omitting Time Pressure	0.923***	(0.068)
$Inflex_o$ omitting Contact w/ Others	0.789***	(0.070)
$Inflex_o$ omitting Maintaining Relationships	0.719***	(0.075)
$Inflex_o$ omitting Contact w/ Others & Relationships	0.625***	(0.081)
$Inflex_o$ omitting Unstructure Work	0.823***	(0.066)
$Inflex_o$ omitting Freedom Over Decisions	0.842***	(0.068)
Top $Inflex_o$ and Overwork	1.117***	(0.143)
$AltInflexIndex_o$	0.225**	(0.110)

Notes: Each row reports the results from a separate regression which regresses a social perceptiveness ($Social_o$) on a proxy for inflexibility. The first column reports the coefficient and the second column reports the robust standard errors. In all cases, the number of observations is 320 except the second row, which limits the sample to the top 80 occupations with the highest income as in Figure 1. $Inflex_o$ refers to the O*NET index discussed in Section 2.1.1. and $AltInflexIndex_o$ are the two alternative measures of flexibility discussed in the same section.

TABLE C.2: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - LIMIT SAMPLE TO COLLEGE EDUCATED

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times Inflexibility_o$	0.192		-0.131	0.195		-0.140
	(0.127)		(0.155)	(0.120)		(0.142)
$Female_g \times Social_o$		0.314***	0.393**		0.346***	0.434***
		(0.119)	(0.152)		(0.115)	(0.143)
$Female_g \times Analytical_o$				-0.289***	-0.310***	-0.311***
				(0.096)	(0.098)	(0.096)
$Female_g \times Routine_o$				-0.009	0.079	0.089
				(0.140)	(0.147)	(0.146)
Obs.	1210	1210	1210	1210	1210	1210
Adj. R-sq.	0.855	0.863	0.864	0.863	0.873	0.873
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.014		-0.053***	-0.021**		-0.047***
	(0.011)		(0.013)	(0.010)		(0.012)
$Female_g \times Social_o$		0.015	0.047***		0.004	0.033***
		(0.010)	(0.010)		(0.010)	(0.009)
$Female_g \times Analytical_o$				-0.012	-0.014	-0.014*
				(0.008)	(0.008)	(0.007)
$Female_g \times Routine_o$				-0.034***	-0.030***	-0.027***
				(0.009)	(0.010)	(0.010)
Obs.	1210	1210	1210	1210	1210	1210
Adj. R-sq.	0.963	0.964	0.966	0.968	0.967	0.970
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, Female_g. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social perceptiveness intensity ($Social_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation. In contrast to the main results, this table limits the sample to college-educated cells.

TABLE C.3: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - USING ALTERNATIVE FLEXIBLITY INDEX ($Inflexibility_o$ AND OVERWORK)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, ln($\frac{Emp_{oegp}}{Emp_{egp}}$)					
$Female_g \times$ top $Inflexibility_o$ and Overwork	-0.252		-0.712***	-0.119		-0.524**
	(0.210)		(0.226)	(0.238)		(0.231)
$Female_g \times Social_o$		0.321***	0.508***		0.481***	0.584***
		(0.105)	(0.093)		(0.105)	(0.104)
$Female_g \times Analytical_o$				-0.099	-0.347***	-0.271**
				(0.113)	(0.116)	(0.112)
$Female_g \times Routine_o$				0.192*	0.338***	0.268**
				(0.108)	(0.117)	(0.111)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.789	0.801	0.815	0.793	0.818	0.825
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times$ top $Inflexibility_o$ and Overwork	-0.033**		-0.053***	-0.038***		-0.054***
	(0.014)		(0.015)	(0.014)		(0.014)
$Female_g \times Social_o$		0.008	0.022***		0.013*	0.023***
		(0.007)	(0.006)		(0.007)	(0.005)
$Female_g \times Analytical_o$				0.001	-0.014	-0.006
				(0.007)	(0.009)	(0.007)
$Female_g \times Routine_o$				-0.010	-0.000	-0.007
				(0.009)	(0.008)	(0.008)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.969	0.968	0.970	0.969	0.968	0.970
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, $Female_g$. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's social perceptiveness intensity ($Social_o$) and a dummy indicating that the occupation is in the top tercile of the $Inflexibility_o$ and Overwork measures. See Appendix A for full details. The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE C.4: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - USING ALTERNATIVE FLEXIBLITY INDEX (OVERWORK, COORDINATED HOURS, AND EARNINGS CONVEXITY)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times AltInflexIndex_o$	-0.135		-0.222**	-0.151		-0.205**
	(0.089)		(0.087)	(0.111)		(0.092)
$Female_g \times Social_o$		0.321***	0.402***		0.481***	0.513***
		(0.105)	(0.100)		(0.105)	(0.098)
$Female_g \times Analytical_o$				-0.015	-0.347***	-0.211*
				(0.134)	(0.116)	(0.121)
$Female_g \times Routine_o$				0.216**	0.338***	0.349***
				(0.105)	(0.117)	(0.111)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.791	0.801	0.811	0.796	0.818	0.824
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times AltInflexIndex_o$	-0.027***		-0.031***	-0.033***		-0.035***
	(0.005)		(0.005)	(0.006)		(0.006)
$Female_g \times Social_o$		0.008	0.020***		0.013*	0.018***
		(0.007)	(0.005)		(0.007)	(0.005)
$Female_g \times Analytical_o$				0.016*	-0.014	0.009
				(0.008)	(0.009)	(0.009)
$Female_g \times Routine_o$				-0.003	-0.000	0.002
				(0.006)	(0.008)	(0.006)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.972	0.968	0.973	0.972	0.968	0.973
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, $Female_g$. Observations are a cells at the occupation o × education e × gender g × parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's social perceptiveness intensity ($Social_o$) and an alternative index proxying for inflexibility, which is the average of three standardized components: long hours, coordinated hours, and elasticity of earnings to hours. See Appendix A for full details. The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE C.5: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - USING BROADER SOCIAL SKILL COMPOSITE

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times Inflexibility_o$	0.199*		0.351*	0.300***		0.288
	(0.106)		(0.186)	(0.104)		(0.184)
$Female_g \times SocialComposite_o$		0.084	-0.178		0.241**	0.016
		(0.105)	(0.183)		(0.120)	(0.205)
$Female_g \times Analytical_o$				-0.263**	-0.285**	-0.268**
				(0.106)	(0.122)	(0.122)
$Female_g \times Routine_o$				0.246**	0.277**	0.249**
				(0.112)	(0.122)	(0.123)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.792	0.788	0.793	0.802	0.798	0.801
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.002		-0.001	0.001		-0.000
	(0.007)		(0.010)	(0.007)		(0.010)
$Female_g \times SocialComposite_o$		-0.003	-0.002		0.001	0.001
		(0.007)	(0.010)		(0.008)	(0.012)
$Female_g \times Analytical_o$				-0.008	-0.009	-0.009
				(0.008)	(0.010)	(0.010)
$Female_g \times Routine_o$				-0.003	-0.003	-0.003
				(0.008)	(0.008)	(0.008)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.967	0.967	0.967	0.968	0.968	0.968
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, Female_g. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social intensity ($SocialComposite_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE C.6: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - USING BROADER SOCIAL SKILL COMPOSITE AND ALTERNATIVE FLEXIBLITY INDEX (OVERWORK, COORDINATED HOURS, AND EARNINGS CONVEXITY)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, ln($\frac{Emp_{oegp}}{Emp_{egp}}$)					
$Female_g \times AltInflexIndex_o$	-0.135		-0.230**	-0.151		-0.247**
	(0.089)		(0.110)	(0.111)		(0.115)
$Female_g \times SocialComposite_o$		0.084	0.231*		0.241**	0.353***
		(0.105)	(0.128)		(0.120)	(0.129)
$Female_g \times Analytical_o$				-0.015	-0.285**	-0.177
				(0.134)	(0.122)	(0.129)
$Female_g \times Routine_o$				0.216**	0.277**	0.310***
				(0.105)	(0.122)	(0.117)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.791	0.788	0.796	0.796	0.798	0.806
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times AltInflexIndex_o$	-0.027***		-0.035***	-0.033***		-0.038***
	(0.005)		(0.005)	(0.006)		(0.006)
$Female_g \times SocialComposite_o$		-0.003	0.020***		0.001	0.018***
		(0.007)	(0.006)		(0.008)	(0.007)
$Female_g \times Analytical_o$				0.016*	-0.009	0.008
				(0.008)	(0.010)	(0.009)
$Female_g \times Routine_o$				-0.003	-0.003	0.002
				(0.006)	(0.008)	(0.006)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.972	0.967	0.973	0.972	0.968	0.973
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, Female_g. Observations are a cells at the occupation o × education e × gender g × parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's social intensity ($SocialComposite_o$) and the alternative flexibility index, which is the average of three standardized components: long hours, coordinated hours, and elasticity of earnings to hours. See Appendix A for full details. The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE C.7: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - USING BROADER SOCIAL SKILL COMPOSITE AND ALTERNATIVE FLEXIBLITY INDEX (OVERWORK, COORDINATED HOURS, AND EARNINGS CONVEXITY)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, ln($\frac{Emp_{oegp}}{Emp_{egp}}$)					
$Female_g \times$ top $Inflexibility_o$ and Overwork	-0.252		-0.545*	-0.119		-0.432
	(0.210)		(0.303)	(0.238)		(0.301)
$Female_g \times SocialComposite_o$		0.084	0.254*		0.241**	0.360**
		(0.105)	(0.132)		(0.120)	(0.145)
$Female_g \times Analytical_o$				-0.099	-0.285**	-0.262**
				(0.113)	(0.122)	(0.119)
$Female_g \times Routine_o$				0.192*	0.277**	0.230*
				(0.108)	(0.122)	(0.118)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.789	0.788	0.795	0.793	0.798	0.802
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times$ top $Inflexibility_o$ and Overwork	-0.033**		-0.047**	-0.038***		-0.051***
	(0.014)		(0.021)	(0.014)		(0.019)
$Female_g \times SocialComposite_o$		-0.003	0.012		0.001	0.015
		(0.007)	(0.009)		(0.008)	(0.009)
$Female_g \times Analytical_o$				0.001	-0.009	-0.006
				(0.007)	(0.010)	(0.008)
$Female_g \times Routine_o$				-0.010	-0.003	-0.009
				(0.009)	(0.008)	(0.009)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.969	0.967	0.969	0.969	0.968	0.969
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, $Female_g$. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's social intensity ($SocialComposite_o$) and a dummy indicating that the occupation is in the top tercile of the $Inflexibility_o$ and Overwork measures. See Appendix A for full details. The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation.

TABLE C.8: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - CONTROLLING FOR NON-ROUTINE MANUAL TASKS

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, ln($\frac{Emp_{oegp}}{Emp_{egp}}$)					
$Female_g \times Inflexibility_o$	0.199*		-0.102	0.021		-0.149
	(0.106)		(0.165)	(0.122)		(0.141)
$Female_g \times Social_o$		0.321***	0.393**		0.175	0.271
		(0.105)	(0.163)		(0.141)	(0.172)
$Female_g \times Analytical_o$				-0.542***	-0.574***	-0.555***
				(0.125)	(0.126)	(0.126)
$Female_g \times Routine_o$				0.193**	0.238**	0.246**
				(0.098)	(0.105)	(0.104)
$Female_g \times NonroutineManual_o$				-0.970***	-0.878***	-0.889***
				(0.139)	(0.157)	(0.156)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.792	0.801	0.801	0.855	0.858	0.859
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.002		-0.019**	-0.001		-0.016*
	(0.007)		(0.009)	(0.008)		(0.009)
$Female_g \times Social_o$		0.008	0.022**		0.014	0.024***
		(0.007)	(0.009)		(0.009)	(0.009)
$Female_g \times Analytical_o$				-0.010	-0.013	-0.011
				(0.009)	(0.009)	(0.009)
$Female_g \times Routine_o$				-0.004	0.000	0.001
				(0.008)	(0.008)	(0.008)
$Female_g \times NonroutineManual_o$				-0.006	0.003	0.001
				(0.010)	(0.010)	(0.011)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.967	0.968	0.968	0.968	0.968	0.968
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, $Female_g$. Observations are a cells at the occupation o $\times$ education e $\times$ gender g $\times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social perceptiveness intensity ($Social_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation. In contrast to the main results, I additionally control for non-routine manual task intensity.

TABLE C.9: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - CONTROLLING FOR PHYSICAL STRENGTH INTENSITY

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times Inflexibility_o$	0.199*		-0.102	0.252**		-0.191
	(0.106)		(0.165)	(0.124)		(0.157)
$Female_g \times Social_o$		0.321***	0.393**		0.475***	0.602***
		(0.105)	(0.163)		(0.119)	(0.156)
$Female_g \times Analytical_o$				-0.405***	-0.531***	-0.524***
				(0.142)	(0.138)	(0.139)
$Female_g \times Routine_o$				0.190*	0.278**	0.284**
				(0.114)	(0.121)	(0.119)
$Female_g \times PhysicalStrength_o$				-0.285	-0.326*	-0.363**
				(0.216)	(0.180)	(0.179)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.792	0.801	0.801	0.808	0.826	0.828
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.002		-0.019**	0.003		-0.013
	(0.007)		(0.009)	(0.007)		(0.009)
$Female_g \times Social_o$		0.008	0.022**		0.013*	0.022***
		(0.007)	(0.009)		(0.007)	(0.008)
$Female_g \times Analytical_o$				-0.001	-0.006	-0.006
				(0.009)	(0.009)	(0.009)
$Female_g \times Routine_o$				-0.001	0.002	0.003
				(0.008)	(0.008)	(0.008)
$Female_g \times PhysicalStrength_o$				0.014**	0.014**	0.011*
				(0.007)	(0.006)	(0.006)
Obs.	2440	2440	2440	2440	2440	2440
Adj. R-sq.	0.967	0.968	0.968	0.968	0.969	0.969
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, $Female_g$. Observations are a cells at the occupation o × education e × gender g × parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social perceptiveness intensity ($Social_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation. In contrast to the main results, I additionally control for an index representing skill demands for physical strength.

TABLE C.10: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - LIMIT SAMPLE TO COLLEGE EDUCATED, HIGH INCOME SAMPLE

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Propensity to sort, $\ln(\frac{Emp_{oegp}}{Emp_{egp}})$					
$Female_g \times Inflexibility_o$	0.329*		-0.146	0.339*		-0.215
	(0.183)		(0.300)	(0.186)		(0.304)
$Female_g \times Social_o$		0.511**	0.629		0.616**	0.767*
		(0.249)	(0.389)		(0.250)	(0.386)
$Female_g \times Analytical_o$				-0.095	-0.068	-0.129
				(0.183)	(0.179)	(0.194)
$Female_g \times Routine_o$				0.152	0.256*	0.260*
				(0.136)	(0.146)	(0.143)
Obs.	316	316	316	316	316	316
Adj. R-sq.	0.850	0.866	0.866	0.852	0.874	0.874
Panel B	Avg. log earnings (resid.), $\overline{ln(earnings_{oegp})}$					
$Female_g \times Inflexibility_o$	-0.024		-0.067***	-0.050***		-0.080***
	(0.014)		(0.019)	(0.014)		(0.020)
$Female_g \times Social_o$		0.002	0.056**		-0.014	0.042**
		(0.017)	(0.022)		(0.019)	(0.019)
$Female_g \times Analytical_o$				-0.023	-0.002	-0.025
				(0.015)	(0.016)	(0.015)
$Female_g \times Routine_o$				-0.036***	-0.032***	-0.030***
				(0.011)	(0.012)	(0.010)
Obs.	316	316	316	316	316	316
Adj. R-sq.	0.933	0.930	0.938	0.943	0.936	0.946
γ_{egp}	X	X	X	X	X	X
γ_o	X	X	X	X	X	X
γ_{oep}	X	X	X	X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. This table reports results from regressions of log employment shares and wages on occupation attributes interacted with a dummy for female, Female_g. Observations are a cells at the occupation $o \times$ education $e \times$ gender $g \times$ parental-status p (at least one child vs no children) level. Observations are weighted by number employed in the cell. The outcome in the top panel is the log employment share for a demographic group in that occupation. The outcome in the bottom panel is the average of log earnings (residualized with respect to hours worked and demographics). The first three columns regress the outcome on an occupation's O*NET index proxying for the price that workers pay flexibility ($Inflexibility_o$) and an occupation's social perceptiveness intensity ($Social_o$). The next three columns additionally include controls for occupation's analytical intensity and routine task intensity. All columns include fixed effects for demographic group FE (γ_{egp}) and all interactions of occupation, parental status, and education FE (γ_{oep}). Standard errors are robust and clustered on occupation. In contrast to the main results, this table limits the sample to college-educated cells and the top income occupations (as in Figure 1).

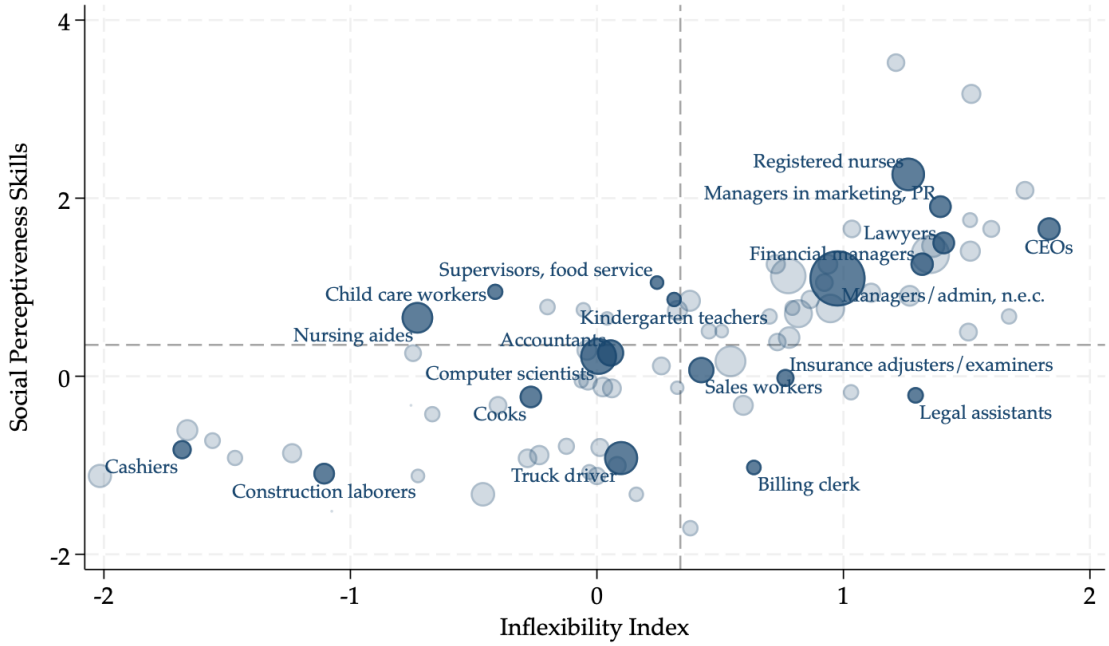
TABLE C.11: RELATIONSHIP BETWEEN OCCUPATION CHARACTERISTICS AND GENDER GAPS IN SORTING AND WAGES - HIGH INCOME, COLLEGE EDUCATED SAMPLE

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	Women's Relative Propensity to Sort					
$Inflexibility_o$	0.504**		-0.175	0.418**		-0.340
	(0.208)		(0.339)	(0.203)		(0.333)
$Social_o$		0.763**	0.899*		0.827***	1.052**
		(0.295)	(0.460)		(0.306)	(0.442)
Obs.	80	80	80	80	80	80
R-squared	0.13	0.28	0.29	0.15	0.34	0.36
Other Task Controls				X	X	X
Panel B	Women's Relative Earnings (Residualized)					
$Inflexibility_o$	-0.035**		-0.085***	-0.056***		-0.093***
	(0.015)		(0.020)	(0.015)		(0.021)
$Social_o$		0.001	0.066***		-0.010	0.052**
		(0.021)	(0.021)		(0.024)	(0.021)
Obs.	80	80	80	80	80	80
R-squared	0.09	0.00	0.22	0.24	0.10	0.32
Other Task Controls				X	X	X

Notes: data are from ACS 2015-2019 and O*NET 2013. As in Figure 1, each observation is a 3-digit occupation in at least the 75th percentile in male earnings, and the sample of workers is limited to the college educated. The outcome in Panel A is the gender wage gap (residualized) and the outcome in Panel B is the log of the relative propensity for women to sort into an occupation (the fraction of women sorting into an occupation divided by the fraction of men sorting into the occupation). Columns 1-3 regress these outcomes on social intensity and the inflexibility index, and columns 4-6 repeat these regressions while additionally controlling for routine task intensity and analytical skill demands. Standard errors are robust and observations are weighted by employment.

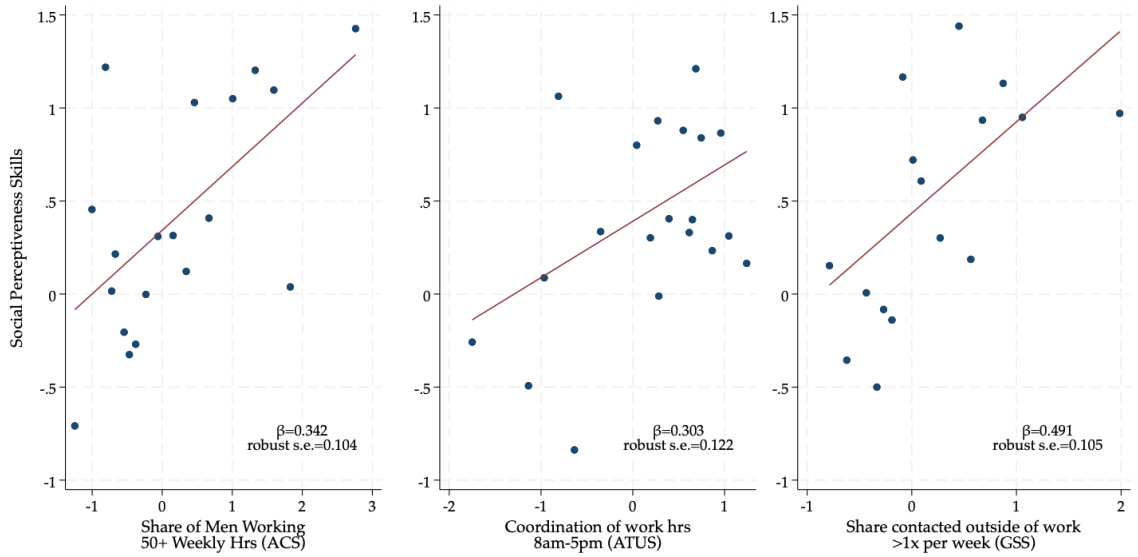
Appendix D Additional Figures

FIGURE D.1: CORRELATION BETWEEN SOCIAL PERCEPTIVENESS AND INFLEXIBILITY INDEX AMONG LARGEST OCCUPATIONS



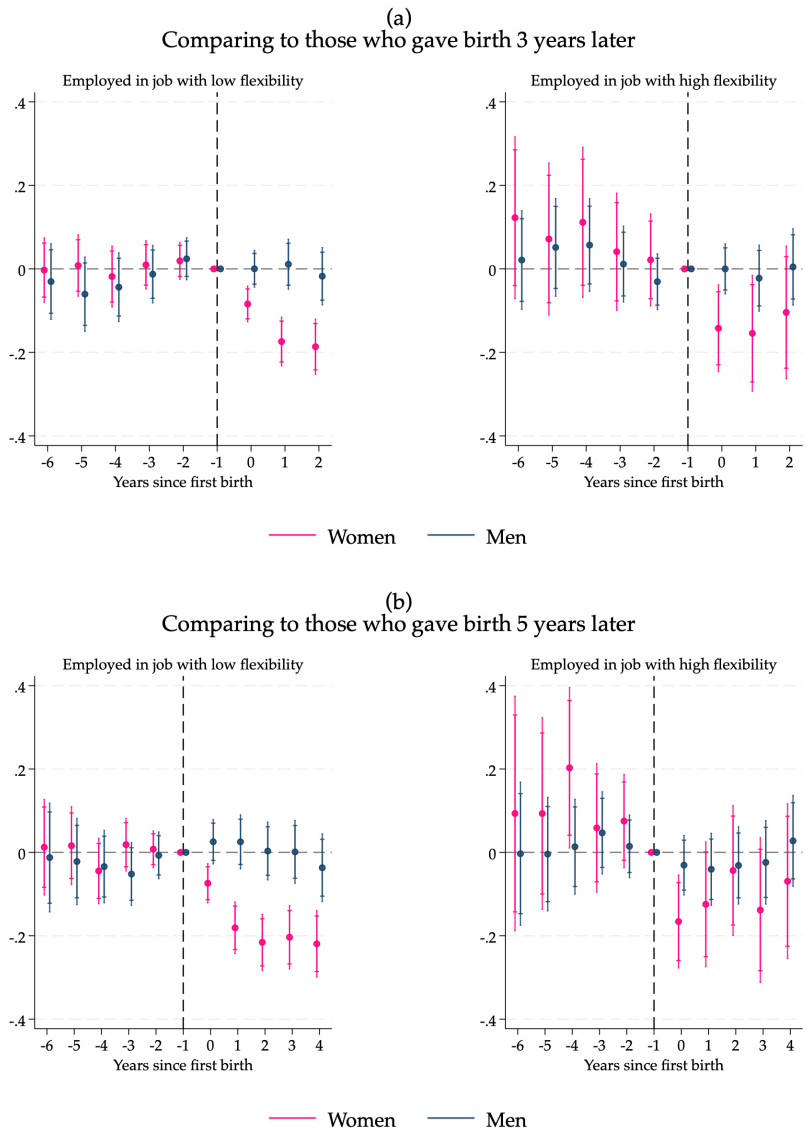
Notes: the scatterplot depicts the relationship between an occupation's social perceptiveness and inflexibility. The plot is restricted to occupations in the top employment quartile. The corresponding β is 0.881 with robust standard errors of 0.100. Illustrative occupations are labeled. Dashed gray lines represent the median in each measure within the limited sample of occupations.

FIGURE D.2: CORRELATION BETWEEN SOCIAL PERCEPTIVENESS AND SPECIFIC DIMENSIONS OF INFLEXIBILITY



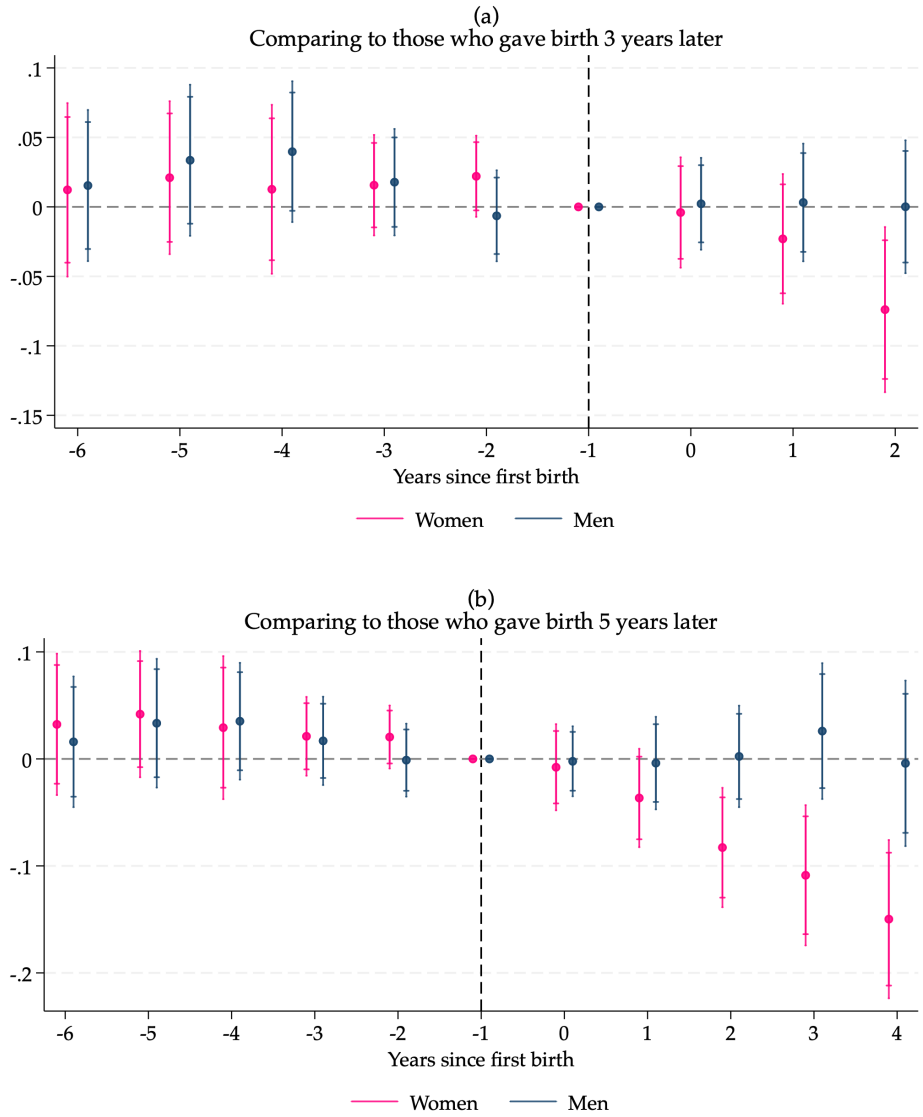
Notes: these binscatters plot an occupation's social perceptiveness against specific dimensions of inflexibility. The first is the share of men in an occupation working over 50 hours per week (measured in the ACS). The second is the concentration of work hours within the 8am-5pm window for a given occupation (as measured in the ATUS). The third is the share of individuals in an occupation who report being contacted outside of work hours (as measured in the General Social Survey in 2014). See data appendix for more details.

FIGURE D.3: CHANGE IN THE PROPENSITY TO BE IN JOBS WITH LOW/HIGH FLEXIBILITY AFTER CHILDBIRTH



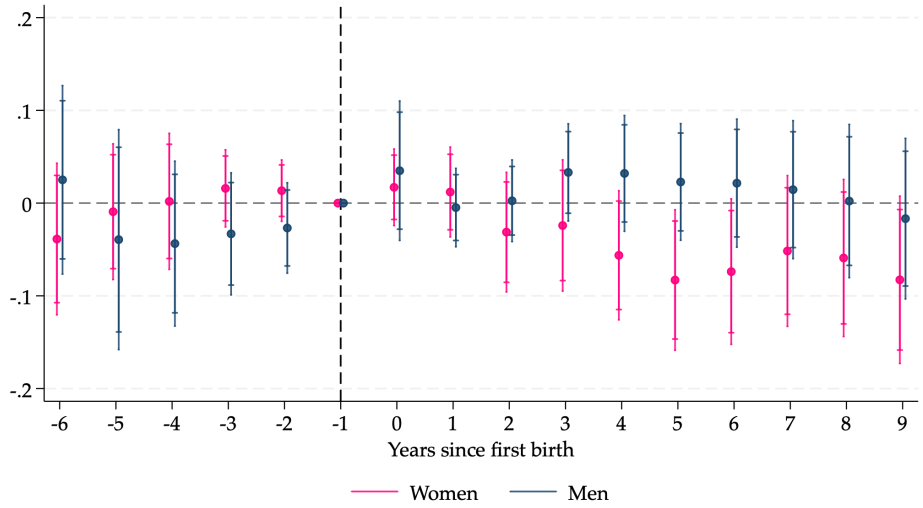
Notes: this figure presents event studies of first child birth separately by gender. In the figures on the left, the outcome is whether the person is employed in a job with low (below median) flexibility (i.e. a high "price" of flexibility according to the index, $Inflexibility_o$). In the figures on the right, the outcome is whether the person is employed in a job with high (above median) flexibility. In both cases, the outcome equals 0 if the person is not employed. The plotted points represent coefficients on dummies indicating the number of years since an individual's first child birth. The estimates in pink represent the results for women, while the estimates in blue represent the results for men. 90% and 95% confidence intervals displayed. Standard errors are robust and clustered on individual. Data are from the NLSY79. As discussed in the text, the "control group" are those who gave birth 3 years later in Panel A and those who gave birth 5 years later in Panel B.

FIGURE D.4: HETEROGENEITY IN WAGE CHANGES AFTER CHILDBIRTH BY FLEXIBILITY OF PRE-BIRTH OCCUPATION



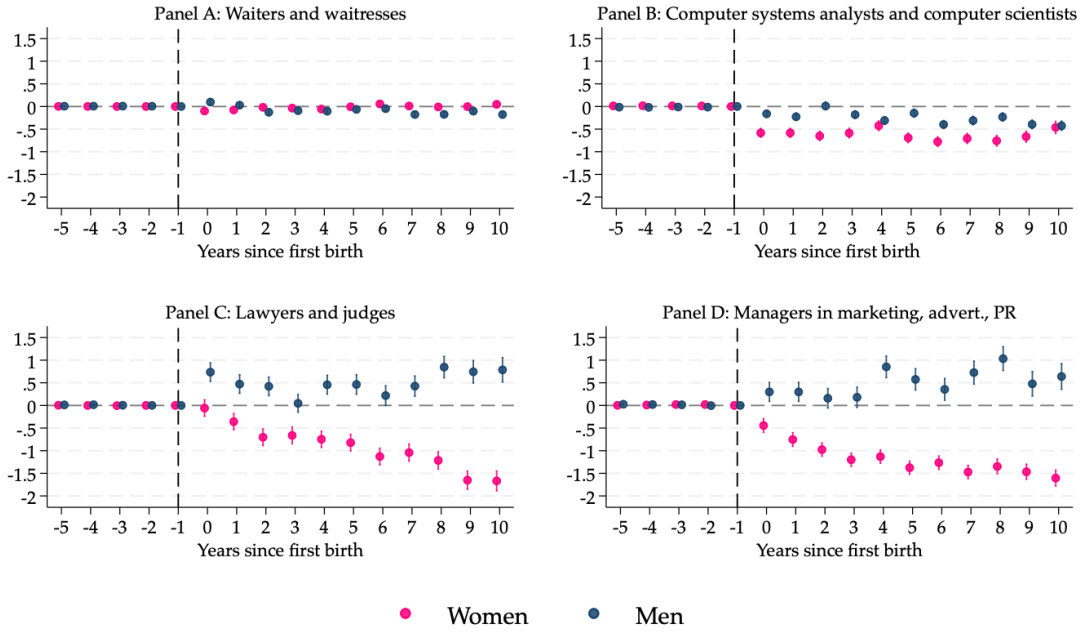
Notes: this figure presents heterogeneity in how wages respond to child birth by the flexibility of the individual's occupation prior to birth. The outcome is an individual's wage (coded as zero if not working). The plotted points represent coefficients on event time dummies interacted with a variable representing the lack of flexibility in the individual's occupation. Specifically, the variable is the average of the index proxying for the "price of flexibility" ($Inflexibility_o$) over all pre-birth observations for the individual. Estimates are divided by pre-birth averages in the outcome and therefore represent percentage changes. The estimates in pink represent the results for women, while the estimates in blue represent the results for men. 90% and 95% confidence intervals displayed. Standard errors are robust and clustered on individual. Data are from the NLSY79. As discussed in the text, the "control group" are those who gave birth 3 years later in Panel A and those who gave birth 5 years later in Panel B.

FIGURE D.5: HETEROGENEITY IN WAGE CHANGES AFTER CHILDBIRTH BY FLEXIBILITY OF PRE-BIRTH OCCUPATION



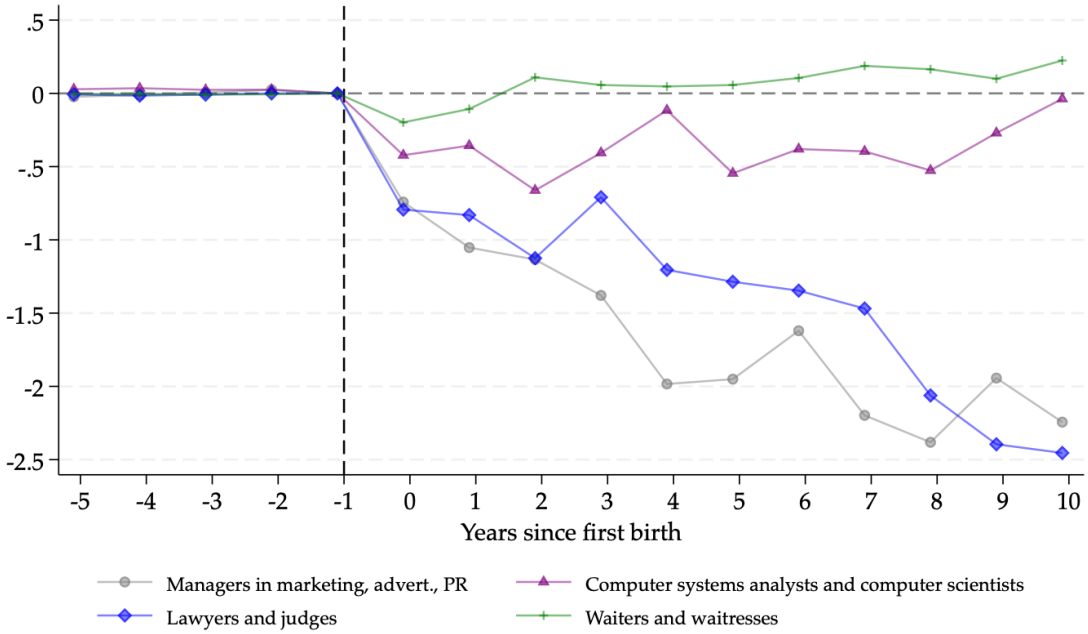
Notes: this figure presents heterogeneity in how wages respond to child birth by the flexibility of the individual's occupation prior to birth. The outcome is an individual's wage (coded as zero if not working). The plotted points represent coefficients on event time dummies interacted with a variable representing the lack of flexibility in the individual's occupation. Specifically, the variable is the average of the index proxying for the "price of flexibility" ($Inflexibility_o$) over all pre-birth observations for the individual. Estimates are divided by pre-birth averages in the outcome and therefore represent percentage changes. The estimates in pink represent the results for women, while the estimates in blue represent the results for men. 90% and 95% confidence intervals displayed. Standard errors are robust and clustered on individual. Data are from the NLSY79. As discussed in the text, the "control group" are individuals who never give birth over the sample period. In contrast to the main results, this controls for a proxy for spousal income interacted with event dummies, age dummies, and year dummies. Due to this variable being missing, there are 33 fewer men and 37 fewer women represented in this sample relative to the main results.

FIGURE D.6: CHILD PENALTY ON SORTING FOR A FEW ILLUSTRATIVE OCCUPATIONS



Notes: this figure presents the pseudo event studies of first child birth separately gender as described in Section 7.2. In Panel A, the outcome is a dummy indicating that the (synthetic) person is employed as a waiter or waitress. The outcome is zero if the person is not employed or is employed in another occupation. Panel B, C, and D present the analogous results for computer scientists, lawyers, and marketing managers. The plotted points represent coefficients on dummies indicating the number of years since an individual's first child birth. The estimates in pink represent the results for women, while the estimates in blue represent the results for men. 95% confidence intervals displayed. Data are from a matched sample in the CPS covering years 1972-2019.

FIGURE D.7: FEMALE-MALE CHILD PENALTY ON SORTING FOR FEW ILLUSTRATIVE OCCUPATIONS



Notes: this figure plots the female-male difference in the coefficients presented in Figure D.6.