Gender norms and labour supply: Identifying heterogeneous patterns across groups of women^{*}

Estefanía Galván[†] Cecilia García-Peñalosa[‡]

Preliminary version: September 2018

Abstract

The reduction in gender gaps observed in the second half of the 20th century seems to have halted and social norms are argued to be a major cause behind the remaining differences in the labour market performance of men and women. This paper seeks to understand heterogeneity in labor market responses to social norms. We focus on two norms: the fact that males should be a household's main breadwinner, and the expectation that women will do most of the childcare. Our results indicate that distinct gender roles have different effects. Using panel data for US couples and an event study approach we estimate the impact of the birth of the first child on labour market trajectories, and we find that motherhood generates very strong responses, in terms of the intensive and the extensive margins of labour supply. Although women who were the main earners have a somewhat smaller reduction in employment than secondary earners, both groups largely conform to the traditional division of roles, pointing towards a misallocation of resources. In contrast, the effect of the male breadwinner norm is highly contextualized. Women of all education levels are affected by it, yet the margin along which they react differs. The results also suggest a hierarchy of norms, as gender roles seem to be less important in households without offspring.

JEL Classification: D10, J16, J22

Key words: gender identity norms, female labour supply, children, relative income

^{*}We are grateful to seminar participants at AMSE.

[†]Aix-Marseille University, CNRS, EHESS, Centrale Marseille, AMSE. Email: estefania.galvan-galan@univ-amu.fr

[‡]Aix-Marseille University, CNRS, EHESS, Centrale Marseille, AMSE and CESifo. Email: cecilia.garciapenalosa@univ-amu.fr

1 Introduction

Gender gaps in labour force participation, hours worked and earnings have declined dramatically over the past half century. This convergence has been largely driven by two factors: the reduction in the pay gap for equivalent jobs and an increase in female education and labour market experience (Blau and Kahn, 2017; Goldin, 2014; Goldin et al., 2006). Nevertheless, substantial gender gaps still remain in pay and employment levels, and the process of convergence seems to have slowed down in the past two decades. For instance, Goldin (2014) finds evidence of a plateau in participation for US women starting in the 1990s. Traditional factors, notably education and experience, seem to have little explanatory power for the remaining gender gaps and recent work has started to focus on the importance of gender identity roles and beliefs. This paper contributes to the literature on the role of social norms by examining how women's or household circumstances affect the intensity of the effect of social norms on labour market outcomes.

Social norms concerning gender roles push women to adhere to certain gender-specific behavioural patterns because not following the norm could be inherently costly for them (Akerlof and Kranton, 2000). Traditionally, social expectations imposed that mothers of children in preschool age should not work and that a man should be the household's main breadwinner, thus creating gendered norms that put pressure on women to conform and behave in a way that affected their labour market outcomes. Such norms are still very relevant, even in rich countries like the US. For example, the World Value Survey (WVS) asks individuals whether they agree with the statement "If a woman earns more money than her husband, it's almost certain to cause problems" (WVS 1994-1998). In the US, 38% of respondents agree or strongly agree, while more recent surveys indicate that 75.5% agree with the statement "being a housewife is as fulfilling as working for pay".¹

These attitudes may, however, differ depending on individual and household circumstances. Figure 1 depicts the differences in the percentages of agreement to gender norms according to educational levels and parental status for the US in 2011. The figures are striking because of the strong prevalence of norms, and in particular the degree of agreement with the statement that "when a mother works for pay the children suffer" suggests that even if the idea of women working in the labour market is well accepted, this may not be so for mothers. The figure also indicates that agreement with gender norms is more likely the less educated an individual is as

¹This is not an exclusively US phenomenon. The agreement rates are 44% in Spain and 33% in Sweden for the first statement, and 49.5% and 42.1%, respectively, for the second.

well as among those with at least one child, as compared to those without children.

This paper seeks to understand heterogeneity in responses to social norms. In their seminal work, Bertrand, Kamenica and Pan (2015) find that in couples in which the wife's "potential earnings" are likely to exceed her husband's, women reduce their labour supply or remain in lower-pay jobs so as to avoid a gender role reversal in earnings, and they conclude that conforming to gender identity roles partly explains gender gaps. We extend their analysis by examining the interaction between different identity norms and the possibility of heterogeneous responses in order to explore the underlying mechanisms through which gender specific behavioural prescriptions associated with male and female traditional roles affect women's labour market outcomes.

We argue that the cost of departing from a norm is likely to depend on individual and household characteristics. Two aspects seem to be particularly relevant. The first is educational attainment. Individuals can get a sense of identity from the job they perform and, at least in the US, this effect seems to be stronger for more educated individuals (Gallup News, 2014). Consequently, more educated women may find themselves in conflict between the identity derived from their career and that associated to conforming to gender identity norms. If dropping out of the labour force has a greater identity cost for them than for those who are less educated and, potentially, less likely to see their job as part of their identity, then we may find that these two groups respond differently to gender norms. The second aspect we explore is the household environment as defined by the presence of children. Social norms on child-rearing roles imply that women should bear the burden of childcare, and hence it is possible that norms about the household division of labour interact with those concerning a couple's relative earnings.

There are two ways in which norms could interact. On the one hand, certain women may have a higher intrinsic cost of deviating from norms than others and hence would tend to conform both in terms of earnings and concerning childcare, and vice versa. On the other hand, the presence of children could act as a trigger for other norms. Women may find it easier to deviate from gender roles in the absence of offspring, while once they have a child they choose to conform either because they (consciously) want to adhere to gender roles in front of their children or because motherhood triggers an unconscious reaction that makes gender roles more salient.

We use panel data for US couples to examine these questions. We start by following Bertrand et al. (2015) and consider the dynamics of women's labour force participation and hours of work. Our key explanatory variable is whether a woman earns more than her husband in year t, and we find that if she does, at t+1 she is less likely to be in the labour force and, conditional on working, works fewer hours than those who do not, in line with earlier work. We then split the sample both across education groups and on the basis of whether the couple has at least one child or not. We next take advantage of the panel nature of the data to further analyse the relationship between the male breadwinner norm and child-rearing roles. Using an event study approach we estimate the impact of the birth of the first child on labour market trajectories for men and women. Although maternity decisions are endogenous, the event of having a child generates sharp changes in the trajectory of the results in the labour market (employment, hours worked, income), just after birth. Therefore, by comparing the results in the environment, just before and after birth it is possible to estimate the effects of having a child. We explore heterogeneity in the responses for women of different educational levels and based on relative earnings within the couple, that is, assessing whether there are differences in the responses of women depending on whether or not they were the main breadwinners in their households before having their first child.

We find important differences across groups in the response to the male breadwinner norm. Women with a college degree reduce their hours but not their probability of participating, while the opposite occurs for those in the middle of the skill distribution. In contrast, women with a high-school degree or less, react in terms of both the intensive and extensive margin. Having at least one child also makes a difference, as for women without children we find small and insignificant coefficients on our variable of interest. These estimates indicate a strong contextualisation of norms as women's willingness to depart from the norm is strongly affected by education and the presence of children.

The event study results show that women react strongly both in terms of the intensive and the extensive margins when they have children, as has been established before for other countries (Kleven et al., 2018). The arrival of the first child has little impact on the father's employment, market hours, and domestic work hours but reduces dramatically market labour supply for women. Contrary to what an efficient allocation of household time would imply, we find sharp changes both for women who were secondary earners and those who were the main breadwinners. Interestingly, the reduction in the probability of employment is initially lower for the latter group, but the estimated effect on employment keeps falling over time so that 10 years after the birth of the first child both groups of women have an equal reduction in the probability of being employed relative to the pre-birth level.

Together with our earlier results, these estimates indicate that distinct gender roles have different effects. Motherhood generates very strong responses along traditional lines, and although women who were the main earner have a somewhat smaller reduction in employment than secondary earners, both groups largely conform to the traditional division of roles. In contrast, the effect of the male breadwinner norm is highly contextualised. Women of all education levels are affected by it, yet the margin along which they react differs. This is important because the long-term consequences (not explored in this paper) can be different depending on whether a woman reduces her hours or drops out of the labour force. The results obtained when we consider childless women also indicate a hierarchy of norms, as gender roles seem to be less important in households without offspring.

This paper contributes to three different strands of literature. First, it adds to the recent work on how social norms can help us understand the gender gaps in the labour market (a comprehensive review of this literature is presented in Bertrand (2011)). Some of this literature has focused on the intergenerational transmission of norms within the family or in society as a whole, providing micro-foundations for differential sorting of men and women across occupations and women's decisions to participate in the workforce.² One of the remaining challenges is to understand why women "choose" to work fewer hours or to be in less-well-paid firms, sectors and occupations, and behavioural responses to gender norms within couples appears as a potential explanation. Bertrand et al. (2015) identify such effects, and their results indicate that the breadwinner norm can partly explain women's labour market choices.³ Their insight has been confirmed for German data by Wieber and Holst (2015), who argue that context can affect the strength of responses to the norm. Comparing West and East Germany, they find that the male breadwinner norm significantly affects labour supply decisions in West Germany but not in the East where the socialist regime had implemented strong policies to erode gender stereotypes. Lippmann et al. (2016) also use Germany's division as a natural experiment and find that women who earn more than their husbands "compensate" it by increasing the number of housework hours in West Germany but not in the East. Our results complement these findings, examining other dimensions along which a contextualisation of norms is important.

Secondly, this paper is related to recent work on the so-called "child penalty", defined as

²This literature has examined, for example, women's under-representation in mathematics and in other scientific disciplines (Nollenberger et al., 2016), the persistence of occupational gender segregation (Breen and García-Peñalosa, 2002), and labour force participation and hours of work (Fortin, 2005; Bertrand et al., 2015; Fernández et al., 2004; Farré and Vella, 2013; Olivetti et al., 2016). According to Blau and Kahn (2017), it has the potential to help explain not only the unexplained gender gap (i.e. the fraction not accounted for gender differences in measured qualifications), but also gender differences in some of the measured factors themselves

³Some authors have criticised parts of the analysis in Bertrand et al. (2015). Notably, Hederos Eriksson and Stenberg (2015) reproduce, using Swedish data, the analysis of the distribution of the wife's earnings share. They find that, as is the case for the US, the data shows a sharp discontinuity to the right of 50% but this discontinuity is mainly explained by bunching of couples exactly at 50%.

the extent to which women fall behind men (in employment rates, hours or earnings) due to having children. Adda et al. (2017) use a dynamic life-cycle model to estimate the child penalty on German data and find that about three quarters of the earnings cost are due to intermittent or reduced labour supply. Alternatively, an event study approach around the birth date of the first child has been used to estimate various outcomes, such as how hours worked change after the birth of the first child, the dynamics of within-couple gaps in earnings, and changes in within-firm remuneration and bonuses; see Paull (2008), Angelov et al. (2016) and Lucifora et al. (2017).⁴ Lundborg et al. (2017) analyse use this approach to identify a causal effect by focusing on differences in actual fertility amongst women who had in vitro fertilization treatment, and find negative, large, and long-lasting effects on earnings of having children. We follow the event approach proposed by Kleven et al. (2018), who find a long-run penalty on earnings in Denmark, and show that since total gender inequality in earnings has fallen, most of the remaining gender disparities in earnings are due to child penalties. Our focus is on how the pre-birth circumstances, notably the relative earnings of the two members of a couple, affect the intensity of the effects of parenthood on labour market trajectories.

Lastly, our research is also related to a growing literature that has focused on the relevance of occupations and firm characteristics to analyse the dynamics on men and women careers (Goldin, 2014; Goldin and Katz, 2016).⁵ Our study is close to Bertrand et al. (2010) who study the careers dynamics in the financial and corporate sectors, which are characterized by long work schedules. They find large and rising gender gaps in earnings, explained by greater career discontinuity and shorter work hours for female MBAs. We contribute to this literature by examining how motherhood results in women shifting towards occupations and industries where females are overrepresented and have lower yearly hours. Our results suggest that the differential sorting of men and women across occupations and industries is, at least partly, due to gender norms related to child rearing.

The paper is organized as follows. Section 2 presents our framework of analysis to interpret the mechanisms through which gender identity norms affect female labour supply. Section 3 describes the data and sample. Section 4 focuses on the analysis of heterogeneous responses

⁴The advantage of the event study approach is that it has the potential to capture the global treatment effect of children, as opposed to only local treatment effect of a second or third child, obtained from the twin or sibling sex mix instruments proposed by previous literature on the child penalty such as Angrist and Evans (1998).

⁵Blau and Kahn (2017) show that while the share of the gender wage gap due to human capital (education and experience) has declined noticeably, the share accounted for by locational factors like occupation and industry actually increased from 27% of the 1980 gap to 49% of the (much smaller) 2010 gap. The continued importance of occupation and industry in accounting for the gender gap and the rise in the relative importance of these factors has led to research arguing that the way in which jobs are structured and remunerated explains part of the gender gaps.

to the breadwinner norm, while Section 5 investigates the responses to child rearing norms, by estimating the effect of the first child born on labour market outcomes. The last section concludes.

2 Gender norms and labour supply

According to standard economic theories of the household (Becker, 1973), an "efficient" allocation of household time would imply a division of market and home production determined by the relative productivities of each member of the couple. Hence, men in two-earner families typically work more in the labour market and do less housework than their wives because they earn more per hour. That is, gender affects the labour supply and housework indirectly, through its effect on relative earnings.

A second view, initially developed by sociologists, emphasizes that gender has pervasive effects at many levels, structuring identities, norms, interaction, and institutions (Bittman et al., 2003). These views are not necessarily exclusive. In the second one, gender not only affects labour market decisions through earnings, but also affects it in other ways, through gendered expectations. West and Zimmerman (1987) pioneered the notion of "doing gender": that individuals' behaviour is affected by the expectations held by others. In this alternative view of social norms, actors do not necessarily internalize gendered identities or norms they really believe are morally right or preferable, but rather they internalize only expectations that others will follow norms along with a need to present themselves as cognitively "making sense" in terms of these norms. This constrains individuals to act in normative ways. Akerlof and Kranton (2000) import these ideas into economics, defining identity as a sense of belonging to a social category (i.e. man and woman), together with a view on how people in that category should behave. They argue that deviating from the prescribed behaviour is inherently costly, leading identity to influence economic outcomes at many levels.

In this paper we focus on two norms typically associated with male and female traditional roles at household and the labour market: the fact that males should be a household's main breadwinner, and the expectation that women will do most of the childcare. We claim that while these gender norms might be relevant for all women, their way to respond to them may differ both because of preferences and constraints. First, responses to gender norms in relation with relative earnings within the couple may depend on the educational attainment of women. Women with more years of education may be more likely to see their careers as part of their identity and be less willing than other women to give it up. In fact, Goldin (2006) describes the

period that began in the late-1970s (when the birth cohorts of the late 1940s were in their early thirties) as a "quiet revolution" characterized by a change in female identity. Women started to work not only because they and their families "need the money" but also because occupation and employment define one's fundamental identity and societal worth. As a result, it is possible that women with more years of education are less likely to drop out of the labour force than less educated ones and instead choose to adjust their income by moving to less demanding firms or positions or simply working fewer hours so as to avoid violating the male breadwinner norm.

Education may also have an effect on the mechanism of adjustment chosen by women because different categories of workers may be more or less constrained in their capacity to adjust hours. Our hypothesis is that adjusting hours of work is hardest for those with middle qualifications. Those at the lower end of the skill distribution are often paid by the hour and hence have flexibility about how many hours to work, while women with at least a college degree may have enough bargaining power (because they are hard to replace, have firm-specific skills, etc.) that allows them to negotiate working time. In contrast, those in the middle of the skill distribution perform tasks that are not paid by the hour, yet do not have enough bargaining power to modify their hours relative to the standard contract. As a result, they may only be able to choose whether or not to work. Lastly, there may be fixed costs associated to being employed (transport, organizing childcare, etc.) that those with high hourly wages will be able to pay but which will make part-time employment unattractive to those lower wages.

Our second hypothesis concerns the male breadwinner norm and parenthood. The cost of deviating from social norms can change with the environment of the individual, and evidence exists indicating that the institutional environment matters.⁶ This cost can be affected by the presence of children in the household, and their impact on the reaction to norms can be conscious or unconscious. The presence of children may trigger a conscious decision to conform to the norm, possibly because parents want to transmit the socially-accepted norm to their offspring, thus making women who were previously unaffected by deviating from the norm choose to conform. Alternatively, the reaction can be unconscious. Certain experiences may make norms more salient, and motherhood could have such an effect. Historically, motherhood has been the key aspect defining female identity, hence it is possible that couples pay little attention to whether they conform to traditional gender roles as long as they do not have children. The arrival of offspring could be a trigger that makes gender norms salient and hence pushes women who are earning more than their husbands to modify their behaviour. We will thus explore

⁶See Fernández (2007)

whether the effect of the male breadwinner norm is stronger for women with than for women without children.

Lastly, we analyse the effect of the child rearing norm and the interaction with relative earnings within the couple. The arrival of the first child demands a large amount of time in caring activities, which implies an increase in the total housework performed by household's members. Standard theories of the household state that an "efficient" allocation of household time would imply that the member of the couple with higher labour earnings remains in the labour market and the other takes the main child-carer role. However, child rearing is affected by social norms which assign this role to mothers. Then, the arrival of the first child comes together with the behavioural prescriptions that the mother should stay at home taking care of the children, while this is not the case for fathers. Our hypothesis is that while female employment and developing a professional career is now acceptable, norms related with children are still strongly gendered. Then, while some women are less constrained in the labour market by the male breadwinner norm, behavioural prescriptions related with female role at child rearing would have a stronger effect on all women's labour market trajectories and increasing housework. Yet, the mechanism of adjusting to this norm and the responses in the labour market can be heterogeneous depending on women pre-birth situation. In particular, the responses can differ by educational levels and relative earnings within the couple. We would expect that those women who were earning more than their partners before child arrival are less likely to leave the labour market when they become mothers.

Our empirical strategy proceeds in two steps. After describing the data, the first part of our analysis follows the approach proposed by Bertrand et al. (2015) and considers the effect of the breadwinner norm on women's labour force participation and hours of work. We split the sample and analyse heterogeneous responses by educational levels and parental status. The second part of our analysis focuses on child rearing norms. Using an event study approach, we estimate the impact of the birth of the first child on labour market trajectories, assessing whether there are differences in the responses of women and men depending on their education level and whether or not they were the main breadwinners before having their first child. We also examine whether these effects are related to transitions to feminized industries and occupations.

3 Data and sample description

Our primary data source is the Panel Study of Income Dynamics (PSID) for the period 1968-2015. The PSID is a longitudinal study of US households that began in 1968 with a nationally representative sample of households, and contains information collected on sampled families and their descendants for nearly 50 years. All persons living in PSID families in 1968 were interviewed yearly through 1997 and every other year since then. The PSID also follows those born into or adopted by a PSID family even after they moved out of the original household. Those who married into PSID families were followed for as long as they lived with a member of the PSID sample.⁷

More specifically, the data that we use comes from the Family files and the Cross-year Individual files. The Cross-year Individual files contain one record for each person ever in a PSID family from the beginning of the study through 2015. The Family files contain family-level information and detailed information for the head and wife, including married and cohabitant couples. The main advantages of using these data are that they allow us to follow couples for a long time and that it contains detailed information on labour market outcomes for household heads and spouses, including labour supply, income and industry and occupation variables, as well as information on hours of housework. This makes it particularly suitable for analysing outcomes for the two members of a couple.

Our estimation sample is composed of 129,265 couple-year observations for women who live in couples, where both, husband and wife are between 18 and 65 years old, and where at least one of the spouses has positive earnings (in the previous year).⁸ The sample is also restricted to family heads and spouses/cohabitors because the PSID only supplies the crucial work history information for these individuals. Information on birth dates is taken from the *Birth and Adoption History File*. The event study analysis consequently includes only those couples who have a register in this data.

Table 1 presents summary statistics for the main variables in our sample. Education is defined as the maximum education level reached by the individual, and comprises three categories: highschool or less, 1-3 years of college and at least 4 years of college (college graduates). All the income related variables correspond to the previous year, and the values of wife's and husband's income in the table are expressed in 2015 dollars. On average, the men in our couple-sample are two years older than the women. The mean of the relative income is 0.29, which means that women earn on average around 29% of the income of the couple. The wife earns more than the

 $^{^{7}}$ The Panel Study of Income Dynamics, public use dataset is produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI (2017). PSID data are publicly available at ht-tps://psidonline.isr.umich.edu/

⁸Our sample includes only individuals belonging to the "core sample". The Latino sample, the immigrant refresher sample, and the low-income over-sample are excluded in order to prevent changes in these samples from affecting the results. The disadvantage of this is that the composition of the core sample is representative of US population in 1968, not of the US population today.

husband in 18 percent of the cases. Also, 85 percent of women have at least one child, and the average number of children is 2.53.

4 The breadwinner norm and labour market outcomes

4.1 Methodological approach

The first part of our analysis considers the effect of the breadwinner norm on women's labour force participation and hours of work. We follow the approach proposed by Bertrand et al. (2015), where the actual realization of the wife's earning more than her husband in the previous period is used as a predictive. The panel nature of the data allows us to include couple fixed effects, and hence investigate whether realizations of earnings that imply that the male breadwinner norm is violated in year t result in a change in the labour supply of women at t+1.

Specifically, we estimate the following lineal probability model:

$$y_{it} = \alpha W e m_{i,t-1} + \beta X_{it} + \mu_i + \gamma_t + \delta_s + \varepsilon_{it}$$
(1)

where y_{it} is the outcome for a woman *i* in couple at time *t*, which is either the wife's labour force participation or the logarithm of the number of hours she worked in year *t*. $Wem_{i,t-1}$ is a dummy that takes value one if the wife earned more than the husband in t - 1, i.e. $Wem_{i,t-1} = 1$ if $RelativeIncome_{t-1} > 0.5$, where $RelativeIncome_{t-1} \equiv WifeIncome_{t-1}$ $/(WifeIncome_{t-1} + HusbandIncome_{t-1})$. The parameter of interest, α , indicates the predicted changes in the likelihood that the wife participates in the labour force and the hours of work when $Wem_{i,t-1}$ changes by one unit, holding the other variables fixed.

The vector X_{it} is a set of controls that include the logarithms of the labour income of the husband $(\ln Husband Income_{t-1})$ and the wife $(\ln WifeIncome_{t-1})$ and couple's income $(lnCoupleIncome_{t-1})$, and a quadratic in both the wife's and the husband's age. As well as individual fixed effect (μ_i) , the regressions include year (γ_t) and state (δ_s) fixed effects. In the regression including all the women in our sample the logarithm of one plus yearly income is used in order to include zeros in addition to an indicator for whether only the wife is working and an indicator for whether only the husband is working. Following Bertrand et al. (2015) we estimate the regression through a linear probability model, pooling all year observations together. The main question we want to address is how the parameter α varies across subgroups of the population for which social norms may be more or less stringent.

4.2 The breadwinner norm and educational attainment

We start by examining the differential responses of women depending on their level of education. Table 2 considers married women's labour force participation as a function of the dummy variable $Wem_{i,t-1}$, which takes value one if the wife earned more than the husband in t-1. The first panel in the table replicates the analysis in Bertrand et al. (2015), with $Wem_{i,t-1}$ having a significant and negative coefficient. The four columns present different specifications as additional controls are added. Column (1) has only couple fixed effects, column (2) adds polynomials for the wife's and the husband's income as well as their relative income, and column (3) adds children controls, which include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, or older than 6. The last column reports results with all controls but removing the couple fixed effects.

The effect is statistically and economically significant. For the entire female population, if a wife earns more than her husband, she is 2 percentage points less likely to be in the labour force the following year. The next three panels consider the impact of $Wem_{i,t-1}$ on participation separately for the three educational groups. The coefficient is negative, large and significant for women with a high-school degree or less as well as for those with "some college", but it is non-significant and of a much smaller magnitude for those with at least 4 years of college, indicating a differential response across educational groups. Women with some college have a reduction of similar magnitude to that of the entire population, 1.8 percentage points, while the effect is substantially stronger for less educated women, amounting to a negative effect of 3.2 percentage points.

Table 3 performs the same analysis looking at hours of work for a subsample of women with positive earnings (conditional on being employed). As before, the various columns sequentially add controls. The top panel indicates that, when we consider all women, earning more than her husband leads women to subsequently reduce their hours of work. The sample is then divided between the three educational categories. Both women at the bottom and those at the top of the skill distribution exhibit a negative and significant coefficient. Interestingly, the magnitude of the effect is similar for the two groups. Holding other variables fixed, if a wife earned more than her husband in t-1, conditionally of remaining employed, she works 2.6% less hours in the next period if she is low educated and 2.9% less hours if she has at least some college. In terms of hours, considering that the average of hours worked during this period for the less educated women is 1475 hours per year and 1600 hours for the most educated, it would imply that, if a wife earned more than her husband, in the next period she works 38 hours less on average if she

is low educated and 46 hours less for those with college. In contrast, for the middle skill group (*some college*), once we include controls the effect becomes statistically insignificant and small in magnitude, although it remains negative.

Our results indicate that the level of education is an important determinant of how women respond to the male breadwinner norm. Less educated women adjust both the intensive and the extensive margin of labour supply, but other educational categories seem to react by changing only one of these margins. For those with some college, gender norms imply mainly an adjustment in terms of participation, while those with at least 4 years of college respond by reducing their hours of work but do not leave the labour market.

A possible explanation for these patterns is a combination of identity traits and labour market constraints. Women at the top of the education distribution are more likely to form their identity around their professional life; see Goldin (2006). As a result, dropping out of the labour force will have a larger cost than for other women and hence they choose to try to conform to the male breadwinner norm by adjusting hours. Moreover, these women are often employed in occupations where adjusting hours is not very difficult, such as school teachers or liberal health professions.⁹

Women in the middle of the skill distribution respond by adjusting participation but not hours worked. There are two possible (not mutually exclusive) explanations. On the one hand, the nature of the jobs these women perform may give them little freedom to choose their hours. In our data, 39.5% of women in this category are employed in jobs classified as "Office and Administrative support", and are also over-represented in the category "Health technologist and technicians". These are to a large extent jobs in which the individual is providing support for more skilled workers and hence their presence is "required" for the latter to be able to perform their job. As a result, it may be difficult for these women to negotiate shorter hours. On the other hand, being employed implies fixed costs both in terms of transport but also concerning childcare. With moderate hourly wages, working few hours may imply that the resulting total earnings are not high enough to make it worth it to work.

Lastly, women with a high-school degree or less adjust both their participation and (conditional on remaining employed) their hours. If these women have a weaker labour market attachment, dropping out of the labour force may not involve a cost in terms of identity. They are employed in both low-flexibility occupations, such as "Office and Administrative support"

⁹The single most important category is "Education, teacher (except university)", were 23.4% of the women with college degrees are employed. Women with a college degree account of a large share of female employment in this category (85%), as well as in "Healthcare practitioners" (75.8%), and "Community and social services" (72.5%).

where 29.9% of them are employed, and in low-skill service occupations with little training and by-the-hour payment where adjusting hours is easy (such as "Building and grounds cleaning and maintenance" and "Maids and housekeepers, cleaners"). As a result, some of them may be able to adjust hours and others not, implying that we observe both ways of adjusting their earnings.

4.3 Gender norms in the presence of children

We consider next the household's context. As we have argued, the presence of children may affect the way in which couples view gender norms, hence we divide our sample into women that have had no children and women with at least one child.

Table 4 presents the results for labour force participation, with the results for childless women being on the left panel and those for women with at least one child on the right panel. The coefficients for the latter are negative, significant and of similar magnitude to those found earlier on. In contrast, the coefficient for childless women is negative, lower, and, once we include the additional controls, not statistically significant. The sample of childless women is considerably smaller than for women with children, raising the question of whether sample size is the main reason behind the lack of statistical power. But even if the absence of significance is driven by sample size, the estimated coefficients for the sample of childless women are about half those obtained for women with children. Our data hence indicates that there is at best inconclusive evidence of an effect on the labour supply of women without children.

Table 4 runs the same regressions for yearly hours of work. The effect of having earned more is not significant for those who never had a child, while for mothers it is negative and significant in all specifications except on that with all the additional controls (possibly due to a less precise estimate). We can see that despite it being not significant, the estimated coefficient, -0.11 is quite similar to the result obtained in the regression for all the women (-0.13). This evidence is suggestive that the male breadwinner norm might be more relevant for women who had at least one child.

5 Gender norms and the child penalty

In the second part of our analysis we investigate the relationship between child-rearing roles and relative earnings within the couple. We focus on the labour market adjustments that women make following the arrival of their first child. We are interested in understanding in which way the pre-motherhood situation of men and women affects their adjustments.

5.1 Event study approach

We follow the methodology proposed by Kleven et al. (2018). For each individual in the sample we denote by z = 0 the year in which the individual has his/her first child and index all years relative to that year. Event time z runs from -5 to +10. We estimate the following regression:

$$y_{itz}^{g,Wem} = \sum_{j \neq -1} \alpha_j^g. I[j = z] + \sum_k \beta_k^g. I[k = age_{it}] + \sum_l \gamma_l^g. I[l = t] + v_{itz}^g$$
(2)

where y_{itz}^g is the outcome of interest for individual *i* of gender *g* in year *t* and at event time *z*. The regression includes a full set of event time dummies (first term on the right-hand side), age dummies (second term) and year dummies (third term). The event dummy goes from -5 to 10 omitting the event time dummy at z = -1, implying that the event time coefficients measure the impact of children relative to the year just before the first child is born. By including a full set of year dummy controls non-parametrically for underlying life-cycle trends, and including a full set of year dummy controls non-parametrically for time trends such as wage inflation and business cycles. The inclusion of age dummies improves the comparison between men and women as women are, on average, a couple of years younger than men when having their first child, and between women of different educational levels, as the less educated tend to have children at a younger age.

The estimated level effects are converted into percentages by calculating:

$$P_z^g = \hat{\alpha}_z^g / E[y_{itz}^{\tilde{g}}|z]$$

where $y_{itz}^{\tilde{g}}$ is the predicted outcome when omitting the contribution of the event dummies, i.e. $y_{itz}^{\tilde{g}} = \sum_k \beta_k^g I[k = age_{it}] + \sum_l \gamma_l^g I[l = t]$. Hence, P_z^g captures the year-z effect of children as a percentage of the counterfactual outcome, i.e that in the absence of children.

The regression is first estimated for men and for women, and then we analyse differences by educational levels and by considering separately women (men) who were the main breadwinner at z = -1 and those who were secondary earners. This allow us to explore if there are heterogeneous responses across the various groups.

Although fertility choices are endogenous, the event of having a child generates sharp changes in labour market outcomes just after the birth, which can be assumed orthogonal to unobserved determinants of those outcomes as they should evolve smoothly over time (Kleven et al., 2018). Therefore, by comparing the results just before and after birth it is possible to estimate the effects of having a child.¹⁰

¹⁰Children may have two conceptually different effects on labour market outcomes. The event study approach

As in quasi-experimental settings, the main identification assumption is that is that the "treatment", that is, the decision on when to have a child should not be induced by unobservable information of a changed direction of the outcome of one of the spouses. This means that the timing of parenthood should not be influenced by expected future shocks to the participation and hours of work that they would have experienced in the absence of entering parenthood. However, if these shocks take place before the birth of the child, they are observed in our data and hence controlled for. Our graphical evidence supports this assumption. In the descriptive figures corresponding to the levels of employment and hours of work before and after the first child born (Figure 2), it is possible to observe that the sharp breaks in labour market outcomes trajectories occur just after the birth of the first child. That is, there is no evidence that the outcomes respond prior to the childbirth.

5.2 Child penalty in labour supply

Before presenting the estimated impacts of children, consider Figure 2 which depicts the average levels of employment, hours of work and housework around the birth of the first child. In Figure 2(a) it is possible to observe that before the birth of the first child the employment trajectories of men and women evolve in parallel trends, suggesting no gender-specific trend over time-events but with higher employment levels for men. Immediately after the birth of the first child, there is a sharp drop in employment levels for women and a slight increase for men. A similar situation is observed with respect to the hours of work and housework. Conditional on remaining employed, there is a sharp drop in the yearly hours of work for women after having their first child (Figure 2(b)), while there is an important increase in the hours of housework performed by women (Figure 2(c)).¹¹ This indicates that the gaps between men and women that exist prior to the birth become much larger after the first child arrives.

Results of the event study for the entire sample of women and of men are reported in Figure 3. As previously defined, these are year-z effects (as a percentage of the counterfactual outcome absent children) relative to the year just before the first child birth. The figure includes 95% confidence bands around the event coefficient. Once life-cycle and time trends are controlled

is designed to identify only post-child effects of realized fertility. It cannot capture pre-child effects of anticipated fertility (for example if women choose certain occupations in anticipation of expected fertility). This effect is incorporated in the pre-event levels. Then, the event study estimates provide a lower bound on the total lifetime impact of children. Taking a structural approach, Adda et al. (2017) estimate that occupational choices at the age of 15–16 due to anticipated fertility represent less than 5% of total earnings loss from children.

¹¹Figures A.1 and A.2 in the Appendix present these trends by educational attainment. These sharp drops in labour market trajectories take place for women in all educational groups. The most noticeable difference among them is that the trends in employment levels for highly-educated women before having the child were almost identical to those of men.

for, men and women have stable paths before parenthood, but the dynamics differ following the arrival of the first child. Women experience an immediate drop in both employment and hours of around 30%. Hours increase slightly and stabilize on year z + 2, while employment stays low for 6 years and increases moderately after that. Ten years after having their first child, women are 30% less likely to be employed and work 28% fewer hours compared with the year just before motherhood. In contrast, men experience a small increase in both their employment probability and in hours (of 5% and 2%, respectively) following the birth of their first child.

Summing up, in line with previous existing work, such as Kleven et al. (2018) for Denmark, we find that the arrival of children has a negative effect on women's trajectories, reducing both labour force participation and their hours of work. These responses can, however, differ across women, and here we explore to what extent education and relative earnings affect the child penalty.

Figure 4 reports the results obtained when we divide the sample across education categories. Although the dynamics are similar across the three groups, the magnitudes differ. The effect on employment is strongest for those with at most a high-school degree, and at its lowest results in employment being 40% lower than the pre-birth rate. For the other two education groups it oscillates at around 30% of its pre-birth level. In contrast, the three groups have a similar behaviour in terms of hours worked. Conditional on remaining employed, the arrival of the first child has similar effects on the yearly hours of work for women of different educational levels, a reduction of about 35%, although by the end of the period the impact is slightly larger for high-skilled women than for the other two groups.

5.3 Child rearing and relative earnings within the couple

We next consider the interaction of child rearing and the male breadwinner norm. An efficient allocation of household resources would imply that when a child arrives the parent with the higher income share remains at work and the one with the lowest share adjusts work hours or employment status in order to take on household responsibilities. Figure 5 presents the estimated impacts when we divide the sample of women into women who were earning more than 50% of the household's labour income and those that were earning a smaller share prior to the birth. Our results do not support an efficient division of labour. Both women who were earning less and those who were earning more than their husband reduce their hours and employment probability dramatically, while men exhibit stable hours and a slight increase in their employment probability as an effect of the first child birth. Conditional on being employed, the reduction in hours of work is roughly the same for women who are and are not the main

breadwinner (right panel).

The magnitude of the effect in employment is, however, different across the two groups of women (left panel). The effect of the first child born is a drop in employment of around 20% for women who were main breadwinners, and about 35% for those who were secondary earners. Interestingly, the dynamics differ. Secondary earners exhibit a sharp initial reduction, but the employment probability recovers after year 6. Their estimated effect after ten years is around 28%. In contrast, for those who were the main breadwinners, the initial drop is much smaller but the employment probability keeps decreasing over time, and 10 years after the birth the estimated effect is similar for the two groups of women.¹²

To further understand the differences between the two groups, Figure 6 presents the estimates for those women who have only one child (23.5% of our sample).¹³ Again, there is a considerable difference in the magnitude of the initial employment reduction between those who earn less and those who earn more than their husbands. The latter have an initial negative effect on employment of around 12% and then fluctuate around 10% for the next nine years. Women who are secondary earners experience a much larger effect of the first child birth (of 30%) but then this estimate slowly increase. Ten years after the birth they are around 12% less likely to work than before the birth, an estimated effect equivalent to that of the breadwinners. These results confirm those in Figure 5 indicating that employment behaviour initially differs across groups, but not in the long run.

We also consider the impact of children on men's employment. Again, we split the sample between men earning more and those earning less than their wives and report the change relative to the year before the birth of the first child. For those who were already main breadwinners, having a child has virtually no effect on the probability of employment and hours worked. However, men that were earning less than their wives before the birth, increase their probability of employment by around 12%, while the estimated effect on their hours of work rises steadily in the 10 years following the birth, and hours are around 20% higher by the end of the decade (Figure 7). There are two possible reasons for these changes. One is simply that the increased

 $^{^{12}}$ The short-run impact is estimated considering event times just before and after time zero. However, when we consider an event time long after zero, for example ten years after the birth of the first child, it is necessary to consider, first, that it can be capturing the effect of total lifetime fertility as opposed to the effect of the first child. Second, the smoothness assumption is no longer sufficient for identification and the long-run child penalty might be a biased estimate of the true post-child impact. However, by comparing standard event study estimates to more sophisticated event study approaches that use control groups or instrument for child births, Kleven et al. (2015) show that the event study approach, once we control non-parametrically for age and time trends, does a good job of identifying child penalties even in the long run.

¹³Given that the number of observations is considerable reduced, the confidence intervals become larger, especially for women with only one child who earn more than their husbands.

labour market activity of these men seeks to compensate the reduction in household income occurring when their wives, who were bringing home a high share of earnings, reduce their employment. Alternatively, men may react to the presence of children by reverting to a traditional division of labour and hence increase their market activity and reduce the amount of time spent in housework.

We consider this possibility by examining the response of hours of housework following the birth of the first child. Average weekly hours of housework before having children are 12 hours for women and 7 for men (Figure 2). Following the birth of their first child, the dynamics of housework for men and women diverge, with men still performing about 7 hours per week and women almost doubling them (to 23 hours).

When we perform our estimations of the effect of the first child on housework we find that the birth does not affect the hours spent in housework for men, and that this is the case for both men that were and were not the main breadwinner (Figure 8, right panel). For women there is a large increase just after the child's birth, with hours roughly doubling one year after the birth. The left panel of Figure 8, reports the estimates for both secondary earners and women who were the main breadwinner prior to the child's arrival. Initially, the estimated effect of the child is an increase by about the same amount in the number of weekly hours of household, and for both groups hours keep increasing. However the increase in the estimated effects is much faster for those who were the main breadwinners, who after 10 years have an effect of 250% compared to the year before motherhood.

Our results contrast to those obtained by Angelov et al. (2016) on Swedish data. They find that wives whose husbands are relatively better paid take greater responsibilities at home after the birth of a child, compared to those whose husbands had a lower income (relative to the wife's), in line with the theory of comparative advantage within the couple. Our results for the US indicate the opposite. Both women who were secondary earners and main breadwinners take the greatest responsibilities at home when the child arrives increasing their hours of housework, while it remains almost unchanged for men. Moreover, those women who were the main breadwinners prior to the child's birth experience a greater change in their housework time compared to the group of secondary earners. The contrasting evidence could be explained by different strength of the social norm in the two countries.

Our results indicate that the effect of the arrival of children on labour market trajectories is large, even for those women who were the main breadwinners in their households prior to the child's birth. Women seem to take on the role of main carer, reducing employment and (for those who remain employed) hours. There are, nevertheless, important differences between women that were the main providers in their households before being mothers and those that were secondary earners. Although, conditional on employment, the groups have similar reactions in terms of hours of work, they differ in the probability of employment and the amount of housework performed. Initially the effect on employment is greater for secondary earners, yet the two groups of women exhibit the same reduction in their employment probability after ten years. In contrast, the pattern for hours of housework indicates that it is those women that were initially the main breadwinner that experience the greatest effect in terms of domestic work (compared with their previous levels) when children arrive.¹⁴

5.4 Career changes

An additional way in which individuals can adjust their earnings and hours of work is through changes in occupation and industries. We hence consider next if the birth of the first child affects men and women's choices in this respect. There is evidence that women are employed in different occupations and industries compared with men (horizontal and vertical segregation), and this is especially so in the case of women with children. It is then conceivable that once women become mothers, they make career choices that are more compatible with becoming the main carer (i.e. that favour family amenities over pecuniary rewards). To analyse this possible way of adjusting to the norm of women as main caregivers, we estimate the effect of the first child's birth on the probability of being employed in a feminized industry or occupation relative to the probability of being in such an industry/occupation the year before the birth of the first child. Furthermore, a sizeable literature provides evidence that feminized occupations pay less than male occupations for workers with similar measured characteristics (e.g., Levanon et al. 2009; Blau and Kahn 2017; Bayard et al. 2003). Then, switching to such a job can be interpreted also as a way for women to reduce their earnings and hence conform to the male breadwinner norm.

We create our industry and occupation variables by recovering information on three different variables defined for the head of household and wife in the PSID Family files. From 1968-1980 we consider the industry and occupation variables that were coded retroactively using original

¹⁴The short-run impact is estimated considering event times just before and after time zero. However, when we consider an event time long after zero, for example ten years after the birth of the first child, it is necessary to consider, first, that it can be capturing the effect of total lifetime fertility as opposed to the effect of the first child. Second, the smoothness assumption is no longer sufficient for identification and the long-run child penalty might be a biased estimate of the true post-child impact. However, by comparing standard event study estimates to more sophisticated event study approaches that use control groups or instrument for child births, Kleven et al. (2018) show that the event study approach, once we control non-parametrically for age and time trends, does a good job of identifying child penalties even in the long run.

PSID reports and the three-digit 1970 Census industry and occupation codes.¹⁵ For the period 1981-2001 the industry and occupation information was obtained using the 3-digit 1970 Census Industry code and, finally, for 2003-2015 the 3-digit 2000 Census code was used. In all the cases the information corresponds to the main job.¹⁶ Based on this, we construct 19 industry categories and 25 occupational categories.

We define feminized industries and occupations as those in which women represent at least 50% of the employed and the average of hours worked by women is less than 1650 hours.¹⁷ Taking these criteria (see Tables A.1 and A.2 in the Appendix), we classify as feminized industries: *Retail trade, Accommodation and food services, Other services excluding public administration, Arts, entertainment and welfare services, and Education and related services.* The feminized occupations are: *Education (excluding university teachers) and library, Personal care and services, Maids and housekeepers, Food preparation and serving related, Arts, design entertainment and media, Sales and related, Healthcare support, and Office and Administrative support.*

As a robustness exercise, we use a broader criteria to classify industries and occupations as feminized. We relax the working hours restriction to include some categories in which women are highly represented but also have a higher average of hours worked. In the case of industries, we add *Health care and Social assistance*, in which women represent 81.4% of the labour force and have a mean of 1724 hours. Regarding the occupational categories, we include also *Health technologist and technicians* and *Community and social services*, where women account for 78.4% and 59.3%, respectively. It is noting that for all occupations the average number of hours worked by men is higher than for women. For example, for *Healthcare practitioners*, the mean for women is 1792 yearly hours, while for men it is 2368.

We define indicator variables that take the value of one if the individual is employed in a feminized industry/occupation and zero otherwise, and estimate the effect of the first child's birth on the probability of being employed in a feminized industry/occupation, using the previous event study methodology (conditional on remaining employed). Figure 9 presents the estimated

¹⁵This retroactive coding was done for a selected sample of PSID heads and wives: (a) Original sample Heads and Wives/"Wives still living by 1992 who reported main jobs in at least three waves during the period 1968-1992, with at least one of those reports prior to 1980. (b) Additionally, original sample Heads and Wives/"Wives" who had reported at least one main job between 1968 and 1980 but were known to have died by 1992. The selection criteria did not include all Heads and Wives/"Wives" who had worked between 1968 and 1980, therefore this variable contains missing information. For detailed information about the Retrospective Coding Project please see the document, "A Panel Study of Income Dynamics: 1968-1980 Retrospective Occupation-Industry Files Documentation", available on PSID website.

¹⁶Since there is no clear correspondence between 1970 and 2000 census industries and occupation codes, a reclassification was needed is some cases based on the description of the 3-digit categories.

¹⁷For reference, note that a worker who works 40 hours per week, with 2 vacation weeks, sum 2,000 hours worked per year. The average number of hours worked (including all employed, full and part time) in the US in 2015 was 1786 hours (OECD Stat: https://stats.oecd.org).

effects of children for industries (left panel) and occupations (right panel) for men and women. As far as industry is concerned, we can see that men and women are on very similar pre-child trends in terms of their probability of working in a feminized industry, but begin to diverge soon after having a child. While men do not change after the birth, the probability that a woman is in such an industry increases steadily after the birth and is almost 50 percent higher 10 years after having the first child than it was before the birth. The pattern for occupations is quite different. Throughout their working life men seem to increase their probability of being in such occupations, as an effect of fatherhood. For women, the probability of being about 40 percent higher than in the reference period ten years on. The estimated effects using the broader criteria to define feminized industries and occupations (Figure A.8 in the Online Appendix), are slightly smaller in size but show the same trends. Ten years after child birth, women are almost 40 percent more likely to be employed in a feminized industry/occupation relative to the year before child birth.

6 Conclusions

The aim of this paper is to explore the heterogeneity in response to gender norms, both in terms of how different types of women respond and in comparing reactions across norms. We consider two types of norms, that according to which mothers should be the main carers of children and that which maintains that the man should be the household's main breadwinner. Our results indicate that distinct gender roles have different effects.

The insights from this study may be summarized in two broad conclusions. First, our event study analysis shows that the relationship between children and labour trajectories is strongly gendered. Motherhood has a strong impact on women's labour supply. Both women who were earning less and those who were earning more than their husband reduce their hours and employment probability dramatically, while men exhibit stable hours and a slight increase in their employment probability as a results of the first child's birth. Although the two groups of women differ in the initial strength of their response, the medium term effect is almost the same. This indicates that the social identity norm strongly shapes responses and prevails over considerations of comparative advantage in the household.

Secondly, the effect of the breadwinner norm is more nuanced, and important differences appear across groups of women. Having children seems to be key to determine whether couples adhere to the norm, with the effect being statistically insignificant for childless couples. When we split the data by education level, we find that the most skilled women reduce their hours of work but do not leave the labor market, consistent with the idea that their job is an important source of identity. For those in the middle of the skill distribution, the norm implies a reduction in employment but not in hours worked, while the least skilled reduce both. Such results are probably a combination of the perceived importance of conforming across groups as well as of constraints on the capacity to modify working hours.

Interestingly, our analysis finds effects that are much stronger than those obtained by similar work focusing on Nordic countries. The institutional context is vastly different, with the US being the only high-income country without paid maternity leave, while countries like Sweden and Denmark have extensive maternal and paternal leave (Olivetti and Petrongolo, 2017). This points to these policies as being an important element in shaping the perceived importance of social norms and hence the cost that they impose on women's careers.

References

- Adda, J., C. Dustmann, and K. Stevens (2017). The career costs of children. Journal of Political Economy 125(2), 293–337.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. The Quarterly Journal of Economics 115(3), 715–753.
- Angelov, N., P. Johansson, and E. Lindahl (2016). Parenthood and the gender gap in pay. Journal of Labor Economics 34(3), 545–579.
- Angrist, J. D. and W. N. Evans (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *American Economic Review* 88(3), 450–477.
- Bayard, K., J. Hellerstein, D. Neumark, and K. Troske (2003). New evidence on sex segregation and sex differences in wages from matched employee - employer data. *Journal of Labor Economics* 21(4), 887–922.
- Becker, G. S. (1973). A theory of marriage: Part i. Journal of Political Economy 81(4), 813-846.
- Bertrand, M. (2011). New Perspectives on Gender, Volume 4 of Handbook of Labor Economics, pp. 1543 – 1590. Elsevier. Editors David Card and Orley Ashenfelter.
- Bertrand, M., C. Goldin, and L. F. Katz (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Econom*ics 2(3), 228–255.
- Bertrand, M., E. Kamenica, and J. Pan (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics* 130(2), 571–614.
- Bittman, M., P. England, L. Sayer, N. Folbre, and G. Matheson (2003). When does gender trump money? bargaining and time in household work. *American Journal of Sociology* 109(1), 186– 214.
- Blau, F. D. and L. M. Kahn (2017). The gender-wage gap: Extent, trends, and explanations. Journal of Economic Literature 55(3), 789–865.
- Breen, R. and C. García-Peñalosa (2002). Bayesian learning and gender segregation. Journal of Labor Economics 20(4), 899–922.

- Farré, L. and F. Vella (2013). The intergenerational transmission of gender role attitudes and its implications for female labour force participation. *Economica* 80(318), 219–247.
- Fernández, R. (2007). Women, work, and culture. Journal of the European Economic Association 5(2-3), 305–332.
- Fernández, R., A. Fogli, and C. Olivetti (2004). Mothers and sons: Preference formation and female labor force dynamics. The Quarterly Journal of Economics 119(4), 1249–1299.
- Fortin, N. M. (2005). Gender role attitudes and the labour-market outcomes of women across oecd countries. Oxford Review of Economic Policy 21(3), 416–438.
- Gallup News, w. (2014). Gallup poll social series: Work and education. http://news.gallup.com/.
- Goldin, C. (2006). The quiet revolution that transformed women's employment, education, and family. *The American Economic Review* 96(2), 1–21.
- Goldin, C. (2014). A grand gender convergence: Its last chapter. The American Economic Review 104(4), 1091–1119.
- 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., L. F. Katz, and I. Kuziemko (2006). The homecoming of american college women: The reversal of the college gender gap. *The Journal of Economic Perspectives* 20(4), 133–133.
- Hederos Eriksson, K. and A. Stenberg (2015). Gender identity and relative income within households: Evidence from sweden. IZA Discussion Paper No. 9533.
- Kleven, H., C. Landais, and J. E. Sogaard (2018). Children and gender inequality: Evidence from denmark. NBER Working Paper No. 24219, January 2018.
- Levanon, A., P. England, and P. Allison (2009). Occupational feminization and pay: Assessing causal dynamics using 1950-2000 us census data. *Social Forces* 88(2), 865–891.
- Lippmann, Q., A. Georgieff, and C. Senik (2016). Undoing gender with institutions. lessons from the german division and reunification. PSE Working Papers No. 2016-06.
- Lucifora, C., D. Meurs, and E. Villar (2017). Children, earnings and careers in an internal labor market.

- Lundborg, P., E. Plug, and A. W. Rasmussen (2017). Can women have children and a career? iv evidence from ivf treatments. *American Economic Review* 107(6), 1611–37.
- Nollenberger, N., N. Rodríguez-Planas, and A. Sevilla (2016). The math gender gap: The role of culture. *The American Economic Review* 106(5), 257–261.
- Olivetti, C., E. Patacchini, Y. Zenou, et al. (2016). Mothers, peers and gender identity. Boston College, Department of Economics, Working Paper, 904.
- Olivetti, C. and B. Petrongolo (2017). The economic consequences of family policies: Lessons from a century of legislation in high-income countries. *The Journal of Economic Perspectives* 31(1), 205–230.
- Paull, G. (2008). Children and women's hours of work. The Economic Journal 118(526), F8–F27.
- West, C. and D. H. Zimmerman (1987). Doing gender. Gender & Society 1(2), 125–151.
- Wieber, A. and E. Holst (2015). Gender identity and womens' supply of labor and non-market work: Panel data evidence for germany. DIW Berlin Discussion Paper No. 1517, November 2015.

Tables

-	Mean	SD	Min	Max	N
In t-1:					
Wife earns more	0.18	0.38	0	1	129.265
Wife's Income	20.785	25.766	0	1,027,758	129.265
Husband's Income	53.779	68.271	0	6,410,906	129.265
Relative Income	0.29	0.28	0	1	129.265
In t:					
Wife's age	37.15	11.12	18	65	129.265
Husband's age	39.55	11.49	18	65	129.265
Number of children at home	1.38	1.39	0	13	129.265
No children at home	0.34	0.47	0	1	129.265
Wife's labor force participation	0.63	0.48	0	1	129.265
Wife's yearly hours of work	1.152	918.9	0	7,980	129.265
Highschool or less	0.49	0.50	0	1	129.265
1-3 years of college	0.26	0.44	0	1	129.265
At least 4 years of college	0.24	0.43	0	1	129.265
Ever had a child	0.85	0.35	0	1	120.789
Number of children	2.53	1.72	0	15	120.789

Table 1: Sample summary statistics

Notes: The data is from the 1968-2015 PSID. The sample correspond to couple-year observations of women who live in couple (wives or non-married cohabitant couples), where both husband and wife are between 18 and 65 years old, and where at least one of the spouses has positive earnings (corresponding to the previous year). Yearly labour income is expressed in 2015 prices (deflated using Consumer Price Index (CPI) - Bureau of Labor Statistics). The number of observations in the variables corresponding to the *Childbirth History File* is smaller because information was not collected for some individuals.

	(1)	(2)	(3)	(4)			
	Dependent variable: female labor force participation						
	(a) All women in the sample						
Wem _{t-1}	-0.020***	-0.021***	-0.021***	-0.023***			
	[0.004]	[0.004]	[0.004]	[0.004]			
Observations	129,265	129,265	129,265	129,265			
R-squared	0.668	0.673	0.674	0.609			
	(b) Highschool or less						
Wem _{t-1}	-0.021***	-0.033***	-0.032***	-0.034***			
	[0.006]	[0.007]	[0.007]	[0.006]			
Observations	63,080	63,080	63,080	63,080			
R-squared	0.687	0.691	0.692	0.630			
	(c) Some college						
Wem _{t-1}	-0.022***	-0.019**	-0.018**	-0.019***			
	[0.007]	[0.008]	[0.008]	[0.007]			
Observations	33,953	33,953	33,953	33,953			
R-squared	0.622	0.626	0.628	0.553			
	(d) At least 4 years of college						
Wem _{t-1}	-0.013**	-0.006	-0.007	-0.008			
	[0.006]	[0.006]	[0.006]	[0.006]			
Observations	31,502	31,502	31,502	31,502			
R-squared	0.625	0.629	0.634	0.571			
Additional controls Couple fixed effects Cubic in InWifeIncome	yes	yes	yes	no			
and InHusbIncome	no	yes	yes	yes			
Relative Income	no	yes	yes	yes			
Children controls	no	no	yes	yes			

Table 2: Female Labour force participation and relative income. 1968-2015 PSID

Notes: The data is from the 1968-2015 PSID. The dependent variable is an indicator variable which takes the value of one if the woman is in the labour force, zero otherwise. Wem_{t-1} is an indicator variable that equals one if the relative income is greater than 0.5 at time t - 1. All regressions include the log of the wife's income, the log of the husband's income, the log of the couple's income, an indicator for whether only the wife is working or only the husband is working, a quadratic in wife's and husband's age, year fixed effects and state fixed effects. *Children controls* include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, or older than 6. The regression is estimated using a linear probability model. Standard errors are clustered at the couple level and are reported in brackets. *** p<0.01, ** p<0.05, * p<0.127

	(1)	(2)	(3)	(4)			
	Dependent variable: In yearly hours of work						
	(a) All women in the sample						
Wem _{t-1}	-0.097***	-0.014**	-0.013**	-0.019***			
	[0.006]	[0.006]	[0.006]	[0.006]			
Observations	89,606	89,606	89,606	89,606			
R-squared	0.787	0.805	0.807	0.718			
-		(b) Highschool or less					
Wem _{t-1}	-0.101***	-0.027***	-0.026***	-0.029***			
	[0.008]	[0.010]	[0.010]	[0.009]			
Observations	38,809	38,809	38,809	38,809			
R-squared	0.801	0.822	0.822	0.743			
_	(c) Some college						
Wem _{t-1}	-0.101***	-0.007	-0.006	-0.014			
	[0.012]	[0.011]	[0.011]	[0.010]			
Observations	25,241	25,241	25,241	25,241			
R-squared	0.770	0.794	0.796	0.712			
		(d) At least 4 y	vears of college				
Wem ₁₋₁	-0.085***	-0.030***	-0.029**	-0.023**			
	[0.010]	[0.012]	[0.011]	[0.011]			
Observations	25,113	25,113	25,113	25,113			
R-squared	0.777	0.791	0.794	0.713			
Additional controls Couple fixed effects Cubic in lnWifeIncome and	yes	yes	yes	no			
InHusbIncome	no	yes	yes	yes			
Relative Income	no	yes	yes	yes			
Children controls	no	no	yes	yes			

Table 3: Hours of work and relative income. 1968-2015 PSID

Notes: The data is from the 1968-2015 PSID. The dependent variable is the logarithm of yearly hours of work. The regressions are conditional on being employed (positive labour income and hours). Wem_{t-1} is an indicator variable that equals one if the relative income is greater than 0.5 at time t - 1. All regressions include the log of the wife's income, the log of the husband's income, the log of the couple's income, an indicator for whether only the wife is working or only the husband is working, a quadratic in wife's and husband's age, year fixed effects and state fixed effects. *Children controls* include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, or older than 6. The regression is estimated using a linear probability model. Standard errors are clustered at the couple level and are reported in brackets. *** p < 2981, ** p < 0.05, * p < 0.1.

	(1)	(2)	(3)	(4)	(5)	(6)		
		(a) Dependent variable: female labor force participation						
	1	No child born			At least one child born			
Wem _{t-1}	-0.013*	-0.004	-0.011	-0.021***	-0.012***	-0.018***		
	[0.007]	[0.008]	[0.008]	[0.004]	[0.004]	[0.005]		
Observations	17,761	17,761	17,761	103,028	103,028	103,028		
R-squared	0.369	0.375	0.375	0.408	0.415	0.415		
	(b) Dependent variable: In yearly hours of work							
	1	No child born			At least one child born			
Wem _{t-1}	-0.075***	-0.006	0.001	-0.101***	-0.031***	-0.011		
	[0.013]	[0.012]	[0.014]	[0.007]	[0.006]	[0.007]		
Observations	14,928	14,928	14,928	69,830	69,830	69,830		
R-squared	0.598	0.627	0.627	0.633	0.665	0.665		
Couple fixed effects	yes	yes	yes	yes	yes	yes		
and InHusbIncome	no	yes	yes	no	yes	yes		
Relative Income	no	no	yes	no	no	yes		

Table 4: Labour force participation and relative income by ever had a child. 1968-2015 PSID

Notes: The data is from the 1968-2015 PSID. In Panel (a) the dependent variable is an indicator variable which takes the value of one if the woman is in the labour force, zero otherwise. In panel (b) the dependent variable is the logarithm of yearly hours of work, and the regressions are conditional on being employed (positive labour income and hours). No child born is a subsample of women who in year t do not have any child born, while At least one child born includes those women who in year t have at least one child born. Wem_{t-1} is an indicator variable that equals one if the relative income is greater than 0.5 at time t - 1. All regressions include the log of the wife's income, the log of the husband's income, the log of the couple's income, an indicator for whether only the wife is working or only the husband is working, a quadratic in wife's and husband's age, year fixed effects and state fixed effects. Children controls include indicator variables for whether the respondent has no children, whether the youngest child is 3 or younger, between 4 and 6, or older than 6. The regression is estimated using a linear probability model. Standard errors are clustered at the couple level and are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Figures



Figure 1: Gender norms in the US

Notes: Author's calculations based on the *World Value Survey* data for the US, 2011. Each column shows the percentage of respondents who agree or strongly agree with the corresponding statement. For the first statement the possible answers were: Agree, Disagree or Neither, and only respondents who answer "Agree" where included. For the two other questions the possible answers were Strongly agree, Agree, Disagree, Strongly disagree, and both respondents who answer "Agree" or "Strongly agree" were included.



Figure 2: Average employment, yearly hours of work and housework before and after having their first child

Notes: The graphs show average levels of employment, yearly hours of work (conditional on being employed) and weekly hours of housework for men and women in each event-time before and after the first child is born (vertical line).



Figure 3: Impact of children on employment and hours worked

Figure 4: Impact of children by educational levels



Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women. The effects on yearly hours worked are estimated conditional on employment. The shaded 95 % confidence intervals are based on robust standard errors.



Figure 5: Impact of children and main breadwinner norm

Figure 6: Impact of children and main breadwinner norm: employment



Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women who were main breadwinners or secondary earners in z = -1. The effects on yearly hours worked are estimated conditional on employment. The shaded 95 % confidence intervals are based on robust standard errors.



Figure 7: Impact of children and main breadwinner norm

Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women who were main breadwinners or secondary earners in z = -1. The effects on yearly hours worked are estimated conditional on employment. The shaded 95 % confidence intervals are based on robust standard errors.



Figure 8: Impact of children on hours of housework

Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women who were main breadwinners or secondary earners in z = -1. The shaded 95 % confidence intervals are based on robust standard errors.



Figure 9: Impact of children on feminized industries and occupations

Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women. Sample is restricted to men and women for whom an occupation and industry code is declared in z = -1. Feminized industries and occupations are described in Sub-section 5.4 The shaded 95 % confidence intervals are based on robust standard errors.

Appendix

	Percentage	Distribution of	Yearly hours of work			
Industry	of women	industries	Women		Men	
(a) Feminized industries			Mean	Std. Dev.	Mean	Std. Dev.
Other services, exc public adm	65.6	9.1	1298	7.6	2100	11.1
Retail Trade	55.6	13.9	1480	5.7	2169	6.5
Education and related services	72.2	13.3	1507	5.1	2005	8.3
Accomodations and Food services	65.9	3.1	1516	11.6	1958	17.5
Arts, Entertainment and recreation and welfare services	67.6	3.9	1559	11.4	1918	15.9
(b) Included as feminized industries in the broader classification						
Health care and social assistance	81.4	16.3	1724	4.5	2102	10.3
(c) Non feminized industries						
Agriculture, Forestry, Fishing and Hunting	17.3	1.0	1231	25.3	2319	12.9
Management, Administrative and support and wast management services	44.1	1.9	1584	16.1	1922	14.9
Profesional, Scientific and techinical services	47.5	4.3	1617	9.7	2086	9.2
Transportation and warehousing	23.3	1.9	1620	14.8	2199	8.4
Real State	46.2	1.4	1663	17.8	2056	17.1
Construction	7.6	1.0	1666	20.0	1960	5.5
Wholesale trade	28.2	2.0	1724	13.1	2217	7.9
Mining	14.4	0.2	1751	39.1	2326	19.0
Manufacturing	34.5	13.4	1764	4.5	2135	3.1
Information, Newspapers, Radio, etc.	49.2	1.7	1781	11.7	2171	12.1
Public Administration and active duty military	40.0	5.6	1785	7.2	2200	6.1
Finance and Insurance	65.2	5.6	1793	6.3	2183	9.7
Utilities	15.9	0.5	1837	21.5	2106	8.9

Table A.1: Feminized industries classification. Descriptive statistics

Notes: The data is from the 1968-2015 PSID. The Industry and occupation variables correspond to main job, and are defined for Head of household (man, or woman if single) and wife (including cohabitors) at the family level. Since there is no perfect correspondence between 1970 and 2000 census industries and occupation codes, a reclassification was needed is some cases based on the description of the 3-digit categories.

	Percentage of women	Distribution of	Yearly hours of work			
Occupations		occupations	Women		Men	
(a) Feminized occupations			Mean	Std. Dev.	Mean	Std. Dev.
Maids and housekeepers, cleaners	97.5	3.2	1116	11.6	1698	92.3
Personal care and services	87.2	5.7	1392	10.0	1881	25.9
Food preparation and serving related	73.1	6.1	1404	8.1	1846	14.5
Sales and related	50.9	5.9	1428	8.8	2175	8.2
Arts, design, entertainment, sports and media	52.6	1.6	1497	18.0	1923	19.1
Education, teachers (except university) and library	78.7	6.6	1525	7.0	1954	14.0
Office and administrative support	81.8	26.5	1641	3.3	1995	7.2
Healthcare support (nurses, etc)	91.0	8.1	1644	6.6	1949	21.0
(b) Included as feminized occupations in the broader classification						
Community and social services	59.3	1.7	1734	13.5	2150	21.0
Health technologist and technicians	78.4	1.8	1780	12.9	2199	27.5
(c) Non feminized occupations						
Farm, forestry and fishing laborers	13.6	0.6	1006	30.5	1922	11.3
Building and grounds cleaning and maintenance	42.3	3.3	1414	10.9	1750	9.9
Transportation and material moving	13.3	1.8	1514	16.0	2098	6.0
Education, university teachers	48.3	0.7	1569	26.4	2083	27.4
Production	41.1	10.4	1658	5.5	2088	4.4
Protective service workers	22.7	0.8	1725	21.4	2210	11.9
Installation, repair and maintenance	3.7	0.4	1748	27.9	2096	5.2
Healthcare practitioners	53.1	1.2	1792	18.3	2368	21.8
Architect and engineering	13.0	0.7	1828	19.8	2157	6.6
Construction and extraction	5.4	0.5	1846	26.3	2015	6.1
Life, physical and social scientist	46.2	1.2	1872	15.3	2195	14.2
Computer and mathematical	29.3	0.8	1926	15.7	2114	10.5
Legal, lawyers and judges	35.1	0.5	1936	27.1	2236	19.1
Managers and administrators, accountants, business and financial operations	34.4	9.7	1960	6.0	2403	4.6
Military	17.0	0.1	2028	83.5	2518	33.6

Table A.2: Feminized occupations classification. Descriptive statistics

Notes: The data is from the 1968-2015 PSID. Industry and occupation variables correspond to main job, and are defined for Head of household (man, or woman if single) and wife (including cohabitors) at the family level. Since there is no perfect correspondence between 1970 and 2000 census industries and occupation codes, a reclassification was needed is some cases based on the description of the 3-digit categories.

NOT FOR PUBLICATION APPENDIX

Gender norms and labour supply: Identifying heterogeneous patterns across groups of women

Online Appendix

A Additional figures



Figure A.1: Average employment level for women and men before and after having their first child

Notes: The graphs show average levels of employment in each event-time before and after the first child is born (vertical line).



Figure A.2: Average yearly hours for women and men before and after having their first child

Notes: The graphs show average levels of yearly hours of work (conditional on being employed) in each event-time before and after the first child is born (vertical line).



Figure A.3: Average weekly hours of housework before and after having their first child

Notes: The graphs show average levels of weekly hours of housework in each event-time before and after the first child is born.



Figure A.4: Impact of children and main breadwinner norm: Hours of work

Figure A.5: Impact of children and main breadwinner norm: Sample of women and men with two or more children born



Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women who were main breadwinners or secondary earners in z = -1. The effects on yearly hours worked are estimated conditional on employment. The shaded 95 % confidence intervals are based on robust standard errors.



Figure A.6: Female relative income before and after having their first child



Figure A.7: Impact of children and main breadwinner norm: Sample of husbands with positive earnings

Figure A.8: Impact of children on feminized industries and occupations, broader criteria.



Notes: The graphs show event time coefficients estimated from equation (2) as a percentage of the counterfactual outcome absent children: $P_t^g = \hat{\alpha}_t^g / E[Y_{ist}^{\tilde{g}}|t]$ for men and women. Sample is restricted to men and women for whom an occupation and industry code is declared in z = -1. Feminized industries and occupations are described in Sub-section 5.4 The shaded 95 % confidence intervals are based on robust standard errors.