

The Mommy Effect: Do Women Anticipate the Employment Effects of Motherhood?

Ilyana Kuziemko, Jessica Pan, Jenny Shen, Ebonya Washington*

June 27, 2018

Abstract

After decades of convergence, the gender gap in employment outcomes has recently plateaued in many rich countries, despite the fact that women have increased their investment in human capital over this period. We propose a hypothesis to reconcile these two trends: that when they are making key human capital decisions, women in modern cohorts underestimate the impact of motherhood on their future labor supply. Using an event-study framework, we show substantial and persistent employment effects of motherhood in U.K. and U.S. data. We then provide evidence that women *do not anticipate* these effects. Upon becoming parents, women (and especially more educated women) adopt more negative views toward female employment (e.g., they are more likely to say that women working hurts family life), suggesting that motherhood serves as an information shock to their beliefs. Women on average (and, again, more educated women in particular) report that parenthood is harder than they expected. We then look at longer horizons—are young women’s expectations about future labor supply correct when they make their key educational decisions? In fact, female high school seniors are increasingly and substantially overestimating the likelihood they will be in the labor market in their thirties, a sharp reversal from previous cohorts who substantially underestimated their future labor supply. Finally, we specify a model of women’s choice of educational investment in the face of uncertain employment costs of motherhood, which demonstrates that our results can be reconciled only if these costs increased unexpectedly across generations. We end by documenting a collage of empirical evidence consistent with such a trend.

JEL codes: J16, J13, J22, D84

*We thank seminar participants at Princeton, Dartmouth, Tinbergen Institute, Aarhus University, GRIPS, NUS, Yale-NUS, UCLA Anderson, Cornell, and SOLE. We benefited from helpful discussions with Raquel Fernandez, Henrik Kleven, Suresh Naidu and Richard Rogerson. Courtney Burke, Elisa Jacome, Rais Kamis, Tammy Tseng and especially Elena Marchetti-Bowick have provided excellent research assistance. Kuziemko: Princeton and NBER (kuziemko@princeton.edu); Pan: National University of Singapore and IZA (jesspan@nus.edu.sg); Shen: Princeton (jenny.shen@princeton.edu); Washington: Yale and NBER (ebonya.washington@yale.edu).

1 Introduction

The entry of women, especially mothers, into the labor force in high-income countries has been one of the most transformative economic and social developments of the past century. However, in the US, this trend has plateaued since the 1990s (see Figure 1), and even in countries where it has continued to increase, it has been driven by part-time work (Blau and Kahn, 2013). This “stalling” of male-female convergence in the labor market occurs *despite* the fact that over this period women have increased their human capital in the form of education (they are now significantly more educated than men, see Figure 2) and in the form of job experience by delaying childbirth (Appendix Figure A.1).

In this paper, we propose a novel resolution to this puzzle. We hypothesize that modern cohorts of women *underestimate the employment costs of motherhood*, which we loosely define as the time, effort or money required for them to raise their children in a manner they deem appropriate, while also working outside the home. Examples of these costs might include the per-hour cost of a nanny or day-care service, the emotional cost of being separated from the child while at work, guilt over (perceived or real) underperformance as an employee or mother, or diminished sleep or other aspects of wellbeing due to working while also performing child-care activities.¹ If women have assumed that these costs have fallen when in fact they remain high, then women would acquire costly human capital (because they assume they will enjoy the returns over a long period of labor-force attachment), but also exhibit diminished labor supply after motherhood (because, in fact, the employment costs of motherhood remain high, leading to withdrawal from the labor market). We build support for this hypothesis in several steps, using a variety of datasets from the US and the UK.

We begin by documenting in an event-study framework the evolution of employment outcomes before and after the “event” of the birth of a first child. Women’s probability of employment declines by roughly thirty to forty percentage points upon motherhood in the US and UK, effects that are substantially larger than the effects recently found in Scandinavian countries (Kleven *et al.*, 2018; Angelov *et al.*, 2016). This decline is essentially fully realized the year of the first child’s birth, and exhibits very little recovery in the five- to ten-year post-period we can observe. Heterogeneity in the employment response to children is also quite similar in the US and UK. For example, a college education has a modest “protective” effect on employment, but even these more educated women exhibit large and significant

¹See Bertrand (2013), who finds that among college-educated mothers, those with a career report being more unhappy, stressed, and tired than stay-at-home mothers.

“mommy effects” of decreased employment.

We argue that women *underestimate* these large employment effects, in both the short- and longer-run, and the bulk of the empirical work presents evidence supporting this claim. If the effect of motherhood on the ability to maintain both work and family commitments is indeed unexpected in the short-run, then women’s *beliefs* about the appropriate balance between home and market work should change discontinuously upon the birth of the first child, as this event allows them to update their beliefs with new information. We thus examine how attitudes on gender roles and work-family balance (e.g., agreement with statements such as “a woman and her family would all be happier if she goes out to work” or “both husbands and wives should contribute financially”) change once a woman has a child, using an event-study framework. Whereas before motherhood most women say that work does not inhibit women’s ability to be good wives and mothers, after the birth of their first child they become significantly more negative toward female employment, consistent with having underestimated the impact of motherhood on employment. While noisier than the employment results, the evolution of gender-role attitudes also roughly approximates a step-function: no meaningful pre-trends, a steep change around the time of the first child’s birth, and no recovery in the medium-run. Interestingly, the heterogeneity results are quite different than those for employment: while educated mothers experience smaller employment declines than their less educated counterparts, they exhibit larger updates in their beliefs (in the direction that working women hurt family life).

We interpret the event-study results for the work-family attitudes as suggesting that women have underestimated the employment effects of motherhood—as upon motherhood they become more pessimistic about women’s ability to both work outside the home and fulfill their perceived responsibilities at home—and that this tendency is especially pronounced for educated women. We find additional evidence for this claim from surveys that ask parents retrospective questions. The large majority of women agree that “being a parent is harder than I thought it would be,” and this opinion is concentrated among college-educated women. For men, a far lower share agree with that sentiment, and there is no difference by education.

The evidence described so far suggests that women underestimate the employment effects of motherhood, but for this underestimation to help explain the rise in female human capital accumulation (as we claim it does), then they would have to underestimate these effects while they are making their key education decisions. We thus examine the expectations of young women around the age of eighteen. Goldin (1990) and Goldin (2006) had earlier

demonstrated that teenage women in the US substantially *underestimated* their future labor supply, when asked in 1968 about their likelihood of employment in their thirties. We show that this result has now completely reversed for more modern cohorts. Even by 1978, only ten percent of female high school seniors said they would be home-makers at age thirty, a dramatic decline over a ten-year period and a roughly ten-point underestimate relative to their *actual* experience twelve years later. From roughly 1985 through the present, no more than two percent of female high school seniors say they expect to be home-makers at age thirty, even though the share of home-makers among thirty-year old women has remained roughly constant at 15-18 percent since 1990. Importantly, these same young women expect to be mothers (of two or three children), but not, apparently, stay-at-home mothers.

That modern cohorts of women seem to systematically underestimate the employment consequences of motherhood—and thus overestimate their future labor market attachment—begs the question of how they could get such an important prediction so substantially wrong. After all, motherhood is a very common event and they could in principle learn from their own mothers' experience and that of peers. We close the paper with our preferred explanation: *that the employment costs of motherhood have risen in an unexpected manner*. While we do not formally test and reject other explanations, we provide a variety of evidence that supports this proposed increase in costs.

We start by presenting a simple model of women's education and employment choices when the future employment costs of motherhood are uncertain. Women form predictions over the level of their own future costs by observing their own mothers, but they inherit the costs their mothers faced with some noise (which may or may not have a mean of zero). We first show that the model yields two of our most striking empirical results: that educated women exhibit smaller employment “mommy effects” than their less educated counterparts, but at the same time appear the most “surprised” by the high employment cost of motherhood (in terms of their greater updating of beliefs about work-family gender roles and their larger retrospective expressions of surprise at how hard parenthood is). Women in our model choose higher levels of education—which raises the return to market work *ex post* and thus results in higher labor supply post-motherhood—in part *because* they underestimate the cost of motherhood *ex ante*. We show that these results hold *regardless* of whether, on average, the employment costs of motherhood have increased or decreased for the current generation relative to the previous one. That these results hold under relatively general conditions not only bolsters the credibility of the model, but also serves as a “proof of

concept” for the empirical results. We then show that some of our other central results—that *unconditional on education* women in modern cohorts appear to underestimate the costs of motherhood—hold if and only if costs have on average *increased* for today’s mothers relative to earlier cohorts.

To complement this theory-based argument that the cost of motherhood must have risen, we provide a collage of empirical evidence suggesting that, while the cost likely fell during much of the twentieth century, it appears to have recently risen along some important metrics, especially since the 1990s. Some of this evidence we take from past literature and some, to the best of our knowledge, we provide for the first time. For example, past work has emphasized the *decline* in the cost of motherhood over the middle decades of the 20th century due to technological advances (e.g., household appliances and advances in infant formula) and links these falling costs to the large rise in female LFPR over the same period.² We argue that those advances may have fully played out and in fact more recent developments (e.g., research advocating the benefits of breastmilk over formula) may have effectively increased the costs of motherhood. Indeed, in both the US and the UK, we show that breastfeeding rates have increased significantly over the past several decades, especially since the 1990s. Similarly, while the underlying cause is still debated, the time cost associated with child care beyond infancy has been rising, especially for educated mothers (see Ramey and Ramey, 2010 and Guryan *et al.*, 2008 for evidence from the US and Borra and Sevilla, 2015 for the UK). Consistent with these trends, we show that large majorities of mothers tell survey takers that motherhood is harder today than for their own mothers and that they are more involved in their children’s lives than their mothers were in theirs. In the conclusion, we briefly take up the question of *why* the cost of motherhood may have increased over recent decades.

Our paper contributes to the literature on the employment effects of motherhood, a topic of long interest among labor economists. Perhaps surprisingly, the number of papers that have performed parallel analysis using the birth of a first child in an event-study analysis of labor supply measures is quite limited.³ Recent event-study analyses by Kleven *et al.*

²We review this literature in more detail in Section 7.

³Papers that have used other methods to examine the question of children and women’s labor supply include, of course, Angrist and Evans (1998), who focus on developing credible instruments for the birth of a *third* child. More recently, Lundborg *et al.* (2017) argue that, conditional on undergoing *in vitro* fertility (IVF) treatment, its success is sufficiently random to serve as a valid instrument for fertility. While event-studies cannot formally address endogeneity of fertility, interestingly, Kleven *et al.* (2018) shows that the event-study estimate of the effect of a *third* birth lines

(2018) and Angelov *et al.* (2016) using rich administrative data from Denmark and Sweden, respectively, uncover persistent declines in women’s relative earnings and labor supply at the onset of parenthood. Similar studies for the US include Chung *et al.* (2017), who use SIPP data linked to SSA earnings records to examine the impact of parenthood on the earnings gap between mothers and fathers. Bertrand *et al.* (2010) focus on female MBA graduates and show that these highly educated women reduce their labor supply in the year of first birth, with further reductions taking place over the next four years.

We also contribute to a large economics literature on gender-role norms, and in particular how they change over time.⁴ For example, Goldin (2006) suggests that innovations in contraception may have contributed to altering women’s identity in the 1960s and 1970s. Fernández *et al.* (2004) argue that men growing up in families with working mothers appear to have developed more liberal gender-role attitudes.⁵ We show that gender-role attitudes appear to exhibit life-cycle effects: women themselves adopt more traditional attitudes after they become mothers.

Of course, the idea that motherhood changes women’s priorities is hardly original, and appears frequently in the popular press.⁶ To the best of our knowledge, we are the first to examine how gender-role norms regarding employment change after parenthood in a formal event-study framework. We are aware of a small number of academic papers, mostly from sociology, that examine similar questions, all using panel data from high-income countries (Grinza *et al.*, 2017; Baxter *et al.*, 2015; Shafer and Malhotra, 2011; Jarrallah *et al.*, 2016; Vespa, 2009). But these panels tend to be short (in some cases, only two waves), and thus these analyses cannot perform event studies to determine if any estimated effect on attitudes

up closely to the IV-based estimate using twins or sex composition as an IV for a third birth. A small set of papers have examined the effect of childbirth using a woman fixed-effects strategy but have not traced out the effects in event-time, so have less ability to untangle coincident trends in labor supply and family planning (Waldfogel, 1997; Wilde *et al.*, 2010).

⁴Recent contributions to this literature have focused on norms in the marriage market or within couples. Bertrand *et al.* (2015) and Kleven *et al.* (2016) provide a variety of evidence that couples follow the norm that husbands should out-earn their wives. Bursztyn *et al.* (2017) show that traditional gender-role norms remain relevant even among career-oriented MBA students—anticipation of the marriage market appears to limit the public signaling of career ambitions of single female MBA students.

⁵On the other hand, Alesina *et al.* (2013) highlight the long-term persistence of gender norms, showing that ethnicities and countries whose ancestors practiced plough cultivation (which is more suited to male labor) have lower rates of female labor force participation even today.

⁶See, e.g., “What Happens to Women’s Ambitions in the Years After College,” in *The Atlantic*, December 2016. <https://www.theatlantic.com/business/archive/2016/12/ambition-interview/486479/>.

is coincident with motherhood or instead driven by trends.

The remainder of the paper is organized as follows. Section 2 introduces our main data sources. Section 3 describes the empirical strategy. Section 4 reports our results on the effect of parenthood on labor supply. Section 5 provides a variety of evidence that, in the short and medium run, women underestimate the employment effects of motherhood. Section 6 examines the accuracy of their longer-run expectations. Section 7 relies on both theoretical and empirical evidence to argue that the reason why modern women have not anticipated the costs of motherhood is because these costs recently increased in an unpredictable manner. Section 8 concludes and poses some questions for future work.

2 Data sources

In this section, we describe our main data sources and our sampling restrictions. A substantial portion of the empirical work involves event-study analysis. We thus adopt sample restrictions with that exercise in mind. Our preferred sample will include only those individuals whom we observe at least once before and once after the “event” of the birth of a first child, so that all subjects help identify a standard difference-in-difference estimate of the effect of parenthood (though in robustness checks we will add back in subjects who remain childless or had children before entering the sample, to serve as controls). We also, whenever possible, we require that we observe respondents over the full period of their most likely child-bearing years (between ages 20 and 40).⁷ Again, in robustness checks, we relax this restriction.

Below, we briefly describe our four main panel data sources. In the interest of space, we relegate greater detail on variable construction, weighting and sampling restrictions to Appendix B.

2.1 British Household Panel Survey

The British Household Panel Survey (BHPS) is a longitudinal survey that runs from 1991 through 2009. Col. (1) of Table 1 shows summary statistics for women in our preferred sample in the BHPS. The sample restrictions described above leave us with a total of 681 women. On average, in our analysis sample, we observe individuals 4.1 times pre-birth and 6.6 times

⁷The one exception is the BHPS. As the BHPS only spans 18 years (with a later additional booster sample spanning only about 10 years), imposing the 20-40 age is impossible. Instead of imposing the 20-40 restriction, we require our main sample of women to be at least 16 years old at the start of our panel. For our control samples of childless women, we require that we observe them at some point at age 29, the median age at first birth among women in our treatment sample.

after. While we do not impose perfect balance for our event-time figures, it is reassuring to know that we observe parents several times before and after their first birth so that trends are unlikely to be explained by composition changes, though we return to this concern in Section 4.2. While not our focus, Appendix Table A.1 provides summary statistics for the men in the BHPS as well for the men in our U.S. panel data sources.

2.2 U.S. panel data sources

We make use of three, commonly used sources of U.S. panel data: the National Longitudinal Surveys of Young and Mature Women (hereafter NLSW68), the National Longitudinal Survey of Youth (NLSY79), and the Panel Study of Income Dynamics (PSID).

We make use only of the “young women” component of the NLSW68.⁸ Women were 14-24 when first interviewed in 1968 and interviews continue through 2003. For the most part, we can only examine employment outcomes with these data, as there are no questions on gender-role attitudes that are asked in a consistent manner. However, we will make use of a variable that asks women in 1968 about their employment and parenthood expectations when they are 35. Our preferred sample includes only those women whom we observe at least once before and after the birth of their first child and who remain in the sample between ages twenty and forty. This sample includes only women from the 1948-1954 cohorts (so they are roughly in their sixties today). Our preferred sampling criteria leave us with 1,413 women, though, as noted above, we will drop some of our preferred sample restrictions in robustness checks, enlarging the sample.

Summary statistics are presented in col. (2) of Table 1 for our preferred NLSW68 sample. Given the longer time period covered by the NLSW68 (and, in fact, all our U.S. panel data sources), we have many more observations pre- and post-baby than in the BHPS.

The NLSY79 begins interviewing young men and women between the ages of 14 and 22 in 1979 and continues through 2014. While we occasionally will show results for men, we focus on the female half of the sample. Again, to ensure that no one in our analysis sample must have their first child either before the age of twenty or after the age of forty to meet our sample restrictions, we only include the 1959-1964 cohorts (so our sample is born roughly ten years after their NLSW68 counterparts). These sampling criteria leave us with 2,256 women. Summary statistics for this preferred NLSY79 sample appear in col. (3) of Table 1.

⁸While the National Longitudinal Surveys of Young and Mature Women include an older (“mature women”) sample, they have for the most part had their first child before entering the sample.

Our final U.S. panel data source is the PSID. While it begins in 1968 (and continues through today), our preferred analysis sample includes all individuals for whom we observe the first birth *after 1979* (because the employment outcomes we most use are not coded consistently before this date). To meet our usual requirement that we observe individuals between the ages of twenty and forty, our preferred sample includes only the 1957-1975 cohorts, leaving us with 1,007 women. Col. (4) of Table 1 displays summary statistics for our PSID sample. On average, the women in our sample are born about five (fifteen) years after those in the NLSY79 (NLSW68).

Due to slightly different definitions, we cannot compare all variables across the BHPS and our U.S. datasets, but where we can, the differences are as expected. Note that even the PSID, our most recent U.S. panel dataset, samples women who are from slightly earlier cohorts than the BHPS. Consistent with having later cohorts in the BHPS and the general tendency of the British to delay childbirth relative to Americans (recall Appendix Figure A.1), individuals in all our U.S. samples have their first child at a younger age than do women in the BHPS and also have more children in the period in which we observe them.⁹ Consistent with the evidence in Figure 2, the summary statistics show that American women are more likely to have completed university education than their British counterparts, despite coming from earlier cohorts. American women in the PSID and NLSY79 samples are also significantly more likely to have obtained a college degree than their counterparts in the NLSW68, consistent with the rapid growth in female college graduation rates in the US between the 1940s and 1970s cohorts (Goldin *et al.*, 2006).

3 Empirical strategy

Much of our analysis examining the short- and medium-run changes associated with parenthood makes use of a basic event-study methodology, defining the first child’s year of birth as the “event.” Specifically, we model a given outcome y_{it} (e.g., current employment, attitudes toward gender roles) for person i in year t as:

$$y_{it} = \sum_{\tau=-5}^{\tau=\tau^{max}} \beta_{\tau} \cdot \mathbb{1}[\tau = t - e^i] + \sum_a \gamma_a \cdot \text{Age}_{it}^a + \delta_t + \alpha_i + \epsilon_{it}. \quad (1)$$

We index event time (time relative to birth of a child) by τ . The variable $\mathbb{1}[\tau = t - e^i]$ is

⁹Our British panel is also shorter than our American panels, so the U.K. women are also less likely to have completed their fertility by the end of the panel.

defined as follows: e^i denotes the calendar year in which person i had their first child, so $\mathbb{1}[\tau = t - e^i]$ is an indicator for person i in year t having had their first child τ years ago. Negative values of τ indicate having a first child $|\tau|$ years in the future. In the summation term, we omit the year before the first birth, and the $\tau = -5$ (or $\tau = \tau^{max}$) term in fact includes all years greater than or equal to five years before (or τ^{max} years after) the first birth. We set τ^{max} to five or ten, depending on the length of the sample period in a given panel dataset. This specification normalizes the year $\tau = -1$ to zero, though in our figures we will often add to all coefficients the raw mean of the outcome variable at $\tau = -1$ to facilitate comparisons across different subgroups in both levels and changes. We control for a vector of individual fixed effects (α_i), a vector of calendar-year fixed effects (δ_t) and a vector of age-in-years fixed effects (Age_{it}^a). The error term is ϵ_{it} . We cluster standard errors at the individual level.

We present our results by plotting the β_τ coefficients, to show the within-person evolution of our outcome variables relative to the event of parenthood, conditional on year and age fixed effects. We also show difference-in-difference results, which we estimate using the same specification as in equation (1), but with the event-time year dummies replaced by a dummy variable, Post-Baby_{it} (indicating that in year t we observe person i *after* the birth of her first child).

4 How does motherhood affect employment?

We begin by showing how labor supply measures evolve before and after the event of parenthood, using both the BHPS and our U.S. panel data sources. For the main empirical results in the paper, we focus on event-study figures, as they provide more transparent and visual evidence on how outcomes evolve before and after parenthood (e.g., pre-trends, persistence) than do the difference-in-difference results. However, performing a comprehensive battery of robustness checks on each event-study graph would require exorbitant space, so in the Appendix we provide the difference-in-difference results corresponding to the event-study graphs and subject them to a battery of checks. In Appendix tables we also show how outcomes evolve after the birth of a *second* child (and whether we can distinguish this effect from that of the first child) as well as how *men's* outcomes react to the birth of a first child. In the interest of space, we only discuss some of these results, but interested readers can refer to the Appendix for a fuller treatment.

4.1 Main results

We begin our analysis by examining the share of women who are currently employed, before and after the birth of a first child. Figure 3 shows the event-study coefficients from equation 1, estimated on our U.K. data, which defines currently employed as having been at work last week. The year before the birth of a child, 87 percent of women are working (recall that the coefficient for $\tau = -1$ is merely the raw mean at $\tau = -1$), a rate only three percentage points lower than for men who will become fathers the next year (not shown). The year of the birth itself exhibits about a forty percentage-point decline in employment. Approximately five to ten percentage points of this decline are recovered in the following year, though no further recovery can be detected beyond that. When we instead collapse event time into a simple before and after the birth of the child in a difference-in-difference estimation (the estimate reported at the top of the figure), we find that the probability of being employed the week of the survey declines by 38 percentage points.

It is worth emphasizing that Figure 3 and our other event-study figures plot event-study coefficients (how employment changes relative to $\tau = -1$, conditional on all the controls in equation 1) and thus except for $\tau = -1$ raw means cannot be read off the figure. For example, for each of our four panel datasets, female employment is on a positive, secular trend, which we capture with the year fixed effects. So, in raw levels, women *do* increasingly return to work as time passes after the birth of their first child, but not any faster than the year fixed effects and other controls would predict. Put differently, Figure 3 implies that after the first child is born there is no catch-up *relative to* the counterfactual no-baby trend. For the sake of exposition we will often refer to an outcome as rising or declining upon motherhood. However, such a statement rarely refers to raw levels, but instead to changes conditional on the controls in the event-study specification.

We now replicate our U.K. employment results using our U.S. panel data sources. For each dataset, we use an extensive margin measure of labor supply that matches as closely as possible that in the BHPS, but see Appendix B for exact variable definitions. Figure 4 plots the event-study coefficients separately estimated in the NLSW68, the NLSY79 and the PSID. In their overall shape, the three event studies in Figure 4 are very similar to each other and to that from the BHPS. Employment reacts to motherhood more or less like a step-function, with little evidence of pre-trends and generally no sign of recovery in the ten-year post-period we observe.¹⁰ Note that given the longer sample periods covered in our U.S.

¹⁰With our data we cannot observe persistence in the *very* long run, and recent evidence suggests

data, we can extend the post-period an additional five years.

There are some differences in levels and magnitudes, though given that the exact definition of employment varies slightly among our four datasets, comparisons should be made cautiously. U.S. women are starting, in the year before parenthood, from a lower employment baseline than women in the BHPS. This difference seems mainly due to many more U.S. women being students in the years right before motherhood, consistent with their being younger at the time of first birth (and thus closer to their high school or early post-secondary years) and attending school longer than their British counterparts.¹¹ Among the U.S. data sources, ordering the results from the oldest to the most recent birth cohorts, we see the “mommy effect” decline somewhat over time: it is -0.42 in the NLSW68 (average birth cohort of 1951), -0.30 in the NLSY79 (average birth cohort of 1962) and -0.26 in the PSID (average birth cohort of 1967). While the decline between the NLSW68 and NLSY79 is large enough so that the difference is unlikely a mere artifact of employment definitions, we find it harder to interpret the smaller difference between the NLSY79 and the PSID. We take away that even for our most modern cohorts, the effect size remains large.

How do our results compare to existing work? Figures 3 and 4 very closely parallel the results in Kleven *et al.* (2018), which uses Danish register data drawing from the 1955-1973 birth cohorts. Like us, they find little evidence of pre-birth pre-trends, a large effect the year of the birth, and little if any long-run recovery. Our effect sizes, however, are much larger—Danish mothers are roughly twelve percentage points less likely to work than they were before motherhood, whereas the decline for British women is roughly forty percentage points and for American women twenty-five to forty percentage points, depending on the cohort. In general, Denmark has a much higher female labor force participation rate than the UK or US (and in fact, higher than essentially any other country), and comparing our results to those using Danish data suggests that a large part of the difference likely comes from women’s different responses to motherhood. Smaller effect sizes in Denmark may be explained by the country’s more generous parental-leave and child-care policy. In addition to that mothers may make up some of their lost prime-age employment by delaying retirement (see, e.g., Goldin and Katz, 2016).

¹¹The PSID has a lower $\tau = -1$ mean than the other two U.S. data sources. We suspect this difference is driven by two factors. First, the PSID makes individuals choose their *main* activity, so one cannot work and be a student at the same time, and indeed the share of women who are students the year before motherhood is substantially higher than in our other U.S. datasets. By contrast, the NLSY79 employment status variable simply indicates whether you currently have a job. Second and related, women in the PSID also have the most education, thus increasing the chances they are classified as students and not workers in the years right before motherhood.

covering 52 weeks of parental leave, the Danish government guarantees all children between the ages of six months and five years a spot in heavily subsidized child-care facilities. By contrast, U.K. and especially U.S. parental leave is much less generous, a point to which we briefly return in Section 7.

4.2 Robustness and other results

As discussed earlier, we present a standard set of robustness tests across our various datasets. These checks include adding, as controls, a person-specific linear time trend, dropping as controls the individual fixed effects, adding in a control group of childless adults, and including or excluding “agers-in,” children of original sample households who become eligible for interviews as they come of age.¹² In no case do these changes have an appreciable effect on the estimated difference-in-difference “mommy effect” for employment, as seen in Appendix Tables A.2 and A.3. Appendix Table A.4 also shows, as we would expect, that employment effects for men in the UK and US are very small and of mixed signs.

Appendix Table A.5 shows that, in all four datasets, the effect of a second child of maternal employment is statistically and economically significant, though also smaller in magnitude than the first (about one-fourth the size of the initial “mommy effect”). The table also shows that both effects can be separately distinguished from each other even after including a person-specific linear time trend.

Finally, one concern with the event-study graphs is that, as we do not impose that every observation helps identify every event-study coefficient, respondents who mostly help identify the pre-period coefficients (e.g., those who have children at older ages) may differ in their reactions to motherhood from those who mostly help identify post-period coefficients. As a check on this concern, for each dataset, we split respondents based on the share of their observations that come from the pre-period and run separate event studies for the two groups. We show results for the BHPS (which is our most imbalanced event study) in Appendix Figure A.2 and the shape of the two event studies are nearly identical. This same exercise repeated for our other datasets and outcomes also show little difference and are available upon request.

¹²While the control group of childless adults exhibit no variation in the *Post-baby* variable, they help identify the year and age fixed effects.

4.3 Heterogeneous employment responses to motherhood

So far, we have shown that women in the US and UK exhibit a sudden, large, persistent and robust drop in employment, coinciding with the birth of their first child. In this subsection, we examine whether certain person-level characteristics predict the magnitude of this response. We are especially interested in heterogeneity by characteristics that are often assumed to increase women’s attachment to the labor force (education, delaying motherhood, and having had a mother who herself worked) or relate to her expectations about the employment costs of motherhood (her pre-baby attitudes toward working mothers or her teenage expectations of whether she would work in her thirties).¹³

In the interest of space, we will generally examine heterogeneous responses to motherhood in a difference-in-differences as opposed to an event-study framework. That is, we take our difference-in-difference equation (based, again, on equation 1) and add an interaction between $Post\text{-}baby_{it}$ and the dummy, person-level variable of interest X_i (e.g., having had a working mother yourself). For each of our X variables of interest, the panels of Figure 5 show the point-estimate and confidence interval for the “mommy” effect for women with $X_i = 0$ (labeled “No” in the figure) versus $X_i = 1$ (labeled “Yes”) and report as well the difference and its level of significance.

The heterogeneous employment responses to motherhood along the dimensions we study are remarkably similar across datasets. In panel (a) of Figure 5, we allow the mommy effect to vary by education. In the BHPS, women with a college degree have a statistically significant but modestly smaller (roughly eight percentage points) employment decline than their less educated counterparts. In the NLSW68, the “mommy effect” is roughly six percentage points smaller for women with a college degree relative to their less educated peers; in the NLSY79 it is ten points smaller; and in the PSID thirteen points smaller. All of the differences in the U.S. data are statistically significant. However, it is worth emphasizing that in all four

¹³Examples of recent studies that examine the impact of these characteristics on women’s employment and wages include: Fitzenberger *et al.* (2013) and Wilde *et al.* (2010) (for education); Miller (2011), Herr (2015), and Herr (forthcoming) (for delaying motherhood); Olivetti *et al.* (2016) and McGinn *et al.* (2015) (for own-mother’s employment). While we do not focus on them as they do not directly relate to predicted labor-force attachment or expectations, we also examine differences by race (in the US only, as the share of non-whites is too small in the BHPS), marital status, and gender of the first-born child. Blacks have smaller mommy effects (mostly driven by lower pre-baby employment levels) than other respondents. The mommy effect does not differ by marital status or gender of the first child for any of our panel datasets. Given that roughly forty percent of our respondents are unmarried at the time of first birth, we do not focus on heterogeneity by characteristics of the partner.

datasets, even more educated women in the US and the UK exhibit very economically and statistically significant declines in employment upon motherhood.

In panel (b), we see that delaying childbirth (which we define in each dataset as having a first child at an age above the female median) offers no such protection, and in fact even the sign of the effect varies across data sources (though magnitudes are uniformly small and insignificant). In panel (c), we find similar “non-results” for having yourself had a working mother.¹⁴

While we can look at these three attributes in all four of our datasets, the remaining panels of Figure 5 do not include all four datasets due to differences in variable availability. For the BHPS, NLSW68 and NLSY79, we can stratify by pre-baby views about how women should divide their time between work and family, and for lack of better terminology we call those women who say women (and mothers) should work as having “pro-work” views. (Note that we use this gender-role variable as an outcome in the next section and describe it in more detail then.) To generate two groups, we split the sample based on the pre-baby median of this variable. Perhaps not surprising, women who have pre-motherhood stronger beliefs that women should work also exhibit smaller “mommy effects” on employment, though this result is not significant at conventional levels in the NLSW68. In the NLSW68 and the NLSY79 we can stratify observations by whether, in their teen years, women predicted they would be in the labor force at age 35. These women tend to have significantly smaller “mommy effects,” but even women who predicted they would work in their thirties are experiencing large declines in employment upon motherhood.¹⁵

5 Do women anticipate the “mommy effect” in the short-run?

Section 4 provided a variety of evidence showing the large, negative employment effects of motherhood. In the UK (US), women’s employment rates decline 38 (28–40) percentage points upon motherhood, with no evidence of recovery. To put this decline in perspective, in the most recent CPS, employment rates are 23 percentage points lower for 65-69 year-

¹⁴This result contrasts with Kleven *et al.* (2018), who find an intergenerational effect of working motherhood. We speculate that the key difference is that Kleven *et al.* (2018) are able to better control for family wealth during childhood. For example, in our data, having had a working mother could correlate with having grown up in a poorer family where the mother had to work. If poorer families instill more conservative gender-role attitudes in their children, then this omitted variable would dampen the observed intergenerational correlation of labor supply between mothers and daughters.

¹⁵We defer discussion of the “family-friendly” occupation results in panel (f) to Section 7.

olds than for 60-64 year-olds.¹⁶ Thus, our estimated “mommy effect” is comparable to a retirement effect.

While a large literature studies whether households fully anticipate and prepare for the effects of retirement, the question of whether women anticipate the employment effects of motherhood has received relatively less interest. As we will discuss in more detail in Section 7, dynamic labor supply models simply assume that women accurately anticipate these effects and rationally incorporate them into human capital and fertility decisions.¹⁷ This section provides several pieces of evidence suggesting that women, in the short- to medium-run, are systematically underestimating the effect of motherhood on employment. Section 6 takes up the question of the accuracy of their expectations in the longer run.

5.1 Effect of parenthood on views toward gender roles and female employment

In this section, we explore whether women’s beliefs about the proper balance between family responsibilities versus market work change upon having a child. If, as we argue, women underestimate the costs of motherhood, then the event of the first birth serves as an information shock and thus beliefs should be updated.

To measure how beliefs change, we make use of six questions in the BHPS. These questions are in the form of statements about the proper role of women in the workplace versus the home, and respondents are asked to state their agreement with each statement on a 1–5 scale (e.g., “All in all, family life suffers when the woman has a full time job,” “Both the husband and wife should contribute to the household income”). If needed, we reverse the order of the responses so that each question is increasing in the pro-female-employment direction. We then standardize each of these variables and sum them. Appendix Table A.6 provides summary statistics for this “omnibus” measure of summed standardized variables as well as each component (and the wording of each component statement). Note that, in contrast to the employment question, which is asked every year, the BHPS asks these six gender-role questions every other year, so we expect point-estimates to be noisier given the smaller

¹⁶See <https://www.bls.gov/cps/cpsaat03.pdf>.

¹⁷An important exception is Fernández (2013), who explicitly models how women form beliefs over the effect of working on family life, using signals they inherit from their own mothers as well as common signals embodied in the current female labor force participation rate. While we adopt a much more empirical approach, we share her view that it is legitimately hard for women to predict the costs of working post-motherhood. As she writes: “It is not an exaggeration to state that, throughout the last century, the consequences of women’s (market) work have been a subject of great contention and uncertainty.”

sample.

We refer to this omnibus variable as our *index of gender-roles regarding work and family*, or simply *gender-roles index*, for expositional ease. To get a sense of magnitudes, note that women are generally 0.77 units more in the pro-work (that is, agreeing that it is good for women to work) direction than are men.

This index has strengths and weaknesses. A strength is that all the questions refer to women’s roles as workers versus mothers and wives. One could imagine instead questions about gender stereotypes, but not about employment (e.g., “On average, women and men have equal innate intelligence”), which would be less germane to our hypothesis. Another strength is that the question is abstract and does not specifically ask about the respondent’s *own* household, as such a question would not lend itself to event-study analysis. The BHPS statements are about households *in general* and thus respondents without children can still answer these questions.

A weakness is one that arises for any type of subjective response: it is impossible to know exactly what the respondent means. But because statements in the BHPS are not about one’s own situation but about the appropriate responsibilities for men and women generally, in principle a woman who perfectly anticipates the costs of motherhood should give the same response before and after becoming a parent. By contrast, a woman who overestimated (underestimated) the employment costs of motherhood should change her answer in a pro- (anti-) work direction once she learns how easy (hard) motherhood actually is and thus how realistic or desirable it is for mothers to work while raising children. Our implicit assumption is that respondents are answering this question about households in general, but are using their own experiences to inform these answers.

5.1.1 Main results

Figure 6 shows the results of our event-study analysis of gender-role attitudes. The birth of the first child is associated with women moving 0.92 units in the *anti-work* direction on our omnibus gender-roles index. The effect is not only statistically but economically significant—in fact, it is slightly larger than the male-female difference. Like the employment effects, there is little if any pre-trend, though as expected the results are noisier, with larger standard errors. Unlike employment, the effects on our gender-roles index take about two years to be fully realized: attitudes become somewhat more anti-work the year of birth and even more so the year after. Beyond that point, they stabilize and show no evidence of recovery. Appendix

Figure A.3 shows no statistically significant effect of fatherhood, though, if anything, men also also adopt more anti-female-employment views upon parenthood as measured by our index.

As noted, this omnibus index is composed of six underlying questions. In Appendix Figure A.4 we show how women’s views evolve for each component of our omnibus measure. Of the six components, five show statistically significant movement in the anti-work direction and one (“pre-school child suffers if mother works”) shows an insignificant movement in the pro-work direction.¹⁸ Some of the individual components of the gender-roles index have a modest pre-trend, but we show in Appendix Table A.7 that all five of the outcomes with a significant “mommy” effect retain their sign and significance even after adding a person-specific linear trend.

We can in principle replicate this analysis in the NLSY79, though we emphasize an important caveat. The battery of gender-role questions are asked only four times: in 1979, 1982, 1987, and 2004. The fact that the questions span a long period means we can observe essentially the entire sample over their child-bearing years, as almost everyone in our NLSY79 sample who will have a child does so between 1979 and 2004. However, the infrequency of the questions means that the event study is highly imbalanced and there are relatively few (100-300) non-missing observations underlying each coefficient that we plot.¹⁹ The wording of and summary stats for the NLSY79 questions can be found in Appendix Table A.8. In general, the questions in the BHPS and NLSY79 are quite similar, and ask agreement with statements relating to two broad themes: (1) the roles a husband and a wife should adopt with regard to splitting paid work vs. housework and (2) the impact on children (and family) if a woman works. Fortunately for us, both sets of questions focus on *employment* in particular, rather than gender more broadly.

In the NLSY79, we find event-time coefficients that roughly support the results from the BHPS: that motherhood is associated with a decline in support for the idea that female employment does not detract from family life. Figure 7 shows these results. It is important to note that standard errors are substantially larger and that there is a non-trivial, negative pre-trend in the years before motherhood.

¹⁸We have no intuition for why this question moves in the opposite direction of the others, but do note it is the only question that refers to *pre-school* children, so is somewhat misaligned in terms of our event-study and difference-and-difference assumption that the event is the birth itself.

¹⁹Although the NLSW68 also contains similar questions, which we exploit in panel (f) of Figure 5, they are not asked in a consistent manner over time so we cannot replicate the event-study analysis in the NLSY79.

As with the employment results, we subject the gender-role results to our usual battery of robustness checks (adding a linear time trend, using additional controls group of non-parents, etc.). For each robustness check and for both datasets, the effect of the first child's birth on moving the mother in the anti-work direction remains statistically significant with a very similar point-estimate to the baseline result, as seen in Appendix Tables A.9 and A.10. The effect of a second child is an additional move in the same direction, large and highly significant in the BHPS. This effect is negative but insignificant in the NLSY79, though given the gender-role questions are only asked four times in the NLSY79, separately identifying the effect of both a first and second birth may be asking a lot of the data. The results for men are similar in that they move in the same direction as women do on the gender-role index after parenthood, but effects are smaller and not always statistically significant.

5.1.2 Heterogeneous effects

We examine heterogeneity in how attitudes change in response to motherhood along the same dimensions we earlier examined for employment. Panels (a) through (c) of Figure 8 show that, for both the BHPS and NLSY79, all of the factors that we hypothesized would make mothers remain attached to the labor force (advanced education, delaying childbirth, and having had a working mother), produce *larger* movements in the anti-work direction upon motherhood, though two of the six estimated differences are not significant at conventional levels. Similarly, panel (d) compares women with relatively more pro-work gender-role attitudes pre-motherhood to those with more anti-work views. Obviously, this regression is complicated to interpret because we are stratifying on the dependent variable, but we include it nonetheless to remain parallel to Figure 5. In fact, the entire anti-work shift in attitudes is driven by women who were more pro-work pre-motherhood. In the NLSY79, we can also stratify the sample on whether they predicted at age 18 that they would be working at age 35, and, again we see larger shifts for those who assumed they would be working.

Note that, for both datasets, patterns of heterogeneity for the employment results do *not* reflect the patterns for the gender-role-attitudes results. For employment, we found that education had the effect of blunting the mommy effect. By contrast, educated women appear the most affected by motherhood in terms of their answers to the gender-role index questions. The same pattern holds (i.e., smaller employment effects but in fact larger attitudinal shifts) for having less traditional views on the gender-role index pre-baby and (in the NLSY79) having predicted working at age 35. These conflicting patterns suggest that, importantly, the

gender-role results *cannot* be explained by simple models of cognitive dissonance, whereby individuals change their beliefs about the proper balance between family and work to match their own changing patterns of behavior.

5.2 Evidence from retrospective questions

We have so far shown robust evidence that in relatively modern cohorts in both the UK and the US, women adopt more traditional views about the appropriate balance between work and family responsibilities upon becoming mothers. We interpret these results as demonstrating that, pre-motherhood, women on average underestimate the employment costs of motherhood. This lack of anticipation appears concentrated among more educated women.

We can, in fact, look more directly at whether parents are underestimating the costs of child-rearing. The PSID asks subjects with children if parenthood is harder than they thought it would be. Obviously we cannot investigate how this question evolves in event time, as it is not asked of individuals *until* they become parents. We can, however, determine whether educated women are more likely to agree with this statement than their less educated counterparts, which would be consistent with our reading that their larger change in responses to the gender-role questions in the previous subsection represents a greater underestimate of the costs of motherhood.

We examine how both mothers and fathers answer this question, stratifying by education. Table 2 shows results from regressing, separately by gender, an indicator variable for parenting being “harder than I thought” on a college-degree indicator and a vector of standard covariates (which we vary to probe robustness). Before discussing the regression coefficients, it is interesting to note the large differences in the mean of the dependent variable by gender: whereas 52 percent of mothers with children age six and under report having underestimated how hard parenting would be, only 37 percent of fathers with similarly aged children respond the same way.²⁰ Most striking, in col. (1) we see that those mothers with a college degree are 12 percentage points more likely to agree than their less educated counterparts, an effect robust to adding controls (col. 2). It remains significant (though smaller) when we include mothers of older children (col. 3). For men (cols. 4–6), education is not a significant predictor

²⁰Respondents were asked to rank the truth of the statement “Being a parent is harder than I thought it would be” on a scale of (1) not at all true to (5) completely true. Responses for our regression sample are as follows: 16 percent of women and 22 percent of men select (1) “not at all true;” 11 percent of women and 21 percent of men select (2); 22 percent of women and 21 percent of men select (3), 18 percent of women and 20 percent of men select (4), and 33 percent of women and 17 percent of men select (5), “completely true.”

of agreement with this sentiment.

5.3 Do women correctly predict the effect of motherhood on labor force participation in the short-run?

In the previous subsection, we argued that, if women do not fully anticipate the costs of motherhood, then the birth of a child is an important informational shock, and we should see the updating of beliefs. And, indeed, we find evidence supporting this claim in both U.K. and U.S. data.

In this subsection, we more directly examine anticipation of employment effects, returning to our usual event-study framework and using as an outcome variable a BHPS question asking respondents whether they foresee giving up paid work in the next twelve months. *Unfortunately, the question is only asked of individuals who currently have a paid job.* Given the evidence in Figure 3 that women in the BHPS substantially reduce their labor supply upon motherhood and the evidence in Figure 5 that this reduction is mediated by covariates such as education, the sample who remains post-motherhood to answer this question is highly selected. Ideally, a question such as “Do you expect to be in paid work twelve months from now” would be posed to all respondents (so, for those without employment, would refer to finding or returning to a job). As such, we view these results as more suggestive and do not make them the core of our argument.

Sample-selection concerns notwithstanding, we conduct an event-study analysis using two related outcomes: the share of women who, in year t , will leave the labor market in year $t + 1$; and the share of women who, in year t , predicted they would leave the labor market in year $t + 1$.²¹ We include only women who work in year t , so both outcomes are affected by the same selection biases. Figure 9 shows that, pre-motherhood, these two series closely track each other and are close to zero: almost no working women predict they will drop out of the labor force the following year, and essentially none do. However, the two series start to diverge post-childbirth, with actual rates of giving up paid work increasing roughly 17 percentage points, while the share of women who *predict* they will leave increases by less than half this amount. Interestingly, mothers continue to overestimate the probability they continue working even several years after the birth of the first child, perhaps due to underestimating the effects of future children (recall that we find a second-child effect

²¹Note that for these purposes we define the *event* of child-birth as corresponding to the prediction in year $t - 1$, so the $t = 0$ observation corresponds to the prediction made in the year before childbirth ($t = -1$), realized in the year of childbirth.

for both employment and our gender-role index outcomes, Section 4.2 and Section 5.1.1, respectively) or underestimating the demands of the first child at older ages. Appendix Table A.11 reports the regression analogue of this figure, using a dependent variable that is equal to one if the respondent expected to work this year but was wrong, and subjects the result to our usual battery of robustness checks. For men (columns (7) and (8)), we find no statistically significant change post-baby.

In summary, post-motherhood, there is substantial increase in the likelihood of a working mother in year t dropping out of the labor force in year $t + 1$. However, less than half of these departures are predicted by the women even one year earlier. We emphasize that sample selection is an issue, but one point in defense of this exercise is that, given that a large share of women drop out of the labor force (and thus the universe for this question), we expect that we differentially observe in the post-period those women most attached to the labor force and thus who would be the most likely to remain in any given year post-motherhood. As such, we feel it is likely that sample selection leads us to understate the tendency of mothers who report that they will remain in the labor force when in fact they do not.

6 Do young women accurately forecast their future labor supply in the longer run?

In the previous section, we presented evidence that women underestimate the costs of motherhood in the relatively short-run. In this section we take a longer-run view. Specifically, we ask whether, at the time they are making key human-capital and early career decisions (i.e., their late high school and college years), young women are able to forecast their prime-age employment behavior.

6.1 Data and variable definitions

Some of the analysis in this section makes use of the NLSW68 and the NLSY79, and some relies on data we have yet to introduce. Summary statistics for these supplementary data sources can be found in Appendix Table A.12 and more details on the data can be found in Appendix B.5. We provide a brief overview below.

Monitoring the Future (MTF) is an ongoing study of the behaviors, attitudes, and values of a nationally representative sample of U.S. students. We examine twelfth-graders' responses to a question on the type of work they expect they will be doing at age thirty. Specifically,

the question asks, “*What kind of work do you think you will be doing when you are 30 years old? Mark the one that comes closest to what you expect to be doing.*” There are fifteen possible occupational categories from which students can choose, with “home-maker” as one possible option. See Appendix B.5 for a list of the occupations provided by the MTF for this question (as well as the list of occupations offered to respondents in the other data sources we describe in this subsection).

The NLSY79 asks respondents what they plan to be doing at age 35. Unlike in the MTF, respondents can only choose among four options: “working at current job,” “working at a different job,” “married, raising a family,” and “other.” Respondents who indicate that they plan to be “married, raising a family” are *then asked a follow-up question* about whether they also plan to be working outside the home at age 35. We define expected home-maker in the NLSY79 to be those who indicate planning to be “married and/or raising a family” *and* not working outside the home at age 35.

For years prior to 1977, we can in principle rely on the NLSW68. However, we are concerned that the NLSW68 may *overstate* the share being categorized as “home-maker,” since respondents are forced to choose between motherhood and work. Specifically, respondents are asked the *first* question from the NLSY79 but then those who choose “married, raising a family” are *not* asked the follow-up question on whether they also plan to work outside of the home.

We thus supplement the pre-1977 period with two additional, cross-sectional surveys: a 1961 Purdue Opinion Panel of high school students grades 9–12 and a 1961 Gallup opinion poll of young high school, college, and non students. The Purdue Opinion Panel asks “What kind of job do you expect to have 20 years from now?” and respondents can choose between twelve possible occupations, including “housewife.” Respondents in the Gallup poll were instead asked the open-ended question “What do you expect to be doing when you are 40 years old?” Responses were organized into 18 categories. We use the “house wife,” “home maker,” “house work,” and “raising children” category as our definition of expected homemaker. In both surveys, we restrict our data to female high school seniors, to match the MTF.

How do we compare the expectations of teenagers to their actual behavior once they are in their thirties? For the longitudinal surveys, we simply compare the expectations and realizations of the same group of individuals (so, we look only at those individuals who *both* answer the expectations question in their teens *and* report their realized employment

status in their mid-thirties).²² For the cross-sectional surveys (the MTF, and the Purdue and Gallup surveys), we compare respondents to their “synthetic selves” in the CPS 12-20 years later (this twelve-to-twenty-year range accounts for the fact that the MTF asks about expectations at age 30, twelve years after age 18; the NLSY79 at age 35, seventeen years after; the Purdue “twenty years later;” and the Gallup at age forty, twenty-two years later). That is, we compare answers regarding expectations at age 18 to realizations averaged over the ages 30 to 40.²³

6.2 Results

Figure 10 depicts several series of expectations and realizations. We describe these results chronologically. The earliest expectation results we have are from the Purdue and Gallup surveys, both from 1961. Reassuringly, results from these surveys agree with each other: fifty-five to sixty percent of women in both surveys predict that in their late thirties they will be home-makers. In general, these shares are quite similar to those based on eighteen-year-olds in the NLSW68 from later in the decade. Consistent with Goldin (1990) and Goldin (2006), we find that, in the NLSW68, roughly two-thirds of young women in 1968 expect to be home-makers at age 35, with this share falling over the following two years.²⁴

Comparing these expectations to the realized labor supply of women twelve to twenty years later, we see that young women in the 1960s vastly underestimated their future labor supply (again consistent with Goldin’s earlier work). Comparing their expectations to the behavior of their future, synthetic selves in the CPS, women in the Purdue and Gallup surveys overestimate the probability they would be housewives by about fifty percent. As the actual housewife share declines somewhat rapidly over the next ten years, the women in the NLSW68 in fact make more severe (by a factor of two) overestimates of their future housewife status. (This result holds whether we compare the NLSW68 women to their future selves in later waves of the study or to their synthetic selves in the CPS.)

²²In fact, the MTF does appear to include a longitudinal component, but the data are not shared and in fact we can find few citations that make use of these data. Not even codebooks appear to be published.

²³In previous versions, we matched each expectation to the exact future age to which it referred. Results are nearly identical, though far more cluttered visually.

²⁴Our panel data sources, the NLSW68 and the NLSY79, ask its expectations questions several times. The points plotted in Figure 10, however, use only the data from the year a respondent turns 18, to match the MTF. A small concern is that some of these 18 year-olds will have seen the question in previous waves of the survey. Readers concerned with this potential anchoring bias can simply ignore the points after 1968 and after 1979 for the NLSW68 and the NLSY79, respectively.

Sometime between the high school class of 1968 and 1978, young women’s tendency to underestimate their future labor supply disappears and in fact reverses. The share of the MTF class of 1978 who predicts being a housewife at age thirty is roughly ten percent, a dramatic decline from their counterparts in the 1960s and early 1970s. Moreover, while the 1960s high school classes were *overestimating* the probability of being a housewife, from 1978 onward, every high school class in our data underestimates it. For the late 1970s high school classes (when we can compare the MTF and the NLSY79 as their samples overlap in terms of birth cohorts) the magnitude of the underestimate is very similar whether we compare MTF respondents to their synthetic selves in the CPS or the NLSY79 respondents to their own future selves: around ten percent of respondents predict they will be housewives, but in reality just under twenty percent will be. Since about 1990, less than two percent of our MTF sample predicts they will be housewives at age thirty, whereas the actual housewife share in the CPS has not declined appreciably since 1980 (it has bounced between 15 and 18 percent).

One natural explanation is that more recent cohorts of young women are underestimating their *fertility*—perhaps they do not expect to have children (or expect to have fewer than they actually will). That is, instead of underestimating the employment costs of motherhood, they simply underestimate the probability of motherhood itself. In fact, Appendix Figure A.5 shows that, if anything, women in the MTF overestimate their future fertility. As such, they plan on being mothers, but, at the same time, do not connect fertility with home-making.

Taken together, these results indicate that there has been a sharp reversal in the accuracy of young women’s predictions of their future labor supply. Throughout the 1960s and early 1970s, we find that young women are largely *overestimating* the probability that they will be a home-maker. From 1978 through the early 21st century, we instead find that young women are *increasingly underestimating* this probability. As such, since the late 1970s, women who are on the cusp of making key human capital investment decisions appear to substantially overestimate their future attachment to the labor market.

It would be interesting to examine whether women’s expectations become more accurate in their twenties, when many of them will be in college, post-baccalaureate studies, or making early career decisions. While we know of no dataset that would allow quantitative analysis of this question, recent ethnographic evidence suggests that, at least for high-achieving women in college, expectations remain quite poor. Schank and Wallace (2018) extensively interview 43 of their Northwestern sorority sisters who graduated, like they did, in the early 1990s.

These women are surely not representative of all U.S. women, but they reflect high-achieving women who have made sizable investments in education.²⁵ Whereas all these women anticipated high-powered careers while in college and most left college for graduate degrees or jobs with long hours, only eight of the thirty-four who became mothers maintained full-time jobs. Fifteen moved to flexible or part-time work and eleven fully dropped out of the labor force.

Almost all the women who moved to flexible work or dropped out express surprise at the outcome. “I wasn’t planning on staying home with my son after he was born,” one mother said, “but once he was born, I realized something had to give and I needed to figure out what.” Said another: “I never wanted to stay home with my kids, ever,” but six weeks after the birth of her first child, the woman decided not to return to work. As the authors write, *only two* of the eleven mothers who dropped out of the labor force planned to do so before the actual birth.

In summary, this section has shown a variety of evidence consistent with the claim that, during the years they make their key human capital decisions, women in the US and the UK do not fully anticipate the effects of motherhood on employment. Given how large the effect remains today and how common an occurrence motherhood is, our results so far raise the question of *why* women would fail to anticipate these effects. We now turn to this question.

7 Why do women today underestimate the effects of motherhood?

Motherhood is associated with large and lasting declines in employment. What is perhaps most puzzling is that these large employment effects occur for cohorts of women who are far more prepared for the labor market than previous cohorts in terms of education levels (substantially higher than men’s) and job experience (substantially higher than in previous cohorts, due to delayed childbirth). The large increase in investments are not surprising, as they are consistent with an enormous change in *expectations* about future labor supply, as we demonstrate in Section 6. And yet when motherhood arrives, these women exhibit substantial and lasting employment declines.

We hypothesize that the employment costs of motherhood are higher for current mothers

²⁵As the authors point out, sororities at Northwestern are quite unlike typical sororities at other U.S. universities with “no hazing and a lot of studying.” Women joined sororities at Northwestern to make friends and to gain a sense of belonging in a large school. As a result, the women in the sample were economically diverse and academically accomplished.

than they would have predicted by projecting from past trends. Such an explanation has a number of attractive features. It explains why young women now (but not in earlier cohorts) overestimate the likelihood they will work at age thirty. It also explains why they have invested so much in education and have delayed childbirth despite, since roughly 1990, seeing no increase in labor force attachment (because they were making education decisions under the mistaken assumption of high, future labor force attachment, due to their underestimation of motherhood costs).

We begin this section with a simple model of women’s educational investment decisions in the face of uncertain employment costs of motherhood. To bolster the credibility of this model, we first show that it yields, under quite general conditions, two of the more striking results from the paper: that, compared to their less educated counterparts, college-educated women exhibit smaller *employment* effects of motherhood (recall Section 4.3, where we show that this result holds across all four of our datasets), but at the same time they are also more “surprised” by how hard motherhood actually is (recall Section 5.1, where we show their attitudes change more in the anti-work direction on work-family balance questions in both the BHPS and the NLSY79 and Section 5.2 where they are more likely to say in the PSID that parenthood is harder than they anticipated). *We show that this result holds whether or not the average cost of motherhood has increased or decreased relative to the previous generation.*

We then show that the model predicts some of our other key results—namely, that women *on average* update their attitudes in the anti-work direction upon motherhood and report that parenthood is harder than they expected —*only* when the costs of motherhood have increased on average across generations.

We complement this theory-based argument with a collage of empirical evidence that while the cost of motherhood fell during much of the twentieth century, by some important metrics, it has recently risen. Some of this evidence we take from existing literature and some, to the best of our knowledge, we provide here for the first time.

7.1 Modeling women’s education and employment as a function of expected employment costs of motherhood

There have been several dynamic labor supply models from the macro literature that attempt to jointly explain women’s fertility, human capital and employment decisions (or some subset of these decisions). Most of these models (e.g., Attanasio *et al.*, 2008; Blundell *et al.*, 2016;

Adda *et al.*, 2017) assume that agents have perfect foresight of how children will affect labor supply. We present a much simpler, static model in which women predict the employment costs associated with their future children based on their observations of their own mother. In particular, when they are making human capital decisions, they can over- or underestimate these future costs (which leads to under- and overestimates, respectively, of their future labor supply). In this sense, our model takes Fernández (2013) as inspiration, but imposes more functional form assumptions and is static (focusing on a single generation whose choices are impacted by the experiences of the previous generation).

7.1.1 Assumptions and set-up

Let utility $u(c, h)$ be a quasi-linear function of consumption c and labor h (for hours worked, say). Specifically, assume that

$$u(c, h) = c - \frac{h^{\gamma+1}}{\gamma + 1},$$

where $\gamma > 0$.

Women’s consumption will be equal to market wages net of employment costs of motherhood (both per hour) times labor supply (in hours). Market wages are $\tilde{w} = w + \beta \cdot e$, where w is a base wage for those with higher education $e = 0$ and $\beta > 0$ is the hourly return to higher education $e = 1$. Gaining education $e = 1$ costs some “tuition” $\alpha > 0$, while $e = 0$ is free.

We view “employment costs of motherhood” very broadly, as any cost mothers endure during work. Recall, from the introduction, that these might include the per-hour cost of a nanny or day-care service, the emotional or psychic cost of being separated from the child while at work, the sense of guilt at underperforming as a mother or colleague, or simply lack of sleep due to working while performing child-care activities. Actual employment costs for a given woman i are μ_i , but she predicts this cost with some error δ_i . In particular, she observes her mother’s employment costs, $\mu_i + \delta_i$, but she only inherits the μ component. The δ component was her mother’s “luck” (good or bad) and is *not* passed down to the daughter. This luck is the sum of a generational component and an idiosyncratic component. For example, her own mother’s idiosyncratic good luck might include having had a very understanding supervisor at work, or having parents or in-laws who live in close proximity and can provide free child care. Generational “good luck” could be, for example, giving birth in a moment with cheap, technological substitutes for child care. But the woman in our

problem cannot distinguish between μ_i and δ_i and instead uses the sum as the best signal of her own, future employment costs of motherhood.²⁶

Assume further that $\mu \sim U[0, 1]$ and $\delta = \{\lambda - \epsilon, \lambda + \epsilon\}$ with equal probability, where λ and $\epsilon > 0$ are both constants. Note that λ can take negative values (for much of the analysis below, λ will drop out, but allowing $\mathbb{E}(\delta) \neq 0$ will be useful later). We assume μ and δ are independent. For simplicity, we further assume no taxes and that w is sufficiently large that $w - \mu - \delta$ is always greater than zero (so mothers' predicted effective wages are always positive).

We abstract away from many possible dimensions of heterogeneity. There is no variation in fertility in our model: women become mothers with probability one and we ignore issues related to delaying or timing childbirth. Unlike many models of human capital, we ignore any variation in “ability” and thus the effective return to and cost of education do not vary along this dimension.

The implicit timing of the model is as follows: A woman makes her education decisions assuming that her future employment costs of motherhood are $\mu + \delta$, but then will make her actual labor supply decisions after this uncertainty is resolved and her actual costs, μ , are revealed to her. While this timing is helpful to keep in mind, in fact the problem can be collapsed to a single decision: the woman's (only) problem is to optimally choose education $e \in \{0, 1\}$. Once she chooses e and once her true costs μ are revealed to her, her labor supply is given by a simple optimization problem and thus can be viewed as deterministic. Note that, as we are assuming away many other dimensions of heterogeneity, variation in optimal e comes entirely from variation in expected employment costs, $\mu_i + \delta_i$.

Working backward, the woman calculates her predicted utility for each $e \in \{0, 1\}$. This comparison requires calculating the optimal predicted labor supply \hat{h}^* , given her assumptions about employment costs, for each e and then determining which value for e yields higher utility. We describe key characteristics of both the predicted labor supply \hat{h}^* and actual labor supply h^* in the Appendix.

7.2 Predictions yielded by the model

Claim 1. *Women who choose $e = 1$ will on average work more (post-baby).*

²⁶Note also that women in our model are *naive* in the sense that they do not take into account that δ is drawn from a distribution and instead just assume employment costs will be $\mu + \delta$ with probability one.

Proof. See Appendix. ■

The intuition is relatively simple. Women who choose $e = 1$ have higher wages, and because the utility function is quasi-linear, there is no income effect of higher wages on labor supply. While some women with $e = 0$ will supply more labor than some women with $e = 1$ (because some women with $e = 1$ will have much higher employment costs than they had expected), on average those with $e = 1$ work more.

Note that to connect this result directly to the empirical work, we need to add slightly more detail to the timing of the model. We have effectively collapsed the model into a single period (women make their education decision based on expected costs of motherhood, and then instantaneously have a child and have their costs revealed to them). If we instead assume that women all have a (perhaps brief) pre-baby period where everyone works some max “full-time” hours (which is roughly what we see in our data), then the claim maps not only into educated women having higher post-baby labor supply, but also exhibiting smaller employment “mommy effects.”

The next result states that, even though more educated women have smaller mommy effects with respect to employment, they are nonetheless the most “surprised” by the demands of motherhood.

Claim 2. *Women with $e = 1$ underestimate the costs of motherhood more than women with $e = 0$. That is, $\mathbb{E}[\delta \mid e = 1] < \mathbb{E}[\delta \mid e = 0]$.*

Proof. See Appendix. ■

Recall that a “smaller” (“less positive” or “more negative”) value for δ means that actual employment costs μ will be *greater* than expected costs $\mu + \delta$, thus leading to “surprise” at how hard child-rearing is. Intuitively, women who draw a low δ will, all else equal, be more likely to invest in $e = 1$, because they assume low employment costs of motherhood and thus high future labor supply.

As is emphasized in the proofs, Claims 1 and 2 hold for any value of λ —that is, they hold if the previous generation had on average lower, higher or the same costs of motherhood relative to the current generation. The model thus yields our results on heterogeneity by education under fairly general assumptions, which we hope bolsters the credibility of the model for readers.

Now, we show that our *unconditional* results on belief-updating (i.e., that, *on average*, women’s gender-role attitudes move in the anti-female-employment direction after baby

and that they report parenthood being harder than they expected) hold *if and only if* the employment costs of motherhood have increased for the current generation relative to the previous.

Claim 3. $\mathbb{E}[\mu + \delta] < \mathbb{E}[\mu]$ *iff* $\lambda < 0$.

Proof. See appendix. ■

In the next subsection, we discuss empirical evidence in support of our claim that maternal employment costs have on average increased. In fact, we describe suggestive evidence for a stronger claim: that while they have decreased over much of the twentieth century, they have increased recently, which would likely have made it even harder for the current cohort of women to have predicted this evolution. Relative to our event-study evidence, the evidence in the subsection below is much less systematic and we view it more as a “collage of evidence,” taken from a variety of sources. We see further work in this area to be an excellent path for future research.

7.3 Empirical evidence on the evolution of the costs of motherhood across cohorts

Several macroeconomists have modeled the *rise* of working mothers over much of the twentieth century as a function of *declining* costs of motherhood. Greenwood *et al.* (2005) argues that electrification liberated women from household chores, chronicling the dramatic decline in time cost associated with tasks such as laundry and cooking. Albanesi and Olivetti (2016) focus more on the biological aspects of motherhood. They show that the health costs of pregnancy and delivery have fallen dramatically since the 1920s, due to pharmaceutical innovations such as sulfa drugs and later antibiotics as well as the modernization of obstetric practices, all of which dramatically reduced mortality and morbidity associated with childbirth. Moreover, the quality of infant formula improved substantially over this time, minimizing the need for mothers to breastfeed.

Maternal employment remains a sensitive subject, suggesting that stigma costs play a role in mothers’ decisions to work. Fernández *et al.* (2004) argue that World War II served as a negative shock to this stigma cost, as mothers were told it was their patriotic duty to work, and that this change in norms may have passed on to future generations.

Some recent work suggests that any decline in the cost of motherhood may have stalled or even reversed. Ramey and Ramey (2010) documents a dramatic increase since the mid-1990s of the time-cost associated with motherhood in the US, especially for educated women

(increases for fathers are similar proportionally, but from a substantially smaller base), trends that we reproduce in Figure 11. As late as 1994, mothers aged 25-34 spent roughly 11 hours a week on child-care activities, with little variation by mothers' education. By the early 2000s, college educated mothers spent 20-22 hours, while their less educated counterparts had increased their total to 16 hours.²⁷

A major time cost associated with infants is breastfeeding. In a reversal of the trends highlighted by Albanesi and Olivetti (2016), mothers today are encouraged to breastfeed exclusively for at least six months and formula is deemed a distant, second-best option.²⁸ Indeed, mothers appear to follow this advice, at least in part. Figure 11 shows that breastfeeding rates decline among U.S. mothers in the 1950s and 1960s (we are unable to find U.K. data from this early period). By 1965, only twenty percent of U.S. mothers report ever having breastfed. But then in the mid-1970s, in both the US and the UK, the share of mothers nursing increases, before temporarily plateauing in the UK and modestly declining in the US during the 1980s. Since 1990, the share of mothers who report nursing their infants has exhibited a steady increase in both countries, especially in the US. Indeed, since 1990, this share has increased from approximately fifty to eighty percent. While we focus on “breastfeeding at all” because it is the measure we can best compare across time and countries, Appendix Figure A.6 shows that alternative measures follow the same *U*-shaped pattern over time.

Finally, we note that the cost of one of the most common substitutes for maternal inputs into child-rearing—professional child-care services—has grown steadily in real terms since the early 1980s. The final series in Figure 11 documents a 65 percent increase from the early 1980s to today.

Do women themselves perceive that the costs of motherhood have risen since they were themselves being parented? It is certainly possible that the increased costs we highlight above (time costs of child-rearing, prevalence of nursing, and monetary cost of child-care services)

²⁷Borra and Sevilla (2015) document similar patterns in U.K. data.

²⁸The U.K. National Health Services (NHS) recommend: “Exclusive breastfeeding (breast milk only) is recommended for around the first six months of your baby’s life.” Moreover, the NHS lists as a “myth” the claim that formula provides comparable benefits to breastmilk. The American Academy of Pediatricians official policy statement on breastfeeding reads: “The American Academy of Pediatrics reaffirms its recommendation of exclusive breastfeeding for about six months, followed by continued breastfeeding as complementary foods are introduced, with continuation of breastfeeding for 1 year or longer as mutually desired by mother and infant.” See <https://www.nhs.uk/conditions/pregnancy-and-baby/your-breastfeeding-questions/> and <http://pediatrics.aappublications.org/content/129/3/e827.full#content-block>, respectively.

are offset in other ways. In fact, recent polls have asked women to make this intergenerational comparison.²⁹ Over 66 percent of current mothers say that being a mother today is harder than it was when they were children. Of course, mothers today are more likely to work than their own mothers were, and thus perhaps they feel they have it harder because they are being asked to both work and mother. But on the specific question of the demands of child-rearing, 56 percent of current mothers say they are more involved in their child’s life than their own mothers were in theirs.

7.4 Is it harder to be a working mother today because workplaces have grown less family friendly?

We have so far provided evidence that the family component of work-family balance has become more difficult for mothers. But perhaps (instead or additionally) workplaces have become more hostile to mothers. While we have been implicitly assuming that maternal employment costs are driven by the difficulty of finding substitutes for maternal child care, it could instead be the case that employment costs are driving by workplaces becoming hostile to women once they become mothers.

In many ways, workplaces are more family-friendly than in the past. In 1975, U.K. maternity leave policy entitled eligible women to twenty-nine weeks of leave, of which only six were paid (Gregg *et al.*, 2007). Since then, eligibility has expanded, as has the length of leave, currently at fifty-two weeks, of which twenty are paid (Glassdoor, 2016). In 1999, the government expanded subsidies and provision of daycare (Gregg *et al.*, 2007). By contrast, the US has much less generous support of working mothers, though it, too, has increased over time. At the national level, the Family and Medical Leave Act, passed in 1993, mandates that employers provide eligible workers with twelve weeks of job-protected *unpaid* leave (previously there had been no federal mandate for any amount of leave). Growth in paid parental leave has been driven primarily by private provision and state-level policies, with nine states currently having programs to fund paid leave (Gault *et al.*, 2014). Today, about fourteen percent of the civilian U.S. workforce has access to paid family leave (which includes maternity leave).³⁰

Of course, there are important countervailing trends. Goldin (2014) documents the rise

²⁹Specifically, we pool data from the following surveys: the ABC News Mother’s Day Poll in 2006 and the CBS Monthly National Poll from April of 2009, 2010, 2011, and 2013.

³⁰<http://www.pewresearch.org/fact-tank/2017/03/23/access-to-paid-family-leave-varies-widely-across-employers-industries/>.

in the returns paid to long hours. Since women remain the chief providers of child care and tend to work fewer hours in the labor market, they are especially penalized in occupations that are characterized by these high rewards to working long hours, many of which are highly remunerated.³¹

While we cannot speak directly to whether workplaces on net have become more or less hostile to working mothers, we provide some suggestive evidence that the “work” component of work-family balance does not seem to be driving the employment “mommy effect.”

The BHPS asks all current workers about their overall satisfaction with their job, as well as satisfaction on a number of specific aspects of their work (e.g., hours, prospects for promotion, pay). Of course, there is serious sample selection into who can answer this question because it is only asked of those who stay in the labor force (though our analysis includes person fixed effects, mitigating some concerns about composition bias), so we view this analysis as only suggestive. That (important) concern notwithstanding, we show in Figure 12 that women exhibit no decline in overall job satisfaction post-baby. Appendix Figure A.7 shows that, in fact, for two of the five components (total pay and hours worked), satisfaction increases significantly and for none does it decrease significantly.³²

As a second piece of suggestive evidence, we return to the final panel of Figure 5. Here, we have divided women into whether their pre-baby occupations were “family-friendly,” which we define as having a pre-baby occupation with an above the median share of women. In fact, we find no difference in the magnitude of the mommy effect on employment.³³ Again,

³¹Cha and Weeden (2014) argue that the rising returns to long hours, coupled with the gender gap in the propensity to work overtime, plays a role in accounting for the slow convergence of the gender wage gap from 1979 to 2009. Indeed, there is abundant empirical evidence that women, particularly those with young children, place a higher value on flexibility at work and shorter work hours relative to men (Flabbi and Moro, 2012, Mas and Pallais, 2017 and Wiswall and Zafar, 2017). Not surprisingly, women respond to the challenges of combining career and family by disproportionately sorting into family-friendly firms or public sector jobs after the birth of their first child (Hotz *et al.*, 2017 and Pertold-Gebicka *et al.*, 2016) and making occupational choices with these considerations in mind (Wasserman, 2018).

³²While again only suggestive, ethnographic evidence supports our interpretation that it is not the work component of work-family balance that is causing motherhood to be “harder” than expected. After describing a long list of tasks associated with child care, one of the working mothers interviewed by Schank and Wallace (2018) says: “Then I go to work, where I can rest.”

³³We found this result surprising, and thus examined several different labor-supply outcomes (e.g., continuous hours instead of a binary employment indicator) and alternative definitions of “family-friendly” (e.g., percent of women in the occupation, percent of moms in the occupation, and percent of moms with children under age five in the occupation). In none of more than a dozen specifications did we find that having pre-parenthood a family-friendly occupation softens the employment effect of motherhood.

while this evidence is indirect, it suggests that the “work” component of work-life balance is not the driving force behind the mommy effect.

8 Conclusion

This paper contributes evidence to a growing literature showing that, despite huge increases in women’s human capital investment, there remains a large and lasting negative effect of motherhood on labor-force attachment. We view our chief contribution, however, as demonstrating that women do not appear to anticipate this “mommy effect” on employment. They now substantially overestimate the likelihood they will be in the labor force during their chief child-rearing years, suggesting that they may be making education and other human capital investment decisions under mistaken assumptions.

We rationalize this behavior with a simple model. One implication of the model is that the costs of motherhood have risen in a manner today’s cohort of mothers was likely unable to predict. In the final part of the paper, we add further evidence to recent work suggesting that such an increase, likely starting in the 1990s, indeed took place.

We offer two suggestions for future work. First, if current cohorts of women indeed overestimate their future labor supply (as we claim), are they in fact *over-investing* in education? Chen and Chevalier (2012) argue that female medical students systematically over-invest in education: given that female doctors work significantly fewer hours than male doctors, they argue that, from a purely pecuniary perspective, these students should have instead become physicians’ assistants. Other papers have instead examined the *social* return of the rise of women’s educational attainment. For example, Currie and Moretti (2003) show that exogenous increases in educational opportunities for adolescent women correlate with improved birth outcomes when these women become mothers. Yet other papers have emphasized the marriage market returns to women’s education (Lefgren and McIntyre, 2006; Oreopoulos and Salvanes, 2011; Kaufmann *et al.*, 2015; Bruze, 2015), though it is important to note that these returns are private, not social, returns, as they are likely zero-sum. Related, women may invest in education as insurance against the risk of divorce, which had been rising during the second half of the twentieth century (Bronson, 2014; DiPrete and Buchmann, 2006). Our work suggests that, unless the employment effects of motherhood decline in the future, non-labor market returns to female education are perhaps just as important to study as labor-market returns.

Second, throughout this paper we have taken as given that the employment costs of

motherhood might change over time due to technological developments or changes in norms, but future work might focus on *why* they change.³⁴ Especially in the recent period, changes in maternal employment costs seem more driven by changing expectations and norms as opposed to changes in technology. On face, it is deeply puzzling that at a moment when women are more prepared than ever for long careers in the labor market, norms would change in a manner that encourages them to spend more time at home.

References

- ADDA, J., DUSTMANN, C. and STEVENS, K. (2017). The career costs of children. *Journal of Political Economy*, **125** (2), 293–337.
- ALBANESI, S. and OLIVETTI, C. (2016). Gender roles and medical progress. *Journal of Political Economy*, **124** (3), 650–695.
- ALESINA, A., GIULIANO, P. and NUNN, N. (2013). On the origins of gender roles: Women and the plough. *The Quarterly Journal of Economics*, **128** (2), 469–530.
- ANGELOV, N., JOHANSSON, P. and LINDAHL, E. (2016). Parenthood and the gender gap in pay. *Journal of Labor Economics*, **34** (3), 545–579.
- ANGRIST, J. and EVANS, W. (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *American Economic Review*, **88** (3), 450–77.
- ATTANASIO, O., LOW, H. and SÁNCHEZ-MARCOS, V. (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review*, **98** (4), 1517–52.
- BAXTER, J., BUCHLER, S., PERALES, F. and WESTERN, M. (2015). A life-changing event: First births and men's and women's attitudes to mothering and gender divisions of labor. *Social Forces*, **93** (3), 989–1014.
- BERTRAND, M. (2013). Career, family, and the well-being of college-educated women. *American Economic Review*, **103** (3), 244–50.
- , GOLDIN, C. and KATZ, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, **2** (3), 228–55.
- , KAMENICA, E. and PAN, J. (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics*, **130** (2), 571–614.
- BLAU, F. D. and KAHN, L. M. (2013). Female labor supply: Why is the United States falling behind? *The American Economic Review*, **103** (3), 251–256.
- BLUNDELL, R., COSTA DIAS, M., MEGHIR, C. and SHAW, J. (2016). Female labor supply, human capital, and welfare reform. *Econometrica*, **84** (5), 1705–1753.

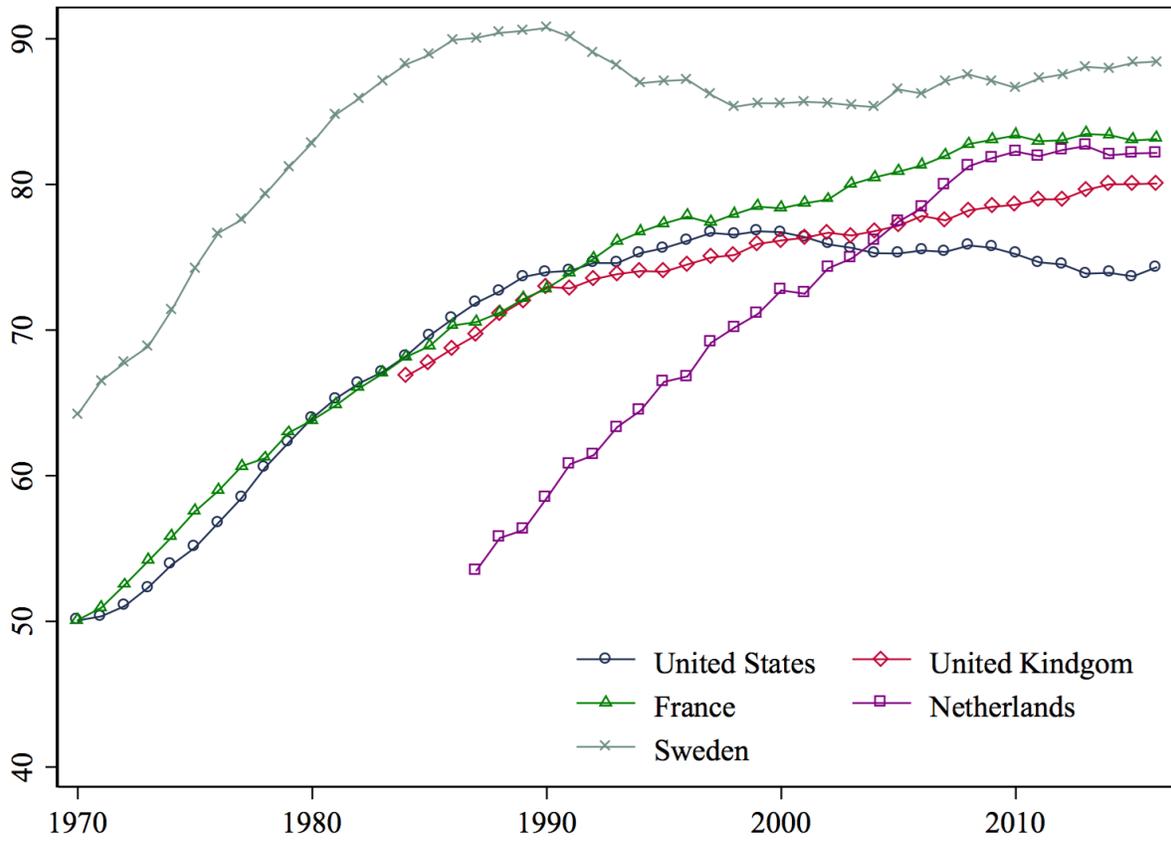
³⁴A small number of recent papers suggest that the rise in involved parenting (and increase in time spent on children), particularly among educated parents, could be due to factors such as the increase in competition for slots in top colleges (Ramey and Ramey, 2010) and increases in inequality (Doepke and Zilibotti, 2017).

- BORRA, C. and SEVILLA, A. (2015). *Parental Time Investments in Children: The Role of Competition for University Places in the UK*. Tech. rep., IZA Discussion Paper.
- BRONSON, M. A. (2014). Degrees are forever: Marriage, educational investment, and life-cycle labor decisions of men and women.
- BRUZE, G. (2015). Male and female marriage returns to schooling. *International Economic Review*, **56** (1), 207–234.
- BURSZTYN, L., FUJIWARA, T. and PALLAIS, A. (2017). ‘Acting wife’: Marriage market incentives and labor market investments. *American Economic Review*, **107** (11), 3288–3319.
- CHA, Y. and WEEDEN, K. A. (2014). Overwork and the slow convergence in the gender gap in wages. *American Sociological Review*, **79** (3), 457–484.
- CHEN, M. K. and CHEVALIER, J. A. (2012). Are women overinvesting in education? evidence from the medical profession. *Journal of Human Capital*, **6** (2), 124–149.
- CHUNG, Y., DOWNS, B., SANDLER, D. H. and SIENKIEWICZ, R. (2017). *The Parental Gender Earnings Gap in the United States*. Tech. rep., U.S. Census Bureau Working Paper, CES 17-68.
- CURRIE, J. and MORETTI, E. (2003). Mother’s education and the intergenerational transmission of human capital: Evidence from college openings. *The Quarterly journal of economics*, **118** (4), 1495–1532.
- DIPRETE, T. A. and BUCHMANN, C. (2006). Gender-specific trends in the value of education and the emerging gender gap in college completion. *Demography*, **43** (1), 1–24.
- DOEPKE, M. and ZILIBOTTI, F. (2017). Parenting with style: Altruism and paternalism in intergenerational preference transmission. *Econometrica*, **85** (5), 1331–1371.
- FERNÁNDEZ, R. (2013). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review*, **103** (1), 472–500.
- FERNÁNDEZ, R., FOGLI, A. and OLIVETTI, C. (2004). Mothers and sons: Preference formation and female labor force dynamics. *The Quarterly Journal of Economics*, **119** (4), 1249–1299.
- FITZENBERGER, B., SOMMERFELD, K. and STEFFES, S. (2013). Causal effects on employment after first birth – a dynamic treatment approach. *Labour Economics*, **25**, 49–62.
- FLABBI, L. and MORO, A. (2012). The effect of job flexibility on female labor market outcomes: Estimates from a search and bargaining model. *Journal of Econometrics*, **168** (1), 81–95.
- GAULT, B., HARTMANN, H., HEGEWISCH, A., MILLI, J. and REICHLIN, L. (2014). *Paid Parental Leave in the United States: What the Data Tell Us about Access, Usage, and Economic and Health Benefits*. Tech. rep., Washington, D.C.: Institute for Women’s Policy Research.
- GLASSDOOR (2016). *Which Countries in Europe Offer the Fairest Paid Leave and Unemployment Benefits?* Tech. rep., Llewellyn Consulting.
- GOLDIN, C. (1990). Understanding the gender gap: An economic history of American women. *Oxford University Press, New York*.

- (2006). The quiet revolution that transformed women’s employment, education, and family. *American Economic Review*, **96** (2), 1–21.
- (2014). A grand gender convergence: Its last chapter. *The American Economic Review*, **104** (4), 1091–1119.
- and KATZ, L. F. (2016). *Women Working Longer: Facts and Some Explanations*. Working Paper 22607, National Bureau of Economic Research.
- , — and KUZIEMKO, I. (2006). The homecoming of american college women: The reversal of the college gender gap. *The Journal of Economic Perspectives*, **20** (4), 133–133.
- GREENWOOD, J., SESHADRI, A. and YORUKOGLU, M. (2005). Engines of liberation. *The Review of Economic Studies*, **72** (1), 109–133.
- GREGG, P., GUTIÉRREZ-DOMÈNECH, M. and WALDFOGEL, J. (2007). The employment of married mothers in great britain, 1974-2000. *Economica*, **74** (296), 842–864.
- GRINZA, E., DEVICIENTI, F., ROSSI, M. and VANNONI, D. (2017). *How Entry into Parenthood Shapes Gender Role Attitudes: New Evidence from Longitudinal UK Data*. Tech. rep., Working Paper.
- GURRYAN, J., HURST, E. and KEARNEY, M. (2008). Parental education and parental time with children. *Journal of Economic Perspectives*, **22** (3), 23–46.
- HERR, J. L. (2015). The labor supply effects of delayed first birth. *American Economic Review*, **105** (5), 630–37.
- (forthcoming). Measuring the effect of first birth timing on women’s wages. *Journal of Population Economics*.
- HOTZ, J. V., JOHANSSON, P. and KARIMI, A. (2017). *Parenthood, family friendly firms, and the gender gaps in early work careers*. Tech. rep., NBER Working Paper No. 24173.
- JARRALLAH, Y., PERALES, F. and BAXTER, J. (2016). *Parenthood and men’s and women’s gender-role attitudes: Does child’s gender matter?* Tech. rep., Working Paper.
- KAUFMANN, K. M., MESSNER, M. and SOLIS, A. (2015). *Elite higher education, the marriage market and the intergenerational transmission of human capital*. Tech. rep., Working paper.
- KLEVEN, H., LANDAIS, C. and SØGAARD, J. (2016). The breadwinner notch: A bunching approach to estimating reference-dependent preferences. *Unpublished manuscript, London School Econ*.
- KLEVEN, H. J., LANDAIS, C. and SØGAARD, J. E. (2018). *Children and gender inequality: Evidence from Denmark*. Tech. rep., NBER Working Paper No. 24219.
- LEFGREN, L. and MCINTYRE, F. (2006). The relationship between women’s education and marriage outcomes. *Journal of Labor Economics*, **24** (4), 787–830.
- LUNDBORG, P., PLUG, E. and RASMUSSEN, A. W. (2017). Can women have children and a career? iv evidence from ivf treatments. *American Economic Review*, **107** (6), 1611–37.
- MAS, A. and PALLAIS, A. (2017). Valuing alternative work arrangements. *American Economic Review*, **107** (12), 3722–59.
- MCGINN, K. L., MAYRA, R. C. and LINGO, E. L. (2015). *Mums the word! Cross-national Effects of Maternal Employment on Gender Inequalities at Work and at Home*. Tech. rep.,

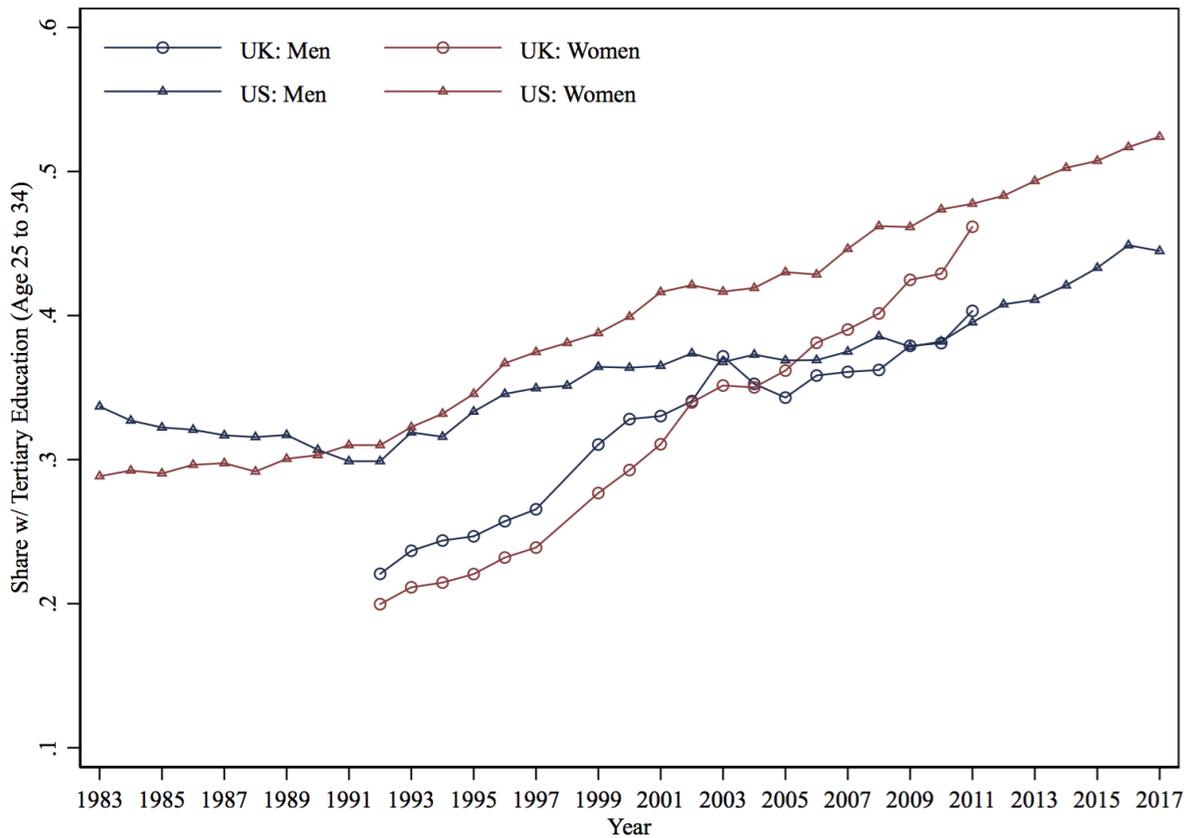
- Harvard Business School Working Paper No. 15-094.
- MILLER, A. R. (2011). The effects of motherhood timing on career path. *Journal of Population Economics*, **24** (3), 1071–1100.
- OLIVETTI, C., PATACCINI, E. and ZENOU, Y. (2016). *Mothers, peers and gender identity*. Tech. rep., Working Paper.
- OREOPOULOS, P. and SALVANES, K. G. (2011). Priceless: The nonpecuniary benefits of schooling. *Journal of Economic perspectives*, **25** (1), 159–84.
- PERTOLD-GEBICKA, B., PETOLD, F. and DATTA GUPTA, N. (2016). *Employment adjustments around childbirth*. Tech. rep., IZA Discussion Paper No. 9685.
- RAMEY, G. and RAMEY, V. A. (2010). The rug rat race. *Brookings Papers on Economic Activity*.
- SCHANK, H. and WALLACE, E. (2018). *The Ambition Decisions: What Women Know About Work, Family, and the Path to Building a Life*. Viking.
- SHAFER, E. F. and MALHOTRA, N. (2011). The effect of a child’s sex on support for traditional gender roles. *Social Forces*, **90** (1), 209–222.
- VESPA, J. (2009). Gender ideology construction: A life course and intersectional approach. *Gender and Society*, **23** (3), 363–387.
- WALDFOGEL, J. (1997). The effect of children on women’s wages. *American sociological review*, pp. 209–217.
- WASSERMAN, M. (2018). *Hours constraints, occupational choice, and gender: Evidence from medical residents*. Tech. rep., Working Paper.
- WILDE, E. T., BATCHELDER, L. and ELLWOOD, D. T. (2010). *The mommy track divides: The impact of childbearing on wages of women of differing skill levels*. Tech. rep., NBER Working Paper No. 16582.
- WISWALL, M. and ZAFAR, B. (2017). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, **133** (1), 457–507.

Figure 1: Female labor force participation rates in select countries (ages 25-54)



Source: OECD (https://stats.oecd.org/Index.aspx?DataSetCode=LFS_SEXAGE_I_R#). Data downloaded March 20th, 2018.

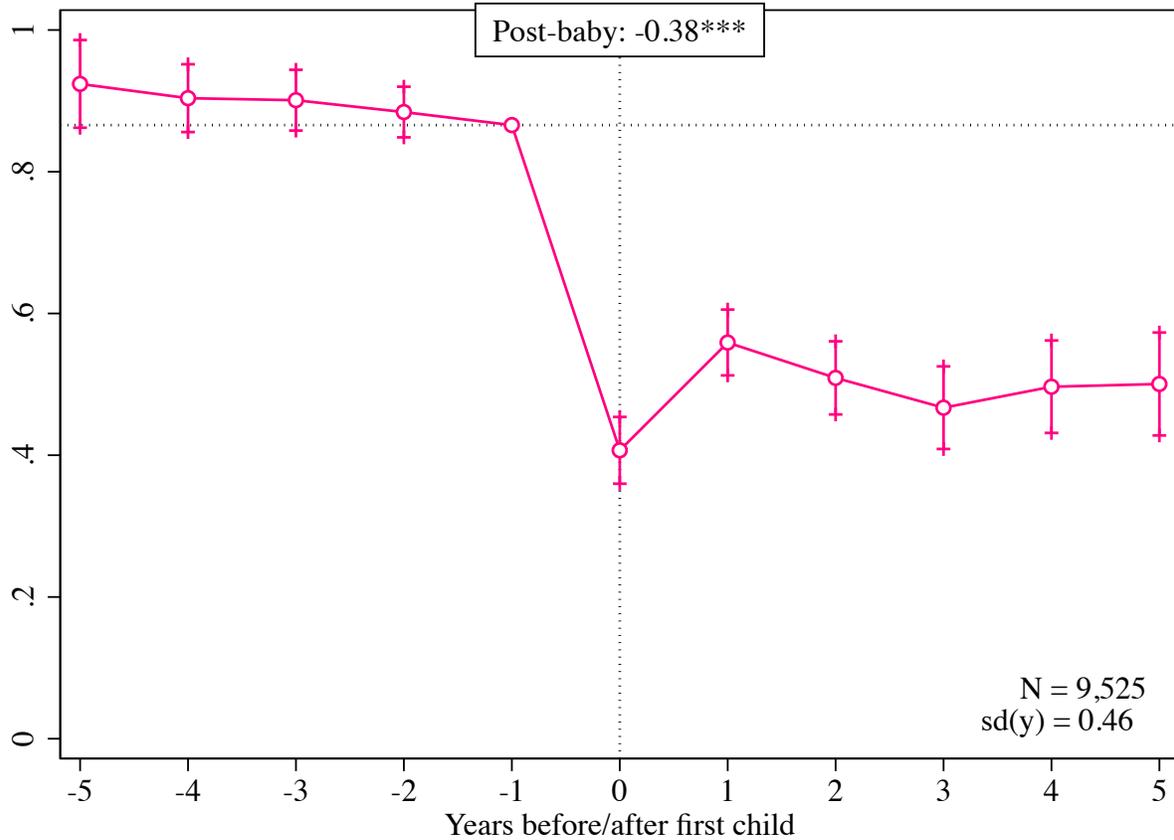
Figure 2: Share of young adults with tertiary education, by gender



Source: CPS and EU-LFS.

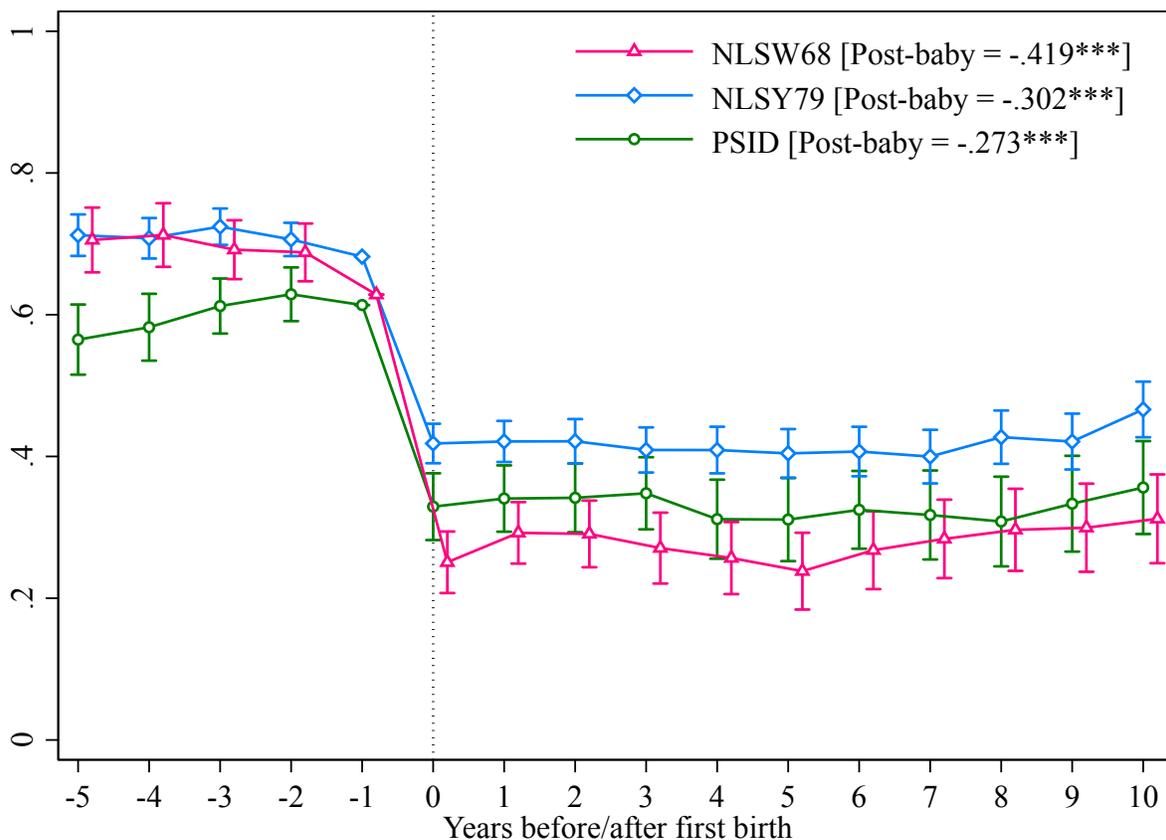
Notes: The sample is restricted to individuals age 25 to 34. In the EU-LFS, tertiary education is defined as having completed ISCED levels 5 or 6. In the CPS, tertiary-education is defined as an associate's degree or more. In 1992, the CPS changed its coding of education. To account for this change, we set the share of people with tertiary education in 1991 to be the same as the share in 1992. For years prior to 1991, we set the share of people with tertiary-education to be the share in 1992 times the percent change between that year and 1991.

Figure 3: Event-study analysis of how probability of employment changes after motherhood, UK



Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.
Notes: The dependent dummy variable is a variable equal to 1 if an individual did paid work (i.e. was physically present at their job) last week. Event years are measured with some rounding error, so that event year 0 (the year of birth) is ± 6 months around the first birth, so the child is either born or in utero. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. Standard errors are clustered at the person level. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. The sample comprises women who were original sample members in the British Household Panel Survey and for whom we observe their first birth.
 *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Figure 4: Event-study analysis of how probability of employment changes after motherhood, US

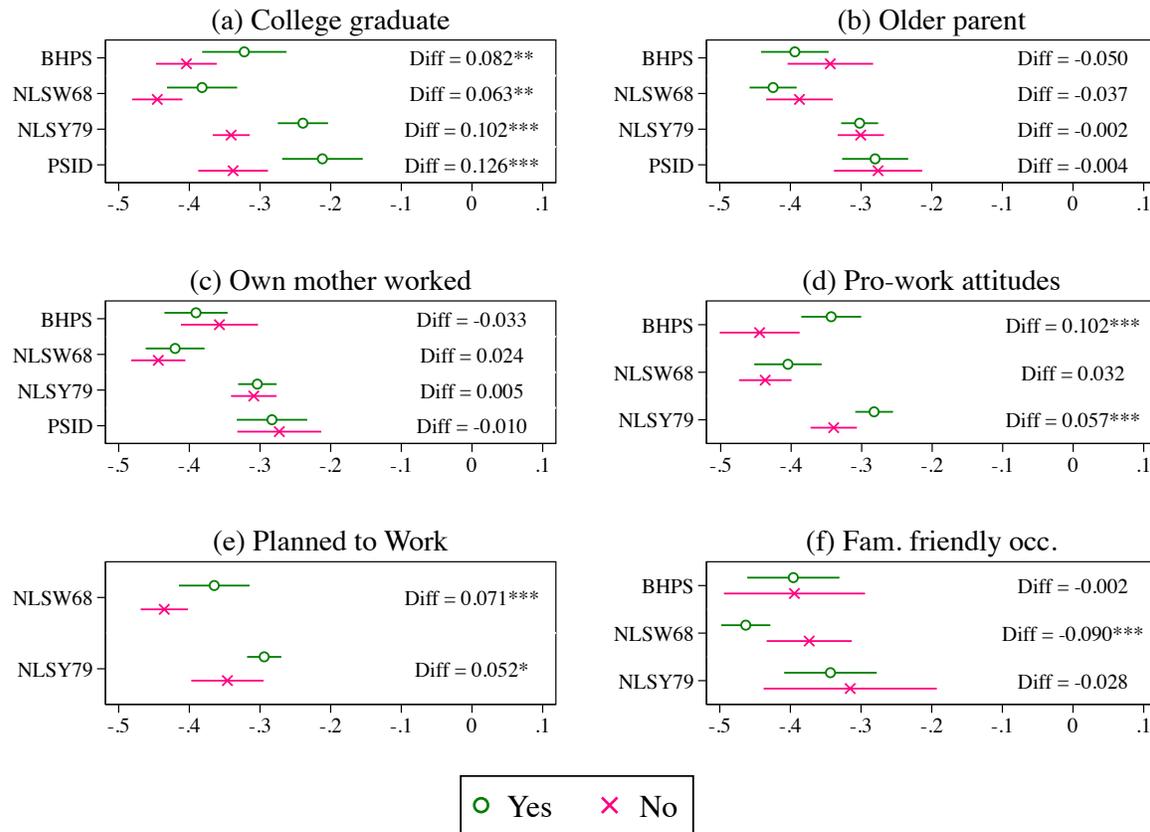


Sources: National Longitudinal Survey of Young Women (NLSW68), National Longitudinal Survey of Youth (NLSY79), and Panel Study of Income Dynamics (PSID). For more detail see Section 2.2 and Appendix B.2-B.4.

Notes: Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 10$) include all person-year observations five or more years before (ten or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

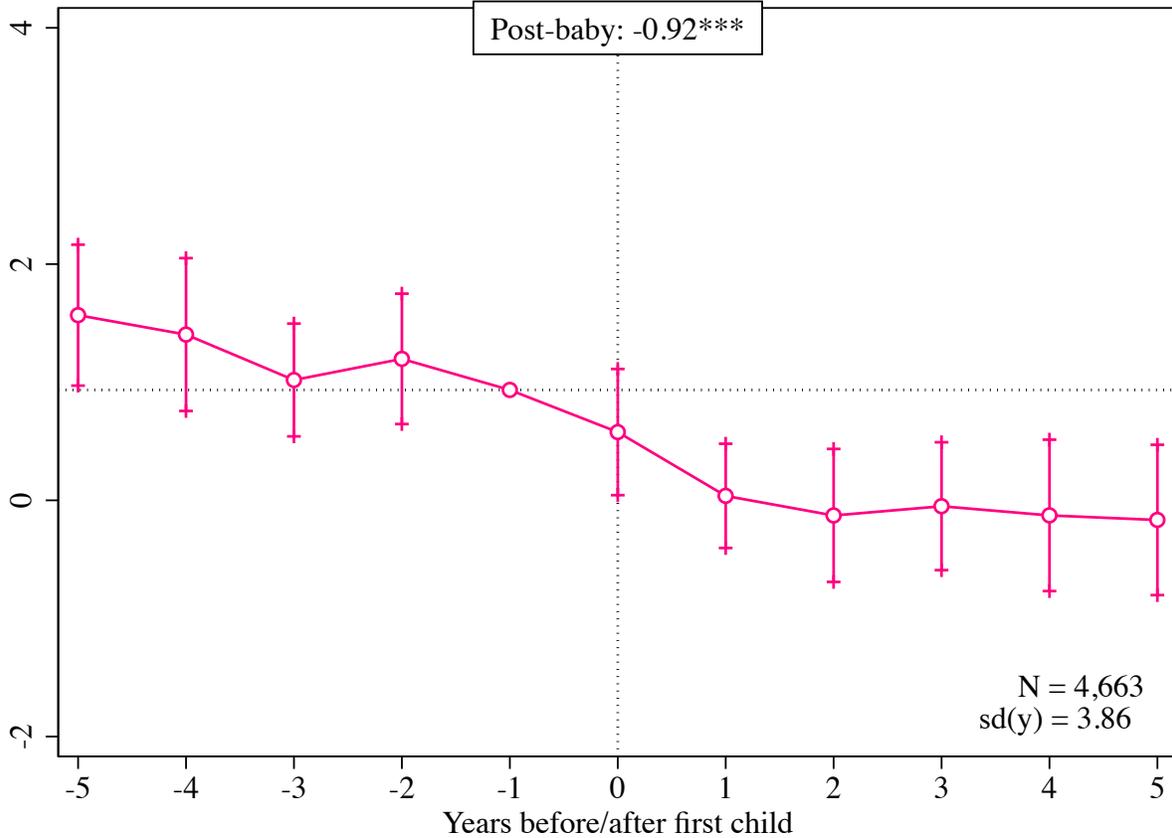
Figure 5: Heterogeneous effects of first birth on current employment, women



Sources: British Household Panel Survey (BHPS), National Longitudinal Survey of Young Women (NLSW68), National Longitudinal Survey of Youth (NLSY79), and Panel Study of Income Dynamics (PSID). For more detail see Section 2.2 and Appendix B.2-B.4.

Notes: Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Estimates are from a regression that is identical to our difference-in-differences specification (i.e., including person, age, and calendar year fixed effects), except that we also include an interaction term of *Post-baby* with our various heterogeneity variables (e.g. completed tertiary education). Estimates corresponding to “No” are from the *Post-baby* term and estimates corresponding to “Yes” are from the sum of the *Post* and interaction term. The “Diff” values listed to the right of each series of estimates are the difference between the “Yes” and “No” terms. Regressions in the NLSW68, NLSY79, and PSID are weighted by the probability of selection. Standard errors are clustered at the person level.

Figure 6: Event-study analysis of how gender-role attitudes change after motherhood (BHPS)

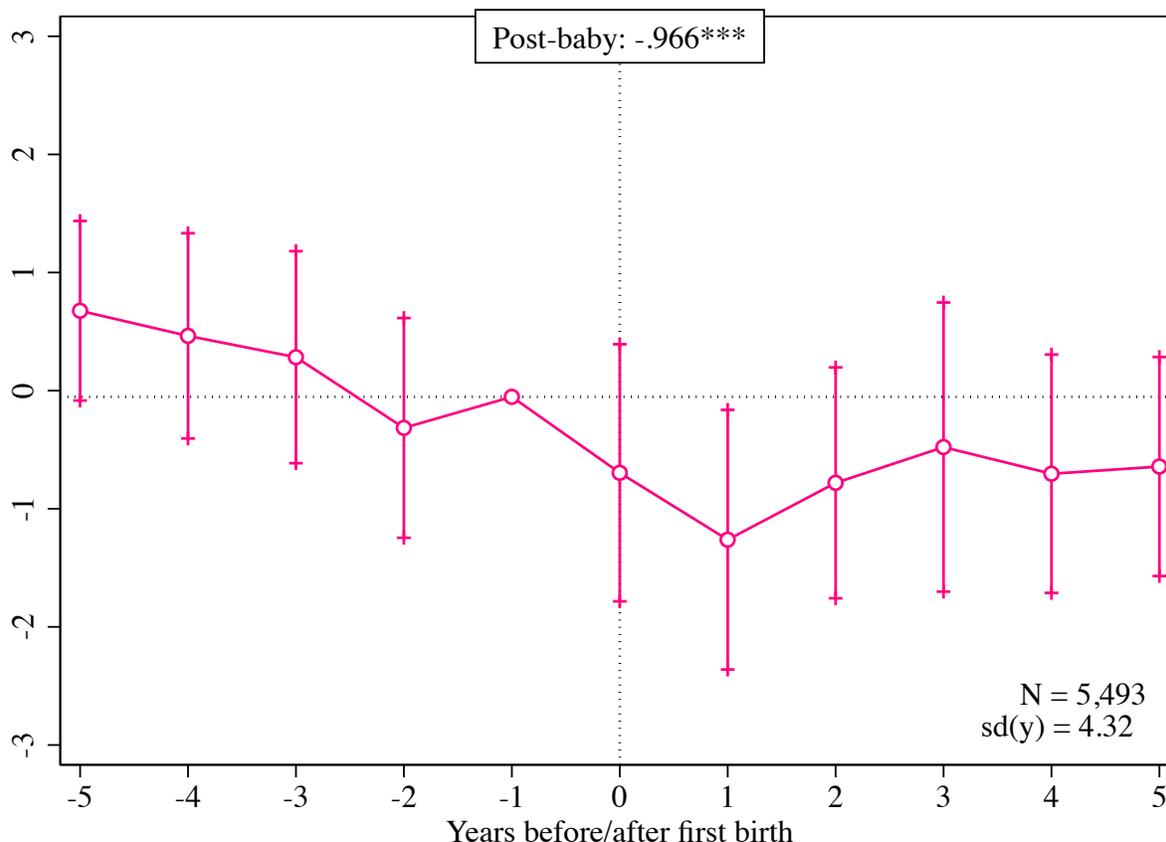


Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: The gender-role index variable is constructed as the (standardized) sum of six questions on the appropriate role of men and women with respect to household, child-care, and labor-market responsibilities. It is increasing in the “pro-female-employment” direction. A list of these questions can be found in Appendix Table A.6 and event studies for each component can be found in Appendix Figure A.4. The sample comprises all female original sample members for whom we observe the event of childbirth. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Standard errors are clustered at the person level.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

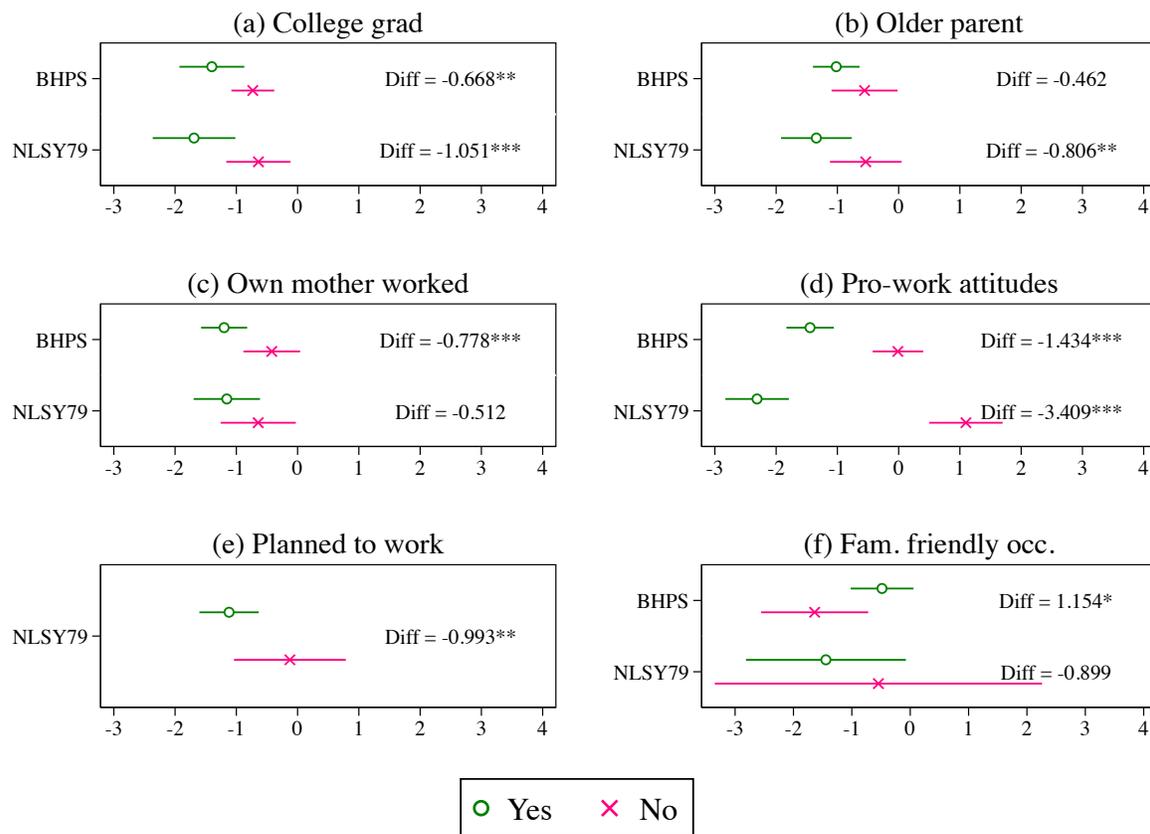
Figure 7: Event-study analysis of how gender-role attitudes change after motherhood (NLSY79)



Source: National Longitudinal Survey of Youth. For more detail see Section 2.2 and Appendix B.3.
Notes: The gender-role index is constructed as the (standardized) sum of seven questions on the appropriate role of men and women with respect to household, child-care, and labor-market responsibilities. It is increasing in the “pro-female-employment” direction. A list of these questions can be found in Appendix Table A.8. The sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. For a more detailed description of sampling criteria see Appendix B. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Standard errors are clustered at the person level.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

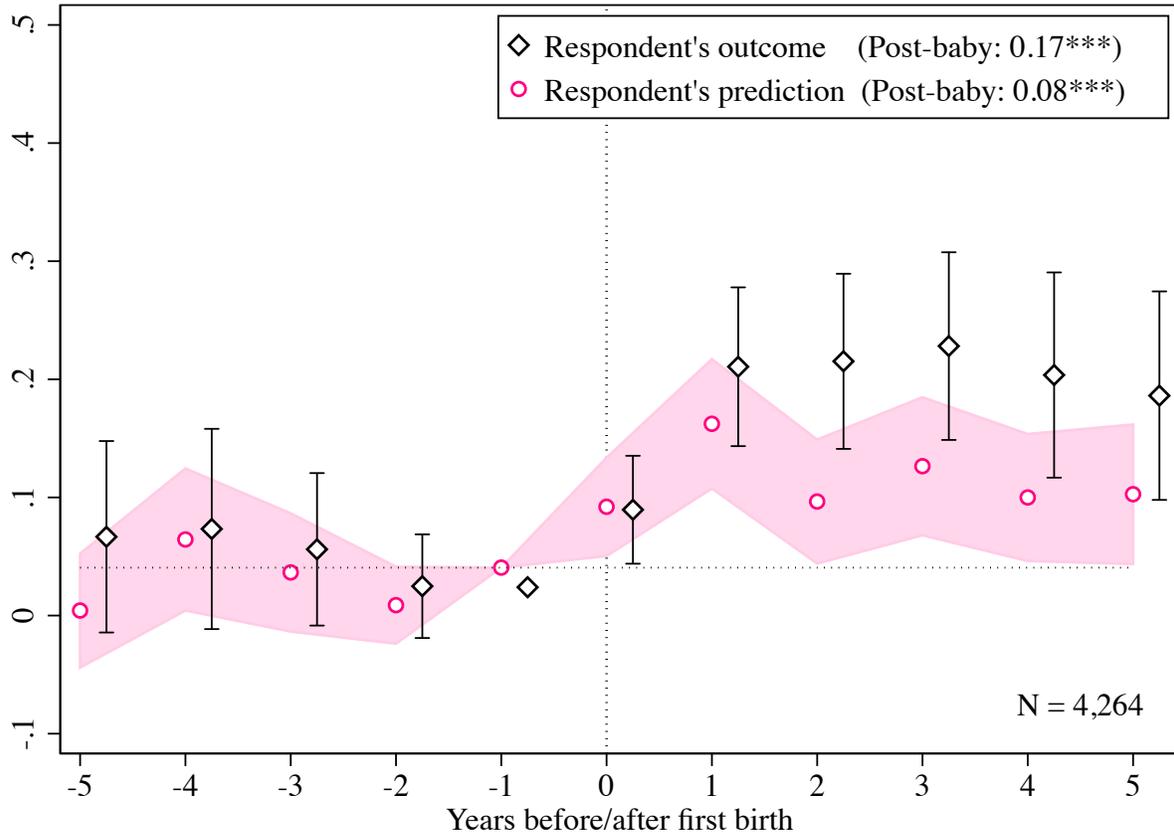
Figure 8: Heterogeneous effects of first birth on women’s gender-role attitudes



Sources: British Household Panel Survey (BHPS) and National Longitudinal Survey of Youth (NLSY79). For more detail see Section 2 and Appendix B.

Notes: Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Estimates are from a regression identical to our difference-in-differences specification (i.e., including person, age, and calendar-year fixed effects), except that we also include an interaction term of *Post-baby* with our various heterogeneity variables (e.g., obtained a college degree). Estimates corresponding to “No” are from the estimated coefficient on the *Post-baby* term and estimates corresponding to “Yes” are from the sum of the *Post-baby* and interaction term. The “Diff” values listed to the right of each series of estimates are the difference between the “Yes” and “No” terms. Regressions in the NLSY79 are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

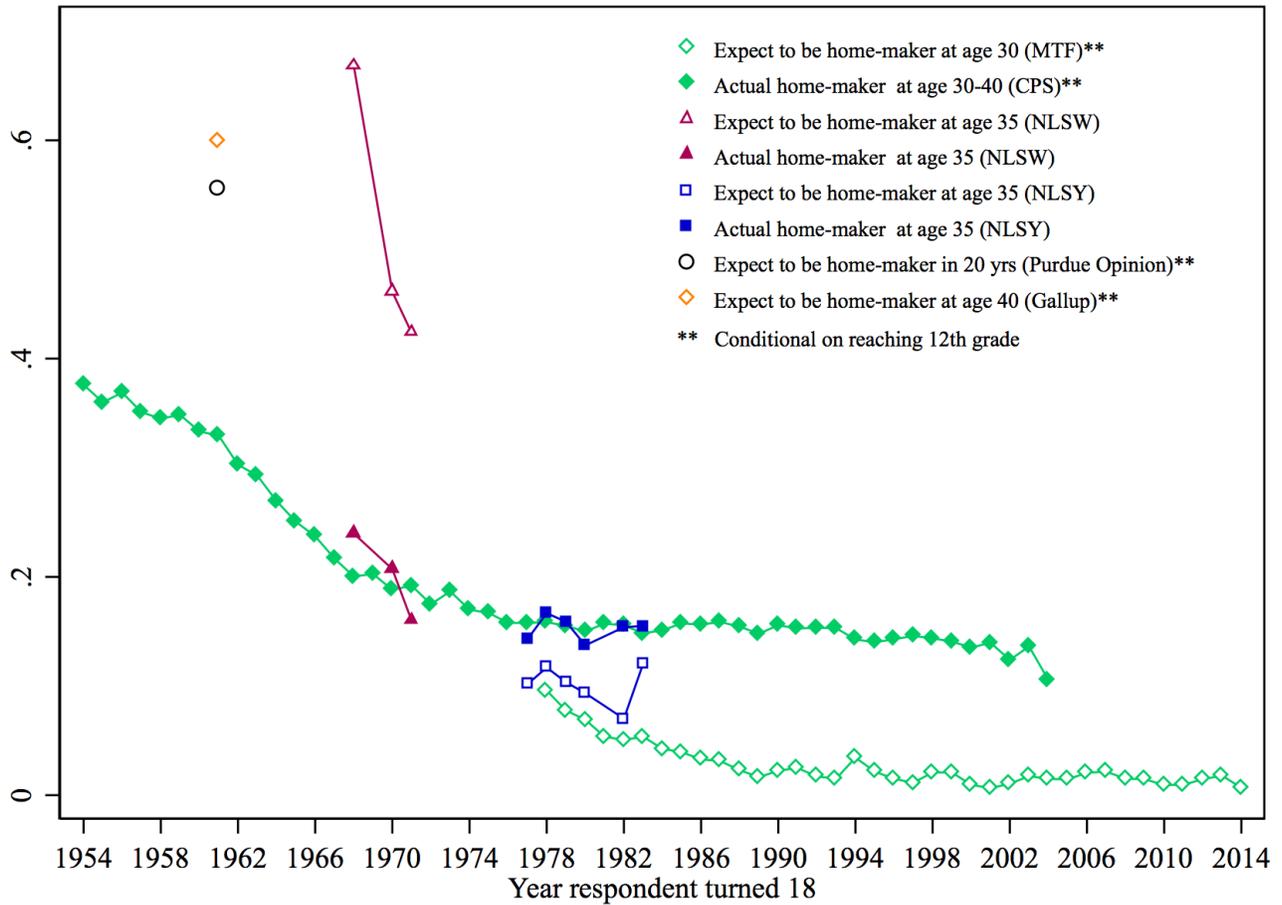
Figure 9: Event-study analysis of how expectations and reality of giving up paid work change upon motherhood, BHPS



Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Sample includes only women who were observed working in the previous year. Predictions are made at $t = -1$ regarding labor supply in year t . Actual labor supply is whether a woman is working this year, conditional on having worked in the previous year. Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Standard errors are clustered at the person level.

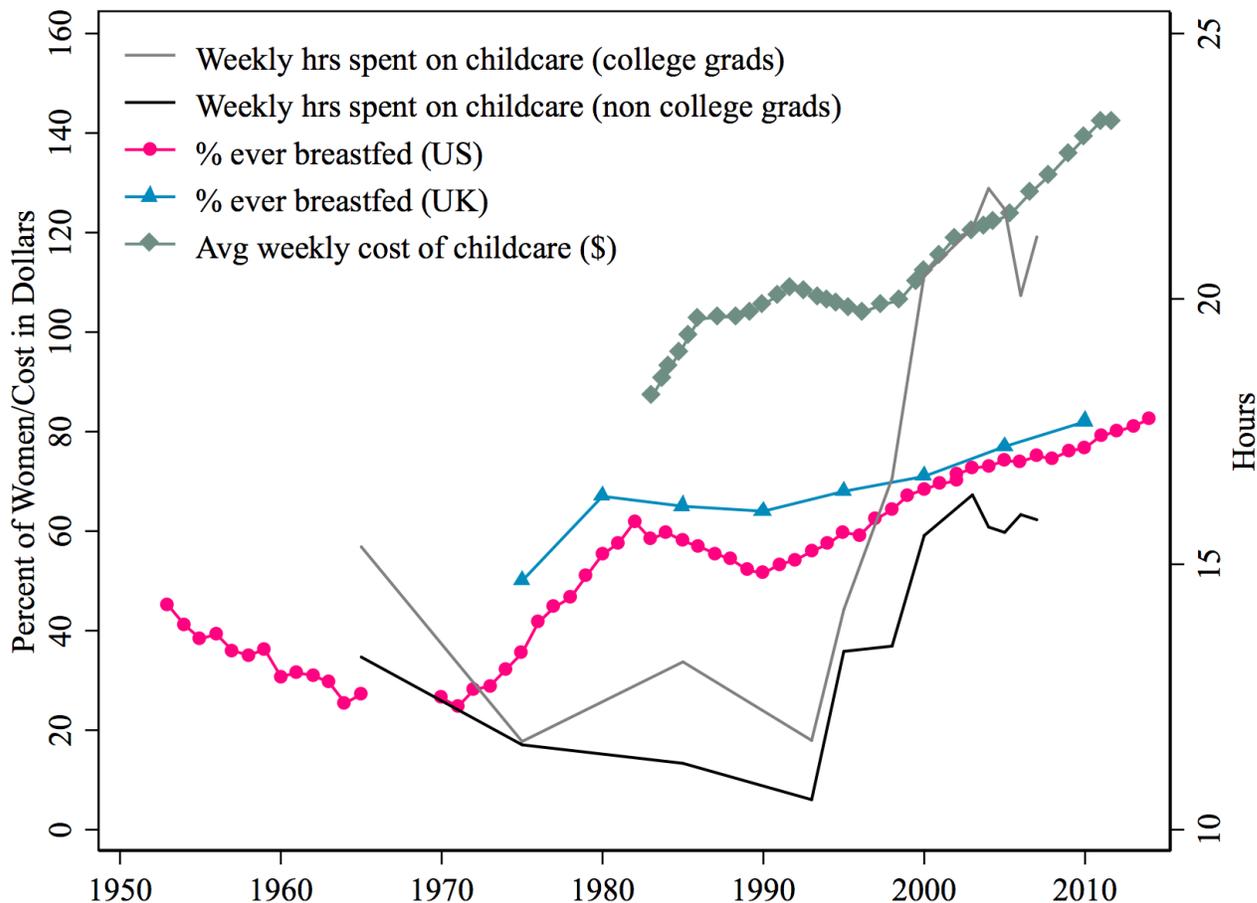
Figure 10: Expectations and realizations of future home-making among U.S. high school seniors



Source: Monitoring the Future Survey (MTF), 1978 to 2014; CPS, 1972 to 2016; NLSW68, 1968 to 1973; NLSY79, 1979 to 1984; Purdue Opinion Panel, 1961; Gallup Opinion Poll, 1961.

Notes: Home-maker is defined in the CPS, NLSW68, and NLSY79 as an individual who is out of the labor force and married to a wage earner. In the MTF, Purdue Opinion Panel, and Gallup Opinion Poll, expected home-maker is defined as someone who reports planning to be a housewife or homemaker at age 30, 20 years from now, and at age 40, respectively. In the NLSW68, expected home-maker is someone who reports planning to be married and raising a family at age 35. In the NLSY79, expected home-maker is someone who reports planning to be married and raising a family *and* not working outside the home at age 35. To match the MTF, the CPS sample is restricted to individuals who have completed at least 11th grade. We similarly restrict our Purdue and Gallup samples.

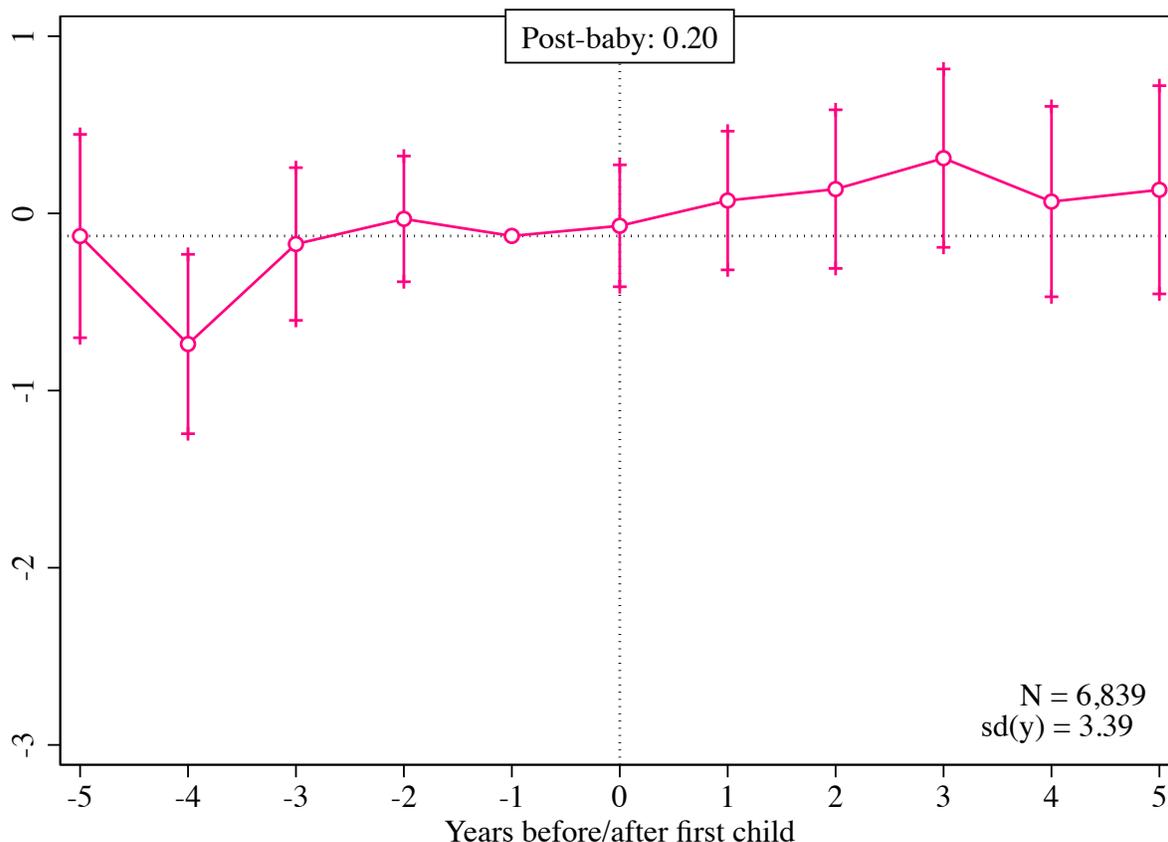
Figure 11: Various proxies for “employment costs of motherhood”



Source: Breastfeeding data for the U.S. come from three different sources: the National Fertility Survey (1953-1965), the Ross Labs Mothers Survey (1970-2002), and the National Immunization Survey (2002-2014). Breastfeeding data for the UK come from the Infant Feeding Survey. Data for weekly hours spent on childcare are taken from Table 1 of (Ramey and Ramey, 2010). Data for average weekly childcare costs were extracted from Census Bureau Report CB13-62, available at <https://www.census.gov/newsroom/press-releases/2013/cb13-62.html>, using an online WebPlotDigitizer tool.

Notes: Estimates of breastfeeding rates from the National Fertility Survey, which was performed in 1965, 1970, and 1975, pool all three surveys and are calculated based on the year of birth of the child, rather than the year each woman was surveyed. We only plot years in which estimates come from 900 or more observations. In the Ross Labs Mothers Survey, the National Immunization Survey, and the Infant Feeding Survey, mothers are asked about their children who are currently infants. In these surveys, the survey year, rather than the year of birth of the child, is used. Estimates for weekly hours spent on childcare, taken from Ramey and Ramey (2010), come from a combination of twelve nationally representative time use surveys. Cost of childcare measures the average weekly expenditures among people making childcare payments, in 2011 dollars.

Figure 12: Event-study analysis of how overall job satisfaction changes upon motherhood (BHPS)



Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Higher values means greater overall job satisfaction. Components of the omnibus measure are: Overall job satisfaction as well as satisfaction with total pay, security, work itself, and hours worked. Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Standard errors are clustered at the person level.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 1: Summary statistics for main panel data sources

	BHPS	NLSW 68	NLSY 79	PSID
<i>–Employment</i>				
Working this week	0.70	0.62	0.69	0.68
In work force this week	0.78	0.72	0.73	0.74
Hours worked last week	24.87	21.79	26.03	.
Expected to give up paid work last wave	0.06	.	.	.
Actually gave up paid work this wave	0.09	.	.	.
<i>–Background</i>				
Year of birth	1968.75	1950.79	1961.70	1965.96
Age at first birth	28.95	23.25	24.48	24.40
Total number of children	1.81	2.33	2.27	2.27
Married pre-baby	0.63	0.63	0.60	0.57
Black	.	0.11	0.14	0.16
College graduate	0.26	0.25	0.29	0.38
Planned to be working at 35	.	0.25	0.84	.
Gender-liberal pre-baby	0.62	0.31	0.60	.
Mom worked when resp. was 14	0.65	0.41	0.56	0.62
<i>–Average # of Pre-/Post- Obs. per Person</i>				
Number of employment observations, pre-baby	4.08	5.77	8.02	7.65
Number of employment observations, post-baby	6.63	12.99	16.04	14.49
Number of respondents	681.00	1413.00	2256.00	1025.00
Observations	10,410	26,376	54,175	22,356

Source: British Household Panel Survey (BHPS), National Longitudinal Survey of Young Women (NLSW68), National Longitudinal Survey of Youth 1979 (NLSY79), Panel Survey of Income Dynamics (PSID). For more detail see Section 2 and Appendix B.1-B.4.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Statistics in the NLSW68, NLSY79, and PSID are weighted by the probability of selection.

Table 2: Relationship between college education and probability of reporting that being a parent is harder than expected

	Dependent variable: Parenthood harder than expected				
	Women			Men	
	(1)	(2)	(3)	(4)	(5)
College degree	0.123*** [0.0258]	0.108*** [0.0282]	0.0412*** [0.0149]	0.0195 [0.0459]	-0.00488 [0.0221]
Mean, dep. var.	0.516	0.516	0.521	0.372	0.411
Observations	1,687	1,687	6,183	576	2,539
Controls?	No	Yes	Yes	Yes	Yes
Age of children at home	0-6	0-6	0-17	0-6	0-17

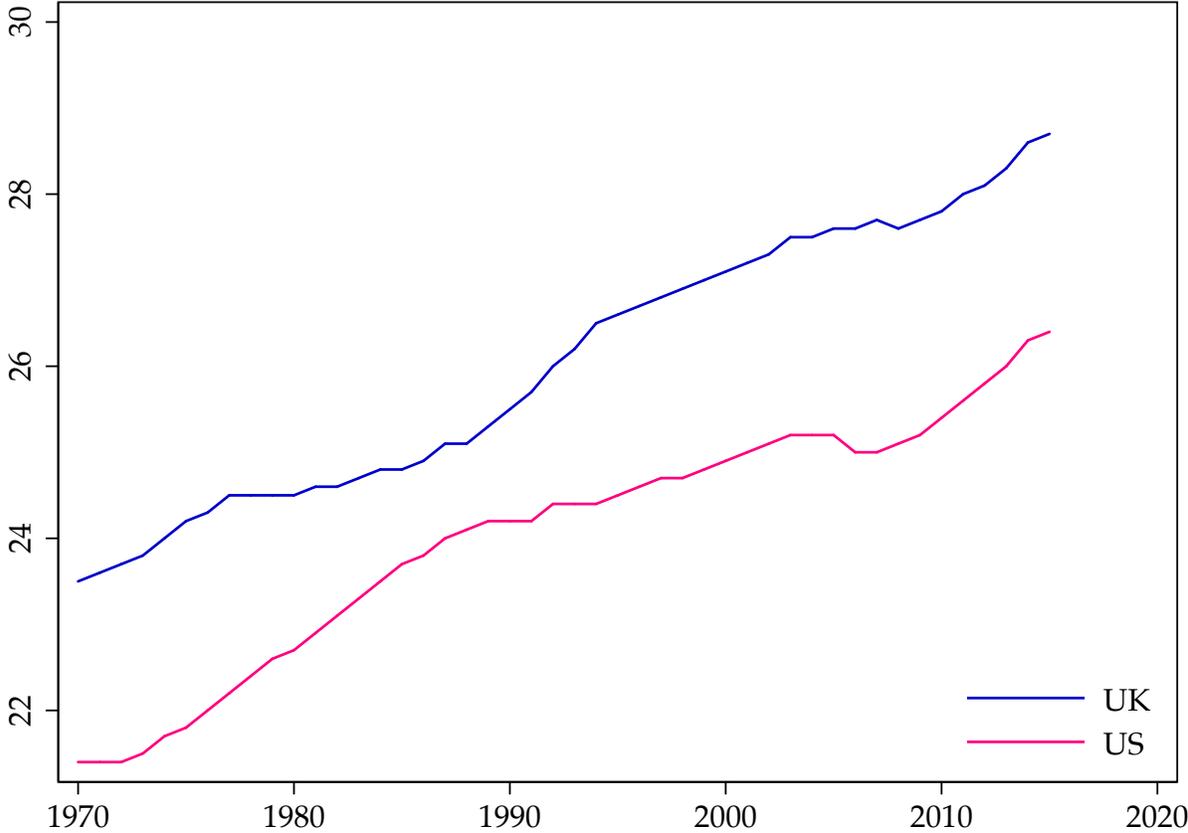
Source: Childhood Development Supplement to the Panel Study of Income Dynamics. For more detail see Appendix B.4.

Notes: Sample comprises all primary and secondary caregivers in the Childhood Development Supplement to the Panel Study of Income Dynamics. Columns (2)-(5) control for birth decile, survey year, race, child age, and number of children fixed effects, marital status fixed effects and log of household income. Columns (1), (2), and (4) restrict to parents with children under 6 whereas columns (3) and (5) restrict to parents with children under 18. Standard errors are reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

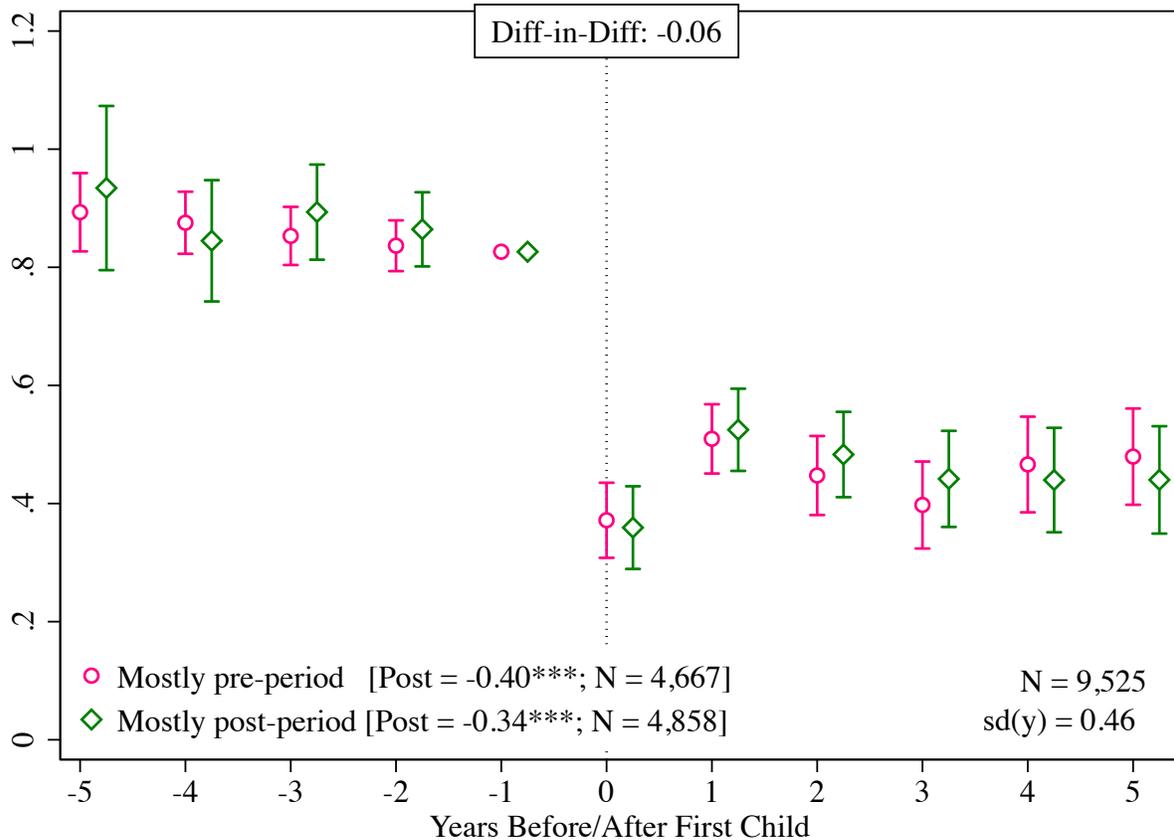
Appendix A. Appendix figures and tables

Appendix Figure A.1: Average age at first birth (U.S. and U.K. women)



Source: Center for Disease Control (US) and Office for National Statistics (UK).

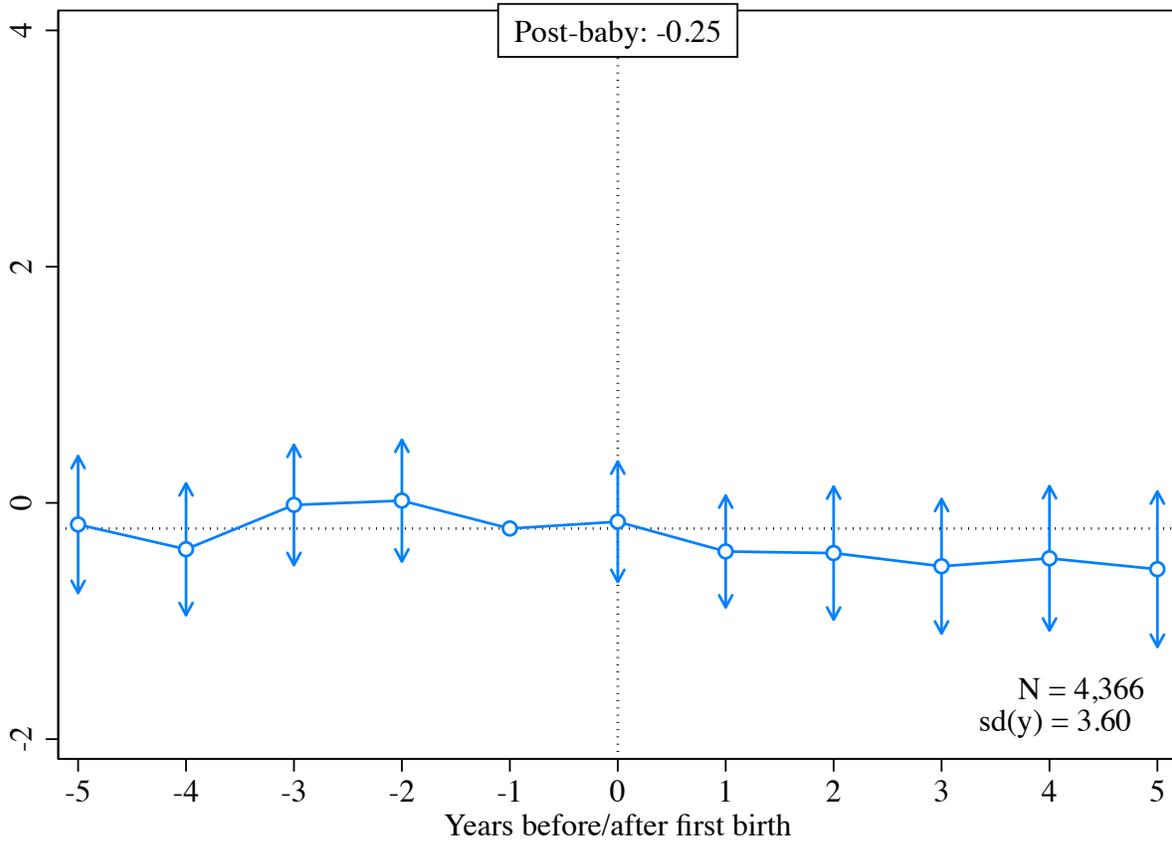
Appendix Figure A.2: Employment event-study analysis, splitting respondents by whether they mostly identify the pre- or post-periods



Source: British Household Panel Survey.

Notes: The above graph plots estimates from an event study regression, which is our standard event study regression but also includes event study dummy variables interacted with a “mostly pre-period” variable. We construct our “mostly pre-period” variable by first calculating the share of total observations that are in the pre-period for each respondent. Then, we count someone as mostly identifying pre-period event-time coefficients if her share of pre-period observations is above the median across our preferred sample of women. In the BHPS, that median is 33%; thus, if a respondent’s pre-period observations are greater than 33% of her total observations, then she is counted as identifying mostly pre-period observations. In the graph, the “mostly pre-period” series plots the event study coefficients plus the interacted (event-time) \times (mostly pre-period) coefficients. The “mostly post-period” series plots just the standard event study coefficients. As with our standard event study regression, we include person, calendar-year, and age-in-years fixed effects and cluster our standard errors at the individual level.

Appendix Figure A.3: Event-study analysis of how gender-role attitudes change after fatherhood (men in the BHPS)



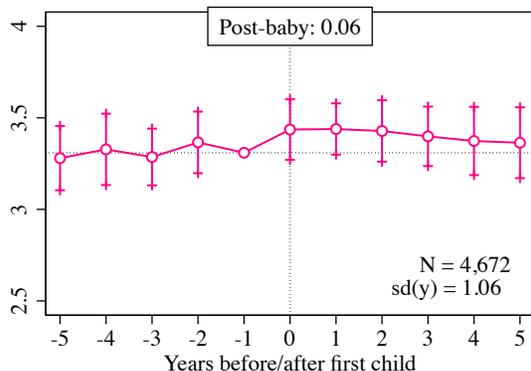
Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: The gender-roles index is constructed as the (standardized) sum of six questions on the appropriate role of men and women with respect to child-care and labor-market responsibilities. It is increasing in the “pro-female-employment” direction. A list of these questions can be found in Appendix Table A.6. Sample comprises all male original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. The year before birth is normalized to the mean of the outcome variable at event time $t = -1$ and $t = -5$ ($t = 5$) include all person-year observations five or more years before (five or more years after) the first birth. The event-study specification controls for age-in-years, calendar-year, and person fixed effects. The estimate reported at the top of the figure corresponds to a parallel difference-in-difference analysis where the event-time is collapsed into a simple before and after the birth of the child. Standard errors are clustered at the person level.

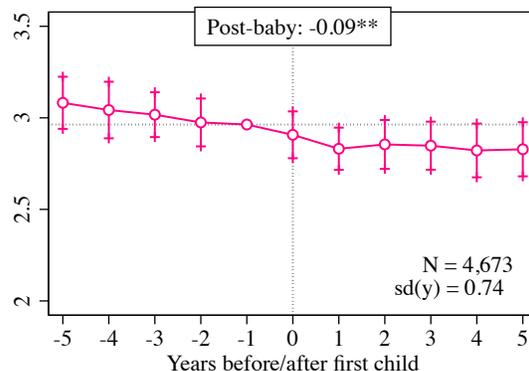
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Appendix Figure A.4: Event-study analysis of how gender-role attitudes change after motherhood (if needed, responses are reversed so that all are increasing in the pro-female-employment position)

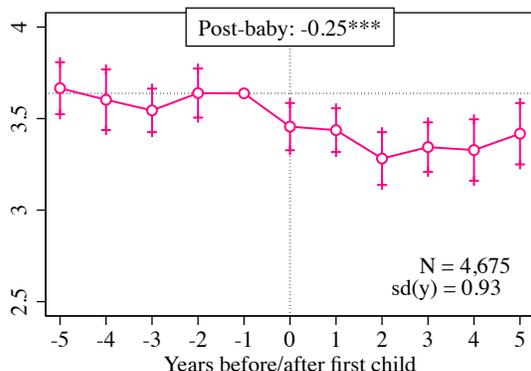
(a) Pre-school child suffers if mom works



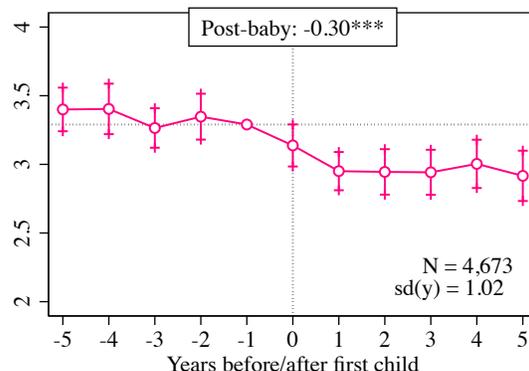
(b) Woman and family happier if she works



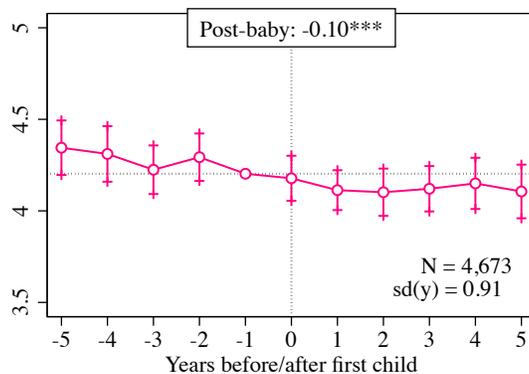
(c) Husband and wife should both contribute



(d) Full time job makes woman independent



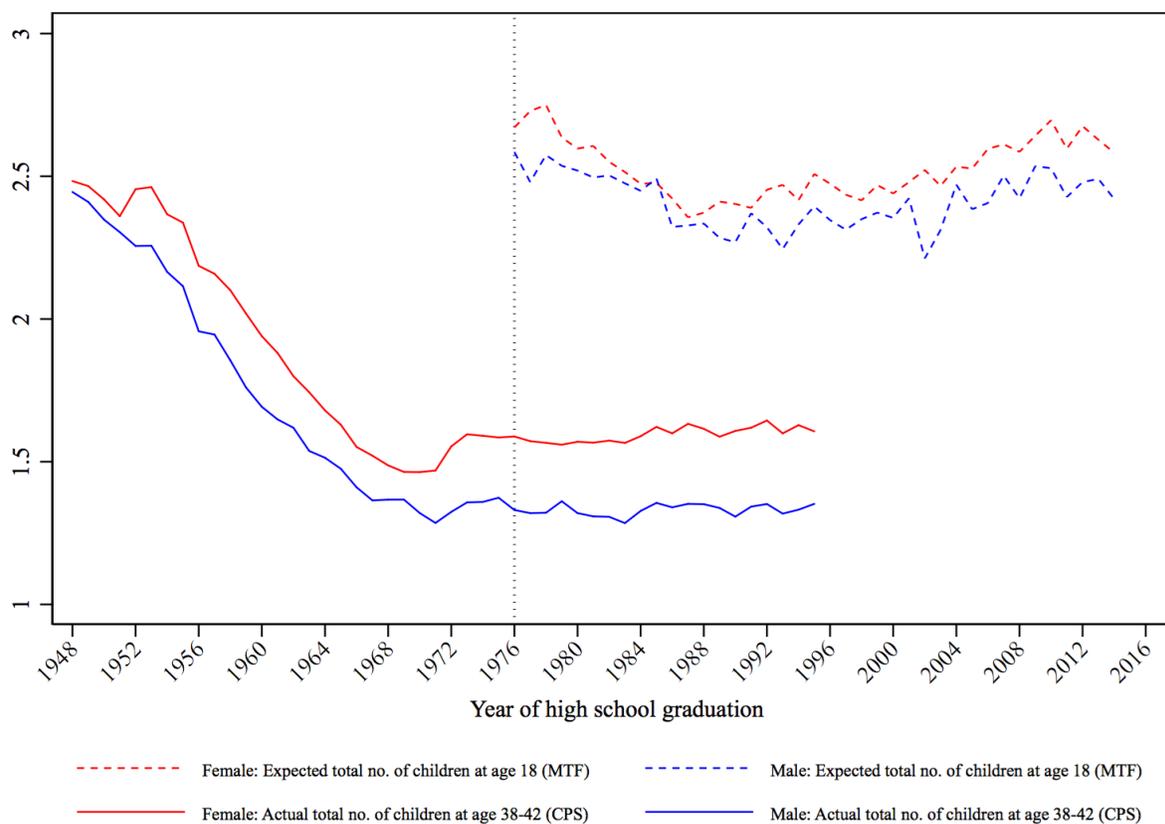
(e) Husband should earn, wife stay at home



Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: See notes to Figure 6. The analysis is identical here except instead of the standardized index of gender-role attitudes, we use as outcomes each component of the index. Unlike the omnibus measure, here we do not standardize the components to have mean zero and variance one.

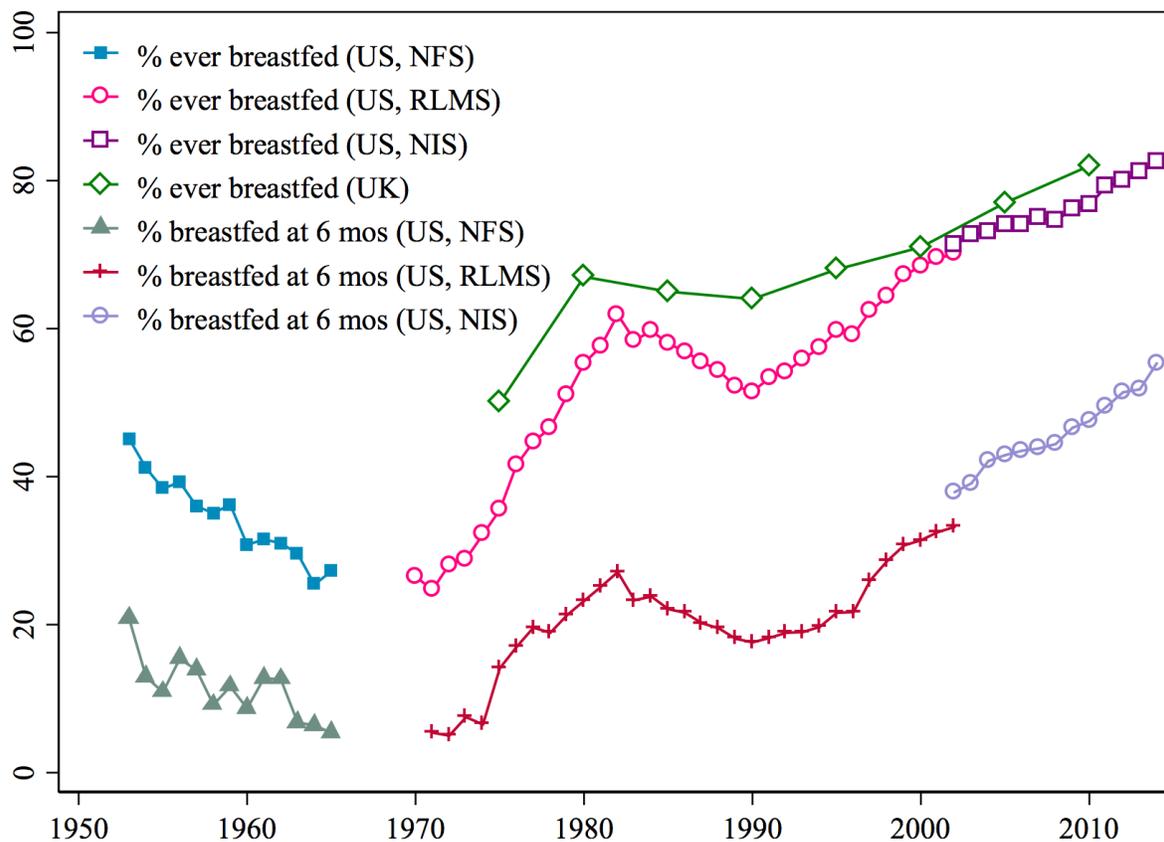
Appendix Figure A.5: Expectations and realizations of fertility (number of children) at age 38 to 42 among U.S. high school seniors



Source: Monitoring the Future Survey, 1976 to 2014; CPS, 1970 to 2017.

Notes: The fertility variable in the MTF is coded from responses to the following question: “All things considered, if you could have exactly the number of children you want, what number would you choose to have?” The responses in the MTF are compared to the realized fertility of the same cohort of high school seniors when they are between the ages of 38 to 42 in the CPS. To match the MTF, the CPS sample is restricted to individuals who have completed at least 11th grade. The variable used to infer completed fertility in the CPS includes biological children, step-children, and adopted children. The observations are weighted using the appropriate sample weights from each survey.

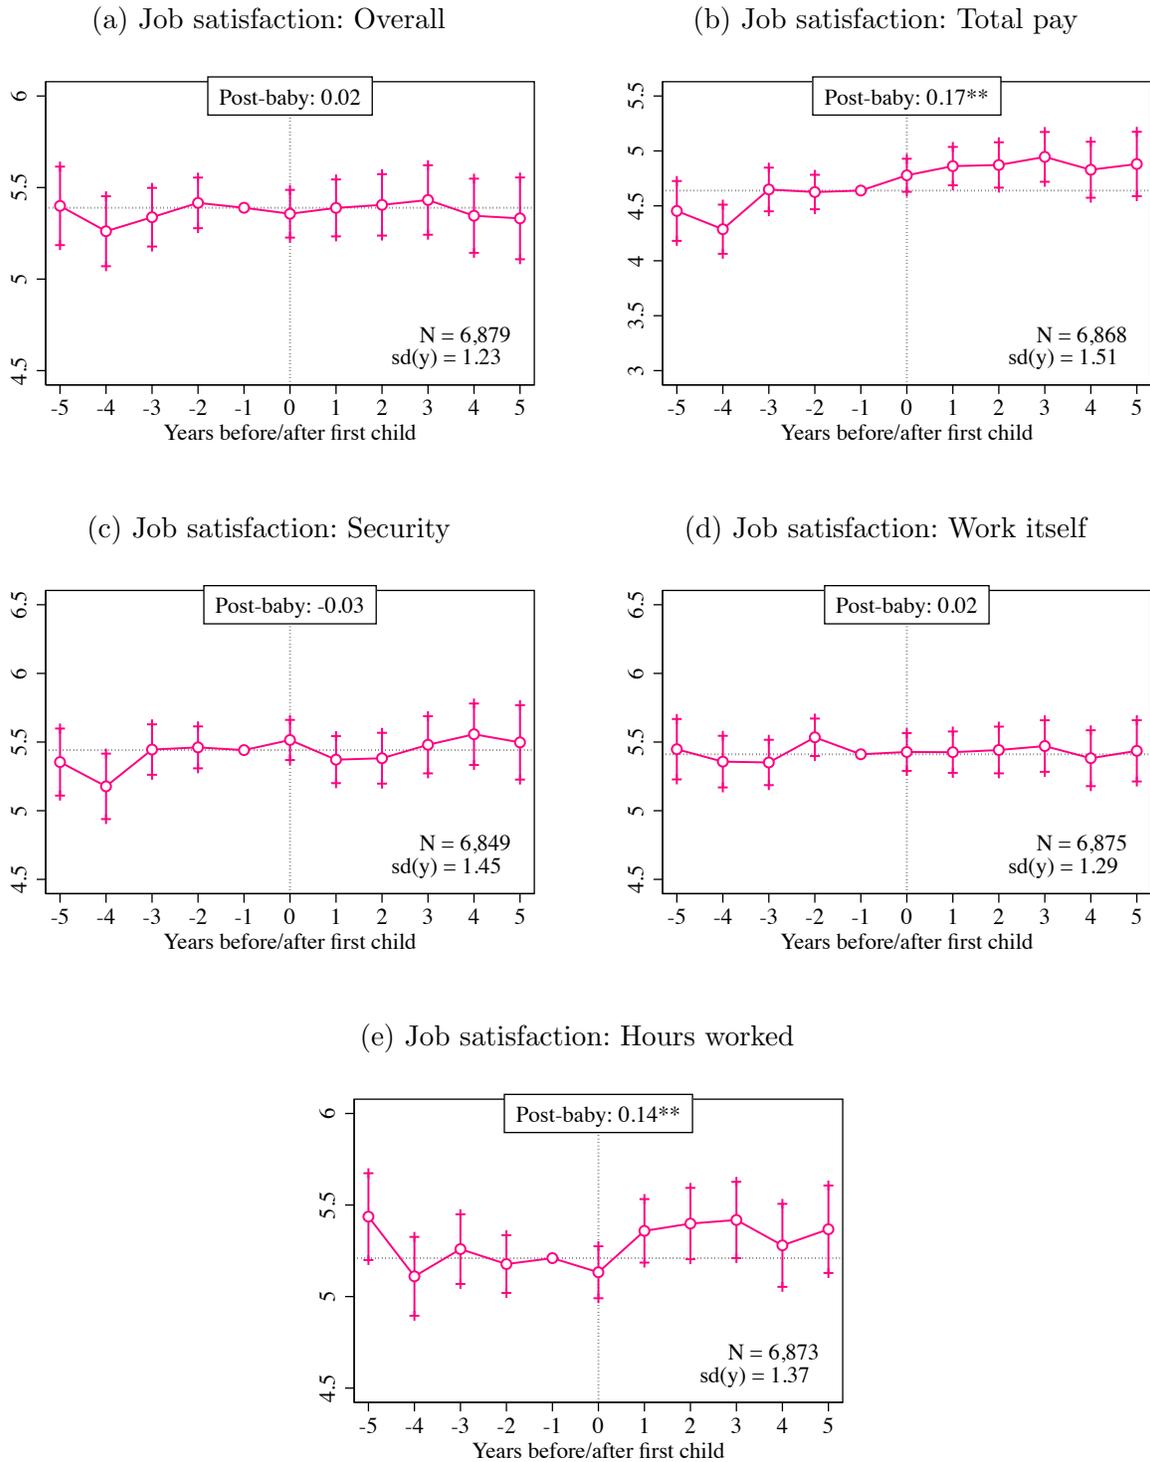
Appendix Figure A.6: Evolution of various measures of breastfeeding prevalence



Source: Data sources for the US: Ross Labs Mothers Survey (RLMS), the National Immunization Survey (NIS), and the National Fertility Survey (NFS). Data source for the UK: Infant Feeding Survey (IFS).

Notes: Estimates of breastfeeding rates from the NFS, which was performed in 1965, 1970, and 1975, pool all three surveys and are calculated based on the year of birth of the child, rather than the year each woman was surveyed. We only plot years in which estimates come from 900 or more observations. In the RLMS, NIS, and IFS, mothers are asked about their children who are currently infants. In these surveys, the survey year, rather than the year of birth of the child, is used.

Appendix Figure A.7: Event-study analysis of how various factors of job satisfaction change after motherhood, BHPS (Higher values indicate greater satisfaction)



Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.
 Notes: The analysis is identical to that in Figure 12, except instead of the “overall” satisfaction variable, we examine questions related to satisfaction of specific attributes of the respondent’s current job.

Appendix Table A.1: Summary statistics for main panel data sources, men

	BHPS	NLSY 79	PSID
<i>–Employment</i>			
Working this week	0.89	0.81	0.82
In work force this week	0.94	0.85	0.88
Hours worked last week	38.42	36.32	.
Expected to give up paid work last wave	0.01	.	.
Actually gave up paid work this wave	0.02	.	.
<i>–Background</i>			
Year of birth	1967.29	1961.70	1966.25
Age at first birth	30.45	25.96	26.29
Total number of children	1.78	2.29	2.24
Married pre-baby	0.60	0.58	0.61
Black	.	0.15	0.15
College graduate	0.27	0.26	0.39
Planned to be working at 35	.	0.88	.
Gender-liberal pre-baby	0.51	0.38	.
Mom worked when resp. was 14	0.67	0.54	0.58
<i>–Average # of Pre-/Post- Observations per Person</i>			
Number of employment observations, pre-baby	4.52	9.48	8.96
Number of employment observations, post-baby	6.46	14.52	12.18
Number of respondents	628.00	2308.00	796.00
Observations	10,034	55,087	16,605

Source: British Household Panel Survey (BHPS), National Longitudinal Survey of Youth 1979 (NLSY79), Panel Survey of Income Dynamics (PSID). For more detail see Section 2 and Appendix B.1-B.4.

Notes: Sample comprises all male original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. For a more detailed description of sampling criteria for each dataset see Appendix B. Statistics for the NLSY79 and PSID are weighted by the probability of selection.

Appendix Table A.2: Effect of first birth on employment last week, robustness checks, BHPS and NLSW68

	Dependent variable: Employed last week								
	BHPS					NLSW68			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-baby	-0.381*** [0.0195]	-0.369*** [0.0179]	-0.369*** [0.0178]	-0.381*** [0.0195]	-0.376*** [0.0204]	-0.419*** [0.0158]	-0.423*** [0.0159]	-0.364*** [0.0127]	-0.419*** [0.0158]
Pre-baby mean, dep. var.	0.826	0.826	0.826	0.826	0.709	0.624	0.624	0.624	0.624
Observations	9,525	9,525	9,525	9,525	8,940	26,308	26,308	26,308	26,308
+ Indiv. lin. time trend	No	Yes	No	No	No	No	Yes	No	No
+ Drop person FE	No	No	Yes	No	No	No	No	Yes	No
+ Control group	No	No	No	Yes	No	No	No	No	Yes
+ Exclude age outliers	No	No	No	No	Yes	n/a	n/a	n/a	n/a

Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar year, and person fixed effects included in all regressions unless otherwise noted. “Exclude age outliers” in the BHPS drops women who have their first child before the age of 20 or after the age of 40. Regressions for NLSW68 are weighted by the probability of selection; regressions for BHPS are not weighted, but inclusion of weights to not substantively change estimated coefficient. (Results available on request.) Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Appendix Table A.3: Effect of first birth on current employment, robustness checks, NLSY79 and PSID

	Dependent variable: Employed last week								
	NLSY79				PSID				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-baby	-0.227*** [0.0153]	-0.224*** [0.0156]	-0.266*** [0.0138]	-0.227*** [0.0153]	-0.260*** [0.0191]	-0.277*** [0.00969]	-0.276*** [0.00973]	-0.314*** [0.00941]	-0.277*** [0.00969]
Pre-baby mean, dep. var.	0.543	0.543	0.543	0.543	0.570	0.675	0.675	0.675	0.675
Observations	22,458	22,458	22,458	22,458	15,802	54,175	54,175	54,175	54,175
+ Indiv. lin. time trend	No	Yes	No	No	No	No	Yes	No	No
+ Drop person FE	No	No	Yes	No	No	No	No	Yes	No
+ Control group	No	No	No	Yes	No	No	No	No	Yes
+ Agers-In	Yes	Yes	Yes	Yes	No	n/a	n/a	n/a	n/a

Source: National Longitudinal Survey of Young Women (NLSW68), National Longitudinal Survey of Youth 1979 (NLSY79), Panel Survey of Income Dynamics (PSID) For more detail see Section 2.2 and Appendix B.2-B.4.

Notes: Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. “Agers-in” refers to children of original sample households who become eligible for interviews as they come of age. Age-in-years, calendar year, and person fixed effects included in all regressions unless otherwise noted. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Appendix Table A.4: Effect of first birth on current employment, U.S. and U.K. men

	Dependent var: Employed last week		
	BHPS	NLSY79	PSID
	(1)	(2)	(3)
Post-baby	-0.0332*** [0.0126]	0.0156* [0.00866]	0.00709 [0.0146]
Pre-baby mean, dep. var.	0.860	0.659	0.586
Observations	8,980	55,087	13,241

Source: British Household Panel Survey (BHPS), National Longitudinal Survey of Youth 1979 (NLSY79), Panel Survey of Income Dynamics (PSID). For more detail see Section 2 and Appendix B.1-B.4.

Notes: Sample comprises all male original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Appendix Table A.5: Impacts of second baby on current employment, women

	Dependent variable: Employed last week							
	BHPS		NLSW68		NLSY79		PSID	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-baby	-0.373*** [0.0193]	-0.376*** [0.0200]	-0.308*** [0.0262]	-0.307*** [0.0262]	-0.287*** [0.0117]	-0.287*** [0.0117]	-0.242*** [0.0220]	-0.243*** [0.0220]
2+ kids	-0.0986*** [0.0207]	-0.0993*** [0.0209]	-0.126*** [0.0263]	-0.132*** [0.0264]	-0.0849*** [0.0115]	-0.0853*** [0.0115]	-0.0729*** [0.0221]	-0.0736*** [0.0221]
Pre-baby mean, dep. var.	0.826	0.826	0.624	0.624	0.675	0.675	0.504	0.504
Observations	9,525	9,525	26,308	26,308	54,175	54,175	17,606	17,606
+ Indiv. lin. time trend	No	Yes	No	Yes	No	Yes	No	Yes

Sources: British Household Panel Survey (BHPS), National Longitudinal Survey of Young Women (NLSW68), National Longitudinal Survey of Youth (NLSY79), and Panel Study of Income Dynamics (PSID). For more detail see Section 2 and Appendix B.1-B.4.

Notes: Sample comprises all female original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Regressions in the NLSW68, NLSY79, and PSID are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Appendix Table A.6: Summary statistics, gender-role attitudes (BHPS)

	Men	Women
<i>–Gender Norms (Higher = More Liberal, Scale: 1-5)</i>		
A pre-school child is likely to suffer if his or her mother works (Rev.)	3.07	3.39
All in all, family life suffers when the woman has a full time job (Rev.)	3.27	3.39
A woman and her family would all be happier if she goes out to work	2.85	2.88
Both the husband and wife should contribute to the household income (Rev.)	3.33	3.41
Having a full-time job is the best way for a woman to be an independent person	3.11	3.03
A husband's job is to earn money; a wife's job is to look after the home and family	3.78	4.06
Omnibus gender norms (higher = more liberal)	-0.48	0.29
<i>–Average # of Pre-/Post- Observations per Person</i>		
Number of omnibus gender observations, pre-baby	2.48	2.22
Number of omnibus gender observations, post-baby	2.80	3.01
Number of respondents	628.00	681.00

Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Individual component attitudes are on a scale of 1 to 5 (1 = strongly agree; 5 = strongly disagree) and reversed so that higher values imply more liberal, less traditional attitudes toward gender roles. The omnibus gender-role index takes each component question, standardizes it to have mean zero and standard deviation one, and sums across all questions.

Appendix Table A.7: Effect of motherhood on individual components of the gender-role index, controlling for a person-specific linear time trend (BHPS)

Dependent variable:	Child Suffers		Family Suffers		Family Happier	
	(1)	(2)	(3)	(4)	(5)	(6)
Post-baby	0.0697 [0.0474]	0.0697 [0.0474]	-0.199*** [0.0480]	-0.200*** [0.0480]	-0.0849* [0.0348]	-0.0850* [0.0348]
Mean, dep. var.	3.391	3.391	3.393	3.393	2.877	2.877
Observations	4,720	4,720	4,722	4,722	4,721	4,721
+ Linear time trend	No	Yes	No	Yes	No	Yes

Dependent variable:	Both Contribute		Full-time Indep.		Husband Earn	
	(1)	(2)	(3)	(4)	(5)	(6)
Post-baby	-0.248*** [0.0420]	-0.248*** [0.0420]	-0.301*** [0.0461]	-0.301*** [0.0461]	-0.0914* [0.0361]	-0.0916* [0.0361]
Mean, dep. var.	3.415	3.415	3.035	3.035	4.059	4.059
Observations	4,723	4,723	4,721	4,721	4,721	4,721
+ Linear time trend	No	Yes	No	Yes	No	Yes

Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: If needed, responses are reversed so that all are increasing in the pro-female-employment position but otherwise are not standardized. Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Appendix Table A.8: Summary statistics, gender-role attitudes (NLSY79)

	Men	Women
<i>–Gender Norms (Higher = More Liberal, Scale: 1-5)</i>		
A woman’s place is in the home, not in the office or shop	3.04	3.37
A wife who carries out her family responsibilities doesn’t have time for a job	2.86	3.10
A working wife feels more useful than one who doesn’t hold a job (Rev.)	2.74	2.51
The employment of wives leads to more juvenile delinquency	2.87	2.98
It’s better if the man is the achiever, the woman takes care of the family	2.72	2.90
Men should share the work around the house with women (Rev.)	3.12	3.35
Women are happier if they stay home and take care of their children	2.75	2.92
Omnibus gender norms (higher = more liberal)	-0.46	0.96
<i>–Average # of Pre-/Post- Observations per Person</i>		
Number of omnibus gender observations, pre-baby	2.28	2.13
Number of omnibus gender observations, post-baby	1.51	1.69
Number of respondents	1303.00	1456.00

Source: National Longitudinal Survey of Youth, 1979. For more detail see Section 2.2 and Appendix B.3.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Statistics are weighted by the probability of selection. Individual component attitudes are on a scale of 1 to 5 (1= strongly agree; 5 = strongly disagree) and reversed so that higher values imply more liberal, less traditional attitudes toward gender roles. The omnibus gender-roles index takes each component question, standardizes it to have mean zero and standard deviation one, and sums across all questions. Question wording is altered for readability, see B.3.5 for exact wording.

Appendix Table A.9: Effect of first and second birth on gender-role attitudes, UK

	Dependent variable: Omnibus gender-roles index							
	Women						Men	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-baby	-0.921*** [0.158]	-1.137*** [0.236]	-1.178*** [0.225]	-0.853*** [0.149]	-0.874*** [0.158]	-0.880*** [0.226]	-0.250 [0.158]	-0.174 [0.157]
2+ kids					-0.460*** [0.170]	-0.832*** [0.249]		-0.735*** [0.188]
Pre-baby mean, dep. var.	1.168	1.168	1.168	1.186	1.168	1.168	-0.0679	-0.0679
Observations	4,663	4,663	4,663	6,414	4,663	4,663	4,366	4,366
+ Indiv. lin. time trend	No	Yes	No	No	No	Yes	No	No
+ Drop person FE	No	No	Yes	No	No	No	No	No
+ Control Group	No	No	No	Yes	No	No	No	No

Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Appendix Table A.10: Effect of first birth and second birth on gender-role attitudes, US (NLSY79)

	Dependent variable: Omnibus gender-roles index							
	Women						Men	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-baby	-0.966*** [0.244]	-0.980*** [0.247]	-1.639*** [0.234]	-0.883*** [0.191]	-0.922*** [0.249]	-0.935*** [0.251]	-0.498** [0.251]	-0.451* [0.255]
2+ kids					-0.198 [0.278]	-0.200 [0.278]		-0.214 [0.285]
Pre-baby mean, dep. var.	0.881	0.881	0.881	0.845	0.881	0.881	-1.077	-1.077
Observations	5,493	5,493	5,493	8,545	5,493	5,493	4,834	4,834
+ Individ. lin. time trend	No	Yes	No	No	No	Yes	No	No
+ Drop person FE	No	No	Yes	No	No	No	No	No
+ Control Group	No	No	No	Yes	No	No	No	No

Source: National Longitudinal Survey of Youth (NLSY79). For more detail see Section 2.2 and Appendix B.3.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Appendix Table A.11: Expected to work this year but wrong, robustness checks, U.K. men and women

	Dependent variable: Expected to work but wrong							
	Women						Men	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-baby	0.160*** [0.0279]	0.144*** [0.0181]	0.148*** [0.0179]	0.143*** [0.0256]	0.161*** [0.0279]	0.141*** [0.0189]	-0.00471 [0.0130]	-0.00508 [0.0129]
2+ kids					0.0314 [0.0297]	0.00866 [0.0147]		-0.00910 [0.0134]
Pre-baby mean, dep. var.	0.0744	0.0744	0.0744	0.0678	0.0744	0.0744	0.0280	0.0280
Observations	3,785	3,785	3,785	5,314	3,785	3,785	4,122	4,122
+ Individ. lin. time trend	No	Yes	No	No	No	Yes	No	No
+ Drop person FE	No	No	Yes	No	No	No	No	No
+ Control Group	No	No	No	Yes	No	No	No	No

Source: British Household Panel Survey. For more detail see Section 2.1 and Appendix B.1.

Notes: Sample comprises all original sample members for whom we observe the event of childbirth. Individuals must also be observed during prime child-bearing years. Age in years, calendar year, and person fixed effects included. Regressions are weighted by the probability of selection. Standard errors are clustered at the person level and reported in brackets. Because this question is only asked of individuals if they currently have a paid job, we restrict the sample to years up to and including the year after the birth of the first child, since women are less likely to be working after having a child, and this creates a selected sample post-baby and makes interpretation difficult. Age-in-years, calendar-year, and person fixed effects included in all regressions unless otherwise noted. Standard errors are clustered at the person level and reported in brackets.

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Appendix Table A.12: Summary statistics for supplementary data sources asking about expectations of future labor supply

	MTF	Gallup	Purdue
White	0.69	0.93	.
Black	0.13	0.07	.
Married	0.10	0.01	.
Expect to be homemaker at 30	0.03	0.60	0.56
Professional w/ PhD or equiv. at 30	0.21	.	.
Professional w/out PhD at 30	0.38	.	.
Def. will get college degree	0.53	0.54	0.21
Choose to have 2 or more kids	0.89	0.97	.
Choose to have 3 or more kids	0.42	0.79	.

Notes: MTF data is from the 12th grade sample of the 1978-2014 Monitoring the Future Surveys. Current marital status and fertility expectations are obtained from Form 3 of the MTF. Education expectations, employment expectations and race are from Form 4 of the MTF. Gallup data is from a 1961 Gallup opinion survey of high school and college students. We retain only respondents in 12th grade. Purdue data is from 1961 and 1962 Purdue Opinion Panels of high school students. Again, we retain only students in 12th grade. Out of worries about anchoring bias, we use only the first wave, in 1961.

Appendix B. Data Appendix

B.1 BHPS

The BHPS started in 1991 with 5,500 households (10,300 individuals), known as the “original sample.” Since then, the sample has grown in various ways. First, the main survey is only administered to adults, so children in original sample households “age in” to the survey upon turning 16. These “agers-in” and all other original sample members are continuously tracked as they form new households, and the individuals in these new households are also interviewed so long as they live with an original sample member. (If these new-comers move out of the household, they are no longer tracked.) One notable exception is that some will “marry in” to the sample and be tracked permanently if they have a child with an original sample member. Children born to original sample members are automatically counted as original sample members.

The BHPS has also grown via the addition of new booster samples. Between 1997 and 2001, the BHPS roughly doubled in size with the addition of 5,000 households (8,600 individuals) out of concerns that the original sample under-represented Scotland, Wales, and Northern Ireland. In 2009, the BHPS became subsumed under the “Understanding Society” dataset, which, including BHPS households, tracked roughly 40,000 households (50,000 individuals) in the UK. While Understanding Society maintained most questions from the BHPS, it notably dropped several of the gender-roles questions we are most interested in, so our main results focus on the BHPS through 2008, for the sake of consistency across our main outcome variables. Results that include data from Understanding Society are not much different from results using BHPS through 2008 only and are available upon request.

Our preferred analysis sample imposes the restriction that we observe each person at least once before becoming a parent and at least once after, so that all observations help identify the main effect of parenthood. Our sample comprises only those who were original sample members, *i.e.* those who were interviewed in the first wave and *not* those who marry into the sample later on. Because our panel is relatively short (18 waves) compared to our U.S. panel data sources, we exclude children who age into the sample because, in order to meet our restriction that we see the event of the birth of their first child, these individuals would disproportionately have children early (and, related, would be disproportionately female, as women have children at a younger age than men).

The BHPS conducts fertility history interviews for all respondents at the second wave after entry into the sample, and from this set of questions, we can identify parents who have had children before the BHPS starts. Unfortunately, this fertility history is not asked again after the second wave, so we look to the household grid, which lists all individuals who live in a household together as well as some basic demographic information, such as date of birth. Importantly for our purposes, this grid also includes relationships between household members, so we can link children to their parents and construct a full fertility history for each parent in our sample. This is largely the same method suggested by the BHPS. Unfortunately, because a parent must be living with their child in order to identify birth dates, we undercount fathers, as (we hypothesize) that they on average are less likely to be living with their children than are mothers. Relative to a BHPS quality profile conducted at wave 13, we slightly over-count mothers (692 versus 673) and undercount fathers (567

versus 653).

Our main employment outcome variables are defined as follows. Employed last week is equal to 1 if an individual did paid work last week (i.e. was physically present at work) and excludes people who, for example, were on maternity leave. Hours worked is the *typical* hours that an individual works within a given year, and is coded as zero if the individual is not working (including not looking for work). Thus, an individual may be on maternity leave in the week prior to a survey, but their usual hours worked might be greater than zero, since a respondent can interpret the question to be about when their hours prior to maternity leave. Other employment variables we use include job last week (which includes individuals who had a paid job but for whatever reason were away from it), in labor force currently (which includes individuals who are unemployed and looking for a job), and employed full/part time (which is defined as \geq and $<$ 30 hours, respectively). Finally, we also examine whether an individual's *spouse* is employed in a given year. Note that these individuals are excluded from our sample, as they are not randomly chosen to be included.

We also examine a range of questions on an individuals' opinions on gender roles in the household. Appendix Table A.6 summarizes these questions. While 9 questions are available in total, we only choose a subset of 6 that pertain to the impacts of a woman working. These statements, asked on a scale of 1 (strongly agree) to 5 (strongly disagree) are: (1) A pre school child is likely to suffer if his or her mother works, (2) All in all, family life suffers when the woman has a full time job, (3) A woman and her family would all be happier if she goes out to work, (4) Both the husband and wife should contribute to the household income, (5) Having a full-time job is the best way for a woman to be an independent person, and (6) A husband's jobs is to earn money; a wife's job is to look after the home and family.

We take the individual gender-roles questions and construct an omnibus measure as follows. First, since some questions are phrased in a gender "liberal" manner (i.e. men and women should take on equal roles in the household) and others are phrased in a gender "conservative" manner (i.e. men and women should take on distinct roles in the household), we reverse some questions as necessary so that all are increasing in the gender liberal direction. We then standardize each variable so that they have a mean of zero and standard deviation of one, and sum them up.

Another set of questions we examine in the BHPS relate to expectations of future labor force. The question that we use about labor force expectations is worded as follows: "Do you think this actually will happen in the coming twelve months? Give up paid work." Note that this question is only asked of individuals who currently have a paid job. We examine this question, combining yes and "don't know," and compare to actual labor force outcomes, conditional on answering the question in the previous wave. These variables are always keyed on the year the expectation was made *about*, so that in event year zero, the expectation variable is made in the year before an individual had their first child. Finally, we construct a variable that asks whether your prediction was correct. This variable (expected to work this year but wrong) is equal to 1 if an individual predicted that she would work this year and was correct and is equal to 0 if she predicted that she would work this year and was wrong. Note that this variable is missing for individuals who expected *not* to work this year. In practice, this group is only 3% of total answers.

A final set of questions we examine in the BHPS relate to job satisfaction. Like the gender-roles questions, we examine each individually and construct an omnibus variable.

The individual questions are asked on a scale of 1 (not satisfied at all) to 7 (completely satisfied), and pertain to: (1) overall satisfaction, (2) total pay, (3) security, (4) the work itself, and (5) hours worked. Like the gender-roles questions, we standardize these variables so that they have a mean of zero and standard deviation of one and sum them all together to form an omnibus measure. Note that more questions were available, but they were not asked after wave 7, so for the sake of consistency in our omnibus measure, we selected only those questions that were asked in all waves.

B.2 NLSW 1968

B.2.1. Survey Design and Sample

We make use only of the “young women” component of the NLSW68. (While the National Longitudinal Surveys of Young and Mature Women include an older “mature women” sample, they have for the most part had their first child before entering the sample.) The “young women” component of the NLSW68 is a sample of 5,159 women born between 1941 and 1954. The women were aged 14-24 when first interviewed in 1968 and aged 49-62 when last interviewed in 2003. During the last wave of the survey 55.4% of the original sample was interviewed.

The sample is represented by a multi-stage probability sample drawn by the Census Bureau from 1,900 primary sampling units. In order to provide reliable statistics for black respondents, the NLSW68 oversampled the black population at about twice the expected rate of the total population. Black respondents make up approximately 28 percent of the sample, compared to 11 percent of the population at the time. Probability-of-selection weights created by the NLSW68 correct for this bias.

B.2.2. Sample Selection

Our sample comprises all people in the National Longitudinal Survey of Women who meet the following criteria: (1) We observe them at least once before and after the birth of their first child, (2) We observe them for the first time by age 20 and for the last time after age 40. We include people who drop out at some point in the sample.

In a given year, it is possible to observe a respondent’s employment status, but not the number of hours they worked last week, or vice versa. Thus, for a respondent to be included in our sample, we must observe each of the three outcomes variables (“working now”, “in labor force now”, “hours worked last week”).

B.2.3. Employment Status

Employment status is obtained from combining two NLSW68-created variables: employment status recode (ESR), available every 1-3 years from 1968 to 1993, and monthly labor recode (MLR), available every other year from 1995 to 2003. Although not exactly the same, the two variables are very similar. Employment status recode specifies whether the respondent was working, with a job but not at work, laid off and looking for work, unemployed, going to school, keeping house, unable to work, or retired. Monthly labor recode specifies whether

the respondent was employed and at work, employed but absent, unemployed and laid off, unemployed and looking for work, retired, or disabled.

We code respondents as “working now” if they report working (employed and at work) and as “not working” if they report any other activity in the ESR (MLR). Negative values of each variable, which are assigned when respondents refuse to answer, do not know the answer, or are not interviewed, are coded as missing.

We code respondents as “in labor force” if they report working, having a job but not working, looking for work, or being laid off (employed and at work, employed but absent, unemployed and laid off, unemployed and looking for work) and as “not in labor force” if they report any other type of activity in the ESR (MLR). Negative values of each variable, which are assigned when respondents refuse to answer, do not know the answer, or are not interviewed, are coded as missing.

B.2.4. Hours Worked

Between 1968 and 1993, the NLSW68 asks respondents the number of hours they worked in the previous week. We assign this variable to be “hours worked.” Respondents deemed to be not working (“working now” = 0) are assumed to have worked zero hours last week.

After 1993, the NLSW68 asks respondents the number of hours worked at their main job and at their other job. We assign our “hours worked” variable to be the sum of these two variables. Again, respondents deemed to be not working (“working now” = 0) are assumed to have worked zero hours last week.

B.2.5. Date of Birth of 1st and 2nd Child

The NLSW68 does not consistently ask respondents for the date of birth of their children. Thus, to determine the date of birth of respondents’ first and second children, we combine several categories of variables pertaining to children’s date of birth and take the minimum and second minimum value among these.

Between 1968 and 1993, the NLSW68 asks respondents directly about the date of birth of their children. In 1973, the NLSW68 records the date of birth of each child a respondent has had in their lifetime. In 1978, 1983, 1985, 1987, 1988, 1991, and 1993 the NLSW68 records the date of birth of each child the respondent has had since their last interview.

Post 1993, the NLSW68 does not explicitly ask for dates of birth of children, but instead asks for a demographic information, including date of birth, for each member of the household. We include these variables in our calculation of minimum date of birth only if the household member is a child of the respondent (relationship to respondent equal to child, son, or daughter).

To determine the first (second) child’s date of birth, we take the minimum (second minimum) date of birth among all date of birth variables described above. Children with missing months of birth but non-missing years of birth are assumed to have been born in September (the most common birth month).

B.2.6. Event Time

The event time variable is calculated by determining the number of months between the date of birth of the respondent's first child and the current interview, dividing by 12, and rounding to the nearest whole number.

B.2.7. Other Demographic Information

Education: Each year, the NLSW68 records the highest completed grade and highest college degree obtained by respondents. We assign the highest grade completed to be the respondent's maximum highest grade reported among all years. Similarly, we assign the highest college degree to be the respondent's maximum reported degree. Respondents who report completing less than 16 years of school are coded as having no college degree.

Age: Age is calculated by determining the number of months between the respondent's date of birth and the current interview, dividing this number by 12, and rounding down. Unknown interview months are assumed to be March.

Mom Worked: In 1968, respondents were asked whether their mother worked when they were 14. Responses of yes are coded as 1, responses of no are coded as 0, and all other values are coded as missing.

Planned to Work at Age 35: In 1968, respondents were asked what they plan to be doing at age 35. Responses of working at a different job and working at the same job are coded as 1. Responses of married, keeping house, other, or do not know are coded as 0. All other values are coded as missing.

B.3 NLSY 1979

B.3.1. Survey Design and Sample

The NLSY79 is a sample of 12,686 men and women born between 1957 and 1964. They were aged 14-22 when first interviewed in 1979 and aged 50-57 when last interviewed in 2014.

The NLSY79 purposefully over samples economically disadvantaged, Hispanic, and black youth. Approximately 25 percent of the sample is black and 16 percent of the sample is Hispanic, compared to their share of population which was around twelve and six percent, respectively, at the time. The inclusion of probability-of-selection weights corrects for this over sampling.

B.3.2. Sample Selection

Our sample comprises all people in the National Longitudinal Survey of Youth who meet the following criteria: (1) We observe them at least once before and after the birth of their first child, (2) We observe them for the first time by age 20 and for the last time after age 40. We include people who drop out at some point in the sample.

In a given year, it is possible to observe a respondent's employment status, but not the number of hours they worked last week, or vice versa. Thus, for a respondent to be included

in our sample, we must observe each of the three outcomes variables (“working now”, “in labor force now”, “hours worked last week”). For regressions on gender-role attitudes, we must observe the respondent’s answers to each of the eight questions to be included in the sample.

B.3.3. Employment Status

The NLSY79 records respondents’ employment status for each week between 1978 and the date of their last interview. To determine a respondent’s current employment status, we use the employment status variable corresponding to week the respondent was interviewed.

We code respondents as “working now” if their employment status is a job number (values of 100 through 3000) as as “not working” if they report not working, being unemployed, being out of the labor force, or being in active military service. All other values are coded as missing.

We code respondents as “in labor force” if their employment status is a job number or they report being unemployed and as “not in labor force” if they report being out of the labor force or in active military service. All other values are coded as missing.

B.3.4. Hours worked

As with employment status, the NLSY79 records the number of hours respondents worked for each week between 1978 and the date of their last interview. We assign this variable, in the week-year corresponding to each interview, to be “hours worked.” Respondents deemed to be not working (“working now” = 0) are assumed to have worked zero hours in the previous week.

B.3.5. Gender-role attitudes

NLSY79 asks respondents about their attitudes towards women working four times over the course of the survey – in 1979, 1982, 1987, and 2004. In particular, they ask if respondents personally agree or disagree with each of the following statements

- a: A woman’s place is in the home, not in the office or shop
- b: A wife who carries out her full family responsibilities doesn’t have time for outside employment
- c: A working wife feels more useful than one who doesn’t hold a job
- d: The employment of wives leads to more juvenile delinquency
- e: It’s much better for everyone concerned if the man is the achiever outside the home and the woman takes care of the home and family
- f: Men should share the work around the house with women, such as doing the dishes, cleaning, and so forth
- g: Women are much happier if they stay at home and take care of their children

We recode each question so that a higher value indicates a more gender-liberal position. Non-interviews as well as responses of “do not know” and “refused” are coded as missing.

B.3.6. Date of Birth of 1st and 2nd Child

The NLSY79 provides the date of birth for each of the respondent’s children. We take the minimum (second minimum) date of birth among all children as the birthday of the first (second) child.

B.3.7. Event Time

NLSY79 provides the date of interview for each wave. The event time variable was calculated by determining the number of months between the birth of the respondent’s first child and the current interview, dividing by 12, and rounding to the nearest whole number.

B.3.8. Other Demographic Information

Education: Each year, the NLSY79 records the highest completed grade. We assign the highest grade completed to be the respondent’s maximum highest grade reported among all years. A respondent is assumed to have a college degree if they have completed 16 or more years of schooling.

Age: Age is calculated by determining the number of days between the respondent’s date of birth and the current interview, dividing this number by 365, and rounding down. Unknown interview months are assumed to be March and unknown interview days are assumed to be the 15th.

Mom Worked: In 1979, respondents were asked whether their mother worked at age 14. Responses of yes are coded as 1, responses of no are coded as 0, and all other values are coded as missing.

Planned to Work at Age 35: In 1979-1984 respondents were asked what they plan to be doing at age 35. Responses of working at a different job and working at the same job are coded as 1. Responses of married, keeping house, other, or do not know are coded as 0. All other values are coded as missing.

B.4 PSID

B.4.1. Survey Design and Sample

The PSID started in 1968 with a national sample of about 4,800 U.S. households. This initial sample was made up of two sub-samples. The first was a cross-sectional national sample conducted by the Survey Research Center (SRC) that produced about 2,900 interviews. The second was a sample of low-income families conducted by the Survey of Economic Opportunity (SEO) that produced about 1,900 interviews. The low income sample had unequal selection probabilities as it was limited to SMSAs and non-SMSAs in the southern region. By construction, the PSID oversamples low-income households.

In 1997 and 1999 the PSID added an immigrant booster sample of about 500 households. Those eligible for this sample must have immigrated to the US after 1968 or have been born 1969 or later to people not in the US in 1968.

The PSID includes probability-of-selection weights that are designed to adjust for the oversampling of low-income households and for differential attrition. They are especially important given the combination of SEO and SRC samples that made up the original PSID sample. Weights are also adjusted every five years (1969, 1974, etc) for cumulative panel attrition.

PSID follows sample members no matter their living arrangements: if they split off and form a new household, then that household is added to the sample and each of its members are interviewed; if they have children, then each of their children are interviewed.

B.4.2. Sample Selection

Our sample comprises all people in the Panel Study of Income Dynamics who meet the following criteria: (1) We observe them at least once before and after the birth of their first child, (2) We observe them for the first time by age 20 and for the last time after age 40, and (3) They either begin in or are born into a household in the “core sample” – that is, a household randomly selected to be included in the sample. We include those who are added later on in “booster” samples that increase representation among immigrant and Latino populations and include people who drop out at some point in the sample.

In a given year, it is possible to observe a respondent’s employment status, but not the number of hours they worked last week, or vice versa. Thus, for a respondent to be included in our employment sample, we must observe each of the two outcomes variables (“working now” and “in labor force now”).

B.4.3. Employment Status

Employment status in the PSID is obtained from the individual employment status variables, available from 1979 to 2015. The variable specifies whether the respondent is working now, temporarily laid off, unemployed and looking for work, retired, permanently disabled, keeping house, or a student.

We code respondents as “working now” if they report working now and coded as “not working” if they report being temporarily laid off, unemployed and looking for work, retired, permanently disabled, keeping house, or a student. All other values are coded as missing.

We code respondents as “in labor force” if they report working now, being temporarily laid off, unemployed and looking for work. They are coded as “not in labor force” if they report retired, permanently disabled, keeping house, or a student. All other values are coded as missing.

B.4.4. Date of Birth of 1st and 2nd Child

To determine the date of birth of respondent’s first and second child, we use the childhood and adoption history file. We limit the data to birth records (as opposed to adoption records), so the date of birth of first and second children correspond only to birth children.

If the child's month of birth is missing, we assume they were born in September. If the child is listed as having been born in winter, spring, summer, or fall we assume they were born in January, April, July, and October, respectively. If the child's birth year is missing, we code the date of birth as missing. The minimum (second minimum) date of birth among all the respondent's births is taken to be the date of birth of the first (second) child.

B.4.5. Event Time

The event time variable is calculated by determining the number of months between the date of birth of the respondent's first child and the current interview, dividing by 12, and rounding to the nearest whole number.

B.4.6. Other Demographic Information

Education: In each survey wave, the PSID records the number of years of education a respondent has received. A respondent is coded as having a college degree if they ever report having completed 16 or more years of schooling.

Age: To determine age, we rely on the PSID variable which asks respondent's for their age at the time of the interview. We then calculated an imputed age – equal to the number of months between the respondent's date of birth and current interview, divided by 12, and rounded down. If the reported age is more than two years from the imputed age or if the current age is lower than a previously reported age, we replace the age variable with our imputed age variable.

B.4.7. Childhood Development Supplement (CDS)

In 1997 PSID began to supplement its main data collection with a survey targeting 0-12 year old children and their parents. Original CDS respondents were interviewed an additional two times, in 2002 and 2007. In 2014, the CDS continued with a new cohort of 0-17 year old children and their parents.

CDS asks children's primary and second caregivers a variety of questions relating to parenthood. Primary caregivers are defined, in order of precedence, as the biological, step, foster, or adoptive mother of the child, the "wife" of a a PSID head who is father to the child, the biological, step, foster, or adoptive father of the child, the legal guardian of the child, and the adult in the household unit who takes primary responsibility for the child. Thus, primary caregivers are almost always women. Secondary caregivers (known as other caregivers after 1997), are defined, in order of precedence, as the biological, step, foster, or adoptive father of the child, the grandmother of the child, the boyfriend or girlfriend of the primary caregiver, another adult relative of the child, and another adult non-relative of the child.

Our analysis uses the "parenting is hard" question in the CDS. The question asks caregivers to rank on a 5 point scale how much they agree or disagree with the statement "being a parent is harder than I thought it would be." A value of 1 corresponds to "strongly disagree" whereas a value of 5 corresponds to "strongly agree." The "parenting is hard" question is

asked to both primary and secondary caregivers in 1997, 2002, and 2007 but only to primary caregivers in 2014.

B.5 Supplementary data sources asking expectations of future labor supply

The Monitoring the Future (MTF) survey is an annual cross-sectional survey of American youth. It was first conducted in 1975 and includes about 16,000 high school seniors from approximately 133 public and private high schools in the contiguous United States. Beginning in 1991, the MTF added nationally representative samples of 8th and 10th graders, although we do not make use of them in this paper.

The MTF uses a multi-stage random sampling procedure to create a nationally representative sample of students from each grade. In the first two stages, researchers select particular geographic areas and (with probability proportionate to size) one or more schools in each area. The last stage involves selecting students within each school. In a given school, up to 350 students can be selected. Typically, in schools with fewer than 350 students, all students are included. In larger schools, entire classrooms are randomly sampled.

Our analysis makes use of a question that asks respondents the type of work they expect they will be doing at age thirty. Specifically, the question asks, “*What kind of work do you think you will be doing when you are 30 years old? Mark the one that comes closest to what you expect to be doing.*” There are fifteen possible occupational categories that students can choose from, with “home-maker” as one possible option. Other possible choices include laborer; service worker; operative or semi-skilled worker; sales clerk in a retail store; clerical or office worker; craftsman or skilled worker, sales representative; protective service; manager or administrator; professional without doctoral degree; professional with doctoral degree or equivalent; owner of small business; farm owner/manager; and military service. Below is a list of example occupations provided in the MTF questionnaire for each occupational category.

Job Group	MTF 1976	MTF2016
Laborer	car washer	custodian
	sanitary worker	material mover
	farm laborer	maid landscape worker farm worker
Service worker		food preparer or food service worker including fast food
	cook	walter/waitress
	waiter	call center worker
	barber	stock clerk
	janitor	order filler
	gas station attendant	nursing aide/orderly
	practical nurse beautician	teacher assistant childcare worker
Operative or semi-skilled worker	garage worker	bus or truck driver
	taxicab	maintenance or repair worker
	bus or truck driver	assembly line worker
	assembly line worker welder	
Sales clerk in a retail store	shoe salesperson	cashier
	department store clerk drugstore clerk	supervisor of retail workers
Clerical or office worker		secretary
	bank teller	receptionist
	bookkeeper	bookkeeper
	secretary	supervisor of office workers
	typist	bank teller
	postal clerk or carrier ticket agent	postal clerk or carrier
Craftsman or skilled worker		carpenter
	carpenter	mechanic
	electrician	machinist
	brick layer	welder
	mechanic	
	machinist	
	tool and die maker telephone installer	

Job Group	MTF 1976	MTF2016
Sales representative	insurance agent real estate broker bond salesman	insurance agent real estate
Protective service	police officer fireman detective	police firefighter paramedic
Manager or administrator	office manager sales manager school administrator government official	office manager government official sales manager
Professional without doctoral degree	registered nurse librarian engineer architect social worker technician accountant actor artist musician	registered nurse school teacher accountant architect artist information technology worker
Professional with doctoral degree or equivalent	lawyer physician dentist scientist college professor	lawyer physician dentist scientist college professor
Owner of small business	restaurant owner shop owner	
Farm owner, farm manager		
Military service		
Full-time homemaker or housewife		

We also make use of the 1961 wave of a two year Purdue Opinion Panel of high school students grades 9-12. We restrict our sample to female high school seniors. The panel asks “What kind of job do you expect to have 20 years from now?” and respondents can choose between twelve possible occupations, including “housewife.” Other possible occupations include salesman; owner of factory or small business; professional-teacher, doctor, musician, musician, scientist; office worker; clerk in a store; farm or ranch owner; farm or construction laborer; factory worker or mechanic; big business management; deliveryman, truck driver; and carpenter, plumber, electrician. Raw data and documentation for this survey can be found at <https://ropercenter.cornell.edu/CFIDE/cf/action/catalog/abstract.cfm?type=&start=&id=&archno=USPOP1961-063&abstract=>.

Finally, we use a 1961 Gallup Opinion Poll of young high school, college, and non students. We again restrict our sample to female high school seniors. The Gallup poll asks respondents the open-ended question “What do you expect to be doing when you are 40 years old?” Responses were organized into 18 categories. We use the “house wife,” “home maker,” “house work,” and “raising children” as our definition of expected homemaker. Other categories include teacher, sports coach, professor; scientist, physicist, biologist, chemist, medical research, psychologist, research; business executive, own business, industry, management, business administration; minister, missionary, social worker; engineering, research engineer, managing engineer; entertainment, actor, broadcasting; medicine, dentist, psychiatry; lawyer; farmer, rancher, agriculture; armed services; government work; sales, clerical: secretary, sales clerk, office worker; nurse; mechanic, machinist, tool and dye maker; skilled trade: electrician, plumber, carpenter, mason, electronics; contracting, building construction, excavating; and artist, cartoonist, designer, draftsman, decorator. Raw data and documentation for this survey can be found at <https://ropercenter.cornell.edu/CFIDE/cf/action/catalog/abstract.cfm?type=&start=&id=&archno=USAIPSPPOS1961-544&abstract=>.

Appendix C. Full model with proofs

C.1 Overview

In our model, a young woman forms an estimate of how hard it will be to take care of children while working, which informs her educational decisions. She forms this estimate by observing her own mother, but this estimate is subject to two types of forecasting errors. First, her mother’s generation might have higher or lower mean costs of motherhood than her generation. Second, relative to the rest of her cohort, her own mother may have had idiosyncratic (mean-zero) higher or lower costs of motherhood. While the daughter inherits part of her mother’s costs, neither of these two error terms gets passed on to her.

C.1.1. Assumptions and set-up

Let utility $u(c, h)$ be a quasi-linear function of consumption c and labor h (for hours worked, say). Specifically, assume that

$$u(c, h) = c - \frac{h^{\gamma+1}}{\gamma + 1},$$

where $\gamma > 0$.

Women’s consumption will be equal to market wages net of employment costs of motherhood (both per hour) times labor supply (in hours). Market wages are $\tilde{w} = w + \beta \cdot e$, where w is a base wage for those with higher education $e = 0$ and $\beta > 0$ is the hourly return to higher education $e = 1$. Gaining education $e = 1$ costs some “tuition” $\alpha > 0$, while $e = 0$ is free.

We view “employment costs of motherhood” very broadly, as any cost mothers endure during work. These might include, for example, the per-hour cost of a nanny or day-care service or the emotional or psychic cost of being separated from the child while at work. Actual employment costs for a given woman i are μ_i , but she predicts this cost with some error δ_i . In particular, she observes her mother’s employment costs, $\mu_i + \delta_i$, but she only inherits the μ component. The δ component was her mother’s “luck” (good or bad) and is *not* passed down to the daughter. This luck is the sum of a generational component and an idiosyncratic component. For example, her own mother’s idiosyncratic good luck might include having had a very understanding supervisor at work, or having parents or in-laws who live in close proximity and can provide free child care. Generational “good luck” could be, for example, giving birth in a moment with cheap, technological substitutes for child care. But the woman in our problem cannot distinguish between μ_i and δ_i and instead uses the sum as the best signal of her own, future employment costs of motherhood.³⁵

Assume further that $\mu \sim U[0, 1]$ and $\delta = \{\lambda - \epsilon, \lambda + \epsilon\}$ with equal probability, where λ and $\epsilon > 0$ are both constants. Note that λ can take negative values (for much of the analysis below, λ will drop out, but allowing $\mathbb{E}(\delta) \neq 0$ will be useful later). We assume μ and δ are independent. For simplicity, we further assume no taxes and that w is sufficiently large

³⁵Note also that women in our model are *naive* in the sense that they do not take into account that δ is drawn from a distribution and instead just assume employment costs will be $\mu + \delta$ with probability one.

that $w - \mu - \delta$ is always greater than zero (so mothers' predicted effective wages are always positive).

We abstract away from many possible dimensions of heterogeneity. There is no variation in fertility in our model: women become mothers with probability one and we ignore issues related to delaying or timing childbirth. Unlike many models of human capital, we ignore any variation in "ability" and thus the effective return to and cost of education do not vary along this dimension.

The implicit timing of the model is as follows: A woman makes her education decisions assuming that her future employment costs of motherhood are $\mu + \delta$, but then will make her actual labor supply decisions after this uncertainty is resolved and her actual costs, μ , are revealed to her. While this timing is helpful to keep in mind, in fact the problem can be collapsed to a single decision: the woman's (only) problem is to optimally choose education $e \in \{0, 1\}$. Once she chooses e and once her true costs μ are revealed to her, her labor supply is given by a simple optimization problem and thus can be viewed as deterministic. Note that, as we are assuming away many other dimensions of heterogeneity, variation in optimal e comes entirely from variation in expected employment costs, $\mu_i + \delta_i$.

Working backward, the woman calculates her predicted utility for each $e \in \{0, 1\}$. This comparison requires calculating the optimal predicted h^* given her assumptions about child-care costs for each e and then seeing which realized utility is higher. For a given e and prediction of child-care cost $\mu_i + \delta_i$, the optimal predicted \hat{h} is given by setting

$$\frac{\partial}{\partial h} \left[(w + \beta e - \mu - \delta)h - \alpha e - \frac{h^{\gamma+1}}{\gamma + 1} \right] = 0,$$

which yields the following *predicted* labor-supply function:

$$\hat{h}^* = (w + \beta e - \mu - \delta)^{\frac{1}{\gamma}}. \quad (2)$$

Note that, as we would expect, predicted hours will increase in wages (and thus education, as education increases wages) and decreases in hourly child-care costs.³⁶

The woman takes her optimal predicted \hat{h}^* for each of the two possible levels of education $e \in \{0, 1\}$ and determines which yields higher utility. She thus compares the following two expressions:

$$u(\hat{h}^*(e = 0)) = \underbrace{(w - \mu - \delta)}_{\substack{\text{effective wage} \\ \text{if } e = 0}} \cdot \overbrace{(w - \mu - \delta)^{\frac{1}{\gamma}}}^{h^*(e=0)} - \frac{\left[(w - \mu - \delta)^{\frac{1}{\gamma}} \right]^{1+\gamma}}{1 + \gamma} \quad (3)$$

and

³⁶As utility is quasi-linear, there is no income effect and thus the substitution effect dominates and labor supply is always a positive function of wages, consistent with a long line of empirical work suggesting that women (relative to men) increase labor supply more in response to increases in effective wages.

$$u(\hat{h}^*(e=1)) = \underbrace{(w + \beta - \mu - \delta)}_{\text{effective wage if } e=1} \cdot \underbrace{(w + \beta - \mu - \delta)^{\frac{1}{\gamma}}}_{h^*(e=1)} - \alpha - \frac{\left[(w + \beta - \mu - \delta)^{\frac{1}{\gamma}} \right]^{1+\gamma}}{1 + \gamma}. \quad (4)$$

After some algebra, it is easy to show that expression (4) will be greater than (3) (i.e., the woman will choose $e = 1$) iff:

$$(w + \beta - \mu - \delta)^{\frac{1+\gamma}{\gamma}} - (w - \mu - \delta)^{\frac{1+\gamma}{\gamma}} > \alpha \frac{1 + \gamma}{\gamma}. \quad (5)$$

Intuitively, a higher tuition cost α will discourage choosing $e = 1$. As $\frac{1+\gamma}{\gamma} > 0$ (in fact, it is greater than 1) by assumption, the LHS of the equation is increasing in β so more women will choose $e = 1$ as the return to education rises (also intuitive).

Less intuitive but important is that the LHS is decreasing in $\mu + \delta$. This claim follows from $\frac{1+\gamma}{\gamma} > 1$ and can be easily shown by taking the derivative of the expression with respect to $\mu + \delta$.

C.2 Predictions yielded by the model

Claim 1. *Women who choose $e = 1$ will on average work more (post-baby).*

Proof. We first prove the following useful lemma:

Lemma 1. *For a given set of parameter values $(w, \beta, \gamma, \alpha, \lambda, \epsilon)$ for which (the technical condition) $\beta^{\frac{1+\gamma}{\gamma}} < \alpha^{\frac{1+\gamma}{\gamma}}$ and a given realization of the error term $\delta = \{\lambda - \epsilon, \lambda + \epsilon\}$, there exists some μ' such that a woman will choose $e = 1$ iff $\mu < \mu'$ when $\delta = \lambda - \epsilon$ and another such cut-off μ'' such that for all $\mu < \mu''$ she will always choose $e = 1$ when $\delta = \lambda + \epsilon$.*

Proof. First, fix $\delta = \lambda + \epsilon$. The LHS of equation (5), $(w + \beta - \mu - (\lambda + \epsilon))^{\frac{1+\gamma}{\gamma}} - (w - \mu - (\lambda + \epsilon))^{\frac{1+\gamma}{\gamma}}$, is continuous and strictly decreasing in μ . It also tends to infinity as μ tends to negative infinity. It reaches its minimum (real) value, $\beta^{\frac{1+\gamma}{\gamma}}$, as $\mu \rightarrow (w - (\lambda + \epsilon))$. As $\beta^{\frac{1+\gamma}{\gamma}} < \alpha^{\frac{1+\gamma}{\gamma}}$ by assumption, the LHS tends to infinity as $\mu \rightarrow -\infty$, then by the intermediate value theorem we know there exists some μ'' such that the LHS and RHS of (5) are equal. Finally, as the LHS is strictly decreasing in μ and the RHS is constant, then this μ'' is unique and equation (5) holds iff $\mu < \mu''$.

A parallel argument establishes a unique value for μ' .

Finally, note that the LHS of equation (5) is everywhere higher when $\delta = \lambda - \epsilon$ than when $\delta = \lambda + \epsilon$, so it will equal $\alpha^{\frac{1+\gamma}{\gamma}}$ for a strictly larger value of μ and thus $\mu'' > \mu'$.

While we have now proved the claim as it is stated, note that we have only shown the existence of some unique $\mu', \mu'' \in \mathbb{R}$, not necessarily $\mu', \mu'' \in [0, 1]$. We will heretofore always assume values for $(w, \beta, \alpha, \gamma, \lambda, \epsilon)$ such that $0 < \mu'' < \mu' < 1$ to rule out situations where everyone or no one gets educated. Figure C.1 provides a set of such parameter values and shows a graphical derivation of μ', μ'' . ■

We now return to the proof of Claim 1. Note that the actual labor supply is made *after* the education decision is made and after true employment costs of motherhood μ are revealed. It follows that, holding e and μ fixed, labor supply is simply:

$$h^* = (w + \beta e - \mu)^{\frac{1}{\gamma}}. \quad (6)$$

The claim requires us to show that $\mathbb{E}(h|e = 1) - \mathbb{E}(h|e = 0) > 0$, which can now be written (using equation 6 and Lemma 1) as:

$$\begin{aligned} & \mathbb{E}(h|e = 1) - \mathbb{E}(h|e = 0) = \\ & \mathbb{E}\left[(w + \beta - \mu)^{\frac{1}{\gamma}} \mid \left((\delta = \lambda - \epsilon) \cap (\mu < \mu')\right) \cup \left((\delta = \lambda + \epsilon) \cap (\mu < \mu'')\right)\right] - \\ & \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \left((\delta = \lambda - \epsilon) \cap (\mu > \mu')\right) \cup \left((\delta = \lambda + \epsilon) \cap (\mu > \mu'')\right)\right] = \\ & \frac{1}{2} \cdot \mathbb{E}\left[(w + \beta - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu'\right] + \frac{1}{2} \cdot \mathbb{E}\left[(w + \beta - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu''\right] - \\ & \left(\frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu > \mu'\right] + \frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu > \mu''\right]\right) > \\ & \left(\frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu'\right] + \frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu''\right]\right) - \\ & \left(\frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu'\right] + \frac{1}{2} \cdot \mathbb{E}\left[(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu''\right]\right) = \\ & \frac{1}{2} \left(\mathbb{E}(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu'\right] - \mathbb{E}(w - \mu)^{\frac{1}{\gamma}} \mid \mu > \mu'\right) + \\ & \frac{1}{2} \left(\mathbb{E}(w - \mu)^{\frac{1}{\gamma}} \mid \mu < \mu''\right] - \mathbb{E}(w - \mu)^{\frac{1}{\gamma}} \mid \mu > \mu''\right) > 0. \end{aligned}$$

The last step follows because as $\frac{1}{\gamma} > 0$, $(w - \mu)^{\frac{1}{\gamma}}$ is decreasing in μ and thus both terms in the large parentheses are positive. ■

Note that to connect this result directly to the empirical work, we need to add slightly more detail to the timing of the model. We have effectively collapsed the model into a single period (women make their education decision based on expected costs of motherhood, and then instantaneously have a child and have their costs revealed to them). If we instead assume that women all have a (perhaps brief) pre-baby period where everyone works some max “full-time” hours (which is roughly what we see in our data), then the claim maps not only into educated women having higher post-baby labor supply, but also exhibiting smaller employment “mommy effects.”

The next result states that, even though more educated women have smaller mommy effects with respect to employment, they are nonetheless the most “surprised” by the demands of motherhood.

Claim 2. *Women with $e = 1$ underestimate the costs of motherhood more than women with $e = 0$. That is, $\mathbb{E}[\delta \mid e = 1] < \mathbb{E}[\delta \mid e = 0]$.*

Proof. : From the Lemma, we know that $e = 1$ iff event a or event b occurs, where event a is $(\delta = \lambda + \epsilon) \cap (\mu < \mu'')$ and event b is $(\delta = \lambda - \epsilon) \cap (\mu < \mu')$. Recall that μ and δ are independent.

We can thus write:

$$\begin{aligned} \mathbb{E}[\delta \mid e = 1] &= \mathbb{E}[\delta \mid a] \frac{P(a)}{P(a) + P(b)} + \mathbb{E}[\delta \mid b] \frac{P(b)}{P(a) + P(b)} = \\ &= (\lambda + \epsilon) \frac{P(\delta = \lambda + \epsilon) \cdot P(\mu < \mu'')}{P(\delta = \lambda + \epsilon) \cdot P(\mu < \mu'') + P(\delta = \lambda - \epsilon) \cdot P(\mu < \mu')} + \\ &= (\lambda - \epsilon) \frac{P(\delta = \lambda - \epsilon) \cdot P(\mu < \mu')}{P(\delta = \lambda + \epsilon) \cdot P(\mu < \mu'') + P(\delta = \lambda - \epsilon) \cdot P(\mu < \mu')} = \\ &= (\lambda + \epsilon) \frac{\frac{1}{2}\mu''}{\frac{1}{2}\mu'' + \frac{1}{2}\mu'} + (\lambda - \epsilon) \frac{\frac{1}{2}\mu'}{\frac{1}{2}\mu'' + \frac{1}{2}\mu'} = \\ &= \lambda \left(\frac{\mu'' + \mu'}{\mu'' + \mu'} \right) + \epsilon \left(\frac{\mu'' - \mu'}{\mu'' + \mu'} \right) = \lambda + \epsilon \left(\frac{\mu'' - \mu'}{\mu'' + \mu'} \right). \end{aligned}$$

Note that the expression multiplying ϵ is negative, as $\mu'' < \mu'$.

We follow a parallel argument to simplify $\mathbb{E}[\delta \mid e = 0]$. Now, $e = 0$ if either event a' or b' occur, where a' is $(\delta = \lambda + \epsilon) \cap (\mu > \mu'')$ and b' is $(\delta = \lambda - \epsilon) \cap (\mu > \mu')$.

$$\begin{aligned} \mathbb{E}[\delta \mid e = 0] &= \mathbb{E}[\delta \mid a'] \frac{P(a')}{P(a') + P(b')} + \mathbb{E}[\delta \mid b'] \frac{P(b')}{P(a') + P(b')} = \\ &= (\lambda + \epsilon) \frac{P(\delta = \lambda + \epsilon) \cdot P(\mu > \mu'')}{P(\delta = \lambda + \epsilon) \cdot P(\mu > \mu'') + P(\delta = \lambda - \epsilon) \cdot P(\mu > \mu')} + \\ &= (\lambda - \epsilon) \frac{P(\delta = \lambda - \epsilon) \cdot P(\mu > \mu')}{P(\delta = \lambda + \epsilon) \cdot P(\mu > \mu'') + P(\delta = \lambda - \epsilon) \cdot P(\mu > \mu')} = \\ &= (\lambda + \epsilon) \frac{\frac{1}{2}(1 - \mu'')}{\frac{1}{2}(1 - \mu'') + \frac{1}{2}(1 - \mu')} + (\lambda - \epsilon) \frac{\frac{1}{2}(1 - \mu')}{\frac{1}{2}(1 - \mu'') + \frac{1}{2}(1 - \mu')} = \\ &= \lambda + \epsilon \left(\frac{\mu' - \mu''}{2 - \mu' - \mu''} \right). \end{aligned}$$

Note that the expression multiplying ϵ is positive, as $\mu'' < \mu'$ and $\mu' + \mu'' < 2$ as $\mu', \mu'' \in (0, 1)$ by assumption.

It thus follows that $\mathbb{E}[\delta \mid e = 1] - \mathbb{E}[\delta \mid e = 0] < 0$. As λ drops out, note that the claim holds for any value of λ . ■

Now, we show that our *unconditional* results on belief-updating (i.e., that, *on average*, women's gender-role attitudes move in the anti-female-employment direction after baby

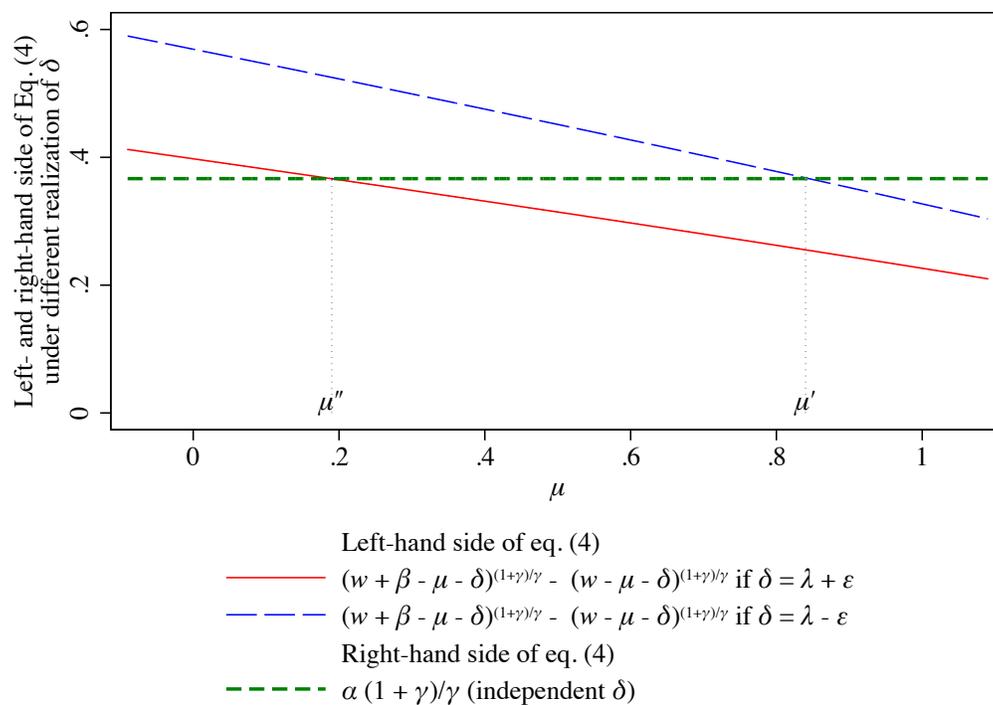
and that they report parenthood being harder than they expected) hold *if and only if* the employment costs of motherhood have increased for the current generation relative to the previous.

Claim 3. $\mathbb{E}[\mu + \delta] < \mathbb{E}[\mu]$ iff $\lambda < 0$.

Proof. As $\delta = \{\lambda + \epsilon, \lambda - \epsilon\}$ with equal probability, the claim follows trivially:

$$\mathbb{E}[\mu + \delta] = \mu + \lambda + \frac{1}{2}\epsilon - \frac{1}{2}\epsilon = \mu + \lambda < \mu \text{ iff } \lambda < 0. \quad \blacksquare$$

Appendix Figure C.1: Graphical derivation of the μ' , μ'' cut-off values described in Lemma 1



Notes: To generate this graph, we use the following parameter values:
 $w = 2, \gamma = 1.2, \lambda = 0, \epsilon = 0.025, \beta = 0.12, \alpha = 0.2$.