The Effect of Maternal Labor Supply on Children: Evidence from unching*

arolina aetano[1], Gregorio aetano[1], Eric Nielsen[2], and Viviane Sanfelice[3]

[1]University of Georgia
[2]Federal Reserve Board
[3]Temple University

March 2023

bstract

We study the effect of maternal labor supply in the first three years of life on early childhood cognitive skills. We pay particular attention to heterogeneous effects by the skill of the mother, by the intensity of her labor supply, and by her pre-birth wages. We correct for selection using a control function approach which uses the fact that many mothers are bunched at zero working hours – skill variation in the children of these bunched mothers is informative about the effect of unobservables on skills. We find that maternal labor supply typically has a significant, negative effect on children's early cognitive skills with more negative effects for higher-skill mothers. By contrast, we do not find significant heterogeneity depending on the pre-birth wage rate of the mother. These findings suggest that there may be more scope to mitigate short-term, unintended consequences of maternal labor supply through policies that promote more flexible work arrangements rather than through policies that increase the financial support to working mothers.

JEL Codes: D13, I21, I2, J01, J22, C24. Keywords: cognitive skills, bunching, maternal labor supply, early childhood, skill development

1 Introduction

This paper estimates the effect of mothers working longer hours during the first three years of a child's life on that child's cognitive skills around age 6. We use data on maternal work histories in the National Longitudinal Surveys of Youth 1979 (NLSY79) linked to childhood skill measures from the Children of the National Longitudinal Surveys (CNLSY), focusing on mothers whose children were born between 1979 and 2008. We aim to understand an important aspect of the trade-off mothers may face when deciding how much to work. On the one hand, maternal labor supply may be detrimental to children's skills because time spent at work is time not spent with children.

We would like to thank Joseph ltonji, ndrew Goodman-Bacon, Joseph Hotz, Josh Kinsler, Luiz Renato Lima, Jonathan Mansfield, Rodrigo Pinto, Christopher Ruhm, David Slichter, Christopher Taber, Hao Teng, and seminar participants at various institutions. The analysis and conclusions set forth here are those of the authors and do not indicate concurrence by other members of the research staff, the Board of Governors, or the Federal Reserve System.

Indeed, there is a wealth of evidence suggesting that an enriching environment with high-quality parent/child interactions in early childhood is important for subsequent skill development (e.g., Todd and Wolpin 2007, Del Boca, Flinn, and Wiswall 2014, Hsin and Felfe 2014, Bono, Francesconi, Kelly, and Sacker 2016). On the other hand, additional work hours will bring in additional income, which may itself have a direct, positive impact on skills (Blau 1999, Milligan and Stabile 2011, Dahl and Lochner 2012, Løken, Mogstad, and Wiswall 2012).

The trade-off between time working and time at home has become increasingly salient as maternal labor supply has increased in recent decades (Eckstein and Lifshitz 2011, Fogli and Veldkamp 2011). Understanding the sign and magnitude of the total effect of maternal labor supply on childhood skill development is critical both for understanding the sources of childhood skill differences and as an input into various policy-relevant analyses. For instance, many public policies – including child allowances and tax credits, subsidized child care, or even the progressivity of the tax code – can alter mothers' labor supply choices.[1] Such policies may have unintended consequences on childhood skill development, with important implications for intergenerational mobility and inequality in general (Blau and Currie 2006, Currie and lmond 2011, Flood, McMurry, Sojourner, and Wiswall 2022).

This paper takes a distinct approach in focusing on heterogeneous effects by the skills of mothers and by the quantity of their labor supply. Both of these dimensions should alter the intensity of the trade-off between maternal labor supply and time at home. More skilled mothers tend to earn higher wages, but their time not working may also be more valuable (in terms of skill production) to their children. It is unclear whether the additional resources (financial or otherwise) earned by skilled working mothers can better offset any detrimental effect of working.[2] Moreover, on the margin, this trade-off is likely to change depending on whether the mother works longer hours (Ettinger, Riley, and Price 2018). These two dimensions may interact: higher-skilled mothers tend to work longer hours (Cortes and Tessada 2011, dda, Dustmann, and Stevens 2017, Chen, Grove, and Hussey 2017), and this may be due in part to a higher return to working longer hours.

Estimating these effects is challenging because maternal labor supply may be correlated with unobservables that are themselves inputs in childhood skill production. Prior research has addressed this endogeneity using standard approaches including family fixed effects and instrumental variables (IVs). We discuss this related work in greater detail in Section 2. Methodologically, we add to the literature by using a novel control function approach that does not require IVs, leveraging instead

[1] The literature on tax credits effects on female labor supply is large; for some references, see work by Eissa and Liebman 1996, verett, Peters, and Waldman 1997, Meyer and Rosenbaum 2001, Grogger 2003, Bosch and Van der Klaauw 2012, Blundell, Costa Dias, Meghir, and Shaw 2016, Bick and Fuchs-Schündeln 2017. For work on child care effects on maternal labor supply at early ages (less than 3 years old) see Baker, Gruber, and Milligan 2008, Goux and Maurin 2010, Givord and Marbot 2015, Carta and Rizzica 2018, Yamaguchi, sai, and Kambayashi 2018, Gathmann and Sass 2018, and ndresen and Havnes 2019.

[2] On the one hand, their spouse or others in her network might be more available to the child (Kalenkoski, Ribar, and Stratton, 2009; Sayer and Gornick, 2012), they may be able to afford higher-quality childcare (Blau and Hagy, 1998; Flood et al., 2022), or they may be better able to substitute market-purchased goods for their own time (nderson and Levine, 1999). On the other hand, the time of a higher-skilled mother may be less substitutable (from the perspective of the child) with any of these options (Ruhm, 2009; Carneiro, Meghir, and Parey, 2013; Polachek, Das, and Thamma- piroam, 2015).

the observation that there is a concentration of mothers who work exactly zero hours (Caetano, Caetano, and Nielsen 2023). This control function approach helps us add value to this literature because it allows us to use the full NLSY79/CNLSY sample – we do not need to restrict our analysis to families with siblings (as in fixed effects models) or to families for whom a particular instrumental variable is available.[3] This enables us to uncover effects broken down along important dimensions of heterogeneity, which enrich our understanding of the relationship between maternal labor supply and child development. It also allows us to recover average treatment effects using an entirely new source of variation relative to prior literature.

Figure 1 illustrates the source of variation in a simple context without controls. The left panel shows hypothetically how the outcome variable (child's skill) varies with the treatment variable (working hours of the mother).[4] The positive slope in this panel combines the treatment effect of maternal hours on the child's skill (what we want to identify) and the endogeneity bias, as confounders are not held constant when the treatment variable varies in the horizontal axis. The left panel also shows a discontinuity in the expected skill at zero: the average skill of children whose mothers work zero hours is sharply lower than the average skill of children whose mothers work just a few hours per year. It is not plausible that this discontinuity represents a treatment effect: working just a few hours over the first three years of the child's life, instead of zero, is unlikely to generate such a sharp increase in the skill of the child at age 6.

Figure 1: Isolating the Effect of Confounders

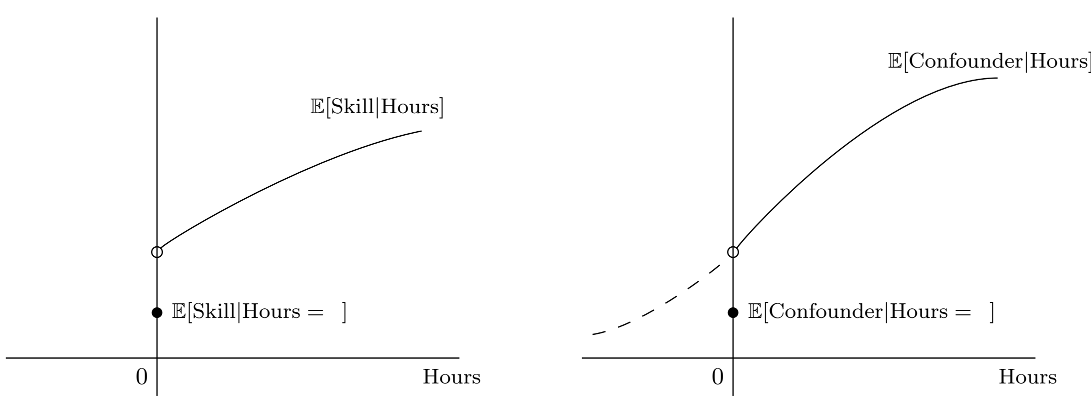

The right panel of Figure 1 offers an explanation for this discontinuity. It illustrates how a hypothetical confounder varies with the treatment variable. For concreteness, we can think of the confounder as positively correlated with the treatment variable, so that mothers with higher values of the confounder tend to work longer hours. This confounder likely reflects a combination of many unobserved factors that correlate to both maternal working hours and the skill of the child, such

[3] nother potential concern with understanding heterogeneous effects using IVs is that the "compliers" of a given IV – who cannot be directly observed – may disproportionately have a certain skill level, making it difficult to compare the estimates across skill levels in a meaningful way. By contrast, we demonstrate empirically that mothers of all skill levels are well-represented in the bunched-at-zero group.

[4] See the left panel of Figure 5 in Section 5 for the empirical version of this figure in our setting.

as the skill of the mother, the quality of her interactions with the child, and the overall quality of the child's environment. The plot shows that the average value of the confounder is comparable for mothers that work a similar number of hours. However, there is a discontinuity at zero hours: the average value of the confounder among the mothers who do not work is very different from the average among the mothers who work marginally positive hours. This happens because the mothers who do not work belong to two different groups: those with confounder values that led them to be exactly indifferent between not working and working, and those with confounder values that led them to be far from indifference, in favor of not working. Because the number of working hours cannot be negative, the first group of mothers choose an "interior solution," while the second group is constrained to choose a "corner solution". s the dashed line in the right plot suggests, the existence of the corner-solution group is what generates the discontinuity.

Summarizing the ideas above, (1) the non-negativity constraint leads confounders to be discontinuous at the bunching point, and (2) this discontinuity in confounders is the sole reason for the discontinuity in the outcome at the bunching point (i.e. the treatment did not cause that discontinuity). The ratio of these two discontinuities (i.e., the discontinuity in the outcome divided by the discontinuity in the confounder) can be interpreted as the effect of confounders at the bunching point. This effect can then be used to correct the endogeneity bias and identify the treatment effect. s this discussion suggests, our approach is somewhat related to "fuzzy" regression discontinuity designs, which also identify a causal effect with the ratio of two discontinuities (e.g. Hahn, Todd, and van der Klaauw 2001). This relationship is informally discussed in ppendix .

Besides bunching, we need two assumptions in order to identify the effect of interest. First, because we do not directly observe the confounder, we need to make an assumption that allows us to obtain the size of the discontinuity in the right panel of Figure 1. Intuitively, this is equivalent to identifying how far from exact indifference between not working and working is the average mother who does not work. Second, because we can only identify the direct effect of confounders on the outcome at zero, we need to be able to extrapolate this effect to mothers who work positive hours. Intuitively, this is equivalent to assuming that the confounder enters linearly in the outcome equation. In practice, we show that both assumptions become much weaker as we add observed controls. Section 6 presents detailed sensitivity analyses of both assumptions and demonstrates that our key empirical conclusions do not depend on them.

We find that maternal labor supply has, on average, negative effects on children's cognitive skills in the short-run. Our estimates imply that an additional 10 hours of maternal labor per week during the first three years of a child's life lowers the child's cognitive skills at age six by about 10% of a standard deviation (s.d.). We also find substantial heterogeneity in these effects by the mother's skill, measured by the rmed Forces Qualifying Test (FQT) score, and some heterogeneity by the total number of hours worked. These heterogeneous results are shown in Figure 2. The hollow circles show the average observed childhood cognitive skills and post-birth maternal work hours for each quartile of the maternal FQT distribution. The lines show the counterfactual skills of children observed at each of these hollow points if their mother worked a different number of hours above

Figure 2: Children's Cognitive Skills – Quartiles of Maternal FQT

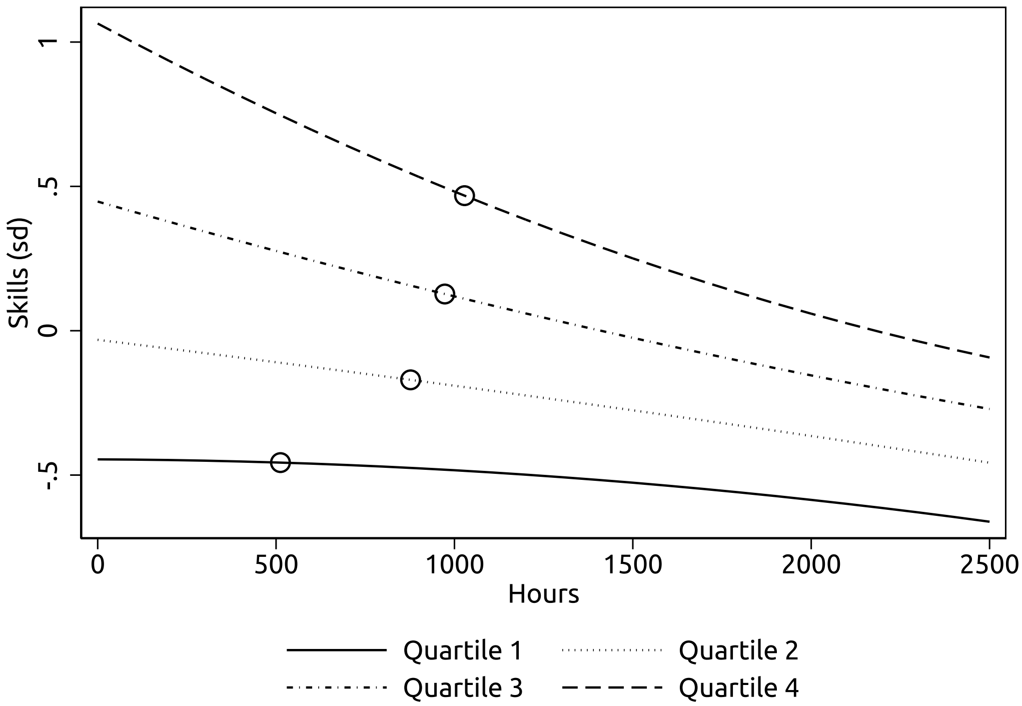

Note: Total effects based on estimates from Table 3. The hollow circles represent the average skills and working hours of all observations in the corresponding quartile of the maternal FQT distribution.

or below the average value in the data. First, we find evidence that the schedule is nonlinear, with the degree of concavity/convexity changing depending on the skill level of the mother. However, although statistically significant, this curvature is economically not very important – the linearity assumption that is typically made in this literature seems to be a good approximation for the range of hours and FQT scores observed in the data. Second, we find substantial heterogeneity depending on the skill of the mother. Labor supplied by higher-skilled mothers tends to have more negative effects, while for lower-skilled mothers the effects are closer to zero.

Of course, maternal labor supply has many other positive effects which may justify the implementation of work-promoting policies. Specifically, the additional income the mother earns may be beneficial to the child in the long run via several channels – better schools and social networks, support for college admissions, reduced levels of stress, etc. Furthermore, the additional income will generally affect all family members, including the mother herself, in various positive ways. There are also important considerations with respect to career timing. lthough a mother may have liked to cut back her hours during her child's early years, she might prefer to remain in the labor force full time, even if doing so is detrimental to the child in the short-run, because of potential negative long-term effects on her career. Better career growth, and the higher resources that come with it, could in turn benefit all family members, including the child, in the long run. Thus, it would be valuable to investigate further whether it is feasible to design and implement work-promoting policies that mitigate the scope for this negative unintended consequence for higher-skilled mothers.

With these considerations in mind, we investigate why working longer hours seems to have particularly detrimental short-run effects on the cognitive skills of the children of higher-skilled mothers. natural candidate explanation is that the last hour worked may be particularly costly for the children of mothers working longer hours, as higher-skill mothers disproportionately work longer hours. However, this potential explanation is ruled out by our finding that the effect is approximately linear, as shown in Figure 2. In fact, to the extent that we find some evidence of nonlinearity, Figure 2 indicates that the last hour might be particularly costly for children of mothers working longer hours only for mothers who have sufficiently low skills.

nother potential explanation is that the additional money earned by these higher-skilled mothers may not be enough to offset their opportunity cost of working, at least in the short-run. It is therefore valuable to consider heterogeneity of the effects by another dimension beyond maternal skills: the mother's wage rate. If the negative effects accrue mostly to the children of high-skilled, *low-wage* mothers, then policies that provide financial support to low-wage mothers could be effective by allowing their families to pay for goods and services to offset their absence. However, if the negative effects accrue also to the families of high-skilled, *high-wage* mothers, suggesting that close substitutes for high-skilled maternal time are unavailable even for this high-earning group, then providing financial support may be ineffective. In this case, it may make more sense to focus on policies aimed at promoting flexible location and work schedules for mothers, thus allowing mothers to maintain (or increase) their labor supply while not reducing the time they spend with their children.

In order to investigate these issues, we study the heterogeneity of the effects of maternal work hours (beyond skill and number of hours) by the hourly pre-birth wage rate of the mother. llowing for this third dimension of heterogeneity leads to more noisy estimates, as expected. Nonetheless, this exercise is still informative. We find some evidence that mothers with higher pre-birth wages may be able to mitigate some of the detrimental effect of their absence, although the degree of mitigation appears to be modest at best. It appears that current institutions and norms do not provide sufficient scope for families with higher-skill working mothers to mitigate any detrimental effects of the mother's work on the child's cognitive skills in the short-run, even if the mother is also a high-earner.

These findings yield a number of important insights about the effects of maternal work during the first years of the children. First, using a new identification strategy leveraging different features of the data, we confirm what has typically been found in the prior literature (see Section 2): maternal labor supply on average has a negative short-run effect on children's cognitive skills. Second, we provide new evidence that this short-run unintended consequence tends to be small for low-skill mothers, even those who work long hours. Third, our analysis gives us a clue about how work-promoting policies could potentially avoid short-run unintended consequences for the children of higher-skill mothers: our results are consistent with the idea that it is safer to focus on policies that encourage greater flexibility in work arrangements for mothers, rather than focusing on policies that provide financial support to working mothers. Policies that increase flexibility in the schedule and location of jobs could allow mothers to spend more time with their children without sacrificing their work hours. dditionally, enhancing spouses' flexibility along the same lines might be complementary, further mitigating any potential cost to children.[5] Of course, it is difficult to draw direct implications for policy without a better understanding of the specific channels through which the causal effects we estimate are operating. More work is thus needed to assess the potential benefits of different policies. It is worth remarking that remote and hybrid work has recently expanded dramatically due to the necessity of social distancing during the COVID-19 pandemic, especially for higher skilled workers (Bartik, Cullen, Glaeser, Luca, and Stanton, 2020; Bick, Blandin, and Mertens, 2020; Dingel and Neiman, 2020). It would be valuable to understand the impact of these changes on the skill of the children, which could complement the findings of our study.

The rest of this paper is organized as follows. Section 2 relates this paper to previous work in the literature, while Section 3 presents our data. Section 4 discusses our empirical approach. Section 5 presents our main empirical findings, while Section 6 provides a detailed sensitivity analysis and various assessments of our key identifying assumptions. Section 7 concludes. The appendix contains additional material that helps provide further context to our study.

2 Related literature

The vast majority of studies on this topic use the NLSY79/CNLSY data and focus on estimating the impact of maternal hours worked during the three first years of a child's life on the child's skills at an early age, as we do. To overcome the endogeneity of maternal labor supply, these studies use either (i) a considerable set of control variables (Desai, Chase-Lansdale, and Michael, 1989; Baydar and Brooks-Gunn, 1991; Vandell and Ramanan, 1992; Parcel and Menaghan, 1994; Hill and O'Neill, 1994; Waldfogel, Han, and Brooks-Gunn, 2002; Baum II, 2003; Ruhm, 2004, 2009), (ii) local labor market conditions as an instrumental variable (Blau, Grossberg, et al., 1992; James-Burdumy, 2005), (iii) family fixed effects (Waldfogel et al., 2002; James-Burdumy, 2005), or (iv) dynamic choice models that simultaneously consider a mother's choice to work and invest in the child's cognitive skill (Bernal, 2008).

The results in these studies vary widely, making it difficult to draw a clear conclusion about the magnitude of the effect of maternal employment. Nonetheless, on balance, this literature finds that maternal labor supply has either a null or a detrimental effect on children's cognitive skills. Our results support these findings. For instance, Ruhm (2004), which adopts a selection-on-observables approach, finds that each additional twenty hours worked per week during the first three years of life is associated with a 0.11 standard deviation decrease on the reading assessment and a 0.08 standard deviation decrease in the mathematics assessment. Similarly, Bernal (2008) finds that working full-

[5] While our results pertain to the effects of maternal labor supply, similar considerations and concerns would in principle apply to the labor supply choices of any parent or caregiver. Our focus on mothers is purely pragmatic – data sources such as the NLSY79/CNLSY do not allow one to connect paternal labor supply to measures of childhood skills.

time and using childcare for one year is associated with a 0.13 standard deviation reduction in test scores. Other papers finding negative effects include Desai et al. (1989), Baydar and Brooks-Gunn (1991), Hill and O'Neill (1994), and Baum II (2003). Using fixed-effects models, James-Burdumy (2005) finds null effects in some cases and negative effects in others. Parcel and Menaghan (1994) similarly find null effects, while Blau et al. (1992) and Waldfogel et al. (2002) find negative effects in the first year of the child's life and offsetting, positive effects subsequently. Finally, Vandell and Ramanan (1992) reports positive effects of early maternal employment on math achievement for children from low-income families, which is consistent with our heterogeneous results.

Our analysis matches the context of this literature: we also focus on the impact of maternal labor supply in the first three years of the child on the child's early outcomes, and we also use the NLSY79/CNLSY data. Because of the similar context, we complement the main findings in this literature in many ways: (a) we confirm the main findings of negative effects with a different approach to control for confounders; (b) we confirm that the linearity assumption made in this literature is a good approximation for the range of hours and skills in the data; (c) we provide new results about heterogeneity by skills; and (d) we investigate the direct vs. income-mediated channel of the effect, providing further context to the findings of this literature while shedding light on the potential impacts of different policies.

We are not the first study to investigate the direct vs. the income-mediated channel of maternal labor supply. Two recent papers investigate such effects, but in contexts different than those in the literature discussed above. gostinelli and Sorrenti (2021) use the NLSY79/CNLSY to estimate time and income effects of maternal labor supply when children are 4-16 years old on the children's contemporaneous outcomes, instrumenting for maternal labor supply with local labor market conditions and for family income with Earned Income Tax Credit (EITC) expansions. They find negative direct effect of maternal hours worked and positive income effect that are not fully offsetting, as we do. Using Norwegian registry data, Nicoletti, Salvanes, and Tominey (2020) estimates the direct and income-mediated effects of maternal labor supply during the first five years of the child on test scores at ages 11 and 15. To handle the endogeneity of maternal work hours and family income, the authors construct instruments for each based on the characteristics of the peers of the parental peers. They find a negative direct effect of maternal labor supply on test scores and a positive income effect that fully offsets the negative direct effect.

Remark 2.1. ***Censoring and Fixed Costs***
This paper is also related to the empirical literature estimating labor supply wage elasticities (see, e.g., Heckman, 1974; Gronau, 1974; Cogan, 1981; Mroz, 1987; Zabel, 1993). In addition to being concerned mainly with the estimation of different quantities, these papers differ from ours in at least two other important dimensions. First, they face a censoring problem: wages are only observed for those who are working. By contrast, we do not face a censoring problem because we observe childhood skills for everyone in our sample, including the children of mothers who work zero hours. Second, these papers often explicitly model both an intensive and an extensive margin of labor supply, and fixed costs are typically a relevant feature of the models they explore. The treatment variable used in

our paper and the whole related literature on child development discussed above is aggregated at the year level or higher. Following Cogan (1980) and Cogan (1981), we argue that, in this instance, a model with only variable costs approximates the distribution of the treatment variable well.[6] *Indeed, Figure 14 in ppendix C shows that we observe no "hole" around zero in the distribution of working hours, which is what would be expected if fixed costs were empirically important. In Section 6.2.2 (Remark 6.2), we consider a sensitivity analysis that allows for fixed costs, and we confirm in an independent way that fixed costs do not appear to be important in our context.*

3 Data

We use data from two linked surveys: the NLSY79, which gives us information about mothers, and the CNLSY, which gives us information about their children. The NLSY79 follows a cohort of young adults aged 14-22 from 1980 through the present, while the CNLSY follows the children born to the women in the NLSY79 sample.[7] Linked together, these surveys provide a unique source of information on children and their parents, including detailed information on maternal labor supply, childhood cognitive development, and household characteristics. Our final sample is a cross-sectional data set of children born from 1979 to 2008 for whom information on cognitive measures, maternal labor supply, and family characteristics are available. Children who were reported not to be living with their mother in the first years of life are dropped from our sample. We also drop observations who report working exactly 40 hours per week for all 52 weeks during each of the three years, as this lack of variation across weeks suggests that these reported hours do not reflect the actual working hours of the mother. However, replicating our analysis using these observations yields nearly the same results in all instances.

Following the economic literature in child development, we measure cognitive skills using the reading recognition and math tests from the Peabody Individual chievement Test (PI T). The reading recognition test is designed to measure reading comprehension based on a child's ability to recognize and pronounce words. The math test assesses attainment in mathematics beginning with early skills, such as recognizing numerals, and progressing to advanced concepts in geometry and trigonometry. The PI T was administered to all children over the age of 5 in each CNLSY wave. Because our focus is on early childhood skill development, we adopt as our outcome a unified score for childhood cognitive skills constructed by applying factor analysis to the age-standardized math and reading PI T scores from the first time each child in the CNLSY is assessed, which happens around age 6.[8] Throughout the analysis, we measure skills in standard-deviation (s.d.) units.

We measure our primary variable of interest, maternal labor supply, using the average number

[6]For instance, Cogan (1980) states, "It may be argued that the major sources of the costs of work are more properly treated as variable costs with respect to annual hours. This is true, especially if most of the variation in annual hours worked was the result of variations in days worked per year."

[7]The NLSY79 interviews are annual from 1979-1994 and biennial thereafter. The CNLSY interviews are biennial starting in 1986.

[8]These age-specific scores are based on a nationally representative sample of children and are normalized to have a mean of 100 and a standard deviation of 15.

of hours worked annually by the mother in the first three years of the child's life. The NLSY79 collects extensive weekly information on employment status and hours worked. This allows us to construct a weekly work history for each mother after giving birth. Some mothers may report that they are working shortly after giving birth when they are actually on paid maternity leave (Baum II, 2003). We can only distinguish these two possibilities – working after birth versus paid maternity leave – in the survey waves from 1988 onward. To avoid losing a large portion of our sample and yet to avoid measurement error due to maternity leave, we begin to measure hours worked in the fourth month following the month of birth.[9] For instance, for a child born in July, we compute hours worked by the mother starting in the first week of November. For this child, maternal labor supply in the first year of life would be computed from the first week of November of the year of birth until the last week of October in the following year. We continue this yearly computation for the next two years in order to measure hours worked by the mother in the second and third year of the child's life. Finally, our treatment variable is computed by taking the average of annual number of hours worked by the mother in these three years.[10]

key explanatory variable in this study is the mother's cognitive skill, which we measure using the rmed Forces Qualifying Test (FQT). The FQT was administered to almost all NLSY79 respondents in the base year of the survey. The FQT is a general measure of achievement in math and reading and is a primary eligibility criterion for service and placement in the United States rmed Forces. Because of its use in U.S. military personnel decisions, the FQT has undergone extensive vetting and has been used in numerous prior economic studies as a proxy for cognitive skill or human capital (Neal and Johnson, 1996; Hirsch and Schumacher, 1998; rcidiacono, Bayer, and Hizmo, 2010).[11]

In addition to maternal FQT, we construct a number of other control variables based on the child, mother, and household characteristics. Unless otherwise specified, control variables such as the mother's education and marital status are computed at the year of birth. We opt for this approach in order to keep our control variables pre-determined.[12]

Table 1 presents summary statistics of our sample. The table first shows the mean and standard deviation of each element used to generate the children cognitive skill measure. These variables are normalized by age and follow a nationally representative sample with a mean of 100 and standard deviation 15. On average, children in our sample score above the national average on the PI T reading recognition, and marginally below the average on math.

[9]The findings in this paper do not change if we start counting hours in the month immediately after the month the child is born.

[10]For some years, the NLSY79 reports weekly employment information over 53 weeks instead of 52 weeks. In order to avoid this type of measurement error, we discard information about hours worked in the 53rd week of a year, if any. In practice, this change turns out to be immaterial for the results.

[11]The FQT is based on a subset of tests from the rmed Services Vocational ptitude Battery (SV B). Throughout, we use the current (post-1989 renormalization) definition of FQT math as the sum of the arithmetic reasoning and mathematics knowledge subscores of the SV B.

[12]For children born after 1994 in odd years, the survey was not conducted in their year of birth. In these cases, we measure control variables in the year before birth, except family size which is measured at the year after birth in order for the child itself to be counted as part of the family.