

Fathers, Time Investment and Human Capital Inequality ^{*}

Jiaqi Li Yaolang Zhong
University of Warwick [†] University of Tokyo [‡]

September 16, 2025

Abstract

This study investigate the impact of fathers' time investment on children's human capital development. Using household data from the PSID time diary and Child Development Supplement, we document significant heterogeneity in paternal time investment. To quantify these effects, we employ a dynamic factor model framework to estimate the production function for children's human capital, focusing on cognitive and health outcomes. To address the endogeneity of parental time investment, we implement a control function approach that leverages local labor demand shocks as instruments. Our counterfactual analysis demonstrates that equalizing paternal time investment across different stages of child development can significantly reduce inequalities in cognitive and health outcomes by adolescence. Specifically, equalizing paternal time during both early and middle childhood leads to a 22 percent reduction in cognitive disparities and a 49 percent reduction in health disparities. These findings underscore the pivotal role of paternal time investment in shaping inequalities in children's human capital development.

Keywords: Intergenerational Mobility, Paternal Time Allocation, Child Development, Inequality, Dynamic Latent Factor Models, Human Capital Production Functions

JEL: D13, J13, J21, J22, J24, I14

^{*}We are grateful to Sonia Bhalotra, Mingli Chen, and Steven Durlauf for the early discussion of the project.

[†]Department of Economics, University of Warwick . Email: j.li.51@warwick.ac.uk

[‡]Department of Economics, University of Tokyo. Email: yaolang.zhong@warwick.ac.uk

1 Introduction

In the economics literature, human capital production functions have become central to understanding how parental inputs shape child outcomes (Cunha and Heckman, 2007, 2008; Cunha et al., 2010; Attanasio, 2015; Attanasio et al., 2020). However, these models often simplify parental inputs, either by focusing solely on monetary investments or assuming homogeneity between parents. As a result, the heterogeneity in paternal time investments is frequently overlooked, despite growing evidence that father-specific inputs play a crucial role in children’s development. On the other hand, intergenerational mobility is another central issue in economics, reflecting the extent to which individuals can improve their economic standing relative to their parents. The persistence of earnings across generations has been linked to various parental characteristics, such as income, education, socioeconomic status, and wealth, as well as the genetic transmission of traits. This persistence is often viewed as evidence of unequal opportunities in society (Chetty et al., 2014; Blanden, 2019). A key determinant of intergenerational mobility is children’s human capital development, which is strongly influenced by parental investments.

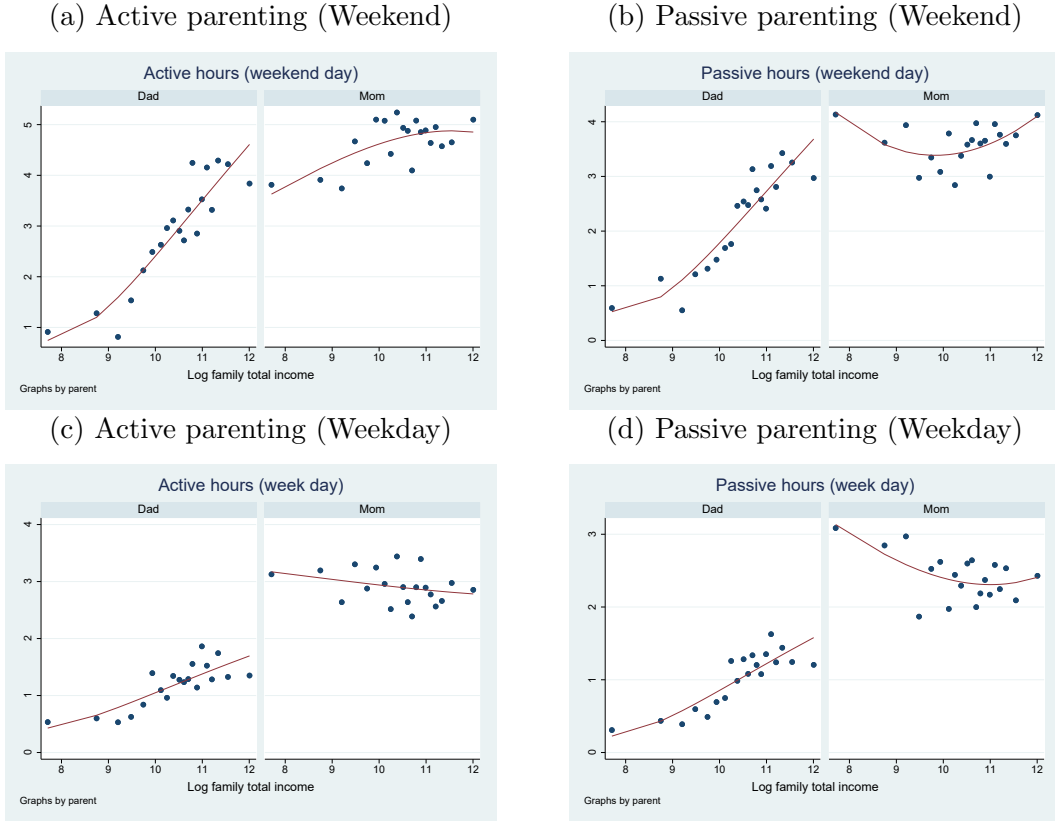
While the economics literature has primarily focused on the relationship between maternal parenting time, female labor supply, and child development, relatively little research has investigated the paternal role in child development. Studies show that mothers often face a trade-off between spending time with their children and participating in the workforce (Baum II, 2003; Ruhm, 2004; Bernal, 2008; Agostinelli and Sorrenti, 2021). While working longer hours can boost household income, it often reduces the time available for child engagement. However, questions about whether this same trade-off applies to fathers, as well as the potential heterogeneity in paternal time investments, remain largely unexplored.

This paper addresses these questions by examining the heterogeneity in parental time inputs, providing new insights into how these investments shape children’s cognitive and health outcomes and, by extension, intergenerational mobility. The remainder of this section first presents stylized facts from the PSID dataset, highlighting distinct patterns in paternal time investment compared to maternal time investment. Next, we briefly introduce the human capital production function framework and discuss the results of our counterfactual exercise in the methodology section, followed by the literature review. Our work emphasizes the crucial role of paternal time investments in explaining disparities in child development, thus filling an important gap in the economics literature.

We first present stylized facts about fathers’ time investment using the 1997 PSID time diary dataset. The fact one is child in household with higher income also receive more time investment from fathers. Such relation is not true for mothers. Figure 1 illustrates the relationship between family income and parental time investment, revealing substantial heterogeneity in the amount of time fathers spend—both actively and passively—with their children on weekdays and weekends, across various income levels. Notably, while maternal

time shows a more nuanced relationship with family income, paternal time investment is significantly and positively correlated with family income. This pattern persists regardless of a child's race, age, or gender, suggesting that paternal behavior may be highly endogenous. For instance, fathers with higher incomes may be more educated and thus more aware of the importance of time investment in their children's development.

Figure 1: Family income versus parenting hours

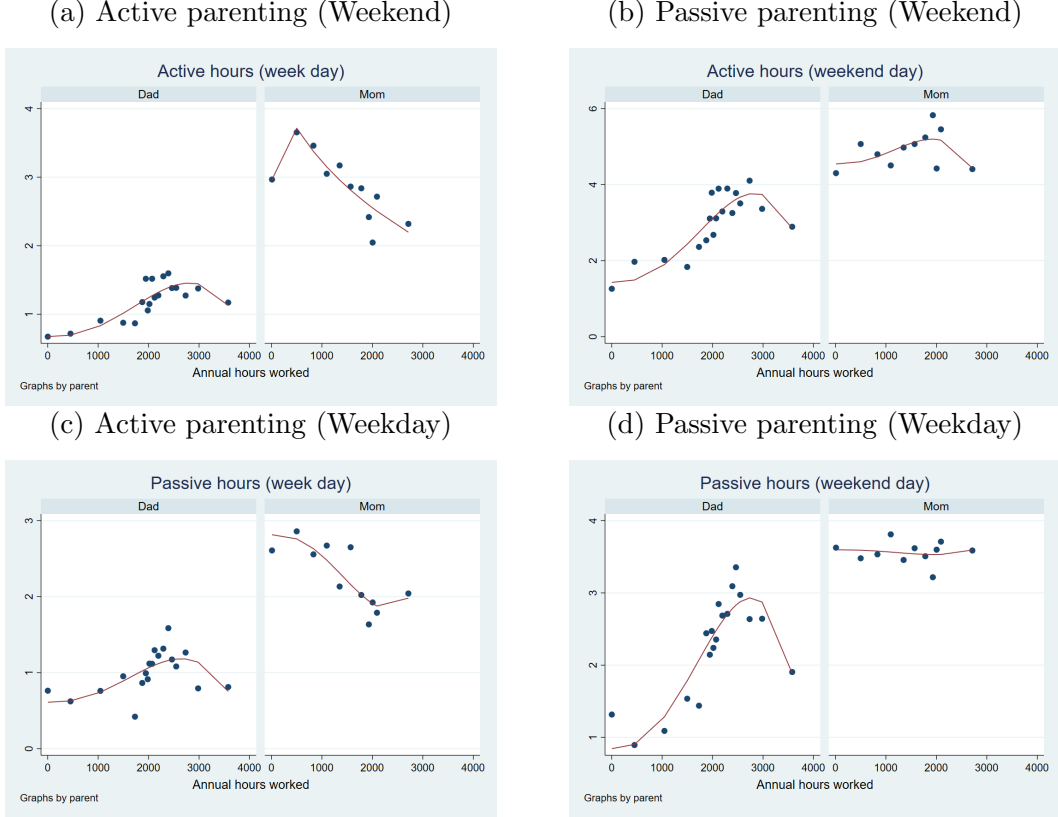


Note: This figure shows the relationship between family income and both active and passive parenting hours on weekdays and weekends.

The second stylized fact is that data reveals a trade off between working hours and parenting hours for mothers, as literature often discussed. However, for father, such trade off is not true. Father who spend more time in labor market also tend to spend more time in time investment with their children. Figure 2 shows the relationship between labor hours and parental time investment—both active and passive—on weekdays and weekends. Maternal time investment exhibits a clear negative correlation with labor hours, consistent with the well-documented trade-off between labor supply and maternal time discussed previously. In contrast, paternal time investment initially has a positive correlation with labor hours, but this relationship becomes negative after reaching a peak. This pattern further supports the idea that paternal time investment decisions are endogenous.

Given these empirical findings, this paper aims to quantify how the heterogeneity in

Figure 2: Labor hours versus parenting hours



Note: This figure shows the relationship between annual labor hours and both active and passive parenting hours on weekdays and weekends.

paternal time investment contributes to inequality in child development. We begin by estimating a dynamic latent factor model of the human capital production function, following a series of earlier works (Cunha and Heckman, 2007, 2008; Cunha et al., 2010; Attanasio, 2015; Attanasio et al., 2020). Specifically, we examine children's human capital across cognitive and health dimensions, where each period's levels depend on prior cognition and health as well as parental time investments, distinguishing maternal from paternal inputs. Our analysis uses the PSID dataset and categorizes children into three age groups: early childhood (ages 0 to 5), middle childhood (ages 6 to 11), and adolescence (ages 12 to 18).

Our estimation strategy consists of three steps. First, we assume a semi-log linear relationship between the observed data and the unobserved latent factors. Under the assumption that both the log of these latent factors and the measurement system follow a mixture Gaussian distribution, we employ the EM algorithm to estimate their joint distribution. In the second step, we use minimum distance estimation to link the measurement system to the latent factors.

Finally, in the third step, we estimate the coefficients of the production function. To address endogeneity in both maternal and paternal time investment decisions—one of the

main challenges in obtaining consistent estimates—we employ a control function approach. This approach uses local gender-specific labor demand shocks as instruments for maternal and paternal time inputs, respectively. Our identification assumption is that, after controlling for family income, these local shocks are uncorrelated with unobserved factors in the child development production function. This strategy follows the methodology of [Attanasio et al. \(2020\)](#), who used local prices for items like food, clothing, notebooks, and medication for worms, and [Attanasio \(2015\)](#), who used local prices of toys and food to account for the endogeneity of parental monetary investment, and [Agostinelli and Sorrenti \(2021\)](#), who used similar strategies for parental time allocation.

Given the estimated coefficients of the production functions, we then perform a counterfactual analysis by equalizing paternal time investment among all children across different developmental stages. The results indicate that applying this treatment during both early and middle childhood would reduce the variation in cognitive outcomes by 22 percent and in health outcomes by 49 percent by adolescence. These significant findings underscore the crucial role that heterogeneity in paternal time investment plays in driving disparities in children’s human capital development.

This paper contributes to the literature in several ways. It adds to the intergenerational mobility literature ([Chetty et al., 2014](#); [Blanden, 2019](#)) by identifying a new mechanism, heterogeneous paternal time input, as a driver of intergenerational persistence. It also contributes to research applying dynamic latent factor models to child human capital development ([Attanasio, 2015](#); [Attanasio et al., 2020](#)), building on the framework of [Cunha and Heckman \(2007, 2008\)](#); [Cunha et al. \(2010\)](#). Our work highlights the critical role of fathers in explaining inequalities in children’s cognitive and health outcomes.

While the economics literature has primarily focused on maternal time investment, research in sociology, developmental psychology, and family studies has long emphasized the importance of paternal involvement in child development. [Lamb \(2004\)](#) shows that father engagement through direct interaction, accessibility, and responsibility affects child outcomes, challenging the traditional maternal focus. [Cabrera et al. \(2007\)](#) finds that paternal involvement in caregiving and play uniquely contributes to children’s development, offering complementary stimulation. [Sarkadi et al. \(2008\)](#) concludes that both the quantity and quality of paternal engagement are essential for better cognitive, emotional, and social outcomes. This paper builds on these findings by incorporating paternal time investment into the dynamic production function of children’s human capital. By examining the heterogeneity of parental time inputs, this study contributes to the literature by recognizing the role of paternal time in shaping intergenerational mobility.

In addition to empirical studies, the econometrics literature has focused on developing methods for identifying and estimating dynamic latent factor models ([Agostinelli and Wiswall, 2016a,b](#); [Freyberger, 2021](#)), while other works have modeled the endogenous parental decision-making process in child development ([Del Boca et al., 2014](#); [Yum, 2023](#); [Agostinelli](#)

and Sorrenti, 2021).

The rest of the paper is structured as follows. Section 2 describes the data used in this study. Section 3 details the dynamic latent factor model that we estimate. Section 4 provides a more detailed discussion of the estimation strategy. Section 5 presents the estimation results and conducts counterfactual analyses. Finally, Section 6 offers concluding remarks.

2 Data, Measurement and Instrumental Variables

In this study, we utilize the PSID Child Development Supplement (CDS), a longitudinal dataset that provides detailed information on child development, health, and parenting within families, drawn from the broader Panel Study of Income Dynamics (PSID). Launched in 1997 and updated in multiple waves (2001, 2007, 2013, and 2019), the CDS has become an invaluable resource for child development research. Its comprehensive and longitudinal design offers deep insights into the factors influencing child development across diverse populations. For example, studies such as Del Boca et al. (2014), Del Boca et al. (2016), Del Boca et al. (2017), Del Boca et al. (2019), and Del Boca et al. (2023) have drawn on this rich dataset to explore various aspects of parental involvement and its effects on child outcomes.

Our dataset comprises two distinct groups of children. The first group includes children who were between the ages of 0 and 17 during the initial waves in 1997, 2001, and 2007. By the 2013 wave, all children in this group had reached adulthood. Consequently, the 2013 and 2019 waves focus on a second group of children who were aged 0 to 17 and lived in PSID families during these later years.

In total, our unbalanced panel consists of 5,463 children. Among these, 2,194 children were measured in both the 1997 and 2001 waves, and 1,086 children were observed in all three waves of 1997, 2001, and 2007. This dataset provides a rich resource for analyzing developmental trajectories and the influences on child outcomes over time across different cohorts.

As shown in Table 1, mothers spend more time on parenting than fathers across all measures—whether active or passive, and during both weekdays and weekends. Family income and childcare expenditure are deflated by the CPI index.

The dataset provides a comprehensive set of attributes covering various aspects of child development, family dynamics, and parental participation. These attributes are organized into several key domains, each offering valuable insights into the factors that influence child outcomes. In what follows, we describe the specific measures used in each domain and introduce the construction of the instrumental variables employed to address the endogeneity of parental time investments.

Table 1: First and second moments of some household statistics and parenting hours

	Mean	SD
Maternal year of schooling	7.73	5.29
Paternal year of schooling	12.94	2.30
Annual total family income (2015 price)	8260.03	7792.47
Annual child care expenditure (2015 price)	163.31	383.45
Girls	0.49	0.50
White	0.50	0.50
Age	6.01	3.66
Mother active hours on a weekend	4.61	3.34
Father active hours on a weekend	2.73	3.19
Mother passive hours on a weekend	3.54	2.88
Father passive hours on a weekend	2.04	2.58
Mother active hours on a weekday	2.90	2.55
Father active hours on a weekday	1.11	1.61
Mother passive hours on a weekday	2.43	2.26
Father passive hours on a weekday	0.95	1.47

Note: Based on PSID-CDS data in 1997. Monetary values are deflated by the CPI index. Parenting is classified as active when a parent participated and passive when a parent participated and passive when the parent was only present in the house. Family total income consists of both father's and mother's labor income, as well as household non-labor income or debt.

2.1 Child Mental and Physical Health

Child mental health is assessed using the Child Depression Inventory in the PSID, which evaluates various dimensions such as appearance, crying, task performance, friendships, irritability, isolation, feelings of love, sadness, self-hate, and optimism. Physical health is measured using indicators such as height, weight status, and Body Mass Index (BMI), adjusted for the child’s age and gender.

2.2 Cognition

Child cognitive abilities are evaluated using the Woodcock–Johnson tests across multiple domains, including applied problems, broad math, broad reading, calculation, letter–word identification, and passage comprehension. In addition, for children aged 7 and older, the Memory for Digit Span assessment—a component of the Wechsler Intelligence Scale for Children-Revised (WISC-R)—is administered to measure short-term memory (?).

2.3 Parenting

Parent–child interactions are measured using indicators of various activities, such as building something together, cleaning the house, discussing books and family matters, engaging in crafts, doing dishes, performing yard work, shopping, playing games or sports, preparing food, using the computer, and washing clothes together.

Parental warmth is assessed through indicators such as discussing TV programs (for children aged 6 and above), expressing appreciation, engaging in favorite activities, attending parenting classes before birth, playing together, saying “I love you”, showing physical affection, and discussing current events, interests, relationships, and daily activities.

2.4 Parental Time Investment

Parental time investment is measured using a 24-hour time diary from the PSID, which records detailed information about a child’s activities on both a randomly selected weekday and a weekend day. The diary captures start and end times, activity types, locations, and social interactions. Parental involvement is further categorized into passive and active participation, with separate records for weekday and weekend activities.

2.5 Family Environment

The family environment is assessed across three main domains. The first domain examines family conflict, including the calmness of family discussions, the frequency of family arguments, physical altercations, and instances of objects being thrown. The second domain evaluates cognitive stimulation at home, as measured by indicators such as the number of books read in the past year, the total number of books in the household, and the presence of cellphones. The third domain captures aspects of the home and neighborhood environment, including noise levels both inside and outside the house, as well as overall cleanliness.

2.6 Local Labor Demand

Following [Agostinelli and Sorrenti \(2021\)](#), to address the potential endogeneity of parental time investment, we propose using local gender-specific labor demand as an instrument for time spent with children. In addition, we adopt the leave-one-out approach proposed by [Borusyak et al. \(2022\)](#) such that when calculating the national industry trend, we leave the own state out.

The instrument is derived from the harmonized microdata provided by the IPUMS CPS, which compiles individual records from the Current Population Survey (CPS) from 1962 to the present. Our measure captures state-level fluctuations in employment opportunities for women and men, by education group, that are plausibly exogenous to family-level determinants of parenting.

The construction of the instrument follows the logic of a shift-share or Bartik-style instrument, as first introduced in [Bartik \(1991\)](#) and developed further in [Goldsmith-Pinkham et al. \(2020\)](#). For each state s , education group e , gender g , and year t , we compute predicted employment growth as

$$\text{LabDemShock}_{s,e,t}^g = \sum_{\text{ind}} (\ln E_{\text{ind},-s,t} - \ln E_{\text{ind},-s,1980}) \cdot \frac{E_{\text{ind},se,1980}^g}{E_{se,1980}^g}, \quad (1)$$

where $E_{\text{ind},-s,t}$ is national employment in industry ind excluding state s , and $\frac{E_{\text{ind},se,1980}^g}{E_{se,1980}^g}$ is the baseline 1980 employment share of group g in industry ind within state s and education group e . This "shift-share" design isolates variation that comes from national industry trends (the shifts) which is leave one out [Borusyak et al. \(2022\)](#), weighted by the pre-determined industry composition of each state and education group (the shares). We construct separate measures for women and men using CPS ASEC data. Education is grouped into four categories: less than high school, high school graduate, some college, and college or more.

The exclusion restriction is that local labor demand shocks affect child outcomes only through their impact on parental employment opportunities and, consequently, parental time

Table 2: First-Stage Regression: Labor Demand Shocks and Parental Time Investment

Dependent Variable:	Univariate		Both Instruments	
	Paternal	Maternal	Paternal	Maternal
Male Shock	-0.39*** (0.15)		-0.36** (0.15)	-0.30 (0.21)
Female Shock		-1.59*** (0.24)	-0.58*** (0.16)	-1.57*** (0.24)
Observations	6,615	6,615	6,615	6,615
F-statistic	7.23	43.55	9.74	22.77

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Regressions include fixed effects. The dependent variables are parental time investment on weekdays. Shocks are standardized to have mean zero and unit variance. Columns (1) and (2) show univariate regressions, while columns (3) and (4) include both instruments to demonstrate distinguishing power.

inputs. While this assumption cannot be tested directly, we present robustness checks that include additional state-level controls and compare female and male shocks as a way to validate the instrument.

To validate the relevance of our instruments, Table 2 presents first-stage regression results showing the relationship between gender-specific labor demand shocks and parental time investment. The results demonstrate that our instruments are statistically significant predictors of parental time allocation. Male shocks negatively predict paternal time investment, while female shocks negatively predict maternal time investment. These first-stage results confirm the relevance of our instrumental variables approach.

The results demonstrate both instrument relevance and distinguishing power between male and female labor demand shocks. In the univariate regressions (columns 1 and 2), male shocks significantly predict paternal time investment, while female shocks strongly predict maternal time investment. When both instruments are included simultaneously (columns 3 and 4), we observe clear gender-specific patterns: male shocks primarily affect paternal time, while female shocks have their strongest effect on maternal time. Notably, cross-effects exist but are weaker—female shocks also significantly affect paternal time, while male shocks have a smaller, non-significant effect on maternal time. This distinguishing power justifies using both instruments in our control function approach, as they provide differential identification for maternal and paternal time investments in the production function.

2.7 Exclusion Restriction and Home Environment Measures

A natural concern is that local labor demand shocks may affect not only parental time inputs, but also other dimensions of the home environment that matter for child development. If labor market fluctuations systematically change household conflict, stimulation, or

material conditions, then our exclusion restriction would be violated.

A second concern is that changes in parental employment opportunities may also affect the time spent with children by relatives or nonrelatives, such as grandparents, other family members, or childcare providers. If labor demand shocks alter these alternative caregiving arrangements, then our instruments would capture broader changes in the child’s social environment beyond direct parental time investment.

To address this, we exploit rich measures of the child’s home environment available in the PSID Child Development Supplement (CDS). These include: (i) parental conflict and warmth scales, and (ii) measures of orderliness such as noise levels and cleanliness. Each of these captures a complementary channel through which labor demand shocks could, in principle, influence child outcomes beyond time investment.

We proceed in two steps. First, we group items into conceptually coherent domains following the HOME inventory and related literature. Second, we apply principal components analysis (PCA) within each domain to extract a standardized index. This reduces dimensionality and ensures that our tests are not sensitive to arbitrary variable choice or multiple-hypothesis testing.

Our exclusion restriction implies that local gender-specific labor demand shocks should not systematically predict these home environment indices, conditional on household income and demographic controls. Empirically, we find that the shocks are strongly predictive of parental time investments, but they have no robust effect on home environment measures. This supports our interpretation of the shocks as exogenous variation in parental time inputs, rather than general changes in household quality.

Table 3 presents formal tests of the exclusion restriction by regressing various omitted variables on both male and female labor demand shocks. The results show that neither instrument significantly predicts the potential omitted variables. For male shocks: non-parental care hours, family conflict, and orderly environment all show non-significant effects. For female shocks: non-parental care hours, family conflict, and orderly environment also show non-significant effects. These null results provide empirical support for the exclusion restriction, as both instruments appear to affect child outcomes primarily through parental time investment rather than through alternative channels.

2.8 Home Environment Measures

To capture dimensions of the child’s home environment, we use information from the PSID Child Development Supplement (CDS) and construct two standardized indices via principal component analysis (PCA). Each measure is designed such that higher values correspond to more adverse or more enriched environments, depending on the construct. The advantage of this approach is that it aggregates multiple related survey items into a pars-

Table 3: Exclusion Restriction Tests: Labor Demand Shocks and Omitted Variables

	Dependent Variable:					
	Non-parental Care Hours		Family Conflict		Orderly Environment	
	(1)	(2)	(3)	(4)	(5)	(6)
Male Shock	0.034 (0.043)		-0.088 (0.051)		0.111 (0.069)	
Female Shock		0.054 (0.049)		-0.004 (0.046)		0.091 (0.054)
Observations	4,059	4,059	882	882	325	325
F-statistic	0.63	1.23	2.98	0.01	2.58	2.82
P-value	0.427	0.267	0.085	0.928	0.109	0.094

Note: Standard errors in parentheses. Regressions include fixed effects. Non-parental care hours measures time spent with children by relatives or non-relatives on weekdays. Family conflict and environment are standardized PCA indices from the PSID Child Development Supplement. All dependent variables are potential omitted variables that could violate the exclusion restriction if systematically affected by labor demand shocks.

monious continuous score while reducing measurement error and collinearity across highly correlated items. We construct two home environment indices—parent–child conflict and household orderliness—by applying principal component analysis (PCA) to the corresponding sets of survey items. In each case, the first principal component explains a sizable share of the variance (between 44–58%) and loads positively on the expected dimensions. We standardize the resulting scores for ease of interpretation. Full PCA loadings, eigenvalues, and variance explained are reported in Appendix Tables [A1–A3](#).

Parent–Child Conflict. We measure the degree of conflict within the family using five items from the CDS wave. Children were asked whether their family: (i) fights a lot, (ii) throws things, (iii) calmly discusses disagreements, (iv) criticizes one another, and (v) hits each other. The item on calm discussion was reverse-coded so that higher values indicate more conflict. We then performed PCA on these five indicators, and retained the first component, which explains 44% of the total variance. This index captures the intensity of conflictual interactions in the household.

Orderly Environment We also measure the degree of household orderliness using three items: (i) whether rooms in the home are clean, (ii) whether the inside of the home is noisy, and (iii) whether the outside environment is noisy. The noise variables were reverse-coded so that higher values indicate a more orderly and less chaotic home. A PCA was then conducted, with the first component retained, explaining 58% of the variance. The resulting index captures the physical and auditory orderliness of the child’s immediate environment.

Summary. Each PCA score was standardized to mean zero and standard deviation one. In robustness checks, we use these measures to ensure that our instrumental variables (local labor demand shocks) predict parental time allocation but not broader aspects of the home environment.

2.9 Falsification Test

A key concern with our identification strategy is that local labor demand shocks might capture persistent state-level trends that are correlated with parental time investment, rather than truly exogenous variation. To address this, we conduct a falsification test using future shocks. Specifically, we re-estimate the first-stage regressions but replace contemporaneous shocks with shocks occurring one, two, and three survey waves ahead. If our instruments capture unanticipated fluctuations in local employment opportunities, then future shocks should have no predictive power for current parental time.

Table 4 reports the results. For paternal time, the coefficient on the contemporaneous male shock is negative and statistically significant (-0.389, $p=0.007$), confirming instrument relevance. However, future male shocks have no predictive power: -0.045 ($p=0.720$) for $t+1$ and 0.037 ($p=0.758$) for $t+2$. Similarly, for maternal time, the contemporaneous female shock is highly significant (-1.587, $p<0.001$), while future shocks are insignificant: -0.119 ($p=0.583$) for $t+1$ and -0.071 ($p=0.696$) for $t+2$. These null results support our interpretation of the instruments as exogenous and unanticipated shocks to local labor demand.

Table 4: Falsification Test: Current vs. Future Labor Demand Shocks

	Paternal Time			Maternal Time		
	Current (1)	$t + 1$ (2)	$t + 2$ (3)	Current (4)	$t + 1$ (5)	$t + 2$ (6)
Male Shock	-0.389*** (0.145)	-0.045 (0.125)	0.037 (0.119)			
Female Shock				-1.587*** (0.240)	-0.119 (0.216)	-0.071 (0.183)
Observations	6,615	6,531	6,070	6,615	6,531	6,070
R-squared	0.001	0.000	0.000	0.007	0.000	0.000
F-statistic	7.23	0.13	0.10	43.55	0.30	0.15
P-value	0.007	0.720	0.758	0.000	0.583	0.696

Note: Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. All regressions include fixed effects. Columns (1) and (4) show contemporaneous effects (first-stage), while columns (2)-(3) and (5)-(6) show future shocks which should have no predictive power if instruments are truly exogenous. The insignificant future coefficients support our identification strategy.

3 Dynamic Latent Factor Model

To understand the process of human development and the role of parental investment, we build upon the seminal works on children’s skill formation by [Cunha and Heckman \(2007\)](#) and [Cunha et al. \(2010\)](#), which employ a dynamic latent factor model. This framework is comprised of three key components: first, the technology of skill formation (or equivalently, the human capital production function); second, the instruments used to address the endogeneity of parental investment decisions, particularly those related to time investment; and third, the measurement system that links observable data to unobservable latent factors. In what follows, we discuss each of these components in detail.

3.1 Human Capital Technology

We begin by specifying that the development of children’s human capital comprises two primary dimensions: cognitive skills (θ_c) and physical health (θ_h), as is commonly addressed in the literature. To account for the different developmental stages of childhood, we categorize children into three age groups: early childhood (ages 0 to 5), middle childhood (ages 6 to 11), and adolescence (ages 12 to 18). From this point forward, we refer to these stages as Stage 1, Stage 2, and Stage 3, respectively. This classification allows us to model the evolution of cognition and health through dynamic production functions that span these stages, resulting in two periods of production ($t = 1, 2$).

In each period, a child’s cognition ($\theta_{c,t+1}$) and health ($\theta_{h,t+1}$) in the subsequent stage depend on several factors: their current cognition ($\theta_{c,t}$) and health ($\theta_{h,t}$), the mother’s time investment ($I_{m,t}$), the father’s time investment ($I_{f,t}$) during that stage, and positive total factor productivity (TFP), denoted by ($A_{c,t}, A_{h,t}$), which captures the efficiency of inputs in producing cognitive skills and health outcomes, respectively.

We assume a Constant Elasticity of Substitution (CES) form for the technology:

$$\begin{aligned}\theta_{c,t+1} &= [\alpha_{c,t}(\theta_{c,t})^{\rho_t} + \alpha_{h,t}(\theta_{h,t})^{\rho_t} + \alpha_{m,t}(I_{m,t})^{\rho_t} + \alpha_{f,t}(I_{f,t})^{\rho_t}]^{\frac{1}{\rho_t}} A_{c,t} \\ \theta_{h,t+1} &= [\beta_{c,t}(\theta_{c,t})^{\zeta_t} + \beta_{h,t}(\theta_{h,t})^{\zeta_t} + \beta_{m,t}(I_{m,t})^{\zeta_t} + \beta_{f,t}(I_{f,t})^{\zeta_t}]^{\frac{1}{\zeta_t}} A_{h,t}\end{aligned}\tag{2}$$

With the input shares α_t and β_t summing to 1 respectively, we are particularly interested in the values of (ρ_t, ζ_t) , which govern the elasticity of substitution among the various inputs in the production process. When these parameters equal 1, the production function described in Equation 2 simplifies to a linear function, implying that all inputs are perfectly substitutable. In contrast, when these parameters are less than 1, the inputs become complementary. In particular, when they equal 0, the function reduces to a Cobb-Douglas form, where the elasticity of substitution is 1.

An alternative functional form to consider is the translog form, as discussed in [Freyberger](#)

(2021), which offers additional flexibility to capture interactions among inputs.

The next equations determine the log of total factor productivity (TFP):

$$\begin{aligned}\log A_{c,t} &= d_{0t} + d'_{Xt}X_t + u_{c,t} \\ \log A_{h,t} &= g_{0t} + g'_{Xt}X_t + u_{h,t}\end{aligned}\tag{3}$$

We assume that log of TFP is a linear function of a set of family characteristics, represented by the covariate X_t . These characteristics include race, the gender of the children, and time-varying family income. Family income, which captures material parental investment, is incorporated as part of TFP rather than as a separate input to human capital. This simplification is motivated by two main reasons: first, our focus in this paper is specifically on the effects of parental time investment, and second, the findings of [Del Boca et al. \(2014\)](#) using the same PSID dataset suggest that monetary inputs are less productive compared to parental time investment.

3.2 Control for the Endogeneity of Parental Time Investments

Next, we introduce additional equations to account for the endogeneity of parental time investments. This step would be unnecessary if parental time investment were exogenous, conditional on the current stage of a child's cognitive and health status as well as household characteristics. However, endogeneity arises because parents may adjust their time investments in response to the evolving human capital of their children. For example, a father might choose to spend more time reading with his child upon recognizing the child's aptitude for learning. In such cases, identification requires instruments that affect parental time investments but are excluded from the child's human capital production function.

As explained in [Section 2](#), we use the residuals from the deviation of gender-, race-, and state-specific linear trends in employment as proxies for local labor demand shocks. Our identification assumption is that these race- and gender-specific variations in local labor demand are uncorrelated with omitted inputs in the child's human capital production function after controlling for family income, which is incorporated into the household characteristics. The empirical specification of parental time investments $(I_{m,t}, I_{f,t})$ is as follows:

$$\begin{aligned}\ln I_{m,t} &= \gamma_0 + \gamma_{c,t} \ln \theta_{c,t} + \gamma_{h,t} \ln \theta_{h,t} + \gamma'_{X,t}X_t + \gamma'_{Z,t} \ln Z_{m,t} + v_{m,t} \\ \ln I_{f,t} &= \eta_0 + \eta_{c,t} \ln \theta_{c,t} + \eta_{h,t} \ln \theta_{h,t} + \eta'_{X,t}X_t + \eta'_{Z,t} \ln Z_{f,t} + v_{f,t}\end{aligned}\tag{4}$$

where $(Z_{m,t}, Z_{f,t})$ are the instruments. Using these instruments, we apply a control function approach ([Gronau, 1974](#); [Heckman, 1979](#); [Heckman and Navarro-Lozano, 2004](#)). Specifically, endogeneity arises because the error terms $(v_{m,t}, v_{f,t})$ in Equation (4) are correlated with the error terms $(u_{c,t}, u_{h,t})$ in the production function (Equation 2). Therefore, we

assume a linear specification conditional on the information set:

$$\begin{aligned} E(u_{c,t} \mid Q_t, Z_t) &= \kappa_{c,m} v_{m,t} + \kappa_{c,f} v_{f,t} \\ E(u_{h,t} \mid Q_t, Z_t) &= \kappa_{h,m} v_{m,t} + \kappa_{h,f} v_{f,t} \end{aligned} \quad (5)$$

where Q_t denotes the full set of variables in the production functions, including parental time investments, and Z_t represents the instruments. The control function approach is implemented by including the regression residuals $(\tilde{v}_{m,t}, \tilde{v}_{f,t})$ from Equation (4) as additional regressors in Equation (2).

3.3 The Measurement System

Since many of the variables in the production function are measured in multiple ways within the dataset, a simple selection of one proxy may introduce bias of unknown direction due to the nonlinearity of the production function (Griliches and Ringstad, 1970). To address this challenge, we apply the latent factor framework proposed by Cunha et al. (2010). Rather than using the flexible nonparametric identification procedure outlined in that work, we assume a specific functional form, as in Attanasio et al. (2020), for simplicity.

We begin by assuming a semi-log mapping between the observed measures and the latent variables:

$$m_{j,k,t} = a_{j,k,t} + \lambda_{j,k,t} \ln(\theta_{k,t}) + \epsilon_{j,k,t} \quad (6)$$

where $m_{j,k,t}$ denotes the j th measure associated with the k th latent variable at time t , and $\lambda_{j,k,t}$ is the corresponding factor loading. Representing the loadings in log form is convenient since the production function incorporates the logarithm of the latent factors. We assume that the measurement error $\epsilon_{j,k,t}$ is separable, normally distributed, and independent of both the latent factor $\theta_{k,t}$ and the other measurement errors.

For compactness, we can express the above equation in matrix form as

$$M = \mathbf{A} + \Lambda \ln \theta + \Sigma \varepsilon \quad (7)$$

where we define I as the number of measures and J as the number of latent factors. In this formulation, M is a vector of length I representing all the measures from the dataset, \mathbf{A} is a vector of length I representing the constant terms, and ε is a vector of measurement errors of length I . Moreover, Σ is the $I \times I$ diagonal matrix, and Λ is the factor loading matrix of dimensions $I \times J$. Finally, $\ln \theta$ is a vector of length J containing the log-transformed latent factors.

We now consider the joint distribution of the log-transformed latent factors. Given

the semi-log mapping in Equation (6), a natural functional form is joint normality, due to its closure under linear combinations. However, assuming joint normality would impose additional restrictions on Equation (2), effectively reducing the CES production function to a Cobb-Douglas form (which is linear in the log form). To avoid this, we instead assume that the joint distribution is a mixture of Gaussians. This approach preserves the desirable property of closure under linear combinations while allowing for greater flexibility.

The next question is determining the appropriate number of clusters for the Gaussian mixture. In our implementation, we experimented with up to 5 clusters and observed that once the number of clusters reached 3 or more, the distance between clusters became insignificant. This indicates that increasing the number of clusters did not substantially improve the approximation of the true underlying distribution. Therefore, we adopt a parsimonious approach and use only 2 clusters in our model:

$$F_\theta = \tau \Phi(\mu_A, \Omega_A) + (1 - \tau) \Phi(\mu_B, \Omega_B) \quad (8)$$

Here, τ represents the mixture weight, i.e., the probability that an observation is drawn from the first Gaussian component. The parameters (μ_A, Ω_A) and (μ_B, Ω_B) denote the mean and variance of the first and second components, respectively. The semi-log mapping between the observed measures and the latent factors then yields a mixture Gaussian distribution for the measurement system:

$$F_M = \tau \Phi(\Pi_A, \Psi_A) + (1 - \tau) \Phi(\Pi_B, \Psi_B) \quad (9)$$

Here, (Π_A, Ψ_A) and (Π_B, Ψ_B) denote the mean and variance of the measures in the first and second clusters, respectively. This formulation enables us to establish the mapping between the parameters characterizing the mixture distribution of the latent factors and those of the measures. The mapping is defined as follows:

$$\begin{aligned} \Psi_A &= \Lambda^T \Omega_A \Lambda + \Sigma, & \Pi_A &= \mathbf{A} + \Lambda \mu_A \\ \Psi_B &= \Lambda^T \Omega_B \Lambda + \Sigma, & \Pi_B &= \mathbf{A} + \Lambda \mu_B \end{aligned} \quad (10)$$

Finally, in addition to the latent variables in our model, there are observable variables used in the production functions and as instruments in the time investment equations. These variables are considered error-free and are included in both the measurement system and the latent variable system. This inclusion results in an augmented system of mixture Gaussian distributions for the latent factors:

$$F_{\theta,X} = \tau \Phi(\mu_A^{\theta,X}, \Omega_A^{\theta,X}) + (1 - \tau) \Phi(\mu_B^{\theta,X}, \Omega_B^{\theta,X}) \quad (11)$$

This also results in an augmented system of mixture Gaussian distributions for the mea-

surement system, along with an expanded mapping between the parameters of the latent variable distributions and those of the measurement system. To avoid unnecessary notational complexity, we omit the detailed equations here.

4 Estimation Strategy

In the following, we describe the estimation strategy we adopt following [Attanasio, Meghir and Nix \(2020\)](#), which consists of three steps:

Step 1: We use the expectation-maximization (EM) algorithm ([Dempster, Laird and Rubin, 1977](#); [McLachlan and Krishnan, 2007](#)) to estimate all parameters characterizing the mixture Gaussian measurement system. These parameters include the weight τ , the mean vectors (Π_A, Π_B), and the covariance matrices (Ψ_A, Ψ_B).

In the context of the Gaussian Mixture Model (GMM), the Expectation-Maximization (EM) algorithm is an iterative procedure that alternates between an Expectation (E) step and a Maximization (M) step, each providing closed-form updates for the parameters.

Based on the notation for the mixture distribution of the measurement system from Equation (9), we define the following notation for our EM algorithm in the context of the Gaussian Mixture Model (GMM):

- N : the number of observations.
- I : the number of dimensions (i.e., measures) in each observation.
- K : the number of clusters; in our context, $K = 2$.
- Y : the observed dataset with shape (N, I) .
- W : the unobserved latent variable representing the true cluster memberships, with each element binary and shape (N, K) ; each cluster is denoted by W_k .
- \hat{W} : the estimated probability matrix of cluster memberships, with shape (N, K) , where each element represents the current estimated probability that an observation belongs to a specific cluster.
- $\hat{\Theta} = \{\hat{\tau}_k, \hat{\Pi}_k, \hat{\Psi}_k\}_{k=1}^K$: the set of current parameter estimates of the GMM, where for each cluster k :
 - $\hat{\tau}_k$ is the estimated mixing coefficient for cluster k , representing the probability that a randomly chosen observation belongs to that cluster. The collection $\{\hat{\tau}_k\}_{k=1}^K$ forms a $K \times 1$ vector that sums to 1.
 - $\hat{\Pi}_k$ is the estimated mean vector for cluster k , with dimensions $I \times 1$.

– $\hat{\Psi}_k$ is the estimated covariance matrix for cluster k , with dimensions $I \times I$.

With this notation, the EM algorithm proceeds as follows:

E-step: For each cluster k , we compute \hat{W} by

$$\hat{W}_k = \mathbb{E}[W_k | Y; \hat{\Theta}] = \frac{\hat{\tau}_k \exp\left(-\frac{1}{2} \text{diag}\left((Y - \hat{\Pi}_k)^T \hat{\Psi}_k^{-1} (Y - \hat{\Pi}_k)\right)\right)}{\sqrt{(2\pi)^I \det(\hat{\Psi}_k)}} \quad (12)$$

M-step: Using the estimated responsibilities \hat{W}_k from the E-step, we update $\hat{\Theta} = \{\hat{\tau}_k, \hat{\Pi}_k, \hat{\Psi}_k\}_{k=1}^K$ by:

$$\begin{aligned} \hat{\Pi}_k &= \frac{Y^T \hat{W}_k}{1_N^T \hat{W}_k}, \\ \hat{\Psi}_k &= \frac{(Y - \hat{\Pi}_k)^T \text{diag}(\hat{W}_k) (Y - \hat{\Pi}_k)}{1_N^T \hat{W}_k}, \\ \hat{\tau}_k &= \frac{1_N^T \hat{W}_k}{N}, \end{aligned} \quad (13)$$

where 1_N is a column vector of ones of length N .

However, while the Gaussian assumption yields closed-form solutions for both the E-step and M-step, our application faces challenges due to a significant amount of missing data. This issue arises from the structure of the PSID dataset, where most children are observed in only one or two waves of data collection. As a result, the updates in the EM algorithm must be conditioned on the observed data for each observation, forcing us to loop over observations rather than employing the more efficient vectorized operations described in Equations 12 and 13. This reliance on loops considerably slows computational time. Additionally, the temporary parameters obtained in each iteration must be managed carefully to prevent the formation of ill-conditioned matrices in subsequent iterations.

Step 2: We then apply minimum distance estimation to recover the parameters that characterize the joint distribution of the latent variables θ . Specifically, we estimate the shift vector \mathbf{A} , the factor loading matrix Λ , the measurement error matrix Σ , as well as the mean vectors μ_A and μ_B and the covariance matrices Ω_A and Ω_B based on the parameters obtained in Step 1.

Before implementing the minimum distance estimation, which is a nonlinear optimization procedure, it is essential to impose a set of restrictions to reduce the degrees of freedom in the parameter space. We now state these restrictions as follows:

1. **Normalization on the latent factor:** For the augmented latent factor (θ, X) , we normalize the time-invariant variables so that the mean of their logarithms is zero. For

the time-varying variables, we normalize only the first period to have a mean log of zero, allowing the log values in subsequent periods to be interpreted as growth relative to the first period.

2. **Restrictions on the factor loadings:** We assume that each measure is linked to only one underlying factor, meaning that only the corresponding factor loading is non-zero. In addition, we set the scale of the latent factor by normalizing one of the measurements, referred to as the primary measurement, with its associated factor loading fixed at one. For example, for child cognition, we normalize the loading on the Woodcock-Johnson letter-word test to one; for child health, we normalize on the z-scores of height; and for parental time investment, we normalize on active hours spent with the child on a weekday. This assignment of primary measurements is maintained consistently across periods.
3. **Restriction on the constant term:** We assume that the growth of the measurements is solely due to the growth of the associated latent variables. Together with the normalization of the loading to the first measurement for each latent variable (as described in point 2), this imposes a restriction on the constant vector \mathbf{A} such that the elements corresponding to the first measurement remain unchanged over time. Consequently, we only need to determine their values in the first period. Point 1 implies that we can set the constant term in \mathbf{A} for the first period equal to the mean of the measurements from the data.

Step 3: We generate a synthetic dataset from the parametric joint distribution of the latent factors estimated in Steps 1 and 2. Next, we estimate the parameters of interest in the production function system (Equation 2) using a control function approach, implemented as a two-stage estimation procedure. In the first stage, we perform OLS regression on Equation 4 to obtain the residuals. These residuals are then included as additional regressors in the linear part of the production function. In the second stage, we estimate the log form of Equation 2 using nonlinear least squares.

5 Results and Counterfactual Analysis

We follow the estimation procedure discussed in Section 4, employing a bootstrap approach with 100 iterations to construct the 90% empirical confidence intervals.

5.1 Estimates

We begin by presenting the estimated joint distribution of the latent factors obtained from Steps 1 and 2 of our estimation procedure. Table 5 reports the results, showing the

mixture weights and the mean vectors for each cluster in the Gaussian mixture model that characterizes the joint distribution of these latent factors. For brevity, the covariance matrices are omitted.

The estimated mixture weights indicate that approximately one-third of the observations belong to the first cluster (Mixture A), while two-thirds belong to the second cluster (Mixture B). This suggests a meaningful division within the population, with Mixture B representing the majority.

The table reflects differences in parental investment between the two mixtures. For example, both maternal and paternal investments in children’s development are consistently higher in Mixture A across all age groups. This suggests that families in Mixture A may prioritize time spent with their children more than those in Mixture B, which could contribute to the observed differences in cognitive and health outcomes.

These results collectively underscore the necessity and effectiveness of employing a Gaussian mixture model to capture heterogeneity in the population. The significant differences in the means of most latent variables between the two clusters indicate a substantial departure from a simple multivariate normal distribution. This finding supports our decision to implement a mixture model, which provides a more nuanced and accurate representation of the underlying data structure.

Table 5: Weights and means of the Gaussian mixture model for latent factors

	Mixture.A	Mixture.B
Weights	0.382 [0.36,0.39]	0.618 [0.61,0.64]
Mean Cognition Stage 3	2.235 [2.17,2.266]	2.254 [2.236,2.32]
Mean Cognition Stage 2	1.846 [1.816,1.871]	1.635 [1.617,1.652]
Mean Cognition Stage 1	0.408 [0.352,0.436]	-0.251 [-0.269,-0.206]
Mean Health Stage 3	2.153 [1.33,2.219]	2.265 [2.21,2.281]
Mean Health Stage 2	1.697 [1.659,1.718]	1.377 [1.361,1.397]
Mean Health Stage 1	-0.067 [-0.106,-0.042]	0.042 [0.025,0.065]
Mean Mom’s Investment Stage 3	1.603 [1.568,1.655]	0.121 [0.073,0.185]
Mean Mom’s Investment Stage 2	1.334 [1.306,1.358]	-0.836 [-0.857,-0.819]
Mean Mom’s Investment Stage 1	0.544 [0.511,0.599]	-0.336 [-0.361,-0.305]
Mean Dad’s Investment Stage 3	2.109 [2.071,2.156]	0.519 [0.471,0.603]
Mean Dad’s Investment Stage 2	1.932 [1.908,1.962]	-0.847 [-0.866,-0.829]
Mean Dad’s Investment Stage 1	0.928 [0.883,1.01]	-0.573 [-0.613,-0.517]

Notes: 90% confidence intervals based on 100 bootstrap replications in square brackets.

We next present the results from Step 3, which show the estimated coefficients of the production function in Table 6. A crucial finding is that the estimated elasticity parameters (ρ_t, ζ_t) all lie between 0 and 1, indicating that the elasticity of the production function is less than 1. This suggests that the inputs in the production of a child’s cognition and health are complementary, consistent with the findings reported in Attanasio et al. (2020). Another important observation is that a child’s initial health status has a dominant impact not only on its subsequent health—indicating a strong self-productivity effect—but also on cognition during childhood (ages 6–12). This influence remains significant as the child transitions from childhood to adolescence. Conversely, the self-productivity effect of cognition is less pronounced during childhood but becomes more significant during adolescence.

Regarding parental time investments, our estimation results indicate that both maternal and paternal time investments have an insignificant direct effect on child health. This finding is plausible, as a child’s physical health is likely more directly influenced by material investments. However, parental time investments are crucial for cognitive development. Specifically, paternal time investment is significant during early childhood, while maternal

time investment becomes significant during adolescence.

Furthermore, we find that the coefficients on the paternal time investment residuals are significantly different from zero for child health during childhood and for cognition during adolescence. This suggests that paternal time investment is endogenous and correlated with unobservable factors in the production function at different stages of child development.

Table 6: Estimated values and confidence interval of the coefficients in the production functions

	Cog, Stage 2	Health, Stage 2	Cog, Stage 3	Health, Stage 3
Cognition	−0.019 [−0.082,0.038]	0.013 [−0.077,0.026]	0.399 [0.288,0.435]	0.198 [0.154,0.315]
Health	0.946 [0.874,0.993]	0.961 [0.9,1.019]	0.406 [0.352,0.595]	0.685 [0.551,0.909]
Mom’s time	−0.113 [−0.211,0.046]	−0.043 [−0.119,0.168]	0.165 [0,0.264]	0.053 [−0.193,0.159]
Dad’s time	0.186 [0.036,0.273]	0.069 [−0.082,0.151]	0.03 [−0.121,0.173]	0.064 [−0.172,0.15]
Mom’s time ctrl func	−0.02 [−0.069,0.094]	−0.037 [−0.083,0.092]	0.02 [−0.034,0.066]	−0.023 [−0.186,0.017]
Dad’s time ctrl func	0 [−0.044,0.088]	−0.159 [−0.22,−0.031]	0.134 [0.006,0.211]	0.107 [−0.104,0.172]
elasticity parameter	0.221 [−0.118,0.868]	0.416 [−0.438,0.736]	0.336 [0,0.635]	0.711 [−0.098,0.884]
Family income	0.128 [0.098,0.138]	−0.003 [−0.027,0.015]	0.082 [0.064,0.1]	0.024 [−0.007,0.035]
White	0.08 [0.036,0.125]	−0.036 [−0.076,−0.004]	0.079 [0.061,0.129]	−0.004 [−0.045,0.042]
Female	0.084 [0.047,0.114]	0.018 [−0.012,0.051]	0.001 [−0.039,0.023]	−0.228 [−0.259,−0.189]
constant	1.568 [1.501,1.634]	1.465 [1.386,1.557]	0.707 [0.612,0.755]	0.785 [0.213,0.866]
Residual standard error	0.451 [0.426,0.479]	0.248 [0.202,0.294]	0.35 [0.32,0.386]	0.34 [0.304,0.599]

Notes: 90% confidence intervals based on 100 bootstrap replications in square brackets.

5.2 Counterfactual Analysis

Using the estimated coefficients from the production function shown in Table 6, we conducted a counterfactual exercise to assess the impact of equalizing paternal time investments on child development inequality. Specifically, we evaluate how much the inequality in child development could be reduced if all children received the same amount of paternal time investment.

We perform the following steps in our counterfactual analysis. First, we simulate a synthetic population based on the estimated joint distribution of parameters. We then focus

on children in Stage 1 and consider three intervention scenarios: (1) equalizing paternal time investment only in Stage 1, (2) equalizing paternal time investment only in Stage 2, and (3) equalizing paternal time investment in both stages. In each scenario, all other inputs and household characteristics in the production function of cognition and health are held constant. This setup creates an artificial experiment in which paternal time investment is treated as the “treatment”. Finally, we compute the standard deviation of cognition and health scores for the population in Stage 3 under the different intervention scenarios.

Table 7: Counterfactual Analysis: Impact of Paternal Time Investment Interventions

Intervention Scenario	Standard Deviation		Relative to Baseline (%)	
	Cognition	Health	Cognition	Health
No Intervention	0.063	0.085	100.0	100.0
Intervention in Stage 1	0.059	0.078	92.7	92.4
Intervention in Stage 2	0.054	0.047	85.0	55.7
Intervention in Both Stages	0.049	0.043	78.0	51.0

Notes: This table presents the results of counterfactual exercises where all children receive equal paternal time investment. Standard deviations measure inequality in child development outcomes. Percentages show the reduction in inequality relative to the no-intervention baseline.

Table 7 compares the standard deviations of cognition and health scores across different intervention experiences. The first column indicates whether an individual experienced an intervention and at which stage. The second and third columns report the standard deviations of cognition and health scores, respectively. The fourth and fifth columns display these standard deviations as percentages relative to the group without any intervention (set at 100%).

From the table, it is evident that children who experienced an intervention at either Stage 1 or Stage 2 exhibit lower standard deviations in both cognition and health scores compared to those with no intervention experience. Specifically, children with intervention experience at Stage 1 show a 7.3% lower standard deviation in cognition and a 7.6% lower standard deviation in health scores relative to the no-intervention group. Similarly, children with intervention experience at Stage 2 exhibit a 15% lower standard deviation in cognition and a 44.3% lower standard deviation in health scores. These findings suggest that equalizing paternal time investment, particularly during Stage 2, could significantly reduce inequality in child development outcomes.

This artificial experiment suggests that, to reduce inequality in children’s development, in both cognition and health, policymakers could consider introducing interventions that influence fathers’ decision-making, thereby balancing the endogenous disparities in paternal time investment across different groups. Further investigation into parental decision-making and the design of such policies may be warranted, although these aspects fall outside the scope of this paper.

6 Conclusion

This paper contributes to the literature on child development and intergenerational mobility by investigating the impact of parental time investment on the development of children’s human capital, with a particular focus on distinguishing between maternal and paternal contributions.

Specifically, we begin by presenting our empirical findings, which reveal significant heterogeneity and endogeneity in paternal time investment. To quantify these impacts, we apply a dynamic factor model framework to estimate the production function of children’s human capital, measured through cognitive and health outcomes, using both maternal and paternal time investments as inputs. In particular, we assume a semi-log linear relationship between the observed data and the unobserved latent factors. Moreover, assuming that the joint distribution of the logarithms of the latent factors, and consequently that of the measurement system, follows a mixture Gaussian distribution, we employ the EM algorithm to estimate this joint distribution. We then use minimum distance estimation to map the measurement system onto the latent factors. Finally, for the estimation of the production function, we utilize a control function approach with local labor demand shocks as instruments to address the endogeneity of parental time investment.

Through a counterfactual analysis, we demonstrate that equalizing paternal time investment across different periods of child development significantly reduces inequality in cognitive and health outcomes by adolescence. Specifically, our results indicate a 22 percent reduction in the variation of cognitive outcomes and a 49 percent reduction in health disparities when paternal time investment is equalized during both early and middle childhood. These findings underscore the critical role of heterogeneity in paternal time investment in shaping the inequality of children’s human capital development.

References

- Agostinelli, Francesco and Giuseppe Sorrenti**, “Money vs. Time: Family Income, Maternal Labor Supply, and Child Development,” Working Paper 273, University of Zurich, Department of Economics February 2021. Revised version.
- **and Matthew Wiswall**, “Estimating the technology of children’s skill formation,” Technical Report, National Bureau of Economic Research 2016.
- **and** —, “Identification of dynamic latent factor models: The implications of re-normalization in a model of child development,” Technical Report, National Bureau of Economic Research 2016.
- Attanasio, Orazio, Costas Meghir, and Emily Nix**, “Human capital development and parental investment in India,” *The Review of Economic Studies*, 2020, 87 (6), 2511–2541.
- Attanasio, Orazio P**, “The determinants of human capital formation during the early years of life: Theory, measurement, and policies,” *Journal of the European Economic Association*, 2015, 13 (6), 949–997.
- Bartik, Timothy J**, “Who benefits from state and local economic development policies?,” 1991.
- Bernal, Raquel**, “The effect of maternal employment and child care on children’s cognitive development,” *International Economic Review*, 2008, 49 (4), 1173–1209.
- Blanden, Jo**, “Intergenerational income persistence,” *IZA World of Labor*, 2019, 129 (4), 176–176.
- Boca, Daniela Del, Christopher Flinn, and Matthew Wiswall**, “Household choices and child development,” *Review of Economic Studies*, 2014, 81 (1), 137–185.
- , —, **and** —, “Transfers to households with children and child development,” *The Economic Journal*, 2016, 126 (596), F136–F183.
- , —, **and** —, “Parental time investments and child development,” *Journal of Human Resources*, 2017, 52 (4), 1063–1096.
- , —, **and** —, “Actors in child development: The role of parents and government,” *European Economic Review*, 2019, 118, 516–549.
- , —, **and** —, “Parenting with style: Altruism and paternalism in intergenerational preference transmission,” *The Review of Economic Studies*, 2023, 90 (3), 1272–1303.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-experimental shift-share research designs,” *The Review of economic studies*, 2022, 89 (1), 181–213.
- Cabrera, Natasha J, Jacqueline D Shannon, and Catherine Tamis-LeMonda**, “Fathers’ influence on their children’s cognitive and emotional development: From toddlers to pre-K,” *Applied Development Science*, 2007, 11 (4), 208–213.

- Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez**, “Where is the land of opportunity? The geography of intergenerational mobility in the United States,” *The Quarterly journal of economics*, 2014, *129* (4), 1553–1623.
- Cunha, Flavio and James Heckman**, “The technology of skill formation,” *American economic review*, 2007, *97* (2), 31–47.
- **and James J Heckman**, “Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation,” *Journal of human resources*, 2008, *43* (4), 738–782.
- , – , **and Susanne M Schennach**, “Estimating the technology of cognitive and noncognitive skill formation,” *Econometrica*, 2010, *78* (3), 883–931.
- Dempster, Arthur P, Nan M Laird, and Donald B Rubin**, “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the royal statistical society: series B (methodological)*, 1977, *39* (1), 1–22.
- Freyberger, Joachim**, “Normalizations and misspecification in skill formation models,” *arXiv preprint arXiv:2104.00473*, 2021.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik instruments: What, when, why, and how,” *American Economic Review*, 2020, *110* (8), 2586–2624.
- Griliches, Zvi and Vidar Ringstad**, “Error-in-the-variables bias in nonlinear contexts,” *Econometrica: Journal of the Econometric Society*, 1970, pp. 368–370.
- Gronau, Reuben**, “Wage comparisons—A selectivity bias,” *Journal of political Economy*, 1974, *82* (6), 1119–1143.
- Heckman, James and Salvador Navarro-Lozano**, “Using matching, instrumental variables, and control functions to estimate economic choice models,” *Review of Economics and statistics*, 2004, *86* (1), 30–57.
- Heckman, James J**, “Sample selection bias as a specification error,” *Econometrica: Journal of the econometric society*, 1979, pp. 153–161.
- II, Charles L Baum**, “Does early maternal employment harm child development? An analysis of the potential benefits of leave taking,” *Journal of labor Economics*, 2003, *21* (2), 409–448.
- Lamb, Michael E**, “The role of the father in child development,” 2004.
- McLachlan, Geoffrey J and Thriyambakam Krishnan**, *The EM algorithm and extensions*, John Wiley & Sons, 2007.
- Ruhm, Christopher**, “Parental employment and child cognitive development,” *The Journal of Human Resources*, 2004, *39* (1), 155–192.

Sarkadi, Anna, Robert Kristiansson, Frank Oberklaid, and Sven Bremberg, “Fathers’ involvement and children’s developmental outcomes: a systematic review of longitudinal studies,” *Acta paediatrica*, 2008, *97* (2), 153–158.

Yum, Minchul, “Parental time investment and intergenerational mobility,” *International Economic Review*, 2023, *64* (1), 187–223.

A Appendix Data

The Child Development Supplement (CDS) to the PSID collected comprehensive data in the following areas: (i) reliable, age-graded assessments of the cognitive, behavioral, and health status of 3,600 children (including approximately 250 immigrant children), obtained from the mother, a second parent or parent figure, the teacher or childcare provider, and the child; (ii) a detailed accounting of parental and caregiver time inputs to children, as well as other aspects of how children and adolescents spend their time; (iii) teacher-reported time use in elementary and preschool programs; and (iv) non-time use measures of other resources—such as the learning environment at home, teacher and administrator reports of school resources, and parent-reported measures of neighborhood resources.

B Construction of Home Environment Indices: PCA Results

Construction of Home Environment Indices (CDS/PSID)

We use data from the Child Development Supplement (CDS) of the PSID to construct three standardized indices of the home environment: (i) parent–child conflict, (ii) cognitive stimulation/resources, and (iii) orderliness. Each index is derived by applying principal component analysis (PCA) to survey items consistently available across CDS waves (1997, 2002, 2007, 2014, 2019). All items are recoded so that higher values represent a worse environment (for conflict) or a more stimulating/organized environment (for resources and orderliness). “Don’t know” and “Refused” responses are coded as missing, and in waves where “0” indicates inapplicability, these values are also set to missing.

Parent–child conflict. For each wave we combine five indicators of family conflict: whether family members fight a lot, throw things, criticize each other, hit each other, and whether they calmly discuss issues (reverse coded). Item labels differ slightly by wave, but the underlying constructs are identical. For example, in 1997 we use Q2A42A, Q2A42C, Q2A42E, Q2A42F, and Q2A42D (reverse coded). The first principal component is extracted separately for each wave and saved as `conflict_year` (e.g., `conflict_1997`).

Cognitive stimulation/resources. This index is based on the number of books in the home, the number of books the child has read, and the number of cell phones in the household. For example, in 2002 we use Q22J38 (cell phones), Q22J39 (books in home), and Q22J40 (books read by child). The first principal component is extracted and saved as `resources_year`.

Orderliness of the home environment. This index is based on interviewer observations of whether rooms are clean and whether the household is noisy inside or outside. Noise measures are reverse-coded so that higher values correspond to a quieter, more orderly household. For example, in 2002 we use Q22K30 (rooms clean), Q22K34 (noisy inside), and Q22K35 (noisy outside). The first principal component is extracted and saved as `environment_year`.

The resulting dataset contains one observation per child (identified by PSID household and person ID) with variables `conflict_year`, `resources_year`, and `environment_year` for each available wave (1997, 2002, 2007, 2014, 2019). These continuous, standardized PCA scores are used in robustness checks to test whether our shift-share instrument predicts underlying home environment quality.

This appendix reports the detailed PCA results underlying the three home environment indices described in Section X. For each construct, we show eigenvalues, proportion of variance explained, and variable loadings. In each case, the first principal component is retained as the index.

B.1 Parent–Child Conflict (1997)

Appendix Table A1: PCA results: Parent–Child Conflict (1997)

	Comp1	Comp2	Comp3	Comp4	Comp5
Q2A42A: Family fights a lot	0.469	0.093	-0.588	-0.512	0.404
Q2A42C: Family throws things	0.495	-0.218	0.351	-0.468	-0.604
Q2A42E: Family criticizes	0.483	-0.102	-0.444	0.665	-0.342
Q2A42F: Family hits each other	0.468	-0.321	0.508	0.256	0.595
Q2A42D_rev: Calmly discuss (rev.)	0.289	0.911	0.273	0.105	-0.015
Eigenvalue	2.21	0.91	0.77	0.59	0.53
Prop. Var.	0.44	0.18	0.15	0.12	0.11

B.2 Cognitive Resources (2002)

Appendix Table A2: PCA results: Cognitive Resources (2002)

	Comp1	Comp2	Comp3
Q22J38: Number of books in home	0.537	-0.686	0.491
Q22J39: Books read last month	0.665	-0.014	-0.747
Q22J40: Access to phone/cellphone	0.519	0.728	0.448
Eigenvalue	1.24	0.95	0.81
Prop. Var.	0.41	0.32	0.27

B.3 Orderly Environment (2002)

Appendix Table A3: PCA results: Orderly Environment (2002)

	Comp1	Comp2	Comp3
Q22K30: Rooms clean	0.565	-0.687	0.457
Q22K34_rev: Noise inside (rev.)	0.607	-0.030	-0.794
Q22K35_rev: Noise outside (rev.)	0.560	0.726	0.400
Eigenvalue	1.75	0.69	0.56
Prop. Var.	0.58	0.23	0.19