Intra-Household Insurance and the Intergenerational Transmission of Income Risk*

Francesco Agostinelli†
University of Pennsylvania

Domenico Ferraro‡
Arizona State University

Xincheng Qiu§
University of Pennsylvania

Giuseppe Sorrenti¶
University of Amsterdam

November 2022

(click here for the latest version)

Abstract

This paper aims to quantify the mechanisms and the extent to which parental wage risk passes through to children's skill development. Through a quantitative labor supply model in which parents choose time spent working, time spent with children, and child-related expenditures, we find that income risk lowers skill accumulation, permanently lowering children's skill levels. To the extent that making up for skill losses during childhood is hard—as the available evidence suggests—income risk with imperfect credit markets can negatively impact the labor market prospects of future generations.

Keywords: Wage risk; Household labor supply; Child development; Social insurance.
JEL Classification: D1; J13; J22.

*We thank Hector Chade and Mariacristina De Nardi for comments and suggestions. All errors are our own.

†Agostinelli: Department of Economics, University of Pennsylvania, 133 S 36th St, Philadelphia, PA 19104, USA. Email: fagostin@sas.upenn.edu.

‡Ferraro: Department of Economics, W.P. Carey School of Business, Arizona State University, PO Box 879801, Tempe, AZ 85287-9801, USA. Email: domenico.ferraro@asu.edu.

§Qiu: Department of Economics, University of Pennsylvania, 133 S 36th St, Philadelphia, PA 19104, USA. Email: qiux@sas.upenn.edu.

¶Sorrenti: Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, Amsterdam, 1018 WB, NL. Email: g.sorrenti@uva.nl.

1 Introduction

Prominent and influential literature highlights the critical role of parental investment in the early stages of childhood for the development of children's skills (see Heckman and Mosso, 2014, for a comprehensive literature review). Because of dynamic complementarity in skill accumulation, the return to parental investment is high early when the child begins acquiring critical cognitive skills. An implication of this theory, hitherto unexplored, is that uninsurable parental income risk can significantly alter how parents allocate time and resources within the family, affecting their children's skill accumulation path. For example, a temporary reduction in family earnings may induce parents to work more, thus subtracting time allocated to their child. To the extent that parental time investment is valuable in producing children's skills, and the adverse earning shock occurs early on in the child's life, life-cycle wage risk can pass through to children. Whether later investments can make up for earlier skill losses critically depends on the degree of dynamic complementarity in skill accumulation. Uninsurable idiosyncratic risk can have a "scarring effect" that permanently reduces children's skill accumulation path.

In this paper, we study how and the extent to which idiosyncratic wage risk of parents affects the intra-household allocation of time and resources, parental investment decisions, and, thereby, children's skill levels. In an ideal world with perfect insurance markets, parents could hedge parental investments in children's skills so that skill development would be, to a large extent, insulated from idiosyncratic risk. However, it is well-known that households face a nontrivial amount of uninsurable risk (Guiso, Jappelli, and Pistaferri, 2002; Blundell, Pistaferri, and Preston, 2008; Jappelli and Pistaferri, 2010; Low, Meghir, and Pistaferri, 2010; Arellano, Blundell, and Bonhomme, 2017; De Nardi, Fella, and Paz-Pardo, 2019), and that the number of existing welfare programs can only partially insure households against lost income (Autor et al., 2019). Thus, the limited ability of families with children to smooth out temporary wage shocks is an essential and realistic aspect of understanding children's outcomes and policy evaluation.

We develop a two-parent life-cycle model of endogenous children's skill formation to study the intergenerational transmission of parental income risk. In the model, both parents make consumption and labor supply decisions in the face of uninsurable, idiosyncratic wage shocks. We model labor supply along the extensive margin, allowing for nonparticipation, part-time and full-time work. This three-state representation of the labor supply decision captures the well-documented discreteness of hours worked in the data (see, e.g., French, 2005; Bick, Blandin, and Rogerson, 2022). In the model, and consistently with US data, individuals do not choose from a continuous menu of hours-wage bundles; instead, each parent has a discrete choice between working short or long hours or not working at all. Such discreetness in hours worked limits parents' ability to self-insure against wage shocks. Also, allowing for nonparticipation is critical to capture the observation that the female labor supply is quite elastic to changes in work incentives (Attanasio, Low, and Sánchez-Marcos, 2005; Attanasio et al., 2018).

The main idea in the model is that children's skills are exposed to parents' idiosyncratic wage risk. We assume that parents cannot borrow, as in Caucutt and Lochner (2020), so labor supply is the only insurance mechanism at work in response to adverse wage shocks. Children's skills accumulate over childhood as a function of parental investment. In our context, investment in a child's skill development consists of time spent with the children, which varies by parent and child-related expenditures. Time spent with children is imperfectly substitutable with child-related expenditures. When deciding how much time to spend with the children and the amount of child-related expenditures, parents also consider that their time spent with children is imperfectly substitutable.

To bring the model to the data, we combine cross-sectional moments from three datasets: the Consumer Expenditure Survey (CE), the American Time Use Survey (ATUS), the National Longitudinal Study of Youth 1979 (NLSY79), and the NLSY79 Children survey (NSLY79-C). These datasets provide information on labor market outcomes, time use, expenditures, wages, earnings, and children's test scores. More specifically, we require the model to reproduce the profile of children's skills (means and standard deviations), parents' accepted wages, earnings, and joint labor supply decisions by child's age. Further, to discipline the initial draw of children's skills at age 5, we use children's test scores and their correlation with the earnings of both parents. A child's initial skill levels correlate with parental income because of intergenerational skill transmission.

Having established that the model replicates key cross-sectional moments of the data, we use it to conduct several quantitative experiments to quantify the pass-through of wage risk on children's skills. Three main results stand out: (1) We find that a mean-preserving spread of the wage offer distributions of both parents permanently reduces children's skill levels at age 14. To the extent that these skill gaps are challenging to close later in life, as the empirical evidence suggests, parental wage risk has scarring effects on children's skills. (2) Increasing fathers' wage risk has a more significant adverse effect on skills than increasing the wage risk of mothers. The distribution of skill changes is wider for fathers than mothers, implying that heightened wage risk for fathers leads to more skill inequality than heightened wage risk for mothers. Hence, when risk shocks hit fathers, mean skill levels fall, and at the same time, children's skills become unequal relative to when the same shock hits the mothers. (3) When the wage risk shock hits both parents, we see a fall in average skill levels combined with increased skill dispersion that is larger than when wage risk rises for each parent in isolation. In this case, intra-household insurance via joint labor supply is limited since both mothers and fathers face higher earnings risk profiles over the child's age.

Turning to the transmission mechanism of income risk, we stress that disposable household income falls when the wage risk of each parent in isolation or both rises. Such a fall in disposable income is associated with a sizable reduction in child-related expenditures, whereas household consumption remains virtually unchanged. Parents facing increased wage risk respond by smoothing their consumption at the expense of child-related expenditures. However, parents raise time spent with children to offset this reduction in child-related expenditures. We also find that within-family specialization by wage risk arises as an endogenous outcome. As one parent's wage risk rises, the other parent specializes in hours toward market work, while the "riskier" parent focuses on child-related time activities. This intra-household reallocation of time results from a self-insurance motive against wage risk and, in turn, impacts child development.

Finally, we evaluate the extent to which a progressive tax system and Universal Basic Income (UBI) can mitigate the adverse effect of income risk on children's skill accumulation. We find that they do. Income risk affects relatively more low-income households with limited ability to insure against adverse wage shocks via labor supply adjustments. A more progressive tax-and-transfer system or the income floor implied by UBI can attenuate the negative consequences for child development due to higher wage risk, particularly in the case of higher life-cycle wage risk for fathers. Our results shed light on how possible future changes in life-cycle wage risk, possibly related to technological progress or structural transformation, can impact child development. Public policies can mitigate these negative consequences for the young generation.

We structure the paper as follows. In Section 2, we discuss our contribution to the literature. In Section 3, we begin with documenting novel evidence on the impact of EITC on parents' time allocations and child-related expenditures. In Section 4, we present the model. We parameterize and estimate the model in Section 5 and conduct quantitative experiments in Section 6. Section 7 concludes.

2 Related Literature

Our paper contributes to understanding how household decisions affect children's skills. We share the literature view that parents are unique actors in the process of child development (Heckman and Mosso, 2014). Previous studies in this literature shed light on the importance of monetary and time resources as investments for the skills of children (see, for example, Del Boca, Flinn, and Wiswall, 2014; Agostinelli and Sorrenti, 2022). In a recent study, Mullins (2022) analyzes the design of optimal social welfare programs in the US while accounting for the endogenous children's human capital costs and benefits of the policies. The paper finds that the optimal insurance policy should provide reduced work incentives for low-income families because of the associated costs for the skill development of disadvantaged children. Our paper builds upon the idea that policies can interfere with child development. We study to what extent the degree of progressivity and generosity of the tax-and-transfer system, which affects the labor supply incentives in the household, impacts the pass-through of a family's income risk to the next generation.

Our study is also closely related to Abbott (2022), who investigates the role of incomplete markets for parental investments. The author develops a life-cycle model of a single-parent household that endogenously allocate time and monetary resources between consumption and educational investments for a child. The study finds that in the presence of uninsured risk, credit constraints can distort human capital investments in children. Although our study abstracts from credit constraints, we show that the intra-household labor supply choices of two parents represent an important channel that affects the extent to which income-risk pass through parental investments and children's outcomes.

Several recent studies, starting with the seminal work of Cunha, Heckman, and Schennach (2010), model and estimate the technology of skill formation in children (Agostinelli and Wiswall, 2016; Del Boca et al., 2019; Attanasio et al., 2020; Attanasio, Meghir, and Nix, 2020; Caucutt et al., 2020). Our work contributes to this strand of the literature by estimating a general nested-CES (constant elasticity of substitution) skill production function that accounts for monetary investments, both parents' time investments, as well as dynamic complementarity. Del Boca, Flinn, and Wiswall (2014) develop and estimate a model of two-parent household choices, with time and monetary investments in children, and a Cobb-Douglas children's skill production function. However, the authors abstract from the role of taxes and income uncertainty on child development.

The second literature is on the role of labor supply as an insurance mechanism for household consumption in the presence of income shocks. Blundell, Pistaferri, and Preston (2008) and Blundell, Pistaferri, and Saporta-Eksten (2016) analyze the impact of income shocks on life-cycle consumption volatility of families. They find that family labor supply interacts with the tax system to provide partial insurance of income shocks to family's consumption. More related to our work, Blundell, Pistaferri, and Saporta-Eksten (2018) extend the analysis to the case of households with children. The authors consider a life-cycle model of household's time allocation between working, leisure and child care, under uncertainty about income. Parents derive utility from spending time with children. Their results suggest that time allocation between children and hours worked interact in the presence of income shocks, as parents' labor supply provides insurance against wage shocks. Our framework extends this previous work in two major directions. First, we model the technology of skill formation, which allows us to study the consequences of income uncertainty on the development of children. Second, the endogenous formation of skills depends on the degree of substitution between parental time and monetary investments in the production of a child's skills, which in turn are affected by the aim to provide insurance to household's consumption from income shocks via labor supply responses.

3 Empirical Facts

In this section, we provide novel empirical evidence on the effects of changes in work incentives on parents' time allocations and child-related expenditures. To do so, we leverage variation induced by Earned Income Tax Credit (EITC) reforms. We start by describing the data sources for the analysis. Then, we discuss the empirical strategy and results.

3.1 Data

Consumer Expenditure Survey (CE). The CE is a household-level data set with detailed information on the expenditures of U.S. households. The CE is representative at the national level and is collected by the U.S. Bureau of the Census and sponsored by the Bureau of Labor Statistics. The data set is a rotating panel that covers four consecutive quarters following about 5,000 households per year. In our analysis, for reasons of data availability and measurement precision, we focus on surveys starting from 1996.

The CE contains an interview section and a diary section. The interview is conducted quarterly with each household unit and tracks large expenditures. The diary section aims to trace smaller everyday expenditures more likely to be forgotten in the interview. The diary is conducted over two consecutive one-week periods with each respondent. The interview and diary sections are independent. Given the focus of our analysis on sizeable child-related expenses, we exclusively consider the interview section of the data.

Child-related expenditures fall into four categories: clothes, childcare, education, and toys. Childcare expenditures include babysitting and childcare at home and in someone else's, expenditures on daycare centers, nurseries, and preschools. Education expenditures include school books, supplies, equipment, tuition, and other expenses, including rentals, test preparation, and tutoring services. We separate tuition from other education expenditures in the analysis.

Other nondurable expenditures include food, utility, services, transport, personal care, as well as semidurables such as clothing (for adults). This definition of nondurable expenditures exclude expenditure on various durables such as housing (furniture, appliances, etc.), vehicle, health, and education. The definition follows Blundell, Pistaferri, and Preston (2008).

American Time Use Survey (ATUS). The ATUS is a data set of U.S. national time-diary samples. In January 2003, ATUS was implemented as an ongoing monthly survey sponsored by the Bureau of Labor Statistics. ATUS is a time-use data set for a representative sample in the U.S., and it collects time-use data within 24 hours for one respondent of each household.

We restrict the sample of interest to households with one or more children. However, data limitation makes it impossible to identify parents and children in the data precisely. Based on the age distribution in the sample, we opted for a simple (and reasonable) rule: we exclude from the sample respondents younger than 18 years old and use the gender information to distinguish the father and the mother.[1]

The activities reported in time use data are aggregated into three main categories: market work, childcare, and residual time. Time spent on market work includes paid work—including a second job or other paid work—either at home or not, work breaks, additional time at the workplace, traveling to and from work, and job-seeking activities. Time spent on childcare includes care of infants, general care of older children, medical care of children, playing with children, supervising children or helping with homework, reading to or talking with children, and other childcare activities.[2] The residual time, labeled as *leisure* time hereafter, is defined as the time in a day not devoted to work or childcare. In other words, leisure time is the complement to 24 hours after considering time dedicated to work and childcare. For instance, it includes sleeping, personal care, education, adult care, civic, voluntary, and religious activities, out-of-home free time and leisure, sports, and outdoor activities, in-home free time leisure, media, and computing.

3.2 Empirical Specification and Results

3.2.1 Child-Related Expenditures

We analyze whether changes in the tax system induce responses in household's child-related expenditures. To do so, we focus on reforms of the EITC program, the largest U.S. tax credit program for families with dependent children. Hotz and Scholz (2003) and Nichols and Rothstein (2016) describe the EITC program, its eligibility rules, and the main empirical findings on the EITC effect on outcomes such as income and labor supply.

[1] Precisely, the data on age distribution show a bunch at respondent's age 15-17 within the sample of married households with child(ren). This bunch in the distribution accounts for about 9 percent of the sample, whereas those for the age range 18-20 is negligible and amount to about 0.5 percent of the sample. This evidence suggests that, with some exceptions, respondents in the 15-17 years old group are usually "children" in the household instead of parents.

[2] We have also experimented with an alternative, slightly different definition of child care by further including travel related to child care as in Aguiar, Hurst, and Karabarbounis (2013). The results are barely changed.

The identification of EITC-induced causal effects is particularly challenging as eligibility to the program is nonrandom.[3] In our empirical analysis, we exploit the quarterly nature of our data and two features of the EITC program to deal with endogeneity concerns. First, one of the characteristics of the EITC is that tax refunds are usually paid in the second quarter of the year.[4] Second, starting in the early '90s, the EITC was considerably expanded up to becoming the largest cash transfer program nationwide. These two features of the program allow us to perform a difference-in-differences (DiD) analysis comparing (i) different individual treatment intensities as a function of the evolution of the program over time; and (ii) within-year expenditures made before receiving the tax credit (first quarter) and after the receipt of the tax credit (quarters two to four). A similar empirical strategy relies on the fact that, starting in the second quarter of a year, expenditures might increase due to the receipt of the tax credit. In other words, the implicit assumption underlying the analysis is that tax returns could be treated as an income shock and households neither plan for it in advance, nor save considerable amounts of the credit for future needs.

Our DiD specification takes the following form:

$$Expenditure_{istq} = \alpha_0 + \alpha_1 MaxEITC_{ist} + \alpha_2 Post_i + \alpha_3 MaxEITC_{ist} \times Post_i + \mathbf{X'}\boldsymbol{\gamma} + \eta_q + \varepsilon_{itq}, \quad (1)$$

where i represents the interviewee's household, s is the state of residence, t is the year, and $q \in [1, 2, 3, 4]$ stays for the quarter of the year. We express child-related expenditures y in year 2016 real dollars. The variable $MaxEITC$ captures the exposure to the EITC program. It measures the maximum tax credit, in year 2016 real dollars, that a household can receive given the number of dependent children in the household, the state of residence, and the year.[5] A similar measure is able to effectively summarize the size of the EITC program and its evolution over time. At the same time, being independent of family income, it helps in dealing with the endogeneity concerns previously mentioned. *Post* is an indicator variable taking the value of one if the interview is held in quarters two to four, and the value of zero for households interviewed in the first quarter of the year. To take into account the source of variation of the variable for EITC exposure, the vector **X** contains a full set of interactions between the indicator *Post* and the number of dependent children in the household, state fixed effects, and year fixed effects. In some specifications, we will also control for interviewee's characteristics, such as education, marital status, race, and age. Finally, the specification also includes fixed effects for quarters of the year to account for seasonal patterns in household expenditures. ε_{itq}

[3]Eligibility criteria and the amount of received benefits depend on several factors such as family composition, family income, and family members' employment status.

[4]The majority of program's beneficiaries receive the tax credit in the year following their tax declaration.

[5]The amount of EITC benefits varies with the number of dependent children (one, two, three or more) in the household. For this reason, $MaxEITC$ is indexed by the household subscript i.

Table 1: EITC Exposure and Household Expenditures

	(1) Child Expend.	(2) Child Expend.	(3) Child Expend.	(4) Non-durables	(5) Non-durables	(6) Non-durables
Policy Effect	0.05***	0.04***	0.05***	0.03	0.02	0.04
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Observations	108,218	108,218	78,038	108,218	108,218	78,038
Mean Dep.Var.	754	754	851	6728	6728	7280
Individual Controls	No	Yes	Yes	No	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Post*Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Post*Num.Children FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Whole	Whole	Married	Whole	Whole	Married

Notes: The table shows the effect of changes in the EITC program on household expenditures. Dependent variables: child-related expenditures (columns 1 to 3) and expenditures on other non-durable goods (columns 4 to 6). The Policy Effect represents the coefficient for the variable $MaxEITC_{ist} \times Post_i$ in (1). The variable is the interaction of the maximum tax credit, in year 2016 dollars, conditional on state of residence, year, and number of dependent children, and an indicator variable taking the value of 1 if the interview takes place in quarter 2, 3, or 4, and 0 if it takes place in quarter 1. Expenditures are measured quarterly and expressed in year 2016 dollars. Individual controls include education background (less than high school completion, high school completion, some college, college graduate), marital status, race, and age. Standard errors clustered at the state level are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. See text for further details on the analysis.

is the error term of the model. We are interested in the coefficient α_3 that captures the effect of exposure to changes in the tax system, specifically in the EITC program, on household's child-related expenditures.

Table 1 shows the OLS estimates of (1) with standard errors clustered at the state level, the level at which policy changes occur. Columns (1) to (3) focus on child-related expenditures and columns (4) to (6) on expenses on other non-durable goods. As our model considers two-parent households, the specifications in columns (3) and (6) focus on the restricted sample of married couples. Given the definition of the variable for EITC exposure, all results should be interpreted as intention-to-treat effects (ITT).

The analysis in the table suggests that changes in the tax system have the potential to shape responses in terms of household's expenditures. An increase in EITC generosity generates a surge in child-related expenditures. By focusing on column (1), a \$100 increase in program generosity induces a statistically significant average increase in quarterly child-related expenditures by \$5. The effect is similar for married couples (column 3). On the other hand, columns (3) to (6) show that the effect on expenditures on non-durable goods is slightly smaller in size (\$3)

Table 2: EITC Exposure and Child-related Expenditures

	(1) Childcare	(2) Child Clothes	(3) Education	(4) Tuition	(5) Toys
Policy Effect	-0.01	0.02***	0.00	0.03***	0.01***
	(0.01)	(0.00)	(0.00)	(0.01)	(0.00)
Observations	78,038	78,038	78,038	78,038	78,038
Mean Dep.Var.	335	231	42	165	78
Individual Controls	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Post*Year FE	Yes	Yes	Yes	Yes	Yes
Post*Num.Children FE	Yes	Yes	Yes	Yes	Yes
Sample	Married	Married	Married	Married	Married

Notes: The table shows the effect of changes in the EITC program on child-related expenditures. Dependent variables: childcare expenditures (column 1), expenditures on children's clothes (column 2), expenditures on children's education (column 3), expenditures on tuition (column 4), and expenditures on children's toys (column 5). The Policy Effect represents the coefficient for the variable $MaxEITC_{ist} \times Post_i$ in (1). The variable is the interaction of the maximum tax credit, in year 2016 dollars, conditional on state of residence, year, and number of dependent children, and an indicator variable taking the value of 1 if the interview takes place in quarter 2, 3, or 4, and 0 if it takes place in quarter 1. Expenditures are measured quarterly and expressed in year 2016 dollars. Individual controls include education background (less than high school completion, high school completion, some college, college graduate), marital status, race, and age. Standard errors clustered at the state level are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. See text for further details on the analysis.

and statistically insignificant. All results are remarkably similar across empirical specifications.

Table 2 further explores the effects on child expenditure by looking at separate child-related expenses. We focus on the sample of married couples for consistency with our theoretical model and we consider expenditures on childcare (column 1), clothes (column 2), education (column 3), tuition fees (column 4), and toys (column 5). Results suggest that some household's expenditures respond to changes in the EITC policy. A surge in EITC benefits does not shape significant changes in the average expenditure on childcare or in education. On the contrary, although the total expenditure on education is unchanged, we find a positive EITC effect on tuition fees. A positive and statistically significant impact of the EITC expansion is also detected for expenditures on clothes and toys.

3.3 Time Use

We examine the possible effect of changes in the tax system on parental time use. We focus again on the expansion of the EITC program, and, to consider the fact that we have yearly data on time use, we propose an alternative identification strategy with respect to the one for expenditure data. Similar to the previous analysis, the strategy aims at identifying prominent changes in the EITC program. Then, we study the effect of these changes on parental time use by separating observations based on their exposure to the EITC program.

The first step of the identification strategy relies on the definition of prominent policy changes. We define as a prominent change every change in the maximum available EITC benefits by at least \$200 dollars. We focus on the period 1998–2012 and the change is defined based on the schedule at the state-year level for couples with two children.[6] Second, we divide sample units according to their exposure to the EITC program. One of the eligibility requirements to obtain the EITC is indeed a family income below a certain threshold. Therefore, we define exposure to the policy based on family income. Individuals with family income below 110 percent of the state-year specific EITC income threshold are with *high* exposure. The remaining sample units, with income exceeding by more than 10 percent of the state-year eligibility threshold, are classified with *low* exposure to the EITC program.[7] Third, we compare time use for observations interviewed before and after prominent EITC changes to unveil possible responses to policy reforms. The analysis is repeated for families exposed to the program and, as a sort of placebo test, for those families with low exposure to the policy due to their (likely) ineligibility to obtain EITC benefits.

More precisely, we estimate the following empirical specification:

$$TimeUse_{ist} = \alpha PolicyChange_{ist} + \gamma X_i + \mu_s + \rho_t + \varepsilon_{ist}, \tag{2}$$

where $TimeUse_{ist}$ is the amount of minutes per day spent on a certain activity by parent i, living in state s, in year t. $PolicyChange$ is an indicator for a prominent, at least \$200, EITC change. The indicator is equal to zero until the first prominent change in state s where individual i resides, and equal to one for all later years. The vector $\mathbf{X}$ contains control variables for education (below high school, high school graduates, some college, college graduates), marital status, age, race, and gender. The specification also includes state (μ_s) and year (ρ_t) fixed effects. ε_{ist} is the error term of the model. We are interested in the coefficient α that represents the effect of a change in the EITC schedule on parental time use. The effect is identified out

[6] Results are robust to using different schedules (e.g., one child) or different definitions for prominent changes (e.g., \$150).

[7] We opt for 110 percent of the EITC income threshold to take into account possible measurement error in family income. Results are robust to both more- and less-conservative criteria.

of the within-state comparison of individuals before and after a sizeable change in the program generosity. It is important to recall that, given the cross-sectional nature of the data, the pre- versus post-prominent change comparison is performed out of different sample units.

Table 3 reports the estimates of (2). We estimate OLS models with standard errors clustered at the state level, the level at which prominent policy changes occur. The upper panel shows the analysis for the whole sample. Columns (1) to (3) restrict the analysis to individuals highly exposed to the EITC program, i.e., those from families reporting an income that does not exceed by more than ten percent the EITC state-year income eligibility threshold. Columns (3) to (6) replicate the analysis for families with low exposure to the EITC program. We perform the analysis on work time (columns 1 and 4), childcare time (columns 2 and 5), and leisure time (columns 3 and 6).

The analysis in the upper panel of Table 3 highlights, at least, two important results. First, for parents highly exposed to the EITC program, prominent policy changes induce significant increases in labor supply. After a prominent change in EITC benefits, the average parental labor supply increases by about 21 minutes per day. This change corresponds to an increase by more than 12 percent with respect to the average of 168 minutes per day worked by parents in our estimation sample. This result is aligned, among others, with the evidence in Agostinelli and Sorrenti (2022) and Agostinelli, Borghesan, and Sorrenti (2021) on maternal labor supply in response to changes in the EITC program. The increase in time devoted to market work is counterbalanced by a reduction in time dedicated to childcare that significantly decreases by, on average, 8 minutes per day (10 percent of the sample mean). The effect on leisure time is smaller if compared to the sample mean and it is statistically insignificant. Second, results seem to support the use of an identification strategy based on exposure to the EITC program. Indeed, none of the effects for families (likely) unexposed to the EITC (columns 3 to 6) is statistically significant. In other words, we find that families likely unexposed to the program do not respond to changes in the program's generosity.

For consistency with our two-parent household model, the bottom panel of Table 3 replicates the analysis by restricting the sample to married couples. The main conclusions of the analysis remain the same despite a reduction in estimates precision likely due to lower sample sizes. On the one hand, individuals with high exposure to the EITC tend to increase their work time at the cost of a reduction in time devoted to childcare. On the other hand, individuals unexposed to the program do not respond to policy changes.

Table 4 looks for evidence of possible within-household specialization in terms of time use by replicating the analysis by gender. The table only reports results for the subsample of individuals highly exposed to the EITC program. The upper panel considers the whole sample, and the bottom panel, in accordance to our theoretical framework, further restricts the sample to

Table 3: EITC Expansion and Time Use

	(1) Work	(2) Childcare	(3) Leisure	(4) Work	(5) Childcare	(6) Leisure
Policy Change	20.61***	-8.20**	-12.41	1.66	-1.28	-0.37
	(7.62)	(3.91)	(7.78)	(6.54)	(3.14)	(7.28)
Observations	16,829	16,829	16,829	29,472	29,472	29,472
Mean Dep.Var.	167.91	80.11	1191.98	203.8	87.31	1148.88
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Policy Exposure	High	High	High	Low	Low	Low
Sample	Whole	Whole	Whole	Whole	Whole	Whole
	(1) Work	(2) Childcare	(3) Leisure	(4) Work	(5) Childcare	(6) Leisure
Policy Change	24.62*	-11.11*	-13.52	-2.09	-0.38	2.47
	(12.55)	(6.16)	(14.00)	(7.84)	(3.67)	(8.62)
Observations	8,523	8,523	8,523	21,934	21,934	21,934
Mean Dep.Var.	186.35	94.61	1159.04	226.28	107.62	1106.1
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Policy Exposure	High	High	High	Low	Low	Low
Sample	Married	Married	Married	Married	Married	Married

Notes: The table shows the effect of changes in the EITC program on parental time use. Dependent variables: work time (columns 1 and 4), childcare time (columns 2 and 5), and leisure time (columns 3 and 6). The upper panel considers the whole sample. The bottom panel restricts the sample to married individuals. The Policy Change is the coefficient for an indicator variable for prominent changes, at least $200 difference in maximum available benefits, in the EITC schedule. The indicator is equal to 0 until the first prominent change in state s where individual i resides, and equal to 1 for all later years. Time use is measured in minutes per day. Individual controls include education background (less than high school completion, high school completion, some college, college graduate), marital status, race, gender, age and age squared. Standard errors clustered at the state level are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. See text for further details on the analysis.

married couples. The analysis unveils some patterns of within-household specialization. Indeed, for both the whole sample and the sample of married couples, men tend to respond more to policy changes than women. If we consider the sample of married men (columns 4 to 6), they tend to significantly increase labor supply and reduce leisure in response to a prominent policy change. Also childcare is reduced. However, the effect on childcare, although non-negligible in size (16 percent of the sample average), is statistically insignificant. The response of married women (columns 1 to 3) to policy changes seem negligible.

Overall, this analysis sheds some light on parental time use responses to changes in the tax system. Specifically, we find suggestive evidence that a tax system providing incentives for individuals' labor supply, one of the requirement to be eligible for EITC benefits, can induce reductions in childcare and leisure time counteracted by a positive and significant surge in work time.

Table 4: EITC Expansion and Time Use by Gender

	(1) Work	(2) Childcare	(3) Leisure	(4) Work	(5) Childcare	(6) Leisure
Policy Change	11.70	-7.78	-3.93	40.79**	-9.70*	-31.08*
	(8.44)	(4.66)	(9.08)	(18.08)	(4.95)	(18.19)
Observations	11,104	11,104	11,104	5,725	5,725	5,725
Mean Dep.Var.	140.38	96.99	1202.63	221.32	47.36	1171.32
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Policy Exposure	High	High	High	High	High	High
Sample	Whole Women	Whole Women	Whole Women	Whole Men	Whole Men	Whole Men

	(1) Work	(2) Childcare	(3) Leisure	(4) Work	(5) Childcare	(6) Leisure
Policy Change	3.73	-12.05	8.32	52.75**	-10.26	-42.49*
	(11.14)	(9.76)	(12.84)	(23.04)	(7.27)	(23.70)
Observations	4,884	4,884	4,884	3,639	3,639	3,639
Mean Dep.Var.	118.38	118.61	1203.01	277.57	62.4	1100.03
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Policy Exposure	High	High	High	High	High	High
Sample	Married Women	Married Women	Married Women	Married Men	Married Men	Married Men

Notes: The table shows the effect of changes in the EITC program on parental time use. Dependent variables: work time (columns 1 and 4), childcare time (columns 2 and 5), and leisure time (columns 3 and 6). The upper panel considers the whole sample. The bottom panel restricts the sample to married individuals. The analysis is only performed for individuals with high exposure to the EITC program. Each panel provides a separated analysis for women (columns 1 to 3) and men (columns (4 to 6). The Policy Change is the coefficient for an indicator variable for prominent changes, at least $200 difference in maximum available benefits, in the EITC schedule. The indicator is equal to 0 until the first prominent change in state s where individual i resides, and equal to 1 for all later years. Time use is measured in minutes per day. Individual controls include education background (less than high school completion, high school completion, some college, college graduate), marital status, race, gender, age and age squared. Standard errors clustered at the state level are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. See text for further details on the analysis.

4 Two-Parent Life-Cycle Model with Children Skill Formation

In this section, we develop the model we later use to quantify how household earnings shocks pass through to children's skills. Within a unitary model of a two-parent household, we study the labor supply of both parents, each parent deciding whether to work short or long hours or not working at all, and the allocation of the two parents' time and goods to the formation of the sole children's skills.[8] The household makes such decisions in the face of the idiosyncratic risk that takes the form of shocks to wages and the technology of children's skill formation. The model features a progressive tax-and-transfer system, which allows us to evaluate the effects of several public policies, including means-tested transfers and changes in tax progressivity.

4.1 Preferences and Budget Set

The household consists of two parents with identical preferences and a child. We model parental decisions over the first T periods of the child's life, which we associate with the "childhood." The child's age is denoted by t. Each parent is endowed with one unit of time per period that can be allocated to three uses: leisure, hours of work, and time spent with the child.

Preferences. Household's preferences over consumption, c_t, leisure, l_{jt}, of parent $j = \{1, 2\}$, and child's skills θ_t are described by

$$\mathbb{E}_0 \left\{ \sum_{t=0}^{T} \beta^t u\left(c_t, l_{1t}, l_{2t}, \theta_t\right) + \beta^{T+1} v\left(\theta_{T+1}\right) \right\}, \tag{3}$$

where u and v are increasing, concave, and twice continuously differentiable. $\mathbb{E}_0$ indicates the mathematical expectation operator based on available information at childbirth, and β is the time discount factor. Expectations are with respect to the wage offers of both parents, and to the shocks to the technology of skill formation, which we describe later.

Following Del Boca, Flinn, and Wiswall (2014), we allow the parents' valuation of the child's skills over the stage of early childhood, $\{\theta_t\}_{t=0}^{T}$, to differ from that of the terminal value θ_{T+1}. We view θ_{T+1} as the initial condition for the next stage of child development which we leave unspecified here.

Budget set. Household's income consists of labor earnings of possibly both parents, net of taxes and transfers, $\mathcal{T}\left(\sum_{j=1}^{2} w_{jt} h_{jt}\right)$, where w_{jt} and h_{jt} denote respectively the wage and the

[8]See Chiappori and Mazzocco (2017) for a survey article on the pros and cons of the unitary model of household behavior.

labor supply of parent $j \in \{1, 2\}$, and the function $\mathcal{T}$, whose parameterization we describe later, encompasses all the details of the tax-and-transfer system. To capture the well-documented discreteness of hours worked in the data (see, e.g., French, 2005; Bick, Blandin, and Rogerson, 2022), we model the labor supply of each parent as a discrete choice between long-hours, $\overline{h}$, short-hours of work, $\underline{h} < \overline{h}$, and not working $h = 0$.

The household's budget constraint is

$$c_t + e_t = \mathcal{T}\left(\sum_{j=1}^{2} w_{jt} h_{jt}\right), \tag{4}$$

where c_t is consumption expenditures and e_t is expenditures related to the child's skill formation.

4.2 Technology of Children Skill Formation

Following Cunha and Heckman (2007) and Cunha, Heckman, and Schennach (2010) we posit a technology of skill formation according to which the child's skills next period θ_{t+1} depend on their current level θ_t, parental time investments, m_{1t}, and m_{2t}, and expenditures on the child, e_t:

$$\theta_{t+1} = \exp\left(z_t\right) f\left(\theta_t, e_t, m_{1t}, m_{2t}\right), \tag{5}$$

where the function f is increasing and concave in each input $(\theta_t, e_t, m_{1t}, m_{2t})$. To capture the inherent uncertainty in the process of skill accumulation, we also assume that the process of child development is hit by stochastic shocks that disturb the return to parental investment for a given child's skill level, akin to random shifts in the production possibility frontier in standard production theory. Such technology shocks follow an AR(1) process:

$$z_{t+1} = \mu_z\left(1 - \rho_z\right) + \rho_z z_t + \sigma_\eta \eta_{t+1}, \quad \text{with} \quad \eta_{t+1} \overset{iid}{\sim} \mathcal{N}\left(0, 1\right), \tag{6}$$

where μ_z is the unconditional mean of the shocks, ρ_z governs the persistence of such shocks, and σ_η is the standard deviation of independent and identically distributed (iid) normal innovations, η_t. The initial level of the child's skills θ_0 is random and drawn from a distribution that allows θ_0 to be correlated with the wage offers of both parents, in a parametric way that we describe in estimation section.

4.3 Wage Processes

As we consider labor supply decisions along the extensive margin, in the model, there is a distinction between the distribution of offered wages, a latent object, that we take as exogenous, and the distribution of accepted wages, that we can instead read from the data and contrast it with that implied by the model.

More specifically, in the model, each parent has three available options at hand: not-working, working short-hours, and working long-hours. Intuitively, whether a parent works or not, and whether he or she works short- or long-hours depends on the labor supply decision of the other parent. In fact, for each parent, there are four reservations wages that fully describe his or her labor supply. To see this more clearly, think of a one-parent family. In that setup, the labor supply decision of the single parent is described by two reservation wages; one defined as the value of the wage that makes the parent indifferent between working short-hours and not-working, another that makes the parent indifferent between working short-hours and working long-hours. Now, in the case of a two-parent family, things change as one parent's indifference among alternative labor supply choices generally depends on whether the other parent is working, and if so, how many hours he or she is working. This interdependence comes from the sharing of resources within the family, which generates income effects that operate through changes in each parent's consumption.

Concretely, we assume that the log of the wage offered to parent $j \in \{1, 2\}$ at time t is

$$\log w_{jt} = a_j + b_j t + \varepsilon_{jt}, \tag{7}$$

where b_j is the growth rate of the wage, which when we turn to estimation we allow to vary by the parent's gender, and ε_{jt} is an innovation that we assume to follow an AR(1) process:

$$\varepsilon_{jt} = \rho_j \varepsilon_{jt-1} + \sigma_{\nu j} \nu_{jt}, \quad \text{with} \quad \nu_{jt} \overset{iid}{\sim} \mathcal{N}(0, 1), \tag{8}$$

where the persistence parameters ρ_j and standard deviations $\sigma_{\nu j}$ are again allowed to vary by the parent's gender. The initial wage offers $\{w_{10}, w_{20}\}$ are random and drawn from a joint Normal distribution that allows for the offers of both parents to be correlated.

Finally, note that conditional on working, the wages of possibly both parents fluctuate over the life-cycle, subject to stochastic shocks. This is consistent with the large and growing evidence that US households face substantial idiosyncratic earning risk (see, e.g., Blundell, Pistaferri, and Preston, 2008; Blundell, Pistaferri, and Saporta-Eksten, 2016, 2018; Heathcote, Storesletten, and Violante, 2014). More specifically, wages grow over time at a constant rate, which is allowed to differ by parent's gender, and are hit by transitory but persistent shocks around their trend. Given the life-cycle nature of the model, such transitory shocks can have nontrivial permanent effects insofar as they occur sufficiently close to T. The extent of such permanent effects depends on the magnitude and persistence of wage shocks, empirical objects that we estimate from the data.

4.4 Household's Problem

We now turn to study the household's problem and its implications for the child's skills. To do so, we begin with a simplified version of the problem in which we shut down borrowing and saving altogether. This allows us to zoom in on the intertemporal tradeoff from the parental investment decision in the child's skill formation. Notably, the problem remains dynamic.

Given the exogenous wage processes described by (7)-(8), the initial wage offers $\{w_{j0}\}_{j=1}^{2}$ for both parents, the initial level of the child's skills θ_0, and the initial level of savings a_0, the household's maximization problem is given by:

$$\max_{\left\{c_t, e_t, \{h_{jt}, l_{jt}, m_{jt}\}_{j=1,2}\right\}_{t=0}^{T}} \mathbb{E}_0 \left\{ \sum_{t=0}^{T} \beta^t u\left(c_t, l_{1t}, l_{2t}, \theta_t\right) + \beta^{T+1} v(\theta_{T+1}) \right\}, \tag{9}$$

subject to

$$l_{jt} + h_{jt} + m_{jt} = 1, \tag{10}$$

$$c_t + e_t = \mathcal{T}\left(\sum_{j=1}^{2} w_{jt} h_{jt}\right) + a_0, \tag{11}$$

$$\theta_{t+1} = \exp\left(z_t\right) f\left(\theta_t, e_t, m_{1t}, m_{2t}\right). \tag{12}$$

To cast the household's problem in recursive form, we let $\mathbf{s} \equiv \{\log w_1, \log w_2, z\}$ denote the vector of exogenous state variables, that includes log wages for each parent and log TFP of the child's skill formation technology. The endogenous state variable is the current level of child's skills θ.

The value function for $t \leq T$ thus satisfies the following Bellman equation:

$$V_t\left(\theta, \mathbf{s}\right) = \max_{c_t, e_t, \{h_{jt}, l_{jt}, m_{jt}\}_{j=1,2}} \left\{ u\left(c, l_1, l_2, \theta\right) + \beta \mathbb{E}_{\mathbf{s}'|\mathbf{s}} V_{t+1}\left(\theta', \mathbf{s}'\right) \right\}, \tag{13}$$

subject to the constraints (10)-(12). The expectation is taken with respect to both parents' wage offer shocks, as well as the child development shock.

5 Bringing the Model to the Data

In this section, we describe the model's parameterization. We begin by laying out the parametric assumptions about preferences, technology, and the tax-and-transfer system. We then describe how we pin down parameter values, which involves exogenously setting the values of a subset of parameters, and estimating the remaining ones by the method of moments. Finally, we discuss the model's goodness of fit.

5.1 Functional Forms

Preferences. We assume the following functional form for the household's utility function:

$$u\left(c,\theta,l_1,l_2\right)=\frac{c^{1-\gamma_c}}{1-\gamma_c}+\alpha_\theta\frac{\theta^{1-\gamma_\theta}}{1-\gamma_\theta}+\alpha_{l_1}\frac{l_1^{1-\gamma_{l_1}}}{1-\gamma_{l_1}}+\alpha_{l_2}\frac{l_2^{1-\gamma_{l_2}}}{1-\gamma_{l_2}}. \tag{14}$$

The triplet $(\alpha_\theta,\alpha_{l_1},\alpha_{l_2})$ governs the relative weight of children's skills and parents' leisure in household's utility, respectively. The parameters $(\gamma_c,\gamma_\theta,\gamma_{l_1},$ and $\gamma_{l_2})$ pin down the elasticities of the marginal utilities of consumption, child's skills, and parents' leisure with respect to c, θ, l_1, and l_2, respectively.

Technology. The technology of skill accumulation is a multi-layer nested constant-elasticity-of-substitution (CES) production function, which allows for various degrees of substitution between child-related expenditures, parental time, and dynamic complementarity in child development:

$$\theta'=\exp\left(z\right)f; \tag{15}$$

$$f=\left(\omega_f\theta^{\frac{\sigma_f-1}{\sigma_f}}+\left(1-\omega_f\right)I^{\frac{\sigma_f-1}{\sigma_f}}\right)^{\frac{\sigma_f}{\sigma_f-1}}; \tag{16}$$

$$I=\left(\omega_I e^{\frac{\sigma_I-1}{\sigma_I}}+\left(1-\omega_I\right)M^{\frac{\sigma_I-1}{\sigma_I}}\right)^{\frac{\sigma_I}{\sigma_I-1}}; \tag{17}$$

$$M=\left(\omega_M m_1^{\frac{\sigma_M-1}{\sigma_M}}+\left(1-\omega_M\right)m_2^{\frac{\sigma_M-1}{\sigma_M}}\right)^{\frac{\sigma_M}{\sigma_M-1}}. \tag{18}$$

The function f captures how current skills, θ, and a composite of "household investment," I – a CES aggregate of child-related expenditures, e, and parental time, M – contribute to child's skill formation. Note that we allow for a flexible degree of substitutatibility between the two parental time inputs. For each of the three CES aggregators, the parameter σ determines the elasticity of substitution between inputs, while the parameter ω governs input shares. As mentioned in Section 4.2, the process of skill development is assumed to be stochastic and subject to persistent shocks, z_t.

Taxes and Transfers. The US tax-and-transfer system is rather complex, involving a mix of means-tested transfers and progressive income taxation with multiple tax brackets and tax deductions and credits. To parsimoniously capture such complexity, the literature typically adopts a simple, yet empirically plausible tax function, which can be readily estimated using TAXSIM data. Specifically, we assume that after-tax household's income is

$$\mathcal{T}(Y) = \chi_1 (\chi_3 + Y)^{1-\chi_2}, \tag{19}$$

where Y is total household earnings, the sum of the two parents' earnings, χ_1 measures the overall level of taxation, χ_2 measures tax progressivity, and χ_3 is the unconditional minimum income available in the case of zero earnings. When $\chi_1 = 1$ and $\chi_2 = \chi_3 = 0$, income is not taxed. On the other hand, when $\chi_2 = \chi_3 = 0$, taxes are proportional to income, $\mathcal{T}(Y) = \chi_1 Y$. Finally, when $\chi_2 > 0$ the tax system is progressive.

Child Skills' Terminal Value. In the model, parents make decisions over four periods ($T = 4$), while child development unfolds over five periods ($T + 1$). Each period corresponds to two calendar years. The first model period coincides with when a child is 5-6 years old and the last with when the child is 13-14 ($T + 1$). This choice is driven by the structure of the main data we use for estimating the model (National Longitudinal Study of Youth 1979), which features waves of data collections every other year.[9]

Finally, we close the model by specifying a final period terminal value for the dynamic problem. This continuation value aims to capture the stream of future utility household receive as a function of a child's end-of-childhood skills. This function captures the altruistic preference of household for their children's adult life. This specification follows a similar approach in Del Boca, Flinn, and Wiswall (2014) and is computationally attractive because it avoids to solve for the entire children's future life-cycle as a function of their inherited skills.[10] Specifically, we assume

$$V_{T+1}(\theta) = \tilde{\alpha}_\theta \frac{\theta^{1-\tilde{\gamma}_\theta}}{1-\tilde{\gamma}_\theta}, \tag{20}$$

and allow the parameters $(\tilde{\alpha}_\theta, \tilde{\gamma}_\theta)$ to differ from the preference parameters specified for all the previous periods.

[9]Interviews were conducted annually before 1994, but children's information were collected only every other year starting from 1986. From 1994 onward, both adults and children data collections occurred biannually.

[10]A similar approach has been used in different context by De Nardi (2004), who model the parents' utility of bequeathing their financial assets to their children.

5.2 Additional Data for Structural Estimation

For the estimation of the model, we also take advantage of the National Longitudinal Study of Youth 1979 (NLSY79), and their Children survey (NSLY79-C). Our sample represents the subsample of mothers of the original representative sample of youth in the 1979. Individuals in this data are followed over time, and their children are survey every two years starting from 1986. Data include labor market information for mothers and their spouses, such as employment status, hours worked and earnings. The accepted hourly wages for both parents were constructed by dividing annual earnings per annual hours worked.

The data set includes relevant information about test scores of children during childhood. We use as our measure of skills the available scores in the Peabody Individual Achievement Test (PIAT) in mathematics and reading. To construct our index of child development, we average the three PIAT raw scores available (Math, Reading-Comprehension and Reading-Recognition) for every child, and we divide by the total maximum points available in PIAT (84). The constructed index varies between 0 (no correct answer) and 1 (the child scored a perfect score in all of the three tests).

5.3 Estimation Procedure

Our estimation algorithm consists of two "steps." In the first step, we calibrate certain model's parameters directly. This allows us to reduce the computation burden of the estimation of the full model. In the second step, we estimate the rest of the structural parameters via the Simulated Method of Moments (SMM) estimator.

5.3.1 Exogenously Set Parameters

We exogenously set the values of a subset of parameters: (i) the discount factor, β; (ii) the triplet of the tax-and-transfer system, χ_1, χ_2, and χ_3; and (iii) the initial distribution of skills at age 5. A summary of our approach is as follows:

- The annual discount factor is set to be 0.96, which implies a one model's period (two calendar years) β of 0.9216 ($\beta = 0.96^2 = 0.9216$).
- The tax-and-transfer system $\mathcal{T}(Y)$ is approximated based on TAXSIM data. TAXSIM is the NBER program that calculates individual-level tax liabilities based on simulating the actual US tax-and-transfer system. We focus on the tax system as it is in 1996. We adopt a nonlinear least-squares procedure to estimate the tax function in (19). Note that in the model, a period is two years, so that Y denotes two-year household labor income,

and $\mathcal{T}(Y)$ is its after-tax counterpart. The obtained tax parameters are: $\chi_1 = 3.41$, $\chi_2 = 0.118$, $\chi_3 = 2391$. Figure 1 shows the goodness of fit.

- The initial distribution of skills at age 5 (first period of our model) is calibrated directly taken from NLSY data. In particular, the initial heterogeneity of skill endowment among 5-6 year olds is parameterized by a normal distribution $\theta_5 \sim N(\mu_{\theta,5}, \sigma^2_{\theta,5})$, and the mean and variance is estimated directly in the data. We find that the initial distribution of skills has a mean of 0.1893 and a standard deviation of 0.0586.

Figure 1: Tax System

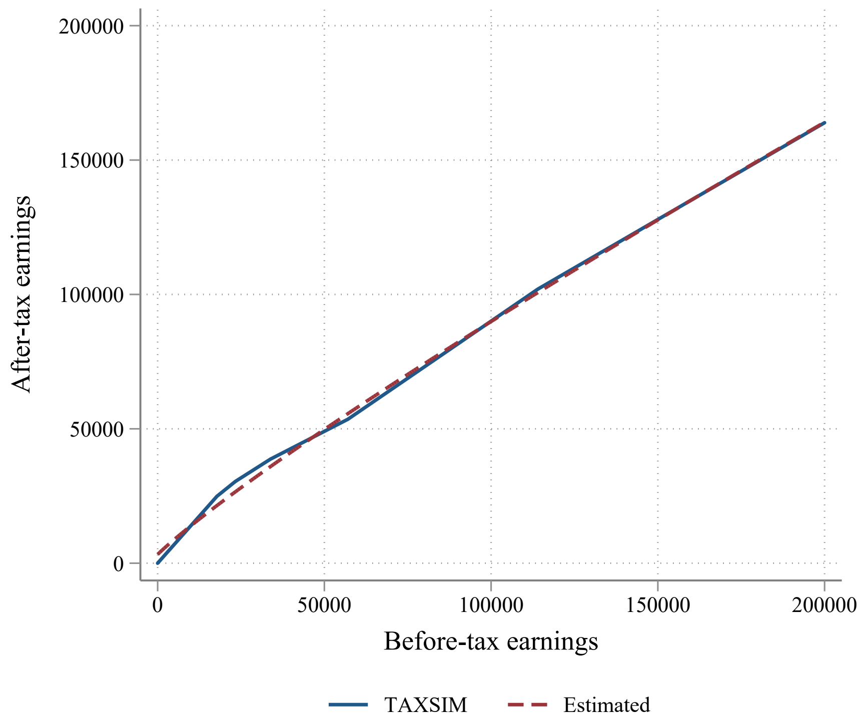

Notes: This figure illustrates the good fit of the parametric tax function. The blue solid line is the data from TAXSIM, and the red dashed line is the estimated tax function.

5.3.2 Estimated Parameters

We estimate the rest of the 29 parameters "internally" by the simulated method of moments (SMM). The structural parameters that we aim to estimate are:

- Flow utility $u(c, l_1, l_2, \theta)$: $\gamma_c, \alpha_\theta, \gamma_\theta, \alpha_{l_1}, \gamma_{l_1}, \alpha_{l_2}, \gamma_{l_2}$.

- Reduced-form terminal condition $V_{T+1}(\theta)$: $\tilde{\alpha}_\theta, \tilde{\gamma}_\theta$.

- Production function $f(\theta, e, m_1, m_2)$: $\omega_f, \omega_I, \omega_M, \sigma_f, \sigma_I, \sigma_M$.
- TFP process: $\mu_z, \rho_z, \sigma_\eta^2$.
- Wage process: $\left\{a_j, b_j, \rho_j, \sigma_{\nu j}^2\right\}_{j=1}^2$.
- Initial joint distribution of wages and skill: $\rho_{w_1,w_2}, \rho_{w_1,\theta}, \rho_{w_2,\theta}$.

We estimate wage processes' parameters and initial joint wages and skill distribution inside the model to account for the fact that we observe wage offers solely for individuals who accepted them. In related work, which also includes the extensive margin of labor supply, Blundell, Pistaferri, and Saporta-Eksten (2018) estimate the wage equations outside the model, correcting for sample selection via Heckman correction method (Heckman, 1979). Our estimation method is internally valid concerning the selection of observed wages induced by the extensive margin of labor supply for both parents.

5.3.3 Simulated Method of Moments

The simulated method of moments requires to solve and to simulate the model to construct the model's counterpart of the empirical moments we observe in the data.

Algorithm 1. *The estimation algorithm, which starts with an initial n=0 guess for the above 29 structural parameters* (Ω^0)*, follows the following steps:*

- ***Step 1.*** *Given the current n-th guess of the vector of parameters* (Ω^n)*, we solve via backward induction for the policy functions for all the endogenous choices*

$$\{c_t(\cdot\,;\Omega^n), e_t(\cdot\,;\Omega^n), m_{1t}(\cdot\,;\Omega^n), m_{2t}(\cdot\,;\Omega^n), l_{1t}(\cdot\,;\Omega^n), l_{2t}(\cdot\,;\Omega^n), h_{1t}(\cdot\,;\Omega^n), h_{2t}(\cdot\,;\Omega^n)\}_{t=0}^T$$

- ***Step 2.*** *We simulate the joint distribution of exogenous state variables in the population:*

$$\{\theta_{i,0}, \{w_{i1t}, w_{i2t}, z_{it}\}_t\}_i\ ,$$

  which includes the realization of the initial heterogeneity in the first period skill endowments,the distribution of life-cycle wage paths for each household, as well as the distribution of life-cycle TFP paths for each household.

- ***Step 3.*** *Given the simulated distribution of exogenous state variables, we use the policy functions from Step 2 to generate the distribution of the endogenous choices and outcomes:*
  - *The endogenous life-cycle path of children's skills* $\left\{\{\theta_{it}\}_{t=1}^{T+1}\right\}_i$

- *labor supply* $\{h_{1t}, h_{2t}\}_{t=0}^{T}$
- *leisure* $\{l_{1t}, l_{2t}\}_{t=0}^{T}$
- *time with kids* $\{m_{1t}, m_{2t}\}_{t=0}^{T}$
- *consumption and expenditures in kids* $\{c_t, e_t\}_{t=0}^{T}$

- ***Step 4.*** *We use the simulated distributions from Step 3 to construct the simulated counterparts* (M^{Sim}) *of the empirical moments estimated in the data* (M^{Data}). *This allows us to construct the SMM objective function, and to update the guess of parameters:*

$$\Omega^{n+1} = \arg\min_{\Omega}(M^{data} - M^{Sim}(\Omega))'(M^{data} - M^{Sim}(\Omega))$$

- ***Step 5.*** *We set n=n+1 and repeat Step 1 until the parameters converge* ($\Omega^n = \Omega^{n+1}$).

Set of Moments. We select several informative moments by combining multiple data sources. Although the estimator uses all the information from the moments jointly, which does not allow to link a particular moment to the identification to a particular parameter, we provide the intuition behind the value for identification of each set of moments. The first set of moments are calculated from the NLSY79-C data:

- A set of four correlations between a child's skills and both spouses' accepted wages and earnings. This set of moments are informative about the joint distribution of spouses' wage offers and the skills of a child.

- A set of ten moments characterizing the distribution of spouses' joint labor supply decisions and the correlation of accepted wages between spouses. These moments are informative about relative importance of leisure in the household utility, the relative importance of parental time in the production function, as well as the complementary of maternal and paternal time investments in the production of skills.

- A set of thirty-two moments describing the life-cycle profile of maternal and paternal hours worked, accepted wages, as well as the part-time and full-time employment rates for both mothers and fathers. This information will help us to understand the consumption-leisure-parenting trade-off, as well as the wage offer functions.

- A set of ten moments about the life-cycle profile of mean and standard deviation of skills, which will be informative about the production of skills and child development.

- A set of two moments from ATUS and CE data which define the relative expenditure and time investments between mothers working full-time and other mothers. This information is informative about the complementarity between monetary and time investments in skill production, as well as on the substitution between hours worked and parental investments in a child's skills.

In total we have 58 moments, which are used to identify 29 parameters. The model's estimates and sample fit of all moments are reported in the following subsection.

5.4 Parameter Estimates and Model Fit

Table 5 summarizes parameters' estimates. Next, instead of discussing each parameter value in isolation, we quantify the implications of the estimated parameters for the model's ability to reproduce many data moments. For example, in our model, the intra-household allocation of the time spent with the child and child-related expenditures are critical determinants of children's skill accumulation. Hence, getting a good model fit of the parents' wage acceptance decisions, joint labor supply, and their relationship with children's skills is paramount for quantifying the income risk pass-through and policy valuation.

To begin, Table 6 shows that the model does reasonably well in generating the observed correlations between child skills and parents' hourly wages (w_1, w_2) and earnings (y_1, y_2), and the correlation between the hourly wages of the two parents. These correlations are unconditional; we calculate them as an average over childhood (children 5-14 years old). However, extensive literature finds that the children's skill development process critically depends on the early stages of childhood, suggesting that the timing of parental investments matters. Figure 2 shows the model successfully replicates the skill profiles in terms of means and standard deviations, parents' accepted wages, hours worked, and employment rates *conditional* on child age.

Figure 3 shows the in-sample fit of the distribution of joint household labor supply decision. The nine different pairs represent the combinations of non-working, part-time, and full time working among parents. The blue bars represent the empirical shares of families for each possible case. The red bars represent the predicted model's counterpart. Overall, the model replicates fairly well the observed joint labor supply decisions among parents: fathers are more likely to work full time, and the mode of the sample is that both parents work. When father does not work full time, the mode is that mothers work full time. These facts are replicated by the estimated model.

Finally, the last two moments wwe matched are from CE and ATUS data. They are not plotted, but reported here. In particular, the model replicates the gaps in both expenditure and time investments between mothers working full-time and the rest of other mothers. We find that mothers who work full-time spend 0.1437 log-points more in monetary investments (0.1367 in the model), and they spend -0.4809 log-points less in time investments (-0.4151 in

Table 5: Parameter Estimates

Symbol	Value	Symbol	Value
A. Preferences			
γ_c	2.5462		
γ_θ	0.2267	α_θ	4.7032
γ_{l_1}	5.7242	α_{l_1}	49.8411
γ_{l_2}	5.7680	α_{l_2}	64.1651
$\tilde{\gamma}_\theta$	2.4733	$\tilde{\alpha}_\theta$	18.3934
B. Technology			
ω_f	0.9069	σ_f	0.9902
ω_I	0.1445	σ_I	2.9109
ω_M	0.5830	σ_M	0.5883
μ_z	0.4206	ρ_z	0.4312
σ_η	0.2255		
C. Wage Process			
a_1	2.9400	b_1	0.0576
ρ_1	0.5345	$\sigma_{\nu 1}$	0.6700
a_2	1.7301	b_2	0.1359
ρ_2	0.6159	$\sigma_{\nu 2}$	0.8744
$\rho_{w_1,\theta}$	0.4508	$\rho_{w_2,\theta}$	0.1931
ρ_{w_1,w_2}	0.0519		

Notes: This table reports the model's parameter estimates. Parameters are estimated by the simulated method of moments.

the model). The information coming from these two moments is important, as it highlights the implicit trade-off between labor supply and both time and monetary investments in children.

Figure 2: Goodness of Fit: Life-Cycle Profiles

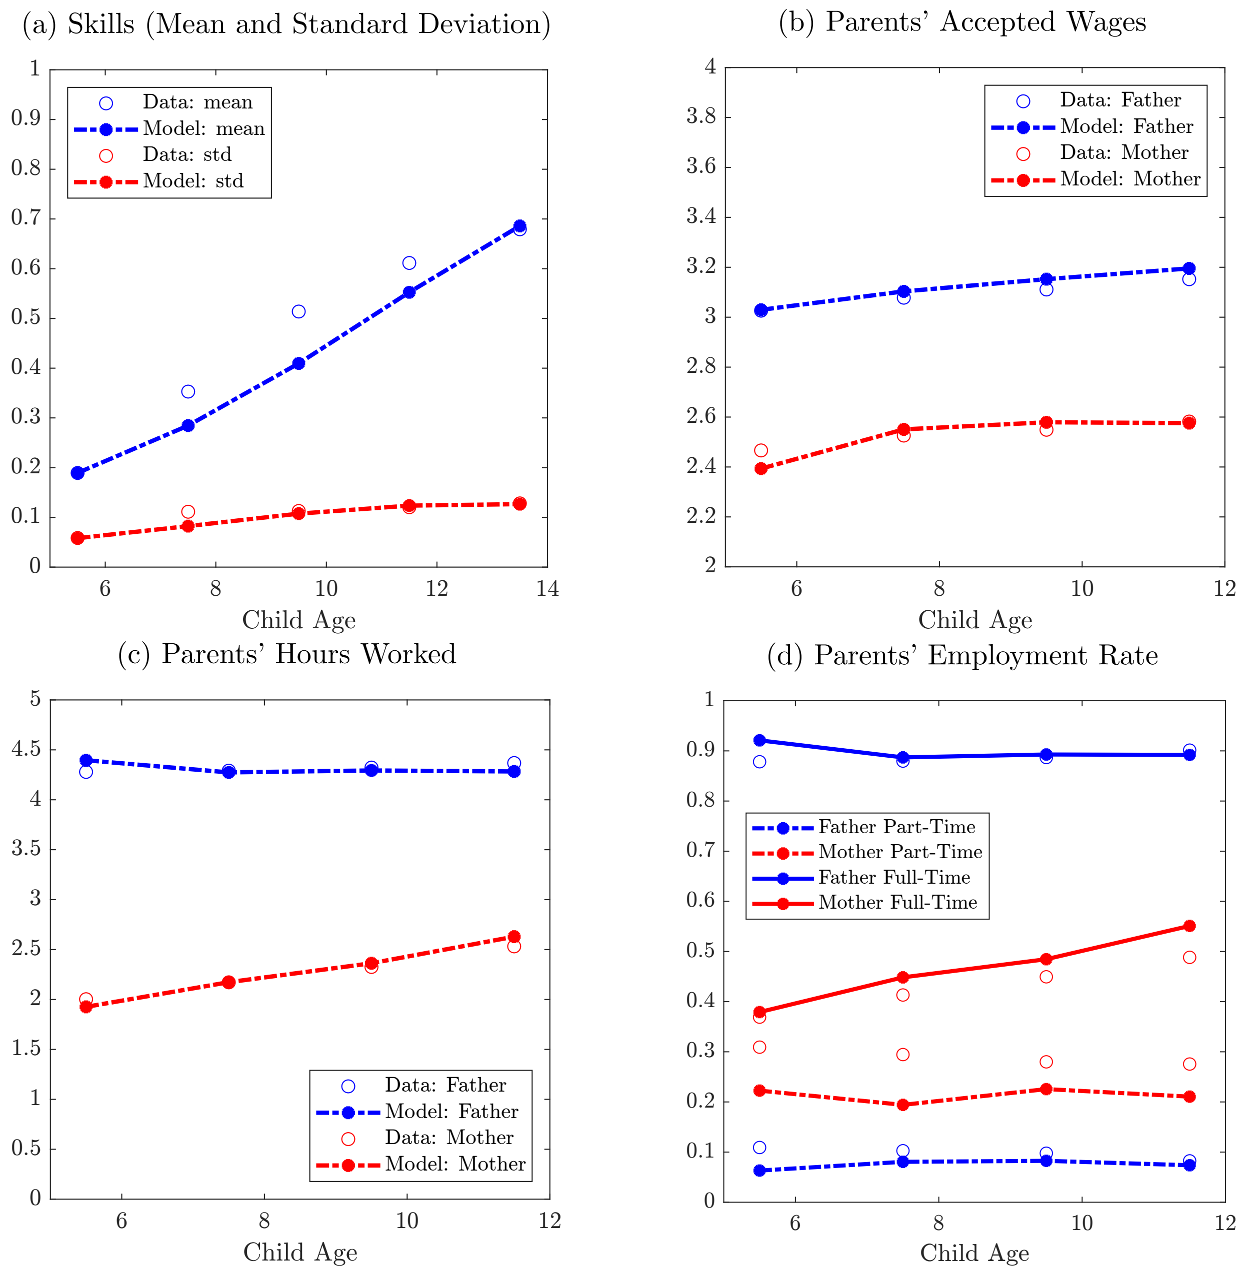

Notes: Panel (a) shows the in-sample fit of the estimated model with respect to the life-cycle profile of means and standard deviations of children's skills. Panels (b)-(d) show the in-sample fit with respect to the moments describing the life-cycle profile of maternal and paternal accepted wages, hours worked, as well as the part-time and full-time employment rates, respectively.

Table 6: Correlation between Parent Wages and Child Skills

Moment	Data	Model
corr (θ, w_1)	0.1721	0.1845
corr (θ, w_2)	0.1080	0.0875
corr (θ, y_1)	0.1556	0.1719
corr (θ, y_2)	0.1324	0.1210
corr (w_1, w_2)	0.2962	0.3552

Notes: This table shows the in-sample fit of the estimated model. The five moments represent the correlations between children's skill levels and accepted wages and earnings of both parents, and the correlation between the accepted wages of the parents.

Figure 3: Distribution (Frequency) of Joint Labor Supply

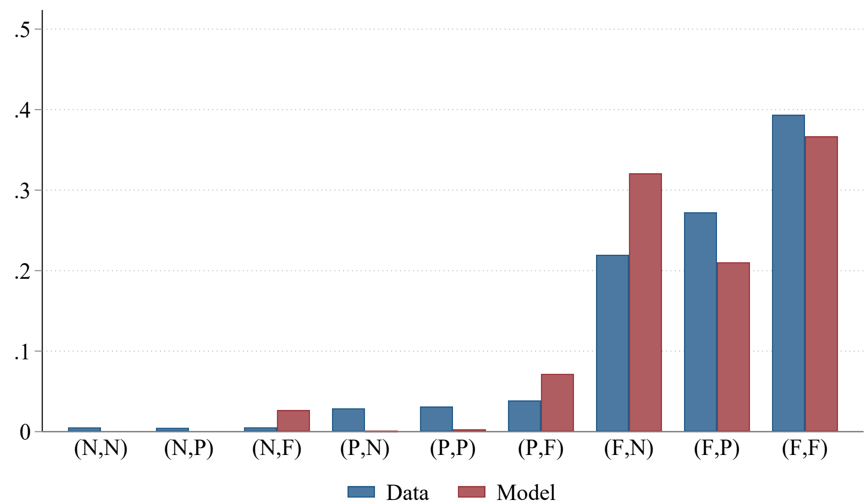

Notes: The figure shows the in-sample fit of the estimated model with respect to the distribution of spouses' joint labor supply. The nine different pairs represent the nine possible combinations of labor supply decisions (non-working, part-time and full-time) between parents. The first element of the pairs is for fathers, while the second element represents the mothers' labor supply decision.

6 Quantifying the Pass-Through of Parental Income Risk

In this section, we implement several experiments aimed at quantifying the mechanisms and the extent to which parental wage risk passes through to children's skill development.

6.1 Wage Risk Shocks and Child Development

We begin by looking at a *how mean-preserving spread* of the parents' wage offer distributions affect child development. We perform three exercises: (1) we vary the standard deviation of father's wage shock from 0.5 to 1.5 its estimated level, while keep everything else fixed; (2) we vary the standard deviation of mother's wage shock from 0.5 to 1.5 its estimated level; (3) we vary the standard deviation of both the father and mother's wage shock from 0.5 to 1.5 their estimated level. The outcome variable of interest is the skill of the child at the final stage. Figure 4 shows the results.

Implementation details. To implement a mean-preserving spread, we are to keep the mean wage offer at baseline level. Recall that we have assumed the wage offer process is

$$\log w_{jt} = a_j + b_j t + \varepsilon_{jt},$$

with

$$\varepsilon_{jt} = \rho_j \varepsilon_{jt-1} + \sigma_{vj} v_{jt}.$$

So the distribution of log wage offer (conditional on t) is

$$\log w_j \sim \mathcal{N}(a_j + b_j t, \sigma_{\varepsilon j}),$$

where

$$\sigma_{\varepsilon j}^2 = \frac{\sigma_{vj}^2}{1-\rho_j^2}.$$

Since w_j exhibits a log-normal distribution, the mean of w_j is

$$\exp(a_j + b_j t + \frac{1}{2}\sigma_{\varepsilon j}^2) = \exp(a_j + b_j t + \frac{1}{2}\frac{\sigma_{vj}^2}{1-\rho_j^2}).$$

In all cases, then, we adjust the mean by $\frac{\sigma_{vj}^2(baseline) - \sigma_{vj}^2(experiment)}{2(1-\rho_j^2)}$ when we change σ_{vj}.

Further, to generate the initial correlation between wages and skills, we have assumed that the initial skill is drawn from a joint distribution that also depends on wage offers. In the exercises here, we are only changing the standard deviation of wage shocks, but the correlation of initial skill and wage offers are kept fixed. As a result, this exercise effectively also changes the distribution of initial skills.

Discussion. Figure 4 plots the distribution of skill changes, measured as a ratio of the standard deviation of skills in the baseline, to a larger (50% more compared to the baseline) father wage offer shock (blue), mother wage offer shock (red), and both (yellow). In all cases we increase the standard deviation of the wage *offer* distributions, as opposed to the the distribution of accepted wage offers, while we keep the mean wage offers fixed at their baseline level. Whether the standard deviations of the accepted wage distributions rise, fall, or stay the same as in the baseline, depends on parents' labor supply choices, partly reflecting a higher option value of working.

Three main results stand out. First, in all three cases the distribution of children skills moves to the left, implying a lower average level of skills. Hence, heightened wage risk is detrimental for children's skill accumulation, permanently reducing skill levels. Second, increasing the wage offer risk of fathers has a larger negative effect on skills than increasing the wage offer risk of mothers. Moreover, the distribution of skill changes is wider for fathers than mothers, implying that heightened wage risk for fathers leads to more skill inequality than heightened wage risk for mothers. Hence, when fathers are hit by risk shocks mean skill levels fall and at the same time children's skills become more unequal relative to when the mothers are hit by the same shock. Third, when the shock hits both parents, we see a sizable fall in average skill levels combined with an increase in skill dispersion.

Figure 5 shows the impact of increasing and decreasing wage offer risk on the mean skill levels relative to the baseline by different shock size. Children mean skill levels are substantially more sensitive to changes in the wage offer risk of the father than the mother. Perhaps not surprisingly, the negative effect of increased risk on skills is the largest when we hit the wage offer process of both parents. Importantly, the combined effect is larger than the sum of mother and father.

Figure 6 zooms in on the transmission mechanisms. We begin by discussing the implications for income and expenditures; see figures in panel (a). In all cases, after-tax income falls as we increase the wage offer risk of the parents. Such a fall in disposable income is associated with a sizable reduction in expenditures on children, whereas parental consumption remains virtually unchanged. Parents facing increased wage risk respond by smoothing their own consumption at the expense of child-related expenditures. To partly offset this reduction in child-related expenditures, parents increase child time; see figures in panels (b) and (c).

Figure 4: Distributions of Skill Changes to a Mean-Preserving Spread in the Wage Offer Shocks

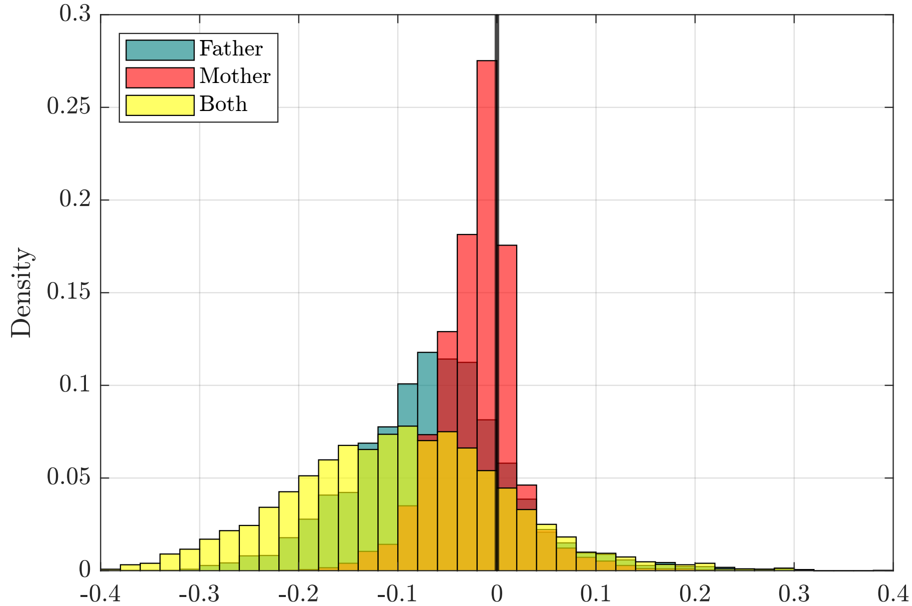

Notes: This figure plots the distribution of skill changes, measured as a ratio of the standard deviation of skills in the baseline, to a larger (50% more compared to the baseline) father wage offer shock (blue), mother wage offer shock (red), and both (yellow).

Figure 5: The Mean Effect of Wage Shock Dispersion

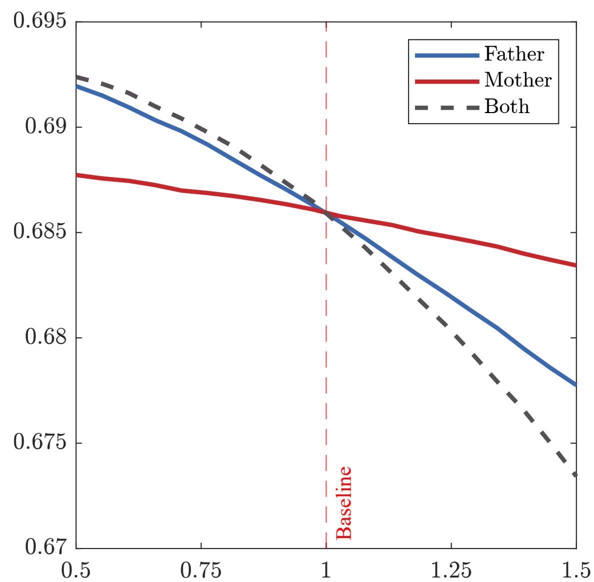

Notes: This figure plots mean skill levels on the y-axis and the size of the change in the standard deviation of the wage offer distribution on the x-axis.

Figure 6: Inspecting the Transmission Mechanisms of Increased Wage Risk

(a) After-Tax Income and Expenditures

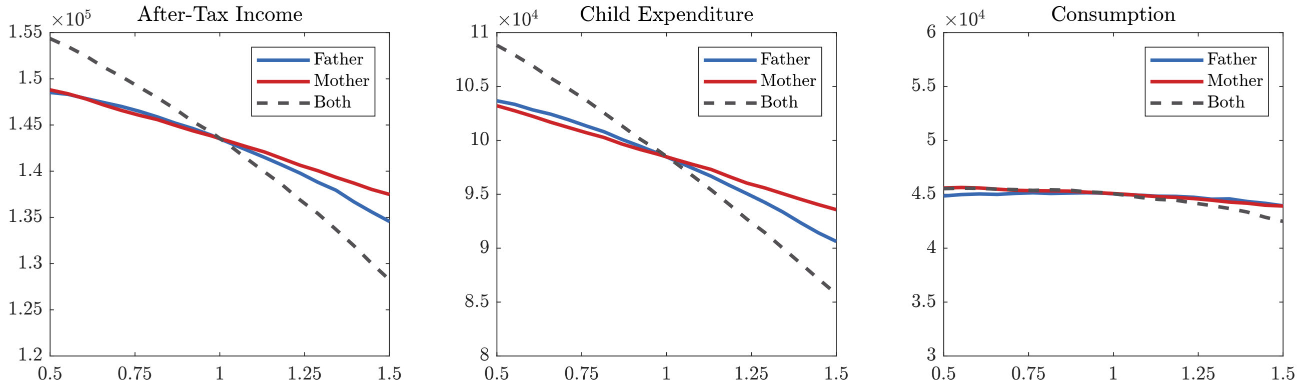

(b) Father's Time Allocation

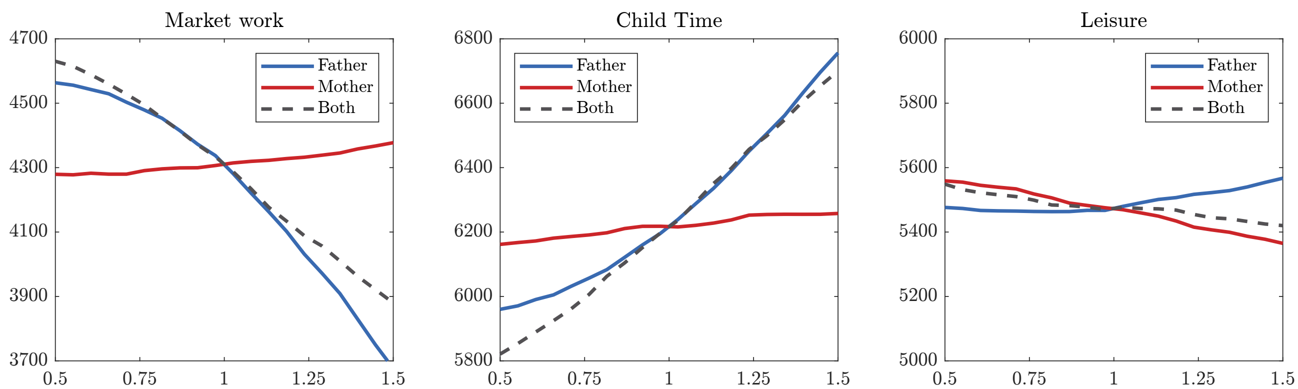

(c) Mother's Time Allocation

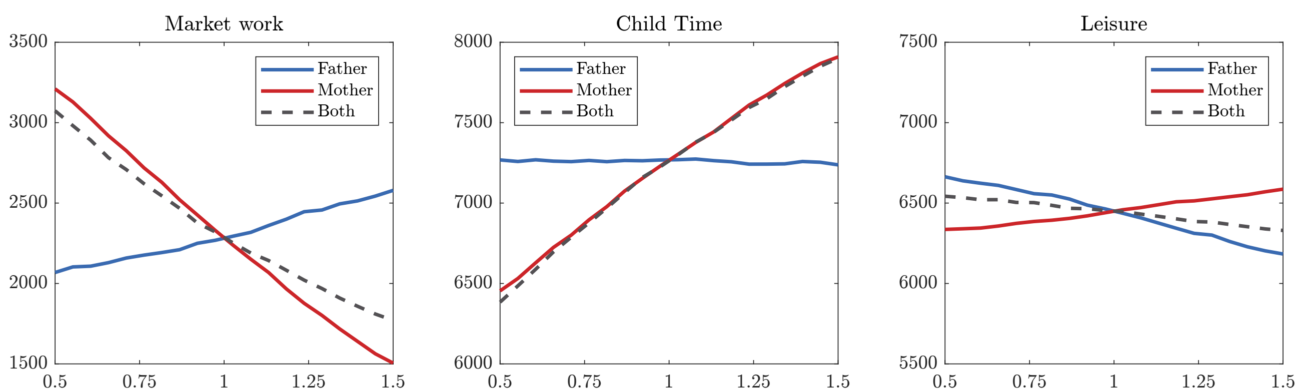

Notes: This figure shows the variable of interest on the y-axis and the size of the change in the standard deviation of the wage offer distributions on the x-axis.

6.2 Can the Social Safety Net Mitigate Children Skill Losses?

We now evaluate the extent to which progressive taxes and Universal Basic Income (UBI) can mitigate the adverse effect of income risk on children's skill accumulation. We begin with an experiment involving an increase in the progressivity of the tax system, then turn to Andrew Yang's UBI proposal, which consists of a $1,000 per month lump-sum transfer to every American adult over the age of 18.

6.2.1 Toward a More Progressive Tax System

Let us think of the following thought experiment. Overnight the government changes the progressivity of the tax-and-transfer system toward more progressive taxes. Is the magnitude of the wage risk pass-through on children's skill levels larger or smaller relative to the current, less progressive system in place? To answer this question, we capitalize on the simple tax function in (19), which parsimoniously describes the complex US system of taxes and transfers with three parameters (χ_1, χ_2, and χ_3).

We are interested in comparing two tax systems that differ in terms of progressivity but have the same overall tax rate level. To implement this exercise, we proceed as follows. We increase the progressivity parameter χ_2 from its benchmark level 0.118, to twice its level, 0.236, while readjusting the level parameter χ_1 so that the effective tax rate for the median household is unchanged. In the baseline economy, the median household (in terms of the average annual income) has an effective tax rate of 16.2%. Figure 9 shows the relationship between pre- and post-tax/transfer income in the baseline (solid line) and after increasing tax progressivity (dashed line). As apparent from the figure, moving toward a more progressive tax-and-transfer system tilts the pre- and post-tax income relationship by giving more post-tax income at the lower end of the pre-tax income distribution and less at the higher end of the pre-tax income distribution.

To begin, we study how the tax progressivity change impacts children's skill levels. Figure 7 shows children's final skill changes compared to the baseline after the increase in tax progressivity. Each dot represents a simulated household. The x-axis plots the log average annual post-tax income at the household level, and the y-axis plots the change in the child final skill in unit of the skills' standard deviation in the baseline. First, the change in progressivity has heterogeneous effect on skills: a negative effect on skills for high-income, a positive one for low-income. Second, quantitatively, at the lower end of the income distribution the positive effect ranges between 5 and 10 percent of the standard deviation of children's skills in baseline. At the higher end of the income distribution, the negative effect is sizable, ranging from 10 to 20 percent of the skills' standard deviation.

Figure 7: Skill Changes from Increasing Tax Progressivity
(Skill Change vs. Income Level)

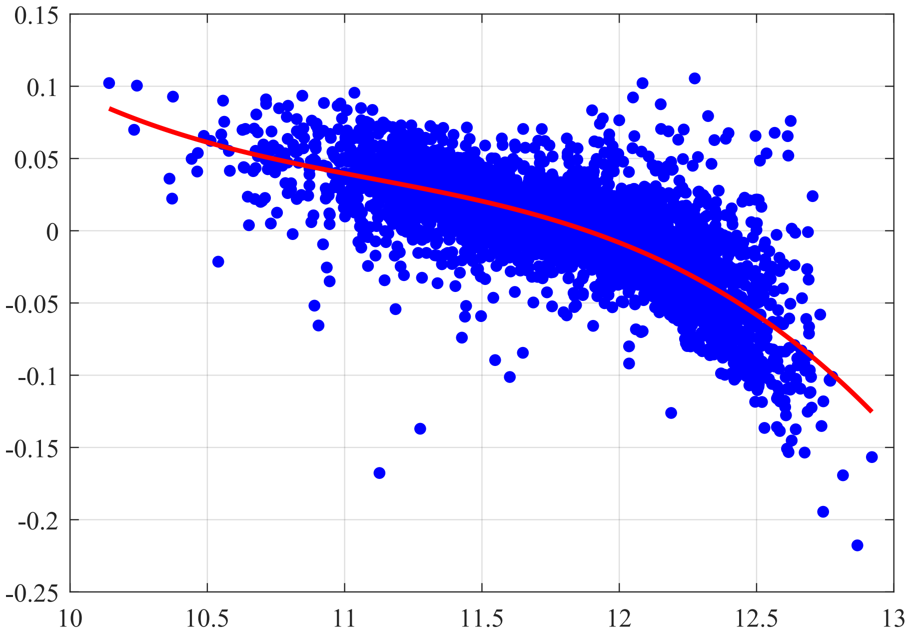

Notes: This figure plots final skill changes compared to the baseline after increasing tax progressivity. Each dot represents a simulated household. The x-axis plots the log average annual post-tax income at the household level, and the y-axis plots the change in the child final skill measured in the unit of the standard deviation in the baseline.

Figure 8b shows the distribution of skill changes for a 50% increase in the standard deviation of the wage offer shock for mother, father, and both, under a more progressive tax system. Figure 8a reproduces the results in Figure 4 to ease comparison. Overall, enhanced tax progressivity reduces the pass through of increased wage risk across the board. The distributions all shift to the left, implying skill losses on average, but by a lesser amount than in the baseline, less progressive tax-and-transfer system. Moreover, increased progressivity attenuates the impact of increased wage risk on children's skill inequality, as the distributions are visibly more concentrated than in the baseline.

Figure 8: Distributions of Skill Changes to a Mean-Preserving Spread in the Wage Offer Shock with a More Progressive Tax-and-Transfer System

(a) Baseline (Reproduced)

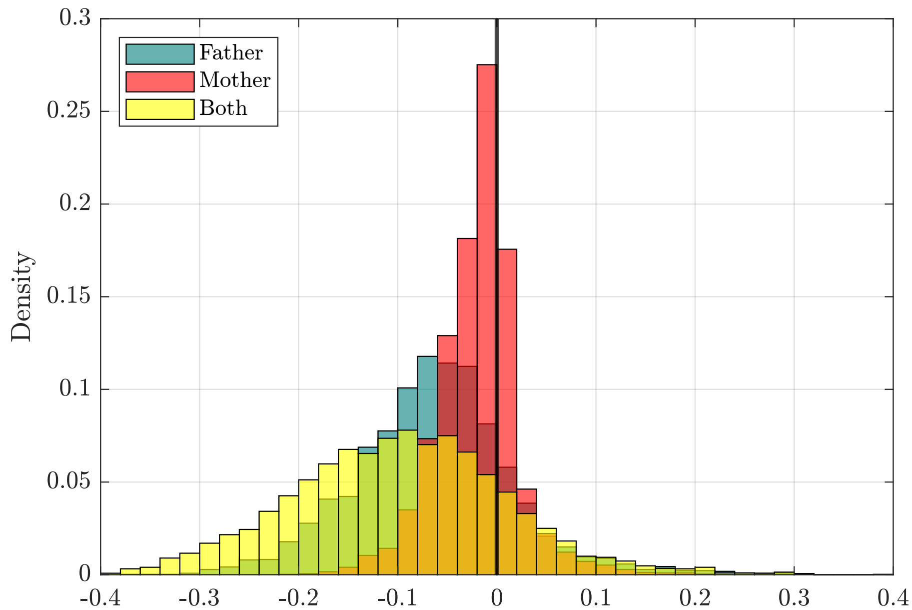

(b) Progressivity Tax Experiment

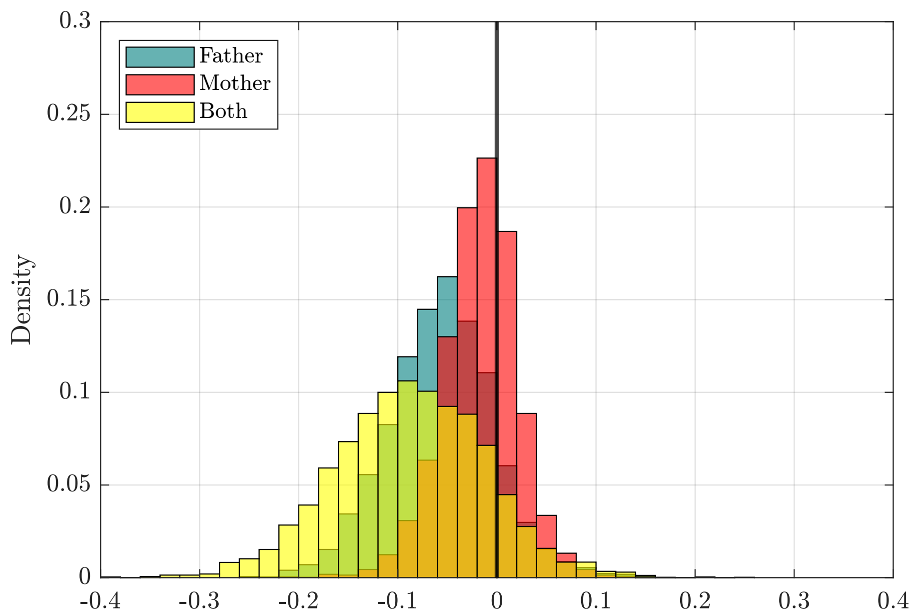

Notes: This figure plots the distribution of skill changes, measured as a ratio of the standard deviation of skills in the baseline, to a larger (50% more compared to the baseline) father wage offer shock (blue), mother wage offer shock (red), and both (yellow).

6.2.2 Universal Basic Income

We implement Andrew Yang's proposal of a universal basic income of $1,000 per month for every American adult over the age of 18. This amount to $24,000 a year for a two-parent household in our economy, or $48,000 for two years (one model period). We introduce UBI while keeping the tax system unchanged. That is, we set χ_3 such that even if a household has no labor income it can get such amount of UBI, while keeping χ_1 and χ_2 unchanged. Figure 9 shows that the introduction of UBI shifts upward the relationship between pre- and post-tax income. (We note that the UBI is unfunded and we do not adjust taxes to finance it.)

Figure 9: Comparing Tax Systems

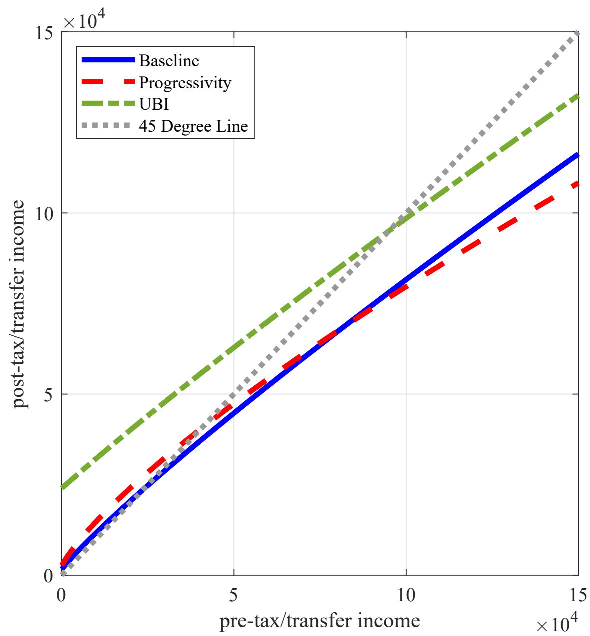

Notes: This figure plots the tax functions in the baseline model that is estimated from data (blue), in the experiment of a more progressive tax system considered in Section 6.2.1 (red), and in the UBI experiment considered in Section 6.2.2 (green), together with a 45 degree line for reference (gray).

Figure 10 plots final skill changes compared to the baseline after the introduction of UBI. Each dot represents a simulated household. The x-axis plots the log average annual post-tax income at the household level, and the y-axis plots the change in the child final skill measured in the unit of the standard deviation in the baseline. In contrast to the results from the tax progressivity experiment in Figure 7, the introduction of UBI raises children skill levels across the income distribution. Such positive effect is larger for lower income households. We stress that such positive effects on children skills are not hard-wired in the model. In fact, children could possibly suffer from the UBI. For example, parents who reduce their labor supply, to work (discretely) less and earn a smaller amount in total could be forced to reduce child-related expenditures.

Figure 10: Skill Changes from Introducing UBI
(Skill Change vs. Income Level)

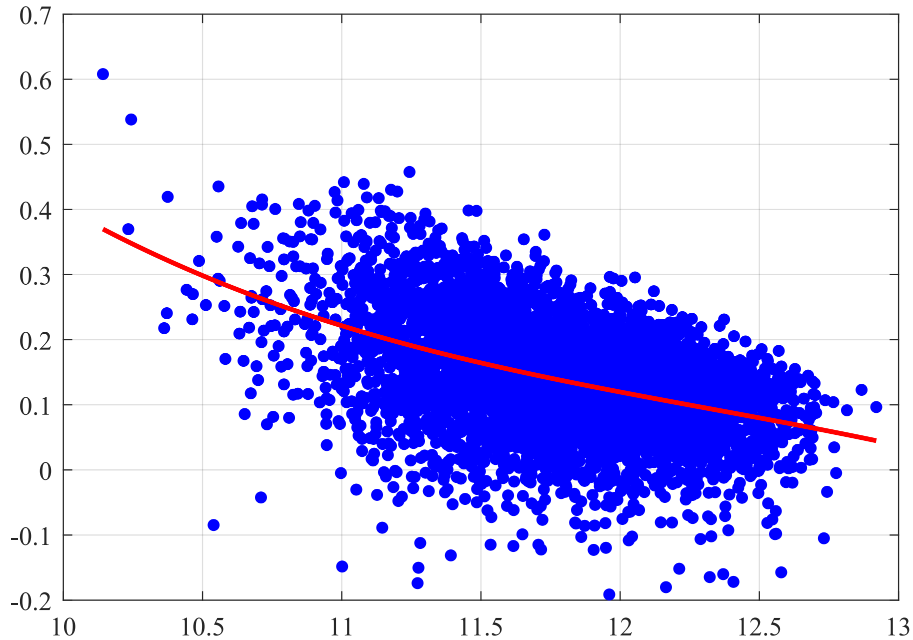

Notes: This figure plots final skill changes compared to the baseline after the introduction of UBI. Each dot represents a simulated household. The x-axis plots the log average annual post-tax income at the household level, and the y-axis plots the change in the child final skill level measured in the unit of the standard deviation in the baseline.

7 Conclusion

This paper aims to quantify the mechanisms and the extent to which parental wage risk passes through to children's skill development. Through the lens of a quantitative model of labor supply in which both parents choose how many hours to work, time spent with children, and child-related expenditures, we find that income risk slows down skill accumulation, permanently lowering children skill levels. Hence, parents' wage risk has a scarring effect on children's skills. To the extent that making up for such skill losses during childhood is hard, as the available evidence suggests, income risk in the presence of imperfect credit markets can negatively impact the labor market prospects of future generations. Income risk affects relatively more low-income households with limited ability to adjust labor supply in the face of negative wage shocks. A more progressive tax-and-transfer system can attenuate children skill losses.

References

Abbott, Brant. 2022. “Incomplete markets and parental investments in children.” *Review of Economic Dynamics* 44:104–124.

Agostinelli, Francesco, Emilio Borghesan, and Giuseppe Sorrenti. 2021. “Welfare, Workfare and Labor Supply: A Unified Evaluation.” HCEO Working Paper Series 2020-083.

Agostinelli, Francesco and Giuseppe Sorrenti. 2022. “Money vs Time: Family Income, Maternal Labor Supply, and Child Development.” HCEO Working Paper Series 2018-017.

Agostinelli, Francesco and Matthew Wiswall. 2016. “Estimating the Technology of Children’s Skill Formation.” Working Paper 22442, National Bureau of Economic Research.

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. “Time Use during the Great Recession.” *American Economic Review* 103 (5):1664–96.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework.” *Econometrica* 85 (3):693–734.

Attanasio, Orazio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina. 2020. “Estimating the Production Function for Human Capital: Results from a Randomized Controlled Trial in Colombia.” *American Economic Review* 110 (1):48–85.

Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos. 2018. “Aggregating Elasticities: Intensive and Extensive Margins of Women’s Labor Supply.” *Econometrica* 86 (6):2049–2082.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2005. “Female Labor Supply as Insurance against Idiosyncratic Risk.” *Journal of the European Economic Association* 3 (2-3):755–764.

Attanasio, Orazio, Costas Meghir, and Emily Nix. 2020. “Human Capital Development and Parental Investment in India.” *The Review of Economic Studies* 87 (6):2511–2541.

Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler. 2019. “Disability Benefits, Consumption Insurance, and Household Labor Supply.” *American Economic Review* 109 (7):2613–54.

Bick, Alexander, Adam Blandin, and Richard Rogerson. 2022. “Hours and Wages*.” *The Quarterly Journal of Economics* 137 (3):1901–1962.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance.” *American Economic Review* 98 (5):1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106 (2):387–435.

———. 2018. "Children, Time allocation, and Consumption Insurance." *Journal of Political Economy* 126 (S1):S73–S115.

Caucutt, Elizabeth M and Lance Lochner. 2020. "Early and Late Human Capital Investments, Borrowing Constraints, and the Family." *Journal of Political Economy* 128 (3):1065–1147.

Caucutt, Elizabeth M, Lance Lochner, Joseph Mullins, and Youngmin Park. 2020. "Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments." Working Paper 27838, National Bureau of Economic Research.

Chiappori, Pierre-Andre and Maurizio Mazzocco. 2017. "Static and Intertemporal Household Decisions." *Journal of Economic Literature* 55 (3):985–1045.

Cunha, Flavio and James J. Heckman. 2007. "The Technology of Skill Formation." *American Economic Review* 97 (2):31–47.

Cunha, Flavio, James J. Heckman, and Susanne M. Schennach. 2010. "Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Econometrica* 78 (3):883–931.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *The Review of Economic Studies* 71 (3):743–768.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2019. "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare." *Journal of the European Economic Association* 18 (2):890–926.

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2014. "Household Choices and Child Development." *Review of Economic Studies* 81 (1):137–85.

Del Boca, Daniela, Christopher J Flinn, Ewout Verriest, and Matthew J Wiswall. 2019. "Actors in the Child Development Process." Working Paper 25596, National Bureau of Economic Research.

French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *Review of Economic Studies* 72 (2):395–427.

Guiso, Luigi, Tullio Jappelli, and Luigi Pistaferri. 2002. "An empirical analysis of earnings and employment risk." *Journal of Business & Economic Statistics* 20 (2):241–253.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review* 104 (7):2075–2126.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1):153–161.

Heckman, James J. and Stefano Mosso. 2014. "The Economics of Human Development and Social Mobility." *Annual Review of Economics* 6:689–733.

Hotz, V. Joseph and John K. Scholz. 2003. "The Earned Income Tax Credit." In *Means-Tested Transfer Programs in the United States*, edited by Robert A. Moffitt. Chicago: University of Chicago Press, 141–197.

Jappelli, Tullio and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annual Review of Economics* 2 (1):479–506.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review* 100 (4):1432–67.

Mullins, Joseph. 2022. "Designing Cash Transfers in the Presence of Children‘s Human Capital Formation." University of Minnesota.

Nichols, Austin and Jesse Rothstein. 2016. "The Earned Income Tax Credit." In *Economics of Means-Tested Transfer Programs in the United States, Volume I*, edited by Robert A. Moffit. University of Chicago Press, 137–218.