It's More Than Time and Money: How the Parent-Child Relationship Shapes Skill Formation

Jacob Hazen*

Click here for the updated version

November 11, 2026

Abstract

I study how the parent-child bond—the emotional connection between parent and child—shapes the development of children's socio-emotional skill. I develop and estimate a dynamic model of household decision-making where parents choose their labour supply and allocate time and money to their children. Motivated by child development research and the economic literature on skill formation, the model treats the parent-child bond as an investable asset that parents can strengthen. It also enters the child's skill production function and parents' utility directly. I find that parental investments of time and money raise skills almost entirely through the parent-child bond, with little direct effect once the bond is controlled for. I also find that mothers' active time is the most productive input into investment, followed by fathers' active time and then expenditures on the child. In counterfactual analysis, I show that policies that raise family income, have little effect on the child's skill, because parents spend the freed resources mostly on consumption and leisure. Instead, I show that policies that could strengthen or change how parents value the parent–child bond, such as a parenting program, are more effective than income transfers in improving child socio-emotional skill.

*McGill University, Department of Economics. Email: jacob.hazen@mcgill.ca.

This research was supported by the Social Sciences and Humanities Research Council of Canada (SSHRC) Doctoral Fellowship and the Canadian Research Data Centre Network (CRDCN) Emerging Scholars Award.

The analyses in this paper were conducted at the Quebec Inter-university Centre for Social Statistics (QICSS), part of the CRDCN. QICSS is funded by the SSHRC, the Canadian Institutes of Health Research, the Canada Foundation for Innovation, Statistics Canada, the Fonds de recherche du Québec, and contributing Quebec universities. The views expressed are those of the authors and do not necessarily reflect those of the CRDCN, QICSS, or their partners.

I thank Fernando Saltiel and Fabian Lange for guidance on this paper. I also thank Daniel Barczyk, Rui Castro, Kentaro Asai, Nicolas Gendron-Carrier, and Doan Hai Ma for their helpful comments and suggestions. I am also grateful for the comments and suggestions from participants in the McGill Labour Economics group, the Université de Montréal CIREQ group, CEA at Vancouver, the McGill CIREQ seminar, and the RES at Newcastle.

1 Introduction

Parents shape their children's development through the time and money they invest in them. A large literature in economics documents how these choices affect children's cognitive skills through well-understood inputs such as books, educational toys, and time spent on homework and instruction (Cunha and Heckman, 2007, 2008; Del Boca et al., 2014). It is much less clear how parental time and money translate into socio-emotional skills such as emotional regulation, prosocial behaviour, or self-control. Yet these skills matter just as much for adult outcomes and, in some contexts, earn higher returns than cognitive skills (Lindqvist and Vestman, 2011; Deming, 2017; Edin et al., 2022; Etheridge et al., 2026).

In this paper, I develop a dynamic model of intra-household allocation to examine how parents' time and money investments, together with the emotional connection between parent and child, shape children's socio-emotional skill. A novel feature of my approach is the use of rich Canadian data that explicitly measure this emotional connection. Such measures are rare in household surveys, which may explain why economists have largely abstracted from it, even though child development experts have long recognized its importance for children's socio-emotional development (Lempers et al., 1989; Stattin and Kerr, 2000; Sanders et al., 2014).

I refer to this emotional connection as the *parent-child bond*: a stock of emotional capital between parent and child.[1] It is built through repeated parenting interactions and behaviours such as nurturance, rejection, and monitoring (Lempers et al., 1989; Darling and Steinberg, 1993; Stattin and Kerr, 2000). Children see this as a sense of trust, security, and closeness with their parents (Bowlby, 1969; Steinberg, 2001). This parent-child bond is malleable and parents can build it, specifically through spending time with their children. Researchers have shown that the parent-child bond is responsive to household income, interventions, and policy changes (Conger et al., 1994; Baker et al., 2008; Conger et al., 2010; Thompson, 2014; Akee et al., 2018; Kim et al., 2018). It is an important and distinct input separate from time and money in the production of children's socio-emotional development (Lizzeri and Siniscalchi, 2008; Fiorini and Keane, 2014; Cobb-Clark et al., 2019; Sevim et al., 2024; Moroni et al., 2025; Barszczewski and Thomas, 2025). However, the economics literature has yet to understand how the parent-child bond

[1] Psychologists and economists use a range of labels for this connection, including attachment, parental warmth, parent–child interaction, parent–child relationship quality, and parent–child relationship. Most of these papers come from parent-reported interactions, while my paper uses children's self-reports. Child-reports describe the emotional connection as the child experiences it.

is formed and how it interacts with parental time and money investments.

To understand how parents' time and money investments shape the bond and their kids' socio-emotional skill, I develop a dynamic model of household decision-making that draws on child development research and the economic literature on skill formation. In each period of childhood, parents receive wage offers and face a tax-and-transfer system which mirrors the Canadian tax system. They choose how to allocate their time between work, leisure, and time with their child, and how to allocate their income between consumption and expenditures on the child. If they choose to work, they earn income that can be spent on consumption and child expenditures. If they choose leisure, they enjoy the time spent relaxing or socializing. They can also choose to spend time with their child. Time and money spent on the child are inputs into a household investment function. This investment is an input into the production of both the parent-child bond and the child's socio-emotional skill.

In my model, parents have two motives to invest. First, they derive per-period utility from the parent-child bond itself: they value a close relationship with their child for its own sake. Second, they get a payoff when their child's socio-emotional skill is higher at the end of childhood, because that skill shapes their child's adult life. The two are linked, because a stronger bond also makes skill investments more productive.

The model's key innovation is to make the parent-child bond an endogenous stock, accumulated through parental investment within the household's optimization problem. Drawing on findings in both child development and economics, the parent-child bond has three roles in my model: it is a target of parental investment, an input into the production of children's socio-emotional skill, and a direct source of utility for parents.

My model builds on the seminal paper by Del Boca et al. (2014), who analyze a similar household model. I differ from their approach in four ways. First, they focus on how time and money investments produce cognitive skill, whereas I study socio-emotional skill, which is less well understood. Second, they do not model the parent-child bond. I model it as an endogenous stock that parents build through investment within the household's optimization problem, to my knowledge the first paper to do so, and the rich data I use allow me to capture its dynamics. Third, their Cobb-Douglas skill technology restricts dynamic complementarity between investment and existing skill, a restriction that has been shown to matter for policy analysis (Agostinelli et al., 2024; Caucutt et al., 2025). I instead use a flexible production function in which the bond and socio-emotional skill can be complements. Fourth, parents in my model derive

utility from the bond itself, consistent with evidence that parents enjoy time with their children (Guryan et al., 2008).

I build a Simulated Method of Moments estimator that allows me to draw on all four Canadian datasets when estimating the model parameters. The primary dataset is the Canadian National Longitudinal Survey of Children and Youth (NLSCY), which provides rich, repeated information on the socio-emotional skill and parental bond of children aged 10 to 15 in two-parent households, together with parental labour supply, wages, and household demographic information. A key feature of this dataset is the measurement of the parent-child bond using a set of child-reported parenting scales created by child development experts. Together, 18 items capture three dimensions of parenting: parental nurturance, rejection, and knowledge. A second key feature is child-reported socio-emotional skill, measured with 29 items that capture four dimensions: prosocial behaviour, hyperactivity and inattention, anxiety and emotional disorder, and conduct problems. The NLSCY has recently been linked with the children's adulthood tax file (T1FF), which provides detailed information on the adult income of the children.

I complement the NLSCY with the Survey of Household Spending (SHS), which contains detailed measures of household expenditures, including child-related spending, and the General Social Survey (GSS) on Time Use, which provides time diaries on how parents allocate time across work, active time with children, and leisure. These two datasets are not linked to the NLSCY. The method of moments estimator allows me to combine the information from all three datasets to estimate the model parameters.

I begin by documenting some stylized facts. (1) I show that socio-emotional skill matters for adult income, about as much as math skill. A one standard deviation increase at ages 14 to 15 predicts 8.3 percent higher adult income, compared with 9.7 percent for math skill. The parent-child bond also predicts adult income on its own, but its predictive power disappears once skill and math are included, so it seems to work through skill. (2) I show that the parent-child bond is an important and distinct input into the production of socio-emotional skill. Investment produces a stronger bond, and a stronger bond produces more socio-emotional skill. Investment has little direct impact on future skill once the bond is accounted for.

These patterns cannot tell me how parents make choices, so I turn to the structural model. Mothers' active child time is the most productive input into investment, followed by fathers' active time and then expenditures on the child. I find that about one-sixth of parental investment is motivated by the direct utility of the parent-child bond, and the remaining five-sixths by the payoff from raising the child's future

skill. When parents place more value on the payoff of future skill, mothers spend more time with their children, and about 70 percent of each additional hour is taken from leisure and 30 percent from work. Fathers take the extra time entirely from leisure and leave their work hours unchanged. Fathers do reduce their work hours, however, in response to positive income shocks. In terms of the technology of bond and skill formation, I find that (1) the bond and skill are persistent, (2) the bond is more cross-productive for skill than skill is for the bond, and (3) the bond and skill are complements in the production of both stocks.

Having established that the model replicates key facts of the data, I use it to conduct several experiments to establish the importance of modelling the parent-child bond to understand its effect on children's skill. I compare two kinds of experiments that a policy maker could implement to target children's socio-emotional skill: (1) changes to the household's budget constraints through the tax-and-transfer system and (2) parenting programs that could increase children's socio-emotional skill.

I find that income transfers raise children's outcomes only slightly. An income transfer of 24,000 CAD raises terminal skill by 0.055 standard deviations and predicted adult income by 0.96 percent. Transfers are weak for two reasons. First, parents spend most of the freed resources on their own consumption and leisure instead of money and time investments. Second, time investments do most of the work in the production of child investments.

A well-known example of a parenting program that targets the parent-child bond with little-to-no direct impact on the household's time and budget constraints is Triple P.[2] It trains parents in consistent discipline, warmth, and realistic expectations, and programs of this kind can be delivered through group classes, home visits, or family counselling.[3] Parents receive no money and spend only a small portion of their leisure time on the program. Therefore, the program has only minimal impact on parents' time constraints.

I use my model to explore two mechanisms through which such a parenting program could work. It

[2]Triple P is a parenting and family support system developed at the University of Queensland, Australia, with five levels of intensity ranging from information campaigns to intensive family intervention. It is delivered to parents of children up to age 12, with a Teen Triple P version for ages 12–16. Triple P is used in 72 countries in 22 languages. Sanders et al. (2014) review and meta-analyze the evidence on the Triple P program.

[3]Kim et al. (2018) find causal evidence, through a randomized intervention, that a group version of the Triple P program affects children's socio-emotional skill. Parents attend four weekly two-hour sessions and children do not take part, and the program changes parenting for at least four years and lowers externalizing behaviour in adolescence.

could raise the bond directly, change parents' preferences for the bond or for their child's future skill, or do both. The first mechanism works through the technology of skill formation. Raising the bond by one standard deviation at ages 10–11 raises terminal skill by 0.425 standard deviations and adult income by 7.73 percent. In comparison, increasing socio-emotional skill by one standard deviation at ages 10–11 does as much for terminal skill and adult income. In other words, improving the parent-child relationship is about as effective as targeting children's socio-emotional skill itself.

The second mechanism works through parents' preferences. A parenting program could shift two preference weights in the model: the weight parents place on the child's future skill, the weight parents place on the parent-child bond, or both. I find that (1) raising the preference weight parents place on the children's future skill by 50 percent raises terminal skill and the bond by 0.155 and 0.143 standard deviations. The response is also concave. Cutting the weight by 50 percent lowers them by 0.262 and 0.242, meaning that caring less about your children's skill hurts more than the gain from caring more. (2) Raising the weight on the relationship by 50 percent raises terminal skill by 0.050 standard deviations. Together, these findings imply that the bond, rather than investment itself, is at the centre of the technology of socio-emotional skill formation.

This paper makes multiple contributions to the literature on child development. First, it bridges the literatures from child development and psychology with the economics literature on household decision-making and child skill formation. It does so by modelling the parent-child bond as an endogenous stock within a household optimization problem, and by using child-reported measures of the bond and socio-emotional skill to estimate the model. This allows me to study the intra-household allocation of time and money investments, together with the parent-child bond, shape children's socio-emotional skill formation.

It contributes to the literature estimating the technology of skill formation (Cunha and Heckman, 2008; Cunha et al., 2010; Attanasio et al., 2020a,b; Bolt et al., 2024; Wiswall and Agostinelli, 2024), which has focused primarily on cognitive skills. I extend this framework to socio-emotional skill by introducing the parental bond. Previous evidence shows that the bond is a production input distinct from parental investment and that it is malleable and responds to household resources, so I allow it to have its own production function. I show that the bond is a distinct input into the production of socio-emotional skill and that constructing latent investment with the bond and investment together, as in Cunha et al. (2010) and Toppeta (2022), produces a significant coefficient that overstates the effect of investment on skill.

Instead, I find that investment builds future bond, which in turn builds skill.

This paper also contributes to the literature on household decision-making and child development (Del Boca et al., 2014; Blundell et al., 2018; Agostinelli et al., 2024). While most of this literature has focused on parental investments of time and money, I model the parent-child bond as an endogenous outcome of household optimization. This allows me to examine how economic conditions (e.g., a higher benefit floor or a change in the tax schedule) and policy interventions (e.g., Triple P) shape the parent-child bond and, through it, children's socio-emotional development. This type of policy analysis on the parent-child bond and socio-emotional skill has been understudied in the literature.

Related Literature

This paper relates to a recent literature on the economics of parenting, which studies how *parenting style* shapes children's human capital formation, starting from the typology of authoritarian, authoritative, and permissive parenting in Baumrind (1966, 1967). Doepke and Zilibotti (2017) classify parents along this typology as paternalistic, acting in the child's best interest, or altruistic, letting the child maximize their own utility, and derive each style as an equilibrium outcome of that objective.[4] A limitation of this approach is that parents are sorted into discrete types, often using arbitrary cutoffs that are imposed by the researcher with no theoretical grounding.[5]

The literature in child development has recognized that Baumrind's typology bundles different components of parenting behaviours that carry distinct, and sometimes opposing, associations with children's outcomes (Darling and Steinberg, 1993). Barszczewski and Thomas (2025) decompose parenting behaviours into warmth, reasoning, hostility, and consistency, and map them onto the typology of Baumrind (1966, 1967).[6] They show that the negative association of authoritarian parenting with socio-emotional skills is driven by hostility rather than by the high consistency also present in that style. The positive asso-

[4] See Doepke et al. (2019) and Agostinelli et al. (2025b). The same authoritarian-or-permissive classification reappears in Agostinelli et al. (2025a), applied to a setting with neighbourhood and peer effects.

[5] For example, in Bruno (2025), caregivers are split into authoritarian or authoritative according to whether a standardized warmth scale falls above or below zero, so two parents with nearly identical warmth are assigned different styles depending only on which side of the mean they fall on. Similarly, Agostinelli et al. (2025a) classify a parent as authoritarian if the child answers "No" to a single survey question asking whether their parents let them make their own decisions about who they spend time with, and as non-authoritarian if the child answers "Yes," so that one binary response fully determines the parent's style.

[6] See Appendix A for the analogous mapping used here.

ciation of authoritative parenting is driven mostly by the low hostility and high consistency. This suggests that this typology, which bundles behaviours with opposing associations into a single classification, may mix components of parenting that drive children's skills and outcomes.

To keep the model tractable and to avoid imposing discrete parenting classifications, with the arbitrary cutoffs needed to define them, I measure the parent-child bond along a single latent continuous dimension, using a set of theoretically grounded, child-reported parenting scales that capture parental warmth, rejection, and knowledge (Lempers et al., 1989; Stattin and Kerr, 2000). I use the term *parental bond* throughout this paper rather than *parenting style* to avoid the association with discrete types.

This measurement choice aligns with the evidence in both child development and economics that the parent-child bond is not a fixed trait and is responsive to household resources, environment, interventions, and policy (Conger et al., 1994, 2010; Baker et al., 2008; Akee et al., 2018; Kim et al., 2018). It is also consistent with treating parenting behaviours as distinct from material and time investment in the production of children's skills (Fiorini and Keane, 2014).

This work also relates to the literature on household decision-making and child skill formation. Del Boca et al. (2014) analyze a cognitive skill production function within a dynamic household optimization model. Parents choose between childcare time and child expenditures under budget and time constraints. They find that both parents' time inputs are important for children's cognitive development, particularly when the child is young, whereas money expenditures are less productive. They perform comparative statics exercises to show that cash transfers to households with children have small effects on child quality. This is because money investments are not very productive in producing these skills, and a sizeable fraction of the transfer goes to other household consumption and parental leisure.

A related literature studies how tax-and-transfer policies can affect children's skills. Blundell et al. (2018) study a life-cycle model in which couples face wage uncertainty, a progressive tax system, and liquidity constraints. They show that the presence of young children changes how responsive spouses' labour supply is to wage shocks, particularly for mothers. They find that an employment subsidy for mothers with young children raises their hours of work but lowers their time with their children by far more than fathers' time rises, so that family childcare hours decline. Agostinelli et al. (2024) further incorporate idiosyncratic income risk and intra-household insurance into a cognitive skill formation model. They show that cognitive skill formation depends not only on household preferences and constraints but also on the household's capacity to insure itself against economic risk. In both of these models, economic

conditions and policy reach the child only through parental time and money. In my paper, I demonstrate how a tax-and-transfer policy may affect the parent-child bond and the child's socio-emotional skill formation.

Section 2 describes the four Canadian datasets used in the analysis. Section 3 presents empirical evidence motivating the model. The model itself is developed in Section 4, followed by the estimation procedure in Section 5. Section 6 presents the structural estimates, and Section 7 first evaluates tax-and-transfer policies and then parenting program policies. Section 8 concludes.

2 Data and Descriptive Statistics

The data I use for this analysis comes from four main sources. The primary dataset is the National Longitudinal Survey of Children and Youth (NLSCY), linked with the Family Tax Files (T1FF). I complement this with the General Social Survey (GSS) on Time Use and the Survey of Household Spending (SHS), which are not linked to the NLSCY and are used separately to construct moments on parental time and expenditure. I use the NLSCY to construct a longitudinal dataset that tracks children from birth to adulthood, allowing me to observe their socio-emotional skill and the parent-child bond over time. It also provides information on wages and hours worked for both parents, and on household income. The GSS provides time diaries on parental time allocation, and the SHS provides household expenditures.

2.1 National Longitudinal Survey of Children and Youth

The NLSCY is a biennial survey that began in 1994 and continued until 2008–2009. It provides both longitudinal and cross-sectional coverage.[7] For my analysis, I use the longitudinal component, which follows a single cohort of children from birth to age 25.[8] This cohort entered the survey at ages 0 to 11 in Cycle 1 (1994/95) and was resurveyed every two years thereafter, so the children in my analysis sample were born between approximately 1983 and 1994/95. I restrict the sample to children aged 10 to 15, as this is the age range in which both socio-emotional skill measures and parental bond measures are self-reported by the child.

[7] The cross-sectional component includes children aged 0–5, with new cohorts added in each cycle.

[8] Appendix A provides details on survey administration, the T1FF linkage, and attrition.

Table 1: Data Sources and Availability

Data source	Time period	Type	Key variables
NLSCY + T1FF	NLSCY: 1994–2009 (biennial) T1FF: 1993–2019	Longitudinal	Child socio-emotional skill; parental bond measures; family structure, income, education; parental earnings & hours, children adult earnings
GSS (Time Use)	2005 (Cycle 19)	Cross-sectional	Time diaries (market work, active time, domestic work); leisure, personal care
SHS	1997–2009 (annual)	Cross-sectional	Household expenditures; child-related (clothing, childcare, education, toys); food, domestic help

Notes: NLSCY is the National Longitudinal Survey of Children and Youth, T1FF the Family Tax Files, GSS the General Social Survey, and SHS the Survey of Household Spending.

Sample Restrictions of the NLSCY The restriction to children aged 10 to 15 is motivated by two considerations. First, the survey stops collecting measures of socio-emotional skill and parental bond after age 15, and children under age 9 do not self-report socio-emotional skill or parental bond measures. As discussed in Del Bono et al. (2024), parental reports of children's skills are subject to systematic bias, which I expect may also affect parental reports of the parental bond.[9] Recently, Barszczewski and Thomas (2025) find that parents who exhibit more hostile behaviour and employ less consistent disciplinary strategies are also more likely to underreport their children's skills relative to teachers' assessments. Second, since the parent who responds to the survey could be either the mother or the father of the child, parental reports may reflect the respondent's own relationship with the child rather than the child's overall experience.[10]

I further restrict the sample to children living in two-parent households with at least one biological parent present. Both parents must have recorded information on hours worked, employment status, and income. Because the survey does not collect information on the working hours of the primary respondent's spouse in the first wave, I exclude this wave from the analysis. I also exclude observations with missing data on

[9]Even conditional on their own socio-emotional skills, parents misreport their child's skills. For example, a parent who rates herself high on grit tends to rate her child similarly, whether or not that reflects the child's true level.

[10]Restricting to mother-PMK households removes one reporting inconsistency, but the father's relationship with the child is still unobserved, given that he is never the direct respondent.

children's socio-emotional skill and parental bond measures.[11]

Finally, because interviews are conducted throughout the year, a child may be, for example, 10 years old at the start of the year and turn 11 by the time of the interview. To account for this age variation, I group children into three age bins, 10–11, 12–13, and 14–15. I restrict the sample to children with skill and bond measures in at least two age bins. After applying these restrictions, the final sample consists of 3,190 children observed over two or three waves.[12]

Parental Labour Market Measures The NLSCY asks parents to report total wage and salary earnings over the past 12 months, before taxes. Since the model uses a weekly time constraint, I need wages and hours measured on a consistent basis. Labour supply is recorded in two categorical variables, usual weekly hours (0; under 9; 10–19; 20–29; 30–39; 40; over 40) and weeks worked over the past year (none; 1–25; 26–47; at least 48).[13] I assign midpoints to convert these to continuous values. These are 4.5 hours for "1–9," 14.5 for "10–19," 24.5 for "20–29," 34.5 for "30–39," 40 for "40," and 50 for "over 40." For weeks, the assignments are 0, 13, 37, and 50, respectively.

The hourly wage is annual earnings divided by annual hours, where annual hours are the product of usual weekly hours and weeks worked. Using weekly hours alone would assume 52 weeks of work, overstating wages for parents employed only part of the year. I deflate wages to real 2001 Canadian dollars using the Canadian Consumer Price Index (CPI) from Statistics Canada.

Panel A of Table 2 reports unconditional means and standard deviations of parents' hourly wages, weekly hours worked, age, and education. Panel B reports means and standard deviations of hourly wages and weekly hours conditional on working, by child age bin.

Child Socio-Emotional Skill Measures Children aged 10 to 15 years completed socio-emotional measures in the self-complete portion of the survey. The measures capture behavioural and emotional indicators, including conduct disorder, hyperactivity, inattention, physical aggression, emotional disorder, anxiety, and prosocial behaviour, and are drawn from the Montreal Longitudinal Survey. These measures

[11] If the child has missing data for one wave, for example observed at ages 10–11 and 12–13 but not 14–15, or at ages 12–13 and 14–15 but not 10–11, the child is nonetheless kept in the sample. I "fill in" the missing wave using the data from the child's other waves, which the model's intertemporal dependencies make possible, a procedure analogous to performing Monte Carlo integration to "integrate out" the missing data. This is also done in Del Boca et al. (2014).

[12] Rounded to comply with the data provider's confidentiality requirements.

[13] This approach follows Molnar (2023), who applies the same midpoint assignment to labour supply categories.

Table 2: Parental Characteristics: Means and Standard Deviations

	Mothers		Fathers	
	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Unconditional				
Hourly Wage	18.79	10.86	23.40	10.54
Weekly Hours	26.97	16.01	41.14	13.34
Age	40.73	4.97	42.95	5.70
Education (years)	13.92	2.11	13.77	2.40
Panel B: Conditional on Working, by Child Age Bin				
Hourly Wage				
Ages 10–11	18.21	10.77	22.60	10.28
Ages 12–13	18.49	10.71	23.24	10.45
Ages 14–15	19.60	11.08	24.30	11.01
Weekly Hours				
Ages 10–11	30.40	12.79	42.21	11.29
Ages 12–13	31.01	13.08	42.77	10.68
Ages 14–15	32.11	12.43	42.84	11.05

Observations (rounded): 9,580 in Panel A and 3,190 per age bin in Panel B.

Rounded to comply with the data provider's confidentiality requirements.

assess emotional intelligence and are on a three-point scale from "Never or not true", "Sometimes or somewhat true", and "Often or very true".[14] Measures that have a negative connotation are reverse-coded so that a higher score indicates better socio-emotional skill. I use the scales to construct a latent factor for socio-emotional skill in each age bin, estimated as a Bartlett score from the measurement system in Appendix C. The full set of measures used is in Table 3, which lists the individual scale items.

Parent-child bond measures I measure the parent-child bond with the *Parenting Scale*, a set of child-reported measures in the NLSCY administered to children aged 10 to 15. The scale has three components, namely parental nurturance, parental rejection, and parental knowledge (Schaefer, 1965; Roberts et al., 1984; Lempers et al., 1989). As like the socio-emotional measures, each measure is on a three-point scale from "Never or not true", "Sometimes or somewhat true", and "Often or very true". Measures that have a negative connotation are reverse-coded so that a higher score indicates a better parent-child bond. I use the scales to construct a latent factor for parent-child bond in each age bin, estimated as a Bartlett score from the measurement system in Appendix C. Table 4 lists the items for each component.[15]

Nurturance reflects parental warmth, support, and involvement. It is associated with lower depression and loneliness in children. Rejection reflects inconsistent, punitive, and mood-dependent discipline, and is associated with greater depression, delinquency, and drug use (Lempers et al., 1989).

Parental knowledge captures how closely parents track their child's whereabouts, activities, and misbehaviour. It is associated with less delinquency, norm-breaking, and police contact (Stattin and Kerr, 2000). I follow Stattin and Kerr (2000) and use the label parental knowledge rather than parental monitoring, the construct these items were originally designed to measure. They argue that monitoring measures mostly capture what parents know about their child. They find that this knowledge comes mainly from the child's voluntary disclosure rather than from the parents' own surveillance. Arim et al. (2011), who evaluate the factor structure of the child-reported parenting scales in the NLSCY, likewise recommend relabelling the monitoring scale as parental knowledge or monitoring efforts.

Parental knowledge can therefore be interpreted as a marker of a closer and more trusting relationship, in which the child chooses to share information with the parent. I include it alongside nurturance and

[14] They were developed by Dr. Reuven Bar-On and Dr. James D.A. Parker to capture the capacity to monitor and interpret one's own feelings and those of others across home, school, and work contexts.

[15] Nurturance, rejection, and parental knowledge can also be mapped onto the authoritarian, authoritative, permissive, and neglectful typology of Baumrind (1966, 1967); see Appendix A.

Table 3: Child Non-Cognitive Scale Items

Hyperactivity/Inattention Score (7 items)	
1	I can't sit still, I am restless.
2	I am easily distracted. I have trouble sticking to any activity.
3	I can't concentrate, I can't pay attention.
4	I am impulsive, I act without thinking.
5	I have difficulty waiting for my turn in games or group activities.
6	I cannot settle to anything for more than a few moments.
7	I am inattentive, I have difficulty paying attention to someone.
Pro-social Score (10 items)	
1	I show sympathy to someone who has made a mistake.
2	I try to help someone who has been hurt.
3	I offer to help clear up a mess someone else has made.
4	If there is an argument, I try to stop it.
5	I offer to help other young people who are having difficulty with a task.
6	I comfort another young person who is crying or upset.
7	I help to pick up things which another young person has dropped.
8	When I am playing with others, I invite bystanders to join in a game.
9	I help other people my age who are feeling sick.
10	I encourage other people my age who cannot do things as well as I can.
Anxiety and Emotional Disorder Score (6 items)	
1	I am unhappy or sad.
2	I am not as happy as other people my age.
3	I am too fearful or nervous.
4	I cry a lot.
5	I am nervous, highstrung or tense.
6	I have trouble enjoying myself.
Conduct Disorder/Physical Aggression Score (6 items)	
1	I get into many fights.
2	When another kid accidentally hurts me, I assume that he/she meant to do it and react with anger and fighting.
3	I physically attack people.
4	I threaten people.
5	I bully or am mean to others.
6	I kick or hit other people my age.

Table 4: Parenting Scale Items

Parental Nurturance (7 items)	*All items begin with "My parents…"*
1	smile at me.
2	praise me (say nice things about me).
3	make sure I know I am appreciated.
4	listen to my ideas and opinions.
5	and I solve a problem together when we disagree.
6	speak of the good things I do.
7	seem proud of the things I do.
Parental Rejection (7 items)	*All items begin with "My parents…"*
1	soon forget a rule they have made.
2	nag me about little things.
3	only keep rules when it suits them.
4	threaten punishment more often than they use it.
5	enforce rules depending on their mood.
6	hit me or threaten to do so.
7	get angry and yell at me.
Parental Knowledge (Parental Monitoring) (4 items)	*All items begin with "My parents…"*
1	want to know exactly where I am and what I am doing.
2	tell me what time to be home when I go out.
3	find out about misbehavior.
4	let me go out any evening I want.

rejection as a measure of the parent-child bond. In the age-invariant specification of the measurement model, parental knowledge loads positively onto the bond factor at 0.42, compared to 0.57 for rejection, which is reverse-coded so that a higher score means less rejection.[16]

Investment measures The NLSCY collects information on parental time use and household expenditures as proxies for investment. Parents report their involvement in school activities such as attending parent-teacher meetings, visiting classrooms, volunteering, and attending school events. I also have measures on how often they discuss school, and the child's friends and relationships, with their child. Parents also report how often they check their child's homework and how much time the child typically spends on it each day. This set of items captures a more direct time dimension. Expenditure proxies are based on children's reports of activities *outside school*, including team sports, instructor-led sports, music lessons, art lessons, and other hobbies that require a financial investment.

While the child reports *in-school* activities, I use *outside-school* activities as a proxy for parental investment expenditure. Outside-school activities typically require direct financial investment, whereas in-school activities could be free or heavily subsidized. While these measures proxy expenditures by the parent, they could also load onto a time investment. For example, a child who plays hockey outside of school has a parent who spends time driving them to and from practice, watching the games, and attending team meetings. This is often why it is difficult to separate time and expenditure investment into latent factors. One has to make a strong assumption that the time and expenditure measures do not cross-load onto each other. Table A.1 summarises the individual items underlying each investment measure.

2.2 Time-Use and Household Expenditures Data

Time-use data The General Social Survey (GSS) on Time Use is a cross-sectional survey conducted by Statistics Canada that collects detailed time diaries from a nationally representative sample of Canadians aged 15 and older, recording how respondents allocated their time across a 24-hour reference day, including work, leisure, and active time with children.[17] I construct three main categories of time use, namely market work, active parenting time, and residual time (leisure). Market work includes all paid employment activities (primary, secondary, and other jobs). Active time covers activities that involve

[16] See Table C.1, which reports the loadings for the latent factors across time.

[17] I use the 2005 cycle. While there is a 1998 GSS Time Use survey, it is not used because it does not collect information on respondents' spouses, which is required to construct household-level time-use measures.

engaging with children, such as helping with homework, transporting children, playing, and supervising. Residual time (leisure) consists of all non-market and non-active parenting time, including personal care, domestic work, education, and other leisure activities. I compute these for parents with children under age 14 as of the reference day.[18]

Table B.1 shows that mothers and fathers divide their time between market work and active parenting in offsetting ways. Fathers report nearly twice the weekly market hours of mothers (46.67 versus 24.09), while mothers report almost twice the active parenting time of fathers (15.16 versus 7.77 hours). Mothers also report roughly 15 more weekly hours of residual time than fathers (128.75 versus 113.56), so that a large part of fathers' additional market work is offset by this residual rather than by active parenting alone. Since the residual includes domestic work, this gap is likely to overstate any gap in leisure time. This is particularly true because mothers typically do more domestic work than fathers.

Figure B.1 breaks down active parenting time by couple employment type. The figure shows that both the level and the division of active time respond to each parent's employment status. Total active time is highest when the mother is not employed and the father works full-time, and lowest when both parents work full-time, as parents who work fewer market hours can provide more active time to children. Mothers account for 76 percent of active time when they are not employed and the father is full-time, falling to 66 percent when the mother works part-time and 61 percent when both parents work full-time. Fathers do the majority of active time (64 percent) only in the one case where the father, rather than the mother, is not employed. Mothers' time is more responsive to the employment status of the household than fathers' time.

Household expenditures The Survey of Household Spending (SHS, previously the Survey of Family Expenditures) is a cross-sectional survey conducted by Statistics Canada that records household spending from a nationally representative sample of Canadian households. I use the SHS from 1997 to 2009 to construct empirical moments on child-related expenditures. The survey records spending categories such as food, shelter, clothing, transportation, and recreation. I use the aggregated measure provided by Statistics Canada on child expenditures based on three categories, namely (i) clothing and footwear, (ii)

[18] In the GSS and SHS, children are grouped into age categories of 0–4, 5–14, and 15–18 years. In the NLSCY, however, the final category includes 14- and 15-year-olds because interviews were conducted at different times of the year. As a result, depending on the child's birth month, some respondents who were 14 at the start of the survey period had turned 15 by the time of the interview.

toys, and (iii) education, including tuition, school supplies, books, and extracurricular activities. I compute them for households with children under 14.[19] All expenditure values are in 2001 Canadian dollars, deflated using the CPI from Statistics Canada. Table B.2 reports the means and standard deviations, and Figure B.2 shows how this spending varies with household income and couple employment type.

Education holds the largest share of child-related expenditure. It averages $1,118.26 annually which is four times more than the average toy spending and over seven times more than the average spending on clothing and footwear. The total of $2,543.04 also includes childcare and other spending, which are not broken out separately here. Households in the sample report mean income of $85,594.57 and average maternal education of 14.22 years. The standard deviation of household income is more than half its mean, so the households differ widely in the resources available for child-related spending. This reflects in the large standard deviations of child-related expenditure across households. This could be due to differences in the number of children they have, their ages, income, and the types of activities they choose to invest in.

Figure B.2 plots mean child-related expenditure against household income decile, broken down by couple employment type. I consider only three employment types, due to the data provider's minimum cell-size requirement for disclosure. The categories are (i) mother not employed, (ii) mother employed part-time, and (iii) mother employed full-time, with the father working full-time in all three categories. Child-related spending rises with household income for all three couple employment types. Households with a mother employed full-time spend slightly more on average throughout the income distribution than households with a mother employed part-time. Households with a non-employed mother spend the least out of the three groups. Paired with Figure B.1, these patterns could suggest a substitution effect, where a non-employed mother substitutes time for money in child investment, while a full-time employed mother substitutes money for time. However, the differences are small, the confidence bands overlap across the three groups, and the figures are just descriptive patterns.

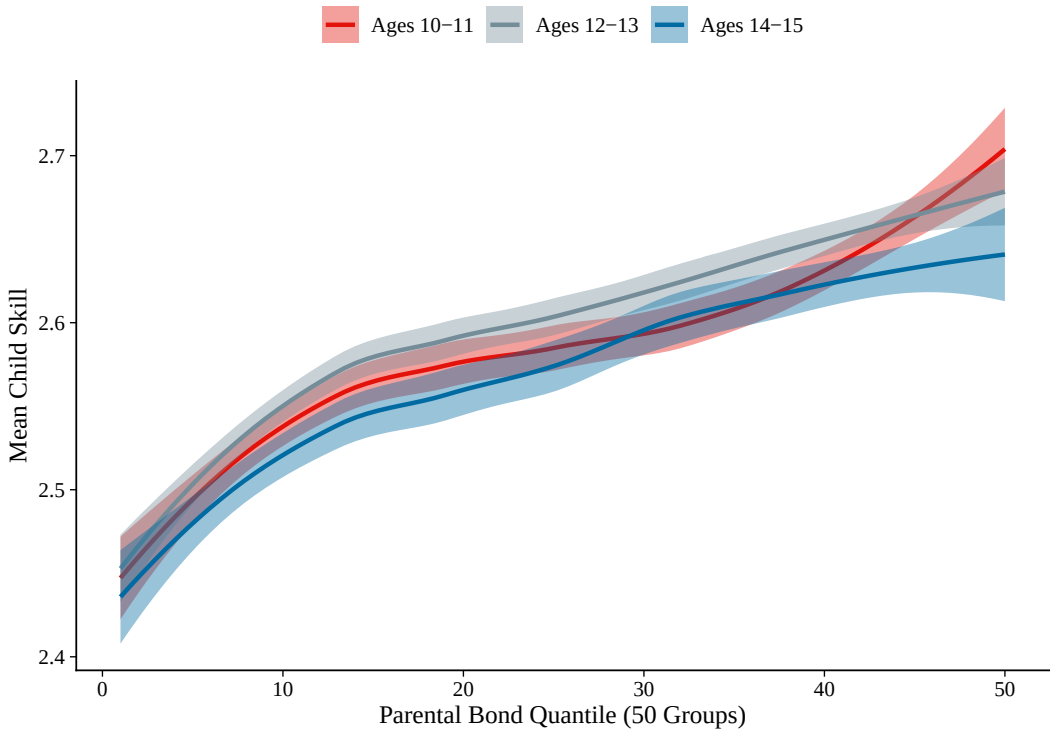
3 Empirical Evidence

I present three sets of empirical evidence. First, I document the association between adolescent socio-emotional skill and parental investment and bond quality, beginning with Figure 1. Second, I present

[19] As with the time-use data, I restrict the sample to households with at least one child under 14 to maintain comparability across data sources and to match the age range over which the bond accumulation process is modelled.

reduced-form regression evidence on the importance of adolescent socio-emotional skill for adult income. Third, I analyze the production functions for child skill and the parental bond.

Figure 1: Child Skill by Parental Bond Quantile (50 Groups), by Age Bin



Notes: Lines are loess-smoothed fits, with 95% confidence bands, of mean log child socio-emotional skill against parental bond quantile (50 groups), computed separately by age bin and shown in the same panel. The quantile grouping is chosen to satisfy the data provider's minimum cell-size requirement for disclosure, not for modelling reasons.

The figure shows a positive relationship between the parental bond and child socio-emotional skill at every age bin. The relationship is concave and steepest at the bottom of the distribution but then flattens over the middle percentiles. However, it continues to remain positive throughout. This suggests that increases in the parental bond at the low end of the distribution are associated with the larger gains in skill.

This pattern holds across time, which suggests that the association between the bond and skill is relatively stable in child development. The correlation between the two latent factor scores, reported in Table C.2, is 0.635 at ages 10–11, 0.593 at ages 12–13, and 0.635 at ages 14–15.[20]

[20] I plot parental latent investment and socioeconomic status in Figure F.2 and show there is a large rise in investment with increasing socioeconomic status.

3.1 Importance of Child Socio-Emotional Skill and Adult Income

I estimate the following regression.

$$\ln \text{Income}_{i,t} = \alpha + \beta_1 \ln \theta_{i,14_15} + \beta_2 \text{Math Score}_{i,14_15} + \beta_3 \ln P_{i,14_15} + Z_i\gamma + \lambda_t + \mu_c + \varepsilon_{i,t}, \quad (1)$$

where $\text{Income}_{i,t}$ is the annual income of individual i at time t. $\theta_{i,14_15}$ is the child's socio-emotional skill, $P_{i,14_15}$ is the parental bond, and $\text{Math Score}_{i,14_15}$ is a standardized math test score at ages 14–15. The child socio-emotional skill and parental bond are Bartlett-scored latent factors constructed from the non-cognitive and parenting scale items in Tables 3 and 4, respectively.[21] Z_i is a vector of household characteristics (father's education, mother's education, household income at ages 14–15, number of siblings, mother's age at the child's birth, and child sex); λ_t is a vector of year fixed effects; and μ_c is a vector of birth-year cohort fixed effects.

Since the T1FF covers 1993 to 2019, the oldest age I observe for adult income in the sample is 36.[22] I report estimates restricting the sample to individuals aged at least 25 years old and above. This is where most individuals have completed formal education and entered the labour market, so that observed income better reflects longer-run earnings potential. Since $\theta_{i,14_15}$ and $P_{i,14_15}$ are estimated with measurement error from the Bartlett scoring procedure, I apply a Croon correction (Croon, 2002) to correct the coefficients for the attenuation bias. I compute 500 bootstraps clustered by child for the standard errors.[23]

[21] A Bartlett score is a generalized least squares estimate of the latent factor from its observed indicators, and it minimizes the score's measurement error relative to the true factor. Appendix C describes the measurement system, and Appendix C.3 derives the Bartlett score estimator in full.

[22] The T1FF records the child's own earnings rather than family earnings. Because children do not begin filing returns until the later years of the panel, income observations are concentrated toward the end of the T1FF window.

[23] Appendix D.1 derives the Croon correction and how it uses the known measurement-error variance of the Bartlett scores from Appendix C.3 to disattenuate the regression coefficients.

Table 5: Croon Corrected Estimates of the Returns to Adolescent Socio-Emotional Skill

	[1]	[2]	[3]	[4]	[5]
Panel A: Age $\geq$ 25					
Log Child Socio-Emotional Skill (θ, per SD)	0.114***	0.106***	0.083**		0.072*
	(0.042)	(0.038)	(0.035)		(0.037)
Log Parental Bond (P, per SD)				0.025**	0.009
				(0.010)	(0.011)
Math Score (standardized)			0.097***		0.096***
			(0.014)		(0.014)
Log Household Income (Ages 14–15)	–	0.082***	0.075***	0.081***	0.074***
		(0.026)	(0.026)	(0.026)	(0.026)
Mother's Education	–	0.001	−0.002	0.001	−0.003
		(0.006)	(0.006)	(0.006)	(0.006)
Father's Education	–	−0.002	−0.004	−0.001	−0.004
		(0.005)	(0.005)	(0.005)	(0.005)
Male (=1)	–	0.258***	0.250***	0.263***	0.251***
		(0.021)	(0.021)	(0.021)	(0.021)
N (individuals, rounded)			3,980		
N (person-years, rounded)			14,800		
Panel B: Age $\geq$ 30					
Log Child Socio-Emotional Skill (θ, per SD)	0.140**	0.130**	0.097*		0.089
	(0.064)	(0.059)	(0.054)		(0.057)
Log Parental Bond (P, per SD)				0.029	0.008
				(0.018)	(0.018)
Math Score (standardized)			0.144***		0.144***
			(0.024)		(0.024)
Log Household Income (Ages 14–15)	–	0.083*	0.071	0.082*	0.071
		(0.047)	(0.045)	(0.047)	(0.046)
Mother's Education	–	0.016	0.012	0.014	0.011
		(0.012)	(0.012)	(0.012)	(0.012)
Father's Education	–	0.005	0.002	0.004	0.002
		(0.009)	(0.009)	(0.009)	(0.009)
Male (=1)	–	0.274***	0.262***	0.282***	0.263***
		(0.037)	(0.036)	(0.037)	(0.036)
N (individuals, rounded)			1,240		
N (person-years, rounded)			4,250		
Controls (round, year FE)			✓		
Controls (siblings, maternal age)	–	✓	✓	✓	✓

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Standard errors in parentheses. Coefficients on $\ln\theta$ and $\ln P$ are per standard deviation of the latent log factor at ages 14–15 (0.112 for skill and 0.248 for the bond), since the latent factors have no natural scale. Column [1] drops Z_i (father's education, mother's education, household income, number of siblings, mother's age at child's birth, child sex), retaining round, year, and cohort fixed effects. Column [2] is the baseline specification with child skill and the full household-characteristics vector Z_i. Column [3] adds the standardized math score. Columns [4]–[5] change the full specification with the parental bond, alone (without child skill or the math score) in [4] and jointly with child skill and the math score in [5]. Panel A restricts the sample to individuals aged 25 and older; Panel B restricts to individuals aged 30 and older. Counts of individuals and person-years are rounded to comply with the data provider's confidentiality requirements.

Table 5 reports the results. Column [1] drops the household-characteristics vector Z_i and shows the skill coefficient is larger once those controls are removed. Since the latent factors have no natural scale, I report the coefficients on $\ln \theta$ and $\ln P$ per standard deviation of each latent factor at ages 14–15 (Bolt et al., 2024). In Panel A, Column [2], a one standard deviation increase in adolescent socio-emotional skill is associated with a 10.6% increase in adult income.

Column [3] adds the standardized math score. The socio-emotional skill coefficient slightly falls to 0.083 (an 8.3% increase in adult income per standard deviation of $\ln \theta$). A one standard deviation increase in the math score is associated with a 9.7% increase in adult income. To note, the math score is likely an imperfect proxy of cognitive ability and it could also be capturing educational attainment, a control I do not observe in the data.

Columns [4]–[5] add in the parental bond to the income equation. On its own in Column [4], without child skill or the math score, the bond is positive and significant (0.025). This is interesting because, even knowing the nurturance, consistency, and warmth that the parent provides to the child at age 14-15 is predictive of adult income. When modelled jointly with child skill and the math score in Column [5], the bond coefficient is small and insignificant, which likely captures the same non-cognitive channel as child skill. The skill coefficient remains positive and significant.

Panel B, restricting the sample to individuals aged 30 and older, reproduces the same pattern with somewhat larger point estimates and wider standard errors because of the smaller sample. The bond enters positively on its own, though it is no longer significant and it remains indistinguishable from zero conditional on skill.

These estimates line up closely with Lindqvist and Vestman (2011), who use Swedish administrative and military enlistment data. They report, using measurement-corrected estimates in standard deviation units, an 8.6% return to socio-emotional skills and an 8.3% return to cognitive skills, magnitudes on the same order as those obtained here despite the differences in country, cohort, and skill measurement. The results in Lindqvist and Vestman (2011) show that the cognitive skills coefficient is sensitive to the inclusion of educational controls, ranging from 2.1% to 8.9% across specifications, a range that lies just below the 9.7% estimate obtained here.[24]

The parental bond coefficient speaks to the results of Barszczewski and Thomas (2025), who examine

[24]The math test score is likely an imperfect proxy for cognitive ability, which could explain the differences with Lindqvist and Vestman (2011).

parenting style along dimensions similar to those underlying the bond measure used here (warmth, reasoning, hostility, and consistency) for children 8 to 15 years old. They find that parenting appears to affect school achievement largely through children's non-cognitive/behavioural development rather than through changes in measured cognitive ability. The bond's explanatory power for adult income is attenuated once child socio-emotional skill is included but survives when skill is excluded.

Socio-emotional skill formed during adolescence therefore carries returns in adulthood even conditional on math score and household income. The next subsection turns to what determines its development, presenting reduced-form evidence on the production functions for child socio-emotional skill and the parental bond, and on the role of parental investment in each.

3.2 Production Functions for Skill and the Parental Bond

I estimate the production function of child socio-emotional skill and the parental bond to answer: do parents investment raise child socio-emotional skill directly or through the parental bond? Continuing with the notation from the previous section, let $\theta_{i,t}$ and $P_{i,t}$ denote the child's socio-emotional skill and the parental bond at age t, respectively. Let $I_{i,t+1}$ denote parental investment realized in the interval between ages t and $t+1$. In these models, each outcome is regressed on its own lag, the lag of the other stock, and parental investment,

$$\ln \theta_{i,t+1} = \gamma_\theta \ln \theta_{i,t} + \delta_\theta \ln P_{i,t} + \beta_\theta \ln I_{i,t+1} + X_{i,t}\eta_\theta + \varepsilon^\theta_{i,t}, \tag{2}$$

$$\ln P_{i,t+1} = \gamma_P \ln P_{i,t} + \delta_P \ln \theta_{i,t} + \beta_P \ln I_{i,t+1} + X_{i,t}\eta_P + \varepsilon^P_{i,t}, \tag{3}$$

where $X_{i,t}$ is a vector of controls (child sex, parental education, number of siblings, mothers age at the childs birth, and round fixed effects). The coefficients γ_θ and γ_P capture the self-productivity, the coefficients δ_θ and δ_P capture the cross-productivity, and the coefficients β_θ and β_P capture the productivity of parental investment.

Several of the underlying items ask about a period of time rather than the interview date itself. Parent-reported school participation asks about behaviour "during this school year," and the child-reported outside-school-activities module asks "during the past 12 months" (Table A.1). The latent factor is best read as investment accumulated up to the interview at which $\theta_{i,t+1}$ and $P_{i,t+1}$ are also recorded, rather than investment that is measured at the same time as $\theta_{i,t}$ and $P_{i,t}$.[25]

[25] This timing of investments follows Attanasio et al. (2020b).

Since $\theta_{i,t}$, $P_{i,t}$, and $I_{i,t+1}$ are estimated with measurement error, I follow the correction proposed by Attanasio et al. (2020b) which has been used to estimate nonlinear production functions for child skills (Attanasio et al., 2017; Bolt et al., 2024). I estimate the residual measurement-error variance of each Bartlett score, recover the error-free joint distribution of the latent factors as a two-component normal mixture, simulate observations from that distribution, and re-estimate the regressions on the simulated data, together with 750 bootstraps to compute standard errors clustered by child.[26] Since the log-latent variables such as socio-emotional skill, bond, and investments have no natural scale, I evaluate them in standard deviation units (Bolt et al., 2024).

Table 6: Reduced Form Evidence

	Cobb-Douglas		Translog	
	Log SE Skill **[1]**	Log Parental Bond **[2]**	Log SE Skill **[3]**	Log Parental Bond **[4]**
Log Investment	0.021	0.110***	0.022	0.108***
	(0.019)	(0.025)	(0.019)	(0.024)
Log SE Skill (Lag)	0.605***	−0.032	0.662***	0.026
	(0.049)	(0.046)	(0.045)	(0.043)
Log Parental Bond (Lag)	0.100*	0.675***	0.163***	0.743***
	(0.052)	(0.069)	(0.051)	(0.074)
Log Investment × Own Lag	—	—	0.014	−0.005
			(0.017)	(0.018)
Skill (Lag) × Bond (Lag)	—	—	0.089***	0.097***
			(0.021)	(0.017)
Mother's Education	−0.007	−0.006	−0.008	−0.007
	(0.009)	(0.010)	(0.009)	(0.010)
Father's Education	0.017*	0.006	0.014	0.004
	(0.009)	(0.009)	(0.009)	(0.009)
Number of Siblings	−0.013	0.005	−0.002	0.016
	(0.020)	(0.022)	(0.020)	(0.022)
Controls		✓		
Observations	3,480	3,480	3,480	3,480

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Measurement-error-corrected estimates, with standard errors from 750 bootstraps clustered by child. The latent factors are standardized, so coefficients are in standard deviation units. Full specification, pooling both age-bin transitions with an age-bin control. Columns [3] and [4] add the interaction terms of Table E.2 (Columns [2] and [4]), with investment and the lagged regressors centred at their sample means. Investment × Own Lag is the interaction of investment with the lag of the dependent variable (skill in Column [3], bond in Column [4]). Controls include mother's education, father's education, child sex, number of siblings, mother's age at the child's birth, and round fixed effects. See Table E.1 for the same specification estimated separately by age-bin transition.

[26] See Attanasio et al. (2020b) for a full derivation of the correction.

Table 6 shows three features worth noting.[27] Investment produces the parental bond but not child skill directly once lagged skill and bond are controlled for. Investment in skill formation appears to run through the bond, since its direct effect on skill is not significant. Both the parental bond and socio-emotional skill are strongly self-productive over this age range.

Both skill and bond are strongly self-productive. A one standard deviation increase in their lagged skill (bond) is associated with a 60.5 (67.5) percent of a standard deviation increase in future skill (bond) in the Cobb-Douglas specification. In the translog specification, they both rise to 66.2 and 74.3 percent of a standard deviation, respectively.

The cross-productivities are asymmetric. The bond raises future skill, such that a one standard deviation increase in the parental bond is associated with a 10 percent of a standard deviation increase in future socio-emotional skill in the Cobb-Douglas specification. This rises with the translog specification to 16 percent of a standard deviation. The cross-productivity of skill on future bond flips sign across specifications and is statistically insignificant.

A one standard deviation increase in investment is associated with an 11 percent of a standard deviation increase in future bond, with a nearly identical coefficient in the translog specification. However, investment does not produce future skill directly, with a small and statistically insignificant coefficient in both specifications.

Columns [3] and [4] of Table 6 add interaction terms to the production functions and show that the bond and skill are complements in production, with an interaction of 0.089 in the skill equation and 0.097 in the bond equation, whereas the investment-skill or investment-bond interactions are not significant. This means that when a child has stronger socio-emotional skills (e.g., pro-sociality, emotional control), the bond becomes more effective at producing future socio-emotional skills, and the same holds for bond production.

The lack of a direct effect of investment on socio-emotional skill differs from Cunha et al. (2010) and Toppeta (2022), who find that investment has a direct effect on socio-emotional skills. Both studies build their investment factor from time and material inputs together with items on the parent-child relationship. Cunha et al. (2010) include how often the child receives encouragement and praise, while Toppeta (2022)

[27] An observation is a child-transition pair, so a child observed in all three age bins contributes two. The pooled sample stacks the two transitions in Table E.1 (1,600 and 1,890 observations), which sum to 3,490 rather than 3,480 because each count is rounded independently to the nearest 10.

include how often a parent shouts at, smacks, or bribes a misbehaving child, alongside a direct measure of maternal closeness. If the bond is a separate channel through which investment affects skill, then including it in the investment factor could credit the direct effect of investment on skill.

I directly test this in Table E.3 by estimating the same production functions with latent investment being constructed from the normal time and material inputs, and the parent-child relationship. For the skill production function, the coefficient on investment rises from 0.021 to 0.469 and is significant (Column [4]). Dropping the lag bond from the original specification (Column [3]) makes the coefficient on investment rise to 0.034 and becomes marginally significant. For both production functions, dropping investment leaves the cross-productivity unchanged. Together these results suggest that investment factors that include the parent-child relationship are likely to overstate the direct effect of investment on future skill. This is consistent with the growing recognition that parenting behaviours capturing the quality of the parent-child relationship should be treated as a distinct input into child development (Cobb-Clark et al., 2019; Moroni et al., 2025).

To relate my results to the literature, I turn to studies whose investment measures exclude the parent-child relationship. Attanasio et al. (2020b) study infants, aged 12 to 42 months, and disaggregate investment into its time and material components, finding that time investment matters more for the accumulation of non-cognitive skills while material investment matters more for cognitive skills, though their results are difficult to compare directly with those reported below because the underlying items are designed for infants rather than adolescents and the age range they study lies well below the age range studied here.

Mitchell et al. (2023) study an age range closer to the one studied here and, estimating a Cobb-Douglas production function, find a positive but statistically insignificant elasticity of socio-emotional skill with respect to investment between ages 12 and 15 and between ages 15 and 19, and a significant positive elasticity between ages 8 and 12. Their setting is in a developing-country context, Peru, which differs substantially from the Canadian setting studied here.

Table 7 reports measurement-error-corrected estimates of the log-linear investment equation

$$\ln I_{i,t+1} = \phi_P \ln P_{i,t} + \phi_\theta \ln \theta_{i,t} + \phi_Y \ln Y_{i,t} + X_{i,t}\eta_I + \varepsilon^I_{i,t}, \tag{4}$$

in which the dependent variable is the latent investment index $I_{i,t+1}$ and the regressors include the parental bond $P_{i,t}$, child socio-emotional skill $\theta_{i,t}$, measured at the start of the investment interval, and household income $Y_{i,t}$, together with parental education, child sex, number of siblings, mother's age at the child's

birth, and round and age-bin fixed effects.

The coefficient on P is positive at 0.217, consistent with a reinforcing pattern in which parents with a stronger bond invest more. In other words, a strong parent-child relationship encourages further parental engagement. The coefficient on θ is negative, which is consistent with a compensating pattern. Parents invest more in children with lower socio-emotional skill. A child that is depressed, anxious, or withdrawn, for example, may require more parental attention and support than a child that is outgoing and socially confident. Income is also positive and significant, which is means that parents with more resources are able to invest more in their children. Household income and mother's education are both positive and significant, which suggest households with more resources and more educated parents are able to invest more in their children.

Table 7: Investment Equation (Measurement-Error Corrected)

	Log Investment
Log Parental Bond (Lag)	0.2174^{***}
	(0.0505)
Log SE Skill (Lag)	-0.0877^{**}
	(0.0443)
Log Household Income	0.0824^{**}
	(0.0354)
Mother's Education	0.0409^{***}
	(0.0141)
Father's Education	0.0167
	(0.0113)
Male (=1)	-0.0936
	(0.0600)
Number of Siblings	-0.0713^{**}
	(0.0286)
Mother's Age at Child's Birth	-0.0018
	(0.0048)
Round FE	✓
Age-bin FE	✓
N (observations, rounded)	3,480

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Measurement-error-corrected estimates, with standard errors from 750 bootstraps clustered by child, as in Table 6. Dependent variable is the latent investment index $I_{i,t+1}$; the bond and skill regressors, $P_{i,t}$ and $\theta_{i,t}$, are measured at the start of the investment interval. The latent factors are standardized, so the dependent variable and the lagged bond and skill are in standard deviation units. Rounding of observations is applied to comply with the confidentiality requirements of the data provider. See Table E.4 for the same specification estimated separately by age-bin transition.

Together, these patterns motivate the technology of the production functions of child socio-emotional skill and the parental bond. However, it cannot tell us the behavioural mechanism on how households choose investment and the bond in response to economic conditions, which I develop in the next section.

4 A Dynamic Model of Household Decisions and the Parental Bond

I develop a dynamic model of two-parent household behaviour. Households face a trade-off among earning income, enjoying leisure, and spending time with their child. Their decisions respond to wage, bond, and skill shocks as well as to a nonlinear tax-and-transfer system. I abstract from household saving as in Del Boca et al. (2014), Agostinelli et al. (2024), and Del Boca et al. (2025). With the skill-bond interaction, the marginal values of skill and the bond depend on the stocks, so I solve the household's dynamic problem numerically.[28]

The household maximizes its unitary utility over childhood, indexed by $t = 0, \ldots, T$. At $T+1$ the child enters new stages of life (such as adulthood), and the household receives a terminal payoff that depends on the child's skill θ_{T+1}. The household's problem can be written recursively as

$$V_t(S_t) \;=\; \max_{x_t \in \Gamma(S_t)} \Big\{ u_t(c_t, \ell_{1t}, \ell_{2t}, P_t) + \beta\, \mathbb{E}\big[V_{t+1}(S_{t+1}) \mid S_t, x_t\big] \Big\}, \qquad t = 0, \ldots, T.$$

Here, $x_t = (c_t, e_t, h_{1t}, h_{2t}, m_{1t}, m_{2t})$ collects the period-t controls, namely consumption, child expenditures, hours worked, and time invested in the child by each parent. Throughout, subscript 1 denotes the mother and subscript 2 the father. Flow utility u_t is a function of consumption, leisure ℓ_{jt} for each parent j, and the parental bond P_t. Each period $t = 0, \ldots, T$, the household observes the state variables:

$$S_t = (\theta_t, P_t, w_{1t}, w_{2t}),$$

where w_{jt} is parent j's wage offer. The skill and bond are latent state variables that evolve endogenously. Wages are observed by the household and follow an estimated stochastic process. Socio-emotional skill can be thought of as an asset that is built up over childhood and pays off at $T+1$, while the bond can be thought of as a dividend-paying asset that pays out in each period of childhood.

The motivation for the bond in the flow utility function is that parents derive warm-glow utility from the relationship with their child. This is consistent with Guryan et al. (2008), who note that well-cared-

[28] Given the marginal values, I solve the within-period choices and then simulate each household forward. I update the marginal values backward from the terminal period along each household's simulated path, and I repeat the two steps until the marginal values converge.

for children may contribute directly to a parent's utility, and who review evidence that parents report recreational and educational time with their children to be among their most enjoyable activities.

Budget and time constraints. The household's constraints are

$$c_t + e_t = \mathcal{T}(Y_t), \quad Y_t = w_{1t}h_{1t} + w_{2t}h_{2t}, \qquad \ell_{jt} + h_{jt} + m_{jt} = 1 \text{ for } j = 1, 2.$$

The first constraint requires parents to allocate after-tax income between consumption and child expenditures. The second requires each parent to allocate a unit of time between leisure, market work h_{jt}, and active time with the child m_{jt}. Y_t is the household's before-tax income, which is the sum of each parent's wage multiplied by hours worked. The function $\mathcal{T}(\cdot)$ maps before-tax income to after-tax income.

Households face a progressive tax-and-transfer system that mirrors the Canadian tax structure, following the nonlinear generalization of the tax system used in Blundell et al. (2016, 2018) and Agostinelli et al. (2024), but adapted for the Canadian context using Canadian Tax and Credit Simulator (CTaCS).[29] After-tax income is given by the function $\mathcal{T}(\cdot)$:

$$\mathcal{T}(Y_t) = \mathcal{X}_1\big(\mathcal{X}_2 + Y_t\big)^{1-\mathcal{X}_3}. \tag{5}$$

$\mathcal{X}_1$ scales the overall generosity of the tax system by shifting the entire budget set up or down. $\mathcal{X}_2$ is a benefit floor that could represent means-tested transfers such as the Canada Child Benefit, GST/HST credit, and provincial assistance that provide resources even when a household has low earnings. Lastly, $\mathcal{X}_3$ governs the progressivity of the tax system. Higher values mean that after-tax income increases more slowly with pre-tax income. High-income households face relatively higher effective tax rates and disposable incomes are more compressed as we increase $\mathcal{X}_3$.[30]

Preferences and functional forms. Per-period household utility is unitary, additively separable, and logarithmic,

$$u_t = \alpha_c \ln c_t + \alpha_P \ln P_t + \alpha_{\ell_1} \ln \ell_{1t} + \alpha_{\ell_2} \ln \ell_{2t},$$

[29] Blundell et al. (2016, 2018) use the NBER TAXSIM program, which was developed for the U.S. tax system. This generalization fits well with the Canadian tax system, which is also progressive and has similar features, while respecting two institutional differences, (i) individual income is taxed separately, but most refundable credits are means-tested on family net income; and (ii) the "floor" is delivered mainly through refundable credits (Canada Child Benefit and GST/HST credit) and provincial social assistance, rather than food stamps.

[30] This formulation is well suited for policy counterfactuals because changing any of the parameters directly corresponds to simple, interpretable reforms, such as increasing benefit generosity, adjusting the income floor, or altering progressivity.

with the normalization $\alpha_c + \alpha_P + \alpha_{\ell_1} + \alpha_{\ell_2} = 1$, which fixes the scale of flow utility and hence the units in which the terminal weight is measured, as in Del Boca et al. (2014). The bond P_t enters directly, so parents draw warm-glow utility from the relationship with weight α_P, as motivated above.[31]

Skill terminal. Following Del Boca et al. (2014), I treat the child's socio-emotional skill at $T+1$ as the initial condition for a later stage of life that parents care about but that the model does not follow.

To illustrate how the terminal value could be derived, suppose the child enters adulthood at $T+1$ with skill θ_{T+1} and lives in a stationary environment, so that the adult value function satisfies $V^{\text{child}} = \ln y^* + \beta V^{\text{child}}$. Let y^* capture some vector of adult outcomes, for example, adult earnings. I assume that $\ln y^* = \psi \ln \theta_{T+1}$ where ψ is the elasticity of adult outcomes with respect to adult skill. Then the fixed point would be $V^{\text{child}}(\theta_{T+1}) = \psi \ln \theta_{T+1}/(1-\beta)$. I introduce κ that captures how much parents care about the child's adult outcomes relative to their own utility that yields:

$$V^{\text{parents}}(\theta_{T+1}) = \kappa V^{\text{child}}(\theta_{T+1}) = \kappa \frac{\psi}{1-\beta} \ln \theta_{T+1}.$$

Only the total $\kappa\psi/(1-\beta)$ enters the household's problem, so κ and ψ are not separately identified.[32] Instead, I define $\lambda_\theta \equiv \kappa\psi/(1-\beta)$ and estimate it jointly with the remaining structural parameters inside the SMM objective, so that the terminal value of skill is $\lambda_\theta \ln \theta_{T+1}$. This follows the approach of Del Boca et al. (2014).

Technologies. Child skill and the parental bond evolve according to according to translog production functions restricted to a skill-bond interaction. The data do not support investment interactions, but do support the skill-bond interaction, so I include it in both laws of motion. This interaction makes optimal choices depend on the stocks.

Laws of motion. The child's skill evolves according to

$$\ln \theta_{t+1} = \ln A(t) + \gamma_1 \ln \theta_t + \gamma_2 \ln I_t + \gamma_3 \ln P_t + \gamma_4 \ln \theta_t \ln P_t + \varepsilon_{t,\theta}, \tag{6}$$

[31] The logarithmic form fixes the elasticity of substitution among consumption, leisure, and the bond at one, so the preference weights are identified from the level and allocation of expenditure and time.

[32] ψ could be a vector of elasticities representing the returns to different adult outcomes. It is true that I could estimate ψ from the data on adult outcomes as my data allows it. However, it will make a strong assumption that parents only care about socio-emotional skill due to its impact on wages. κ is also not identified from the data because it is a free parameter that captures how much parents care about their child's adult outcomes relative to their own utility. Thus I follow the more flexible and standard approach of the literature.

while the parental bond evolves according to

$$\ln P_{t+1} = \ln \tilde{A}(t) + \tilde{\gamma}_1 \ln P_t + \tilde{\gamma}_2 \ln I_t + \tilde{\gamma}_3 \ln \theta_t + \tilde{\gamma}_4 \ln \theta_t \ln P_t + \varepsilon_{t,P},$$

where $A(t) = \exp(a + \delta_\theta\, t)$ and $\tilde{A}(t) = \exp(\tilde{a} + \delta_P\, t)$ are age-varying productivity terms. The coefficients γ_1 and $\tilde{\gamma}_1$ capture self-productivity in skill and bond, while γ_2 and $\tilde{\gamma}_2$ measure the direct role of contemporaneous investment. The cross-stock terms $\gamma_3 \ln P_t$ and $\tilde{\gamma}_3 \ln \theta_t$ allow the parental bond and child skill to be cross-productive. Because γ_4 and $\tilde{\gamma}_4$ multiply the uncentred product $\ln \theta_t \ln P_t$, the persistence of each stock and its cross-productivity depend on the level of the other stock. The shocks $\varepsilon_{t,\theta}$ and $\varepsilon_{t,P}$ are innovations to skill and bond formation.

Investment aggregator. Investment I_t combines spending on the child with each parent's active childcare time in a single latent index,

$$\ln I_t = \alpha_e \ln e_t + \alpha_{m_1} \ln m_{1t} + \alpha_{m_2} \ln m_{2t} + \alpha_{12} \ln m_{1t} \ln m_{2t}.$$

The coefficients α_e, α_{m_1}, and α_{m_2} are the elasticities of I_t with respect to spending and each parent's time, and α_{12} allows maternal and paternal time to be complements or substitutes.[33] Because I_t is latent and has no natural scale, I normalise these elasticities to sum to one, $\alpha_e + \alpha_{m_1} + \alpha_{m_2} = 1$, so that each measures an input's relative contribution to the composite.[34]

Why do parents invest in the bond? Given the preference, terminal, and technology specifications above, a parent facing binding time and income constraints directs resources toward the parent-child bond, rather than toward own consumption alone, for two reasons, one for each term in which P_t appears. First, the bond enters flow utility directly through $\alpha_P \ln P_t$, so that parents obtain current-period utility from a stronger relationship with their child, independently of any later payoff. Second, the bond raises the productivity of skill formation through the cross-stock terms $\gamma_3 \ln P_t$ and $\gamma_4 \ln \theta_t \ln P_t$ in the skill law of motion (6), so that even a parent who places little direct weight on the bond still has a skill-driven

[33] Unlike Cunha et al. (2010) and Attanasio et al. (2020b), where current skill enters the investment equation so that input returns depend on the child's developmental level, I_t here is a pure input aggregator, with current skill entering the production function separately.

[34] The adding-up condition is a scale normalisation, not an economic restriction, and $\alpha_{12} \neq 0$ rules out constant returns to scale in any case. Since $\ln I_t$ has no natural unit, scaling α_e, α_{m_1}, α_{m_2}, and α_{12} by a common $\lambda > 0$ multiplies $\ln I_t$ by λ, which leaves $\gamma_2 \ln I_t$ and $\tilde{\gamma}_2 \ln I_t$ unchanged once γ_2 and $\tilde{\gamma}_2$ are divided by λ. Only the products $\gamma_2\alpha_e$, $\gamma_2\alpha_{m_1}$, $\gamma_2\alpha_{m_2}$, and $\gamma_2\alpha_{12}$ are identified, and the adding-up condition breaks this indeterminacy.

reason to invest in it. Only the first of these channels introduces a free preference parameter, α_P. The skill-spillover channel is disciplined by γ_3 and γ_4.

Wage process. For each parent $j \in \{1, 2\}$, the log wage offer in age bin t is

$$\ln w_{jt} \;=\; b_{0j} + b_{1j}\,\mathrm{edu}_j + b_{2j}\,\mathrm{age}_{jt} + \varepsilon_{jt}, \tag{7}$$

where edu_j is years of schooling and age_{jt} is parental age when the child is in age bin t. The wage shocks are drawn as i.i.d. mean-zero normals with standard deviations $(\sigma_{\mathrm{mom}}, \sigma_{\mathrm{dad}})$ and contemporaneous correlation ρ_{wage}:

$$\begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \;\sim\; \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix},\; \begin{pmatrix} \sigma^2_{\mathrm{mom}} & \rho_{\mathrm{wage}}\,\sigma_{\mathrm{mom}}\sigma_{\mathrm{dad}} \\ \rho_{\mathrm{wage}}\,\sigma_{\mathrm{mom}}\sigma_{\mathrm{dad}} & \sigma^2_{\mathrm{dad}} \end{pmatrix} \right).$$

The cross-parent correlation ρ_{wage} captures assortative mating on unobservable wage determinants and the fact that both parents face the same local labour market. I treat the shocks as serially uncorrelated, so wage offers are independent across age bins conditional on observables. The parameters $(b_{0j}, b_{1j}, b_{2j}, \sigma_j)$ and ρ_{wage} are estimated jointly with the structural parameters inside the SMM objective.[35]

Identification of the model parameters is complicated by two econometric issues. The first arises because the measures underlying the latent factors, socio-emotional skill and the parental bond, are subject to measurement error. I address this by using a measurement model to extract the latent factors from the observed indicators. The second arises because the latent factors are not directly observed at the individual level, which complicates the construction of moments that involve both the latent factors and covariates. I address this by exploiting the properties of Bartlett factor scores to construct such moments. Appendix C presents the full measurement system, including the age-invariant specification used for the anchoring restriction in Section 5.3, and describes how latent moments are constructed from Bartlett scores.

[35] This specification, combined with the absence of saving and borrowing, implies that wage risk operates through a purely static reallocation of time and expenditure within each age bin rather than through a precautionary or intertemporal-insurance margin.

5 Estimation

I use a simulated method of moments (SMM) approach to estimate the model parameters. The SMM approach is suitable for this model because it accommodates complex dynamics and because the data come from independent sources. The estimation procedure involves two steps, (1) estimating the initial parameters *outside* the model and (2) estimating the remaining parameters *inside* the model using SMM. The two-step approach reduces the complexity of the estimation.

Wage offers are observed only for those who accept employment, which creates a selection problem. I address it by estimating the wage process parameters and the initial joint distribution of wages and skills jointly with the rest of the model. Estimating wages within the model ensures that the estimates are internally consistent with the selection induced by both parents' participation decisions.[36]

There are five sets of parameters to estimate, (1) preferences, (2) skill technology, (3) parental bond law of motion, (4) wage offers, and (5) tax parameters. I estimate the tax parameters in the first step, while the wage, preference, skill technology, and parental bond law of motion parameters are estimated inside the SMM in the second step.

5.1 Step 1: Estimate Initial Parameters outside the model

1. Set the discount factor to $\beta = 0.96^2$, since each period covers two years.
2. Estimate the tax-and-transfer system using the Canadian Tax and Credit Simulator (CTaCS).[37][38] Using nonlinear least squares (NLS), I estimate the parameters of the tax function $\mathcal{T}(Y) = \mathcal{X}_1(\mathcal{X}_2 + Y)^{1-\mathcal{X}_3}$. The estimates are $\mathcal{X}_1 = 3.058$, $\mathcal{X}_2 = 742.54$, and $\mathcal{X}_3 = 0.122$.[39]
3. I draw conditional on observed mother and father education and household income. The initial skill level of the child, θ_0, and the initial parental bond, P_0, are drawn directly from the moments in the empirical data. I assume that each is normally distributed with mean μ_θ and standard deviation σ_θ.

[36] This follows the same approach as Del Boca et al. (2014), Agostinelli et al. (2024), and Del Boca et al. (2025).

[37] See Milligan and Stabile (2011) and Milligan (2007) for this tax benefit calculator.

[38] This is similar to the NBER TAXSIM program, which calculates the tax liability of a household based on its income and family structure.

[39] Since children's families are observed in different years, I keep household income in nominal values and estimate after-tax income using the CTaCS tax simulator, because tax rates differ across years and provinces.

5.2 Step 2: Estimate Remaining Parameters inside the model

From Step 1 I take the tax function parameters $(\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3)$ and the initial joint distribution of (θ_0, P_0). The parameter vector Ω estimated inside the SMM has four blocks.

1. *Utility parameters.*

$$\{\alpha_P, \alpha_{\ell_1}, \alpha_{\ell_2}, \lambda_\theta\},$$

subject to the normalization $\alpha_c = 1 - \alpha_P - \alpha_{\ell_1} - \alpha_{\ell_2}$. The first three set the intratemporal allocation between consumption, the parental bond, and parental leisure. The skill terminal λ_θ is the payoff that parents receive at $T + 1$.

2. *Investment aggregator.*

$$\{\alpha_{m_1}, \alpha_{m_2}, \alpha_{12}\},$$

subject to the normalization $\alpha_e = 1 - \alpha_{m_1} - \alpha_{m_2}$. These determine how child expenditures and each parent's time produce the aggregate investment index $\ln I_t$. The term α_{12} is the interaction between maternal and paternal time.

3. *Skill and bond technologies.*

$$\{a, \delta_\theta, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \sigma_\varepsilon^\theta\} \quad \text{and} \quad \{\tilde{a}, \delta_P, \tilde{\gamma}_1, \tilde{\gamma}_2, \tilde{\gamma}_3, \tilde{\gamma}_4, \sigma_\varepsilon^P\}.$$

In the laws of motion for child skill and for the parental bond, these parameters set the level and age trend of productivity (a, δ_θ and $\tilde{a}, \delta_P$), self-productivity ($\gamma_1, \tilde{\gamma}_1$), the returns to aggregate investment ($\gamma_2, \tilde{\gamma}_2$), the cross-stock effects of the bond on skill and of skill on the bond ($\gamma_3, \tilde{\gamma}_3$), the skill-bond interactions ($\gamma_4, \tilde{\gamma}_4$), and the variances of the two innovations.

4. *Wage parameters.*

$$\{b_{01}, b_{11}, b_{21}, \sigma_{\text{mom}},\ b_{02}, b_{12}, b_{22}, \sigma_{\text{dad}},\ \rho_{\text{wage}}\}.$$

For each parent $j \in \{1, 2\}$, this is the intercept, education return, age profile, and wage shocks in Equation (7). The cross-parent correlation ρ_{wage} reflects assortative mating or parents in the same local labor market (Del Boca et al., 2025).[40]

[40] I assume the shocks are serially independent for convenience. The model does not imply this, and identification of the main parameters does not depend on it. Relaxing it is feasible at a higher computational cost, and Del Boca et al. (2025) estimate a similar model with serially correlated shocks.

For a given Ω, I generate the distribution of wage offers from the wage parameters, solve the household problem for the optimal policy functions of all endogenous variables, and simulate the joint distribution of next period's skill and parental bond implied by those policy functions. These simulations give the vector of simulated moments $\widehat{m}(\Omega)$. The SMM estimator then chooses Ω to minimize the weighted distance between $\widehat{m}(\Omega)$ and the vector of empirical moments m^{data},

$$\widehat{\Omega} = \arg\min_{\Omega} \left[m^{\text{data}} - \widehat{m}(\Omega)\right]' W \left[m^{\text{data}} - \widehat{m}(\Omega)\right], \tag{8}$$

where $W = \text{diag}(\widehat{\Sigma}_B)^{-1}$ is a diagonal weighting matrix and $\widehat{\Sigma}_B$ is the sample covariance matrix of the moment vectors computed on 1,000 bootstrap samples drawn from the data. Moments estimated more precisely therefore receive more weight. I minimize the objective with the Nelder-Mead simplex algorithm.

5.3 Identification

Table 8 maps each parameter block to the moments that discipline it.[41] A single moment does not necessarily pin down a single parameter, since several parameters can affect the same moment. Identification comes from how the parameters jointly affect the full set of moments.

Persistence, cross-stock effects, and innovation variances. The persistence and cross-stock coefficients $\gamma_1, \tilde{\gamma}_1, \gamma_3, \tilde{\gamma}_3$ are disciplined jointly by the age-bin variances, autocovariances, and cross-lag covariances of $\ln \theta_t$ and $\ln P_t$ and by the measurement-error-corrected production function estimates (Table 8). The same estimates provide the skill-bond interactions γ_4 and $\tilde{\gamma}_4$. The productivity terms $a, \tilde{a}, \delta_\theta, \delta_P$ come from the age profile of the latent means once the emotional-skill and parental-nurturance scales are held fixed across age bins,[42] and the innovation variances $\sigma^\theta_\varepsilon, \sigma^P_\varepsilon$ follow as the residual variation left after accounting for persistence and cross-stock effects.

Investment aggregator composition. The direct-investment coefficients $\gamma_2, \tilde{\gamma}_2$ come from NLSCY regressions of each latent state on its own lag and on lagged investment $\ln I_{t-1}$, which separate the direct

[41] Identification requires an order condition (at least as many moments as parameters) and a rank condition (parameters must have differential impacts on moments). The SMM step estimates 30 parameters in total after imposing the two adding-up restrictions in Section 5, comprising 3 utility weights, the skill terminal λ_θ, 3 investment-aggregator weights, 7 skill-technology parameters, 7 bond-technology parameters, and 9 wage parameters.

[42] This assumption is known in the literature as age-invariant measurement (Wiswall and Agostinelli, 2024; Freyberger, 2026; Del Bono et al., 2022).

Table 8: Parameter Identification Summary

Parameter block	Key parameters	Primary identifying moments
Skill and bond technologies	$\gamma_1, \tilde{\gamma}_1, \gamma_2, \tilde{\gamma}_2, \gamma_3, \tilde{\gamma}_3, \gamma_4, \tilde{\gamma}_4, a, \tilde{a}, \delta_\theta, \delta_P$	Measurement-error-corrected production function estimates; age-bin variances, autocovariances, and cross-lag covariances of $\ln \theta_t$ and $\ln P_t$; covariance of latent states with investment; age profile of latent means
Innovation variances	$\sigma_\varepsilon^\theta, \sigma_\varepsilon^P$	Latent variances and autocovariances, net of persistence and cross-stock effects
Investment aggregator composition	$\alpha_e, \alpha_{m_1}, \alpha_{m_2}, \alpha_{12}$	SHS expenditure-to-income ratio by employment status; GSS level and division of parental time by couple type and employment state
Leisure and participation	$\alpha_{\ell_1}, \alpha_{\ell_2}$	NLSCY employment states by age bin (mother, couple); hours worked conditional on employment
Consumption weight	α_c	After-tax household income moments; GSS active-time–market work trade-off slopes; SHS income elasticity of the expenditure share
Bond weight and skill terminal	α_P, λ_θ	Age-bin profiles of employment, hours, and household income
Wage offers	$b_{0j}, b_{1j}, b_{2j}, \sigma_j, \rho_{\text{wage}}$	Accepted-wage mean, education covariation, and age profile by parent; wage dispersion conditional on employment; spousal wage correlation

Notes: $j \in \{1, 2\}$ indexes mother and father. The table lists the primary source of identifying variation for each block; several parameters are also disciplined jointly with others, as discussed in the text.

effect of investment from the state dependence already captured by $\gamma_1, \tilde{\gamma}_1, \gamma_3, \tilde{\gamma}_3$. SHS expenditure moments and GSS time-use moments by employment regime pin down $\alpha_e, \alpha_{m_1}, \alpha_{m_2}$ and the interaction term α_{12}. This composition margin is distinct from the SHS income elasticity of the expenditure share below, which disciplines total spending on the child rather than how that spending splits between goods and parental time.

Preference weights, participation, and wage parameters. Employment and hours moments by age bin pin down α_{ℓ_1} and α_{ℓ_2}, while the GSS active-time–market work trade-off slopes and the SHS income elasticity of the expenditure share discipline α_c, jointly with the after-tax household income moments that fix the budget set. The bond weight α_P and the skill terminal λ_θ both act through the marginal value of investment and so cannot be separated within a single age bin. As in Del Boca et al. (2014), who recover the terminal valuation from labour supply decisions, they are separated by the age profile of the employment, hours, and income moments, since the bond enters flow utility in every period whereas skill enters only through the terminal value, so that the two motives weaken at different rates away from the terminal bin, with loadings that depend only on β and the technology parameters disciplined above.[43] The wage-offer parameters are identified from the distribution of accepted wages by age bin, parent, and education in the NLSCY, with the dispersion terms σ_j net of the employment selection already disciplined above and ρ_{wage} from the correlation of spousal wage residuals.

[43]The data discipline most directly the household's combined motive for investing, the marginal value of investment, which governs how much leisure, market work, and consumption parents give up for the child. The split of that motive between α_P and λ_θ rests on its age profile alone and is identified less sharply.

6 Results

6.1 Model Fit

Tables 9 and 10 show the model fits of the model-simulated moments alongside the empirical moments. The mean of the skill and bond is drawn at ages 10–11 from the empirical distribution, thus coincides with the data by construction. The standard deviations at ages 10–11 are those of the latent distribution, which are below the standard deviations of the factor scores in the data. For labour supply, the model reproduces the rising profile of mothers' hours and the high employment of fathers. It estimates mothers' hours up to about two hours above the data and mothers' employment rate up to five percentage points below it. Fathers' hours rise slightly with the child's age in the model, from 41.4 to 42.3, and are roughly flat in the data, at 42.2 to 42.8. These differences could be due to the fact that hours tend to be more noisy in the empirical data, thus the block-bootstrap procedure downweights these moments in the estimation. Appendix G reports the fit of additional moments.

Table 9: Model Fit: Child Skill and Parental Bond

	Data		Model	
Child Age	Mean	SD	Mean	SD
Panel A: Child Skill				
10–11	2.600	0.129	2.600	0.105
12–13	2.561	0.106	2.557	0.150
14–15	2.505	0.112	2.567	0.190
Panel B: Parental Bond				
10–11	3.158	0.241	3.159	0.194
12–13	3.101	0.254	3.048	0.299
14–15	3.028	0.248	3.064	0.359

Notes: Empirical moments are from the NLSCY. Model moments are averages over 3,480 simulated households and 50 replications. Parental bond is measured on the scale implied by the loadings in Table C.1. Simulated households draw initial skill and bond at ages 10–11 from the latent distribution, whose standard deviations (0.105 and 0.195) are below the factor-score standard deviations in the Data columns.

Table 10: Model Fit: Lifecycle Moments by Parent

	Mother				Father			
	Data		Model		Data		Model	
Child Age	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Panel A: Intensive Margin (Hours Working)								
10–11	30.40	12.79	32.67	14.97	42.21	11.29	41.41	13.46
12–13	31.01	13.08	32.59	15.07	42.77	10.68	41.87	13.39
14–15	32.11	12.43	32.90	15.16	42.84	11.05	42.33	13.26
Panel B: Extensive Margin (Employment Rate)								
10–11	0.829	—	0.847	—	0.967	—	0.960	—
12–13	0.863	—	0.844	—	0.966	—	0.964	—
14–15	0.895	—	0.845	—	0.961	—	0.966	—
Panel C: Accepted Wage (Conditional on Employment)								
10–11	2.789	0.641	2.845	0.479	3.072	0.485	3.004	0.532
12–13	2.829	0.646	2.918	0.482	3.098	0.460	3.095	0.529
14–15	2.833	0.618	2.976	0.473	3.115	0.495	3.166	0.540

Notes: Empirical moments are from the NLSCY. Model moments are averages over 3,480 simulated households and 50 replications. Intensive-margin hours are conditional on employment. Accepted wages are in log points. Employment rate has no empirical SD reported.

6.2 Parameter Estimates

Child skill and the parental bond are latent, so neither stock has a natural unit. This makes some parameters in Table 11 hard to interpret and compare. Consider the flow weight on the bond at 0.122.[44] This number cannot be compared with the other weights, because its size depends on how the bond is normalized. Suppose instead that I normalized the bond on parental rejection, whose loading is 0.57. The log bond would then be $\ln P_t^{\text{rej}} = 0.57 \ln P_t$, so that $\alpha_P \ln P_t = (\alpha_P / 0.57) \ln P_t^{\text{rej}}$. Rescaling the four weights to sum to one would give 0.196 as the bond weight.[45] Following the same logic for parental knowledge, whose loading is 0.42, would give 0.249. A similar issue arises for the cross-productivity and investment parameters.

Two comparisons do not depend on these units. The weights on consumption and on each parent's leisure can be compared with each other, because utility has no units of its own and only their ratios matter. The weights in the investment aggregator apply to time and money, which have natural units. Using the point estimates alone, I cannot compare how much investment raises skill with how much it raises the bond, or how much the bond raises skill with how much skill raises the bond. Nor can I compare the weight on the bond with the weights on consumption and leisure.

Instead, I use the experiments in Section 7.3 to make sense of the policy-relevant parameters through effects that are free of the unit issue. Whether treatment effects are expressed in standard-deviation units or anchored to an adult outcome does not affect the implied marginal effects on long-run outcomes of interest (Cunha et al., 2021; Del Bono et al., 2022; Freyberger, 2026). For the remainder of this section, I compare only the parameters that do not depend on the normalization of the latent stocks, and defer the others to Section 7.3.

Starting with the preference parameters, parents put more weight on consumption than on either parent's leisure. However, the total weight on leisure is larger than the weight on consumption. Fathers' leisure

[44] Del Boca et al. (2014) and Del Boca et al. (2025) report weights on child skill of 0.36 and 0.15, against consumption weights of 0.26 and 0.31. These are not directly comparable to α_P, because their child stock is cognitive skill, which is the argument of the household's value function and hence the outcome it ultimately cares about, whereas α_P prices only the enjoyment of the relationship within flow utility each period.

[45] The calculation is $\alpha_P^{\text{rej}} = (0.122/0.57)/(0.340 + 0.256 + 0.282 + 0.122/0.57) = 0.214/1.092 = 0.196$.

gets more weight than mothers' leisure, at 0.282 against 0.256.[46]

In terms of the investment function, mothers' active child time carries the largest weight at 0.461, then fathers' time at 0.311, and then child expenditures at 0.227. Active parental time does most of the work in the production of child investment. The interaction between the two parents' time is small and slightly positive.

The wage offer depends on each parent's schooling and age. Table 12 reports the estimates. An additional year of schooling raises a father's mean wage offer by about 6.3 percent and a mother's by about 8.5 percent, so the return to education is larger for mothers.[47]

The two age profiles differ, at about 2.8 percent per year of parental age for mothers against 3.9 percent for fathers, so the gap in offers widens as parents age. I caution the reader in interpreting these estimates, for two reasons. First, the wage equation is linear in age; thus, it cannot capture the life-cycle of wages. This is commonly modelled as a quadratic term in parents' ages. Second, the parents observed in my dataset fall within a narrow age range, with means of 40.7 years for mothers and 43.0 for fathers and standard deviations of about five years. Fathers face somewhat more wage risk than mothers, with an offer shock standard deviation of 0.483 against 0.442. The two parents' shocks are strongly correlated, at 0.747.

[46] Flow utility follows Del Boca et al. (2014) in adopting a logarithmic form, with weights that are positive and sum to one. The adding-up condition is a normalization. Since only relative utility matters for the household's choices, it fixes the scale of utility without restricting behaviour. The unit elasticity of substitution among consumption, each parent's leisure, and the bond is a substantive restriction. The policy responses in Section 7.1 should be read as conditional on this restriction.

[47] The ordering is the reverse of Del Boca et al. (2014), who estimate roughly 4.8 percent for mothers against 7.8 percent for fathers.

Table 11: SMM Parameter Estimates: Preferences, Technology & Investment

Parameter	Description	Estimate	SE
Preferences and terminal			
α_P	Bond utility weight	0.122	(0.0012)
α_{ℓ_1}	Mom leisure weight	0.256	(0.0006)
α_{ℓ_2}	Dad leisure weight	0.282	(0.0005)
λ_θ	Skill terminal weight	2.781	(0.0168)
α_c	Consumption weight (implied)	0.340	—
Skill technology			
a	Intercept	-0.161	(0.0038)
δ_θ	Age trend	0.076	(0.0018)
γ_1	Persistence	0.748	(0.0019)
γ_2	Investment	0.041	(0.0005)
γ_3	Bond $\to$ skill	0.179	(0.0013)
γ_4	Skill $\times$ bond interaction	0.024	(0.0003)
$\mathrm{sd}(\varepsilon_{t,\theta})$	Productivity shock sd	0.086	(0.0026)
Bond technology			
$\tilde{a}$	Intercept	0.209	(0.0041)
δ_P	Age trend	0.128	(0.0026)
$\tilde{\gamma}_1$	Persistence	0.701	(0.0015)
$\tilde{\gamma}_2$	Investment	0.088	(0.0009)
$\tilde{\gamma}_3$	Skill $\to$ bond	0.068	(0.0015)
$\tilde{\gamma}_4$	Skill $\times$ bond interaction	0.057	(0.0009)
$\mathrm{sd}(\varepsilon_{t,P})$	Productivity shock sd	0.234	(0.0138)
Investment technology			
α_e	Expenditure elasticity (implied)	0.227	—
α_{m_1}	Mom time elasticity	0.461	(0.0053)
α_{m_2}	Dad time elasticity	0.311	(0.0028)
α_{12}	Mom $\times$ dad time interaction	0.022	(0.0006)

Notes: Estimates are from SMM with a bootstrapped diagonal weighting matrix. Standard errors (in parentheses) are computed using 500 bootstrap replications by household. Implied parameters: $\alpha_c = 1 - \alpha_P - \alpha_{\ell_1} - \alpha_{\ell_2}$ and $\alpha_e = 1 - \alpha_{m_1} - \alpha_{m_2}$.

Table 12: SMM Parameter Estimates: Wage Process

Block	Parameter	Description	Estimate	SE
Wage process				
	b_{01}	Mom log wage intercept	0.587	(0.0040)
	b_{11}	Mom education slope	0.081	(0.0003)
	b_{21}	Mom age slope	0.028	(0.0001)
	b_{02}	Dad log wage intercept	0.564	(0.0034)
	b_{12}	Dad education slope	0.061	(0.0002)
	b_{22}	Dad age slope	0.039	(0.0001)
	σ_{mom}	Mom wage shock sd	0.442	(0.0041)
	σ_{dad}	Dad wage shock sd	0.483	(0.0273)
	ρ_{wage}	Cross-parent wage shock correlation	0.747	(0.0088)

Notes: Estimated jointly with the parameters in Table 11, using the bootstrapped SMM estimator for standard errors (in parentheses). Wage shocks are serially uncorrelated, so no persistence parameters are estimated. Indexing follows $j = 1$ (mother) and $j = 2$ (father).

7 Policy Counterfactuals and Mechanisms

In Section 7.1 I start with policy counterfactuals, where I play as the policymaker. I consider the tax-and-transfer system through two reforms: a Universal Basic Income (UBI) of 24,000 CAD and a more progressive tax schedule that raises taxes on higher earners and lowers them for lower earners, while leaving the median household's average tax rate unchanged.

In Section 7.2 I turn to parenting programs and the mechanisms through which they could work. Gains from a parenting program could come from an exogenous increase in the parent-child bond directly or from changing the household's preferences in how they value the parent-child bond or their child's future skill. I consider four exercises. First, I vary the preference weight on the parent-child bond. Second, I do an analogous exercise by varying the weight parents place on the child's future skill. Third, I consider a boost to the initial parental bond. Lastly, I calibrate the size of each mechanism to the treatment effects of the randomized evaluation of Triple P in Kim et al. (2018), and I ask how large an income transfer would have to be to match each one. Table 14 summarizes the results of every exercise.[48]

7.1 Tax-and-Transfer Policies

Family income shocks affect child outcomes substantially (Carneiro et al., 2021; Agostinelli et al., 2024; Bingley et al., 2026), but existing evaluations of the tax-and-transfer system route this effect entirely through parental time and money. In my model, the parental bond is a second channel. It is a stock of relational capital with its own law of motion. A transfer that relaxes the budget constraint can raise child skill by improving the relational stock (parental bond) through the increase of resources. Following the policy experiments done in Agostinelli et al. (2024) but for the Canadian tax system, I vary the level and progressivity of the tax-and-transfer system and trace the effects on the bond, child skill, and parental time and expenditure.

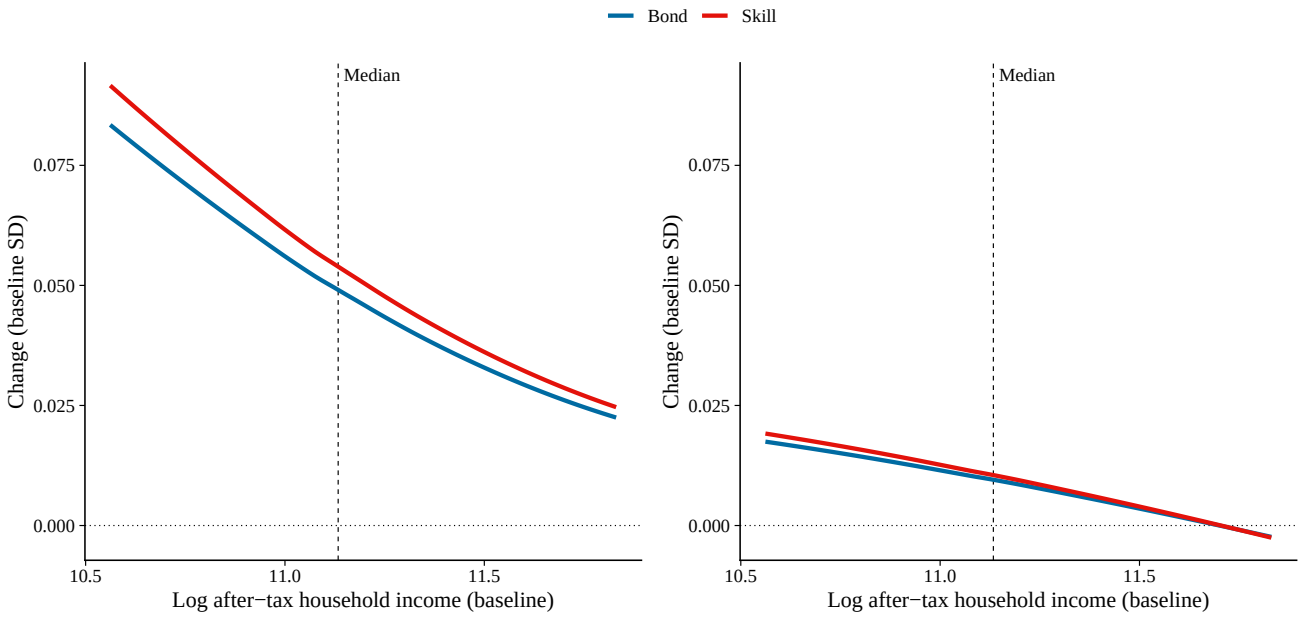
Using equation (5), which provides a summary of the Canadian tax system, I consider two policy exper-

[48] As explained in Section 6, child skill and the parental bond are latent, and their units are set by the normalization of the measurement system; thus, interpretation is awkward and difficult. I follow Del Bono et al. (2022), who show that interpreting treatment effects expressed in standard deviations of the log stocks does not depend on the normalization, since the standard deviation is identified up to the same factor. Thus, I report effects on skill and the bond in standard deviations of the baseline log stock at ages 14–15. Adult income has natural units, so I report it in percent changes.

iments. The first raises the benefit floor by a UBI of 24,000 CAD.[49] The second experiment doubles the progressivity parameter $\mathcal{X}_3$ while readjusting the level parameter $\mathcal{X}_1$. This is done so that the average tax rate faced by the median household remains fixed.[50]

In the baseline economy, the progressivity parameter is 0.122, and the median household (by average annual income) faces an average tax rate of 21.6 percent. In the counterfactual, $\mathcal{X}_3$ is doubled to 0.245 and $\mathcal{X}_1$ is readjusted from 3.058 to 12.06. Appendix I.2 illustrates the mapping from pre-tax to post-tax household income under each of the three tax systems.

Figure 2: Skill and Bond Changes under the Tax-and-Transfer Policies, by Household Income



Notes: Each panel reports the change, relative to the baseline economy, in log child socio-emotional skill and in the log parental bond at the end-line. The left panel is a $24,000 UBI raising the benefit floor. The right panel is a doubling of the progressivity parameter $\mathcal{X}_3$ that holds the median household's average tax rate fixed. The horizontal axis reports log after-tax household income in the baseline economy, and the dashed vertical line marks its median. Lines are non-parametric (loess) smooths.

Figure 2 traces each policy's effect across the household income distribution. In the left panel, the UBI

[49]The amount of 24,000 CAD corresponds to Andrew Yang's Universal Basic Income proposal, which consists of a $1,000 monthly lump-sum transfer to every American adult over the age of 18.

[50]I also do a third policy, Appendix H, which readjusts $\mathcal{X}_1$ to hold the median household's marginal tax rate fixed instead. This version raises the median household's after-tax income, so skill and the bond rise by more at every income level shown. The gains still fall with income and skill still moves more than the bond, but the size of the effects depends on which tax rate is held fixed.

raises both child skill and the parental bond throughout the distribution, with the largest gains at the bottom. The gains are roughly 3.7 times larger at the bottom of the income distribution than at the top, and about 1.7 times larger in the bottom quartile than in the top quartile (Table 15). An important finding is that skill is more elastic to the UBI transfer than the bond is, moving about 1.1 times as much at every income level. However, as we move across the distribution, the gap between the two effects narrows.

The right panel shows the progressivity policy experiment, where $\mathcal{X}_3$ is doubled while the median household's average tax rate is held fixed. Like the UBI, this reform raises both child skill and the parental bond throughout the distribution. Both outcomes decline as income rises, and the decline gets steeper toward the top rather than flattening out, so the gains are close to zero at the top.

This happens because $\mathcal{X}_1$ is readjusted to hold the median household's tax rate fixed, so the added progressivity falls almost entirely on above-median households. The mechanism here differs from the UBI panel. There, poor households gained resources with no change in their incentive to work, an income effect. Here, rich households lose more in after-tax income and face worse marginal trade-offs between working, spending time with children, and leisure. The income and substitution effects work together.

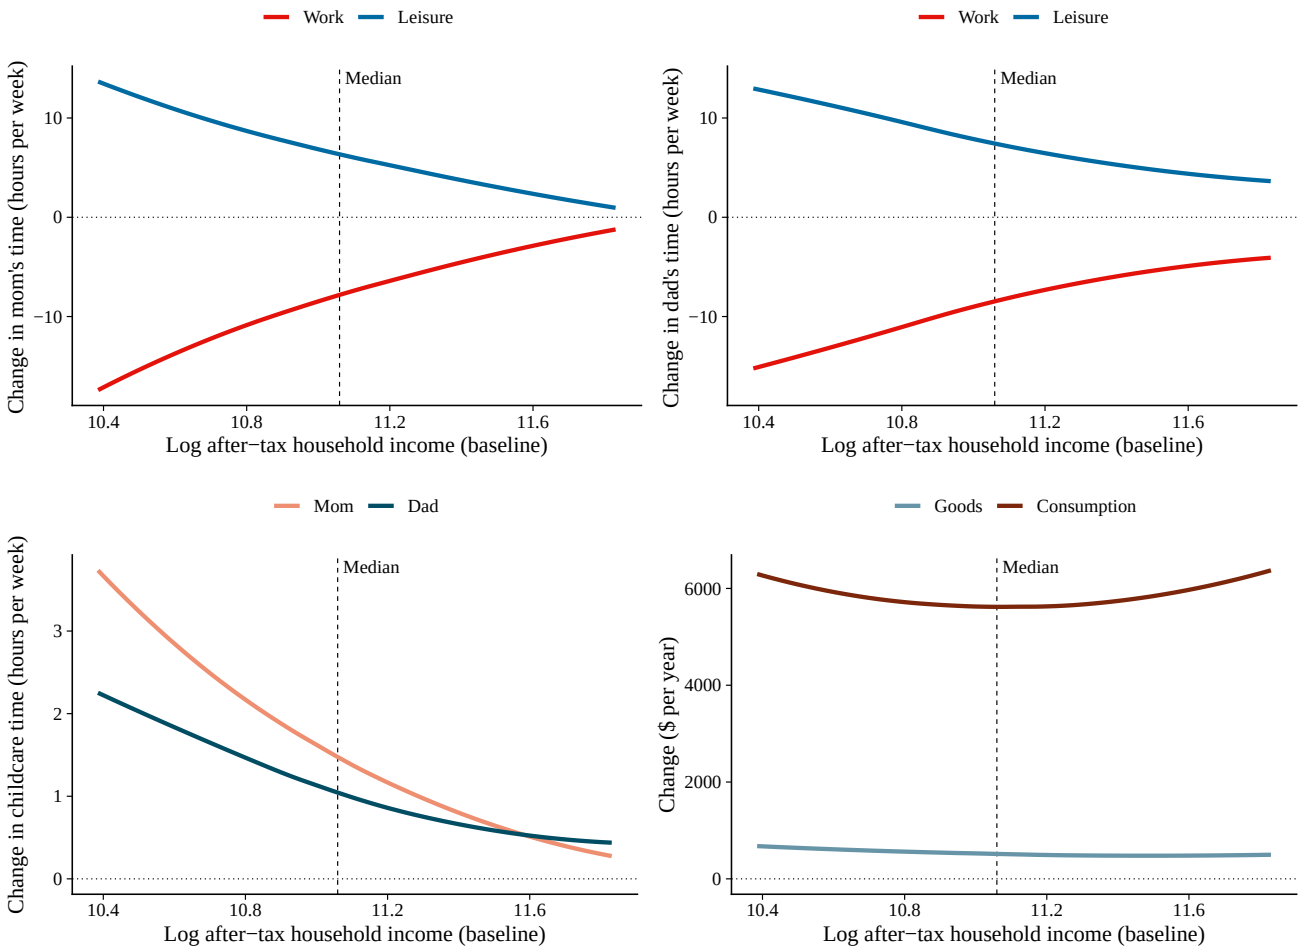
Figures 3 and 4 show the reallocation of parental time and household resources under each policy. Under the UBI, both parents reduce their market hours and reallocate their time towards childcare and leisure, with the largest reallocation among low-income households. For mothers, the reduction in market hours and the increase in active time and leisure fall steadily as income rises. For fathers, the reduction in market hours increases rapidly at the bottom of the distribution and diminishes as income rises. On the leisure and active time side, the increase falls off more quickly as income rises and flattens out at the top of the distribution. For goods and consumption, the UBI raises both, and consumption rises far more than goods expenditure across the whole distribution.

Putting these results together, the parental bond increase under the UBI for low-income households is driven mainly by the time channel. Both parents work less, and the freed time is reallocated to childcare and leisure. This is consistent with the fact that active time with the child produces most of the latent investment. The money channel only moves a few hundred dollars and moves far less than consumption. The UBI buys time for the children by relaxing the labour-supply constraint.

For the progressive tax reform, both parents reduce their market hours and reallocate their time towards childcare and leisure. An asymmetry emerges between the two parents as we move along the income distribution. At the top of the income distribution, mothers rapidly increase their market hours (although it

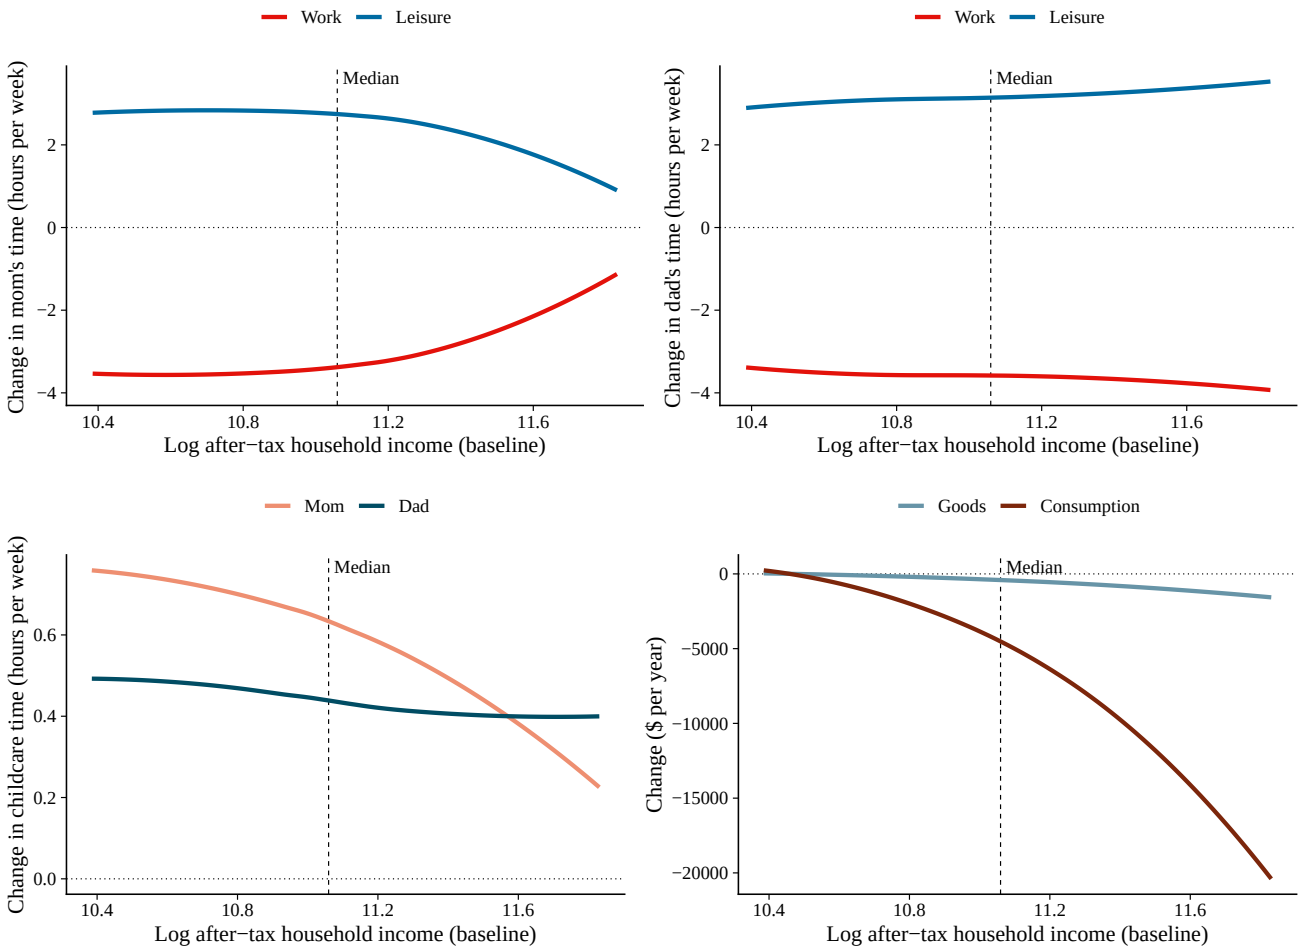
remains negative) and reduce their active time and leisure. Fathers, on the other hand, reduce their market hours and increase their active time and leisure by about the same amount throughout the distribution. For goods and consumption, the progressive tax reform reduces both as income rises, but consumption falls far more than goods expenditure.

Figure 3: Parental Time and Household Resources under the UBI, by Household Income



Notes: Each panel reports the change, relative to the baseline economy, in the mother's work and leisure (top-left panel), the father's work and leisure (top-right panel), the mother's and father's childcare time (bottom-left panel), and household spending on child goods and on consumption (bottom-right panel), under a $24,000 UBI raising the benefit floor. Sample means for mothers (fathers) are 15.2 (7.8) hours of childcare, 27.0 (41.1) of work and 69.9 (63.1) of leisure, and $2,543 of child expenditure. The horizontal axis reports log after-tax household income in the baseline economy. Lines are non-parametric (loess) smooths.

Figure 4: Parental Time and Household Resources under the Progressive Tax Reform, by Household Income



Notes: Each panel reports the change, relative to the baseline economy, in the mother's work and leisure (top-left panel), the father's work and leisure (top-right panel), the mother's and father's childcare time (bottom-left panel), and household spending on child goods and on consumption (bottom-right panel), under a doubling of the progressivity parameter $\mathcal{X}_3$ that holds the median household's average tax rate fixed. Sample means for mothers (fathers) are 15.2 (7.8) hours of childcare, 27.0 (41.1) of work and 69.9 (63.1) of leisure, and $2,543 of child expenditure. The horizontal axis reports log after-tax household income in the baseline economy. Lines are non-parametric (loess) smooths.

7.2 Parenting Program Policies

A parenting program works on the relationship instead. The Strengthening Families Program, for children aged 10 to 14, is one example. Sixth graders and their families attend a seven-week course on parent-child communication, supervision, and warmth. Another well-known parenting program is the Triple P (Positive Parenting Program). The program teaches parents consistent discipline, warmth, and realistic expectations. Both programs teach what maps closely onto the components of the bond I measure in the NLSCY.

Kim et al. (2018) evaluate a randomized trial of its group version in which parents of children aged 2 to 6 attend four weekly two-hour sessions and children do not take part. Parents receive no money and give up little time. The program teaches parents how to properly discipline and engage with their child. They find that harsh discipline falls and positive engagement rises for at least four years, and that at ages 12 to 16 the children of treated parents report externalizing behaviour lower by 0.31 standard deviations and wellbeing higher by 0.2 standard deviations.

In the following section, I explore the mechanisms through which such a program could work. It could raise the bond directly, change what parents value, either the relationship or their child's future skill, or do both.

7.2.1 The relationship motive

The parenting program described above could change how much parents value the relationship itself by teaching parents the importance of warmth and consistency in their interactions with their child. I capture this by raising the flow weight on the bond, α_P, by 50 percent.[51] Recall that parents invest in the relationship for two reasons: they enjoy it every period (flow utility $\alpha_P \ln P_t$), and a stronger bond raises subsequent skill formation through the cross-stock channel.

Table 14 shows that this raises terminal child skill by 0.050 standard deviations, the parental bond by 0.045 standard deviations and adult income by 0.89 percent. The household raises child expenditure and active time with the child by about 10 percent. The added time comes mostly out of leisure. Maternal leisure falls by 1.5 percent and paternal leisure by 1.4 percent. Mothers also reduce market work by 1.5

[51] The normalization $\alpha_c + \alpha_P + \alpha_{\ell_1} + \alpha_{\ell_2} = 1$ rules out changing α_P on its own, so α_c, α_{ℓ_1} and α_{ℓ_2} are rescaled in equal proportion, leaving their weights unchanged relative to one another. All technology parameters are held at their baseline values.

percent, while fathers leave their market hours essentially unchanged. The household income falls by 0.40 percent. Household consumption falls by 1.3 percent because the household earns less and spends more on the child.

For contrast, I also lower the weight. Halving α_P lowers skill by 0.050 standard deviations and the bond by 0.045 standard deviations, so the response is close to symmetric. At the extreme, suppose we had a completely selfish parent, who obtains no direct utility from the bond but still values the child's future skill. This exercise quantifies how much of parental investment reflects the relationship motive, meaning the direct utility parents obtain from the bond, as opposed to a pure investment motive. I simulate a counterfactual in which $\alpha_P = 0$. Under the counterfactual, the objective reduces to

$$u_t = \alpha_c \ln c_t + \alpha_{\ell_1} \ln \ell_{1t} + \alpha_{\ell_2} \ln \ell_{2t}, \qquad V(\theta_{T+1}) = \lambda_\theta \ln \theta_{T+1},$$

so parents' only incentive to build the bond is because it increases skill, acting as a pure investment motive.

Average child skill falls by 0.100 standard deviations and the parental bond by 0.090 standard deviations relative to the baseline. The losses in skill and the bond are small because both are persistent stocks and parents still build the bond for its return through the cross-productivity of the relationship. When $\alpha_P = 0$, the household reduces child expenditure and active time with the child by roughly 17 percent. The fall in active time and expenditure therefore measures the share of baseline investment that comes from enjoying the relationship itself.

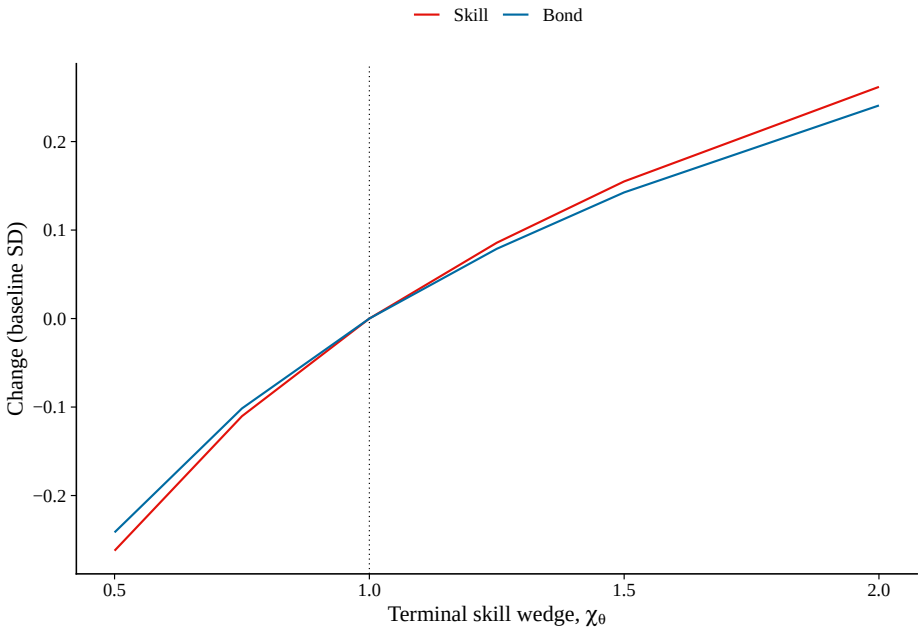
7.2.2 The weight on the child's future success

Another mechanism that the parenting program could work through is changing how much parents value their child's future skill. The program could teach parents the importance of socio-emotional skill for their child's future success. I vary the weight parents place on their child's future skill and trace the effect of such a change. I introduce a wedge χ_θ that scales the estimated terminal skill coefficient, holding all else fixed. The terminal value function becomes:

$$V(\theta_{T+1}) = \chi_\theta \, \lambda_\theta \ln \theta_{T+1}, \qquad \chi_\theta = 1 \text{ at baseline.}$$

Figure 5 shows how the two stocks, at end-line, respond to how parents value the child's future skill. Doubling the terminal skill wedge from its baseline value of one raises mean log skill by about 26 percent

Figure 5: Sensitivity of Child Skill and Parental Bond to the Terminal Skill Wedge χ_θ



Notes: Each line reports the change in the terminal-age (14–15) mean of the log of child skill (θ) or of the parental bond (P) relative to the baseline economy, as the terminal skill wedge χ_θ is varied.

of a standard deviation and mean log bond by about 24 percent of a standard deviation, and halving the wedge lowers the two stocks by about 26 percent and 24 percent of a standard deviation.

The response is also concave in χ_θ, where halving the wedge reduces both stocks by more than raising it by half increases them. In other words, caring less costs more than caring more gains. Furthermore, skill responds slightly more than the bond in both directions: giving more weight to the child's future socio-emotional skill moves skill and the parent-child relationship by similar amounts.

Figure 6 shows the input side of this response. Doubling the terminal skill wedge raises maternal and paternal childcare time, by approximately 11 hours per week, compared to a sample mean of 15 hours for mothers, and by about 6.6 hours, compared to a sample mean of 8 hours for fathers.

An interesting pattern emerges in how parents reallocate their time between work and leisure. For mothers, as we move along the wedge (parents value the child's future skill more highly), about 70 percent of every hour added to childcare is coming out of leisure, with the remaining 30 percent coming from market work. So mothers are willing to give up leisure to invest in their children, with market work as a secondary lever.

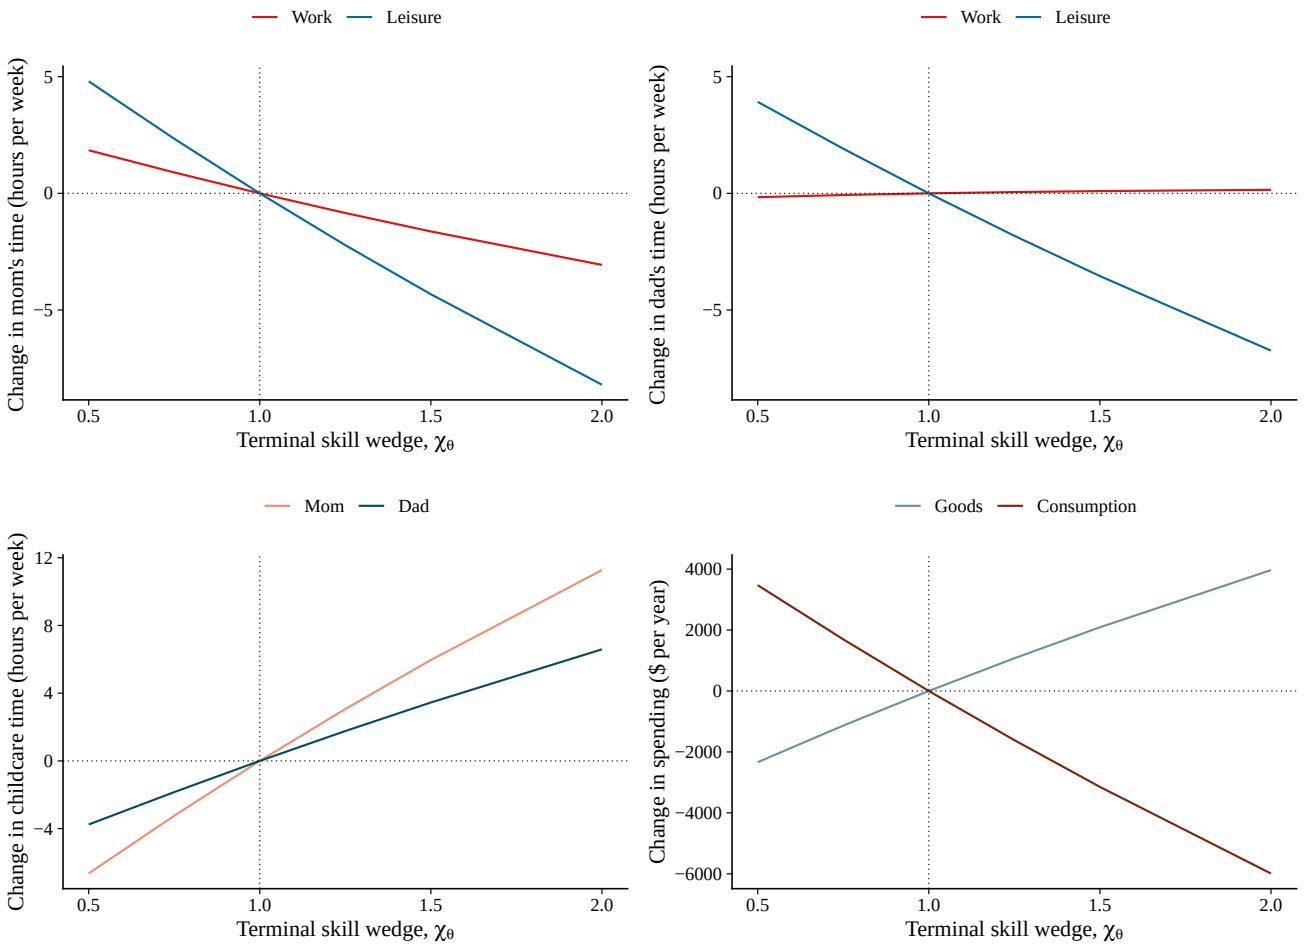
Fathers, on the other hand, follow the same pattern as mothers with the exception that they do not reduce

their market hours. All additional childcare time comes out of leisure, so fathers are willing to give up leisure to invest in their children, but they do not reduce their market work. The model therefore implies that mothers are more flexible in their labour supply than fathers.

Compared to the UBI in Section 7.1, this pattern is reversed. Under the UBI, dad's work and leisure reallocation is larger than mom's. Put together, dad's labour supply responses are elastic to income shocks, but inelastic to the value placed on the child's future skill. The flexible margin for dads' labour supply is income based, not returns-to-skill based.

Finally, I consider that this parenting program could change both how much parents value the relationship and how much they value their child's future skill. I raise both α_P and χ_θ by 50 percent where the results are shown in Table 14. The combined effect is smaller than the sum of the two separate effects, which is consistent with the concavity in χ_θ shown in Figure 5. The combined effect raises terminal skill by 0.197 standard deviations, the parental bond by 0.181 standard deviations, and adult income by 3.5 percent.

Figure 6: Parental Time and Household Resources as the Terminal Skill Wedge χ_θ Varies



Notes: Each panel reports the change, relative to the baseline economy ($\chi_\theta = 1$, dotted vertical line), in the simulated mean of the mother's work and leisure (top-left panel), the father's work and leisure (top-right panel), the mother's and father's childcare time (bottom-left panel), and household spending on child goods and on consumption (bottom-right panel). Sample means for mothers (fathers) are 15.2 (7.8) hours of childcare, 27.0 (41.1) of work and 69.9 (63.1) of leisure, and $2,543 of child expenditure.

7.2.3 Exogenous increase to the initial skill and bond

Lastly, the parenting program could also deliver a direct boost to the parent-child relationship. In this exercise, I perform a boost to the initial stocks of socio-emotional skills to compare the effects of a shock to the bond versus a shock to skill. I raise the child's initial skill θ_0 or the initial parental bond P_0 by one standard deviation of its baseline distribution at ages 10–11.[52]

Table 14 shows that a one standard deviation higher P_0 raises the terminal bond by 0.42 standard deviations and, through the cross-productivity, terminal skill by 0.425. A one standard deviation higher θ_0 also raises terminal skill by 0.41 standard deviations, so a shock to the bond does as much for terminal skill as a shock of the same size to skill itself. Panel A converts these changes in terminal skill into predicted adult income, measured as annual income between ages 25 and 36. A higher θ_0 raises adult income by 7.42 percent and a higher P_0 by 7.73 percent.

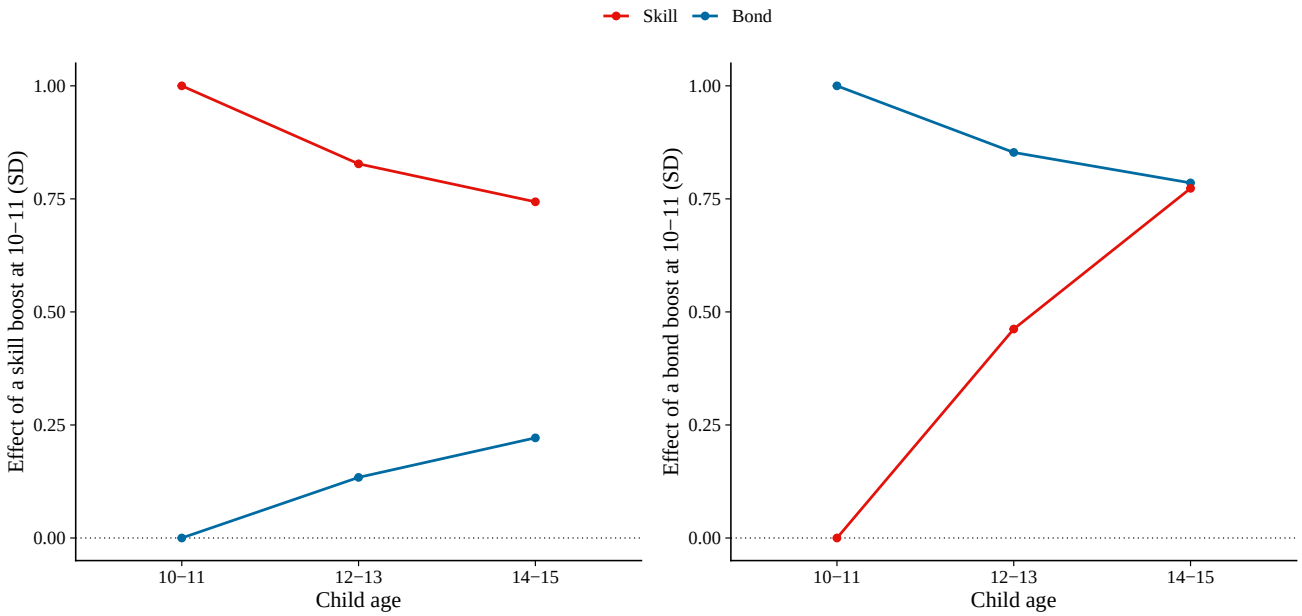
The two shocks differ in the cross-productivity they generate. Raising skill moves the bond by 0.118 standard deviations, while raising the bond moves terminal skill by 0.425. The cross-productivity therefore runs more strongly from the bond to skill than from skill to the bond, as in the reduced-form evidence of Table 6.

Figure 7 shows the trajectories of the endowment shocks of Table 14 over time. Each child receives a one standard deviation boost to skill (left panel) or to the bond (right panel) at ages 10–11. The figure reports how much of that boost remains at later ages, so that 1 is the full boost. The bond boost fades slowly. At ages 14–15, 0.79 of it remains, and it keeps raising skill. About three quarters of the skill boost remains at ages 14–15, and it spills over to the bond from the cross-productivity and complementarity. These findings suggest that a program that strengthens the parent-child relationship raises terminal skill and adult income as much as a program that targets the child's skill directly.[53]

[52] I use one standard deviation changes in the initial stocks because they are latent and their units are set by the normalization of the measurement system, so a percentage change in either makes interpretation difficult (Del Bono et al., 2022).

[53] These results leave out the welfare gains from the higher bond and skill. Raising the bond at ages 10–11 benefits parents in two ways. First, parents enjoy the relationship more, which enters through the flow utility. Second, because the bond raises skill, they also receive a higher terminal payoff from the child's skill, while raising the child's skill directly gives parents smaller welfare gains outside of the terminal payoff from skill due to the weaker cross-productivity.

Figure 7: How a Boost to Skill or the Bond at Ages 10–11 Carries Forward



Notes: Each line reports the change in the mean log stock relative to the baseline after a one standard deviation increase in initial log skill (left panel) or the initial log bond (right panel) at ages 10–11, the endowment shocks of Table 14. Changes are in standard deviations of the log stock at ages 10–11 (0.102 for skill and 0.193 for the bond). Red lines are skill and blue lines the bond.

7.2.4 A Triple P-Sized Program and Cash Equivalents

I calibrate the parenting program to the randomized evaluation of Triple P in Kim et al. (2018). Four years after the program, parent-reported harsh discipline is about 0.2 to 0.3 standard deviations lower and positive engagement about 0.2 to 0.3 higher. I therefore apply a shock at ages 10–11 that leaves the bond 0.2 to 0.4 standard deviations higher four years later, at ages 14–15. Panel A of Table 13 shows that socio-emotional skill then rises by 0.20 to 0.41 standard deviations. Ten years after their program, at ages 12 to 16, the intent-to-treat effects in standard deviations of their sample are a fall of 0.12 in total negative behaviours, a fall of 0.31 in externalizing behaviours, and essentially no change in internalizing behaviours. The comparison I do in this exercise is only suggestive, since their program starts at a younger age and the change in parenting is reported by the parents who took the training.

Table 13 also asks how large an income transfer, paid to every household each year from ages 10 to 15, would have to be to match each experiment. A 50 percent higher bond weight takes roughly $22,000 per year, a 50 percent higher skill weight is equivalent to $87,000, and both together about $129,000. Matching the skill gain of the smallest Triple P-sized program, which raises the bond by 0.2 standard deviations at ages 14–15, takes about $136,000, and the two larger programs exceed the largest transfer I simulate. Panel B sizes the preference levers to the same 0.2 standard deviation gain in the bond. This takes a 213 percent higher bond weight, a 77 percent higher skill weight, or a 57 percent increase in both. Each raises skill by about 0.22 standard deviations, slightly more than the bond shock. These comparisons say how large a transfer would have to be, not what a program would cost.

Table 13: A Triple P-Sized Bond Shock, the Preference Levers and Their Income Transfer Equivalents

Experiment	Bond (SD)	Skill (SD)	Adult Income	Time (hr/wk) Mother	Time (hr/wk) Father	Spending ($/yr)	Income Transfer, Same Skill Gain All	Income Transfer, Same Skill Gain Q1	Income Transfer, Same Skill Gain Q4
Panel A: Bond Shock Sized to the Change in Parenting in Kim et al. (2018)									
Bond shock of 0.48 SD at ages 10–11	0.200	0.203	3.62%	0.04	0.02	13	136.1	101.9	185.2
Bond shock of 0.72 SD at ages 10–11	0.300	0.305	5.49%	0.06	0.03	19	> 288	262.7	> 288
Bond shock of 0.96 SD at ages 10–11	0.400	0.407	7.38%	0.07	0.04	25	> 288	> 288	> 288
Panel B: Preference Levers Sized to the Same Bond Gain, and a Larger Bond Shock									
Bond weight $\alpha_P \uparrow 213\%$	0.200	0.221	3.94%	7.82	4.57	2,716	159.7	120.9	219.8
Terminal skill wedge $\chi_\theta \uparrow 77\%$	0.200	0.217	3.88%	8.93	5.20	3,137	154.8	117.4	212.9
Both, α_P and $\chi_\theta \uparrow 57\%$	0.200	0.218	3.89%	8.71	5.07	3,054	155.9	118.2	214.4

Notes: Bond and skill report the change in the mean log stock at ages 14–15, in standard deviations of its baseline distribution. Panel A leaves the bond 0.2, 0.3 and 0.4 standard deviations higher at ages 14–15. Panel B raises α_P, χ_θ, or both by the same percentage, until the bond is 0.2 standard deviations higher at ages 14–15, as in the first row of Panel A. Time and spending report the change in the mean of the mother's and father's active time with the child, in hours per week, and of child expenditure, in dollars per year, over ages 10–15.Income transfer columns report the transfer to every household, in thousands of dollars per year at every age from 10 to 15, that gives the same skill gain, by interpolation between simulated transfers from $1,500 to $288,000. Q1 and Q4 are the bottom and top quartiles of baseline household income.

7.3 Comparing the Counterfactuals

This section compares the counterfactuals above, complementing Section 6. Table 14 reports outcomes, parental labour market, and investment inputs for each counterfactual. Table 15 reports the same results by quartile of baseline household income for only the tax-and-transfer experiments.

In Section 7.2.1, I separate the direct utility parents obtain from the relationship from its return as an investment, setting $\alpha_P = 0$ to isolate the investment motive. The motivation for this is to understand the role of household preferences in shaping the parent-child relationship and child socio-emotional skill. My results show that about a sixth of parental investment comes from enjoyment of the relationship rather than from the expected payoff from the child's end-line.

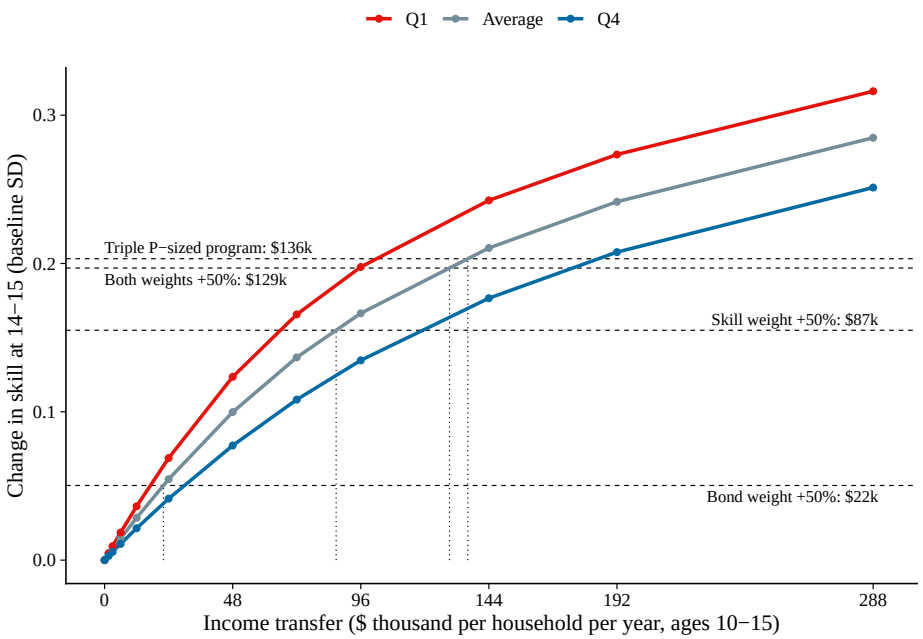
A counterpart experiment raises the weight parents place on the child's future skill by 50 percent. This raises the child's skill (bond) at ages 14–15 by 0.155 (0.143) standard deviations and the child's anchored adult income by 2.75 percent. Raising the weight on the relationship by the same 50 percent moves all three outcomes less than the first experiment. However, for the weight on skill, a 50 percent decrease lowers skill and the bond by more than a 50 percent increase raises them, whereas for the weight on the relationship the two are symmetric.

Turning to the technology parameters, I raise the initial skill stock by one standard deviation. This

raises the child's skill at ages 14–15 by 0.409 standard deviations, the end-line bond by 0.118 standard deviations, and anchored adult income by 7.42 percent. In the same exercise, raising the initial bond by one standard deviation has a similar effect on skill and adult income, at 0.425 standard deviations and 7.73 percent, but a much larger effect on the end-line bond, at 0.418 standard deviations. The cross-productivity of the bond thus raises skill and adult income even when the initial skill stock is unchanged, whereas the cross-productivity of initial skill on the bond is smaller.

Lastly, a UBI of $24,000 raises the child's skill at ages 14–15 by 0.055 standard deviations, the end-line bond by 0.050 standard deviations, and anchored adult income by 0.96 percent, far less than raising the initial stock of either skill or bond. Parents direct freed resources mostly to consumption and leisure. Only a little goes to the investment inputs. Figure 8 plots the experiments considered on a common scale. I ask how large an income transfer would have to be to raise skill by the same amount. The skill gain from a transfer flattens with its size because parents direct most of the added resources to consumption and leisure.

Figure 8: Income Transfer Needed to Match the Skill Gain of an Experiment



Notes: The solid lines report the gain in skill at ages 14–15, in standard deviations of its baseline distribution, from an income transfer of the size on the horizontal axis, paid at every age from 10 to 15, on average and in the bottom (Q1) and top (Q4) quartiles of baseline household income. Dashed lines mark the gains of four experiments, the Triple P-sized program in the first row of Panel A of Table 13 and the 50 percent increases in α_P, χ_θ, or both in Table 14, and dotted lines the transfer that gives the same gain on average. The gains of these experiments are about the same in every quartile.

Two mechanisms drive these results. First, parents gain utility from the bond as well as from its return to skill, so it is not just an investment good. Changing how parents value the relationship with their child or skill largely determines how much they invest in the relationship and in skill. Second, the bond is more productive for skill than skill is for the bond. Policies such as a parenting program that could directly raise the bond, or that could change how much parents value it, raise skill and adult income. Welfare of the household also sees the largest gains from this type of policy due to the high cross-productivity of the bond.

Table 14: Results of Experiments Across Counterfactuals

	Relationship Motive			Terminal Skill Wedge		Both Weights	Initial Endowments		Tax-and-Transfer	
	Investment Only (α_P=0)	$\alpha_P \downarrow 50\%$	$\alpha_P \uparrow 50\%$	$\chi_\theta \downarrow 50\%$	$\chi_\theta \uparrow 50\%$	$\alpha_P, \chi_\theta \uparrow 50\%$	$\theta_0 \uparrow 1$ SD	$P_0 \uparrow 1$ SD	UBI \$24k	$\mathcal{X}_3$ Doubled
Panel A: Terminal Child Skill, Parental Bond, and Adult Income										
Child Skill (θ_{T+1}, SD)	-0.100	-0.050	0.050	-0.262	0.155	0.197	0.409	0.425	0.055	0.010
Parental Bond (P_{T+1}, SD)	-0.090	-0.045	0.045	-0.242	0.143	0.181	0.118	0.418	0.050	0.009
Adult Income	-1.73%	-0.87%	0.89%	-4.49%	2.75%	3.51%	7.42%	7.73%	0.96%	0.18%
Panel B: Parental Labour Market										
Mother Hours	2.53%	1.33%	-1.48%	6.68%	-5.90%	-7.60%	-0.05%	-0.08%	-28.68%	-11.98%
Father Hours	-0.13%	-0.07%	0.07%	-0.41%	0.24%	0.29%	0.00%	0.00%	-21.34%	-8.91%
Household Income	0.67%	0.35%	-0.40%	1.75%	-1.63%	-2.11%	-0.01%	-0.02%	9.53%	-8.18%
Panel C: Child Expenditure and Active Time										
Child Expenditure	-16.31%	-8.58%	9.59%	-43.62%	39.00%	50.26%	0.33%	0.49%	9.68%	-8.06%
Mother Active Time	-17.03%	-8.96%	10.01%	-44.83%	40.16%	51.80%	0.35%	0.53%	10.20%	4.21%
Father Active Time	-17.74%	-9.36%	10.52%	-46.05%	42.29%	54.79%	0.37%	0.56%	13.13%	5.38%

Child skill and parental bond entries report the change in the mean log stock at ages 14–15 in standard deviations of its baseline distribution; all other entries report the percentage change from the estimated baseline. The terminal skill wedge columns set χ_θ to 0.5 and 1.5, scaling the estimated terminal skill coefficient λ_θ with all else held fixed. The both weights column applies the 50 percent increases in α_P and λ_θ together. The initial endowment columns raise the log of θ_0 or P_0 at ages 10–11 by one standard deviation of its baseline distribution. The predicted percentage change in adult income reflects the change in mean predicted annual income observed between ages 25 and 36, by the estimated relationship between child skill at ages 14–15 and later-life earnings in column [2], Panel A, of Table 5, taken per unit of $\ln \theta$ (0.943) rather than per standard deviation.

Table 15: Results of the Tax-and-Transfer Experiments by Baseline Household Income Quartile

	UBI $24k				$\mathcal{X}_3$ Doubled			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Panel A: Terminal Child Skill, Parental Bond, and Adult Income								
Child Skill (θ_{T+1}, SD)	0.069	0.057	0.051	0.041	0.014	0.012	0.010	0.006
Parental Bond (P_{T+1}, SD)	0.063	0.052	0.046	0.038	0.013	0.010	0.009	0.006
Adult Income	1.21%	1.01%	0.89%	0.73%	0.25%	0.20%	0.17%	0.11%
Panel B: Parental Labour Market								
Mother Hours	−32.55%	−28.91%	−26.59%	−25.27%	−10.93%	−11.72%	−12.14%	−13.71%
Mother Employment Rate	−11.64%	−11.54%	−11.36%	−12.74%	−3.05%	−4.07%	−4.64%	−6.60%
Father Hours	−29.01%	−22.96%	−19.71%	−15.10%	−9.74%	−9.11%	−8.82%	−8.12%
Father Employment Rate	−8.79%	−5.32%	−3.96%	−1.98%	−1.86%	−1.42%	−1.31%	−0.78%
Household Income	12.17%	10.20%	9.09%	7.71%	−5.14%	−7.16%	−8.51%	−10.61%
Panel C: Child Expenditure and Active Time								
Child Expenditure	12.36%	10.36%	9.22%	7.83%	−5.02%	−7.04%	−8.39%	−10.50%
Mother Active Time	14.22%	11.03%	9.35%	6.69%	4.71%	4.40%	4.20%	3.58%
Father Active Time	15.63%	13.61%	12.24%	10.80%	5.15%	5.31%	5.39%	5.71%

All entries report the change from the estimated baseline within each quartile of baseline household income (Q1 is the lowest), in standard deviations of the pooled baseline distribution for child skill and the parental bond and in percent otherwise. Households are assigned to quartiles by their average baseline after-tax household income across age bins, and membership is held fixed across scenarios, so that each entry compares the same households before and after the reform. The employment rate is the share of mother-periods (father-periods) with strictly positive market hours, pooled across age bins, so that its entry is the percentage change in that share, not a change in percentage points. The remaining definitions are as in Table 14.

8 Conclusion

I estimate a dynamic model of household decision-making in which the parent-child bond, the quality of the relationship linking parents to their child, is an endogenous outcome of household optimization.

The first fact about the bond is descriptive. The warmth, consistency, and knowledge of the child that the bond captures in adolescence predict the child's income in adulthood, even after controlling for household characteristics. The predictive power disappears once the child's own socio-emotional skill and cognition proxy is included. This is consistent with the finding that the bond works through skill rather than as a separate determinant of income.

Empirically, I estimate translog production functions for the child's socio-emotional skill and the parent–child bond. I show that most of what investment does for skill is achieved indirectly by strengthening the parental bond. The structural estimates follow the same story as the empirical findings. Both skill and the bond are persistent, and the bond raises skill by much more than skill raises the bond, so a stronger bond leads to higher skill later on while higher skill does less to strengthen the bond.

Using the estimated model, I conduct policy experiments to understand how parents' behaviour and decisions can build the child's socio-emotional skill and the parent–child bond. Tax-and-transfer reforms have limited impacts on skill and the bond. A UBI raises terminal skill and the bond, and with them adult income. The gains are largest in the bottom quartile of the income distribution, but they are small throughout. The UBI result mirrors Del Boca et al. (2014), who find that cash transfers to households with children have small effects on child quality. They show that this is due to the money investments being less productive than parental time and to the share of the transfer that goes to household consumption and parental leisure.

I next show the mechanism through which a parenting program like Triple P could raise the bond and skill. My results suggest that policies which change how parents value their child's skill or the relationship, or which strengthen the bond directly, could raise socio-emotional skill by more than policies that raise family income.

Future work calls for longitudinal surveys of children to collect measures of the parent-child relationship, as these are important for understanding how socio-emotional skill is formed. In addition, this research suggests that interventions that strengthen this relationship, such as parenting programs, should

be evaluated for their cost and effectiveness in raising children's socio-emotional skill and improving their long-term outcomes. This model does not include the roles of peers, schools, and neighbourhoods in the formation of socio-emotional skill. Future work could extend the model to include these factors.

References

Agostinelli, Francesco, Domenico Ferraro, Xincheng Qiu, and Giuseppe Sorrenti (2024) 'Intra-Household Insurance and the Intergenerational Transmission of Income Risk.' Technical Report w32096, National Bureau of Economic Research, Cambridge, MA, January

Agostinelli, Francesco, Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti (2025a) 'It Takes a Village: The Economics of Parenting with Neighborhood and Peer Effects.' *Journal of Political Economy* p. 739334

_ (2025b) 'A stairway to success: How parenting shapes culture and social stratification.' Technical Report, National Bureau of Economic Research

Akee, Randall, William Copeland, E. Jane Costello, and Emilia Simeonova (2018) 'How Does Household Income Affect Child Personality Traits and Behaviors?' *American Economic Review* 108(3), 775–827

Arim, Rúbab G., Jennifer D. Shapka, V. Susan Dahinten, and Brent F. Olson (2011) 'Evaluation of the factor structure of the child-reported parenting questionnaire in the National Longitudinal Survey of Children and Youth.' *Health Reports* 22(1), 55

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

Attanasio, Orazio, Costas Meghir, Emily Nix, and Francesca Salvati (2017) 'Human capital growth and poverty: Evidence from Ethiopia and Peru.' *Review of economic dynamics* 25, 234–259

Attanasio, Orazio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina (2020b) 'Estimating the production function for human capital: Results from a randomized controlled trial in Colombia.' *American Economic Review* 110(1), 48–85

Baker, Michael, Jonathan Gruber, and Kevin Milligan (2008) 'Universal child care, maternal labor supply, and family well-being.' *Journal of Political Economy* 116(4), 709–745

Barszczewski, Jacek, and Katherina Thomas (2025) 'Parenting Style and Children's Skill Development.' *Working Paper*

Baumrind, D. (1967) 'Child care practices anteceding three patterns of preschool behavior.' *Genetic psychology monographs* 75(1), 43–88

Baumrind, Diana (1966) 'Effects of authoritative parental control on child behavior.' *Child development* 37(4), 887–907

Bingley, Paul, Lorenzo Cappellari, and Marco Ovidi (2026) 'When It Hurts the Most: Timing of Parental Job Loss and a Child's Education,.' *Journal of the European Economic Association*

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

Bolt, Uta, Eric French, Jamie Hentall-MacCuish, and Cormac O'Dea (2024) 'The intergenerational elasticity of earnings: Exploring the mechanisms.' Technical Report, IFS Working Papers

Bowlby, J. (1969) *Attachment and Loss* Attachment and Loss (Pimlico)

Bruno, August (2025) 'Parenting Decisions and Child Skill Development.' *Working Paper*

Carneiro, Pedro, Italo López García, Kjell G. Salvanes, and Emma Tominey (2021) 'Intergenerational Mobility and the Timing of Parental Income.' *Journal of Political Economy* 129(3), 757–788

Caucutt, Elizabeth, Lance John Lochner, Joseph Mullins, and Youngmin Park (2025) 'Child Skill Production: Accounting for Parental and Market-based Time and Goods Investments.' *Journal of Political Economy* p. 738480

Cobb-Clark, Deborah A, Nicolas Salamanca, and Anna Zhu (2019) 'Parenting style as an investment in human development.' *Journal of Population Economics* 32(4), 1315–1352

Conger, Rand D., Katherine J. Conger, and Monica J. Martin (2010) 'Socioeconomic Status, Family Processes, and Individual Development.' *Journal of Marriage and the Family* 72(3), 685–704

Conger, Rand D., Xiaojia Ge, Glen H. Elder, Frederick O. Lorenz, and Ronald L. Simons (1994) 'Economic Stress, Coercive Family Process, and Developmental Problems of Adolescents.' *Child Development* 65(2), 541

Croon, Marcel A. (2002) 'Using predicted latent scores in general latent structure models.' In *Latent Variable and Latent Structure Models,* ed. George A. Marcoulides and Irini Moustaki (Mahwah, NJ: Lawrence Erlbaum Associates, Inc.) pp. 195–224

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

Cunha, Flavio, and James J. Heckman (2008) 'Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation.' *Journal of Human Resources* 43(4), 738–782

Cunha, Flavio, Eric Nielsen, and Benjamin Williams (2021) 'The Econometrics of Early Childhood Human Capital and Investments.' *Annual Review of Economics* 13(1), 487–513

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

Darling, Nancy, and Laurence Steinberg (1993) 'Parenting style as context: An integrative model.' *Psychological Bulletin* 113(3), 487–496

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall (2014) 'Household choices and child development.' *Review of Economic Studies* 81(1), 137–185

Del Boca, Daniela, Christopher Flinn, Ewout Verriest, and Matthew Wiswall (2025) 'Parenting with Patience: Parental Incentives and Child Development.' *Journal of Political Economy* p. 738481

Del Bono, Emilia, Ben Etheridge, and Paul Garcia (2026) 'The economic value of childhood socio-economic skills.' Technical Report, The IFS, July

Del Bono, Emilia, Josh Kinsler, and Ronni Pavan (2022) 'Identification of dynamic latent factor models of skill formation with translog production.' *Journal of Applied Econometrics* 37(6), 1256–1265

_ (2024) 'Skill Formation and the Trouble with Child Non-Cognitive Skill Measures.' *SSRN Electronic Journal*

Deming, David J. (2017) 'The Growing Importance of Social Skills in the Labor Market*.' *The Quarterly Journal of Economics* 132(4), 1593–1640

Doepke, Matthias, and Fabrizio Zilibotti (2017) 'Parenting With Style: Altruism and Paternalism in Intergenerational Preference Transmission.' *Econometrica* 85(5), 1331–1371

Doepke, Matthias, Giuseppe Sorrenti, and Fabrizio Zilibotti (2019) 'The Economics of Parenting.' *Annual Review of Economics* 11(1), 55–84

Edin, Per-Anders, Peter Fredriksson, Martin Nybom, and Björn Öckert (2022) 'The Rising Return to Noncognitive Skill.' *American Economic Journal: Applied Economics* 14(2), 78–100

Etheridge, Ben, Emilia Del Bono, and Paul Garcia (2026) 'The Economic Value of Childhood Socio-Emotional Skills.' *Journal of Labor Economics*

Fiorini, Mario, and Michael P. Keane (2014) 'How the Allocation of Children's Time Affects Cognitive and Noncognitive Development.' *Journal of Labor Economics* 32(4), 787–836

Freyberger, Joachim (2026) 'Normalizations and Misspecification in Skill Formation Models.' *Review of Economic Studies* 93(4), 2574–2604

Guryan, Jonathan, Erik Hurst, and Melissa Kearney (2008) 'Parental Education and Parental Time with Children.' *Journal of Economic Perspectives* 22(3), 23–46

Hazen, Jacob (2025) 'Human Capital Development of Socio-Emotional Skills in Children.' *Working Paper*

Heckman, James, Rodrigo Pinto, and Peter Savelyev (2013) 'Understanding the mechanisms through which an influential early childhood program boosted adult outcomes.' *American Economic Review* 103(6), 2052–2086

Kim, Jun Hyung, Wolfgang Schulz, Tanja Zimmermann, and Kurt Hahlweg (2018) 'Parent–child interactions and child outcomes: Evidence from randomized intervention.' *Labour Economics* 54, 152–171

Lempers, Jacques D., Dania Clark-Lempers, and Ronald L. Simons (1989) 'Economic Hardship, Parenting, and Distress in Adolescence.' *Child Development* 60(1), 25

Lindqvist, Erik, and Roine Vestman (2011) 'The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment.' *American Economic Journal: Applied Economics* 3(1), 101–128

Lizzeri, Alessandro, and Marciano Siniscalchi (2008) 'Parental Guidance and Supervised Learning.' *The Quarterly Journal of Economics* 123(3), 1161–1195

Milligan, Kevin (2007) 'Canadian tax and credit simulator.' *Database, software and documentation, Version 2*

Milligan, Kevin, and Mark Stabile (2011) 'Do Child Tax Benefits Affect the Well-being of Children?

Evidence from Canadian Child Benefit Expansions.' *American Economic Journal: Economic Policy* 3(3), 175–205

Mitchell, Mark, Marta Favara, Catherine Porter, and Alan Sánchez (2023) 'Human Capital Development: New Evidence on the Production of Socio-Emotional Skills.' *Journal of Human Resources* pp. 1120–11342R1

Molnar, Timea Laura (2023) 'Costs of daycare, complementarities, and heterogeneous productivity of parenting time in child skill formation.' Technical Report, Working paper

Moroni, Gloria, Cheti Nicoletti, and Emma Tominey (2025) 'Stressful Home Environment and the Child's Socio-Emotional Development.' *Journal of Human Capital* 19(3), 474–514

Roberts, Gail C., Jeanne H. Block, and Jack Block (1984) 'Continuity and Change in Parents' Child-Rearing Practices.' *Child Development* 55(2), 586–597

Sanders, Matthew R., James N. Kirby, Cassandra L. Tellegen, and Jamin J. Day (2014) 'The Triple P-Positive Parenting Program: A systematic review and meta-analysis of a multi-level system of parenting support.' *Clinical Psychology Review* 34(4), 337–357

Schaefer, Earl S. (1965) 'Children's Reports of Parental Behavior: An Inventory.' *Child Development* 36(2), 413–424

Sevim, Dilek, Victoria Baranov, Sonia Bhalotra, Joanna Maselko, and Pietro Biroli (2024) 'Trajectories of Early Childhood Skill Development and Maternal Mental Health.' *Journal of Human Resources* 59(S), S365–S401

Stattin, Håkan, and Margaret Kerr (2000) 'Parental Monitoring: A Reinterpretation.' *Child Development* 71(4), 1072–1085

Steinberg, Laurence (2001) 'We Know Some Things: Parent–Adolescent Relationships in Retrospect and Prospect.' *Journal of Research on Adolescence* 11(1), 1–19

Thompson, Ross A. (2014) 'Stress and Child Development.' *Future of Children* 24(1), 41–59

Toppeta, Alessandro (2022) 'Skill formation with siblings.' Technical Report

Wiswall, Matthew, and Francesco Agostinelli (2024) 'Estimating the Technology of Children's Skill Formation.' *Journal of Political Economy* p. 733396

A NLSCY Survey Design and Attrition

Survey administration The National Longitudinal Survey of Children and Youth (NLSCY) was administered by Statistics Canada in partnership with Human Resources and Skills Development Canada. Each sampled household is assigned a Person Most Knowledgeable (PMK). The PMK is usually the mother, who answers questions on the household environment, parenting practices, and information about the child. The measures about the children are reported by the PMK from ages 0 to 9. Children aged 10 and older complete a separate self-administered questionnaire covering socio-emotional outcomes, parental bond measures, and school experiences. However, at the age of 10, both PMK and child report measures on the same topics. The household component was conducted by a computer-assisted personal interviewing program. The child component was a paper questionnaire completed at the time of the interview.

Linkage to administrative tax data The NLSCY records are linked to the T1 Family File (T1FF), an administrative tax dataset drawn from personal income tax returns filed with the Canada Revenue Agency. The linkage tracks child respondents into adulthood, so I observe their adminstrative earnings directly. Match rates rose from below 1% in 1994 to roughly 87% by 2015, although approximately 98% of longitudinal respondents were linked in at least one year. Match rates are low since children were to young to file their own tax returns in the early years of the survey. The T1FF covers the years from 1993 to 2019. Thus, children in the sample are observable as adults at ages 25 to 36 in 2019.

Attrition and sample selection Attrition is non-random according to the NLSCY. Children in single-parent and lower-income households exit the panel at higher rates. I sample choice I make is that children are kept with complete parental labour market data across all three age bins. This choice favours those with continuous survey participation, so the estimation sample may over-represent stable two-parent families. Children in the sample then come from households with higher income and more educated parents. Thus my parameter estimates should be interpreted with respect to children in stable two-parent households and that it may not carry over to single-parent or more disadvantaged families.

Child cognitive measures The NLSCY assessed children's cognitive development using math and reading comprehension tests for grades 2 to 10 (ages 7–15). The math test was a shortened version of the Canadian Achievement Test, Second Edition (CAT/2) Mathematics Computation Test. It covers

arithmetic skills and problem-solving. The math test is available in all cycles. The reading test based on the CAT/2 was also administered and tested comprehension through passage analysis and information recall. Due to the survey's time constraints, the reading test was dropped after cycle 3, so later cycles contain only the math test. In the latent factor literature, it is possible to create a latent cognitive factor from two measures. However, both cognitive measures are available only in cycles 2 and 3 for children in grade 2 or older. This paper only considers the math test as a proxy for cognitive skill.

Construction of wages and weekly hours Letting E_{jt}^{ann} denote annual earnings, h_{jt}^{usual} usual weekly hours, and W_{jt} implied weeks worked, the accepted wage is

$$w_{jt}^{\text{acc}} \;=\; \frac{E_{jt}^{\text{ann}}}{h_{jt}^{\text{usual}} \times W_{jt}} \qquad \text{for } h_{jt}^{\text{usual}} \times W_{jt} > 0.$$

Average weekly market hours are then

$$h_{jt} \;=\; h_{jt}^{\text{usual}} \left(\frac{W_{jt}}{52} \right).$$

Mapping to the classic parenting typology Following Barszczewski and Thomas (2025), the parenting dimensions above map onto the authoritarian, authoritative, permissive, and neglectful typology of Baumrind (1966, 1967) as follows.

- *Authoritative:* High positive interaction and nurturance, low hostility, and consistent, non-punitive discipline.
- *Authoritarian:* High positive interaction and nurturance, high hostility, and rejection-oriented discipline.
- *Permissive:* High positive interaction and nurturance, low consistency and monitoring.
- *Neglectful:* Low positive interaction, nurturance, consistency, and monitoring.

This mapping is descriptive. Households are not assigned to one of the four types, and only the underlying dimensions enter the parental bond and its accumulation technology in Section 4.

Table A.1: Investment Measure Items

Parental School Participation (6 items, parent-reported binary)	*"During this school year, have you done any of the following?"*
1	Spoken to, visited, or corresponded with your child's teacher.
2	Visited your child's class.
3	Attended a school event in which your child participated (e.g. a play, sports competition, or science fair).
4	Gone on a class trip with your child.
5	Helped elsewhere in the school, such as in the library or computer room.
6	Attended a parent-school association, home and school liaison committee, parent advisory committee, or parent council meeting.
Active Parenting Time (3 items, parent-reported ordinal)	
1	How often do you and your child talk about school work or behaviour in class?
2	How often do you and your child talk about his/her school friends or activities?
3	How often do you check his/her homework or provide help with homework?
	Scale: Daily; A few times a week; Once a week; A few times a month; Once a month; Rarely.
Outside-School Activities (4 items, child-reported ordinal)	*"Outside of school, during the past 12 months, how often have you…"*
1	Taken part in dance, gymnastics, karate, or other groups or lessons (always organised outside of school)?
2	Taken part in art, drama, or music groups, clubs, or lessons (organised outside of school)?
3	Played sports WITH a coach or instructor (swimming lessons, baseball, hockey, etc.)?
4	Played sports or done physical activities WITHOUT a coach or instructor (biking, skateboarding, etc.)?
	Scale: Never; Less than once a week; 1–3 times a week; 4 or more times a week.
Additional Investment Measures (child- or parent-reported)	
1	*Homework time* (parent-reported): On days when he/she is assigned homework, how much time does he/she usually spend doing homework? *Scale: 0–15 min; 16–30 min; 31 min to less than 1 hr; 1 to less than 1.5 hrs; 1.5 to less than 2 hrs; 2 to less than 3 hrs; 3 to less than 4 hrs; 4 hrs or more.*
2	*Private school enrolment* (parent-reported binary): whether the child is currently enrolled in a private school.

B GSS and SHS Sample Statistics

Table B.1: GSS Time-Use Sample: Means, All Respondents and by Parent

	All	Mothers	Fathers
Time Use (weekly hours)			
Market Work	34.57	24.09	46.67
Active Parenting Time	11.73	15.16	7.77
Leisure (Residual)	121.70	128.75	113.56
Demographics			
Female (= 1)	0.54	–	–
Observations	3,400	1,820	1,580

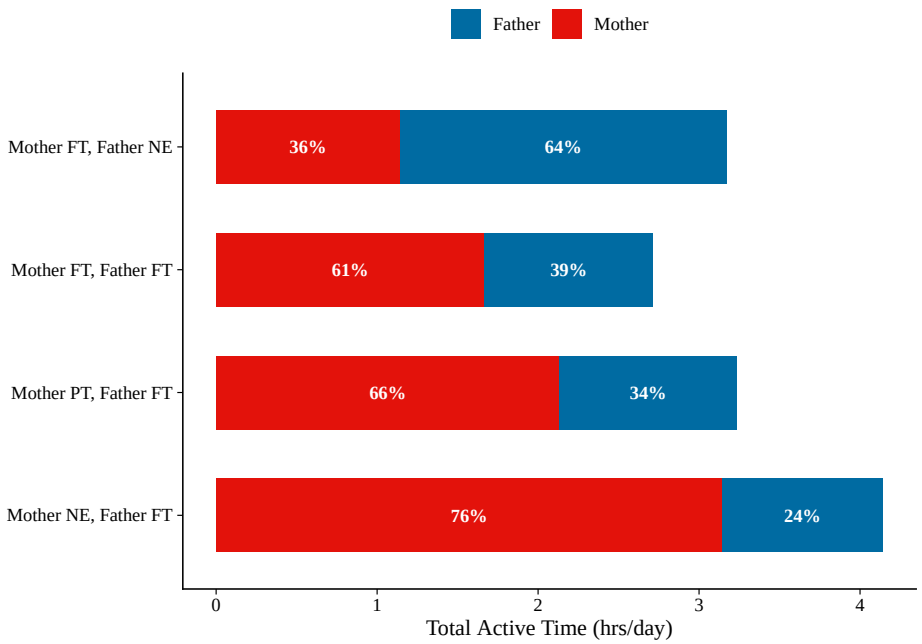
Rounded to comply with the data provider's confidentiality requirements.

Table B.2: SHS Household Expenditure Sample: Means and Standard Deviations

	Mean	Std. Dev.
Child-Related Expenditures (annual, $CAD)		
Total Household Expenditure	2,543.04	3,168.54
Clothing and Footwear	158.72	312.94
Toys	274.58	416.25
Education	1,118.26	3,158.96
Demographics		
Household Income ($CAD)	85,594.57	47,584.50
Mother's Education (years)	14.22	–
Observations	7,130	

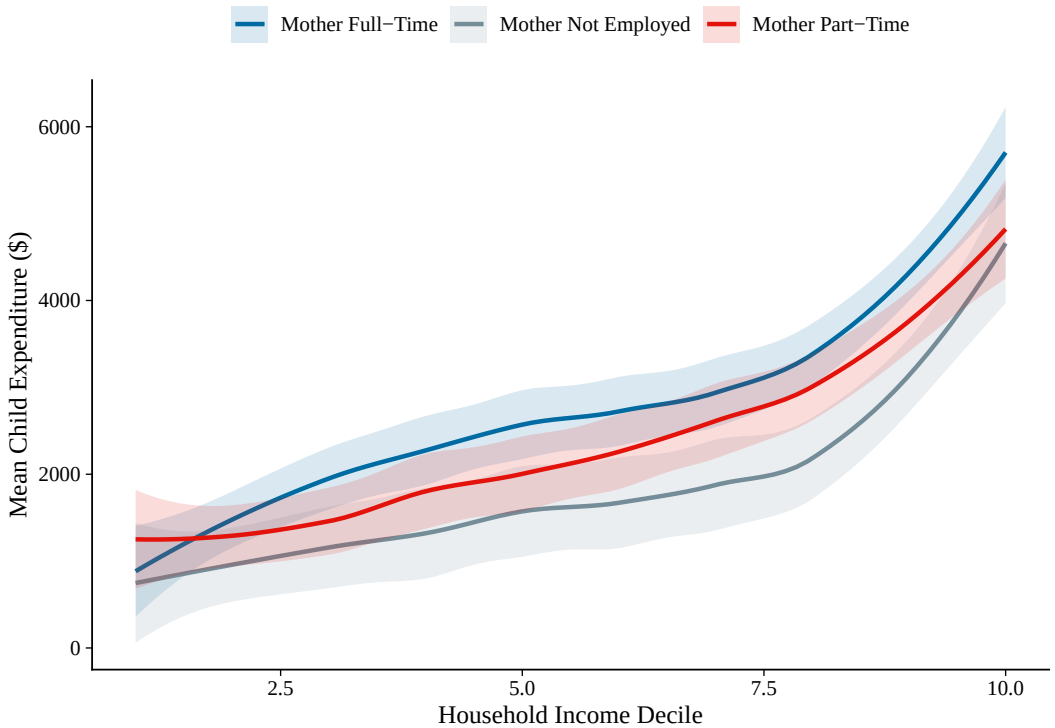
Rounded to comply with the data provider's confidentiality requirements.

Figure B.1: Total Active Time and Its Division between Parents, by Couple Employment Type



Notes: Bar length is the couple's total active parenting time (hours/day, GSS time diaries), stacked by parent; labels report each parent's share of that total. Couple employment type is Mother_Father (e.g., "Mother NE, Father FT" is a mother not employed and a father working full-time). Part-time father categories are not shown, since the corresponding cells do not meet the data provider's minimum cell-size requirement for disclosure.

Figure B.2: Mean Child Expenditure by Household Income Decile and Couple Employment Type



Notes: Lines are loess-smoothed trends, with a 95% confidence band, of mean child-related expenditure by household income decile. The father works full-time in all three groups shown; the groups differ only in the mother's employment status (full-time, part-time, or not employed). The decile grouping is chosen to satisfy the data provider's minimum cell-size requirement for disclosure, not for modelling reasons.

C Measurement Model and Construction of Latent Moments

C.1 Measurement Model

Parental bond and child skill are not directly observed but are inferred from a set of observed measures, namely parenting and behavioural scales. For identification of the measurement model, I refer the reader to Freyberger (2026), who provides a detailed discussion of the identification conditions for measurement models. Following Freyberger (2026), I define

$$Z^{\theta}_{t,m} = \exp(\mu^{\theta}_{t,m})\theta_t^{\lambda^{\theta}_{t,m}}$$
$$Z^{P}_{t,m} = \exp(\mu^{P}_{t,m})P_t^{\lambda^{P}_{t,m}}$$

so that

$$\ln Z^{\theta}_{t,m} = \mu^{\theta}_{t,m} + \lambda^{\theta}_{t,m} \ln \theta_t + \varepsilon^{\theta}_{t,m}, \tag{9}$$
$$\ln Z^{P}_{t,m} = \mu^{P}_{t,m} + \lambda^{P}_{t,m} \ln P_t + \varepsilon^{P}_{t,m}. \tag{10}$$

Here, $Z^{\theta}_{t,m}$ and $Z^{P}_{t,m}$ denote observed measure m at time t associated with child skill and parental bond, respectively. The parameters $\mu^{\theta}_{t,m}$ and $\mu^{P}_{t,m}$ are measure-specific intercepts, $\lambda^{\theta}_{t,m}$ and $\lambda^{P}_{t,m}$ are factor loadings, and $\varepsilon^{\theta}_{t,m}$ and $\varepsilon^{P}_{t,m}$ are measurement errors. The intercepts capture the average level of each observed measure when the corresponding latent log-factor is zero, while the loadings measure the strength of the relationship between the observed indicators and the latent log-factor.

For the child skill and bond measures, I also estimate a measurement-invariant counterpart of this system in which the loading and intercept of each measure are constrained to be equal across the three age bins, $\lambda^{\theta}_{t,m} = \lambda^{\theta}_{m}$ and $\mu^{\theta}_{t,m} = \mu^{\theta}_{m}$ for all t, and analogously for P. The anchoring restriction used in Section 5.3 to identify the age profile of the latent means relies on this invariant specification. Because the anchor measure's loading and intercept are fixed at their pooled, age-invariant value rather than free to vary by age, movement in the estimated latent mean reflects development in the underlying factor rather than drift in the measurement instrument. Table C.1 reports both specifications side by side, with the age-varying loadings and intercepts (λ_t, μ_t) in the first three column pairs and the pooled invariant estimates (λ, μ) in the final column pair; investment measures are not subject to this restriction and are therefore estimated freely by age group only.

C.2 Constructing Empirical Moments from Latent Factors

I have two latent factors in the model, namely the parental bond P_t and the child's socio-emotional skill θ_t.[54] The measures that make up each latent factor are taken from the scales developed by the NLSCY.

Since the latent log-factors $\ln \theta_t$ and $\ln P_t$ are not directly observed, it is difficult to model them jointly with covariates. A common approach is to use *Bartlett factor scores*, which are linear combinations of observed indicators designed to minimize measurement error. Each score equals the latent log-factor plus a noise term,

$$\hat{f}_t = f_t + u_t, \qquad u_t \equiv W_t \varepsilon_t, \tag{11}$$

where $f_t \in \{\ln \theta_t, \ln P_t\}$, W_t depends on estimated factor loadings and error variances Ψ_t, ε_t collects the item-level disturbances at time t, and u_t is the resulting measurement error of the score.[55] Bartlett scores can be interpreted as a generalized least squares (GLS) estimator of the latent log-factor.

Because u_t is uncorrelated with f_t, Bartlett scores reproduce the correct covariances with any covariate X_t independent of item errors,

$$\mathrm{Cov}(X_t, \hat{f}_t) = \mathrm{Cov}(X_t, f_t), \tag{12}$$

and the same holds across time if measurement errors are independent.[56] If $\hat{f}_t^{\theta}$ and $\hat{f}_t^{P}$ denote the Bartlett scores for $\ln \theta_t$ and $\ln P_t$, this cross-time property implies

$$\mathrm{Cov}(\hat{f}_t^{\theta}, \hat{f}_{t+s}^{\theta}) = \mathrm{Cov}(\ln \theta_t, \ln \theta_{t+s}), \quad s > 0, \tag{13}$$

$$\mathrm{Cov}(\hat{f}_t^{P}, \hat{f}_{t+s}^{P}) = \mathrm{Cov}(\ln P_t, \ln P_{t+s}), \quad s > 0. \tag{14}$$

Bartlett scores do not preserve scale, however, since their variance is inflated by the variance of the measurement error u_t (equation (22)), so correlations based on scores are attenuated. In practice, I therefore use Bartlett scores to construct latent *covariances* of the log-factors and rely on model-implied variances for correct standardization.

[54] In a closely related paper, Hazen (2025) explores the multi-dimensional nature of socio-emotional skills. I model socio-emotional skills as a single latent factor to limit the complexity of the estimation. Since Hazen (2025) finds a five-factor structure, allowing for it here would require estimating five times the number of parameters in the skill formation function.

[55] See Appendix C.3 for details on the derivation.

[56] This assumption is standard in the literature; see, for example, Cunha et al. (2010), Attanasio et al. (2020b), and Wiswall and Agostinelli (2024).

C.3 Derivation of Bartlett Scores with Mean Structure

Consider the measurement system at time t,

$$\ln Z_t = \mu_t + \Lambda_t f_t + \varepsilon_t, \qquad \mathbb{E}[\varepsilon_t] = 0, \quad \mathrm{Var}(\varepsilon_t) = \Psi_t, \tag{15}$$

where $\ln Z_t$ is the $p \times 1$ vector stacking the log measures $\ln Z_{t,m}$ of equations (9) and (10) (and, analogously, of the investment measures), f_t is the $k \times 1$ vector of latent log-factors (for example, $\ln \theta_t$, $\ln P_t$, and $\ln I_t$, taken separately or jointly), Λ_t is the $p \times k$ matrix of loadings $\lambda_{t,m}$, μ_t is the $p \times 1$ vector of intercepts $\mu_{t,m}$, and ε_t stacks the measurement errors $\varepsilon_{t,m}$.

The Bartlett score estimator minimizes the squared distance between the observed and model-implied measures,

$$\hat{f}_t = \arg\min_f \, (\ln Z_t - \mu_t - \Lambda_t f)' \Psi_t^{-1} (\ln Z_t - \mu_t - \Lambda_t f). \tag{16}$$

The first-order condition with respect to f is

$$\Lambda_t' \Psi_t^{-1} (\ln Z_t - \mu_t - \Lambda_t \hat{f}_t) = 0, \tag{17}$$

which leads to

$$(\Lambda_t' \Psi_t^{-1} \Lambda_t) \hat{f}_t = \Lambda_t' \Psi_t^{-1} (\ln Z_t - \mu_t). \tag{18}$$

Hence, the Bartlett estimator and its associated weighting matrix are

$$\hat{f}_t = W_t (\ln Z_t - \mu_t), \tag{19}$$

$$W_t = (\Lambda_t' \Psi_t^{-1} \Lambda_t)^{-1} \Lambda_t' \Psi_t^{-1}. \tag{20}$$

Expanding using the true model $\ln Z_t = \mu_t + \Lambda_t f_t + \varepsilon_t$ yields

$$\hat{f}_t = f_t + u_t, \qquad u_t \equiv W_t \varepsilon_t. \tag{21}$$

The weighting matrix W_t determines how measurement errors enter the score, assigning higher weight to reliable items (small error variances in Ψ_t) and lower weight to noisy ones. Because $W_t \Lambda_t = I_k$, the estimator is unbiased for the factor mean,

$$\mathbb{E}[\hat{f}_t] = \mathbb{E}[f_t].$$

Moreover, u_t is orthogonal to f_t, which delivers the covariance result of Appendix C.2 for any covariate X_t uncorrelated with the measurement errors. Finally, the score carries the additional noise term u_t, so that its variance is

$$\mathrm{Var}(\hat{f}_t) = \mathrm{Var}(f_t) + \Sigma_t, \qquad \Sigma_t \equiv W_t \Psi_t W_t' = (\Lambda_t' \Psi_t^{-1} \Lambda_t)^{-1}, \tag{22}$$

where the second equality follows from substituting the expression for W_t above. Correlations based on Bartlett scores are therefore attenuated. Appendix D uses Σ_t to correct regressions on the scores for this attenuation.

C.4 Correlations Among Latent Factor Scores

Table C.2 reports the correlations among the latent factor scores for child socio-emotional skill (θ), the parental bond (P), and parental investment (I) at each age bin, obtained from the measurement-invariant confirmatory factor model underlying the production function and investment equation estimates.

Table C.2: Correlations Among Latent Factor Scores by Age Bin

	Ages 10–11	Ages 12–13	Ages 14–15
θ–P	0.635	0.593	0.635
θ–I	0.076	0.053	0.068
P–I	0.227	0.219	0.254

Notes: Pairwise correlations of the latent factor scores for child socio-emotional skill (θ), the parental bond (P), and parental investment (I), by age bin.

Table C.1: Measurement System Loadings and Means by Age Group

	Ages 10–11		Ages 12–13		Ages 14–15		Invariant	
Measure	λ_t	μ_t	λ_t	μ_t	λ_t	μ_t	λ	μ
Child Socio-Emotional Skill								
Pro-social	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]
Conduct Disorder/Physical Aggression	1.15	−0.59	0.90	0.16	1.94	−2.48	1.07	−0.31
Hyperactivity/Inattention	2.50	−4.28	1.83	−2.30	3.52	−6.46	1.93	−2.61
Anxiety and Emotional Disorder	2.21	−3.51	1.45	−1.29	2.73	−4.45	1.58	−1.68
Parent-Child Bond								
Parental Nurturance	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]
Parental Knowledge	0.53	1.07	0.38	1.57	0.35	1.64	0.42	1.44
Parental Rejection	0.47	1.53	0.43	1.60	0.49	1.36	0.57	1.18
Parental Investment								
Outside-School Activities	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]	1.00[a]	0.00[a]
Parental School Participation	1.38	0.07	1.87	−0.49	1.69	−0.13	–	–
Talks About School	2.09	−0.40	2.96	−1.25	2.84	−0.67	–	–
Talks About Friends/Activities	2.12	−0.42	2.23	−0.46	2.08	0.03	–	–
Checks Homework	2.33	−0.79	3.70	−2.37	3.53	−2.04	–	–
Homework Time	0.89	0.05	1.45	−0.37	0.78	0.53	–	–

Notes: [a]Marker indicator: loading and intercept fixed for identification, not estimated. "–" denotes rows whose loadings/intercepts are freely estimated by age group in both the constrained and unconstrained model.

D Measurement-Error Corrections for the Reduced-Form Regressions

D.1 The Croon (2002) Correction

Section 3 estimates reduced-form regressions using Bartlett scores in place of the unobserved true latent factors. However, the Bartlett scores do not preserve the variance of the underlying factors. Thus, the OLS coefficients on the Bartlett scores are attenuated toward zero. I correct for this attenuation using the approach of Croon (2002), which Heckman et al. (2013) and Del Bono et al. (2026) also use for regressions on estimated factor scores.

Let f_i be the true log-factors $\ln \theta_{i,14_15}$ and $\ln P_{i,14_15}$ for child i, and let $\hat{f}_i$ denote the be the estimated Bartlett scores. As derived in Appendix C.3,

$$\hat{f}_i = f_i + u_i, \qquad \mathrm{Var}(u_i) = \Sigma_u, \tag{23}$$

where u_i is the Bartlett-score measurement error and Σ_u is the age 14_15 error variance of equation (22). Now consider the regression of interest,

$$Y_i = \alpha + \beta' f_i + \gamma' X_i + \epsilon_i, \qquad \mathbb{E}[\epsilon_i \mid f_i, X_i] = 0, \tag{24}$$

where X_i is a vector of controls measured without error. Since f_i is unobserved, I estimate this regression with $\hat{f}_i$ in its place, which equivalent to running OLS on

$$Y_i = \alpha + \beta' \hat{f}_i + \gamma' X_i + \epsilon_i - \beta' u_i. \tag{25}$$

The regressor $\hat{f}_i$ is correlated with the error term $\epsilon_i - \beta' u_i$ through u_i, so OLS on the observed scores is inconsistent for (β, γ).

To see the source of this bias, write out the probability limit of the OLS estimator,

$$\mathrm{plim} \begin{pmatrix} \hat{\beta}^{OLS} \\ \hat{\gamma}^{OLS} \end{pmatrix} = \begin{pmatrix} \mathrm{Cov}(\hat{f}, \hat{f}) & \mathrm{Cov}(\hat{f}, X) \\ \mathrm{Cov}(X, \hat{f}) & \mathrm{Cov}(X, X) \end{pmatrix}^{-1} \begin{pmatrix} \mathrm{Cov}(\hat{f}, f) & \mathrm{Cov}(\hat{f}, X) \\ \mathrm{Cov}(X, f) & \mathrm{Cov}(X, X) \end{pmatrix} \begin{pmatrix} \beta \\ \gamma \end{pmatrix}. \tag{26}$$

Since u_i is orthogonal to f_i and the controls X_i are measured without error and uncorrelated with the measurement-level errors, u_i is also orthogonal to X_i, so

$$\mathrm{Cov}(\hat{f}, f) = \mathrm{Cov}(f, f), \qquad \mathrm{Cov}(\hat{f}, X) = \mathrm{Cov}(f, X),$$

while the added noise inflates the variance of the observed scores,

$$\mathrm{Cov}(\hat{f}, \hat{f}) = \mathrm{Cov}(f, f) + \Sigma_u.$$

The two matrices on the right-hand side of equation (26) therefore agree everywhere except in the latent-factor block, where the first carries an extra Σ_u that the second does not. The entire bias in the OLS estimator comes from this one term.

We can now correct for this bias via the Croon correction. Let

$$R_i = (\hat{f}_i', X_i')', \qquad \Sigma_{\tilde{u}} = \begin{pmatrix} \Sigma_u & 0 \\ 0 & 0 \end{pmatrix}.$$

Subtracting an estimate of $\Sigma_{\tilde{u}}$ from the observed covariance matrix of regressors before inverting it gives a consistent estimator,

$$\hat{b}^{\,\mathrm{Croon}} = \left(\hat{\Sigma}_{RR} - \hat{\Sigma}_{\tilde{u}}\right)^{-1} \hat{\Sigma}_{RY}, \qquad \hat{b}^{\,\mathrm{Croon}} = \begin{pmatrix} \hat{\beta}^{\,\mathrm{Croon}} \\ \hat{\gamma}^{\,\mathrm{Croon}} \end{pmatrix}. \tag{27}$$

I construct $\hat{\Sigma}_{\tilde{u}}$ by placing the estimate of Σ_u in the latent-factor block and zeros elsewhere.

I obtain standard errors by bootstrapping the entire procedure (factor estimation, Bartlett scoring, and the Croon correction) 500 times, resampling at the child level.

E Additional Estimates for the Production Functions and Investment Equation

Table 6 and Table 7 in Section 3 pool observations across both transitions between adjacent survey rounds. This appendix reports the same specifications separately for the two age-bin transitions, specifically ages 10–11 to 12–13 and ages 12–13 to 14–15. I add the Cobb-Douglas and Translog specifications to the analysis. Appendix C reports the correlations among the underlying latent factor scores by age bin.

E.1 Production Functions for Skill and the Parental Bond

Table E.1: Reduced Form Evidence by Age Bin (Full Specification)

	Cobb-Douglas				Translog			
	Log SE Skill		Log Parental Bond		Log SE Skill		Log Parental Bond	
	[1]	**[2]**	**[3]**	**[4]**	**[5]**	**[6]**	**[7]**	**[8]**
Log Investment	0.033	0.005	0.114***	0.101***	0.030	0.012	0.106***	0.103***
	(0.033)	(0.027)	(0.035)	(0.035)	(0.033)	(0.028)	(0.035)	(0.033)
Log SE Skill (Lag)	0.637***	0.581***	−0.039	−0.026	0.695***	0.638***	0.014	0.040
	(0.074)	(0.055)	(0.070)	(0.058)	(0.071)	(0.052)	(0.068)	(0.053)
Log Parental Bond (Lag)	0.038	0.160***	0.532***	0.834***	0.089	0.239***	0.579***	0.927***
	(0.081)	(0.058)	(0.097)	(0.085)	(0.081)	(0.059)	(0.103)	(0.087)
Log Investment × Own Lag	—	—	—	—	0.030	0.005	−0.006	−0.020
					(0.022)	(0.023)	(0.021)	(0.029)
Skill (Lag) × Bond (Lag)	—	—	—	—	0.077**	0.107***	0.074***	0.131***
					(0.031)	(0.025)	(0.022)	(0.028)
Mother's Education	−0.012	−0.002	0.000	−0.009	−0.011	−0.006	0.001	−0.014
	(0.017)	(0.013)	(0.015)	(0.013)	(0.017)	(0.012)	(0.015)	(0.013)
Father's Education	0.018	0.013	0.007	0.000	0.015	0.011	0.004	−0.001
	(0.016)	(0.012)	(0.015)	(0.012)	(0.016)	(0.012)	(0.015)	(0.012)
Number of Siblings	−0.032	0.000	0.009	−0.001	−0.024	0.013	0.018	0.013
	(0.036)	(0.026)	(0.035)	(0.028)	(0.035)	(0.026)	(0.035)	(0.029)
Controls				✓				
Observations	1,600	1,890	1,600	1,890	1,600	1,890	1,600	1,890

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Measurement-error-corrected estimates, with standard errors from 750 bootstraps clustered by child, as in Table 6. Full specification includes both the own lag of each outcome and the cross lag of the other stock, mirroring the pooled specification in Table 6, estimated separately by age-bin transition. Columns [1], [3], [5], and [7] are ages 10–11 to 12–13, and Columns [2], [4], [6], and [8] are ages 12–13 to 14–15. Columns [5] to [8] add the interaction terms of Table E.2 (Columns [2] and [4]), with investment and the lagged regressors centred at their sample means. Investment × Own Lag is the interaction of investment with the lag of the dependent variable (skill in Columns [5] and [6], bond in Columns [7] and [8]). Observations are child-transition pairs, rounded to the nearest 10. The pooled sample of 3,480 in Table 6 stacks the two transitions, and the counts sum to 3,490 only because each count is rounded independently.

E.2 Interaction Effects

Table E.2 adds interaction terms between investment, lagged skill, and the lagged bond to the pooled specification of Table 6. This is also commonly known as the Translog specification in the literature (Del Bono et al., 2022). Interpretation of the interaction investment term captures the extent to which investment is more or less productive when the child already has a strong bond or skill. The interaction on the lagged bond and lagged skill captures whether skill and the bond reinforce each other in production.

Columns [1] and [3] add the investment interaction to the respective equations. Columns [2] and [4] add the bond-skill lag interaction.[57]

Table E.2: Interaction Effects in the Production Functions

	Log SE Skill		Log Parental Bond	
	[1]	**[2]**	**[3]**	**[4]**
Log Investment	0.024	0.022	0.111***	0.108***
	(0.020)	(0.019)	(0.024)	(0.024)
Log SE Skill (Lag)	0.613***	0.662***	−0.032	0.026
	(0.050)	(0.045)	(0.047)	(0.043)
Log Parental Bond (Lag)	0.100*	0.163***	0.677***	0.743***
	(0.051)	(0.051)	(0.072)	(0.074)
Log Investment × Own Lag	0.032**	0.014	0.005	−0.005
	(0.015)	(0.017)	(0.016)	(0.018)
Skill (Lag) × Bond (Lag)	—	0.089***	—	0.097***
		(0.021)		(0.017)
Mother's Education	−0.007	−0.008	−0.006	−0.007
	(0.010)	(0.009)	(0.010)	(0.010)
Father's Education	0.016*	0.014	0.006	0.004
	(0.009)	(0.009)	(0.009)	(0.009)
Number of Siblings	−0.012	−0.002	0.005	0.016
	(0.020)	(0.020)	(0.022)	(0.022)
Controls	✓	✓	✓	✓
Observations	3,480	3,480	3,480	3,480

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Measurement-error-corrected estimates following Attanasio et al. (2020b), with standard errors from 750 bootstraps clustered by child, as in Table 6. Full specification, pooling both age-bin transitions with an age-bin control. Investment × Own Lag is the interaction of investment with the lag of the dependent variable (skill in Columns [1] and [2], bond in Columns [3] and [4]). Controls include mother's education, father's education, child sex, number of siblings, mother's age at the child's birth, and round fixed effects.

The skill-bond interaction is the stable result. It is positive and significant in both the skill and the bond equation, so skill and the bond are complements in production beyond their additive cross-lags. A higher lagged bond raises the marginal product of lagged skill, and vice versa. The investment interactions are less reliable. The investment-skill interaction is significant in Column [1] and insignificant once the bond-skill interaction enters in Column [2]. The investment-bond interaction is insignificant and changes sign between Columns [3] and [4]. The skill-bond interaction is also positive and significant within each transition (Columns [5] to [8] of Table E.1), and the investment interactions are insignificant in both.

[57] Investment and the lagged regressors are centred at their sample means before interacting, so each main effect is the effect at the mean of the variables it is interacted with, not the unconditional effect reported in Table 6.

E.3 Robustness: Investment Measure Including the Parental Bond

Section 3 attributes the large, significant investment coefficients in Cunha et al. (2010) and Toppeta (2022) to a conflation of the parent-child bond with material and time investment. Both latent factors include measures, such as parental encouragement and warmth, that this paper separates to the bond rather than to investment. This exercise, I re-estimate the investment index without excluding the bond, so that I_{t+1} combines material inputs, time inputs, and the parenting behaviours. This is to simulate the results in latent factor construction of Cunha et al. (2010) and Toppeta (2022). I re-run the skill regression (Column [4]). Two other exercises test the bond's mediating role. The first one I omit I_{t+1} (Column [2]) and the second one I omit P_t (Column [3]). Column [1] reproduces the baseline θ_{t+1} specification of Table 6.

Table E.3: Mediation-Style Robustness: Skill and Bond Regressions Excluding I_{t+1} or the Cross-Lag, and Combined Measure

	[1]	[2]	[3]	[4]
Panel A: Dependent Variable is Log SE Skill (θ_{t+1})				
Log Investment	0.021	—	0.034*	0.469***
	(0.019)		(0.019)	(0.038)
Log SE Skill (Lag)	0.605***	0.604***	0.659***	0.542***
	(0.049)	(0.049)	(0.038)	(0.046)
Log Parental Bond (Lag)	0.100*	0.105**	—	—
	(0.052)	(0.051)		
Panel B: Dependent Variable is Log Parental Bond (P_{t+1})				
Log Investment	0.110***	—	0.111***	
	(0.025)		(0.024)	
Log SE Skill (Lag)	−0.032	−0.041	—	
	(0.046)	(0.048)		
Log Parental Bond (Lag)	0.675***	0.699***	0.657***	
	(0.069)	(0.069)	(0.051)	
Controls	✓	✓	✓	✓
Observations	3,480	3,480	3,480	3,480

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Panel A reports the skill regression (θ_{t+1} as dependent variable); Panel B reports the bond regression (P_{t+1} as dependent variable). Pooling both age-bin transitions with an age-bin control. Column [1] is the baseline specification, reproducing Table 6; Column [2] omits I_{t+1}; Column [3] omits the lagged value of the other latent factor (P_t in Panel A, θ_t in Panel B); Column [4] re-estimates I_{t+1} with the parental bond included alongside time and material inputs for the skill regression in Panel A, and correspondingly omits P_t. Controls include mother's education, father's education, child sex, number of siblings, mother's age at the child's birth, and round fixed effects.

Dropping I_{t+1} leaves the coefficients on θ_t and P_t unchanged, so investment has no independent explana-

tory power for skill once the bond is controlled for directly. Dropping P_t raises the persistence coefficient on θ_t. Also, I_{t+1} becomes marginally significant in that specification. Adding back the bond into I_{t+1}, in Column [4], reproduces the pattern in Cunha et al. (2010) and Toppeta (2022). The coefficient on investment becomes large and significant once the measure is no longer restricted to only material and time inputs. A investment latent factor that includes the parent-child relationship could be crediting the investment with a return that belongs to the bond (Columns [2] and [3]). This suggests that investment measures do not separate the two are likely to overstate the productivity of investment.

The bond regression in Panel B is also stable. Dropping I_{t+1} (Column [2]) leaves the bond's own persistence and the near-zero cross-lag from skill unchanged. Investment does not confound either coefficient, as in Panel A. Dropping the skill lag leaves the investment coefficient and the bond's persistence stable.

E.4 Investment Equation

Table E.4 reports the investment equation of Section 3 estimated separately for each age-bin transition.

Table E.4: Investment Equation (Measurement-Error Corrected) by Age Bin

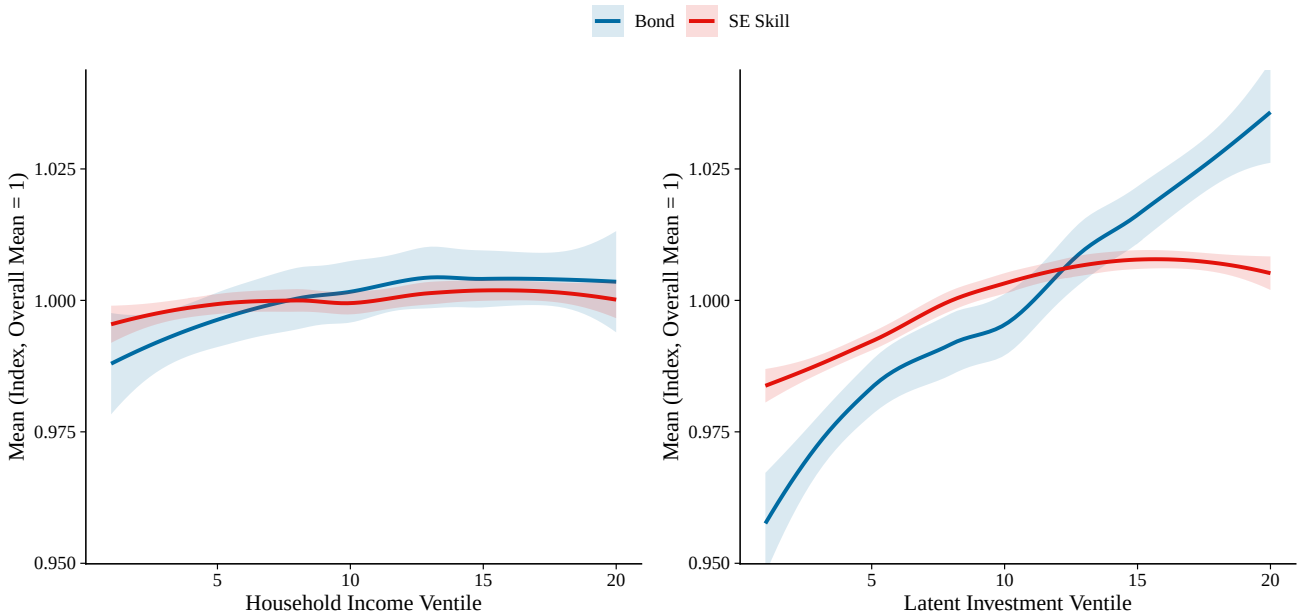
	Log Investment	
	Ages 10–11 to 12–13	Ages 12–13 to 14–15
Log Parental Bond (Lag)	0.2225^{***}	0.2117^{***}
	(0.0806)	(0.0595)
Log SE Skill (Lag)	-0.1165^{*}	-0.0620
	(0.0654)	(0.0595)
Log Household Income	0.0648^{*}	0.0883^{*}
	(0.0335)	(0.0485)
Mother's Education	0.0320^{*}	0.0480^{**}
	(0.0165)	(0.0205)
Father's Education	0.0220^{*}	0.0119
	(0.0130)	(0.0158)
Male (=1)	-0.1504^{**}	-0.0582
	(0.0724)	(0.0841)
Number of Siblings	-0.0855^{**}	-0.0660
	(0.0334)	(0.0409)
Mother's Age at Child's Birth	0.0047	-0.0076
	(0.0058)	(0.0065)
Round FE	✓	✓
N (observations, rounded)	1,600	1,890

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Measurement-error-corrected estimates, with standard errors from 750 bootstraps clustered by child, as in Table E.1. Dependent variable is the latent investment index $I_{i,t+1}$; the bond and skill regressors are measured at the start of the investment interval. Standard errors in parentheses.

F Additional Descriptive Evidence: Child Skill and the Parental Bond by Household Income and Investment

Section 3 documents the association between the parental bond and child socio-emotional skill by bond quantile (50 groups; Figure 1). This appendix reports the gradients along household income and the latent investment index I constructed in Section C.2. Figure F.1 plots, the mean of child socio-emotional skill and of the parental bond, each expressed as an index relative to its own sample mean, against ventile bins of household income (left) and of latent investment (right). For example, a value of 1.08 for either plot can be interpreted as eight percent above its own sample mean.

Figure F.1: Mean Child Skill and Parental Bond by Household Income and Investment Ventile

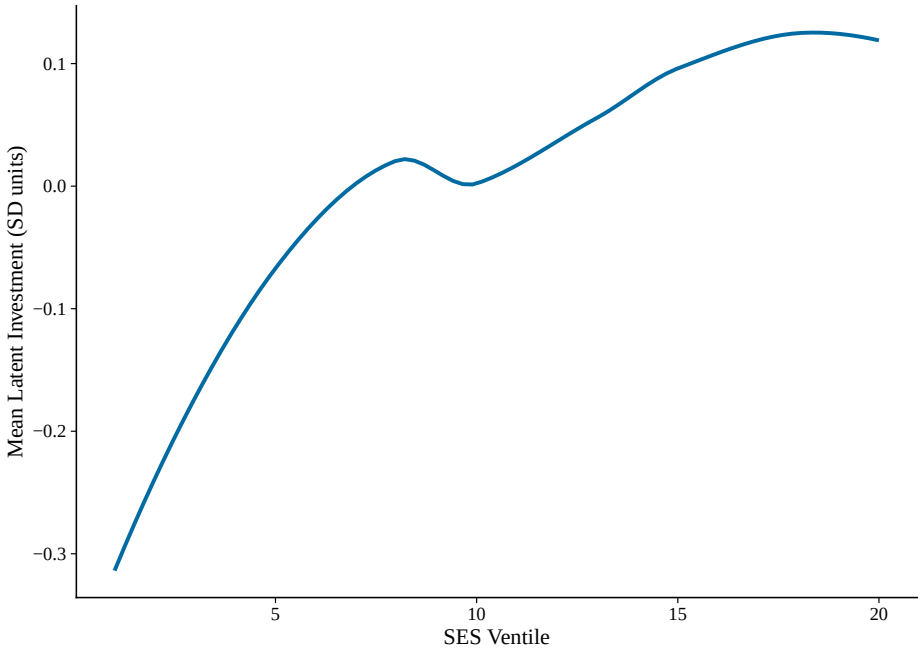


Notes: Lines are loess-smoothed fits, with 95% confidence bands, of mean child socio-emotional skill and mean parental bond against household income ventile (left panel) and latent investment ventile (right panel), each series normalized to an index with its own sample mean equal to 1. The confidence bands reflect the uncertainty of the loess fit across ventile bins rather than a respondent-level standard error. The ventile grouping satisfies the data provider's minimum cell-size requirement for disclosure.

The gradient in household income is nearly flat. Both the bond and child skill vary by less than one percentage point of their respective means across the income distribution. The gradient in latent investment is much steeper, particularly for the parental bond, which increases by roughly eight percentage points of its mean from the bottom to the top investment ventile. In contrast, socio-emotional skill increases by

only about two percentage points. Household income relaxes the budget constraint but need not translate into a stronger bond. Figure F.2 plots the mean latent investment index against ventile of a latent SES factor, constructed from mother's education, father's education, and household income, at ages 14–15. Investment increases steadily with SES, from roughly a quarter of a standard deviation below the mean at the bottom of the SES distribution to just above the mean at the top.

Figure F.2: Latent Investment by SES Ventile, Ages 14–15



Notes: SES is a latent factor constructed from mother's education, father's education, and household income. The line is a loess-smoothed fit of the standardized latent investment index against SES ventile. The ventile grouping is chosen to satisfy the data provider's minimum cell-size requirement for disclosure, not for modelling reasons.

G Additional Model Fit

Table G.1 reports the fit of moments outside Tables 9 and 10.

Table G.1: Model Fit: Additional Moments

	Data	Model
Panel A: Mother's Employment and Couple Employment Shares		
Mother not employed, 10–11	0.171	0.153
Mother not employed, 12–13	0.137	0.156
Mother not employed, 14–15	0.105	0.155
Mother full-time, 10–11	0.535	0.499
Mother full-time, 12–13	0.566	0.495
Mother full-time, 14–15	0.612	0.501
Both parents employed, 10–11	0.809	0.807
Both parents employed, 12–13	0.836	0.808
Both parents employed, 14–15	0.867	0.811
Father only employed, 10–11	0.158	0.153
Father only employed, 12–13	0.129	0.156
Father only employed, 14–15	0.095	0.155
Panel B: Active Childcare Time, Hours per Day, by Couple Employment (GSS)		
Mother, both full-time	1.666	1.872
Mother, mother not employed	3.145	2.791
Mother, mother part-time	2.133	2.366
Father, both full-time	1.049	1.191
Father, mother not employed	0.996	0.927
Father, mother part-time	1.103	1.016
Panel C: Child Expenditure (SHS)		
Mean child spending ($ per year)	2,543	5,357

Notes: Model moments are averages over 3,480 simulated households and 50 replications. Panel A data are from the NLSCY; full-time and not-employed shares are for mothers. Panel B data are means from the GSS time-use diaries by couple type. Panel C compares mean annual child expenditure in the SHS with the simulated mean over all ages.

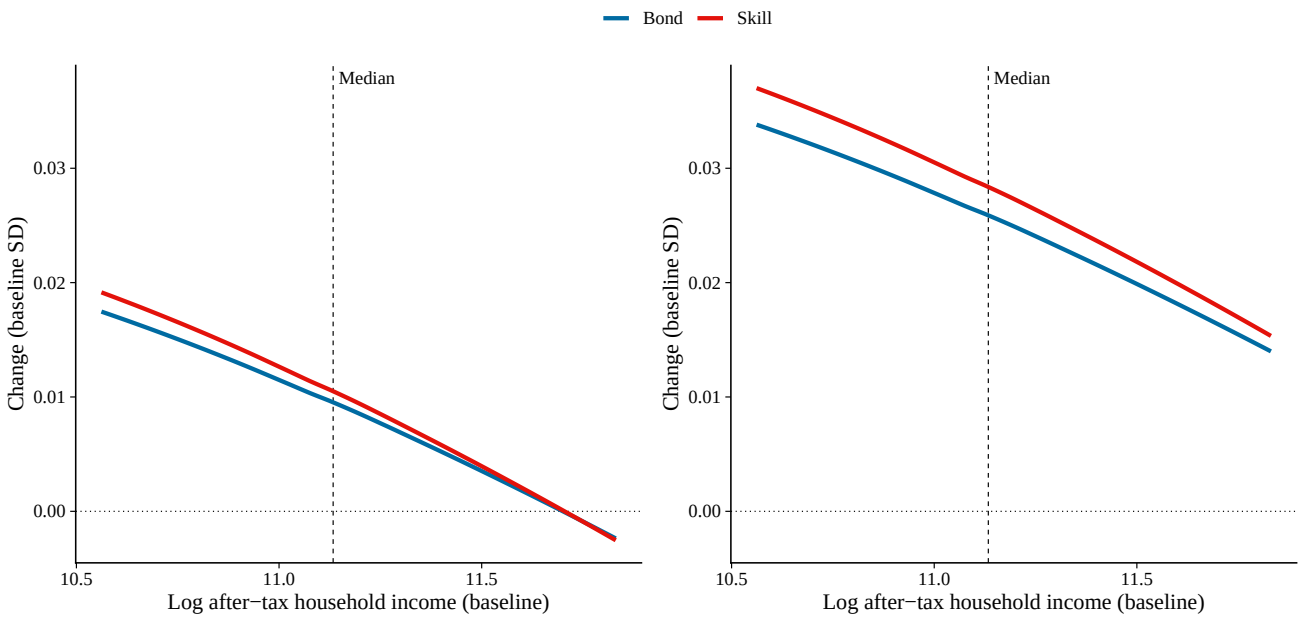
H Progressivity Reform Holding the Median Household's Marginal Tax Rate Fixed

In Section 7.1 I double the progressivity parameter $\mathcal{X}_3$ and readjust $\mathcal{X}_1$ so that the median household's average tax rate stays at its baseline value of 21.6 percent. The median household then pays the same share of its income in tax, but it keeps less of an extra dollar of earnings, since its marginal tax rate rises. In this appendix I readjust $\mathcal{X}_1$ to hold the median household's marginal tax rate fixed instead, following Agostinelli et al. (2024). This requires $\mathcal{X}_1$ to increase, which gives every household roughly 16 percent more after-tax income than in the average-rate version.[58] At the median household's baseline income, the average tax rate falls from 21.6 to about 9 percent. The two versions have the same progressivity and differ only in how generous the system is.

In Figure H.1 I compare the two versions. The marginal rate is held fixed (right panel), skill and the bond rise by more at every income level shown. The gains at the bottom are roughly twice as large as in the average-rate version (left panel). As before, the gains fall with income and skill moves more than the bond. Figure H.2 shows that parents' work, leisure and childcare time are the same as in the average-rate version (Figure 4). Consumption and child expenditures are higher than in the average-rate version and fall below baseline only at the very top percentile. Parents also spend more time with their children at every income level in both versions.

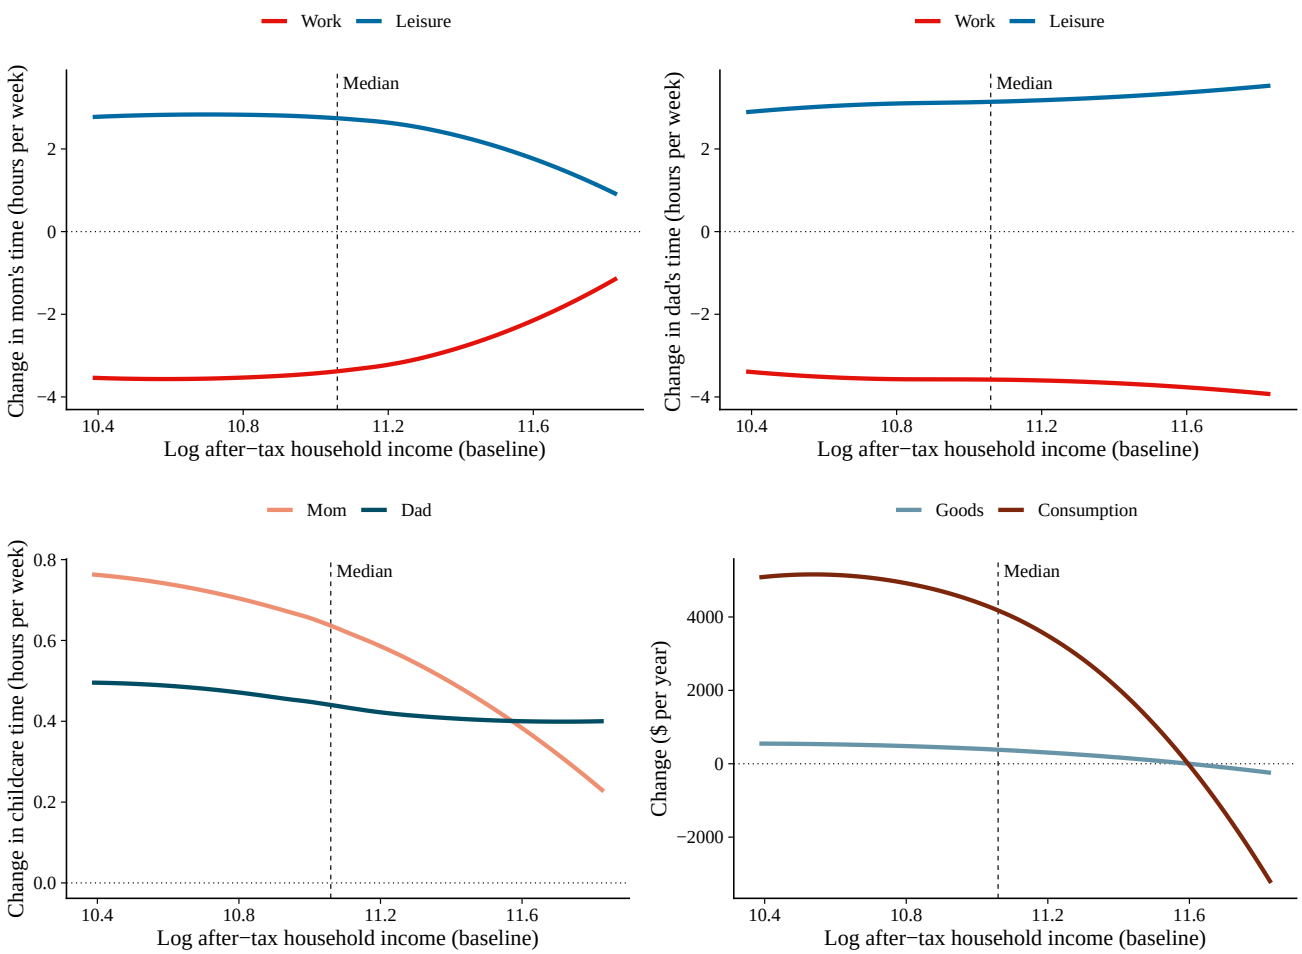
[58] Holding the median household's marginal tax rate fixed sets $\mathcal{X}_1$ to $(1-\mathcal{X}_3)/(1-2\mathcal{X}_3) \approx 1.16$ times its value in the average-rate version, where $\mathcal{X}_3 = 0.122$ is the baseline value. This ratio does not depend on $\mathcal{X}_2$ or on the median household's income, and since $\mathcal{X}_1$ multiplies after-tax income, it applies to every household.

Figure H.1: Skill and Bond Changes under the Progressive Tax Reform, Average versus Marginal Tax Rate Held Fixed, by Household Income



Notes: Each panel reports the change, relative to the baseline economy, in log child socio-emotional skill and in the log parental bond at the end-line, under a doubling of the progressivity parameter $\mathcal{X}_3$. The left panel readjusts $\mathcal{X}_1$ to hold the median household's average tax rate fixed, as in Figure 2. The right panel readjusts $\mathcal{X}_1$ to hold the median household's marginal tax rate fixed instead. The horizontal axis reports log after-tax household income in the baseline economy, and the dashed vertical line marks its median. Both panels share the same vertical scale. Lines are non-parametric (loess) smooths.

Figure H.2: Parental Time and Household Resources under the Progressive Tax Reform, Marginal Tax Rate Held Fixed, by Household Income



Notes: Each panel reports the change, relative to the baseline economy, in the mother's work and leisure (top-left panel), the father's work and leisure (top-right panel), the mother's and father's childcare time (bottom-left panel), and household spending on child goods and on consumption (bottom-right panel), under a doubling of the progressivity parameter $\mathcal{X}_3$ that holds the median household's marginal tax rate fixed. Sample means for mothers (fathers) are 15.2 (7.8) hours of childcare, 27.0 (41.1) of work and 69.9 (63.1) of leisure, and $2,543 of child expenditure. The horizontal axis reports log after-tax household income in the baseline economy. Lines are non-parametric (loess) smooths. Figure 4 reports the same outcomes when the median household's average tax rate is held fixed.

I The Canadian Tax-and-Transfer System

I.1 Canadian Tax and Credit Simulator (CTaCS)

I construct the after-tax income measure used in estimation with the Canadian Tax and Credit Simulator (CTaCS). The CTaCS is an open-source microsimulation tool for the Canadian personal income tax and transfer system developed by Kevin Milligan and collaborators at the University of British Columbia.[59]

The simulator takes the tax year, the province of residence, and the individual's and spouse's pre-tax labour income as inputs. It returns estimated federal and provincial income taxes together with estimated transfers. These transfers could include the Canada Child Benefit (CCB) and the GST/HST credit. I deflate the after-tax income it returns to 2001 dollars and use it to estimate the parameters of the tax-and-transfer function $\mathcal{T}(Y_t)$ in equation (5) by nonlinear least squares.

Calculating after-tax income is shown in the following where I use variables from the documentation. `famtotinc`, defined as the sum of total income across spouses. I construct after-tax income as

$$\mathcal{T}(Y_t) = \texttt{famtotinc} - \texttt{famtotincometax} + \texttt{famtotbens},$$

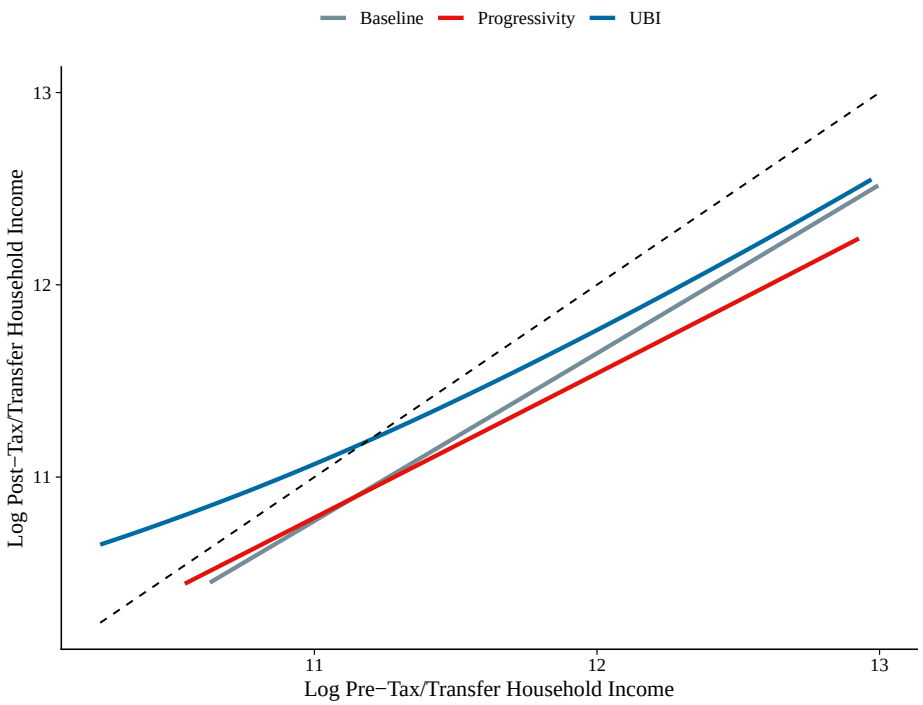
where `famtotincometax` denotes the sum of federal and provincial income tax liability, net of nonrefundable credits, across spouses. `famtotbens` is the sum of federal and provincial refundable transfers, including the Canada Child Benefit, the GST/HST credit, and province-specific child and family benefits.

I.2 Pre- and Post-Tax Household Income under Alternative Tax-and-Transfer Systems

Section 7.1 evaluates two counterfactual reforms to the tax-and-transfer function in equation (5). The first, I raise the benefit floor with a $24,000 UBI. Second, I double the progressivity parameter $\mathcal{X}_3$ while holding the median household's average tax rate fixed. Figure I.1 shows how each tax system maps pre-tax into post-tax household income before any behavioural response. Figure 2 and Table 14 report the simulated outcome effects.

[59] See https://sites.google.com/view/kevin-milligan/home/ctacs for details.

Figure I.1: Pre- and Post-Tax Household Income under the Baseline, Progressivity, and UBI Tax Systems



Notes: The horizontal axis reports log pre-tax household income. The vertical axis reports log post-tax household income under the baseline tax-and-transfer function, the doubled-progressivity counterfactual ($\mathcal{X}_3$ doubled, $\mathcal{X}_1$ readjusted to hold the median household's average tax rate fixed), and the $24,000 UBI counterfactual. The dotted 45-degree line marks post-tax income equal to pre-tax income.