Queen's Economics Department Working Paper No. 1452

Incomplete Markets and Parental Investments in Children

Brant Abbott

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

2-2021

Incomplete Markets and Parental Investments in Children

Brant Abbott
Queen's University
abbottbrant@gmail.com

This draft: February 10, 2021

Abstract. The effect of incomplete markets on parental investments is investigated. Uninsured risk and credit constraints can distort the timing of parental investments, causing them to be delayed relative to what would occur under full-insurance. Age-dependent subsidies, taxes or transfers can all possibly correct this. Analytical results are provided, and a numerical life-cycle model provides quantitative results. Data on ability and parental investment dynamics are used to calibrate the model. A sequence of optimal policy experiments is conducted beginning with a simple reform of the tax and transfer schedule and ending with more complex parent-specific tax parameters and investment subsidies, all of which vary with child age. The final experiment generates substantial improvements in the ability distribution and a consumption equivalent welfare gain that is 2.5 times as large as simply reforming the tax schedule. About 1/4 of this incremental gain results from including child-age dependent policies that alleviate distortions of parental investment timing.

JEL Classification: E2, E24, J24
Keywords: Incomplete Markets, Human Capital, Skill Formation

1 Introduction

Since the seminal work of Cunha and Heckman (2007), a large body of literature has explored issues related to the dynamic process of skill formation in children. One motivation for estimating models of the skill production function has been to help interpret the large estimated effects of early childhood policies reported in the literature (see Heckman and Mosso (2014) for a survey). Estimated models of the skill formation technology are usually embedded in broader models of parental decision making in order to simulate the effects of policy reforms.[1] One approach has been to extend the standard incomplete markets life-cycle model to include multiple periods of parental investment in children's human capital. A common finding in these analyses is that policy interventions targeted to younger children, i.e. preschool aged, are the most effective, much like what is estimated in the empirical policy evaluation literature. Examples of analyses within an incomplete markets framework include Cunha (2013), Caucutt and Lochner (2020), Lee and Seshadri (2019) and Daruich (2018), as well as the seminal work by Restuccia and Urrutia (2004).

An important result regarding skill production functions is that parental investments made at different stages of a child's development are complements (Cunha, Heckman, and Schennach, 2010). This dynamic complementarity means that human capital investments made early in a child's life positively affect the productivity of later investments, and vice versa. While this result would support early childhood policy interventions when parents are myopic, the argument may not follow when parents are optimizing agents. In the standard incomplete markets framework parents are indeed rational agents, and their decisions take dynamic complementarity into account. As such, it is not clear why incomplete financial markets, per se, result in early policy interventions being more effective than general ones. This paper provides analytical and quantitative results that help clarify the key mechanisms embedded within incomplete markets models that lead to these policy conclusions.[2]

[1]This is a burgeoning literature including contributions by Todd and Wolpin (2003), Cunha, Heckman, and Schennach (2010), Del Boca, Flinn, and Wiswall (2013), Agostinelli and Wiswall (2016a), Attanasio, Meghir, and Nix (2015), Attanasio, Meghir, Nix *et al.* (2017), and Pavan (2015), among many others.

[2]Another branch of literature addressing this question asks whether parents are in fact rational (see Cunha, Elo, and Culhane, 2013; Boneva and Rauh, 2018). Empirically, Carneiro and Ginja (2016) and Carneiro, Lopez Garcia, Salvanes *et al.* (2015) show that investments in children are not well-insured against parental income shocks, and that such shocks seem to matter.

This paper develops a model in which parents make annual investments in their child's human capital in the form of both time and expenditure. Throughout the investment process parents face possibly binding credit constraints and uninsurable wage risk. Parents are altruistic and may transfer wealth to their child directly, in addition to making human capital investments. These wealth transfers are restricted to be non-negative, which implies that human capital investments may be inefficient.[3] The model yields an expression describing the dynamics of parental investment decisions. By comparing these dynamics to those that arise under full insurance, novel analytical results on how incomplete markets distort the timing of parental investment are derived. When the model has the feature that good realizations of wage shocks lead to both increased parental consumption and increased investment in children, parents tend to delay those investments relative to what they would choose under full insurance.

The main analytical results are derived by borrowing methods from asset pricing. In particular, the dynamics of investment in children under the counterfactual assumption that parents have access to Arrow-Debreu securities serves as a basis for comparison. Investment dynamics in this full-insurance model are similar to those with uninsured risk, but the analysis shows that there may be differences. A wedge that measures differences in investment timing with and without full-insurance is derived and decomposed into two terms. The first is associated with the Lagrange multiplier that enters the consumption Euler equation when a household is credit constrained. Whenever the borrowing constraint binds investment in the child is delayed compared to the full-insurance model. The second wedge term is associated with uninsured wage risk, which induces uncertainty about future investment levels. This second term may be positive or negative depending on how the realization of wage shocks will affect investment choices. In the case that wage shock realizations and investments are positively correlated, the second wedge term will be negative and also contribute to delayed investment in children (relative to full-insurance). The reason for delayed investment in such a case is that parents take precautionary action, purchasing more of a riskless asset to insure against future investment risk.

[3] See Becker and Tomes (1979, 1986), or the discussion by Mogstad (2017).

Data on parental investment and child ability from the Panel Study of Income Dynamics Child Development Supplement (PSID-CDS) are used to calibrate the model.[4] The method of Agostinelli and Wiswall (2016a; 2016b) is employed in order to make use of several noisy measures of unobserved latent skills. Because the dynamics of parental investment are the focus of the paper, the model is calibrated to match data on parental time and goods investment dynamics, estimates of the dynamics of abilities, as well as other characteristics of households. Because the skill production function cannot be directly estimated, one of the ways the data are utilized is by fitting misspecified auxiliary models of the skill production function, and then matching these to their model counterparts to achieve indirect identification. Calibration also involves estimating a stochastic wage process. This process includes individual fixed effects, which subsume both ability and the elasticity of wages with respect to ability. However, restrictions on the distribution of skills within the quantitative model allow the elasticity to be separately identified as part of the calibration.

The calibrated model is then used to carry out several optimal policy experiments, in which the richness of policy tools available to the planner increases successively across experiments. The welfare criterion considered by the planner accounts for transitional dynamics, and a long-run government budget constraint must be satisfied. First, a standard non-targeted reform of the tax and transfer system is considered in order to establish a baseline with optimal redistribution. This reform increases welfare by the equivalent of a 2.4% consumption increase for all households. When the planner can also subsidize investments in children and target tax policies to parents the consumption equivalent welfare gain is more than twice as large. Much of the additional welfare gain is derived from a large improvement in the ability distribution, where mean ability increases by 37.7% and proportional variation in skills decreases by 14.5%. Finally, policies targeted to parents are allowed to vary with the age of their child, which yields substantial further welfare gains, equivalent to a 5.8% increase in consumption over the benchmark. These further welfare gains are again attributable to improvement

[4] Although the data are similar to those used in Del Boca, Flinn, and Wiswall (2013); their estimates are not directly useful in this paper as their model does not allow for any intertemporal consumption smoothing. Their model also assumes unit elastic substitution between production inputs and non-altruistic parental utility.

of the ability distribution, and, importantly, this incremental improvement in skills is associated with a reduction in the investment wedge described above. Thus, age-dependent tax, transfer and subsidy policies can lead to superior skill, output and welfare outcomes by reducing distortions in the timing of parental investment induced by incomplete markets.

Two closely related papers are those of Cunha (2013) and Caucutt and Lochner (2020), who also study investments in children when markets are incomplete. Caucutt and Lochner focus on a two-period model of goods investment that includes an intergenerational borrowing constraint, parental credit constraints and transitory earnings risk.[5] They provide important results on how binding credit constraints at different stages of life, both for the parent and the child, can lead to underinvestment. The current paper builds on these results by showing how investment timing relates to possibly binding credit constraints and uninsured wage risk. Analytical and quantitative results here show that uninsured parental wage risk can distort investments away from the early years even if credit constraints never bind, provided that the intergenerational transfer constraint could possibly bind. The Caucutt and Lochner results also hold in the model as binding credit constraints exasperate the distortion of parental investment timing. From this we gain new insight into why early policy interventions are more effective than later ones within this type of model.

The remainder of this paper proceeds as follows. Section 2 introduces the child relevant parts of the model, and section 3 derives associated analytical results. Section 4 provides the remaining details of the quantitative model, while the identification strategy and calibration are presented in section 5. Section 6 explores how the incomplete markets frictions distort parental investment timing. Section 7 details the numerical optimal policy experiments, and section 8 concludes.

[5] One way the current paper builds upon Caucutt and Lochner (2020) is by including parental time in a way that does not rule out corner solutions for labor supply, and does not require that efficiency units of time are the same in work and child care. The current paper shows that Caucutt and Lochner's credit constraint results would hold even with this more general incorporation of a parental time decision. For evidence on the importance of parental time in cognitive development see Del Boca, Flinn, and Wiswall (2013), Bernal (2008), Mullins (2016), Moschini (2019), Blandin, Herrington *et al.* (2018), Schoonbroodt (2018), Bono, Francesconi, Kelly *et al.* (2016), Fiorini and Keane (2014), Hsin and Felfe (2014), Erosa, Fuster, and Restuccia (2010).

2 Child Development Model

In this section the part of the model that is relevant for parental investment decisions is presented. In section 4 the remainder of the life-cycle model is presented.

Consider a household i with a parent of age t and child of age t'. Variables are indexed dynamically by the age of the parent, and the child's age is treated as a state variable. The household gains utility from consumption c_{it} and disutility from non-leisure time $\ell_{it} + h_{it}$, where ℓ_{it} is labor supply and h_{it} is time invested in their child. Naturally, these choices are restricted to be non-negative, i.e. $c_{it} \geq 0$, $\ell_{it} \geq 0$ and $h_{it} \geq 0$. Periodic utility is $u(c_{it}, \ell_{it} + h_{it})$ where $u_\ell = u_h < 0$. Labor supply earns a wage w_{it} per unit, and parents pay taxes on earnings according to a function $\tau(w_{it}\ell_{it})$.

Time invested in the child combines with goods investments, denoted x_{it}, according to the following CES aggregator:

$$I_{it} = \left(\alpha_{it} x_{it}^{\delta} + (1 - \alpha_{it}) h_{it}^{\delta}\right)^{\frac{1}{\delta}}. \tag{1}$$

The elasticity of substitution between time and goods investments depends on the parameter δ. The CES weight α_{it} reflects the relative productivity of goods versus time investment, which varies by age and household according to $\alpha_{it} = \alpha(t', \eta_i)$, where η_i is permanent heterogeneity.

The household's decision problem is dynamic. The state variables are the household's assets, a_{it}, the current skills of the child, θ^c_{it}, the current wage offer, w_{it}, the age of the child, t', and parental heterogeneity η_i. The parent's own ability, θ, affects their wage offers, and this dependence is made explicit when the full model is detailed below.

The purpose of investing in the child is to improve their future skills, which enter the continuation value of the household. Skills evolve according to the following production function:

$$\theta^c_{it+1} = \gamma_0 \left(\gamma {\theta^c_{it}}^{\psi} + (1 - \gamma) {I_{it}}^{\psi}\right)^{\frac{1}{\psi}}. \tag{2}$$

The degree of dynamic complementarity between investments at different ages is determined by the parameter ψ. This production function thus combines a flexible degree of dynamic complementarity as in, e.g., Cunha, Heckman, and Schennach (2010) or Caucutt and Lochner (2020), while also including parental time as an explicit choice as in Del Boca, Flinn, and Wiswall (2013).

The recursive decision problem of households during the periods of child-development is:

$$
\begin{aligned}
V_t(a_{it}, \theta^c_{it}, w_{it}, \eta_i, t') &= \max_{\{a_{it+1}, \ell_{it}, x_{it}, h_{it}\}} \Big\{ u(c_{it}, \ell_{it} + h_{it}) + \\
&\quad \beta \mathbb{E}\left[V_{t+1}(a_{it+1}, \theta^c_{it+1}, w_{it+1}, \eta_i, t'+1)\right] \Big\} \ \ s.t. \\
c_{it} &= (1+r)a_{it} - a_{it+1} + w_{it}\ell_{it} - x_{it} - \tau(w_{it}\ell_{it}) \\
a_{it+1} &\geq \underline{a}, \ \text{ eq. (1)}, \ \text{ eq. (2)}.
\end{aligned} \tag{3}
$$

β is the time discount factor, r is the real interest rate, and $\underline{a}$ is the household's borrowing limit.

The final year for investing in a child is when they are $t' = T$ years old because the child will be an independent adult at age $t' = T + 1$. When the child is T years old the parent also decides on a non-negative inter vivos financial transfer, denoted a^c_i. Let $\tilde{V}(a^c_i, \theta^c_{it+1})$ be the expected value function of the child as a new adult. This value depends on both the financial transfer from parents and the child's final skill level. Parents are altruistic with a preference over their child's wellbeing $\kappa\tilde{V}(a^c_i, \theta^c_{it+1})$, where κ is the altruism weight. When the child is age $t' = T$ this altruistic preference enters the continuation value of parents, who solve the following problem at that point:

$$
\begin{aligned}
V_t(a_{it}, \theta^c_{it}, w_{it}, \eta_i, t' = T) &= \max_{\{a_{it+1}, \ell_{it}, x_{it}, h_{it}, a^c_i\}} \{ u(c_{it}, \ell_{it} + h_{it}) + \\
&\quad \beta\kappa\tilde{V}(a^c_i, \theta^c_{it+1}) + \beta\mathbb{E}\left[V^W_{t+1}(a_{it+1}, w_{it+1})\right] \Big\} \ \ s.t. \\
c_{it} &= (1+r)a_{it} - a_{it+1} - a^c_i + w_{it}\ell_{it} - x_{it} - \tau(w_{it}\ell_{it}) \\
a_{it+1} &\geq \underline{a}, \ \ a^c_i \geq 0, \ \text{ eq. (1)}, \ \text{ eq. (2)}.
\end{aligned} \tag{4}
$$

$V^W(\cdot)$ is the notation for the value function of parents after children have come of age, which no longer depends on child specific variables.[6] The assumption that parents can only make non-negative financial transfers ($a_i^c \geq 0$) is standard in the literature, but it is worth emphasizing that it has a non-trivial role in the analysis to follow. Implicitly, this constraint prevents parents and their children from entering into an agreement wherein the parent would invest efficiently in the child's human capital, and the child would repay the associated debt later in life. The fact that a young child cannot commit to such an agreement is a common argument for including such a constraint. Mogstad (2017) provides a great discussion of the seminal work on this mechanism by Becker and Tomes (1979, 1986).

2.1 Simplified Example

In parts of the analysis where the full model is too complicated to allow intuition about the key forces to be understood easily, a simplified version of the model is used. The simplified model assumes that all investments are in the form of expenditure, and that the skill production function is Cobb-Douglas. In terms of the model parameters, the simplified model is equivalent to restricting $\alpha = 1$ and $\psi = 0$. In this case equations 1 and 2 can be replaced by a single equation $\theta_{it+1}^c = \gamma_0 {\theta_{it}^c}^\gamma x_{it}^{1-\gamma}$.

3 Implications of Incomplete Markets

This section demonstrates that the dynamics of investments in children are systematically different when markets are incomplete compared to a world with full-insurance. To be precise, the concept of full-insurance is an extension of the model, in which a full set of state-contingent claims is introduced. For every w_{it+1} that may occur in the next period, a corresponding state-contingent bond may be purchased. The choice of the quantity of bonds associated with a particular future state is denoted $b_{it+1}(w_{it+1})$. This extension is overviewed in detail in online Appendix A. As is standard in Arrow-Debreu equilibrium with only idiosyncratic risk, a risk-free discount factor R can be constructed as the sum of the prices of all state-contingent contracts, and $R = 1/(1+r)$. Furthermore, state-

[6] See online Appendix B for a formal definition of $V^W(\cdot)$.

prices are "actuarially fair," in the sense that the price associated with any w_{it+1} equals $R\pi(w_{it+1})$, where $\pi(w_{it+1})$ is the probability of a future wage occurring. The baseline incomplete-markets model is referred to as the IM model, and the extension including full-insurance through Arrow-Debreu securities is referred to as the FI model. Derivation of the equations presented in this section is explained in online Appendix A.

3.1 IM Model Dynamics

For the IM model presented in section 2, the dynamics of investment satisfy a relatively simple expression whenever there is some degree of dynamic complementarity, i.e. $\psi < 1$, and some degree of time-goods complementarity, i.e. $\delta < 1$.

To derive these dynamics, it is useful to denote the unit cost of I_{it} as $p^I(\tilde{w}_{it}, \alpha_{it}, \delta)$, which depends on the costs of the underlying goods and time investments. Here $\tilde{w} = -u_h(c, \ell + h)/u_c(c, \ell + h)$, which equals the net wage rate (net of the marginal tax rate) whenever parents are in the labor market. When x_{it} and h_{it} are chosen in optimal proportions, the unit cost of investment is:

$$p^I(\tilde{w}_{it}, \alpha_{it}; \delta) = \frac{1}{\alpha_{it}} \left(\alpha_{it} + (1 - \alpha_{it}) \left(\frac{\alpha_{it}}{1 - \alpha_{it}} \tilde{w}_{it} \right)^{\frac{\delta}{\delta - 1}} \right)^{\frac{\delta - 1}{\delta}} . \tag{5}$$

To lighten notation, hereafter let $p^I_{it} = p^I(\tilde{w}_{it}, \alpha_{it}, \delta)$, where the dependence on underlying prices and parameters is implicit. Furthermore, let $u_{c,it} = u_c(c_{it}, \ell_{it} + h_{it})$ be shorthand notation for the marginal utility from consumption of i at time t. Given this notation, the dynamics of investment satisfy:

$$1 = \mathbb{E}\left[\beta \frac{u_{c,it+1}}{u_{c,it}} \times \frac{p^I_{it+1}}{p^I_{it}} \times \gamma_0^{\psi} \gamma \left(\frac{I_{it+1}}{I_{it}} \right)^{1-\psi} \right] . \tag{6}$$

The first two terms inside the expectation are (i) the marginal rate of substitution (MRS) between current and future consumption, and (ii) the relative price of future investment in terms of current investment. Together, these terms measure the relative utility price of future investment in terms of current investment. The third term inside the expectation is the marginal rate of technical substitution (MRTS) between future and current investment, measuring the relative productivity of future investment compared to current, in the sense that

$$\text{MRTS} = \frac{\partial \theta^c_T / \partial I_t}{\partial \theta^c_T / \partial I_{t+1}} = \frac{\partial \theta^c_{t+2} / \partial I_t}{\partial \theta^c_{t+2} / \partial I_{t+1}} = \gamma_0^{\psi} \gamma \left(\frac{I_{it+1}}{I_{it}} \right)^{1-\psi}.$$

Thus, equation 6 says that in expectation the marginal benefit of shifting investments across time equals the marginal cost of this substitution.

3.2 Comparison to FI Model Dynamics

When the full set of Arrow-Debreu securities is available, the dynamics of investment satisfy a similar equation:

$$1 = R \times \mathbb{E} \left[\frac{p^I_{it+1}}{p^I_{it}} \times \gamma_0^{\psi} \gamma \left(\frac{I^*_{it+1}}{I^*_{it}} \right)^{1-\psi} \right]. \tag{7}$$

The notation I^*_{it} indicates investment choice in the FI model, the growth rate of which will be compared to that of I_{it} from the IM model (equation 6).

We would like to understand how the timing of parental investment compares between the IM and FI models, or, in other words, whether incomplete markets cause parents to delay investing in their child. To accomplish this, take an expansion of the expectation in equation 6, as follows:

$$\begin{aligned} 1 =& R \mathbb{E} \left[\frac{p^I_{it+1}}{p^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{I_{it+1}}{I_{it}} \right)^{1-\psi} \right] \\ &+ \left(\mathbb{E} \left[\beta \frac{u_{c,it+1}}{u_{c,it}} \right] - R \right) \mathbb{E} \left[\frac{p^I_{it+1}}{p^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{I_{it+1}}{I_{it}} \right)^{1-\psi} \right] \\ &+ \text{Cov} \left(\beta \frac{u_{c,it+1}}{u_{c,it}}, \frac{p^I_{it+1}}{p^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{I_{it+1}}{I_{it}} \right)^{1-\psi} \right). \end{aligned} \tag{8}$$

The first term of 8 corresponds to the RHS of the FI dynamics equation 7, but the second and third terms are novel to IM investment timing. The second term of 8 reflects the possibility of binding credit constraints. It is negative in the event that the household is currently borrowing constrained, but is otherwise zero. The third term of 8 arises because of uninsured wage risk in the IM model, and could be positive or negative depending on many features of the model. Taken together, the three terms of the expansion indicate that investment timing in the IM and FI models will be equal only if the second and third terms of the expansion sum to exactly zero, otherwise investment timing is distorted by the incomplete markets frictions.

It is useful to define an investment wedge Υ_{it}, which reflects how much investment timing in the IM model deviates from the FI model. Given the discussion in the paragraph above, it makes sense to define this wedge as the gap created in the FI optimality condition 7 when I^*_{t+1}/I^*_t is replaced with I_{t+1}/I_t. This difference will be zero when investment timing is undistorted. The wedge is constructed as:

$$\begin{aligned}
\Upsilon_{it} =& 1 - R\mathbb{E}\left[\frac{p^I_{it+1}}{p^I_{it}}\gamma_0^\psi\gamma\left(\frac{I_{it+1}}{I_{it}}\right)^{1-\psi}\right] && (9)\\
=& R\mathbb{E}\left[\frac{p^I_{it+1}}{p^I_{it}}\gamma_0^\psi\gamma\left(\left(\frac{I^*_{it+1}}{I^*_{it}}\right)^{1-\psi} - \left(\frac{I_{it+1}}{I_{it}}\right)^{1-\psi}\right)\right] && (10)\\
=& \left(\mathbb{E}\left[\beta\frac{u_{c,it+1}}{u_{c,it}}\right] - R\right)\mathbb{E}\left[\frac{p^I_{it+1}}{p^I_{it}}\gamma_0^\psi\gamma\left(\frac{I_{it+1}}{I_{it}}\right)^{1-\psi}\right] && (11)\\
&+ \mathrm{Cov}\left(\beta\frac{u_{c,it+1}}{u_{c,it}}, \frac{p^I_{it+1}}{p^I_{it}}\gamma_0^\psi\gamma\left(\frac{I_{it+1}}{I_{it}}\right)^{1-\psi}\right). && (12)
\end{aligned}$$

The first rearrangement of Υ_{it} in line 10 shows that the sign of the wedge will reflect investment timing differences between the FI and IM models. Unfortunately, the sign is ambiguous, as lines 11 and 12 indicate. Line 11, which reflects possibly binding credit constraints, is non-positive, indicating possibly delayed investment. However, line 12 could be positive or negative because wage shocks affect both parental resources and the unit cost of investment. To gain some intuition it is helpful to examine the simplified case, in which the covariance term has a clear sign, and then return to discussing the general investment wedge.

Under the parameter assumptions of the simplified model (of section 2.1) $I_{it} = x_{it}$ and the unit cost of investment is $p^I_{it} = 1$, in which case the investment wedge simplifies to

$$\Upsilon_{it}|_{\alpha=1,\psi=0} = R\gamma \left(\frac{x^*_{it+1}}{x^*_{it}} - \mathbb{E}\left[\frac{x_{it+1}}{x_{it}}\right]\right) \tag{13}$$

$$= \gamma \left(\mathbb{E}\left[\beta \frac{u_{c,it+1}}{u_{c,it}}\right] - R\right) \mathbb{E}\left[\frac{x_{it+1}}{x_{it}}\right] + \mathrm{Cov}\left(\beta \frac{u_{c,it+1}}{u_{c,it}}, \gamma \frac{x_{it+1}}{x_{it}}\right). \tag{14}$$

The simplified investment wedge will be zero if the expected investment growth rate in the IM model equals the actual (deterministic) rate in the FI model. As with the full model, the wedge can be decomposed into a non-positive term due to binding credit constraints, and a covariance term due to uninsured risk. The sign of the latter term tends to be negative with purely goods investments, as opposed to the ambiguity of this term in the general model. For example, a good wage realization at $t+1$ would lead to a lower marginal utility from consumption and more investment. The covariance term reflects the extent to which future investments might be affected by uninsured wage risk. Such exposure enters the investment wedge because parents take precautionary action, which includes delaying some investment, and thereby insuring the level of resources that will be available at $t+1$.

Return now to the investment wedge in the general model, in particular the covariance term in 12. There are two parameter restrictions that lead to the differences between 12 and 14, $\psi = 0$ and $\alpha = 1$. The sign of the covariance in the simplified model would remain negative if we allow a flexible degree of dynamic complementarity (flexible ψ). The second argument of the covariance in 14 would be the state-contingent MRTS realizations $\gamma_0^{\psi}\gamma(x_{it+1}/x_{it})^{1-\psi}$, which are monotonically increasing in investment growth realizations. If ψ were larger than zero the magnitude of the covariance would be smaller (tending to zero as ψ approaches the perfect substitutes case), and vice-versa, but the sign of the covariance would be unchanged. Thus, the negative covariance term in the simplified case must arise from the exclusion of time investments, i.e. the $\alpha = 1$ restriction. Allowing for time investment introduces the possibility that a bad wage shock would actually lead to increased investment as the (opportunity) cost of investing in children might fall. Whether or not this occurs depends on the parameterization, and hence the covariance term 12 in the general model could be positive or negative.

The possibility of a positive covariance term in the general investment wedge implies that parents might front-load investments relative to the FI model, as opposed to delaying them as they do in the simplified case. What is it about including parental time in the model that might make parents act in a way that appears to be the opposite of precautionary behavior? To understand this, it is helpful to think of the short-run rate of return on I_t as being the realized marginal product of investment in producing skills at $t+2$, i.e. $\partial\theta^c_{t+2}/\partial I_t$. This depends on the realization of I_{t+1} because of dynamic complementarity, and hence there is uncertainty at t. In the case of goods investment only, I_{t+1} is positively correlated with wage realizations, and hence the risk in return on I_t is positively correlated with risk to household resources. Parents then reduce their early investment in favor of holding more of the risk-free bond. However, with time investment added to the model, we have the possibility that I_{t+1} could be negatively correlated with wage realizations. In this case the return on investment is negatively correlated with risk to household income, and risk-averse parents choose to shift investment somewhat towards the earlier period (compared to what is chosen in the FI model).

For the full model we see that the effect of incomplete markets on investment timing is ambiguous overall, and should be studied quantitatively. The effect of incomplete markets on the overall *level* of investment is also ambiguous. Caucutt and Lochner (2020) show this in a somewhat simpler model with possibly binding borrowing constraints. Another example is provided in online Appendix C of the current paper, and in that stylized version of the model early investments are smaller and late investments larger than in the FI model, thus the effect of incomplete markets on the overall level of investment is ambiguous.

3.3 Policy Motives

There are several motives for policy intervention in the model described above. The main policy focus of this paper relates to correction of the investment wedge 9; however, other reasons for policy intervention also exist, in particular the presence of distortionary taxes and the wealth transfer constraint.

Timing wedge correcting policies are discussed first, followed by other policy motives.

3.3.1 Timing Distortions

Age-Dependent Subsidies: Let $S(t')$ be the marginal subsidy rate on goods investments, conditional on the age of the child, for which S_{it} is short-hand notation. The budget constraints in 3 and 4 are adjusted by replacing x_{it} by $(1-S_{it})x_{it}$. Optimal investment timing, described by equation 6, changes because the unit cost of investment becomes

$$\tilde{p}^I(\tilde{w}_{it}, \alpha_{it}; \delta) = \frac{1-S_{it}}{\alpha_{it}} \left(\alpha_{it} + (1-\alpha_{it}) \left(\frac{\alpha_{it}}{1-\alpha_{it}} \frac{\tilde{w}_{it}}{1-S_{it}} \right)^{\frac{\delta}{\delta-1}} \right)^{\frac{\delta-1}{\delta}}. \tag{15}$$

To understand how the subsidy program can change investment timing, consider an alternative version of the wedge in 9 with I_{it} in the baseline IM model replaced by $\tilde{I}_{it}$, which is investment under the subsidy program. This wedge, denoted $\tilde{\Upsilon}_{it}$, is constructed as:

$$\tilde{\Upsilon}_{it} = 1 - R\mathbb{E}_{it+1|it} \left[\frac{p^I_{it+1}}{p^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{\tilde{I}_{it+1}}{\tilde{I}_{it}} \right)^{1-\psi} \right]. \tag{16}$$

As before, the logic is that the wedge will be zero if investment timing equates to that in the FI model. A helpful decomposition of the wedge is as follows:

$$\tilde{\Upsilon}_{it} = \left(\mathbb{E}_{it+1|it} \left[\beta \frac{u_{c,it+1}}{u_{c,it}} \right] - R \right) \mathbb{E}_{it+1|it} \left[\frac{\tilde{p}^I_{it+1}}{\tilde{p}^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{\tilde{I}_{it+1}}{\tilde{I}_{it}} \right)^{1-\psi} \right] \tag{17}$$

$$+ \text{Cov}_{it+1|it} \left(\beta \frac{u_{c,it+1}}{u_{c,it}}, \frac{\tilde{p}^I_{it+1}}{\tilde{p}^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{\tilde{I}_{it+1}}{\tilde{I}_{it}} \right)^{1-\psi} \right) \tag{18}$$

$$+ R\mathbb{E}_{it+1|it} \left[\left(\frac{\tilde{p}^I_{it+1}/\tilde{p}^I_{it}}{p^I_{it+1}/p^I_{it}} - 1 \right) \frac{p^I_{it+1}}{p^I_{it}} \gamma_0^{\psi} \gamma \left(\frac{\tilde{I}_{it+1}}{\tilde{I}_{it}} \right)^{1-\psi} \right]. \tag{19}$$

The first two terms 17 and 18 are the same credit constraint and uninsured risk terms as appear in the IM investment wedge 9. The third "policy term" is novel and results from dynamic variation in the unit cost of investment induced by the subsidy. If the marginal subsidy schedule causes the unit cost of investment to grow faster than in the baseline IM case, the third term will be positive, meaning that the subsidy program induces parents to invest earlier than in the baseline. In the case that the baseline IM investment wedge 9 is negative, subsidies that generate a positive policy term can correct the parental tendency to delay investments.

It is not obvious that a marginal subsidy schedule that decreases with age would induce the unit cost of investment to grow at a faster rate than in the baseline. While the unit cost within a given period is strictly decreasing in S_{it}, changes in α_{it} and $\tilde{w}_{it}$ across time periods complicate matters. Because of this, the simplified model is useful for building intuition about the design of policy. In the simplified case $\tilde{p}^I_{it} = 1 - S_{it}$, and the policy term 19 reduces to

$$R\left(\frac{S_{it} - S_{it+1}}{1 - S_{it}}\right)\gamma\mathbb{E}\left[\frac{\tilde{x}_{it+1}}{\tilde{x}_{it}}\right], \tag{20}$$

which is positive iff $S_{it} > S_{it+1}$. Thus, for the simplified case, we can conclude that a decreasing sequence of marginal subsidy rates could offset the investment wedge (which is negative) unambiguously.

Age-Dependent Marginal Tax Rates: Marginal tax rates on parental earnings that decrease with the age of the child are another possibility to generate $\tilde{p}^I_{it+1}/\tilde{p}^I_{it} > p^I_{it+1}/p^I_{it}$ in the policy term 19. At any given age, a lower marginal tax rate increases $\tilde{w}_{it}$, and thereby increases the opportunity cost of parental time. As was the case with age-dependent S_{it}, changes in α_{it} and w_{it} across time periods mean that $\tilde{p}^I_{it+1}/\tilde{p}^I_{it} > p^I_{it+1}/p^I_{it}$ is not automatically implied by a decreasing sequence of marginal tax rates. In the simplified model, decreasing marginal tax rates would not distort the timing of parental investments, and so age-dependent taxes are only effective for correcting the investment wedge when parental time is a form of investment.

Cash Transfers: Yet another way to reduce the investment wedge is to provide cash transfers to parents. Such transfers would act as a source of insurance for the household, reducing the probability of binding credit constraints and/or fluctuations in the marginal utility from consumption. Reduced consumption variability would directly reduce the covariance term for households that are not well insured. Cash transfers would not affect behavior by changing relative prices, rather they only affect the MRS and likelihood of the credit constraint binding.

3.3.2 Distortions due to taxation

Mispriced parental time: The shadow price of parental time is implicitly reduced when the marginal tax rate is applied to labor income. This distortion only applies to families that are in the labor market. If parents do work, the intratemporal first-order condition for the optimal mix of time and goods investments is,

$$\left(\frac{x_{it}}{h_{it}}\right) = \left(\frac{(1-\tau'(w_{it}\ell_{it}))w_{it}}{1-S_{it}}\frac{\alpha_{it}}{1-\alpha_{it}}\right)^{\frac{1}{1-\delta}}, \tag{21}$$

where the marginal subsidy rate S_{it} has been allowed for. In the baseline IM model $S_{it} = 0$, so the positive marginal tax rate $\tau'(w_{it}\ell_{it})$ distorts the x_{it}/h_{it} ratio downwards. Two ways to reduce this distortion are (i) introduce a positive marginal subsidy rate $S_{it} > 0$, or (ii) adjust the tax schedule so that parents face lower marginal tax rates. If the marginal tax rate is progressive then the marginal subsidy rate would need to increase with income in order to offset distortions in the investment bundles of high earning households.

Distorted Child Value Functions: The expected value function of a child $\tilde{V}(a_i^c, \theta_i^c)$ is also affected by the tax system. Taxation of labor income reduces the return to investing in a child's skills, in the sense that $\partial\tilde{V}(a_i^c, \theta^c)/\partial\theta^c$ falls as marginal tax rates increase. Following the logic presented in Bovenberg and Jacobs (2005), this feature of the model motivates subsidies on forms of investment that are not tax deductible. In the model here, x_{it} would not be tax deductible, implying a motive for goods investment subsidies. However, h_{it} would be classified as tax-deductible because the opportunity cost of parental time is affected by the tax distortion in a similar way as the child's future earnings will be (unless the parent is out of the labor force). Bovenberg and Jacobs also show that if an investment is tax deductible then it would not be subsidized under optimal policy, as long as non-deductible investments are subsidized. Thus, to the extent that their theoretical results hold in the current model, the distortionary earnings tax that a person will face as an adult motivates a subsidy on the goods investments their parents make in them when they are a child.

3.3.3 Intergenerational Transfer Constraint

As in the classic result of Becker and Tomes (1986), whenever the intergenerational borrowing constraint binds the rate of return on human capital investment exceeds that on financial assets. This motivates policy to reduce the price of investment in general, which could be through a subsidy on purchased investments, or higher marginal income tax rates. Lump-sum cash transfers to parents might also redistribute wealth in such a way as to reduce the effect of the constraint.

4 Quantitative Model

The remainder of the life-cycle model is now specified. A key feature is that the distribution of skills is endogenous and evolves over time through intergenerational connections. Several parametric assumptions are required in order to operationalize the model, and identification of those parameters based on investment and ability dynamics data is discussed in detail. Finally, the calibrated model parameters and the fit of the model are discussed.

4.1 Full Life-Cycle Model

Demographic Structure: Each agent lives for 85 years, and there is a unit mass of agents in each age cohort. From their first year to age $T = 17$ agents are dependents being cared for by their parents. From age 18 to 64 agents are workers, and for part of working life they care for their own child, specifically from age 30 to 46. From age 65 to 85 agents are retired and collect a pension τ_{pen}. The skill distribution is an endogenous object in the model as the investment decisions of current parents determine the skills/wages of future workers, who will in turn decide how much to invest in their own children. For tractability, it is assumed that each agent has one child. Recall that t' denotes the age of a child and t denotes the age of a parent. Given the assumption that agents become parents at age 30, it will always be the case that $t = t' + 30$.

Preferences: Each mature agent (18 or older) has preferences over the consumption and time allocation of their household according to the utility function $u(c, \ell + h) = \frac{1}{1-\sigma}\left(c - \nu\frac{(\ell+h)^{1+\chi}}{1+\chi}\right)^{1-\sigma}$. The risk aversion parameter is σ, the Frisch elasticity of labor supply is $1/\chi$, and the time-preference discount factor is β.

Utility from a child is altruistic, with weight κ applied to the expectation of the child's value function when they turn 18: $\kappa\tilde{V}(a_i^c, \theta_{it+1}^c)$. This utility term enters a parent's decision problem when the child reaches age $t' = T = 17$, as in equation 4. The expected value function at age 18 is an endogenous object in the model, and the altruistic utility a parent earns depends on the investment and financial transfer decisions they make.

Wages and Endowments: All agents begin their adult life with an ability level θ_i, a wealth endowment a_{i18}, and an initialized wage shock $z_{i18} = 0$ (abilities and wealth endowments are determined by parents). After age 21, abilities and shocks combine to determine wages according to the specification $\ln w_{it} = \overline{w} + \lambda\ln(\theta_i) + z_{it} + \xi_t$. λ is the elasticity of wages with respect to ability, z_{it} follows an $AR(1)$ process, and ξ_t is an age-earnings profile. The $AR(1)$ process generating z_{it} has persistence ρ and shock variance σ_u^2, and is approximated by a finite-state Markov chain using Rouwenhorst's method. In the first four years of adult life, i.e. ages 18-21, wages are determined according to $\ln w_{it} = \overline{w} + (\lambda/2)\ln(\theta_i) + z_{it} + \xi_t$, which means the full effect of ability on wages is delayed until after the college years. This serves to roughly approximate how the return to ability evolves over the early part of adult life without formally including college in the model.

In order to make optimal policy analysis tractable, a simple tax function is utilized: $\tau(w_{it}\ell_{it}) = \tau_1 w_{it}\ell_{it} - \tau_0$. Here τ_1 is a flat marginal tax rate and τ_0 is a lump-sum transfer. This is progressive in the sense that the average tax rate is an increasing function of earnings.

Parenting Productivity: The relative productivity of goods versus time investments, α_{it}, is parameterized so that there is an age-dependent component, $\overline{\alpha}_t$, and a permanent shock component, ζ_i. The age-dependent component is deterministic and evolves according to the following function:

$$\overline{\alpha}_t = \frac{\exp(a_0 + a_1 \cdot t')}{1 + \exp(a_0 + a_1 \cdot t')}, \tag{22}$$

where a_0 and a_1 are parameters to be calibrated. It is assumed that $\ln \zeta_i$ is normally distributed with mean zero and variance σ_ζ^2. The two components combine in a way that ensures time and goods weights in the CES aggregator remain on the unit interval and sum to one:

$$\alpha_{it} = \frac{\overline{\alpha}_t}{\overline{\alpha}_t + (1 - \overline{\alpha}_t)\zeta_i}.$$

A household's draw of ζ_i positively affects the relative productivity of parental time, while the sign of a_1 determines whether the relative productivity of goods investment increases or decreases as a child gets older. At the time a child is born parents learn what their draw of ζ_i is.

Skill Production Productivity: The overall productivity of the skill production process is determined by γ_0, which is allowed to vary with the age of the child. Specifically, the model includes $\gamma_0 = \mathcal{G}_0 + \mathcal{G}_1 \times t'$, which allows skill production productivity to rise or fall with the age of the child.

Initial Ability: There are two levels of initial ability that a child may be born with: $\theta^c_{low,0}$ or $\theta^c_{high,0}$. To capture possible genetic endowment effects I allow the probability of drawing high initial skills to be a logistic function of parental ability:

$$Pr(\theta^c_{high,0}|\theta_i) = \frac{\exp(\varphi_0 + \varphi_1 \times \theta_i)}{1 + \exp(\varphi_0 + \varphi_1 \times \theta_i)}, \tag{23}$$

where φ_0 and φ_1 are parameters to be calibrated. According to this function, the probability that a child will have high initial ability increases with parental ability if $\varphi_1 > 0$. Parents learn the initial ability of their child when the child is born.

Decision Problems: Other than the periods of child development, the decision problems of agents are standard. The details of these problems are provided in online Appendix B. The life-cycle problem can be solved recursively, given an initial guess of the expected value function of a new adult, $\tilde{V}$. This solution generates a new guess for the function $\tilde{V}$, and the process repeats until a fixed point is found.

5 Calibration

5.1 PSID-CDS Data and Ability Measurement:

Data are drawn from the 2002 and 2007 waves of the Panel Study of Income Dynamics - Child Development Supplement (PSID-CDS), as well as the 1997-2007 main PSID family data. The unit of observation in the PSID-CDS is the child, and the data include measures of ability, parental expenditure on the child, and a time-diary that includes information on time spent actively engaged with parents. The supplement can be linked to the main PSID family data, which includes information on the labor supply and earnings of parents.

The main advantage of the PSID-CDS is the time-use survey, which provides comprehensive information on the activities of children and the inclusion of parents in those activities. The time diaries include 24 hours of information for two representative days, one a weekday and the other a weekend day. For each activity a child does during the survey day information is collected on the duration of the activity, who was participating with the child, who was around but not participating, and where the activity took place. The time inputs of parents are defined as the total duration of activities in which they are actively participating with the child. An estimate of weekly totals is based on five times the weekday time allocation plus two times the weekend time allocation.

Observed child-specific expenditures were limited in 1997; however, in the 2002 and 2007 waves many additional expenditure items were added, resulting in a more reasonable total expenditure variable. Items include tuition costs, tutoring expenses, lessons (e.g. music lessons), sports, community/religious groups, toys/books, vacations, school supplies, clothing, transportation and daycare expenses. The 2002 wave also includes spending on food specifically for the child, but this item was not included in 2007 and so it is excluded in 2002 as well for consistency. All dollar valued data are converted to real 2002 dollars.

While the model includes is a single wage rate for each household, in the data there are often two parents working. In these cases, the wage at the margin of adjustment is approximated by taking the lowest wage between the two spouses. The idea is that the spouse with the lower wage would be the most likely to scale back their labor supply in order to provide more time investment to their child. Labor supply is measured at the household level. These assumptions and the fact that other variables are observed at the level of an individual child, allow for interpretation of the data in the single parent/child model, even though households in the data often have multiple parents/children.

Measurement of Ability: The data include multiple measures of latent ability. The strategy proposed by Agostinelli and Wiswall (2016a) is employed to use these noisy measures to identify estimates of mean log-ability by age, i.e. the age profile $\{E[\ln(\theta^c_{it})|t']\}_{t'=6}^{16}$, and to estimate the regressions based on equation 26. First, three age-invariant measures $Z_{m,it}$ of θ^c_{it} are selected, where m indexes the measure. Each measure is assumed to relate to latent ability as follows:

$$Z_{m,it} = \mu_m + \Lambda_m \ln \theta_{it} + \epsilon_{m,it}. \tag{24}$$

μ_m and Λ_m are measurement parameters, and $\epsilon_{m,it}$ is measurement error. Also define residual measures $\tilde{Z}_{m,it} = (Z_{m,it} - \mu_m)/\Lambda_m$ and $\tilde{\epsilon}_{m,it} = \epsilon_{m,it}/\Lambda_m$, such that $\ln(\theta_{it}) = \tilde{Z}_{m,it} - \tilde{\epsilon}_{m,it}$. The three measures are Letter-Word (LW), Applied Problems (AP) and Paragraph Comprehension (PC) raw scores.

Begin by normalizing mean log-ability to zero at the first observed age in the sample, in this case age 6. Next, normalize one of the ability loadings to unity, in this case the Letter-Word loading is normalized to $\Lambda_{LW} = 1$. These two assumptions identify the intercept μ_{LW} from $E[Z_{LW,it}|t' = 6]$. Furthermore, given the normalization of Λ_{LW}, the other loading parameters Λ_{AP} and Λ_{PC} can be recovered by

$$\Lambda_m = \frac{\text{Cov}(Z_{m',it}, Z_{m,it})}{\text{Cov}(Z_{LW,it}, Z_{m,it})}, \tag{25}$$

where m and m' are either PC or AP, depending on the case. Given these Λ_{PC} and Λ_{AP}, μ_{LW} and μ_{PC} can be recovered using data on children at age 6, just like μ_{LW} was recovered. Finally, because the measurement parameters have now been recovered, we can use averages of $\tilde{Z}_{m,it}$ to estimate $\{E[\ln(\theta^c_{it})|t']\}_{t'=6}^{16}$ (recall that $\tilde{Z}_{m,it} = \ln(\theta_{it}) + \tilde{\epsilon}_{m,it}$).

Summary Statistics: The final sample includes all observations with complete data on the variables described above across both the 2002 and 2007 waves of the PSID-CDS. Table 1 provides summary statistics on child abilities and investments by sample year.

Variable	2002 Mean	2002 Median	2007 Mean	2007 Median
Child's Age	9.22	9	14.08	14
Letter-Word	37.3	39	46.8	48
Paragraph Comprehension	20.8	22	27.7	28
Applied Problems	31.4	32	40.1	40
Goods Investment (2002 dollars)	$3360	$2400	$4037	$2701
Time Investment (annual hours)	1525	1361	1145	936
Sample Size	763		763	

Table 1: Summary Statistics from the PSID-CDS Sample.

5.2 Wages Shock Process

To estimate the wage shock process, observed parental log-wages are first regressed on a quadratic function of age (ξ_t).[7] The associated residuals, w_{it}^r, are then used to estimate the variance of $\lambda \ln(\theta_i)$, as well as the parameters of the $AR(1)$ process that generates z_{it}, i.e. ρ and σ_u^2. By assumption we have $\text{Var}(z_{i18}) = 0$, and then can compute $\text{Var}(z_{it}) = \rho^2 \text{Var}(z_{it-1}) + \sigma_u^2$ at any later age. Connecting this to the wage residuals, the following relationship holds: $\text{Var}(w_{it}^r) = \text{Var}(z_{it}) + \lambda^2 \text{Var}(\ln(\theta_i)) + \sigma_m^2$, where σ_m^2 is the variance of iid measurement error. One can also compute covariances between wage residuals at different ages as $\text{Cov}(w_{it}^r, w_{it-s}^r) = \rho^s \text{Var}(z_{it-s}) + \lambda^2 \text{Var}(\ln(\theta_i))$. It can easily be shown that with enough age variation in the panel ρ, σ_u^2, $\text{Var}(\lambda \ln(\theta_i))$ and σ_m^2 are identified.[8] Using data on 25-65 year-old PSID-CDS parents in the 1997-2007 waves of the main PSID data,[9] I estimate $\rho = 0.987$, $\sigma_u^2 = 0.0173$ and $\text{Var}(\lambda \ln(\theta_i)) = 0.1302$. These estimates are based on an equally weighted minimum distance estimator using all available variances and covariances under the specified sample and age range. The ρ and σ_u^2 estimates are directly employed in the model, and the estimate of $\text{Var}(\lambda \ln(\theta_i))$ is used as a target in the internal part of the calibration.

5.2.1 Externally Calibrated Parameters

The following preference parameters are externally calibrated: the time discount factor is set to $\beta = 0.975$, the risk aversion parameter is set to $\sigma = 2$, and the Frisch elasticity of labor supply is set to $1/\chi = 1/2$ (at the intensive margin).[10] The exogenous real interest rate is set to 2.5%.[11] The marginal tax rate on earnings is set to $\tau_1 = 0.27$.[12] No unsecured credit is allowed in this economy, thus $\underline{a} = 0$.

[7] Recall that if both parents work the lower wage spouse is assumed to be the margin of adjustment, and their wage data is utilized. Naturally, it is also their age data that is utilized.

[8] See Tonetti (2011) for proof of identification and a short literature review.

[9] I use these years as the PSID-CDS runs from 1997 to 2007.

[10] These parameters are similar to those assumed in other recent literature, e.g. Heathcote, Storesletten, and Violante (2020).

[11] Treasury Inflation Protected Security yields (a measure of real returns) were close to this value over the sample period, ranging from about 1.75% to 3.5%. Source: FRED Data.

[12] As suggested in Domeij and Heathcote (2004).

5.2.2 Internally Calibrated Parameters: Identification

The internally calibrated parameters are listed in Table 2, and the moments that identify these parameters are listed in of Table 3. These moments include the estimated parameters of a misspecified translog skill production function; age patterns of child ability, goods investment and time investment; covariances between relevant features of the data; and means of some basic variables. Although identification is described in terms of particular moments being identified by particular parameters, all of the moments jointly identify all of the internally calibrated parameters. The final paragraph of this subsection discusses the objective function that is formally minimized in order to determine the calibrated parameters.

First, some basic features of the data are matched by the model. Units of time are normalized such that one unit in the model corresponds to 11,680 real world hours (which is the annual total based on 16 hours per parent per day). The labor supply parameter ν is calibrated such that average labor supply is 0.262, which corresponds to average annual household labor supply of 3,060 hours.

Second, consider the parameters of the production function in equation 2, which include $\mathcal{G}_0$, $\mathcal{G}_1$, γ and ψ. Given observation of investments and ability at an annual frequency, one could directly estimate the production function, or, given enough measures of ability, follow Cunha, Heckman, and Schennach (2010). However, the PSID-CDS provides measures of ability and investment only every five-years, thus an indirect approach is more suitable. The main identification comes from matching estimates of the parameters of two misspecified translog production functions to those estimated in the data. The first version is

$$\ln(\theta_{it+5}) = b_0 + b_1 \times t'_t + b_2 \ln(\theta_{it}) + b_3 \ln(x_{it}) + b_4 \ln(\theta_{it}) \ln(x_{it}), \tag{26}$$

while the second version replaces x_{it} terms with h_{it}.[13] In estimating this model on data, $\tilde{Z}_{LW,it}$ is

[13] Difficulties with multicollinearity underly the decision to incorporate time and goods investments in separate auxiliary models, as opposed to estimating a single auxiliary production function with both types on inputs.

used to replace $\ln(\theta_{it})$, while $\tilde{Z}_{AP,it}$ is used as an instrument to address measurement error in $\tilde{Z}_{LW,it}$. b_0 and b_1 are informative about $\mathcal{G}_0$ and $\mathcal{G}_1$, b_2 and b_3 are informative about γ, while the importance of complementarities, measured by b_4, is informative about ψ. Some additional over-identifying variation is also included help pin down the scale of abilities, specifically the age profile of average log-ability, $\{E[\ln(\theta^c_{it})|t']\}_{t'=6}^{16}$.[14]

Next, consider the aggregator for time and goods investments in equation 1, which depends on the parameters a_0, a_1, δ and σ^2_ζ. These parameters are largely identified by exploiting the following relationship, which holds in the model for households that are in the labor force:

$$\ln\left(\frac{x_{it}}{h_{it}}\right) = \frac{1}{1-\delta}\ln\left(w_{it}(1-\tau_1)\right) + \frac{1}{1-\delta}\left(a_0 + a_1 \times t'\right) - \frac{1}{1-\delta}\ln\zeta_i. \tag{27}$$

One might consider a regression with net-wages and child-age as dependent variables; however, this equation will not hold for households outside the labor force, leading to a selection bias. With large enough ζ_i and/or low enough w_{it} the household will supply no labor, and equation 27 will not hold. However, this selection bias arises in both the data and model, and so the true parameters can be identified indirectly by matching the biased regression coefficients in the model to those computed using the data.[15] Next, use $\widetilde{\ln}(x_{it}/h_{it})$ to denote residual variation in $\ln(x_{it}/h_{it})$, after subtracting the variation explained by child-age and net-wages. This residual variation is largely generated by heterogeneity in parental productivity (ζ_i). The autocovariance $\text{Cov}(\widetilde{\ln}(x_{it}/h_{it}), \widetilde{\ln}(x_{it+5}/h_{it+5}))$ is used to gauge the variance of these persistent effects, and σ^2_ζ is identified by matching the model counterpart of this moment to the data.[16] Lastly, as over-identifying information, the age-profile of the ratio of goods to time investments from age $t' = 6$ to 16 is included.

The remaining wage process parameters are λ and $\overline{w}$. Recall that the wage process estimation

[14] There are a useful number of observations at each age between 6 and 16, but not younger.

[15] In both the data and the model, $w_{it}(1-\tau_1)$ is replaced by w_{it}, generating the same biases in both regressions.

[16] I use the autocovariance so that any transitory measurement error is excluded from the variance measure.

yielded an estimate of Var$(\lambda \ln(\theta_i))$, but did not isolate λ.[17] However, the model counterpart of this variance can be computed, and λ can be calibrated in order to match it to the estimate from data. This works well because Var$(\ln(\theta_i))$ is largely pinned down by other features of the model, for example the initial distribution of ability and the skill production function. Next, the parameter $\overline{w}$ is identified by matching the median wage rate to what is observed in the data. Given the scale of ability is pinned down by the production function parameterization, and λ is disciplined as just described, $\overline{w}$ determines the scale of wages.

Each household receives a lump-sum tax rebate τ_0 each period. This parameter captures the progressiveness of the tax schedule by compressing the distribution of post-tax income relative to pre-tax. Heathcote, Perri, and Violante (2010) report that the variance of log post tax/transfer income is about 62% as large as the variance of log earnings in the Current Population Survey in the early 2000s. The parameter τ_0 is calibrated in order to replicate this statistic.

Next, consider the altruism parameter, κ. This parameter affects both the level of parental investment in children and the extent to which parents transfer wealth to their children. As both of these are important, the calibration targets both the age profile of goods investments from age $t' = 6$ to 16,[18] and an external data point on the proportion of children who receive positive inter vivos transfers, reported in Abbott, Gallipoli, Meghir *et al.* (2019, Appendix G).[19] The proportion of positive transfers is important to match because the rate at which the non-negative intergenerational transfer constraint binds greatly influences the scope for policy intervention.

Finally, consider the distribution of initial ability. Any increase in φ_1, i.e. any increase in the probability that a high ability parent has a child with high initial ability, will increase the correlation between parental income and child ability. Thus, φ_1 can be identified by matching the model counterpart of the child-ability elasticity with respect to parental earnings, i.e. $\text{Cov}(\ln(\theta^c_{it}), \ln(y_{it})|t' =$

[17]Note that we do not observe measures of ability for parents, so a simple IV regression of wages on an ability measure cannot be used to estimate λ directly.

[18]In order to make investment values comparable between data and model, average investment at each age is normalized by overall average income, which creates unitless ratios to match.

[19]That work uses the NLSY97 data, in which transfers from parents are reported annually. The measure includes any transfer received from parents up to age 22, and finds that 75.1% of children receive some transfer from their parents.

$16)/\text{Var}(\ln(y_{it})|t' = 16)$, to what is measured in the data. The support of initial abilities is set to $\left\{\theta^c_{low,0}, \theta^c_{high,0}\right\} = \left\{0.5 \times \overline{\theta}^c_0, 1.5 \times \overline{\theta}^c_0\right\}$, where $\overline{\theta}^c_0$ is a parameter that determines the scale of initial abilities.[20] Measures of ability are not observed in the data until age 6, thus $\overline{\theta}^c_0$ is calibrated to match the earlier part of the age-profile of ability $\{E[\ln(\theta^c_{it})|t']\}^{16}_{t'=6}$. Finally, φ_0 is calibrated so that the overall probability of $\theta^c_{low,0}$ being drawn is 0.5.

The model is formally calibrated by selecting the set of parameters that minimizes the distance between the data moments described above and their model counterparts. An identity weight matrix is used in doing this; however, some of the moments are scaled so that their order of magnitude is consistent with the remaining parameters. For example, the median wage is scaled by 100, which ensures that this moment is on the $(0.1, 1)$ interval. This is important as it ensures no particular moment receives a large implicit weight due to its order of magnitude being large. The model is overidentified with a total of 55 moments to identify 16 parameters.

5.2.3 Internally Calibrated Parameters: Results

Calibration results are presented in Tables 2 and 3, as well as Figure 1. Although the fit is not perfect due to the overidentification, the model fits the data reasonably well. The ratio of goods to time investment increases with the data as illustrated in Figure 1, but, of course, the randomness in the data is not captured by the model, which exhibits an inherently smooth profile. Very much like Del Boca, Flinn, and Wiswall (2013), the productivity of goods investment increases with the age of the child. This results in goods investments that increase with the age of the child like in the data. Investments in different time periods are found to be somewhat more substitutable than the Cobb-Douglas case, similar to Caucutt and Lochner (2020). This positive value of ψ arises from targeting the negative point estimates of b_4 in both versions of equation 26. The calibrated value of $\mathcal{G}_1$ is negative, as are the targeted point estimates of b_1 in equation 26, which means that the overall productivity of the investment process falls as children get older. This decrease in γ_0 also allows the model to replicate

[20] The choice to model the support of initial ability as $+/-$ 50% of the mean is an arbitrary assumption, but without any data on variation in abilities at birth such an assumption is necessary.

Parameter	Description	Value
ν	Disutility weight on non-leisure time	7.516
$\overline{w}$	Wage equation constant	0.793
λ	Wage elasticity of ability	0.592
κ	Altruism weight	0.402
$\mathcal{G}_0$	Skill production function TFP constant	2.492
$\mathcal{G}_1$	Skill production function TFP age loading	-0.033
γ	Skill production function weight on θ_{it}^c	0.751
ψ	Dynamic complementarity parameter	0.396
δ	Time-goods aggregation complementarity	-0.561
a_0	Constant in α_{it} equation	-1.901
a_1	Age parameter in α_{it} equation	0.142
σ_ζ^2	Parental time productivity distribution	1.915
τ_0	Lump-sum tax rebate	0.880
φ_1	Initial child ability dependence on parental ability	0.017
φ_0	Initial child ability constant	-0.326
$\overline{\theta}_0^c$	Mean initial child ability	0.086

Table 2: Internally Calibrated Parameters.

the curvature of the age-profile of mean ability in the bottom panel of Figure 1.

The remaining parameters also lead the model to fit the data described above well. The elasticity of substitution between time and goods investments is 0.64, indicating a moderate degree of complementarity. A relatively large variance for parental time productivity shocks is needed in order to explain residual heterogeneity in time-goods ratios. The modelling of the initial skill distribution does a good job to explain the correlation between parental income and child skills at age 16. As reported in Table 4, the parameters of the initial skill distribution generate a strong relationship between parental ability and the initial skill draw, which can be interpreted as the innate component of skills.

6 Quantifying Timing Distortions

This section quantitatively assesses how much investment timing differs between the IM and FI models, and what these differences imply for skills. First, age-profiles of investments and skills are plotted for the two models. Second, to help assess the contributions of binding credit constraints vis-a-vis

Moment Matched	Data Value	Model Value
Age-profile of child ability	Reported in Figure 1	
Age-profile of goods investment	Reported in Figure 1	
Age-profile of goods-time ratio	Reported in Figure 1	
Equation 27 regression: constant	-2.841	-2.691
Equation 27 regression: wage coeff.	0.638	0.616
Equation 27 regression: child-age coeff.	0.100	0.098
Equation 26 reg. with x_{it}: b_0	2.533	2.213
Equation 26 reg. with x_{it}: b_1	-0.0475	-0.0367
Equation 26 reg. with x_{it}: b_2	0.0739	0.0907
Equation 26 reg. with x_{it}: b_3	0.456	0.528
Equation 26 reg. with x_{it}: b_4	-0.0466	-0.0008
Equation 26 reg. with h_{it}: b_0	2.228	2.303
Equation 26 reg. with h_{it}: b_1	-0.0289	-0.0093
Equation 26 reg. with h_{it}: b_2	0.153	0.127
Equation 26 reg. with h_{it}: b_3	0.394	0.515
Equation 26 reg. with h_{it}: b_4	-0.0074	-0.0145
Residual investment ratio autocovariance: $\text{Cov}(\widetilde{\ln}(x_{it}/h_{it}), \widetilde{\ln}(x_{it+5}/h_{it+5}))$	0.271	0.288
Ability-earnings elasticity: $\text{Cov}(\ln(\theta^c_{it}), \ln(y_{it})\|t' = 16)/\text{Var}(\ln(y_{it})\|t' = 16)$	0.293	0.314
Proportion of positive inter vivos transfer	0.751	0.750
Wage fixed-effect variance: $\text{Var}(\lambda \ln(\theta_i))$	0.130	0.127
Average labor supply: $E[\ell_{it}]$	0.262	0.260
Median wage: $\text{Med}(w_{it})$	16.02	16.27
Probability of low initial ability: $Pr(\theta^c_{low,0})$	0.500	0.504
Progressivity Ratio: $\text{Var}(\ln(w_{it}\ell_{it}(1 - \tau_1) + \tau_0))/\text{Var}(\ln(w_{it}\ell_{it}))$	0.620	0.623

Table 3: Moments targeted in calibration.

	$Pr(\theta^c_{low,0}\|\theta^{quart})$	$Pr(\theta^c_{high,0}\|\theta^{quart})$
$\theta^{quart} = 1$	0.6398	0.3602
$\theta^{quart} = 2$	0.5465	0.4535
$\theta^{quart} = 3$	0.4566	0.5534
$\theta^{quart} = 4$	0.3758	0.6242

Table 4: Initial Skill Distribution: The importance of the parameter φ_1 is illustrated by showing how the probability of a low skill initial draw varies by parental ability quartile (θ^{quart}).

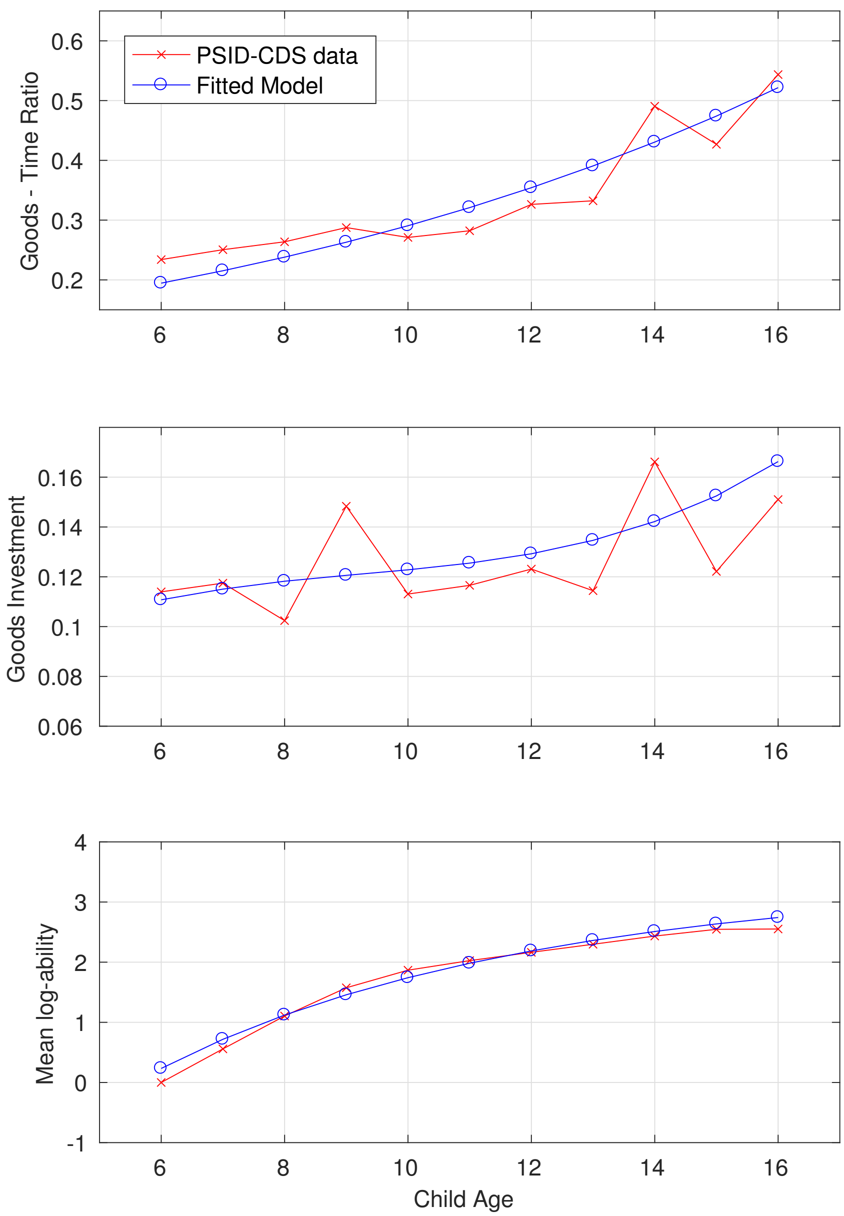

Figure 1: The fit of the model to age profiles of (i) the ratio of mean goods investment to mean time investment, (ii) mean goods investment, and (3) mean log-ability. For comparability, in both data and model goods investment is normalized by mean earnings and time investment is normalized by mean labor supply.

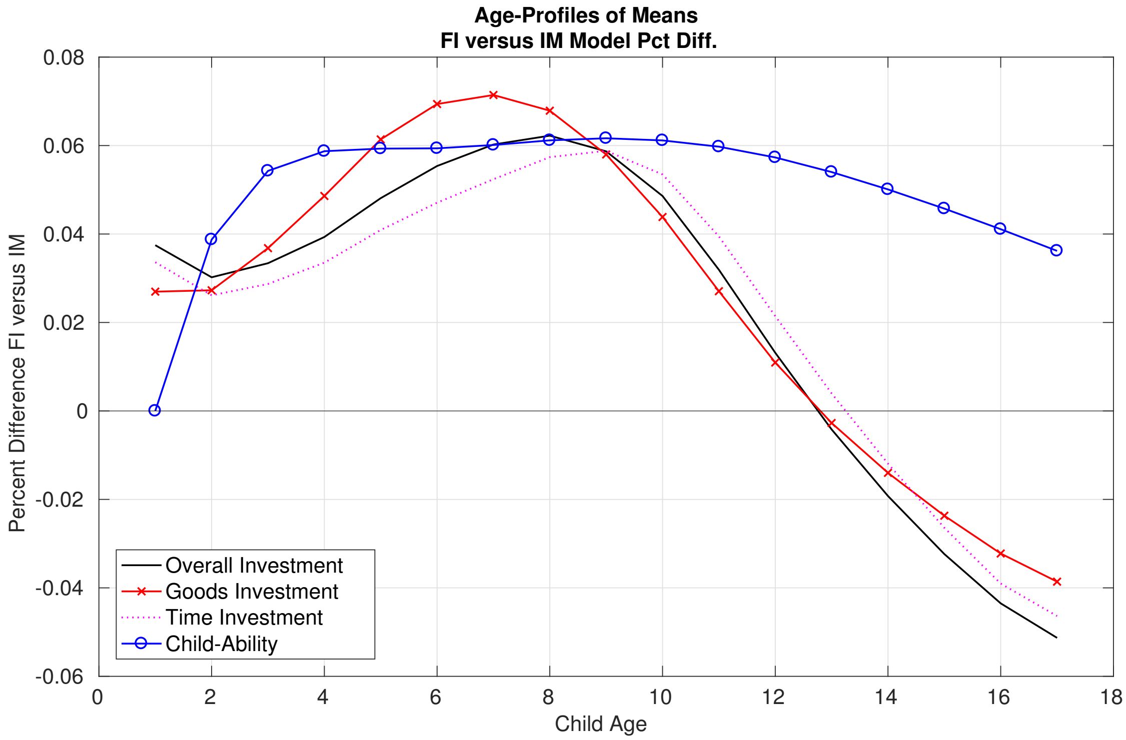

Figure 2: Percentage differences in means of investments and child-ability by age. For example, the overall investment age-profile plots $E[I_t^*]/E[I_t] - 1$ at each child-age. Importantly, the initial distribution of state-variables is identical across the two versions, but they progress differently as different decisions are made.

uninsured risk, the two terms of the investment wedge in equations 11 and 12 are computed for the baseline model, and their distributions are analyzed.

Figure 2 plots the percentage differences in age-profiles of mean investment and child-ability between the FI and IM models. A positive number indicates a larger value under FI model decisions, for example the peak child-ability difference of just over 0.06 indicates that average skill is at most just over 6% higher in the FI version. The IM versions of the age-profiles are from the calibrated benchmark model. The FI versions are constructed by simulating a single cohort of parents, among whom the distribution of state-variables when the child is born is the same as in the benchmark IM model. In this way, initial conditions are held constant, but investment dynamics differ as described in section 3.2.

Figure 2 shows that differences in mean time and goods investments follow similar patterns over

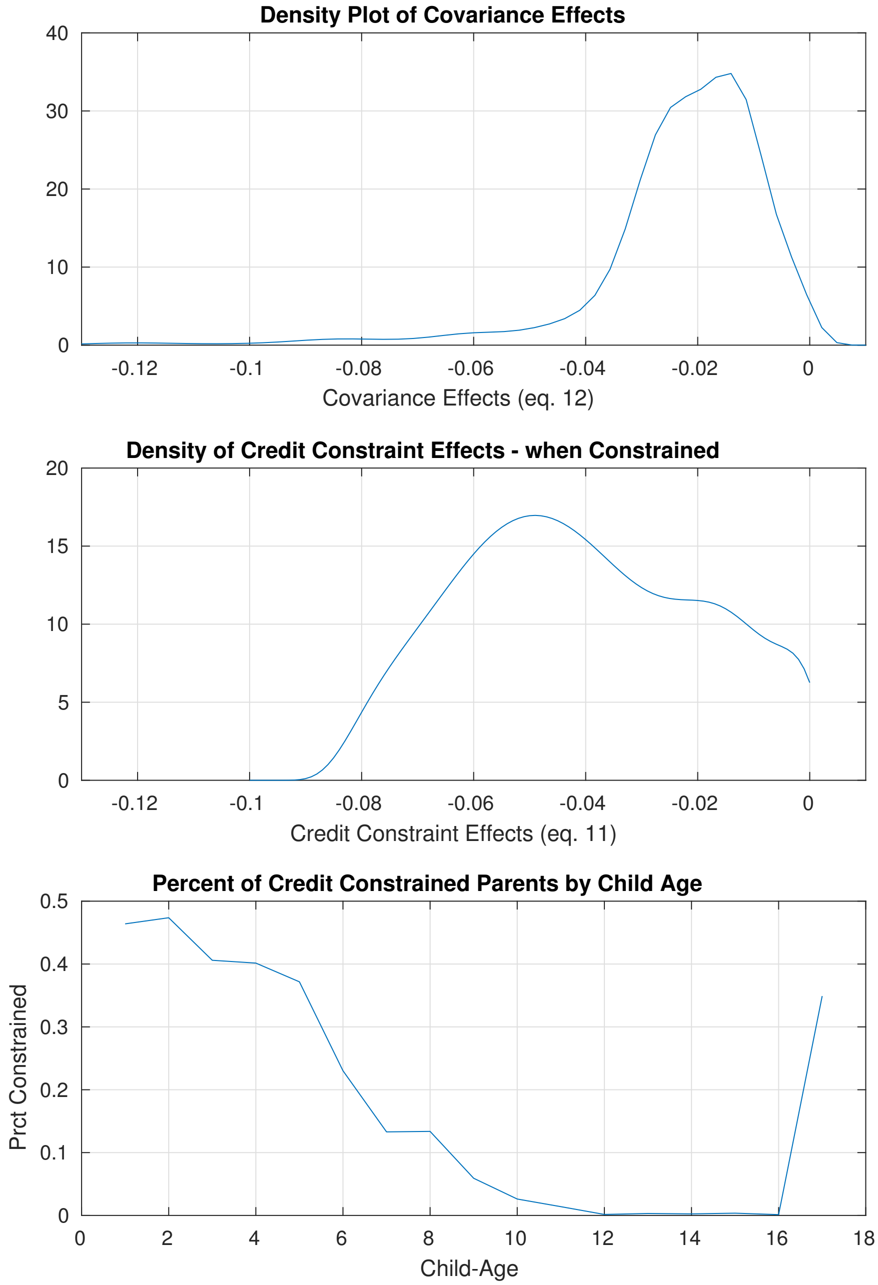

Figure 3: Top panel: This distribution of covariance effects as specified in equation 12. Mid Panel: The conditional distribution of credit constraint effects as specified in 11, where the conditioning is on the borrowing constraint binding. If the constraint is not binding the value of 11 is zero. Bottom Panel: The probability of the parental borrowing constraint binding by the age of the child.

the age of children, and as consequence the overall composite investment also follows that pattern. Initially, investments are about three percent higher in the FI version. This causes average ability in the FI version to increase relative to the IM version, at least after the first period when the distributions of initial skills are identical. The difference between FI and IM overall investment age-profiles peaks at age 8 at about 6%, and then begins to fall, becoming negative at age 13. This means that although FI parents invest more in their children when they are young, they invest less when they are older. As a result, the difference between average ability profiles exhibits a form of 'catch-up effect,' in the sense that the IM version of average ability grows more quickly in the teenage years than the FI model. At age 17 ability is just under 4% higher in the FI version, which reflects the shifting of investments from later ages to earlier when better insurance possibilities are added to the model. Given the calibrated wage elasticity λ, expected lifetime earnings increase by about 2.5% after this increase in ability.

The wedge in equation 9 indicates that investment timing differences between the IM and FI models can be decomposed into effects arising from credit constrained households and effects arising from uninsured wage risk. The former effect is captured by 11, while the latter is captured by 12. The distributions of these terms in the benchmark IM model are presented in Figure 3. The distribution of covariance effects is illustrated in the top panel, which indicates that the vast majority of such effects are between -0.04 and 0, and only a small proportion of covariance effects are positive. Credit constraint effects tend to be about twice as large as covariance effects when the constraint is binding, as seen in the second panel of Figure 3, which shows that the majority of credit constraint effects are between -0.08 and 0 in cases where households are constrained. Although credit constraint effects tend to be larger than covariance effects, they exist for only about 18% of parents (those whose credit constraint binds), thus the average credit constraint effect is about half as large (in absolute value) as the average covariance effect. For certain age groups, particularly very young children or children at the last stage of development (when intergenerational transfers are also decided), the average magnitude of credit constraint effects roughly equals that of covariance effects. The distribution of credit constraint effects, conditional on the constraint binding, is very similar for different age groups.

7 Policy Experiments

7.1 Setup

This section quantifies the value of policies designed to alleviate distortions in parental investment. As section 3.3 illustrated, there are several distortions and frictions in the model that motivate policy interventions. The approach here is to begin with less specialized interventions, and then add more complexity sequentially. The first policy reform changes the existing tax function parameters to achieve an optimal degree of redistribution (given the linear tax structure). Next, child development oriented policies are introduced, allowing tax function parameters to differ between parents and non-parents, as well as adding a subsidy for goods investments. This step indicates the welfare gains from addressing general intervention motives, for example the intergenerational transfer constraint. Last, age-dependent policies are introduced with the intention of correcting investment timing distortions.

The policy experiments consider the transition path of the economy, and years along the transition are indexed by j. Policy parameters are adjusted to their new permanent levels at time $j = 0$, and the path to the new steady-state is simulated. The policy maker faces a dynamic budget constraint in which outlays and net tax revenue changes are discounted at the financial rate r, and must sum to zero in the long-run. This allows the policy maker to borrow up front to finance reforms, but ensures that the present discounted value of incremental future tax revenue covers debt servicing costs.

To formalize the government budget constraint, let ϕ_t be the vector of household state-variables at age t, and let $\mu(\phi_t|t, j)$ be the probability measure of households over the state-space at age t in period j. Net tax revenue at time j, denoted G_j, is computed as follows:

$$\begin{aligned} G_j = &\sum_{t=18}^{64} \int (\tau_1(t) y(\phi_t, t) - \tau_0(t)) d\mu(\phi_t | t, j) \\ &+ \sum_{t=65}^{85} \int (\tau_1(t)\tau_{pen} - \tau_0(t) - \tau_{pen}) d\mu(\phi_t | t, j) \\ &- \sum_{t=30}^{46} \int (S(t)\chi(\phi_t, t)) d\mu(\phi_t | t, j), \end{aligned} \tag{28}$$

where $\chi(\phi_t, t)$ is the optimal decision rule for x at age t, and $y(\phi_t, t)$ is the optimal decision rule for earnings. The first line is net tax revenue from workers, the second line is net revenue from retirees (including pension costs), and the third line is spending on investment subsidies. Tax parameters and marginal subsidy rates vary with age in order to allow for age-dependent policies. The government budget constraint in the policy experiments can then be written as

$$\sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j (G_j - \overline{G}) \geq 0, \tag{29}$$

where $\overline{G}$ is net tax revenue, defined in 28, evaluated at the initial steady-state. The idea is that $\overline{G}$ equates to aggregate non-valued government spending, which is held constant in the experiments.

The welfare criterion values all living adults at time $j = 0$. This implicitly values future generations as well, because the welfare of altruistic parents depends on the welfare of their children. Formally, the welfare criterion is

$$\begin{aligned} W(\tau_0(t), \tau_1(t), S(t)) = &\sum_{t=18}^{29} \int V_t^W(\phi_t) d\mu(\phi_t | t, j) + \sum_{t=30}^{46} \int V_t(\phi_t) d\mu(\phi_t | t, j) + \\ &\sum_{t=47}^{64} \int V_t^W(\phi_t) d\mu(\phi_t | t, j) + \sum_{t=65}^{85} \int V_t^R(\phi_t) d\mu(\phi_t | t, j). \end{aligned} \tag{30}$$

V_t^W and V_t^R are the value functions of age t workers and retirees, respectively, defined in online Appendix B. All agents who are potentially affected by a policy reform are accounted for, either

	Benchmark	(I)	(II)	(III)
Policy Parameters				
τ_0	0.88	1.47	1.42	1.59
τ_1	0.27	0.46	0.47	0.51
τ_0^c			0.50	0.69
τ_1^c			-0.015	-0.017
S				0.47
Policy Effects ($\%\Delta$ between new steady-state and benchmark)				
Average Labor Income	–	-14.30%	-12.52%	-3.77%
Average Ability	–	-0.06%	12.19%	37.70%
Average Labor Supply	–	-17.05%	-17.52%	-18.61%
$Var(\ln(\theta))$	–	-6.45%	-12.76%	-14.48%
$Var(\ln(c))$	–	-27.98%	-27.54%	-21.58%
Welfare Criterion (includes transition path)				
Consumption Equivalent Variation	–	2.37%	2.87%	5.19%

Table 5: Age-Independent Policy Experiment Results. (I) is an optimal reform of the existing tax function parameters, (II) allows tax function parameters to differ between parents and other households, where e.g. the marginal tax rate of a parent (ages 30 through 46) is $\tau_1 + \tau_1^c$. (III) adds a subsidy on goods investments at marginal rate S.

directly or indirectly through altruism.

7.2 Results

(I) Tax Function Reform: To get a sense of the welfare gains from redistribution per se, the first policy experiment allows τ_0 and τ_1 to be reformed. By increasing the marginal tax rate τ_1, a planner is able to finance an increase in the lump-sum transfer τ_0, and thereby equalize marginal utilities from consumption to some extent. However, labor supply and human capital investment respond to the disincentive implied by higher marginal tax rates, which implies an equity-efficiency tradeoff for the planner. The results of this policy experiment are reported in column (I) of Table 5. The planner's optimal choice is to increase the marginal tax rate by about 19 percentage points, which finances an increase in lump-sum transfers by about 67%, compared to the benchmark economy. Such a substantial increase in the marginal tax rate leads to a large 17% reduction in aggregate labor supply, and a similar 14% reduction in aggregate labor income in the new-steady state. Although higher marginal

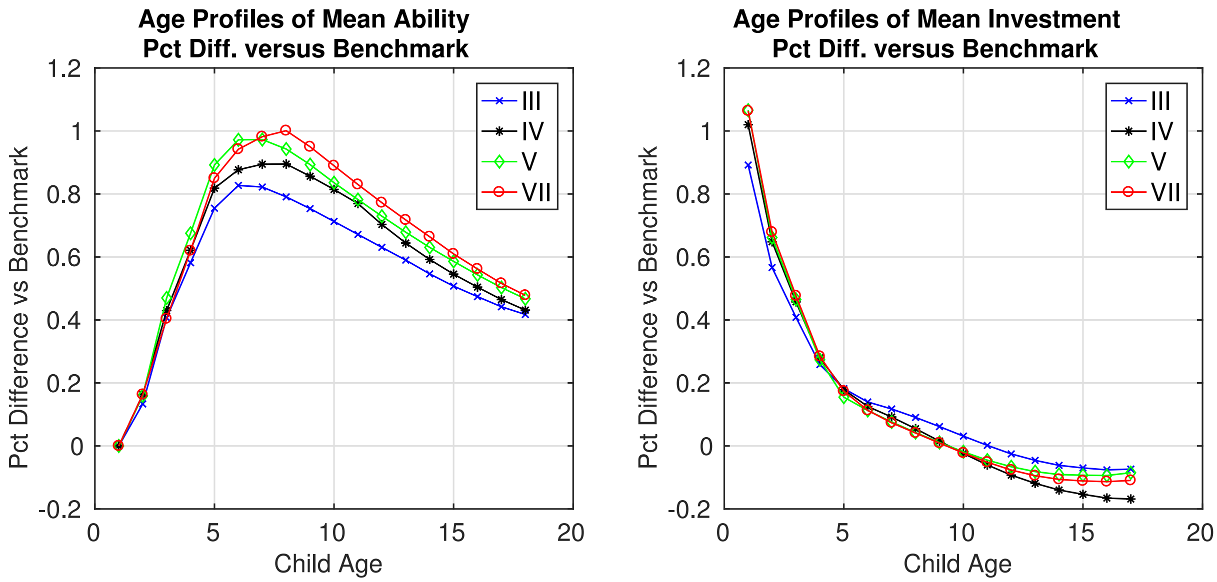

Figure 4: Age-profiles of ability and investment for the first treated cohort of children, relative to the benchmark economy, for experiments III, IV, V, and VII. Given these are for the first treated cohort, the distribution of parental state variables when the child is born is identical between the benchmark and experiment, thus the direct effect of policy is captured.

tax rates are also a disincentive to invest in human capital, there is only a very small decrease in average ability in the new steady-state. This is because the larger lump sum transfers lead to increased investment in certain parts of the distribution. The reduction in $Var(\ln(\theta))$, i.e. proportional variation in ability, by 6.45% captures the associated decrease in skill inequality. Consumption inequality is substantially lower in the new steady-state, with proportional variation in consumption decreasing by about 28%. The welfare gain from implementing this policy is equivalent to increasing every household's consumption by 2.37% in every year along the transition path and in the new steady-state.[21] Much of this gain comes from more stable and equally distributed consumption, despite overall output in the economy being smaller.

(II) Tax Function Reform with Parent Specific Parameters: Column (II) of table 5 allows the parameters of the benchmark tax function to vary between parents and non-parents. The marginal tax

[21]The consumption equivalent calculation holds non-leisure time constant in the benchmark economy, and thus can be thought of as a leisure-compensated consumption equivalent.

rate at ages 30-46 becomes $\tau_1 + \tau_1^c$, and the transfer becomes $\tau_0 + \tau_0^c$. Compared to the first experiment, marginal tax rates for parents are slightly lower, and marginal tax rates for non-parents are slightly higher, but these changes are small compared to how transfers optimally change. Transfers increase by just over 30% for parents, and fall by about 3.5% for non-parents, compared to experiment (I). An important product of larger cash transfers to parents is an increase in skills, with the average skill level now being about 12.2% higher than in the benchmark economy. This increase in skills leads to an increase in aggregate labor income relative to (I). At the same time, skill inequality falls, with $Var(\ln(\theta))$ being almost 13% smaller than in the benchmark. This is nearly double the reduction in skill inequality as experiment (I) accomplished. Average labor income is almost 2% higher than in experiment (I), and welfare is higher by the equivalent of a permanent 0.5% increase in consumption.

(III) Goods Investment Subsidization: Next, goods investments are allowed to be subsidized at marginal rate S. In general, such a subsidy can be desirable in order to reduce distortion in the time-goods investment ratio, induced by taxation of parental wages. It can also be desirable because it offsets the investment disincentive implied by taxation of a child's own future income, as described in Bovenberg and Jacobs (2005). Finally, it can help alleviate parental constraints, for example the intergenerational transfer constraint. In this case, the planner sets the marginal subsidy rate at 0.47. This is just slightly lower than the marginal tax rates, which are set to 0.51 for non-parents and just over 0.49 for parents (ages 30-46), implying that most of the investment distortions due to labor income taxation are eliminated. Introducing such a large subsidy leads to very substantial increases in average ability and, thus, labor income. Average ability is about 25% higher than in experiment (II) without the subsidy, and average labor income is almost 9% higher. The combination of higher incomes and higher marginal tax rates implies an increase in government revenue, and a part of this is spent on increasing transfer amounts τ_0 and τ_0^c. Skills are more equally distributed, but consumption is less equal. The latter is due to an increase in proportional dispersion of intergenerational wealth transfers, at least in part. Overall, the welfare gain from adding the subsidy to the policy reform leads

	(III)	(IV)	(V)	(VI)	(VII)
Policy Parameters					
τ_0	1.59	1.68	1.67	1.63	1.67
τ_1	0.51	0.51	0.51	0.51	0.51
τ_0^c	0.69	0.74	1.05	0.70	1.14
$\tau_0^{c,t'}$			-0.035		-0.034
τ_1^c	-0.017	-0.017	-0.016	-0.010	-0.010
$\tau_1^{c,t'}$				-5.4e-4	-4.8e-4
S	0.47	0.48	0.47	0.47	0.48
$S^{t'}$		-0.0020			-0.0013
Policy Effects ($\%\Delta$ between new steady-state and benchmark)					
Average Worker Income	-3.77%	-1.42%	-0.93%	-3.80%	0.01%
Average Ability	37.70%	42.71%	44.12%	37.83%	45.37%
Average Labor Supply	-18.61%	-18.14%	-18.14%	-18.69%	-17.80%
$Var(\ln(\theta))$	-14.48%	-16.52%	-19.71%	-15.38%	-18.56%
$Var(\ln(c))$	-21.58%	-20.39%	-19.87%	-22.13%	-18.97%
Welfare Criterion (includes transition path)					
Consumption Equivalent Variation	5.19%	5.49%	5.68%	5.21%	5.79%

Table 6: Age-Dependent Policy Experiment Results. (III) repeats from the previous table. (IV) allows the marginal goods investment subsidy rate to depend on child age according to $S + S^{t'} \times t'$, (V) allows the lump sum transfer to depend on child age according to $\tau_0 + \tau_0^c + \tau_0^{c,t'} \times t'$, (VI) allows the marginal tax rate to vary with child age according to $\tau_1 + \tau_1^c + \tau_1^{c,t'} \times t'$, and (VII) allows all three parameters to vary with child age simultaneously.

to a large increase in welfare, this policy being worth 2.3% more in consumption equivalent terms than the policy in experiment (II). To provide more context, Figure 4 plots age profiles of ability and investment for the first affected cohort relative to the benchmark economy, for several policy experiments. Under policy (III) investments are initially much higher, but then are lower for older children. This results in much higher abilities among younger children after a period of build up, but these differences decline to levels close to the means in the new steady state. For policy (III) these changes in the dynamics are less severe than under age-dependent policies (described next), but are quite pronounced nonetheless.

(IV) Age-Dependent Goods Investment Subsidization: Table 6 presents a series of experiments that include age-dependent policies intended to reduce distortions in parental investment timing, over and above what the age-independent policies accomplished. The first version (IV) allows the marginal subsidy rate to vary with child age, such that the marginal subsidy rate for a family with a t' year-old child is $S + S^{t'} \times t'$. The planner chooses the parameters S and $S^{t'}$ such that marginal subsidy rate decreases from 0.48 for the youngest children, to 0.44 for 17 year-olds. Compared to experiment (III) with age-independent subsidies, average ability is about 5% higher with age-varying subsidies, and this leads to a commensurate increase in labor income. In large part, the additional government revenue from the larger tax base is spent on increasing transfers, both τ_0 and τ_0^c. Adding the age-dependent subsidy increases welfare by the equivalent of 0.3% of consumption above what was accomplished with age-independent subsidies.

(V) Age-Dependent Transfers to Parents: Next, the parent specific transfer is allowed to vary with child age, such that each parent receives transfer $\tau_0 + \tau_0^c + \tau_0^{c,t'} \times t'$ (the marginal subsidy rate is constant). There is a substantial increase in transfers for parents of young children under this policy, but parents of older children receive less than under age-independent policies in experiment (III). Providing funds up front to parents who can save gives them more flexibility in the timing of using these funds; however, if parents are credit constrained at any point then investments after that time will suffer from lower future transfers. This policy outperforms the age-dependent subsidy policy, generating a welfare increase equivalent to about 0.5% of consumption over and above what age-independent policy accomplishes. These additional gains largely arise from additional increases in average ability and labor income.

(VI) Age-Dependent Marginal Tax Rates for Parents: This experiment allows the planner to set marginal tax rates between ages 30-46 according to $\tau_1 + \tau_1^c + \tau_1^{c,t'} \times t'$. The planner chooses the tax schedule such that the marginal tax rate falls from 0.50 to 0.49 between birth of the child and when they are age 17. This change has relatively little effect compared to the age-independent policies in (III), though there are small increases in ability and welfare. Examining equation (15), it is apparent that the marginal tax rate does not have a multiplicative term like the marginal subsidy rate does, which explains why the marginal tax rate is a less effective tool.

(VII) Joint Age-Dependent Policies: Finally, marginal subsidy rates, transfers and marginal tax rates are all allowed to vary with the age of the child. The intercept of the transfer function increases by nearly 10% compared to (V), while the age slope is only slightly smaller. Marginal subsidy rates decrease with child age at about 2/3 of the rate that they do in (IV), and the dependence of marginal tax rates on child age continues to be weak as in (VI). Overall, the planner achieves a welfare improvement equivalent to 0.6% of consumption above what is accomplished with child-age independent policies in (III), and about 0.1% of consumption above what is accomplished with age-dependent transfers alone. On net, the policy in experiment (VII) results in virtually the same aggregate labor income as in the benchmark economy, but uses about 18% less labor supply to do so, and exhibits consumption inequality about 19% lower (as measured by proportional variation). Of course, to generate the same income with less labor supply requires greater labor productivity, which is generated under the policy by a 45% increase in average ability. Figure 4 plots ability and investment age profiles relative to the benchmark for (VII), showing that age-dependent policies lead to more early investment and less late investment, which in turn generates higher ability levels, much like the IM versus FI comparison in Figure 2.

7.3 Discussion of Policy Experiments

It is instructive to make comparisons between policy experiments in terms how changes in the policy parameters interact with the incomplete markets frictions. Changes in the components of the investment wedge in equation (16) are a useful way to measure these effects. The main results reported above often refer to the new steady-state at the end of the transition path; however, these depend on both direct policy effects and long-run changes in the distributions of wealth and ability. One way to isolate direct policy effects is to examine the first cohort of children born under the new policy, whose parents will have the same wealth and ability levels as children in the initial steady-state. Table 7

	Benchmark	(I)	(III)	(IV)	(V)	(VII)
Timing Wedge Components						
Credit Constraint Effect (Eq. 17)	−0.012	−0.012	−0.010	−0.007	−0.004	−0.003
Covariance Effect (Eq. 18)	−0.018	−0.014	−0.014	−0.013	−0.012	−0.013
Policy Effect (Eq. 19)	0	≈ 0	≈ 0	0.004	≈ 0	0.003
Total Timing Wedge	−0.030	−0.026	−0.024	−0.016	−0.016	−0.013

Table 7: Investment timing wedge component averages for the first affected cohort in selected experiments.

reports averages of the components of the investment wedge in (16), as well as the total wedge.

To understand how age invariant policies affect the investment wedge, compare experiments (I) and (III) to the benchmark in Table 7. In both cases the covariance effect is considerably smaller as more generous transfers serve to reduce consumption and investment volatility. The credit constraint effect becomes smaller in experiment (III) as parent-specific transfers reduce the likelihood of parental credit constraints binding. Although there is a tiny change in the policy term 19, it is effectively still zero because the policies are age-invariant.

Moving next to age-varying policies, Table 7 reports results from experiments (IV), (V) and (VII).[22] When age-varying subsidies are added in experiment (IV), the first big difference to note relative to age-invariant policies is the substantial policy term, defined in equation (19). This is a significant force in reducing the investment wedge to nearly half the size it was in the benchmark economy. The other big difference compared to (I) and (III) is a big reduction in the credit constraint effect, which results from reallocating resources to younger parents whose credit constraints are most likely to bind. In experiment (V) the reduction in credit constraint effects is even stronger as direct transfers to younger parents is more effective at alleviating their borrowing constraints than subsidizing their expenditure investments is. However, in (V) the policy term is effectively zero again because transfers are not distortionary. Experiment (VII) puts these together (the changes in marginal tax rates are very minor) and thereby generates a substantial reduction in the credit constraint effect and a sizeable positive policy term, which together generate the smallest investment timing wedge among the

[22] Experiment (VI) is effectively identical to (III).

experiments.

8 Conclusion

The goal of this paper has been to understand how market incompleteness leads to the conclusion that subsidizing early investments in children is more beneficial than later ones. Dynamic complementarity in the skill production function alone is not enough to explain this because parents are fully rational optimizing agents in such a framework. Analytical results showed that the combination of dynamic complementarity and uninsurable risk can cause parents delay investments in their child relative to a model with full insurance, as long as there is some possibility of the non-negative intergenerational transfer constraint binding. Parental credit constraints that bind can exacerbate these delays, though the results hold even if parental credit constraints never bind. The analytical results also showed that marginal subsidy or tax rates that decrease as a child gets older, as well as transfers to parents, can alleviate these timing distortions.

Several policy experiments were simulated, where the welfare criterion included the transition path of the economy and a long-run government budget constraint held. The policy tools available to the planner became sequentially richer, starting from a standard reform of the tax and transfer program and finishing with a policy that included child-age dependent transfers, marginal income tax rates and marginal subsidy rates. Optimal redistribution achieved through the first experiment yielded welfare gains in line with the literature, for example Floden and Lindé (2001). The most complex policy experiment with a full suite of child-age dependent policies yielded welfare gains that were almost 2.5 times as large as the first experiment. These gains arose from an improved ability distribution, characterized by a larger mean and smaller proportional variation. About one-quarter of the incremental welfare gain was achieved by allowing transfers and marginal subsidy rates to be larger for younger children, as opposed to constant for all children. This was because such policy formulations alleviated some of the distortion in parental investment timing that results from incomplete markets.

In terms of concrete policy conclusions, the analyses in this paper suggest that it would be best to combine cash transfers with subsidies on purchased investments, and that these policies should be more generous for younger children. Adjusting marginal tax rates for parents would be less beneficial, even if done in an age-varying way. Implementation of the transfer policy could be through a refundable child tax credit that declines as children get older. This could be combined with direct subsidies for purchased investments, such as preschool tuition or music lessons, for younger children. A simpler way to implement the subsidy component might be an age-varying refundable tax credit for certain types of child related purchases.

References

Abbott, B., Gallipoli, G., Meghir, C., and Violante, G. L. (2019). 'Education Policy and Intergenerational Transfers in Equilibrium.' *Journal of Political Economy*, vol. 127(6):pages 2569–2624.

Agostinelli, F. and Wiswall, M. (2016a). 'Estimating the Technology of Children's Skill Formation.' Tech. rep., National Bureau of Economic Research.

Agostinelli, F. and Wiswall, M. (2016b). 'Identification of dynamic latent factor models: The implications of re-normalization in a model of child development.' Tech. rep., National Bureau of Economic Research.

Attanasio, O., Meghir, C., and Nix, E. (2015). 'Human capital development and parental investment in india.' Tech. rep., National Bureau of Economic Research.

Attanasio, O., Meghir, C., Nix, E., and Salvati, F. (2017). 'Human capital growth and poverty: Evidence from Ethiopia and Peru.' *Review of economic dynamics*, vol. 25:pages 234–259.

Becker, G. S. and Tomes, N. (1979). 'An equilibrium theory of the distribution of income and intergenerational mobility.' *Journal of political Economy*, vol. 87(6):pages 1153–1189.

Becker, G. S. and Tomes, N. (1986). 'Human capital and the rise and fall of families.' *Journal of labor economics*, vol. 4(3, Part 2):pages S1–S39.

Bernal, R. (2008). 'The effect of maternal employment and child care on children's cognitive development.' *International Economic Review*, vol. 49(4):pages 1173–1209.

Blandin, A., Herrington, C. *et al.* (2018). 'Family Structure, Human Capital Investment, and Aggregate College Attainment.' *Unpublished Manuscript, Virginia Commonwealth University*.

Boneva, T. and Rauh, C. (2018). 'Parental Beliefs about Returns to Educational Investments—The Later the Better?' *Journal of the European Economic Association*, vol. 16(6):pages 1669–1711.

Bono, E. D., Francesconi, M., Kelly, Y., and Sacker, A. (2016). 'Early maternal time investment and early child outcomes.' *The Economic Journal*, vol. 126(596).

Bovenberg, A. L. and Jacobs, B. (2005). 'Redistribution and education subsidies are Siamese twins.' *Journal of Public Economics*, vol. 89(11-12).

Carneiro, P. and Ginja, R. (2016). 'Partial insurance and investments in children.' *The Economic Journal*, vol. 126(596).

Carneiro, P., Lopez Garcia, I., Salvanes, K. G., and Tominey, E. (2015). 'Intergenerational Mobility and the Timing of Parental Income.' Tech. rep., IZA Discussion Papers.

Caucutt, E. M. and Lochner, L. (2020). 'Early and late human capital investments, borrowing constraints, and the family.' *Journal of Political Economy*, vol. 128(3).

Cunha, F. (2013). 'Investments in Children When Markets Are Incomplete.' *unpublished*.

Cunha, F., Elo, I., and Culhane, J. (2013). 'Eliciting Maternal Expectations about the Technology of Cognitive Skill Formation.' Tech. rep., National Bureau of Economic Research.

Cunha, F. and Heckman, J. (2007). 'The Technology of Skill Formation.' *The American Economic Review*, vol. 97(2):pages 31–47.

Cunha, F., Heckman, J. J., and Schennach, S. M. (2010). 'Estimating the technology of cognitive and noncognitive skill formation.' *Econometrica*, vol. 78(3):pages 883–931.

Daruich, D. (2018). 'The Macroeconomic Consequences of Early Childhood Development Policies.' *FRB St. Louis Working Paper 2018-29*.

Del Boca, D., Flinn, C., and Wiswall, M. (2013). 'Household choices and child development.' *Review of Economic Studies*, vol. 81(1):pages 137–185.

Domeij, D. and Heathcote, J. (2004). 'On the distributional effects of reducing capital taxes.' *International economic review*, vol. 45(2):pages 523–554.

Erosa, A., Fuster, L., and Restuccia, D. (2010). 'A general equilibrium analysis of parental leave policies.' *Review of Economic Dynamics*, vol. 13(4):pages 742 – 758. ISSN 1094-2025.

Fiorini, M. and Keane, M. P. (2014). 'How the allocation of children's time affects cognitive and noncognitive development.' *Journal of Labor Economics*, vol. 32(4):pages 787–836.

Floden, M. and Lindé, J. (2001). 'Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?' *Review of Economic dynamics*, vol. 4(2):pages 406–437.

Heathcote, J., Perri, F., and Violante, G. L. (2010). 'Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006.' *Review of Economic dynamics*, vol. 13(1):pages 15–51.

Heathcote, J., Storesletten, K., and Violante, G. L. (2020). 'Optimal progressivity with age-dependent taxation.' *Journal of Public Economics*, page 104074.

Heckman, J. J. and Mosso, S. (2014). 'The Economics of Human Development and Social Mobility.' *Annual Review of Economics*, vol. 6.

Hsin, A. and Felfe, C. (2014). 'When does time matter? Maternal employment, children's time with parents, and child development.' *Demography*, vol. 51(5):pages 1867–1894.

Lee, S. Y. and Seshadri, A. (2019). 'On the intergenerational transmission of economic status.' *Journal of Political Economy*, vol. 127(2):pages 855–921.

Mogstad, M. (2017). 'The human capital approach to intergenerational mobility.' *Journal of Political Economy*, vol. 125(6):pages 1862–1868.

Moschini, E. (2019). 'Child Care Subsidies with One- and Two-Parent Families.' *Working Paper*.

Mullins, J. (2016). 'Improving child outcomes through welfare reform.' *Working Paper*.

Pavan, R. (2015). 'On the production of skills and the birth order effect.' *Journal of Human Resources*.

Restuccia, D. and Urrutia, C. (2004). 'Intergenerational persistence of earnings: The role of early and college education.' *American Economic Review*, vol. 94(5):pages 1354–1378.

Schoonbroodt, A. (2018). 'Parental child care during and outside of typical work hours.' *Review of Economics of the Household*, vol. 16(2):pages 453–476.

Todd, P. E. and Wolpin, K. I. (2003). 'On the specification and estimation of the production function for cognitive achievement.' *The Economic Journal*, vol. 113(485):pages F3–F33.

Tonetti, C. (2011). 'Notes on Estimating Earnings Processes.' Tech. rep., Citeseer.

A Deriving Dynamics

A.1 Monetary Investments Only - Incomplete Markets

$$
\begin{aligned}
V(a,w,\theta) = \max_{a',x} & \Big\{ u((1+r)a + w - a' - x) + \beta \mathbb{E}\left[V(a',w',\theta')\right] \\
& s.t.\;\; \theta' = \gamma_0 \left(\gamma \theta^{\psi} + (1-\gamma) x^{\psi}\right)^{\frac{1}{\psi}}, \;\; a' \geq \underline{a} \Big\} \\
& foc\, x:\;\; u_c(c) = \beta \mathbb{E}\left[V_\theta(a',w',\theta')\right] \gamma_0 \left(\gamma \theta^{\psi} + (1-\gamma) x^{\psi}\right)^{\frac{1-\psi}{\psi}} (1-\gamma) x^{\psi-1} \\
& B.S.\,\theta:\;\; V_\theta(a,w,\theta) = \beta \mathbb{E}\left[V_\theta(a',w',\theta')\right] \gamma_0 \left(\gamma \theta^{\psi} + (1-\gamma) x^{\psi}\right)^{\frac{1-\psi}{\psi}} \gamma \theta^{\psi-1}
\end{aligned}
\tag{31}
$$

Re-write the last two equations as:

$$
\begin{aligned}
foc\, x:\;\; u_c(c) &= \beta \mathbb{E}\left[V_\theta(a',w',\theta') \theta'^{(1-\psi)}\right] \gamma_0^{\psi} (1-\gamma) x^{\psi-1} \\
B.S.\,\theta:\;\; V_\theta(a,w,\theta) &= \beta \mathbb{E}\left[V_\theta(a',w',\theta') \theta'^{(1-\psi)}\right] \gamma_0^{\psi} \gamma \theta^{\psi-1}
\end{aligned}
$$

Combine the last two equations:

$$
V_\theta(a,w,\theta) = u_c(c) \frac{\gamma}{1-\gamma} \frac{x^{1-\psi}}{\theta^{1-\psi}} \tag{32}
$$

Plug back into f.o.c. to derive:

$$
1 = \mathbb{E}\left[\beta \frac{u_c(c')}{u_c(c)} \gamma_0^{\psi} \gamma \left(\frac{x'}{x}\right)^{1-\psi}\right], \tag{33}
$$

Note that the above equations hold regardless of whether the borrowing constraint binds or not.

A.2 Monetary Investments Only - With Arrow-Debreu Securities

For simplicity, let s be the state of the world, and let $\pi(s'|s)$ be the conditional probabilities of future states. Implicitly, the state of the world determines the wage through a function $w = \omega(s)$. Let $b(s')$ be the quantity of assets paying off in state s' that the household buys for price $q(s')$.

$$V(b(s), w, \theta, s) = \max_{\{b(s')\},x} \Big\{ u(b(s) + w - \sum_{s'} q(s')b(s') - x) + \beta \mathbb{E}_{s'|s} \left[V(b(s'), w', \theta', s') \right] \quad (34)$$
$$s.t. \;\; \theta' = \gamma_0 \left(\gamma\theta^{\psi} + (1-\gamma)x^{\psi} \right)^{\frac{1}{\psi}} \Big\}$$

Firstly, note that all of the derivations from the baseline incomplete markets case 32 still hold as they did not involve the savings decisions, so we have

$$x(s)^{1-\psi} = \gamma_0^{\psi} \gamma \beta \sum_{s'} \pi(s'|s) \frac{u_c(c'(s'))}{u_c(c(s))} x'(s')^{1-\psi}. \quad (35)$$

With only idiosyncratic risk in the model, actuarially fair pricing arises and $q(s') = R\pi(s'|s)$. Therefore, substituting each Arrow-Debreau first-order condition

$$q(s') = \beta \pi(s'|s) \frac{u_c(c'(s'))}{u_c(c(s))}$$

into equation 37 yields the following:

$$x(s)^{1-\psi} = \gamma_0^{\psi} \gamma R \sum_{s'} \pi(s'|s) x'(s')^{1-\psi}, \quad (36)$$

or equivalently

$$x_{it}^{1-\psi} = \gamma_0^{\psi} \gamma R \mathbb{E} \left[x_{it+1}^{1-\psi} \right]. \quad (37)$$

A.3 General Case - Incomplete Markets

In this general case the dynamic program can be written

$$V(a,w,\theta,\alpha) = \max_{a',x,\ell,h} \Big\{ u((1+r)a + w\ell(1-\tau(w\ell)) - a' - x, \ell + h) + \beta\mathbb{E}\left[V(a',w',\theta',\alpha')\right] \tag{38}$$

$$s.t.\;\; \theta' = \gamma_0 \left(\gamma\theta^{\psi} + (1-\gamma)I^{\psi}\right)^{\frac{1}{\psi}}, \;\; I = \left(\alpha x^{\delta} + (1-\alpha)h^{\delta}\right)^{\frac{1}{\delta}}, \;\; a' \geq \underline{a}, \; ... \Big\}$$

$$foc\, x: \;\; u_c(c,\ell+h) = \beta\mathbb{E}\left[V_\theta(a',w',\theta',\alpha')\theta'^{1-\psi}\right]\gamma_0^{\psi}(1-\gamma)I^{\psi-1}I^{1-\delta}\alpha x^{\delta-1}$$

$$foc\, h: \;\; -u_h(c,\ell+h) = \beta\mathbb{E}\left[V_\theta(a',w',\theta',\alpha')\theta'^{1-\psi}\right]\gamma_0^{\psi}(1-\gamma)I^{\psi-1}I^{1-\delta}(1-\alpha)h^{\delta-1}$$

$$B.S.\,\theta: \;\; V_\theta(a,w,\theta,\alpha) = \beta\mathbb{E}\left[V_\theta(a',w',\theta',\alpha')\theta'^{1-\psi}\right]\gamma_0^{\psi}\gamma\theta^{\psi-1}$$

Combine the $foc\, x$ and $B.S.\,\theta$ equations:

$$V_\theta(a,w,\theta,\alpha) = u_c(c,\ell+h)\frac{\gamma}{1-\gamma}\frac{I^{1-\psi}I^{\delta-1}\alpha x^{1-\delta}}{\alpha\theta^{1-\psi}} \tag{39}$$

Plug back into f.o.c. to derive:

$$\frac{I^{-\psi}I^{\delta}x^{1-\delta}}{\alpha} = \gamma_0^{\psi}\gamma\mathbb{E}\left[\beta\frac{u_c(c',\ell'+h')}{u_c(c,\ell+h)}\frac{I'^{-\psi}I'^{\delta}x'^{1-\delta}}{\alpha'}\right], \tag{40}$$

which has the same flavor as the dynamics of the expenditure only case, but now a complicated ratio involving x, I, and α satisfies what x satisfied in the $\alpha = 1$ case.

To understand the dynamics of total investment (I), derive an expression for the unity cost of investment, which will allow x to be substituted out of equation 40. To do this, first write the total resource cost of a given level of investment (in units of consumption) as the sum of the costs of underlying time and money investments.

$$\text{cost}(I) = x + \tilde{w}h. \tag{41}$$

Here $\tilde{w} = -u_h/u_c$, which may or may not equal the actual net wage rate $w(1-\tau'(w\ell))$, depending on whether household labor supply is at an interior solution. Time investments are valued at their (marginal) opportunity cost in terms of consumption in this equation.

Given this cost accounting, the unit cost of investment $p^I(\tilde{w}, \alpha; \delta)$ is the solution to the following equation:

$$\text{cost}(I) = p^I(\tilde{w}, \alpha; \delta) \times I. \tag{42}$$

That is, $p^I(\tilde{w}, \alpha; \delta)$ is the marginal cost in terms of consumption of an additional unit of total investment. If one divides $foc\ x$ to $foc\ h$, the resulting intratemporal first-order condition can be used to substitute time investments out of $\text{cost}(I)$. Furthermore, if I is replaced by its CES specification, the effective price of investment can be solved for in terms of exogenous parameters and prices:

$$p^I(\tilde{w}, \alpha; \delta) = \frac{1}{\alpha}\left(\alpha + (1-\alpha)\left(\frac{\alpha}{1-\alpha}\tilde{w}\right)^{\frac{\delta}{\delta-1}}\right)^{\frac{\delta-1}{\delta}}. \tag{43}$$

The usefulness of this price is that one can show that at an optimal interior solution the following relationship holds:

$$\frac{I^{\delta}x^{1-\delta}}{\alpha} = Ip^I(\tilde{w}, \alpha; \delta).$$

Using this result, the general investment dynamics can be written as:

$$1 = \gamma_0^{\psi}\gamma\mathbb{E}\left[\beta\frac{u_c(c', \ell' + h')}{u_c(c, \ell + h)}\frac{p^I(\tilde{w}', \alpha'; \delta)}{p^I(\tilde{w}, \alpha; \delta)}\left(\frac{I'}{I}\right)^{1-\psi}\right]. \tag{44}$$

General Case - with Arrow-Debreu securities Applying the same approach as in A.2, one can derive:

$$I^{1-\psi} = \gamma_0^{\psi}\gamma R\mathbb{E}\left[\frac{p^I(\tilde{w}', \alpha'; \delta)}{p^I(\tilde{w}, \alpha; \delta)}I'^{1-\psi}\right]. \tag{45}$$

B Quantitative Model Value Functions

Working recursively, the decision problems are presented backwards from the terminal period to the first period of adult life.

Retirement: $t = 65, \ldots, 85$**.** At the terminal age $t = 85$ the continuation value is zero, i.e. $V_{86}^R(\cdot) = 0$, where V_t^R denotes the value function of a retired agent at age t. Otherwise, during the retirement periods agents solve the following problem:

$$\begin{aligned} V_t^R(a_t) &= \max_{\{a_{t+1}\}} \left\{ u(c_t) + \beta V_{t+1}^R(a_{t+1}) \right\} \\ & s.t. \\ c_t &= (1+r)a_t - a_{t+1} + \tau_{pen}(1-\tau_1) + \tau_0 \\ a_{t+1} &\geq 0. \end{aligned} \tag{46}$$

Retired households cannot borrow in this model.

Work After Children: $t = 48, \ldots, 64$**.** Let V_t^W be the age t value function of a working agent. Such agents solve a classic consumption-savings problem:

$$\begin{aligned} V_t^W(a_t, z_t, \theta) &= \max_{\{a_{t+1}, \ell_t\}} \left\{ u(c_t, \ell_t) + \beta \mathbb{E}\left[V_{t+1}^W(a_{t+1}, z_{t+1}, \theta) \right] \right\} \\ & s.t. \\ c_t &= (1+r)a_t - a_{t+1} + \exp(\overline{w} + \lambda \ln(\theta) + z_t + \xi_t)\ell_t(1-\tau_1) + \tau_0 \\ a_{t+1} &\geq 0 \\ z_{t+1} &\sim \mathit{\Gamma}_z(z_t). \end{aligned} \tag{47}$$

Note that θ is the ability of the agent themself, as opposed to that of their child. $\Gamma_z(z_t)$ is the conditional distribution of future wage shocks. In practice $\Gamma_z(z_t)$ is a row of a Markov Chain transition matrix. One slight abuse of notation occurs at age $t = 64$, where $\mathbb{E}\left[V_{65}^W(a_{65}, z_{65}, \theta)\right] = V_{65}^R(a_{65})$.

Child Matures: $t = 47$**.** Age 47 is a unique period for parents because this is the point that the utility payoff from investing in the human capital of their child is realized. The value function of a parent when they are age t is denoted V_t^p. This value function depends on the parent's wealth, earnings shock and own ability, as well as the ability of their child, θ_t^c, and their parenting-time productivity, ζ. When a parent is age $t = 47$ (and their child is age $t' = T = 17$) the parental decision problem is:

$$\begin{aligned}
V_t^P(a_t, z_t, \theta, \theta_t^c, \zeta, t' = T) \quad &= \quad \max_{\{a_{t+1}, \ell_t, x_t, h_t, a^c\}} \Big\{ u(c_t, \ell_t + h_{it}) + \beta \mathbb{E}\left[V_{t+1}^W(a_{t+1}, z_{t+1}, \theta)\right] + \\
&\qquad \beta \kappa \tilde{V}(a^c, \theta_{t+1}^c) \Big\} \\
&s.t. \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (48) \\
c_t \quad &= \quad (1+r)a_t - a_{t+1} + \exp(\overline{w} + \lambda \ln(\theta) + z_t + \xi_t)\ell_t(1 - \tau_1) \\
&\qquad - x_t - a^c + \tau_0 \\
\theta_{t+1}^c \quad &= \quad \gamma_0 \left(\gamma \theta_t^{c\psi} + (1-\gamma) I_t^{\psi}\right)^{\frac{1}{\psi}} \\
I_t \quad &= \quad \left(\alpha_t x_t^\delta + (1-\alpha_t) h_t^\delta\right)^{\frac{1}{\delta}} \\
a_{t+1} \quad &\geq \quad \underline{a}, \quad a^c \geq 0 \\
z_{t+1} \quad &\sim \quad \Gamma^z(z_t) \\
\alpha_t \quad &= \quad \frac{\overline{\alpha}_{t'}}{\overline{\alpha}_{t'} + (1 - \overline{\alpha}_{t'})\zeta}.
\end{aligned}$$

Parenting Ages: $t = 30 \ldots 46$**.** Given the payoff from investing in a child's human capital is built into V_{47}^P, parent are motivated to invest in their child from ages 30, when the child is born, to age 47, when the final opportunity to invest occurs. The decision problem of a parent during the periods of investment is:

$$\begin{aligned}
V_t^P(a_t, z_t, \theta, \theta_t^c, \zeta, t') &= \max_{\{a_{t+1}, \ell_t, x_t, h_t\}} \left\{ u(c_t, \ell_t + h_t) + \beta \mathbb{E}\left[V_{t+1}^P(a_{t+1}, z_{t+1}, \theta, \theta_{t+1}^c, \zeta, t'+1) \right] \right\} \\
& \quad s.t. \qquad (49) \\
c_t &= (1+r)a_t - a_{t+1} + \exp(\overline{w} + \lambda \ln(\theta) + z_t + \xi_t)\ell_t(1-\tau_1) + \tau_0 - x_t \\
\theta_{t+1}^c &= \gamma_0 \left(\gamma \theta_t^{c\psi} + (1-\gamma) I_t{}^{\psi} \right)^{\frac{1}{\psi}} \\
I_t &= \left(\alpha_t x_t^{\delta} + (1-\alpha_t) h_t^{\delta} \right)^{\frac{1}{\delta}} \\
a_{t+1} &\geq \underline{a}, \quad a^c \geq 0 \\
z_{t+1} &\sim \Gamma^z(z_t) \\
\alpha_t &= \frac{\overline{\alpha}_{t'}}{\overline{\alpha}_{t'} + (1-\overline{\alpha}_{t'})\zeta}.
\end{aligned}$$

This is very much equivalent to the problem described in equation 3, except that the less general structure of the quantitative model is introduced.

Working Prior to Children: $t = 18, \ldots, 29$**.** Prior to the arrival of their child, an agent in the model works and solves the same problem as in equation 47, with two exceptions: (i) At age $t = 29$, just before the child is born, the continuation value is

$$V_{30}^W(a_{30}, z_{30}, \theta) = \sum_{j=low,high} Pr(\theta_{j0}^c | \theta) \int V_{30}^P(a_{30}, z_{30}, \theta, \theta_{j0}^c, \zeta)\phi(\ln(\zeta); \sigma_\zeta^2) d\zeta.$$

Here $\phi(\ln(\zeta); \sigma_\zeta^2)$ is the normal PDF for (log) parenting-time productivity, which is revealed to the parent when their child is born. All children are born with initial ability that is either high or low. (ii) From ages 18-21 $\lambda \ln(\theta)$ is replaced by $(\lambda/2)\ln(\theta)$ in the budget constraints.

Initial Value Function $\tilde{V}$**:** The expected value function of a child, which enters the parental decision problem at age $t = 47$, is simply $\tilde{V}(a^c, \theta_{48}^c) = V_{18}^W(a_{18} = a^c, 0, \theta = \theta_{48}^c)$.

C A Stylized Analytical Example

Consider a simple two-period model of parental decisions, where parents make purely expenditure investments in their child. Assume that parents have $ln(c)$ preferences over their own consumption, and value the log of their child's lifetime income. Parents have an initial endowment of wealth a_1, and face uncertainty about their income y_2 in the second period. With probability π their period 2 income will be $y_2 = \tilde{y}$, and with probability $1 - \pi$ it will be $y_2 = 0$. The dynamic skill production function is simplified to a Cobb-Douglas form with $\theta_{t+1}^c = \theta_t^{c\gamma} {I_t}^{1-\gamma}$.

Full-Insurance: First consider the case in which a household has access to Arrow-Debreu securities. In this simple model there are only two state-contingent bonds as there are only two states of the world related to the realization of y_2. Let $b(y_2)$ be the quantity of state-contingent bonds the household previously purchased corresponding to a given y_2 realization. In this case the period 2 parental decision problem is:

$$\begin{aligned} V_2(b(y_2), y_2, \theta_2^c) &= \max_{\{c_2, I_2, a^c\}} \{\ln(c_2) + \kappa \ln(a^c + \varpi \theta_3^c)\} \\ & s.t. \\ c_2 &= b(y_2) + y_2 - a^c - I_2 \\ \theta_3^c &= \theta_2^{c\gamma} {I_2}^{1-\gamma} \\ a^c &\geq 0. \end{aligned} \tag{50}$$

In this problem, the child's lifetime income is the sum of their transfer a^c, and their lifetime earnings, $\varpi \theta_3^c$, which are proportional to their attained ability. In this example, as in the main model, investments will be efficient whenever the non-negative transfer constraint is not binding. As this example is intended to illustrate cases where uninsurable risk can distort investment timing, assume now that the non-negative transfer constraint is indeed binding, i.e. $a^c = 0$.

The reason the particular preferences described above were chosen is that when the non-negative transfer constraint binds the model has the same preferences as the Del Boca, Flinn, and Wiswall (2013) model, in that parents simply have $\ln$ preferences over their child's attained ability. Like Del Boca, Flinn, and Wiswall, simple closed form solutions are the goal of this stylization:

$$c_2^{FI} = \frac{1}{1+\kappa(1-\gamma)}\left(b(y_2)+y_2\right) \tag{51}$$

$$I_2^{FI} = \frac{\kappa(1-\gamma)}{1+\kappa(1-\gamma)}\left(b(y_2)+y_2\right). \tag{52}$$

Continuing with the Arrow-Debreu case, the parental problem in the first period will be

$$\begin{aligned} V_1(a_1,\theta_1^c) &= \max_{\{c_1,I_1,\{b(y_2)\}\}} \{\ln(c_1)+\beta\mathbb{E}_{y_2}\left[V_2(b(y_2),y_2,\theta_2^c)\right]\} \\ & s.t. \\ c_1 &= a_1 - I_1 - R\pi b(y_2=\tilde{y}) - R(1-\pi)b(y_2=0) \\ \theta_2^c &= \theta_1^{c\gamma} {I_1}^{1-\gamma}, \end{aligned} \tag{53}$$

where actuarially fair state-prices have been applied. As with the period-two problem, the solutions have simple forms:

$$c_1^{FI} = \frac{1}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}\left(a_0+\pi\tilde{y}\right) \tag{54}$$

$$I_1^{FI} = \frac{\gamma\beta\kappa(1-\gamma)}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}\left(a_0+\pi\tilde{y}\right) \tag{55}$$

$$b^{FI}(y_2=0) = (1+r)\frac{\beta(1+\kappa(1-\gamma))}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}\left(a_0+\pi\tilde{y}\right) \tag{56}$$

$$b^{FI}(y_2=\tilde{y}) = b^{FI}(y_2=0)-\tilde{y}. \tag{57}$$

Incomplete-Markets: With incomplete markets, the parent chooses a single savings position a_2 in the first period, as opposed to the full set of state-contingent bond positions. In the second period their solutions will be functions of a_2 rather than $b(y_2)$, but otherwise the second-period problem above remains in tact. Thus, with incomplete markets the parental decisions are:

$$c_2^{IM} = \frac{1}{1+\kappa(1-\gamma)}(a_2+y_2) \tag{58}$$

$$I_2^{IM} = \frac{\kappa(1-\gamma)}{1+\kappa(1-\gamma)}(a_2+y_2). \tag{59}$$

The initial period parent's problem has fewer choices, and is as follows:

$$\begin{aligned} V_1(a_1,\theta_1^c) &= \max_{\{c_1,I_1,a_2\}} \{\ln(c_1)+\beta\mathbb{E}_{y_2}[V_2(a_2,y_2,\theta_2^c)]\} \\ & s.t. \\ c_1 &= a_1 - I_1 - Ra_2 \\ \theta_2^c &= \theta_1^{c\gamma} {I_1}^{1-\gamma}. \end{aligned} \tag{60}$$

We now face a complication because the optimal choice of saving a_2 satisfies the intertemporal Euler equation

$$\frac{1}{c_1^{IM}} = \beta(1+r)\mathbb{E}_{y_2}\left[\frac{1}{c_2^{IM}(a_2,y_2)}\right]. \tag{61}$$

Del Boca, Flinn, and Wiswall (2013) do not face this issue because they assume no saving/borrowing. Allowing for saving, as is done here, prevents exact solutions from being derived under the incomplete markets model. However, bounds on consumption and investment in the first period can be derived, which turns out to be enough to prove the result.

Define $\tilde{c}_1$ as the solution to

$$\frac{1}{\tilde{c}_1} = \beta(1+r)\frac{1}{\mathbb{E}_{y_2}[c_2^{IM}(a_2,y_2)]}. \tag{62}$$

Applying Jensen's Inequality, it is the case that $c_1^{IM} < \tilde{c}_1$. Noting that the first-order condition for optimal investment is $I_1 = \gamma\beta\kappa(1-\gamma)c_1$, also define

$$\tilde{I}_1 = \gamma\beta\kappa(1-\gamma)\tilde{c}_1, \tag{63}$$

where it is also the case that $I_1^{IM} < \tilde{I}_1$. Lastly, define $\tilde{a}_2$ as the solution to equation 62. Using these definitions and the budget constraint

$$\tilde{c}_1 = a_0 - \tilde{I}_1 - R\tilde{a}_2, \tag{64}$$

solutions for these counterfactuals can be derived:

$$\tilde{c}_1 = \frac{1}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}(a_0+\pi\tilde{y}) \tag{65}$$

$$\tilde{I}_1 = \frac{\gamma\beta\kappa(1-\gamma)}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}(a_0+\pi\tilde{y}) \tag{66}$$

$$\tilde{a}_2 = (1+r)\frac{\beta(1+\kappa(1-\gamma))}{1+\gamma\beta\kappa(1-\gamma)+\beta(1+\kappa(1-\gamma))}(a_0+\pi\tilde{y}) - \pi\tilde{y}. \tag{67}$$

Now use these bounds to prove the main result. Notice that $\tilde{c}_1 = c_1^{FI}$, and therefor $c_1^{IM} < c_1^{FI}$. This is not terribly surprising as it is the classic precautionary savings result. In the current model we also have the corollary result that $I_1^{IM} < I_1^{FI}$. This shows that parents invest less in younger children in the IM model relative to the FI model.

In this case it can also be shown that $\mathbb{E}_{y2}\left[I_2^{IM}\right] > I_2^{FI}$, or in other words that later investments are higher on average when markets are incomplete. This follows from $\mathbb{E}_{y2}\left[a_2+y_2\right] > \mathbb{E}_{y2}\left[\tilde{a}_2+y_2\right] = \mathbb{E}_{y2}\left[b(y_2)+y_2\right]$. This relates to the ranking of terminal investment decision rules reported in section 3.2. There it was reported that, under incomplete markets, the terminal investment rule is weakly smaller than that under full-insurance. Here the decision rules are strictly equal because parents have no precautionary motive at period 2. However, it was also reported in 3.2 that actual investments depended on wealth, which is a function of prior decisions. Here early precautionary saving leads to increased later investment. In the quantitative model, this extra saving competes against the precautionary motive that parents still posses when their child is age T, and the precautionary motive dominates. The upshot is that while we can prove that timing is distorted in general, whether or not investments are lower at all ages is a quantitative question.