The Effects of Family Structure on Children's Outcomes

Preliminary

Rebecca Lessem* and Carl Sanders†

July 8, 2025

Abstract

Children born in low-income households to unmarried parents develop lower levels of human capital, which leads to worse outcomes in adulthood. In this paper, we study the role of family structure on children's outcomes, focusing on an urban sample of children born in low-income households. To quantify the role of parental relationship status, we create a structural dynamic model of mothers' relationship choices, labor market choices, and child development, and estimate it using measurements on children's outcomes and parents' relationships drawn from the Fragile Families panel dataset. Our counterfactuals consider the effects on children of financially incentivizing marriage between birth parents, and we find that even large increases in marriage rates are associated with only small benefits in child outcomes. Part of the results are due to parent quality, as incentivizing marriage causes more relationships with low quality partners to persist, which can have detrimental effects on child outcomes.

JEL Codes: J12, J13

*Carnegie Mellon University, rlessem@andrew.cmu.edu

†We thank participants at the University of Rochester, Bocconi University, Penn State, Washington University in St. Louis, the University of Colorado-Boulder, the Society for Economic Dynamics, the Barcelona GSE summer forum, Workshop on Labor Supply and Inequality within Families (University of Copenhagen), SITE, Southern Economics Association, and the Midwest Economics Association for helpful comments.

1 Introduction

It is well known that children born into low-income households have worse outcomes later in life (Heckman (2008)). There is a large literature showing that human capital development early in life is crucial to improve outcomes as an adult,[1] and that parents play an important role in this development process (Agostinelli and Wiswall (2025), Cunha et al. (2010), Boca et al. (2014), Boca et al. (2025)). In many low-income households, the parents are unmarried, meaning that children may not fully access these human capital gains from interactions with their parents. Policymakers on both the left and the right sides of the political spectrum have pushed for programs that promote two-parent married households, believing that marriage gives children the most beneficial family structure. However, the evidence on the effect of family structure on children's outcomes is limited.[2]

In this paper, we estimate the effects of family structure on children's academic, physical health, and mental health outcomes. Our primary thought experiment involves subsidization of marriage or cohabitation: can providing financial incentives to mothers to stay with the children's father improve the outcomes for children? Policies incentivizing or disincentivizing marriage exist today – for example, in the US the income tax system creates different incentives for parents to become or remain married depending on their relative incomes. There are a priori arguments that encouraging mothers to stay with birth fathers may be good or bad: staying with a partner may increase resources available to children, but at the same time we may be inducing low-quality relationships to survive, and these low-quality marriages may lead to conflict that hurts children.

Past work has studied how divorce or parental death affect children's outcomes; our work builds on this literature by additionally considering the impacts both when the parents are never married or when they re-partner, which are common phenomena in low-income areas with high rates of unmarried births. Tartari (2015) finds that if parents remained together rather than dissolving a marriage, their children would have performed better on academic achievement tests. Brown et al. (2025) also analyzes the impact of divorce, focusing on endogenous parental investment decisions using time diaries. Lang and Zagorsky (2001) find that early parental death has small impacts on economic outcomes later in life. Our research examines a broader set of relationship outcomes: parents can be unmarried and cohabiting, unmarried and unpartnered, the mother can be remarried or repartnered with another individual, etc., and each of these may have different effects on children. Given the

[1]See, for example, Heckman (2008), Almond and Currie (2011), Heckman et al. (2013), among many others.

[2]See Haveman and Wolfe (1995) for a summary of the reduced form literature that measures the association between child outcomes and parent's marital status.

high rates of children born to unmarried parents, with unmarried births over-represented in urban and lower income areas, understanding the importance of family structure for children even outside of the context of divorce is an important question to address inequality in children's outcomes. Furthermore, Tartari (2015) only looks at only one dimension of child outcomes, whereas we broaden this to look at mental and physical health along with academic outcomes, an important contribution given the results in Heckman et al. (2006) showing that noncognitive skills, such as motivation and persistence, impact schooling decisions and wages. Our model also allows the mother to choose labor market status each period, which is a decision correlated with marital decisions and also affects children's outcomes (Bernal (2008), Bernal and Keane (2011)).

To estimate the effects of family structure on children, we develop and estimate a dynamic model where mothers choose relationship and labor market status, and her choices affect the path of her child's development. We use a three-dimensional vector to describe the child's outcomes, encompassing academic success, physical health, and mental health. Decisions and outcomes depend on the quality of the father. This is important because father quality impacts the likelihood that a relationship persists, which affects the child's outcomes, but father quality also directly impacts the child. We estimate the model using data from the Fragile Family and Child Well-Being Study, a survey that starts from a sample of mothers and children at birth living in urban areas. The survey follows the mother and child up to child's age 15, allowing us to track the child's academic, physical health, and mental health outcomes as well as the relationship status of their parents.

Estimation of the model is technically complex because child outcomes and father quality are not observed directly in the data; instead, we have a wide array of noisy measures of these latent variables. We combine a standard dynamic discrete choice problem with unobserved state variables with the "measurement model" approach to the estimation of latent factors. In our measurement model, we treat observed data as correlated but noisy measures of the child's outcomes in the model, and jointly estimate the choice parameters of the model with the "signal" and "noise" parameters of our measures.

Our counterfactuals show that sizable cash transfers tied to parents being married have small effects on children's academic outcomes, but almost no effect on children's health of mental health outcomes. Put another way, our model explains the vast majority of different outcomes across children of different parental relationships via non-causal factors. Given the substantial costs of this policy, unless a policymaker had intrinsic interest in keeping couples together, the cost/benefit calculation from our results would favor other approaches to helping children of unmarried or single parents. Additionally, we use machine learning techniques to test if there are demographic sub-groups that are more affected from this policy;

if this is the case, then targeted policies to these groups could be effective. However, we find no evidence that any demographic groups are more affected by this policy than others.

We actually find stronger effects on the reverse side of the relationship between family structure and children's outcomes in second counterfactual: exogenously increasing child academic, health, and mental health quality in the initial period increase relationship stability between birth parents and lead to more marriages persisting. In the counterfactual, mothers effectively put less weight on children's quality since the children are now exogenously doing well, and mothers can now make the decision to stay with the birth father since the potential downsides to the children are mitigated.

Our research analyzes the effects of family structure on multiple dimensions of child quality. This contributes to a large literature that has continuously demonstrated the importance of looking at cognitive as well as non-cognitive skills; for example, see Attanasio et al. (2020b), Attanasio et al. (2020a), Cunha and Heckman (2008), Cunha et al. (2010), Heckman and Rubinstein (2001) and Heckman et al. (2006), Heckman et al. (2013), Heckman et al. (2006). We might be concerned that mothers' decisions to leave the birth father, for example, show up in children's mental health without necessarily showing up in academic test scores. Similarly, single mothers may have difficulty providing financial resources for healthy children's eating, even if the mental health environment may be better without a partner in the household. Rather than combining these different aspects of children into one "quality" measure, we allow for three different components of children that parents care about: academic outcomes, physical health, and mental health. In the model, these three components can be correlated within children and potentially evolve independently, and mothers may have different attitudes toward these components.

Most of the past work on this topic studies the impact of divorce on children's outcomes; we extend this by also considering what happens when the parents are never married. Additionally, we extend the literature by considering the effects of non-marriage relationships, such as cohabitation or single parenting, in children's outcomes. Cohabitation is increasingly observed in the data, but we have a limited understanding on how it affects children's outcomes. This approach of considering non-marriage relationships (alongside married relationships) complements papers like Tartari (2015) and Brown et al. (2025), which focus on the decision of married parents to dissolve the marriage. Cohabitation has been explored in past work, such as Brien et al. (2006) who find that couples use cohabitation to learn about their partner before becoming married or choosing to separate. This mechanism fits cleanly into our model, which allows for different transition rates out of cohabitation versus marriage, and we extend this work by examining how these decisions affect child outcomes. This is important because the less-stable cohabitation relationship could differentially impact children's outcomes, and our research allows us to examine this empirical impact.

To learn about the child and father quality, both of which are unobserved, we use multiple noisy measures of each dimension of children's qualities as well as the quality of the father. To do this, we use the "measurement model" approach that treats the quality measures as latent variables which can only be imperfectly measured. We use a large number of measures of father quality and children's quality, and the measurement approach allows us to directly estimate the informativeness of each measure. This measurement model approach has been used extensively in the child development literature, such as Cunha et al. (2010), Agostinelli and Wiswall (2025), Boca et al. (2014), and Agostinelli (2018), amongst many others. We estimate the measurement parameters within the model, which allows for selection bias in the measurement parameters, and derive approximation results that allow for using a large number of measurements within a single-stage maximum likelihood framework.

2 Data and Descriptive Statistics

Our sample is drawn from the Fragile Family and Child Well-being Study (hereafter Fragile Families), a longitudinal survey of parents and children administered jointly by research groups at Princeton and Columbia Universities. This survey is based in 20 urban locations, and is representative of urban areas but not the country as a whole. Given our interest in how marital status at birth affects outcomes, this sample restriction does not seem too costly given that we have coverage of the group most at risk.

Parents were surveyed at the time of the birth of the child, responding to a range of questions on relationship status, household structure, and labor market status, and were re-surveyed at repeated intervals up to when the child turned 15. We have information about the child's physical health at each round of the survey. We also have information on their academic outcomes. This comes through tests to assess cognitive development that were administered as part of the survey, as well as grades in school at later ages. The survey also asks questions of both the parents and children to assess psychological health. We will use the panel data structure to link together relationship status and changes in children's outcomes over time. The surveys were done at birth (year 0), as well as at ages 1, 3, 5, 9, and 15.

Table 1 shows some basic descriptive statistics on the demographic variables we use in our analysis. First we look at the mother's education, where we split the sample into two groups, low (high school degree or less) and high (at least attended college) education. About 60% of the sample is in the high education group. We next look at the racial composition of the sample; about a third of the respondents are white, half are black, and the remainder are other races. The average age of the mother's at childbirth is just over 25.

Crucially for our analysis, we see the mother's relationship status in each wave of the survey. We see if she is married to the child's father, cohabiting with the child's father, married or cohabiting with a new partner[3], or single[4]. A quick look at the data highlights the importance of not just looking at divorce but also analyzing family formation. About 15% of the parents in our sample divorce in the time frame of our sample. On the other hand, almost 60% of the parents are never married in the sample period. Therefore, to understand the effects of family structure on child development in this sample, it is important to consider the role of family formation as well as divorce. Table 2 shows this in more detail, displaying the share of the sample with each relationship status at each age.[5] We additionally split the sample by the labor market status of the mother. First, we see that about 75% of our sample had parents who were unmarried when the child was born. Cohabitation is common, particularly when the children are younger; however, as the children age, cohabitation rates decrease as the parents tend to get married or split up. Looking at employment decisions, we do not see strong patterns across relationship states.

In Table 3, we show the distributions of current period relationship status, conditioning on prior period relationship status. The relationship status changes significantly over time; for example, of the single households at age 0, almost 20% of the parents cohabit with the birth father at age 1. Of those who cohabit at birth, 55% are married at age 1. We continue to see transitions between states each period, meaning that children in these households experience transitions in their household structure over the first 15 years of their life. Marriage and cohabitation seem fairly persistent; however, we see that married parents are less likely to transition into the single state than those who are cohabiting. For example, less than 5% of the couples who are married at age 0 are single by age 1, as compared to almost 30% of couples who are cohabiting at age 0. We use the observed trends in the data to motivate our model structure. Since Table 3 shows that movement between states is fairly common, it is important to use a dynamic model to allow for the chosen path to affect the child's outcomes. In addition, compared to most previous work that does not separately look at cohabitation, we see that this is fairly common in the data. Therefore, it is important to include this state to understand how it affects children's outcomes.

[3]We combine these groups due to sample size limitations.

[4]This includes being single or being in a relationship but not living together.

[5]At age 0, we only see the mother's relationship with the child's birth father. In all other periods, we also see her relationship status with other individuals.

2.1 Child outcomes

In our model, family structure affects child outcomes. In the data, we see a wide variety of information about the child, encompassing academic outcomes, physical health, and mental health. In this section, we explain the measures that we use to learn about the child outcomes.

Table 4 shows shows the measures that we use for academic outcomes as well as summary statistics split by the relationship status of the child's parents. As part of the survey, various academic tests were administered to the children. In particular, we use the Peabody Picture Vocabulary Test (PPVT) at ages 3, 5, and 9, and the Woodcock Johnson vocabulary tests at ages 5 and 9. At age 15, children self-report their grades in Math, English, Science, and History high school classes. Across the board, we see that the means are higher when the parents live together, but the differences are not always statistically significant, and these are unconditional means. The model will inform us to how the parental decisions impact these outcomes over time.

Tables 5 and 6 show our physical health measures and summary statistics. For each period, we have self-reported health, given on a scale of 1-5. We also have other measures of physical health, such as physical disabilities, being overweight, or other medical problems. Again, we mostly see indicators of better physical health when the child is living with both parents.

We also have information on mental health outcomes, including information about depression, behavior problems, and social skills. Table 7 lists the measures we use to learn about each child's mental health outcomes as well as summary statistics. Similar to the other measures, it appears the unconditional outcomes are better when the parents live together, but the model will inform us as to how family structure impacts these outcomes.

2.2 Father quality

One unique feature of the Fragile Families data is that we have information on the quality of the fathers. This is particularly informative for our research given that, in our model, the quality of the father both affects the likelihood of the relationship persisting and the child's outcomes. At each round of the survey, we have information on the father's drug or alcohol abuse, emotional or physical abuse, and jail time or criminal convictions. We use this information to infer about the likelihood that the father is "high quality" in the estimation of our model.[6] Table 8 shows the summary statistics on the measures we use for the dad's

[6] This information also exists on the mother. However, the quality of mothers and fathers is highly correlated, so in this version of the estimation we are only using the information on the father quality.

quality, which demonstrate relatively high frequencies of some of these bad qualities.[7]

3 Model

We use a discrete choice model where the mother chooses a relationship and employment status at each period. Decisions are made to maximize lifetime utility, which depends on relationship and labor market status, income, as well as the child's quality and the father's quality.

Choices The mother chooses relationship status and whether or not to work each period. Specifically, she chooses relationship status rel_t from the set

$$J_R = \{Married_f, Cohab_f, MC_n, Single\},$$

where $Married_f$ is married to the child's birth father, $Cohab_f$ is cohabiting with the child's birth father, MC_n is either married or cohabiting with a new partner,[8] and $Single$ is neither married or cohabiting with anyone. Mothers also choose whether to work or not work from the set J_W, where

$$J_W = \{Work, NotWork\}$$

We are assuming that if she works, she works full time.[9] To simplify notation, we write the relationship rel_t and employment choices emp_t each period as $s_t = \{rel_t, emp_t\}$.

Initial conditions The first initial condition is the child's quality q_0^c, which we assume is 3 dimensional to allow for different dimensions of quality. The second initial condition is the father's quality q^f, which is binary and time invariant. These are both unobserved by the econometrician. Additionally, we see time invariant demographic characteristics X and the relationship and labor market status at the child's birth s_0.

State variables The state variables are prior relationship and labor market status s_{t-1}, as well as child quality q_t^c, which is 3-dimensional and evolves as a function of decisions.

Utility function Utility depends on child quality q_t, as well as your current and prior relationship and employment status. This allows for different utility levels for being married versus cohabiting, for example. Utility depends on your prior period relationship status to

[7] The measures reported are time invariant to simplify the model computation. We construct these time invariant measures by reporting the dad to have the bad characteristic if he ever reports engaging in any of these behaviors.

[8] We combined these 2 categories given that these occur at a relatively low rate in the data

[9] This is abstracting from full versus part time work, which we acknowledge could be important; however, we do not see enough information on full versus part time work in the data to make this distinction.

allow for transition costs. Labor market status affects utility through a disutility of work term. Utility also depends on your demographic characteristics and the father's quality q^f. This can affect your utility of being with the father rather than another person or being single. Combining all of these components, we write the utility function as $U\left(q_t^c, s_t, s_{t-1}, X, q^f\right)$.

Timing At the start of the period, the mother sees her state variables, which include the child's quality at that time. The mother then chooses her relationship and labor market status for that period. Her utility depends on the child quality, which is already known for that period. Decisions, however, affect the child's quality in the next period, meaning it also affects her utility in the next period.

Value function At the start of the period, the mother knows her child's quality q_t^c, her prior relationship and labor market status s_{t-1}, her demographic characteristics X, the father's quality q^f, and her payoff shocks η_t, which we assume follow the extreme value type I distribution and are iid across choices and time. A mother picks both relationship and labor market status each period. We write the value function as

$$V_t\left(q_t^c, s_{t-1}, X, q^f, \eta_t\right) = \max_{j \in J_R \times J_W} v_t\left(j, q_t^c, s_{t-1}, X, q^f\right) + \eta_{jt}. \tag{1}$$

The value function has both a deterministic component, $v(\cdot)$, and a random component, η_t. We write the deterministic component as

$$v_t\left(j, q_t^c, s_{t-1}, X, q^f\right) \quad = \quad U\left(q_t^c, j, s_{t-1}, X, q^f\right) + \beta E\left[V_{t+1}\left(q_{t+1}^c, j, X, q^f, \eta_{t+1}\right)\right] \tag{2}$$

Child quality evolves deterministically, as a function of the mother's choices, as follows:

$$q_{t+1}^c = h\left(q_t^c, s_t, X, q^f\right). \tag{3}$$

When the mother chooses a relationship and labor market status, she does this while considering how the decision will affect the evolution of her child's quality. The evolution of each dimension of child quality depends on the prior quality, relationship and labor market choices in the period, a person's characteristics, and the father's quality, as given by equation (3). We allow the quality to depend on prior quality most obviously to allow for persistence. In addition, the quality growth could depend on the prior level. For example, for academic outcomes, students with higher quality measures could learn faster and then therefore have higher quality growth. On the other hand, there could be decreasing returns to scale. In the estimation, we use a flexible functional form so that the data will inform us as to the sign of this effect. Child quality growth also depends on relationship status of the mother, allowing, for example, for different quality evolutions for children of married versus separated parents. Child quality transitions also depend on the labor market status of the mother. For one, mothers who work may have less time to invest in their child. On the other hand, labor market participation increases the household income, which directly affects child quality. We allow child quality evolution to depend on father quality given that the environment in the household can affect the child's development.

We solve for the expected continuation values. The only unknown future value, conditional on choices, are the extreme value payoff shocks. We compute the expected continuation values as

$$E\left[V_{t+1}\left(q_{t+1}^{c}, s_t, X, q^f \eta_{t+1}\right)\right] = \log\left(\sum_{j \in J_R \times J_W} \exp\left(v_t\left(j, q_{t+1}^{c}, s_t, X, q^f\right)\right)\right) + \gamma \qquad (4)$$

In equation (4), γ is Euler's constant.

We can use the properties of the extreme value distribution to compute the probability that a person chooses a given relationship and labor market status in a period. This takes a logit form as follows:

$$P\left(s_t | s_{t-1}, X, q_t^c, q^f\right) = \frac{\exp\left(v_{t+1}\left(s_t, q_t^c, s_{t-1}, X, q^f\right)\right)}{\sum_{j \in J_R \times J_W} \exp\left(v_t\left(j, q_t^c, s_{t-1}, X, q^f\right)\right)} \qquad (5)$$

Using equation (5), we can calculate the probability of a given choice each period, conditional on latent child and father quality.

Terminal Period The model is solved in the form described above from the birth of the child until the child is 15. Child quality is fixed after age 15, but the mother continues solving the relationship choice problem for another 35 periods. This approach ensures that the continuation value associated with each final relationship state is not overweighted relative to the remainder of the mother's life, as well as allowing the child quality from age 15 onwards to matter to the mother. The final period of the mother's problem has a terminal value of 0.

4 Measurement model for latent variables

The model is written conditional on child and father quality, both which are not directly observed in the data. Instead, they are observed with some error. In the data, we see various measures of the child's academic, physical health, and mental health outcomes. We use these to learn about our 3 dimensional quality vector q_t^c. We also see a variety of measurements about the father's quality which we use to learn about q^f.

We use a large number of measures of child quality in each period, as listed in Section 2.1. We have some measures (self reported health, for example) in all periods, but most of the other measures are unique to a given wave of the survey. Our specification is very flexible, as it allows for a different number of outcomes in each period. Additionally, we do not need to assume that the the same-name tests (e.g. Woodcock-Johnson academic tests) in different periods measure the same skills or measure skills with the same parameters.

Assume that we have a set of M^c measures of child quality. For measure of child quality j which is observed at time t, we write

$$m_{jt} = \alpha_{1jt} + \alpha_{2jt}q^c_{1t} + \alpha_{3jt}q^c_{2t} + \alpha_{4jt}q^c_{3t} + \alpha_{5jt}\varepsilon_{jt},\ j = 1, 2, \ldots, M^c. \tag{6}$$

We write outcome j at time t as a function of all dimensions of child's quality. The α_{2jt}, α_{3jt}, and α_{4jt} parameters describe the relationship between each dimension of child quality and a given outcome. We assume that the ε_{jt} terms are drawn from the standard normal distribution. Therefore the standard deviation of each measure is α_{5jt}. Our methodology estimates a set of parameters for each of the outcomes. However, we do have to make some normalizations. There are three dimensions of quality, which in practice will refer to academic outcomes, physical health, and mental health. For identification, we anchor one measure to each of the 3 dimensions of quality. The first dimension, which we call the academic quality, is anchored to the WJ test at age 9. The second dimension, which is physical health, is anchored to self-reported health at age 3. Mental health is the third dimension, and is anchored to the measure of the youth being impulsive at age 15. For each of the anchored measurements, we set the α parameters equal to 0 for the other dimensions of quality.[10]

Given the model parameters, relationship and labor market decisions, and given initial q^c_0, we can predict quality for each individual and period. In particular, as shown in equation (3), quality evolves deterministically as a function of relationship status, labor market status, demographics, and previous quality. This allows us to predict each quality measure, which we denote as $\hat{m}_{jt}$. In particular,

$$\hat{m}_{jt} = \alpha_{1jt} + \alpha_{2jt}q^c_{1t} + \alpha_{3jt}q^c_{2t} + \alpha_{4jt}q^c_{3t}$$

We also know that the standard deviation of each quality measure is α_{5j}. Since the ε terms

[10]These normalizations do not affect the counterfactuals (other than the interpretation), as the remainder of the model has all components linear in parameters, so the full set of parameters could be rotated if new normalizations were chosen, and would generate the same distribution of observed data.

follow the normal distribution, we can calculate the likelihood of an observed measure. We will use this to construct the likelihood function in Section 5.2.

The model has an additional latent variable, which is the quality of the father q^f. We assume that the underlying variable is discrete and equal to 0 or 1, to indicate a low or high quality father, respectively. We make this assumption because most of the measurements on father quality are yes/no questions, making the normality assumption inappropriate here. We additionally assume that the father quality is static and does not evolve over time. Assume we have a set M^f of measures of father quality. For each measure $m_f \in M^f$, we estimate two parameters:

$$\Pi_k \quad = \quad \{\pi_{1kf} \equiv \Pr(m_{fk} = 1|\xi_f = 0), \pi_{2kf} \equiv \Pr(m_{fk} = 1|\xi_f = 1)\}. \tag{7}$$

The term π_{1km} tells is the probability that a measurement indicates a high quality father (is equal to 1) when the true father is low quality, which we could call the "false positive" rate of measure k. Similarly, π_{2km} gives the "true positive" rate of measure k. The true and false negative rates can be calculated as 1 minus these respective rates. For a good measure of father quality, we would expect that $\pi_{1km} \ll \pi_{2km}$, which would indicate that the measurement being equal to 1 is highly informative that the underlying quality is high and equal to 1. These measurements will be incorporated into the likelihood function in Section 5.2.

5 Estimation

We jointly estimate the choice model developed in Section 3 and the measurement model explained in Section 4 using maximum likelihood. In this section, we first explain the parameterization of the model, and then derive the likelihood function.

Each period, a mother picks her relationship status and labor market status. There are four relationship states (married to birth father, cohabiting with birth father, married or cohabiting with someone else, and single) and two labor market outcomes (working, not working). We denote $s_t \in \{1, 2\}$ as the states for married to the birth father, $s_t \in \{3, 4\}$ as the states cohabiting with the birth father, $s_t \in \{5, 6\}$ as married or cohabiting with a new partner, and $s_t \in \{7, 8\}$ as the single states. For each of these pairs, the first number indicates the state where she is working, and the second is when she is not working. For example, $s_t = 1$ is a woman who is married to her child's birth father and works, and $s_t = 2$ is a woman who is married to her child's father and does not work.

When estimating the model, we assume an annual discount factor $\beta = 0.95$.[11]

5.1 Parameterization

Single index: To simplify computation and the precision of our estimation, we first create a single index term that combines all of the demographics. In particular,

$$SI(X) = X'\nu$$

where X is the demographics we use in our analysis. This includes mother's education, race, the child's gender. We also include information on the mother's wages, calculated as take the average of the mother's wage for all periods where she has recorded wages, converted into percentiles.[12] The parameters ν allow us to combine the effects of the demographics into a single index, reducing the number of parameters associated with a higher-dimensional X.

We allow the paths of child quality and the utility function to depend on the single index instead of all of the demographic characteristics, which simplifies computation. As explained below, we allow for the quality transition function for each dimension of quality to vary for each relationship and labor market status combination. This means that we estimate 8 separate functions for each of the 3 dimensions of quality, meaning that we would have to add 24 parameters to just include education in the process (if we made education a binary variable). The same would be true for every other characteristic in the single index. Our dataset is not large enough to estimate these many parameters. The single index allows us to exploit the heterogenity in the data without facing these computational problems or data limitations.

[11]Our data are spaced out at different intervals. For example, families are surveyed when the child is 1 and 3. For a discount rate between those 2 surveys, we use β^2. We similarly adjust the other periods to account for the gap in the data.

[12]For some mothers, we do not have any wage information for any period. For these mothers, the wage percentile is set to 0, but we additionally include a dummy variable that equals 1 if the mother has no recorded wage information in our data.

Utility function We parameterize the utility function as

$$
\begin{aligned}
U\left(q_t^c, s_t, s_{t-1}, X, q^f\right) &= \tilde{u}\left(rel_t, rel_{t-1}\right) \\
&+ \lambda_1 q^f \mathbb{1}\left\{rel_t = Married_f \text{ or } rel_t = Cohab_f\right\} \\
&+ \lambda_2 SI(X) \mathbb{1}\left\{rel_t = Single\right\} \\
&+ \lambda_3 \mathbb{1}\left\{emp_t = NotWork\right\} \\
&+ \lambda_4 SI\left(X\right) \mathbb{1}\left\{emp_t = NotWork\right\} \\
&- \lambda_5 \mathbb{1}\left\{emp_t = Work \text{ and } emp_{t-1} = NotWork\right\} \qquad (8) \\
&+ \kappa' q_t^c \qquad (9)
\end{aligned}
$$

The first term $\tilde{u}\left(\cdot\right)$ gives the net utility of each relationship transition, meaning it gives the utility of a given state minus the cost of moving to that state. Since there are 4 relationship choices, this is a 4x4 matrix. The term λ_1 gives the utility gain from being in a relationship relationship with a high quality father. This only matters if you are married or cohabiting with the child's father. The term λ_2 tells the relationship between the single index and utility for single people. This allows for a relationship between demographic characteristics and the utility of being single, which adds empirical flexibility. The next set of terms relate to labor market status. The term λ_3 is the utility bonus from not working. The term λ_4 gives the relationship between the single index and the utility of not working. The term λ_5 gives the cost of transitioning from out of the labor market into the labor market. The last term tells the relationship between each dimension of child quality and utility.

Evolution of child quality: We estimate a separate child quality evolution function for each relationship and labor market combination s. There are 4 relationship choices, and 2 labor market choices, so there are 8 total combinations. Additionally, since there are 3 dimensions of child quality, we estimate 24 child quality evolution functions. For each relationship and labor market combination s, we estimate the following quality transition functions

$$
q_{t+1}^c = \delta_{1s} + q_t^c \left(1 + \delta_{2s}\right)' + SI(X)\delta_{3s} + \delta_{4s}\mathbb{1}\left\{Divorce_t\right\} + \delta_{5s} q^f \ . \qquad (10)
$$

The first component of quality in the next period is an intercept term δ_{1s}. Next, child quality growth depends on the child quality in that dimension in the prior period. If $\delta_{2s} > 0$, higher-quality children have faster quality growth, while $\delta_{2s} < 0$ implies a "catch-up" effect where lower quality children have faster quality growth. The next term allows the single index term to affect child quality growth, where the term δ_{3s} gives the relationship between the single index and quality growth. We allow for a shift in quality by δ_{4s} in the period after the parents get divorced. The last terms allows for a relationship between child quality growth and father quality.

Initial child quality: Equation (10) gives the child quality transition function. We need to normalize the distribution of child quality in the first period. We assume that the initial quality is drawn from the normal distribution with a mean that is a function of the single index term and variance-covariance matrix C_0,

$$q_0^z \sim N\left(\kappa^z \cdot SI\left(X\right), C_0\right)$$

We normalize the variance of each dimension of quality to be 1 and estimate the covariance terms.

5.2 Likelihood

The likelihood has 3 components: the choice probabilities, the measures of child quality, and the measures of father quality. The choice probabilities have a logit form and were derived in equation (5) and can be written as $P\left(s_{it}|s_{i,t-1}, q_{it}^c, X, q_i^f\right)$ for household i. Denote $\mathbb{S}_i$ as the series of all relationship and labor market choices over all periods. For measure m_{it}^c of child quality observed at time t, we write the density function as $g\left(m_{it}^c|q_{it}^c\right)$. This comes from the normal distribution as explained in Section 4. We write M_i^c as the full set of child quality measures for the child in household i. The last component of the likelihood function is the the probability of each dad quality measure, which is derived in equation (7).

We start by deriving the likelihood over all periods for household i, conditional on initial q_0 and a given value of q^f:

$$\tilde{L}\left(\mathbb{S}_i, M_i^c, M_i^f|q_0^c, X, q^f, ; \Theta\right) = \prod_{t=1}^{T} P\left(s_t|q_t^c, s_{t-1}, X, q^f\right) \times \prod_{m^c \in M^c} g\left(m^c|q_t^c\right) \times \prod_{m^f \in M^f} \Pr\left(m^f|q^f\right) \quad (11)$$

We assume that a father is high quality with probability ξ. We also integrate over the initial distribution of q_0, which we denote as $h(q_0|X)$. Then we can write the likelihood function as follows:

$$\begin{aligned} L\left(s, M^c, M^f|X; \Theta\right) &= \int \Big(\xi \tilde{L}\left(s, m^c, M^r|q_0^c, X, q_f = H; \Theta\right) \\ &+ (1-\xi)\tilde{L}\left(s, m^c, M^r|q_0^c, X, q_f = K; \Theta\right)\Big) h\left(q_0^c|X\right) dq_0^c \quad (12) \end{aligned}$$

The difficulty in estimation of this model comes through the required integration over the initial unobserved quality. In Appendix A, we describe our computational method to deal with a large number of measures and these unobserved initial qualities, similar to a full-information version of the Conditional Choice Estimator proposed in Hwang (2024).

6 Results

In this section, we present and explain our parameter estimates.

Single Index Table 9 shows the single index parameters. For race, the excluded group is whites, showing that white children have lower single index values than other racial groups. We also allow the mother's wages to affect the single index value. To do this, we take the average of the mother's wages for all periods where it is reported. We then calculate percentile values of these average wages. These parameter estimates indicate that higher wages increase the single index. For some mothers, we do not have any wage data. We set the wage percentile 0 for this group. To allow for a level difference for women who do not report any wage outcomes, we include a dummy variable that equals 1 if we do not have any wage information on a mother. This coefficient is small and not statistically significant. We also include the mother's age and the baby's sex in the single index. Older mothers have higher values of the single index, and baby's sex does not have a statistically significant effect on the single index. Figure 1 shows the distribution of the estimated single index values.

Utility Table 10 shows the first set of utility parameters, the net utility of each relationship transition. Recall that there are 4 relationship states, and we report the net utility from transitioning between each pair of states. Because these are only identified relative to each other, we normalize the utility from starting the period as married to the child's birth father and then staying in that state at 0. Therefore each of the other utilities are relative to that state. The estimated results are intuitive. For example, compare the net utility of being married with the child's father and then becoming single with the net utility of cohabiting and then becoming single. The net utility when you start married is negative, indicating that this is relatively uncommon in the data. This reflects the fact that marriages tend to be relatively stable. On the other hand, the coefficient for cohabiting to single is close to 0, indicating that it is more common, all else equal.

Table 11 shows the remaining utility parameters. As father quality increases, utility increases. The cost of re-entering the labor market is positive, reflecting the empirical fact that people do not transition from not working to working at a high frequency. We also estimate the the relationship between each dimension of child quality and utility. Academic quality appears to have the strongest impact on mother's utility.

Child quality We assume the initial quality of the child follows the normal distribution. Table 12 show the parameters of this distribution. The mean of each dimension of quality depends on the single index, although the parameter estimates show that none of these are statistically significant.[13] The bottom half of the table shows the covariance between each of the dimensions of initial quality.

Next, we look at the evolution of child quality, as given in equation (10). We show the constant term, the effect of the single index, prior quality, divorce, and relationship quality in Tables 13-16 for each relationship status.[14] Because these terms are only defined relative to one another, we set the constant terms for children whose mothers are single and not working to 0.

Measurement of child quality The measurement parameters for child quality are shown in Table 17-19. We use a wide variety of measures to learn about the children's outcomes. A strength of our approach is that we can combine all of these factors to learn about quality, and let the variation in the data tell us the relative importance of each one (as compared to simply taking an average, which would be difficult given that these are all measured on different scales). The measurements are split into 3 tables, based on our judgment on which dimension of quality most impacts that measurement. However, the estimation is flexible, allowing all dimensions of quality to impact each measurement.[15]

To interpret these parameters, we compute the reduction in variance in each dimension of child quality if we were just to learn just the value of one measure. If the reduction of variance for a given measure is large, this means that the measure is informative in learning about that dimension of quality. We do this for each of the measures and dimensions of quality. These results are shown in Figures 2-4. Figure 2 looks at academic quality, and we see that the test scores are the most informative for this dimension. The grades at age 15 are less informative than the tests in earlier periods, suggesting that test scores are better measures of academic outcomes than self-reported grades. Figure 3 looks at the health dimension of child quality.

[13]For identification, we had to fix the impact of the single index on one dimension of child quality. We set the impact of single index on academic quality to 0.28 as a normalization.

[14]There is no divorce term when married, since that state does not exist. For all other relationship states, due to power limitations, we set the impact of divorce on child quality evolution to be the same across states.

[15]For each dimension of quality, we have to index it to 1 measurement. We do this by setting the impact of the other dimensions of quality for that measurement. For academic quality, the indexed measure is PPVT test scores at age 9. For health quality, the indexed measure is self-reported health at age 1. For mental health quality, the indexed measure is bullying at age 15.

We clearly see that the health measurements are strongly informative for this dimension of quality while the others are not. Figure 4 shows this for the mental health measures. In this case, we see that the mental health measures are informative for this dimension of quality. Additionally, we also see that grades at age 15 are somewhat informative, which is interesting given their small role in the academic dimension of quality.

Measurements of father quality For each measure of father quality, we estimate (1) the probability that the measure indicates a high quality father if the father is low quality (false positive) and (2) the probability the measure indicates a high quality father if the father is high quality (true positive). We also estimate the probability that a father is high quality. These parameters are shown in Table 20. We see that all of the measures seem informative in learning about the father's quality. The probability of the dad being high quality is parameterized using the logistic function and depends on the single index. To show the informativeness of the father quality measures, Figure 5 shows the posterior probability that each father is high quality given the observed measure. We see a large spike at 0 and 1, showing we could confidently assign high or low quality status to many fathers given the measures; but there is a significant share without complete confidence in the father's type.

7 Counterfactuals

In this section, we perform counterfactuals to understand the implications of policies that help to support family formation and stability. The counterfactuals are compared with of a baseline set of relationship and labor market choices and child qualities, simulated from the model at our estimated parameters. To generate the counterfactuals, we simply adjust the estimated parameters and re-simulate, holding the unobserved shocks to utility and measurement errors constant.

One important thing to note for the counterfactuals is that we cannot identify the absolute level of child qualities, but only relative to some reference group, since all our child quality measurements have no objective location or scale. In particular, our reference group at time 0 is white women with a low education. In later periods, the reference group is a women with the aforementioned demographics who has been single and not working every period. In the counterfactuals, we will use the baseline distribution of child qualities in each period as a comparison group for the counterfactual children.

How does increased marriage rates affect child outcomes? In the first counterfactual, we increase the utility from being married to the child's birth father by 0.4. This change can be interpreted as any policy that would lead to higher marriage rates and lower divorce rates but does not directly affect the ability of parents to raise their children. Potential real-world interpretations of this counterfactual include increased societal pressure for couples with children to marry, or financial incentives for marriage such as marriage tax credits (to the extent the additional income wouldn't directly affect child quality.) This counterfactual has large increases on marriage rates: 26% and 54% of the sample is married in the terminal period in the baseline and counterfactual, respectively.

The increased utility from marriage in the counterfactuals does not have a direct effect on child quality evolution, holding choices constant, but instead induces changes in child quality via changes in optimal relationship and labor market choices. As more mothers choose to be married, we do not see a meaningful change in the health or mental health dimensions of quality, but we do see an increase in academic quality. In Table 21, we show the child quality in each dimension in the baseline and in each counterfactual. This policy that increases marriage rates has a moderate impact on academic outcomes and smaller impacts on health and mental health quality.

To help interpret the results, we examine how the changed quality outcomes from the counterfactual impact the expected normalized measures in the data. As marriage rates increase due to the counterfactual, the average WJ test scores at age 9 of 0.09 standard deviations. The impacts on the health and mental health measures are smaller: self-reported health at age 3 increases by 0.008 standard deviations, and the measure of bullying behavior improves by 0.01 standard deviations. This shows that the counterfactual seems to impact the academic test scores more than the other measures.

This policy also has a potential adverse outcome – by increasing marriage durations, it leads to more low quality relationships persisting. In the terminal period, 18% of the married or cohabiting relationships are with low quality fathers in the baseline. This increases to 25% in the counterfactual. This puts downward pressure on child quality, partially explaining the small counterfactual impacts on child quality, since persisting relationships with low quality fathers are bad for child outcomes.

The average impacts shown in Table 21 mask heterogeneity within the sample. For each child, we compute the impact of the counterfactual on all 3 dimensions of quality. In Figures 6-8, we show the distribution of these impacts across the sample. When looking at academic quality, we see a wide variation in outcomes, with the estimated impact ranging from 0 to 0.4. For health and mental health, the range of possible impacts is even wider, and we even see some negative estimated impacts for some individuals.

Given this heterogeneity, we next seek to determine which sub-groups of the population benefit most from this counterfactual. This is important from a policy perspective in order to understand which groups would benefit most from higher marriage rates; see Kitagawa and Tetenov (2018) and Athey and Wager (2021). To do this, we use an idea developed in Athey and Wager (2021) in a non-structural context and extend it to our structural model. In the structural model, selection is no longer a concern when it comes to the treatment effects, since under model assumptions we can consistently calculate the distribution of treatment effects at the individual level. For each individual, we simulate their choice and outcome path 10 times under the baseline (control) and ten times with the changed parameters (treatment), and form the individual-level ATE by averaging the difference.

With the individual-level ATEs, we then run a shallow Gradient-Boosted Machine model, restricted to a maximum depth of 3, with the target variable being the ATEs and the independent variables being the mother's demographics and children's quality.[16] We use the LightGBM algorithm of Ke et al. (2017), which will split across the observables to maximize the variance of the across-group ATEs.

Figures 9-11 show the trees detailing which groups benefit the most from the counterfactual. The right-most leaves of the tree give the estimated ATE for that group from the policy as a function of the demographic sub-groups that are formed by the yes/no answers starting from the left of the tree. For example, in Figure 6, the estimated ATE for the group "Child with initial academic quality worse than -1 standard deviations; Low earning mother" is 0.147, while for the group "Child initial academic z-score $>$ -0.47; High quality father; young mother" the estimated ATE is 0.175.

Across all 3 dimensions of quality, we do not see any sub-groups that disproportionately benefit from the policy. This indicates that the average impact of the policy, most importantly for academic outcomes where we see a moderate impact of this policy, is mostly evenly distributed across recognizable demographics in the and is not impacting some groups more than others. This limits the possibility of effective policy targeting. To the limited extent there is heterogeneity in treatment effects for academic quality, higher-SES families tend to see larger benefits.

How does initial child quality affect marriage rates? In the next counterfactual, we consider how child ability affects family formation and stability. It is often thought that struggling children can exert pressure on parents, which could lead parents to separate. On the other hand, it is possible that children with lower outcomes could benefit more from a stable family unit, providing an incentive for the parents to remain together. We empirically examine the role of child quality on family structure in this next counterfactual. To do this,

[16]Another advantage of the structural simulation is that we can use the latent child quality as an independent variable even though it is not actually observed in the data.

we compare the baseline model simulation to a counterfactual in which each child begins life at the 90th percentile of the child quality distribution. Given this new starting point, mothers make optimal choices and the quality of the child evolves as a function of those choices.

We consider a family as "together" if the child's birth parents are married or cohabiting, which occurs 33% of the time when the child is 15 in the baseline. On average, when we increase initial academic quality, we see an increase of 2 percentage points in the rate of parents who are together when the child is 15. We see an increase of 1 percentage points when we increase initial health quality and an increase of 2 percentage points when we increase initial mental health quality.

As before, these average impacts mask heterogeneity within the sample. For each household in the sample, we compute the probability of being married or cohabiting at age 15. We look at how this probability changes as we move from the baseline to the counterfactual. Figures 12-14 show the distribution of the impacts of this counterfactual on marriage or cohabitation rates at age 15. When we increase academic quality, some households have no change in marriage rates, yet in some households the rate goes up by 5 percentage points. When we increase initial health quality, at the top end we see increased marriage rates of 6 percentage points. However, at the left tail of the distribution, we see a negative impact for a substantial portion of the sample. When we increase initial mental health quality, we see the widest distribution of impacts. Some households have a decrease in marriage rates, yet for some households the marriage rates increase by 10 percentage points.

To understand the distribution of these impacts, we again construct a tree to see which sub-groups are most impacted by changes in initial child quality. These are shown in Figures 15-17 for each dimension of child quality. In this case, we are able to identify some variations in outcomes with characteristics when we consider increased initial academic quality. Consider households where the father is low quality and the mother is relatively young. If the mothers wage is relatively high, we see a much larger impact on marriage rates when we increase initial academic quality (almost 5 percentage points versus about 2 percentage points). This suggests that for this sub-group, the impact of child quality on marriage decisions is larger for the households where the mother's wage is high. When we look at increased initial health quality, we do see some sub-groups with larger impacts than others, but overall the average impacts for each sub-group are smaller. For mental health, we do not see any substantial variation aross sub-groups.

8 Conclusion

In this paper, we study how family instability causally affects child outcomes, focusing on cognitive and physical development as well as mental health. We develop and estimate a dynamic model of marriage and child development, where parental marital status affects the evolution of child outcomes over time. We use our results to quantify how much child outcomes would improve if there were policies put in place to encourage parents to stay together. In a counterfactual, we raise the utility from marriage so that there is a 20 percentage point increase in the share of parents that are married in the terminal period. This has a small impact on child outcomes. Although these results are still preliminary, they suggest that policies to increase marriage rates and durations do not have a substantial effect on children's outcomes.

References

Agostinelli, F. (2018). Investing in children's skills: an equilibrium analysis of social interactions and parental investments.

Agostinelli, F. and Wiswall, M. (2025). Estimating the technology of children's skill formation. *Journal of Political Economy*, 133(3).

Almond, D. and Currie, J. (2011). Human capital development before age 5. *Handbook of Labor Economics*, 4b:1315–1486.

Athey, S. and Wager, S. (2021). Policy learning with observational data. *Econometrica*, 89(1):133–161.

Attanasio, O., , Meghir, C., and Nix, E. (2020a). Human capital development and parental investment in india. *Review of Economic Studies*, 87:2511–2541.

Attanasio, O., Cattan, S., Fitzsimons, E., Meghir, C., and Rubio-Codine, M. (2020b). Estimating the production function for human capital. *The American Economic Review*, 110(1):48–85.

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

Bernal, R. and Keane, M. (2011). Child care choices and children's cognitive achievement: The case of single mothers. *Journal of Labor Economics*, 29(3):459–512.

Boca, D. D., Flinn, C., and Wiswall, M. (2014). Household choices and child development. *The Review of Economic Studies*, 81(1):137–185.

Boca, D. D., Flinn, C., and Wiswall, M. (2025). Household choices and child development. Forthcoming in The Journal of Political Economy.

Brien, M. J., Lillard, L. A., and Stern, S. (2006). Cohabitation, marriage, and divorce in a model of match quality. *International Economic Review*, 47(2):451–494.

Brown, M., Flinn, C. J., and Mullins, J. (2025). Family law effects on divorce, fertility, and child investment. *Journal of Labor Economics*, 43:S351–S397.

Cunha, F. and Heckman, J. J. (2008). Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of Human Resources*, 43(4):738–782.

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

Evans, M. and Swartz, T. (2000). *Approximating integrals via Monte Carlo and deterministic methods*, volume 20. OUP Oxford.

Haveman, R. and Wolfe, B. (1995). The determinants of children's attainments: A review of methods and findings. *Journal of Economic Literature*, 33(4):1829–1878.

Heckman, J., Pinto, R., and Savelyev, P. (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *The American Economic Review*, 103(6):2052–2086.

Heckman, J. and Rubinstein, Y. (2001). The importance of noncognitive skills: Lessons from the ged testing program. *American Economic Review*, 91(2):145=149.

Heckman, J. J. (2008). Schools, skills, and synapses. *Economic Inquiry*, 46(3):289–324.

Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24:411–482.

Hwang, Y. (2024). Identification and estimation of a dynamic discrete choice model with time-varying unobserved heterogeneity using proxies. *Available at SSRN 3535098*.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.

Kitagawa, T. and Tetenov, A. (2018). Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616.

Lang, K. and Zagorsky, J. L. (2001). Does growing up with a parent absent really hurt? *The Journal of Human Resources*, 36 (2):253–273.

Tartari, M. (2015). Divorce and the cognitive achievement of children. *International Economic Review*, 56 (2):597–645.

Tables and Figures

Table 1: Descriptive statistics

Percent of mothers with high education	57.58%
Race: white	30.88%
Race: black	49.97%
Race: other	19.15%
Average age of mom	25.21

High education defined as some college or more.

Table 2: Relationship and labor market status of mother

	Married to father		Cohabiting with father		Married/cohabiting with other		Single	
Age	Working	Not Working	Working	Not Working	Working	Not Working	Working	Not Working
0	17.85%	6.50%	25.47%	10.63%			25.85%	13.68 %
1	16.23%	13.66%	14.18%	13.05%	2.43%	2.45%	20.14%	17.88%
3	17.74%	14.31%	10.45%	8.83%	5.23%	4.56%	22.77%	16.10%
5	18.00%	12.90%	7.36%	5.42%	8.91%	7.21%	24.81%	15.39%
9	19.15%	9.98%	5.61%	3.71%	12.47%	8.79%	25.03 %	15.26%
15	19.05%	6.57%	3.30%	1.52%	18.12%	7.13%	30.14%	14.06%

Table 3: Relationship transitions

	Prior relationship status	Current relationship status: Married to father	Cohabiting with father	New partner	Single
Age 0 to 1	Married to father	94.18%	1.07%	0.87%	3.92%
	Cohabiting with father	14.26%	55.39%	3.00%	27.34%
	Single	4.01%	17.82%	9.09%	69.07%
Age 1 to 3	Married to father	90.26%	1.45%	1.71%	6.58%
	Cohabiting with father	13.96%	53.68%	4.34%	28.02%
	New partner	2.22%	2.78%	55.00%	40.00%
	Single	3.32%	11.16%	14.55%	70.98%
Age 3 to 5	Married to father	86.52%	1.93%	2.43%	9.13%
	Cohabiting with father	14.94%	49.79%	5.67%	29.60%
	New partner	0%	0.82%	64.01%	35.17%
	Single	2.41%	6.41%	19.37%	71.81%
Age 5 to 9	Married to father	82.09%	1.72%	4.45%	11.74%
	Cohabiting with father	19.55%	46.54%	9.65%	24.26%
	New partner	1.22%	1.22%	55.99%	42.57%
	Single	2.74%	6.42%	24.67%	66.17%
Age 9 to 15	Married to father	78.18%	1.07%	4.25%	16.51%
	Cohabiting with father	14.79%	45.75%	7.79%	32.69%
	New partner	0.70%	0.18%	61.44%	37.67%
	Single	1.90%	1.82%	26.92%	69.36%

Table 4: Academic outcomes

		Parents live together			Parents do not live together		
Measure	Age	Mean	Standard deviation	Number of observations	Mean	Standard deviation	Number of observations
PPVT	3	87.97	17.46	1,079	83.88	15.72	1,224
PPVT	5	95.89	16.34	930	90.87	15.53	1,330
WJ	5	52.29	28.64	935	46.98	28.32	1,342
PPVT	9	96.26	16.14	1,193	90.79	13.82	1,908
WJ9	9	40.90	25.72	1,187	33.65	23.67	1,902
WJ10	9	53.67	28.85	1,190	43.75	27.80	1,908
English grades	15	3.16	0.79	846	2.84	0.88	1,919
Math grades	15	3.00	0.89	851	2.67	0.96	1,915
History grades	15	3.17	0.86	790	2.85	0.92	1,795
Science grades	15	3.10	0.89	830	2.77	0.92	1,883

PPVT and WJ are standardized tests given to survey respondents at different ages. Grades are on a 1-4 scale.

Table 5: Self-reported health outcomes

Age	Parents live together	Excellent	Very good	Good	Fair	Poor	Number of observations
1	Yes	67.4%	20.6%	9.6%	2.2%	0.3%	2,417
	No	63.6%	21.3%	11.6%	3.2%	0.3%	1,796
3	Yes	63.9%	25.7%	8.6%	1.7%	0.2%	2,092
	No	59.5%	26.6%	11.4%	2.4%	0.2%	1,982
5	Yes	64.4%	25.2%	9.1%	1.3%	0%	1,735
	No	59.6%	27.8%	9.9%	2.5%	0.2%	2,204
9	Yes	59.7%	26.7%	11.9%	1.8%	0%	1,291
	No	53.8%	29.0%	13.4%	3.7%	0.15%	2,033
15	Yes	58.6%	26.9%	12.2%	1.8%	0.4%	925
	No	49.1%	33.3%	13.5%	3.9%	0.3%	2,104

Table 6: Other physical health outcomes

	Parents live together				
Measure	Age	Yes	Somewhat	No	Number of observations
Physical disabilities	1	1.7%		98.3%	2,418
Asthma	1	10.5%		89.5%	2,082
Lacks energy	3	1.1%	4.1%	94.8%	1,598
Physical disabilities	3	2.4%		97.6%	1,598
Asthma, allergies	5	34.7%		65.3%	1,223
Stomach or head issues	5	2.9%		97.1%	1,223
Underactive	5	1.0%	5.1%	93.9%	1,216
Overweight	5	2.0%	4.0%	94.0%	1,217
Allergies	9	27.7%		72.4%	1,291
Head issues/diabetes	9	6.7%		93.3%	1,291
Overweight	9	4.1%	12.8%	83.2%	1,182
Underactive	9	0.9%	7.7%	91.3%	1,190
Limited activities	15	5.8%		94.2%	924
Asthma	15	34.9%		65.1%	925
	Parents do not live together				
Measure	Age	Yes	Somewhat	No	Number of observations
Physical disabilities	1	3.7%		96.3%	1,791
Asthma	1	17.6%		82.4%	1,570
Lacks energy	3	1,598	1.2%	5.9%	92.3%
Physical disabilities	3	3.9%		96.1%	1,554
Asthma, allergies	5	40.1%		60.0%	1,673
Stomach or head issues	5	4.1%		95.9%	1,672
Underactive	5	1.8%	6.0%	92.2%	1,671
Overweight	5	1.9%	4.7%	93.4%	1,671
Allergies	9	30.7%		69.3%	1,409
Head issues/diabetes	9	8.8%		91.2%	2,033
Overweight	9	4.3%	12.6%	83.2%	1,889
Underactive	9	1.2%	6.7%	92.1%	1,899
Limited activities	15	10.6%		89.4%	2,103
Asthma	15	45.2%		54.9%	2,104

Table 7: Mental health outcomes

		Parents live together				Parents do not live together		
Measure	Age	Max value	Mean	Standard deviation	Num. of obs.	Mean	Standard deviation	Num. of obs.
Defiant	3	2	1.40	0.45	1,600	1.33	0.50	1,551
Plays with others	3	2	1.46	0.35	1,598	1.38	0.36	1,550
Sulks	3	2	1.53	0.30	1,599	1.47	0.32	1,550
Disobedient	5	2	1.53	0.41	1,491	1.45	0.45	1,908
Friendly	5	2	1.77	0.25	1,568	1.73	0.27	2,015
Violent	5	2	1.95	0.13	1,222	1.91	0.21	1,672
Excluded	9	5	4.53	0.51	1,210	4.38	0.59	1,935
Destructive	9	3	2.88	0.25	1,210	2.78	0.33	1,935
Argues	9	3	2.76	0.28	1,193	2.70	0.32	1,904
Bullies	15	3	2.95	0.16	925	2.87	0.26	2,104
Disobedient	15	3	2.75	0.34	925	2.60	0.45	2,104
Social	15	5	4.49	0.50	925	4.43	0.52	2,104

All measures are the average value from multiple survey questions. All questions have a minimum value of 0, and the maximum value as indicated in the table. Higher numbers indicate better mental health outcomes.

Table 8: Measures of father quality

Measure	Share	Number of observations
Drug/alcohol use	20.6%	4,288
Physical abuse	15.6%	4,474
Emotional abuse	38.6%	3,837
Incarceration	44.6%	4,472
Depressed	39.4%	3,932

Table 9: Single index parameters

Characteristic	Estimate
Education	0.078
	(0.078)
Black	0.50
	(0.096)
Other race	0.63
	(0.052)
Wage percentile	1.24
	(0.19)
Dummy for missing wages	0.0051
	(0.050)
Mother's age	0.16
	(0.045)
Baby male	-0.028
	(0.035)

We report the effects of each factor on the single index. The wage percentile is based on a person's average wage over all periods, and is set to 0 for people with no wage information. The dummy for missing wages is a dummy variable that equals 1 if we have no wage information for that person. Standard errors in parentheses.

Table 10: Utility from relationship status

Prior relationship	Current relationship status			
status	Married to father	Cohabiting with father	New partner	Single
Married to father	0	-4.01	-2.76	-1.89
		(0.55)	(0.70)	(0.47)
Cohabiting with father	-0.83	0.37	-1.18	0.05
	(0.60)	(0.11)	(0.71)	(0.53)
New partner	-3.31	-3.48	1.03	0.32
	(0.74)	(0.70)	(0.15)	(0.44)
Single	-2.40	-1.26	0.16	0.99
	(0.53)	(0.52)	(0.44)	(0.14)

We report the net utility from each relationship transition. The net utility of remaining married is set to 0 as a normalization. Standard errors in parentheses.

Table 11: Utility parameters

Father quality	0.35 (0.13)
Utility from leisure	0.004 (0.10)
Cost to re-enter labor market	1.28 (0.03)
Effect of academic quality on utility	0.16 (0.05)
Effect of physical health quality on utility	0.063 (0.034)
Effect of mental health quality on utility	0.00004 (0.034)

Standard errors in parentheses.

Table 12: Initial child quality distribution

Dimension of quality	Impact of single index		
Academic	0.28		
Health	-0.094 (0.081)		
Mental health	0.12 (0.076)		
Covariance matrix			
	Academic	Health	Mental health
Academic	1	0.082 (0.23)	-0.27 (0.26)
Health	0.082 (0.23)	1	-0.33 (0.25)
Mental health	-0.27 (0.26)	-0.33 (0.25)	1

Standard errors in parentheses. In the top half of the table, we report the impact of the single index on the mean of the initial quality distribution. We normalize the distribution by setting the impact of the single index on the mean of the academic dimension of initial quality to 0.28. The bottom half of the table shows the covariance between the different dimensions of initial quality. For identification, we set the variances of each dimension of initial quality to 1.

Table 13: Child quality transition parameters, married

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	-0.019	-0.036	0.074	0.062	0.019	0.26
	(0.062)	(0.068)	(0.061)	(0.064)	(0.075)	(0.075)
Prior academic	0.13	-0.067	-0.000021	0.14	-0.068	0.022
quality	(0.056)	(0.076)	(0.068)	(0.055)	(0.083)	(0.063)
Prior health	0.33	0.11	0.064	0.21	0.14	0.11
quality	(0.078)	(0.059)	(0.073)	(0.063)	(0.062)	(0.070)
Prior mental health	-0.046	-0.0081	-0.048	-0.087	-0.081	-0.024
quality	(0.057)	(0.084)	(0.038)	(0.055)	(0.085)	(0.035)
Single index	0.051	0.035	0.031	-0.016	0.036	-0.027
	(0.025)	(0.025)	(0.015)	(0.022)	(0.027)	(0.017)
High quality	0.012	0.23	0.11	0.14	0.14	0.11
father	(0.053)	(0.071)	(0.061)	(0.060)	(0.075)	(0.069)

Standard errors in parentheses.

Table 14: Child quality transition parameters, cohabiting

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	-0.10	0.044	0.025	-0.0030	-0.063	0.20
	(0.056)	(0.071)	(0.060)	(0.058)	(0.084)	(0.076)
Prior academic	0.098	-0.17	0.11	0.10	-0.075	0.16
quality	(0.058)	(0.11)	(0.086)	(0.062)	(0.115)	(0.083)
Prior health	0.26	0.20	0.30	0.22	0.26	0.32
quality	(0.074)	(0.080)	(0.10)	(0.069)	(0.087)	(0.10)
Prior mental health	-0.12	-0.30	-0.030	-0.16	0.28	-0.12
quality	(0.073)	(0.086)	(0.052)	(0.069)	(0.099)	(0.061)
Single index	0.036	0.038	0.036	-0.061	-0.017	0.024
	(0.024)	(0.031)	(0.023)	(0.27)	(0.035)	(0.30)
Separation from	-0.031	0.076	0.031	-0.031	0.076	0.031
father	(0.061)	(0.057)	(0.049)	(0.061)	(0.057)	(0.049)
High quality	0.048	0.079	0.17	-0.0062	0.19	0.085
father	(0.052)	(0.069)	(0.067)	(0.062)	(0.10)	(0.088)

Standard errors in parentheses.

Table 15: Child quality transition parameters, other partner

	Working			Not working		
	Academic	Health	Mental health	Academic	Mental Health	health
Constant	-0.27	0.40	-0.25	-0.25	0.29	0.042
	(0.088)	(0.13)	(0.010)	(0.086)	(0.12)	(0.079)
Prior academic quality	0.11	0.068	0.22	0.10	-0.0036	0.15
	(0.059)	(0.11)	(0.092)	(0.057)	(0.10)	(0.085)
Prior health quality	0.29	0.17	0.14	0.24	0.13	0.14
	(0.080)	(0.083)	(0.090)	(0.072)	(0.081)	(0.094)
Prior mental health quality	-0.11	-0.22	0.024	-0.14	-0.12	-0.018
	(0.053)	(0.096)	(0.061)	(0.058)	(0.10)	(0.058)
Single index	0.032	0.048	0.13	-0.018	0.021	-0.024
	(0.028)	(0.049)	(0.028)	(0.026)	(0.046)	(0.032)
Separation from father	-0.031	0.076	0.031	-0.031	0.076	0.031
	(0.061)	(0.057)	(0.049)	(0.061)	(0.057)	(0.049)
High quality father	-0.41	0.026	0.0086	-0.037	-0.037	-0.011
	(0.089)	(0.17)	(0.11)	(0.14)	(0.25)	(0.16)

Standard errors in parentheses.

Table 16: Child quality transition parameters, single

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	-0.077	0.12	0.061	0	0	0
	(0.044)	(0.063)	(0.050)			
Prior academic quality	0.12	-0.070	0.093	0.083	-0.0061	0.17
	(0.054)	(0.084)	(0.071)	(0.060)	(0.097)	(0.079)
Prior health quality	0.26	0.13	0.15	0.23	0.22	0.17
	(0.069)	(0.063)	(0.081)	(0.067)	(0.080)	(0.087)
Prior mental health quality	-0.081	-0.12	-0.018	-0.14	-0.17	-0.031
	(0.055)	(0.079)	(0.043)	(0.056)	(0.093)	(0.056)
Single index	0.028	0.014	0.015	-0.057	-0.0078	-0.052
	(0.024)	(0.024)	(0.014)	(0.024)	(0.028)	(0.025)
Separation from father	-0.031	0.076	0.031	-0.031	0.076	0.031
	(0.061)	(0.057)	(0.049)	(0.061)	(0.057)	(0.049)
High quality father	0.13	0.070	0.087	0.015	0.13	0.18
	(0.051)	(0.052)	(0.047)	(0.058)	(0.085)	(0.071)

Standard errors in parentheses. The constant term is set to 0 for the not-working state when single as a normalization.

Table 17: Measurement parameters for child outcomes (academic)

Measurement	Age	Constant term	Coefficient on quality dimension: Academic	Physical health	Mental health	Error term
PPVT score	3	-0.020	0.46	-0.034	0.097	0.84
		(0.030)	(0.050)	(0.037)	(0.044)	(0.019)
PPVT score	5	0.016	0.59	-0.030	0.090	0.73
		(0.035)	(0.057)	(0.043)	(0.050)	(0.018)
WJ score	5	-0.013	0.43	0.025	0.15	0.85
		(0.029)	(0.051)	(0.040)	(0.045)	(0.012)
PPVT	9	-0.051	0.59	0	0	0.57
		(0.054)	(0.071)			(0.015)
WJ9 score	9	-0.066	0.52	0.036	0.034	0.69
		(0.049)	(0.065)	(0.026)	(0.024)	(0.011)
WJ10 score	9	-0.086	0.48	0.064	0.075	0.72
		(0.046)	(0.063)	(0.027)	(0.024)	(0.011)
English grade	15	-0.13	0.13	0.067	0.28	0.93
		(0.037)	(0.045)	(0.057)	(0.042)	(0.014)
Math grade	15	-0.12	0.099	0.072	0.27	0.94
		(0.034)	(0.041)	(0.056)	(0.043)	(0.013)
History grade	15	-0.13	0.18	0.087	0.25	0.91
		(0.036)	(0.046)	(0.052)	(0.041)	(0.012)
Science grade	15	-0.12	0.14	0.049	0.23	0.93
		(0.033)	(0.041)	(0.049)	(0.038)	(0.013)

PPVT and WJ are standardized tests given to survey respondents at different ages. The academic dimension of child quality is indexed to the PPVT score at age 9, so for that measure the impact of health and mental health on the test scores are set to 0.

Table 18: Measurement parameters for child outcomes (health)

Measurement	Age	Constant term	Coefficient on quality dimension: Academic	Physical health	Mental health	Error term
Overall health	1	-0.12	0	0.45	0	0.83
		(0.066)		(0.099)		(0.018)
Physical disabilities	1	-0.056	0.0068	0.26	-0.038	0.95
		(0.054)	(0.018)	(0.060)	(0.026)	(0.038)
Asthma	1	-0.12	0.014	0.31	0.018	0.94
		(0.049)	(0.020)	(0.070)	(0.028)	(0.016)
Physical disabilities	3	-0.15	0.000051	0.45	0.015	0.96
		(0.068)	(0.023)	(0.10)	(0.03)	(0.015)
Lacks energy	3	-0.063	0.048	0.14	0.049	0.98
		(0.033)	(0.019)	(0.037)	(0.026)	(0.036)
Overall health	3	-0.061	-0.00060	0.25	-0.0098	0.85
		(0.043)	(0.015)	(0.057)	(0.021)	(0.048)
Overall health	5	-0.14	-0.0047	0.51	0.00052	0.80
		(0.073)	(0.019)	(0.11)	(0.027)	(0.016)
Asthma	5	-0.16	-0.13	0.36	0.12	0.91
		(0.056)	(0.026)	(0.085)	(0.036)	((0.013)
Stomach issues	5	-0.070	-0.026	0.14	0.043	0.99
		(0.030)	(0.018)	(0.037)	(0.025)	(0.050)
Underactive	5	-0.069	0.028	0.17	0.050	0.98
		(0.034)	(0.019)	(0.042)	(0.27)	(0.037)
Overweight	5	-0.090	-0.50	0.12	0.086	0.99
		(0.036)	(0.021)	(0.043)	(0.033)	(0.039)
Overall health	9	-0.25	-0.091	0.45	0.16	0.86
		(0.081)	(0.031)	(0.13)	(0.05)	(0.014)
Allergies	9	-0.16	0.14	0.19	0.18	0.96
		(0.050)	(0.31)	(0.063)	(0.047)	(0.0086)
Headaches	9	-0.16	-0.052	0.17	0.13	0.98
		(0.040)	(0.019)	(0.054)	(0.036)	(0.025)
Overweight	9	-0.094	-0.042	0.13	0.010	0.99
		(0.038)	(0.020)	(0.046)	(0.036)	(0.017)
Underactive	9	-0.15	-0.036	0.15	0.14	0.98
		(0.038)	(0.019)	(0.051)	(0.039)	(0.033)
Overall health	15	-0.25	-0.083	0.41	0.19	0.89
		(0.075)	(0.030)	(0.12)	(0.054)	(0.015)
Activity limited	15	-0.20	-0.077	0.24	0.18	0.95
		(0.050)	(0.024)	(0.075)	(0.048)	(0.024)
Asthma	15	-0.24	-0.14	0.31	0.21	0.92
		(0.067)	(0.032)	(0.10)	(0.054)	(0.011)

Standard errors in parentheses. The health dimension of child quality is indexed to the overall health at age 1, so for that measure the impact of academic quality and mental health on this health measure is set to 0.

Table 19: Measurement parameters for child's outcomes (mental health)

Measurement	Age	Constant term	Coefficient on quality dimension: Academic	Physical health	Mental health	Error term
Defiant	3	-0.40	-0.10	0.26	0.47	0.82
		(0.083)	(0.032)	(0.088)	(0.10)	(0.017)
Social skills	3	-0.35	-0.030	0.27	0.38	0.86
		(0.074)	(0.029)	(0.088)	(0.083)	(0.014)
Sulks	3	-0.46	-0.084	0.25	0.43	0.82
		(0.10)	(0.033)	(.095)	(0.11)	(0.014)
Disobedient	5	-0.57	-0.12	0.24	0.49	0.82
		(0.10)	(0.038)	(0.09)	(0.13)	(0.015)
Social skills	5	-0.43	-0.055	0.22	0.34	0.88
		(0.078)	(0.029)	(0.084)	(0.091)	(0.022)
Aggressive	5	-0.39	-0.064	0.13	0.40	0.89
		(0.085)	(0.026)	(0.055)	(0.10)	(0.033)
Social skills	9	-0.36	-0.0046	0.068	0.31	0.88
		(0.076)	(0.017)	(0.033)	(0.088)	(0.015)
Destructive	9	-0.34	-0.0094	0.50	0.34	0.89
		(0.074)	(0.018)	(0.031)	(0.089)	(0.018)
Argumentative	9	-0.49	-0.074	0.063	0.53	0.77
		(0.11)	(0.026)	(0.039)	(0.14)	(0.021)
Bullies	15	-0.42	0	0	0.44	0.81
		(0.10)			(0.12)	(0.030)
Disobedient	15	-0.50	-0.036	0.058	0.51	0.78
		(0.11)	(0.025)	(0.041)	(0.13)	(0.018)
Social	15	-0.21	-0.0087	0.11	0.15	0.97
		(0.045)	(0.017)	(0.045)	(0.045)	(0.013)

Standard errors in parentheses. The mental health dimension of child quality is indexed to the measurement of bullying at age 15, so for that measure the impact of academic quality and health on bullying is set to 0.

Table 20: Father quality measures

	False positive	True positive
Alcohol/drugs	0.63	0.95
	(0.014)	(0.0070)
Violence/abuse	0.70	0.98
	(0.012)	(0.0054)
Emotional abuse	0.42	0.78
	(0.016)	(0.012)
Incarceration/arrests	0.26	0.82
	(0.015)	(0.014)
Depression	0.46	0.75
	(0.015)	(0.012)
Probability(high quality)		
Constant	0.091	
	(0.071)	
Single index	0.40	
	(0.069)	

The top half of the table shows the measurement parameters. False positive is the probability the measurement reports a good outcome if the dad is low quality, and true positive reports the probability a measurement reports a good outcome if the dad is high quality. The bottom half shows the probability the dad is high quality. We use a logistic framework for the probability the dad is high quality, where the probability depends on the single index. Standard errors in parentheses.

Table 21: Counterfactual effects on child quality

	Academic		Physical health		Mental health	
Age	Baseline	Counterfactual	Baseline	Counterfactual	Baseline	Counterfactual
0	-0.00067	0.0074	-0.0049	-0.00067	0.0074	-0.0049
1	0.012	0.012	0.12	0.012	0.15	0.15
3	0.062	0.078	0.21	0.078	0.34	0.35
5	0.16	0.21	0.29	0.21	0.55	0.57
9	0.31	0.40	0.38	0.40	0.76	0.79
15	0.51	0.68	0.47	0.49	0.96	0.98

We increase the utility from transitioning into marriage by 0.4. Each cell reports the average simulated child quality.

Figure 1: Estimated single index values

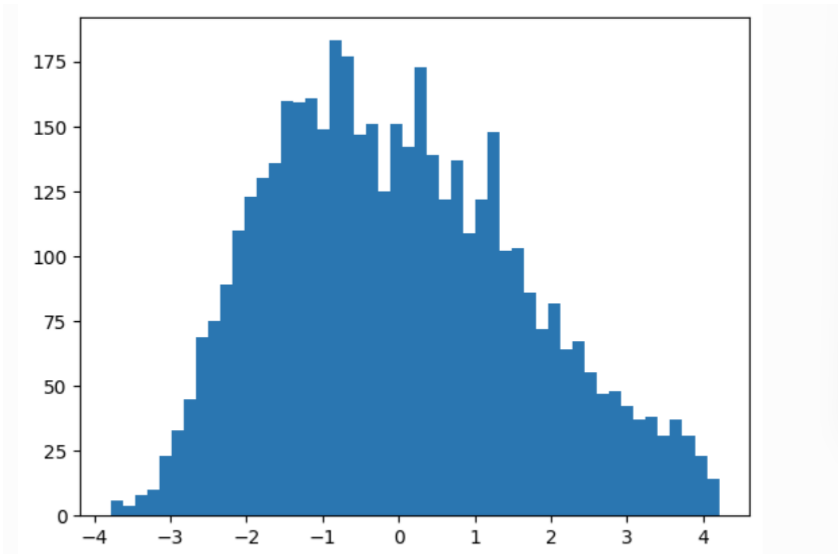

Figure 2: Reduction of variance (academic quality)

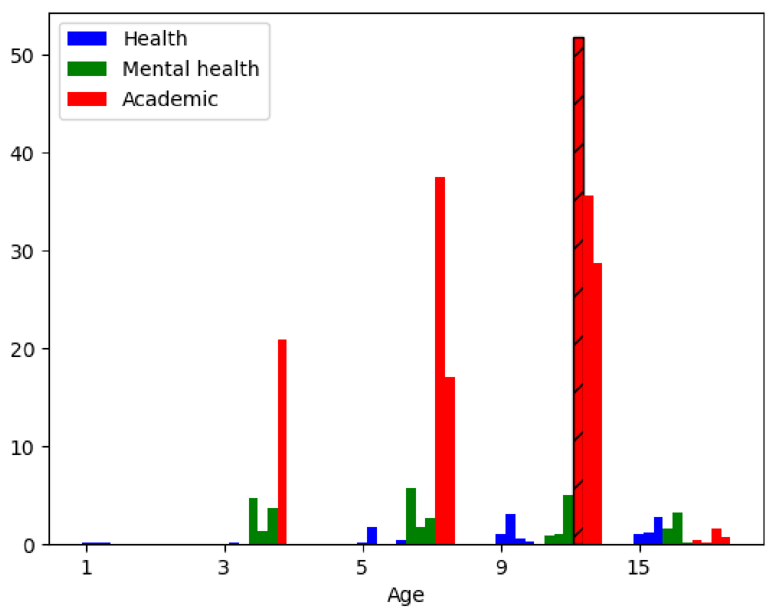

Figure 3: Reduction of variance (health quality)

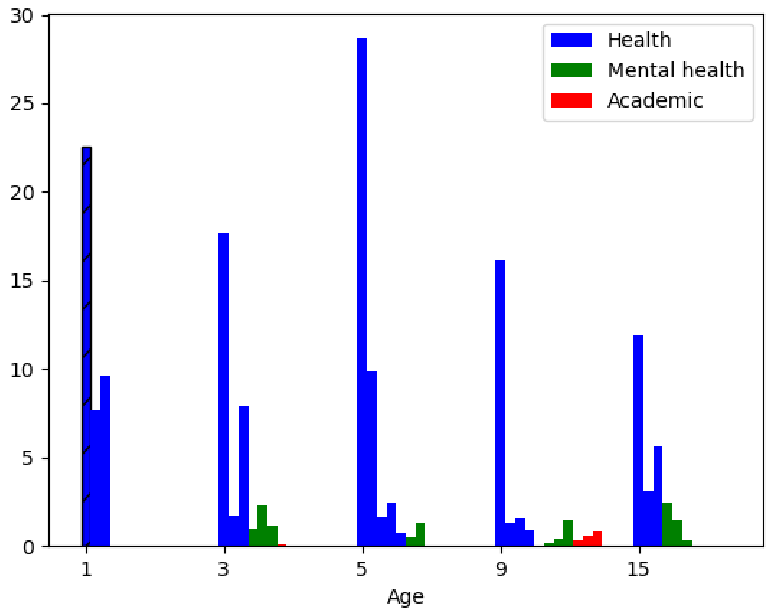

Figure 4: Reduction of variance (mental health quality)

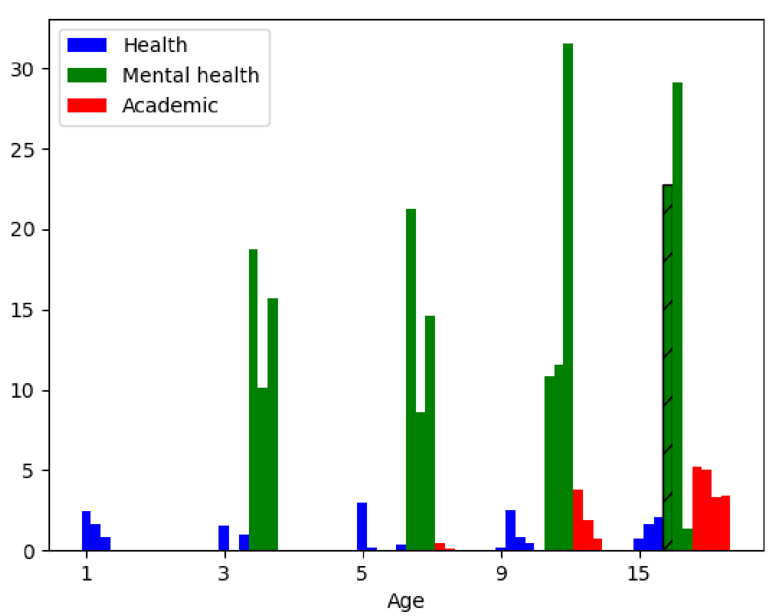

Figure 5: Posterior probability of high quality relationship

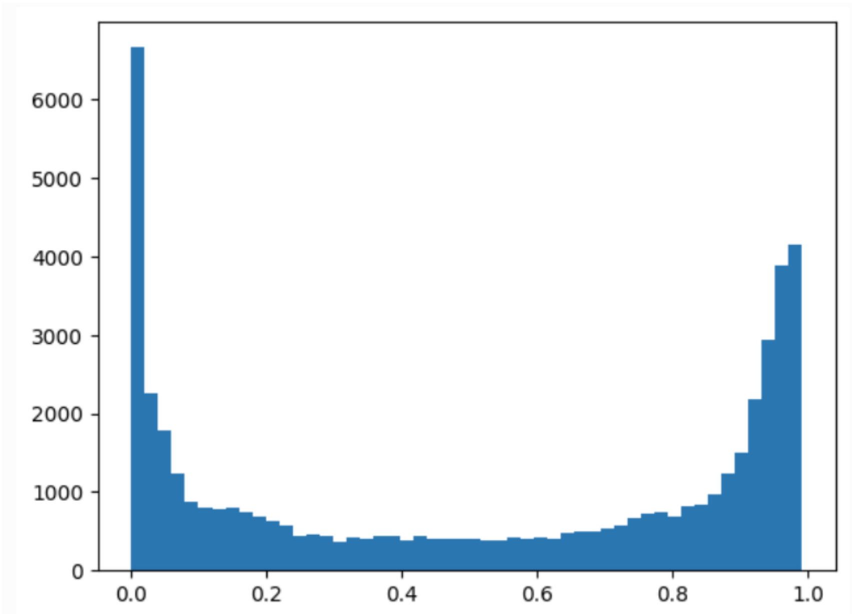

Figure 6: Increased utility from marriage: Distribution of impacts on academic quality

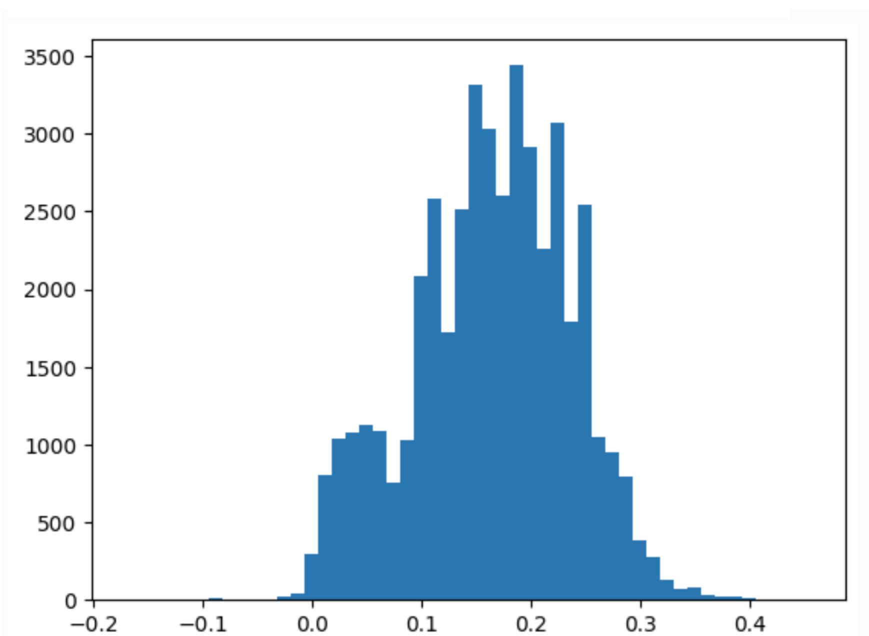

Figure 7: Increased utility from marriage: Distribution of impacts on health quality

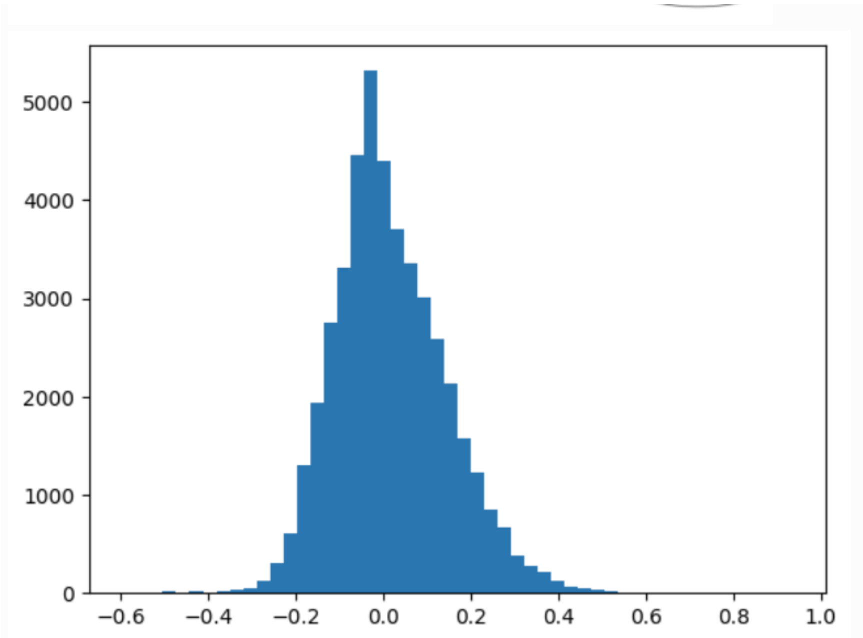

Figure 8: Increased utility from marriage: Distribution of impacts on mental health quality

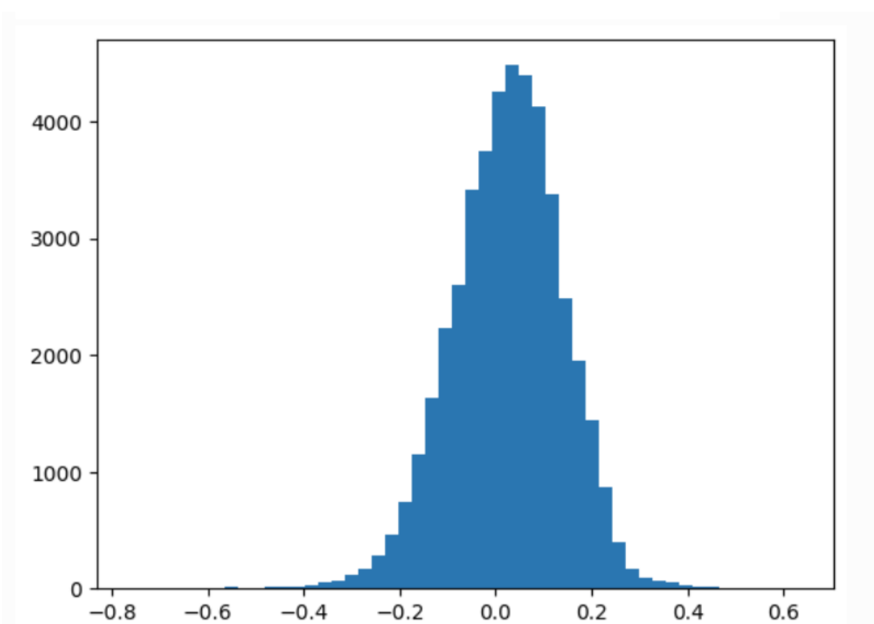

Figure 9: Increased utility from marriage: Heterogeneous impacts on academic quality

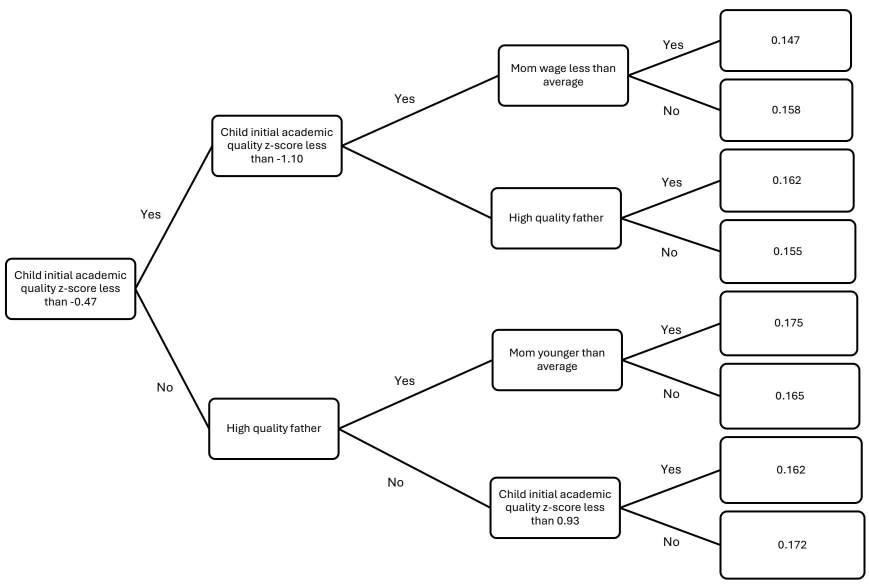

Figure 10: Increased utility from marriage: Heterogeneous impacts on health quality

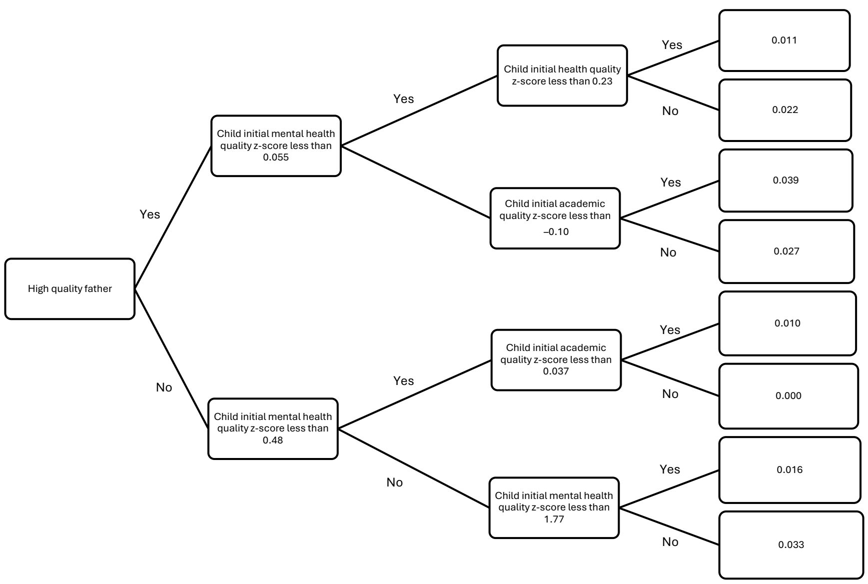

Figure 11: Increased utility from marriage: Heterogeneous impacts on mental health quality

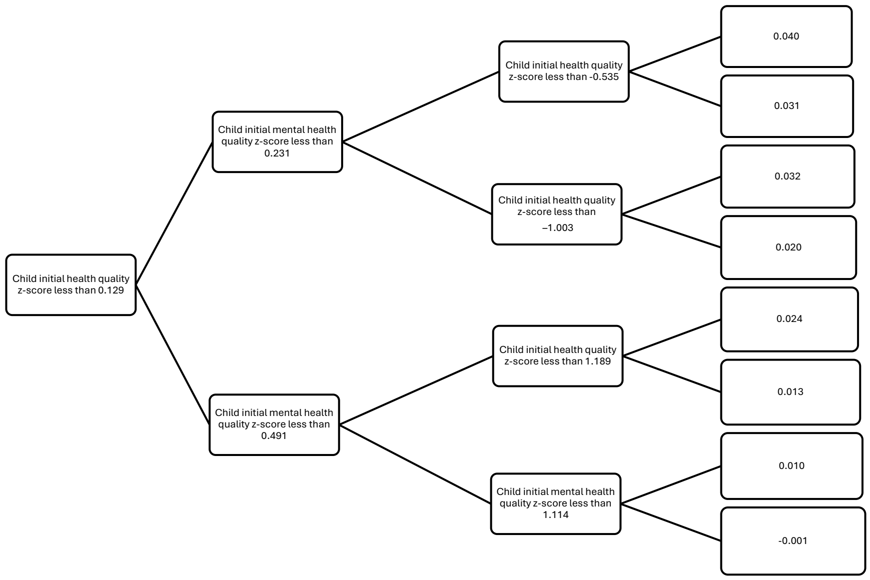

Figure 12: Increased initial academic child quality: Distribution of impacts on marriage rates

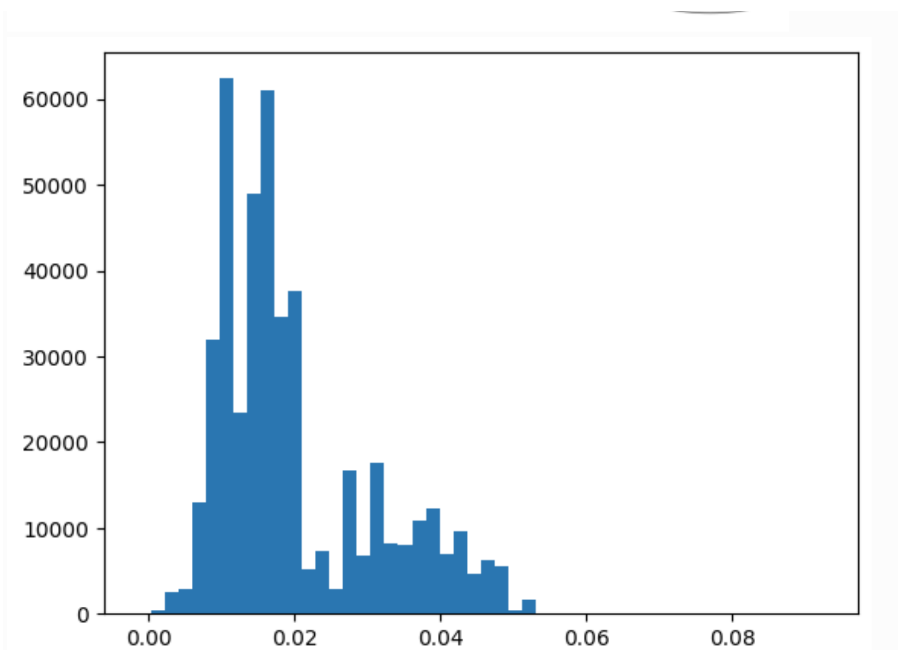

Figure 13: Increased initial health child quality: Distribution of impacts on marriage rates

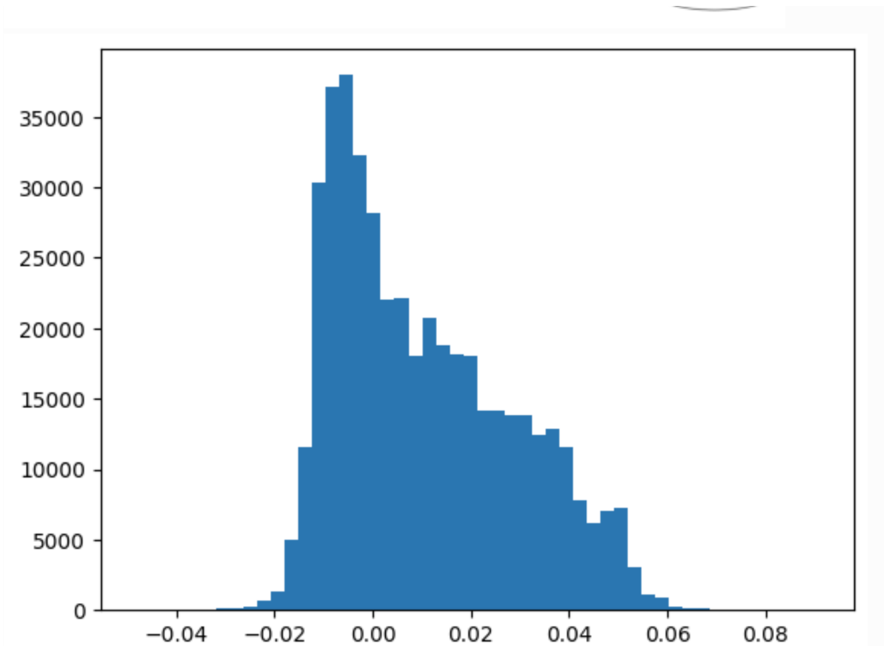

Figure 14: Increased initial mental health child quality: Distribution of impacts on marriage rates

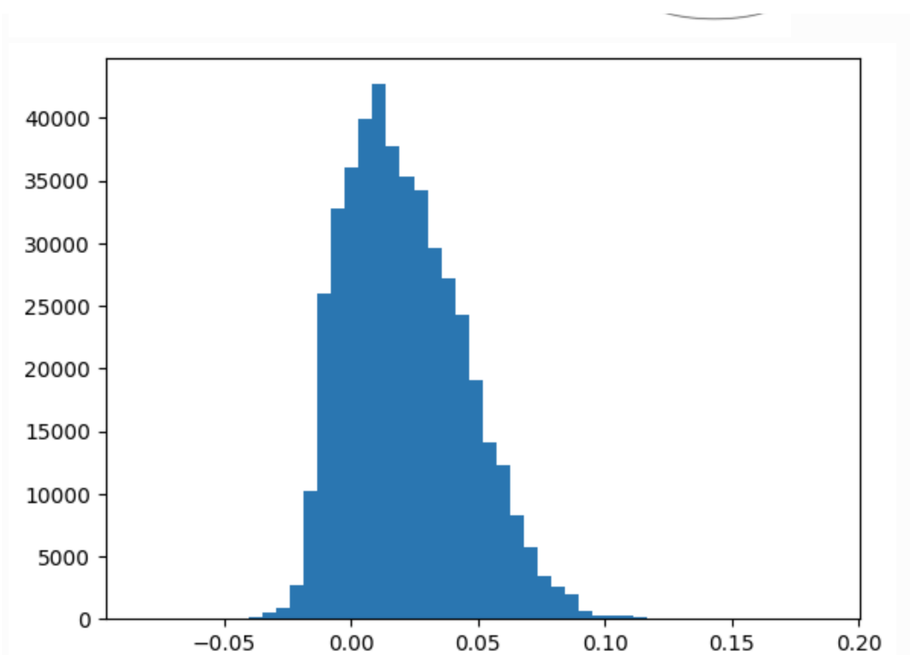

Figure 15: Increased initial academic quality: Heterogeneous impacts on marriage rates

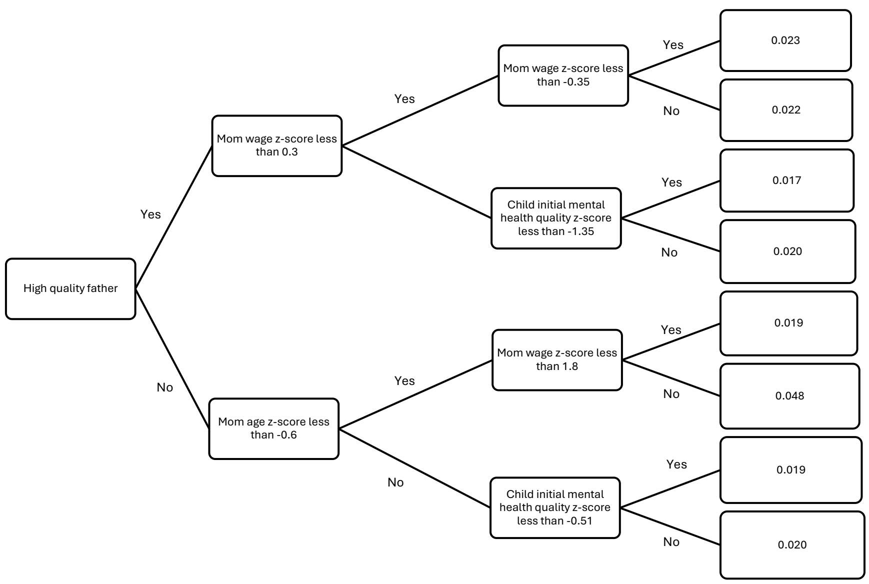

Figure 16: Increased initial health quality: Heterogeneous impacts on marriage rates

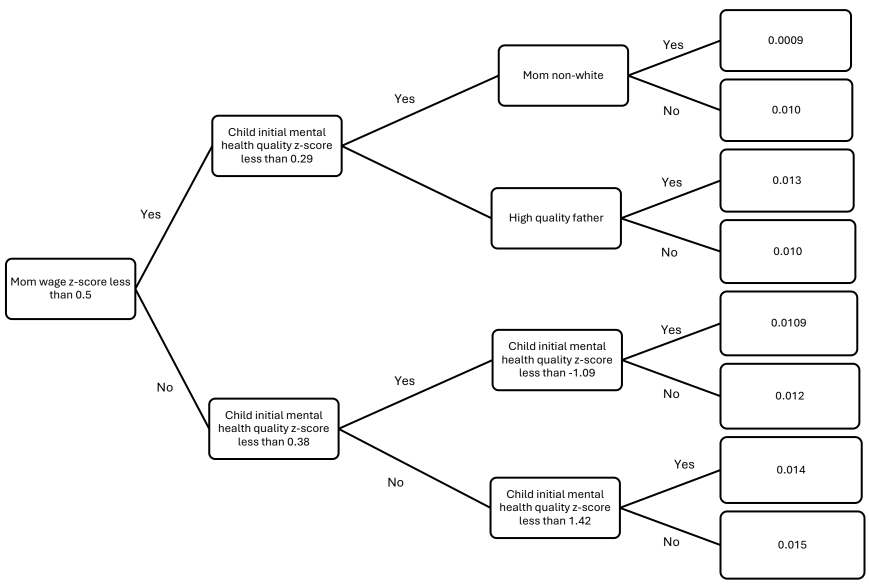

Figure 17: Increased initial mental health quality: Heterogeneous impacts on marriage rates

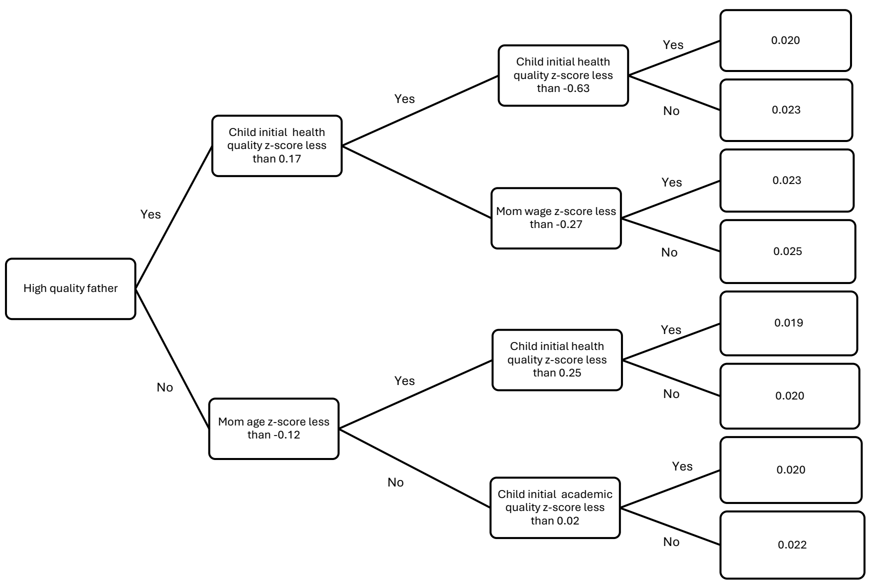

A Estimation Appendix

In this section, we derive the computational method used to implement our one-step Full-Information MLE estimation procedure.

The derived likelihood function in our Estimation section above is

$$\begin{aligned} L\left(s, M^c, M^r | X; \Theta\right) &= \int \Big(\xi \tilde{L}\left(s, m^c, M^r | q_0^c, X, q_r = H; \Theta\right) \\ &+ \left(1-\xi\right) \tilde{L}\left(s, m^c, M^r | q_0^c, X, q_r = K; \Theta\right)\Big) g\left(q_0^c | X\right) dq_0^c \end{aligned} \tag{13}$$

Moving the father quality measures to the outside, we can write this likelihood for one individual as

$$L_i = E_f[E_{q_0^c}[\prod_{t=1}^{T} \Pr(s_t | q_t, s_{t-1}, X, q^f) \prod_{m=1}^{M} g(m^c | q_t^c)]], \tag{14}$$

that is, the likelihood is the expected likelihood of the observed choices times the likelihoods of the observed child quality measures. Here, two assumptions allow us to dramatically speed up calculation of this integral. First, q_t^c is a deterministic linear function of only unobserved q_{t-1}^c and observed variables (given the unknown parameters), and second, the measurement error terms are normally distributed. Using the first assumption we can write $q_t = \Gamma_t q_0 + \Gamma_0$ for some 3×3 Γ_t and 3×1 Γ_0, where both Γ are functions of observables and parameters. We can then see that $g(m^c | q_t^c) = C \exp(-\frac{1}{2\sigma_{m_c}^2}(m_c - \alpha_0 - \alpha' q_t)^2) = C \exp(-\frac{1}{2\sigma_{m_c}^2}(m_c - (\alpha_0 + \Gamma_0) - (\alpha' \Gamma_t) q_0)^2)$.

The final term above, while written as a distribution in m_c, is proportional to some multivariate normal distribution in q_0, as it is quadratic with a negative leading term. We collect this distribution in with the initial distribution of q_0 in the expectation, writing

$$L_i = C_{\tilde{q}} \cdot E_f[E_{\tilde{q}_0}[\prod_{t=1}^{T} \Pr(s_t | \tilde{q}_t, s_{t-1}, X, q^f)]], \tag{15}$$

$$PDF(\tilde{q}_0) \propto \exp(-\sum_{m=1}^{M} \frac{1}{2\sigma_{m_c}^2}(m_c - (\alpha_0 + \Gamma_0) - (\alpha' \Gamma_t)\tilde{q}_0)^2) - \frac{1}{2}\tilde{q}_0' C_0^{-1} \tilde{q}_0). \tag{16}$$

with a known constant $C_{\tilde{q}}$. $\tilde{q}$ is then a multivariate normal in q_0 given the parameters and measures, with a mean and variance we derive computationally.

This can be viewed as an importance sampling approach, where the integral for each individual is not taken over the initial distribution of q_0, but rather over a distribution proportional to the posterior over q_0 given the observed measures and current parameter values.

Calculation of the multivariate PDF of $\tilde{q}$ and the associated constant is extremely fast at any given set of parameter estimations, since for any given value of q_0 we simply need to calculate out the implied path of q evolution given the actual choices and then calculate a quadratic function of this path and the observed measures. Evaluating this function at the set of $q_0 = \{(0,0,0),(0,0,1),(0,1,0),(1,0,0),(0,1,1),(1,0,1),(1,1,0),(1,1,1)\}$ uniquely pins down the 6 parameters plus constant of a 3-dimensional quadratic function.

This approach allows for a significant reduction in the number of evaluation points needed in a Gauss-Hermite integration scheme; see Evans and Swartz (2000). We implement the Gauss-Hermite method with 5 points per dimension.