The Effects of Family Structure on Children's Outcomes

Rebecca Lessem* and Carl Sanders†

June 29, 2026

Abstract

We study the effect of family structure on children's outcomes, using a sample of children born in low-income households in urban areas, where they are most at-risk for adverse outcomes. To do this, we create a structural dynamic model where relationship and labor market choices impact the child's academic, physical, and mental health outcomes. We estimate the model using measurements on children's outcomes and parent relationships drawn from the Fragile Families panel dataset. In our main counterfactual, we incentivize marriage, and find that this modestly increases academic outcomes but has little impact on physical or mental health. These effects are highly heterogeneous and children from advantaged backgrounds benefit most from the policy change.
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 documented that children born into low income households experience worse outcomes later in life (Heckman, 2008). A large body of research emphasizes the importance of early life human capital development for improving adult outcomes,[1] as well as the central role of parents in this process (Agostinelli and Wiswall (2025), Cunha et al. (2010), Boca et al. (2014), Boca et al. (2025)). In many low income households, parents are unmarried, which may limit children's ability to fully benefit from parental investments and interactions. Policymakers across the political spectrum have therefore advocated for programs that promote two-parent married households, based on the view that marriage provides the most advantageous family structure for children. However, despite this policy interest, the evidence on the causal effect of family structure on children's outcomes remains limited.[2]

In this paper, we estimate the effects of family structure on children's academic, physical health, and mental health outcomes. We do this using data from an urban low-income sample, which consists of many children born to unmarried parents; this high-risk population is of particular policy interest. Our primary counterfactual examines the subsidization of marriage or cohabitation: can financial incentives that encourage mothers to remain with the children's father improve child outcomes? Policies that implicitly incentivize or discourage marriage already exist - for example, the U.S. income tax system creates marriage penalties or bonuses depending on parents' relative earnings. Inducing marriage can move children into more stable environments with more resources, improving children's outcomes. On the other hand, the impacts of the policy depend on the marginal marriages, and it is unclear how the children in these households would be impacted by married parents. For example, if the marginal marriages have fathers who are abusive or have substance abuse problems, which is common in our sample, the child's outcomes may not improve in the counterfactual.

A large literature has examined how divorce or parental death affects children's outcomes. Our work builds on this research by considering a broader set of family structures, including cases in which parents are cohabiting, never married, or subsequently re-partner. Given the high incidence of nonmarital births, particularly in urban and lower income communities, understanding how family structure shapes child development outside the context of divorce is crucial for addressing inequality in outcomes. Tartari (2015) finds that children would have performed better academically if their parents had remained together rather than divorcing, while Brown et al. (2025) finds effects of divorce with a focus on endogenous parental time

[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 examining the relationship between child outcomes and parents' marital status.

investments. In contrast, Lang and Zagorsky (2001) finds that early parental death has relatively small effects on later life economic outcomes. Most recently, Bang and Borghesan (2026) estimate a dynamic econometric model relating cognitive skill development to marriage decisions of parents.

In addition, whereas most prior work focuses primarily on academic achievement, we examine a broader set of outcomes, including mental and physical health. This extension is important in light of evidence from Heckman et al. (2006) highlighting the role of noncognitive skills, such as motivation and persistence, in shaping educational attainment and labor market success. Finally, our framework explicitly incorporates impacts of father well-being on children, allowing fathers who have problems such as substance abuse issues to have differential effects on their offspring. Allowing for this heterogeneity is important in understanding the quality of marriages that are induced by our counterfactuals.

To estimate the effects of family structure on children, we develop and estimate a dynamic model in which mothers make decisions about both relationship status and labor market participation,[3] and these choices shape the trajectory of their child's development. We summarize child outcomes using a three-dimensional vector capturing academic achievement, physical health, and mental health. Both decisions and outcomes depend on the well-being of the father. This channel is important for two reasons: father well-being influences the likelihood that a relationship forms and persists, thereby affecting the child's environment, and it also has a direct impact on the child's development. We estimate the model using data from the Fragile Families and Child Well-Being Study, which follows an urban cohort of mothers and children from birth through age 22. This longitudinal structure allows us to track the evolution of children's academic, physical health, and mental health outcomes, alongside detailed information on parental relationship status over time.

Estimation of the model is technically complex because child outcomes and father well-being 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. This approach is similar to Hwang (2024), while we use a full-information approach instead of estimating using conditional choice probabilities.

Our first counterfactual evaluates a policy that provides sizable cash transfers conditional

[3] Bernal (2008), Bernal and Keane (2011) highlight the importance of accounting for labor market decisions in analysis of children's outcomes.

on parents being married. We find that such incentives generate modest improvements in children's academic outcomes and have little effect on average physical or mental health outcomes. From a policy perspective, this suggests that marriage subsidies alone may not be an effective tool for broadly improving child well-being across all domains. The average effects mask substantial heterogeneity, with a substantial share of the sample having large treatment effects from this policy in each of the dimensions of child outcomes. Using machine learning methods to uncover differential impacts across subgroups, we find that children with more favorable initial conditions experience gains in both academic and health outcomes. In contrast, children with poorer initial conditions see the largest improvements in mental health. These findings highlight that the effectiveness of marriage-based incentives depends critically on the target population, suggesting that more narrowly targeted policies may yield larger benefits than universal subsidies.

In a second counterfactual, we examine the reverse relationship: how do improvements in child outcomes affect parental relationship stability? This counterfactual highlights the "stay together for the kids" motive potentially facing parents: if a child is on the margin of being helped by a current marriage or lack of marriage, this could induce changes in parents' choices. We find that exogenously increasing children's academic or physical health leads to a 1-2 percentage point reduction in the share of parents who are married or cohabiting in the terminal period; that is, if parents didn't need to worry about their children's academic or physical outcomes, they would be more likely to separate. We find the reverse effect for mental health, which is that some separations are driven by low mental health outcomes of children: initial exogenous improvements in children's mental health increase the share of parents who are married or cohabiting by approximately 1 percentage point.

Our final counterfactual speaks to the effects of marriage versus cohabitation on children. Cohabitation is less stable than marriage, and we examine how this instability impacts child development. Our thought experiment makes cohabitation as "stable" as marriage, leaving all other parameters the same, but increasing the costs of separation from cohabiting with the father. We find that this counterfactual, despite leaving the marriage parameters unchanged, actually increases marriage rates, as cohabiting mothers move into marriage in the long run rather than their baseline separation from the child's father. Child quality is also improved, with effects from this counterfactual being about 40% of the effect sizes of the marriage incentivization counterfactual. This result is indicative of the fact that it is not cohabitation itself that is bad for child development; rather, the instability of the state means an increased future probability of separation from the father.

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. (2020a), Attanasio et al. (2020b), 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), and Jackson (2018). We might be concerned that mothers' decisions to separate from 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 nutrition, 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.

Most of the past work on this topic studies the impact of divorce on children's outcomes; we extend this by also considering a broader set of relationship choices: married or cohabiting with the child's father or a new partner, or being single. Cohabitation is increasingly observed in the data, but we have a limited understanding on how it affects children's outcomes. Brien et al. (2006) 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.

Child and father well-being are both unobserved, and we treat them as latent variables in the model. We use a measurement model approach to use measures in the data on both the child and father to learn about the underlying latent variables. 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 controls 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

We use the Fragile Family and Child Well-Being Study (hereafter Fragile Families), a longitudinal survey of parents and children. The survey covers families in 20 urban areas and is representative of urban populations, though not of the United States as a whole. Given our focus on how marital status at birth affects subsequent outcomes, this restriction is not limiting, as the sample provides strong coverage of the population most at risk.

Parents were surveyed at the time of the child's birth and answered a wide range of questions on relationship status, household structure, and labor market activity. They were then re-surveyed at regular intervals until the child is 22 years old. At each wave, we observe detailed information on the child's physical health. We also observe academic outcomes, including standardized assessments of cognitive development administered as part of the survey and, at later ages, school grades. In addition, both parents and children respond to questions designed to measure psychological well-being. Surveys were conducted at birth (year 0) and again when the child was ages 1, 3, 5, 9, 15 and 22.

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

We show the share of the sample with each characteristic. High education defined as some college or more.

Table 1 presents basic descriptive statistics for the key demographic variables used in our analysis. We begin with maternal education, dividing the sample into two groups: low education (a high school degree or less) and high education (at least some college). Approximately 60 percent of mothers fall into the high education category. We then examine the racial composition of the sample. Roughly one-third of respondents are white, about half are black, and the remainder identify with other racial groups. Finally, the average age of mothers at the time of childbirth is just over 25.

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	17.9%	6.5%	25.5%	10.6%			25.9%	13.7%
1	16.2%	13.7%	14.2%	13.1%	2.4%	2.5%	20.1%	17.9%
3	17.7%	14.3%	10.5%	8.8%	5.2%	4.6%	22.8%	16.1%
5	18.0%	12.9%	7.4%	5.4%	8.9%	7.2%	24.8%	15.4%
9	19.2%	10.0%	5.6%	3.7%	12.5%	8.8%	25.0%	15.3%
15	19.1%	6.6%	3.3%	1.5%	18.1%	7.1%	30.1%	14.1%

We show the share of the sample with each relationship and labor market combination at each age. The data does not report any information on new partners at age 0.

Crucially for our analysis, we observe the mother's relationship status in every survey wave. We see whether she is married to the child's father, cohabiting with the child's father, married or cohabiting with a new partner,[4] or single.[5] A quick look at the data underscores the importance of examining not only divorce but also family formation. Roughly 15 percent of parents in our sample divorce during the observation period. In contrast, nearly 60 percent are never married at any of the survey waves. Thus, to understand how family structure affects child development in this context, it is essential to consider both the formation and dissolution of partnerships. Table 2 provides more detail, showing the distribution of relationship status at each survey wave.[6] We also split the sample by the mother's labor market status. About 75 percent of children in our sample are born to unmarried parents. Cohabitation is common in early childhood, but rates decline as children age, with parents who start in cohabitation either marrying or separating. In contrast, we do not observe strong or systematic patterns in maternal employment across relationship states.

Table 3: Relationship transitions

		Current relationship status			
	Prior relationship status	Married to father	Cohabiting with father	New partner	Single
Age 0 to 1	Married to father	94.2%	1.1%	0.9%	3.9%
	Cohabiting with father	14.3%	55.4%	3.0%	27.3%
	Single	4.0%	17.8%	9.1%	69.1%
Age 1 to 3	Married to father	90.3%	1.5%	1.7%	6.6%
	Cohabiting with father	14.0%	53.7%	4.3%	28.0%
	New partner	2.2%	2.8%	55.0%	40.0%
	Single	3.3%	11.2%	14.6%	71.0%
Age 3 to 5	Married to father	86.5%	1.9%	2.4%	9.1%
	Cohabiting with father	14.9%	49.8%	5.7%	29.6%
	New partner	0.0%	0.8%	64.0%	35.2%
	Single	2.4%	6.4%	19.4%	71.8%
Age 5 to 9	Married to father	82.1%	1.7%	4.5%	11.7%
	Cohabiting with father	19.6%	46.5%	9.7%	24.3%
	New partner	1.2%	1.2%	56.0%	42.6%
	Single	2.7%	6.4%	24.7%	66.2%
Age 9 to 15	Married to father	78.2%	1.1%	4.3%	16.5%
	Cohabiting with father	14.8%	45.8%	7.8%	32.7%
	New partner	0.7%	0.2%	61.4%	37.7%
	Single	1.9%	1.8%	26.9%	69.4%

In each row, we show the transition between relationship states at a given age.

Table 3 reports the distribution of relationship status in each survey wave, conditional on the relationship status in the prior period. The data show substantial movement across

[4] We combine these categories due to sample size limitations.

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

[6] At age 0, we observe only the mother's relationship with the child's birth father. In all subsequent waves, we also observe her relationship status with other partners.

states over time. For example, among mothers who are single at the child's birth, nearly 20 percent are cohabiting with the birth father by age 1. Among those who are cohabiting at birth, 15 percent are married one year later. Transitions continue throughout childhood, indicating that many children experience changes in household structure over their first 15 years. Marriage and cohabitation are both relatively persistent states; however, married parents are considerably less likely to become single than cohabiting parents. Fewer than 5 percent of couples who are married at birth are single by age 1, compared with nearly 30 percent of those who are cohabiting at birth. These observed patterns motivate our modeling choices. Because movement between relationship states is common, a dynamic framework is essential for capturing how the sequence of family structure transitions shapes children's outcomes. Moreover, in contrast to much of the prior literature that does not distinguish cohabitation from marriage, the data show that cohabitation is both prevalent and distinct. Including this state is therefore important for understanding how different family structures influence child development.

2.1 Child outcomes

In our model, family structure impacts child outcomes. The data provide a rich set of measures on children, spanning academic performance, physical health, and mental health. In this section, we describe the specific measures we use to capture these dimensions of child well-being.

Table 11 in Appendix A shows summary statistics on the academic outcome measures used in our analysis, split by parental relationship status. As part of the survey, children completed several standardized assessments. In particular, we use the Peabody Picture Vocabulary Test (PPVT) at ages 3, 5, and 9, as well as the Woodcock-Johnson (WJ) vocabulary tests at ages 5 and 9. At age 15, children self report their math, English, science, and history grades in high school. Across all measures, mean outcomes tend to be higher for children whose parents live together, although the differences are not always statistically significant, and these comparisons reflect unconditional means. We additionally have information on wages and education at age 22, which we will use as an additional measure of the child's outcomes. The structural model will allow us to assess how parental relationship decisions shape these academic outcomes over time.

Tables 12 and 13 in Appendix A present our physical health measures and corresponding summary statistics. At each survey wave, we observe a self-reported health rating on a 1-5 scale. We also use additional indicators of physical health, including the presence of physical disabilities, whether the child is overweight, and other reported medical conditions. As with the academic outcomes, children living with both parents generally exhibit better physical health indicators, although the differences vary across measures.

We also observe a range of mental health outcomes, including measures of depression, behavior problems, and social skills. Table 14 in Appendix A lists the specific measures used to capture mental health, along with summary statistics. As with the academic and physical health measures, unconditional outcomes tend to look better for children whose parents live together. The structural model will allow us to assess how family structure contributes to these differences.

2.2 Paternal well-being

A unique feature of the Fragile Families data is that it includes detailed information on fathers that allows us to learn about their well-being. This is particularly valuable in our context, as the father's well-being affects both the likelihood that the relationship persists and the child's outcomes. Across survey waves, we observe whether the father abuses drugs or alcohol, engages in emotional or physical abuse, or has jail time or criminal convictions. We use these indicators to learn about the father's well being in the estimation of our model.[7] Table 4 reports summary statistics for the measures used to learn about paternal well-being; we observe relatively high frequencies of several adverse characteristics.[8]

Table 4: Measures of paternal well being

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

For each measure, we look across all waves of the survey to see if there is any indication of the father having engaged in a given action. If so, the measure is a 1, otherwise it is coded as a 0.

[7] Comparable information is available for mothers. However, because maternal and paternal quality are highly correlated, we rely only on the measures of father quality.

[8] To simplify computation, we treat these measures as time-invariant. A father is coded as having a given adverse characteristic if he reports it in any survey wave.

3 Model

We model the mother's relationship and labor market decisions using a dynamic discrete choice framework.[9] In each period, she chooses both her relationship status and employment status to maximize expected lifetime utility. Relationship and labor market choices directly impact utility, and also impact the evolution of the child's quality.

Choices In each period t, the mother selects a relationship status

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

where $Married_f$ and $Cohab_f$ denote being married to or cohabiting with the child's father, MC_n denotes being married to or cohabiting with a new partner,[10] and $Single$ indicates that she is not living with any partner. She also chooses an employment status

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

where working corresponds to full time employment.[11] We denote the joint choice as

$$s_t = (rel_t, emp_t).$$

Initial conditions The first initial condition is the child's quality at birth, q_0^c, a three-dimensional vector capturing academic, physical, and mental health dimensions. The second is the father's well-being $fw \in \{0, 1\}$, which is binary and time invariant. Both the father and any potential new partner have income types $\rho = (\rho^f, \rho^n)$ indicating whether the father f and a new partner n is a high or low income earner. We also observe time invariant demographics X and the mother's relationship and employment status at the child's birth, s_0.

State variables At the start of period t, the mother observes the state vector $(q_t^c, s_{t-1}, X, fw, \rho)$. Child quality q_t^c evolves endogenously with her choices.

[9] Our model allows the mothers to make all decisions, abstracting from a household bargaining model of family. We believe that our specification approximates the true model, and the counterfactual outcomes would be similar if we allowed for father's to be part of the bargaining framework. On the data side, the father is not surveyed or tracked unless he is currently with the mother.

[10] We combine marriage and cohabitation with a new partner due to their low frequency in the data.

[11] The data do not allow us to distinguish full from part time work, so we abstract from this distinction.

Utility function Period utility depends on: (1) the child's current quality q_t^c, (2) the mother's current choice and prior relationship/employment status (s_t, s_{t-1}), (3) demographic characteristics X, (4) the father's well-being fw, and (5) male income types ρ. We write period utility as $U\left(q_t^c, s_t, s_{t-1}, X, fw, \rho\right)$. This structure allows for different utilities across relationship states, transition costs (via dependence on s_{t-1}), a disutility of work, income effects through ρ, and preferences for being with a high well-being father.

Timing At the start of each period, the mother observes the full state vector, including the realized child quality q_t^c. She then chooses her relationship and labor market status for this period s_t, and then utility for period t is realized. Between periods, child quality evolves to q_{t+1} according to a child quality transition function.

Value function Let η_{jt} denote i.i.d. extreme value type I shocks associated with each choice $j \in J_R \times J_W$. The value function is

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

where the deterministic component is

$$\begin{aligned} v_t\left(j, q_t^c, s_{t-1}, X, fw, \rho\right) &= U\left(q_t^c, j, s_{t-1}, X, fw, \rho\right) \\ &+ \beta E\left[V_{t+1}\left(q_{t+1}^c, j, X, fw, \rho, \eta_{t+1}\right)\right], \end{aligned} \tag{2}$$

and child quality evolves deterministically according to

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

This specification allows for state dependence in each dimension of quality. Relationship status and maternal employment decisions can impact the growth of child quality. The father's well being affects the child's quality, and we allow for direct income effects for the father or other partner through male income types ρ.

Using the log-sum formula for extreme value shocks,

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

where γ is Euler's constant. The probability of choosing s_t is the standard multinomial logit expression:

$$Pr^s\left(s_t|s_{t-1}, X, q_t^c, fw, \rho\right) = \frac{\exp\left(v_t\left(s_t, q_t^c, s_{t-1}, X, fw, \rho\right)\right)}{\sum_{j\in J_R\times J_W}\exp\left(v_t\left(j, q_t^c, s_{t-1}, X, fw, \rho\right)\right)}. \tag{5}$$

Terminal Period The model is solved from the child's birth through age 15. After age 15, child quality is held fixed, but the mother continues to solve the relationship choice problem for 35 additional periods. This ensures that continuation values at the age 15 boundary properly reflect the mother's anticipated lifetime relationship and work choices. The terminal value in the final period is zero.

4 Measurement model for latent variables

The model is written conditional on *child quality* q_t^c, *paternal well-being* fw, and *male income types* ρ, all of which are unobserved by the econometrician. Instead, these latent variables are measured with error in the data. For the child, we observe a set of academic, physical health, and mental health outcomes, which we use to infer the three dimensional quality vector q_t^c. For the father, we observe multiple indicators of his behavior, which inform our estimate of father well-being fw. Finally, we observe the income of both the father and other partners of the mother, which provides information about the income types of the father and new partners ρ.

Child quality We use a large set of child quality measures in each period, as detailed in Section 2.1. Some measures are available in every wave, while many others are specific to particular survey rounds. Our specification accommodates this unbalanced structure by allowing the number and type of observed outcomes to vary freely across periods. In addition, we do not require that similarly named tests administered in different waves measure the latent quality vector with the same parameters.

Denote M^c as the set of measures of child quality. For any measure m_{jt} observed at time t, we specify

$$m_{jt} = \alpha_{1jt} + \alpha_{2jt}q_{1t}^c + \alpha_{3jt}q_{2t}^c + \alpha_{4jt}q_{3t}^c + \alpha_{5jt}\varepsilon_{jt}. \tag{6}$$

Each outcome m_{jt} loads on all three dimensions of child quality. The coefficients α_{2jt}, α_{3jt}, and α_{4jt} capture how each latent dimension contributes to outcome j. The measurement errors ε_{jt} are assumed to be standard normal, implying that the standard deviation of outcome j is α_{5jt}. Because the latent quality vector is three-dimensional, corresponding in practice to academic skills, physical health, and mental health, we impose normalizations for identification. We anchor one observed measure to each latent dimension. Academic quality is anchored to the PPVT test at age 9; physical health is anchored to self-reported health at age 1; and mental health is anchored to measure of bullying at age 15. For each anchored measure, we set the loadings on the other two dimensions to zero.[12]

Child quality evolves deterministically as a function of relationship status, labor market status, demographics, and prior quality. Therefore, given the model parameters, the mother's relationship and labor market decisions, and the initial value of q_0^c, we can predict the child's latent quality in every period. This allows us to construct predicted values for each observed quality measure, which we denote by $\widehat{m}_{jt}$:

$$\widehat{m}_{jt} = \alpha_{1jt} + \alpha_{2jt} q_{1t}^c + \alpha_{3jt} q_{2t}^c + \alpha_{4jt} q_{3t}^c.$$

The standard deviation of each measure is given by α_{5jt}. Because the measurement errors ε_{jt} are normally distributed, we can compute the likelihood contribution of each observed outcome directly, which will be used in the full likelihood function described in Section 5.2.

Father well-being We assume that father well-being fw is discrete, taking the value 0 for poor well-being and 1 for high well-being. This specification is motivated by the nature of the available data: most father well-being indicators are binary yes/no questions, making a continuous normality assumption inappropriate. We also assume that father well-being is time-invariant. Let M^f denote the set of father well-being measures. For each measure $m_k \in M^f$, we estimate two parameters:

$$\Pi_k = \{\pi_{1k} \equiv \Pr(m_{ik} = 1 | fw = 0), \pi_{2k} \equiv \Pr(m_{ik} = 1 | fw = 1)\}. \tag{7}$$

Here, π_{1k} is the probability that measure k indicates a high well-being father (equals 1) when the true father well-being is poor - a "false positive" rate. Conversely, π_{2k} is the "true positive" rate: the probability that the measure equals 1 when the father has high well-being. The corresponding false negative and true negative rates follow directly as one minus these probabilities. For an informative measure of father quality, we expect $\pi_{1k} << \pi_{2k}$, indicating that a value of 1 is much more likely when the underlying well-being is good. These measurement parameters enter the likelihood function described in Section 5.2.

[12] These normalizations affect only interpretation. Because the remainder of the model is linear in parameters given the latent variables (although not unconditionally), any alternative set of normalizations would simply rotate the latent space and yield an observationally equivalent distribution of outcomes.

Male income types The last latent variable is the income type of the father and possible future partner, ρ^f and ρ^n. Both of these take on a fixed number of discrete values, representing the father or partner's earning potential.[13] We observe the father's earnings Y_{it}^f and the new partner's earnings Y_{it}^n (in the cases the mother ever marries or cohabits with a new partner), and assume it depends on the income types as follows:

$$\begin{aligned} Y_{it}^f &= \gamma_0 + \gamma_1 \rho^f + \gamma_2 \varepsilon_{it}^f \\ Y_{it}^n &= \gamma_0 + \gamma_1 \rho^n + \gamma_2 \varepsilon_{it}^n. \end{aligned} \tag{8}$$

The term γ_1 gives the relationship between the latent income type and earnings, and γ_2 gives the variance of earnings. We assume the same γ terms for the father and partner to allow for larger sample sizes when estimating these parameters, which is especially important because we have limited observations of earnings for new partners. As for the measurements of child quality, we assume the ε terms follow the normal distribution, meaning the likelihood of a realized income, conditional on type, comes from the normal pdf.

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 $\beta = 0.95$.[14]

[13] In implementation, they can both take on 2 possible values.

[14] 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.

5.1 Parameterization

Single index: To simplify computation and increase 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, and 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.[15] 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.

Using a single index to reduce the number of parameters is necessary for estimating the model. To see why, consider the child quality evolution functions, which depend on demographic variables. As described in the parameterization section below, we estimate a separate quality transition function for each combination of relationship status and labor market status (8 in total). Since child quality has three dimensions, this requires estimating 24 functions. If we were to include education directly (even as a binary variable), we would need to add 24 additional parameters. The same issue would arise for every other characteristic included in the index. Given the size of our dataset, estimating such a large number of parameters is not feasible. The single index approach allows us to capture heterogeneity in the data while avoiding these data limitations.

Total family resources We define total family resources (TFR_t), which is a measure of the resources available to the household. We parameterize this as follows

$$TFR_t\left(s_t, \rho\right) = SI\left(X\right) \mathbb{1}\left(s_t \in \{1,3,5,7\}\right) + \rho^f \mathbb{1}\left(s_t \leq 4\right) + \rho^n \mathbb{1}\left(s_t \in \{5,6\}\right).$$

Total family resources depend on whether or not the mother is working, which we interact with the single index to allow for different effects of labor market participation depending on the characteristics of the mother. When the mother is married or cohabiting with the father or a new partner, the men in the household also impact the total family resources. The impact of the men on total family resources depends on their income types.

Total family resources will enter both the utility function and the child quality evolution

[15] 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.

functions. This allows for increased utility with higher income, as well as income effects on the child quality evolution function.

Utility function We parameterize the utility function as

$$\begin{aligned}
U\left(q_t^c, s_t, s_{t-1}, X, fw, \rho\right) &= \tilde{u}\left(s_t, s_{t-1}\right) + \lambda_1 fw \mathbb{1}\left(s_t \leq 4\right) + \lambda_2 SI(X) \mathbb{1}\left(s_t \geq 6\right) \\
&+ \lambda_3 \mathbb{1}\left(s_t \in \{2,4,6,8\}\right) + \lambda_4 SI\left(X\right) \mathbb{1}\left(s_t \in \{2,4,6,8\}\right) \\
&- \lambda_5 \mathbb{1}\left(s_t \in \{1,3,5,7\}\right) \mathbb{1}\left(s_{t-1} \in \{2,4,6,8\}\right) \\
&+ \kappa' q_t^c + TFR_t\left(s_t, \rho\right) \qquad (9)
\end{aligned}$$

The first term $\tilde{u}(\cdot)$ 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 with a father with high well-being. 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, which allows for demographic characteristics to affect the utility from leisure. The term λ_5 gives the cost of transitioning from out of the labor market into the labor market. The κ terms give the relationship between each dimension of child quality and utility. Lastly, utility depends on total family resources. This allows for an income effect of mother's labor market choices, since choosing to work will increase total family resources and therefore increase utility. It additionally impacts relationship choices, as living with a high income type partner increases family resources.

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 function:

$$q_{t+1}^c = \delta_{1s} + \delta_{2s} q_t^c + SI(X)\delta_{3s} + \delta_{4s} \mathbb{1}\left(s_t \geq 5\right) \mathbb{1}\left(s_{t-1} \leq 4\right) + \delta_{5s} fw + \delta_{6s} TFR_t. \qquad (10)$$

Recall that q_{t+1}^c is a 3×1 matrix. Each of the δ parameters (excluding δ_{2s}) in equation (10) are also 3×1, allowing for an impact for each dimension of quality. The first component of quality in the next period is an intercept term δ_{1s}. Next, child quality in the next period depends on the child quality in all dimensions in the prior period through the 3×3 matrix δ_{2s}. This allows for current academic quality to affect the evolution of mental health, for example. For components of $\delta_{2s} > 0$, 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 array 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 are separated.[16] The term δ_{5s} gives the relationship between child quality growth and father well-being; this allows for fathers with high well-being to differentially impact their child's outcomes. The last term δ_{6s} gives the relationship between total family resources and child quality. This gives a mechanism for mother's labor market decisions to have an income effect on child quality: if a mother works, the household has more resources which can impact the child's quality evolution. Additionally, this allows for the male income types to impact child quality, as living with a high income partner gives the household more resources.

Initial child quality Initial child quality is drawn from the normal distribution with a mean that is a function of the single index term, a 3×1 parameter τ that gives the loading of the single index on the initial mean of that group, and a 3×3 variance-covariance matrix C_0.

$$q_0 \sim N\left(\tau \cdot SI\left(X\right), C_0\right)$$

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

5.2 Likelihood

The likelihood has 4 components: the choice probabilities, the measures of child quality, the measures of father well-being, and the measures of male income types.

Denote the parameters as Θ. Denote $\mathbb{S}_i$ as the series of all relationship and labor market choices over all periods for household i, M_i^c as the full set of child quality measures, M_i^f as the set of father well-being measures, and Y_i as the set of income measures for the men. We start by deriving the likelihood over all periods for a given household, conditional on initial q_0 and a given value of fw and ρ,

[16]This includes exiting marriage or cohabitation and moving to a state where the child's parents do not live in the same home.

$$\tilde{L}\left(\mathbb{S}_i, M_i^c, M_i^f, Y_i | q_0^c, X_i, fw, \rho; \Theta\right) =$$
$$\prod_{t=1}^{T} Pr^s\left(s_{it}|q_t^c, s_{i,t-1}, X_i, fw, \rho\right) \times \prod_{m_i^c \in M_i^c} l^c\left(m_i^c|q_t^c\right) \times \prod_{m_i^f \in M_i^f} Pr^f\left(m_i^f|fw\right) \times \prod_{y_i \in Y_i} l^y\left(y_i|\rho\right) \tag{11}$$

The first component is the choice probabilities $Pr^s\left(s_{it}|s_{i,t-1}, q_t^c, X, fw, \rho\right)$, which have a logit form and are defined in equation (5). For measure m_{it}^c of child quality observed at time t, we write the density function as $l^c\left(m_{it}^c|q_t^c\right)$. This comes from the normal distribution as explained in Section 4. The next component of the likelihood function is the probability of each paternal well-being measure, which is derived in equation (7) and written as $Pr^f\left(m_i^f|fw\right)$. The last component is the likelihood of each realized income for the fathers and new partners of the mother $l^y\left(y_i|\rho\right)$, again coming from the normal distribution as defined in equation (8) .

We integrate over the initial distribution of q_0, which we denote as $h(q_0|X)$. The other latent types are paternal well-being and male income types, which are all discrete and assumed independent. We write the probability of a given type of both father well-being and male income type as $g\left(fw, \rho\right)$. We can write the likelihood function for observation i as

$$L\left(\mathbb{S}_i, M_i^c, M_i^f, Y_i|X_i; \Theta\right) = \int \sum g\left(fw, \rho\right) \tilde{L}\left(\mathbb{S}_i, M_i^c, M_i^f, Y_i|q_0^c, X, fw, \rho; \Theta\right) h\left(q_0^c|X_i\right) dq_0^c. \tag{12}$$

The difficulty in estimation of this model comes through the required integration over the initial unobserved quality: standard numerical integration techniques can have a difficult time since different observations with many different measurements will have mass in nearly distinct regions of the support. In Appendix B, 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 5 presents the parameter estimates for the single index. We normalize the coefficient on white children to 1; therefore, the less-than-1 coefficients on the other racial

Table 5: Single index parameters

Characteristic	Estimate
Education	0.358
	(0.078)
White	1.0
	(.)
Black	0.773
	(0.072)
Other race	0.681
	(0.047)
Mother's age	0.061
	(0.032)
Baby male	-0.037
	(0.029)
Wage percentile	0.973
	(0.197)
Dummy for missing wages	0.017
	(0.033)

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.

indicators imply that non-white children have, on average, lower single index values. The single index further accounts for maternal age and the child's sex. Older mothers tend to have higher single index values, while having a male child is associated with a statistically-insignificant and small lower single index value. To incorporate maternal wages, we calculate each mother's average reported wage across all periods in which wage data are observed and assign her to the corresponding percentile of the wage distribution. The results indicate a positive relationship between wage percentile and the single index. For mothers without wage information, we set the wage percentile to zero and include an indicator for missing wage data. The coefficient on this indicator is small and statistically insignificant, suggesting no meaningful systematic difference for this group. Figure 9 in Appendix A illustrates the distribution of estimated single index values, revealing substantial heterogeneity that allows decisions and outcomes to vary with demographic characteristics.

Utility Table 6 reports the first set of utility parameters, which capture the net utility associated with each possible relationship transition. Because these utilities are identified only up to a normalization, we set the utility of remaining married to the child's birth father equal to zero. All other coefficients should therefore be interpreted relative to this baseline. The estimated utilities are consistent with observed patterns in the data. For instance, the net utility of transitioning from marriage to singlehood is negative, reflecting the empirical stability of marriages. More broadly, transitions that are rarely observed in the data tend to carry strongly negative utility estimates, while more common transitions are associated with relatively higher utilities.

Table 6: Utility from relationship status

Prior relationship status	Current relationship status: Married to father	Cohabiting with father	New partner	Single
Married to father	0	-4.277	-1.787	-1.353
		(0.626)	(0.798)	(0.554)
Cohabiting with father	-0.705	0.190	-0.367	0.476
	(0.667)	(0.117)	(0.786)	(0.597)
New partner	-2.730	-3.154	1.276	0.625
	(0.792)	(0.727)	(0.233)	(0.444)
Single	-2.317	-1.505	0.666	1.144
	(0.593)	(0.579)	(0.487)	(0.161)

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 7 presents the remaining utility parameters. We find a negative utility from leisure, indicating that all else equal we see the women choosing work in the labor market. We interact the single index term with leisure, and find that individuals with higher index values place greater value on not working. The estimated cost of re-entering the labor market is positive, consistent with the low empirical frequency of transitions from non-employment to employment. To allow for demographics to affect relationship decisions, we interact the single index with the utility from being single and find that individuals with higher index values derive greater utility from being single. Utility is increasing in father well-being when the mother is living with the father, consistent with the notion that higher well-being fathers contribute positively to household well being. Finally, we estimate how each dimension of child quality enters the mother's utility. Among these, academic quality exhibits the strongest association with utility, suggesting that improvements in children's academic outcomes are particularly salient for maternal well-being.

Child quality We assume that the child's initial quality vector follows a multivariate normal distribution. Table 8 reports the estimated parameters of this distribution. The top panel presents how the mean of each quality dimension varies with the single index.[17] The bottom panel reports the estimated covariance matrix across the three dimensions of initial quality.

[17]For identification, we fix the effect of the single index on one dimension of child quality, normalizing its impact on academic quality to 0.275.

Table 7: Utility parameters

Utility from leisure	-0.465
	(0.107)
Single index × leisure	0.313
	(0.139)
Cost to re-enter labor market	1.247
	(0.035)
Single index × Single	0.123
	(0.062)
Father well-being × Married/cohabiting	0.879
	(0.172)
Effect of academic quality on utility	0.228
	(0.084)
Effect of physical health quality on utility	0.049
	(0.039)
Effect of mental health quality on utility	0.000042
	(0.040)

We report remaining parameters of the utility function. Standard errors in parentheses.

Table 8: Initial child quality distribution

Mean of distribution			
	Academic	Health	Mental health
Single index	0.275	-0.186	0.233
		(0.121)	(0.111)
Covariance matrix			
	Academic	Health	Mental health
Academic	1	0.169	-0.446
		(0.221)	(0.243)
Health	0.169	1	-0.075
	(0.221)		(0.281)
Mental health	-0.446	-0.075	1
	(0.243)	(0.281)	

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.

Next, we examine the evolution of child quality as specified in equation (10). Tables 15-18 in Appendix A report the estimated coefficients on the constant term, prior quality, the single index, parental separation, father well-being, and total family resources for each relationship and labor market state.[18] Because these parameters are identified only up to

[18]We define separation from the father as a transition from being married or cohabiting to not living together. This effect is therefore only relevant when parents begin the period married or cohabiting. Due to

a normalization, we set the constant term for children whose mothers are single and not working equal to zero. This group thus serves as the reference category for interpreting the remaining coefficients. The estimates reveal substantial heterogeneity across relationship and labor market states, highlighting the value of the model's flexible parameterization. We also find strong interdependence across dimensions of child quality: prior academic quality, for example, not only predicts future academic outcomes but also sometimes influences health and mental health. Father well-being often, though not uniformly, contributes positively to child quality growth. In contrast, total family resources have a statistically significant effect only on health quality.[19]

Measurement of child quality The measurement parameters for child quality are presented in Tables 19-21 in Appendix A. We draw on a wide range of observed outcomes to infer each child's underlying quality. A key strength of our approach is that it provides a coherent framework for combining these heterogeneous measures, allowing the data to determine their relative importance rather than relying on an arbitrary aggregation. For ease of presentation, we group the measures into three tables based on the dimension of quality they most directly capture. However, the underlying estimation is fully flexible: each measurement is allowed to load on all three dimensions of child quality.[20]

To interpret these parameters, we compute the reduction in posterior variance for each dimension of child quality that an econometrician would calculate from observing only that single measurement. A larger reduction in variance indicates that the measure is more informative about that particular dimension of quality. We carry out this calculation for every measure-dimension pair, and present the results in Figures 1-3. Figure 1 shows the results for the academic dimension of quality. Standardized test scores emerge as the most informative measures, while grades at age 15 contribute substantially less. At age 22, we observe wages and educational attainment, which are surprisingly uninformative about academic quality. For wages, this likely reflects the young age at which they are measured: individuals who attended college are just entering the labor market, while those who did not have already accumulated more work experience. Figure 2 examines the health dimension of child quality. Here, health-related measures are clearly the most informative, with other measures

limited statistical power, we impose the restriction that the effect of separation on child quality evolution is the same for married and cohabiting parents.

[19]We restrict the effect of total family resources on child quality to be constant across relationship and labor market states. This restriction is imposed due to the limited sample size and the already large number of parameters to be estimated.

[20]For each dimension of quality, one measurement must serve as an anchor. We impose this by setting its loadings on the other dimensions to zero. Specifically, the anchor measures are the PPVT score at age 9 for academic quality, self-reported health at age 1 for physical health, and bullying at age 15 for mental health.

Figure 1: Reduction of variance (academic)

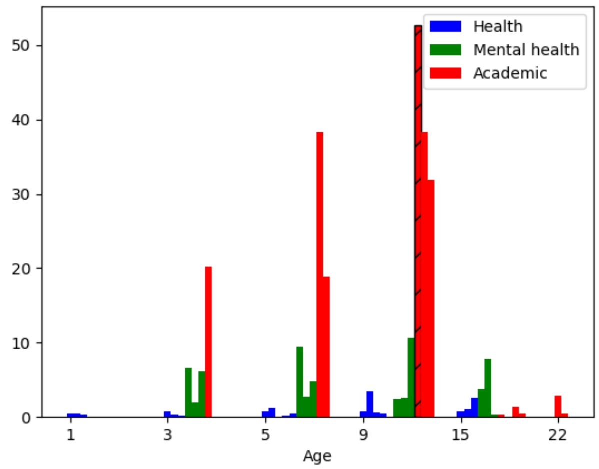

We plot reduction in variance for academic quality for each of the measures. The bar with diagonal lines indicates the measure that was used as a normalization for academic quality.

contributing little to reducing uncertainty about health quality. Interestingly, the age 22 measures of wages and education, while uninformative for academic quality, provide some information about health, suggesting a link between later-life socioeconomic outcomes and physical well-being. Figure 3 presents the results for mental health. As expected, mental health specific measures are highly informative for this dimension. Notably, grades at age 15 also contain some information about mental health, despite playing a relatively limited role in identifying academic quality.

Measurements of father well-being For each measure of father well-being, we estimate (i) the probability that the measure indicates a father has high well-being when the father in fact has poor well-being (a false positive), and (ii) the probability that the measure indicates a father has high well-being when the father truly has high well-being (a true positive). We also estimate the unconditional probability that a father has high well-being. These parameters are reported in the top half of Table 9. Overall, the measures appear to be informative about father well-being, as the true positive rates are much higher than the false positive rates across all measures. The final columns report the determinants of the probability that a father has high well-being, indicating that it is positively correlated with the single index. To further illustrate the informativeness of these measures, Figure 4 plots the posterior probability that each father has high well-being given his observed characteristics. The distribution exhibits substantial mass near both 0 and 1, indicating that for many fathers the data allow

Figure 2: Reduction of variance (health)

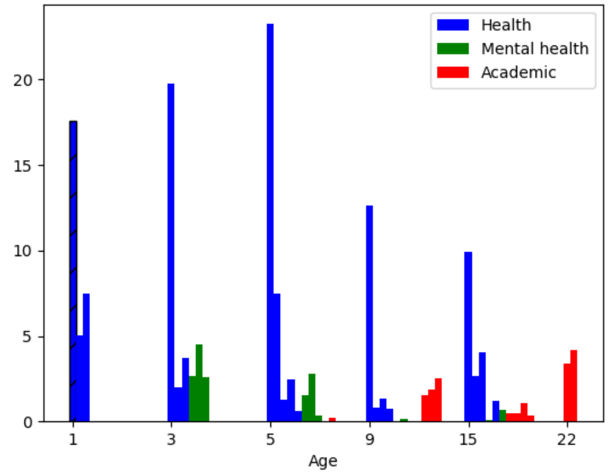

We plot reduction in variance for health quality for each of the measures. The bar with diagonal lines indicates the measure that was used as a normalization for health quality.

Figure 3: Reduction of variance (mental health)

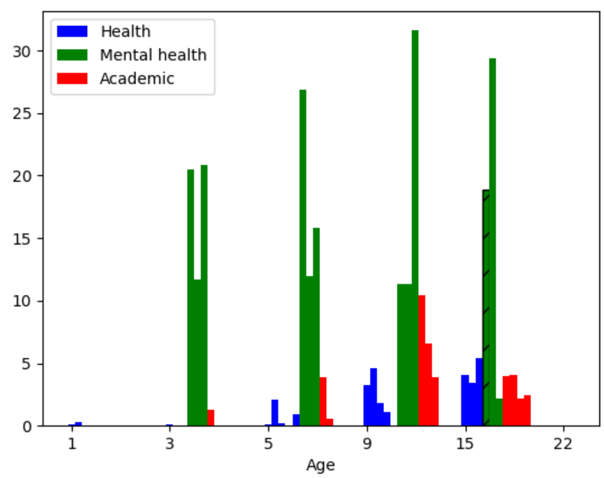

We plot reduction in variance for mental health quality for each of the measures. The bar with diagonal lines indicates the measure that was used as a normalization for mental health quality.

for classification with a high degree of confidence. At the same time, a nontrivial share of fathers have posterior probabilities in the interior, reflecting cases in which the available measures do not fully resolve uncertainty about underlying well-being.

Table 9: Measures of father well-being and male income types

Father well-being measures				
	False positive	True positive	Probability (high quality)	
Alcohol/drugs	0.63	0.94	Constant	0.064
	(0.013)	(0.007)		(0.065)
Violence/abuse	0.71	0.97	Single index	0.50
	(0.012)	(0.005)		(0.10)
Emotional abuse	0.43	0.77		
	(0.016)	(0.012)		
Incarceration/arrests	0.29	0.80		
	(0.015)	(0.013)		
Depression	0.46	0.74		
	(0.015)	(0.012)		
Male income types				
Constant	0.32			
	(0.0020)			
Income type	0.41			
	(0.0029)			
Standard deviation	0.21			
	(0.0013)			

The top half of the table shows the measurement parameters for father quality. 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 last column shows the probability the dad is high quality, where we use a logistic framework. The bottom of the table shows the measurement parameters for father income types. Standard errors in parentheses.

Figure 4: Posterior probability of high well-being father

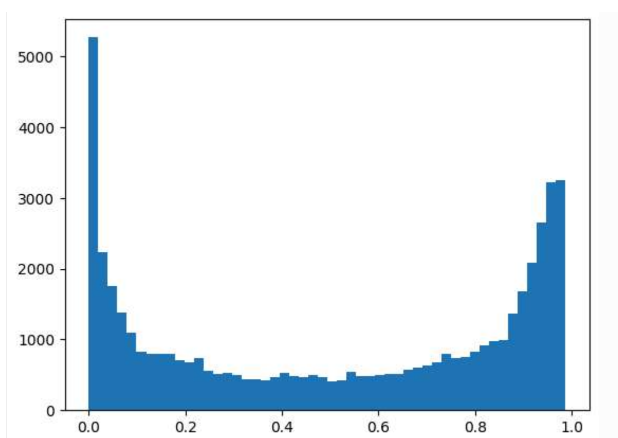

We plot the posterior distribution of the probability that the father is high well-being.

Measurements of male income types We allow fathers and new partners to take one of two income types, coded as 0 or 1. We use reported income as a measurement of income type, and the associated parameters of this measurement model are reported in the bottom half of Table 9. We also estimate the probability that each individual belongs to the high or low income type. For fathers, the estimated probability of being a low income type is 0.63 (standard error 0.01). For new partners, this probability is substantially lower, at 0.50 (standard error 0.028). This pattern suggests positive selection into new partnerships, with mothers more likely to match with higher income type partners when forming new unions.

7 Counterfactuals

In this section, we conduct counterfactual analyses to assess the effects on child outcomes of policies aimed at supporting family formation and stability. These counterfactual outcomes are compared to a baseline set of relationship and labor market choices, as well as child quality trajectories, simulated from the model using the estimated parameters. To generate the counterfactuals, we modify the relevant parameters and re-simulate the model, holding fixed the realizations of unobserved utility shocks.

An important consideration in interpreting these results is that the model does not identify the absolute level of child quality, but only quality relative to a reference group, since the underlying quality does not have a natural scale or location. At time 0, the reference group consists of white, low-education mothers. In subsequent periods, the reference group is defined as a woman with these same characteristics who remains single and non-employed in every period.

7.1 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 utils. This change, chosen at a level to generate a meaningful increase in marriage rates, can be interpreted as capturing policies that promote marriage and reduce divorce without directly affecting parents' ability to invest in their children. Potential real-world analogues include increased social pressure for couples with children to marry or financial incentives such as marriage tax credits, to the extent that any additional income does not directly enter the child quality production function. This counterfactual leads to a near-doubling of marriage rates. In the terminal period, 22% of the sample is married in the baseline, compared to 40% under the counterfactual.

Average impacts on child quality The increase in the utility of marriage in this counterfactual does not directly affect the evolution of child quality, holding choices fixed. Instead, its effects operate indirectly through changes in optimal relationship and labor market decisions. Table 10 reports the average changes in child quality across each dimension in the baseline and counterfactual scenarios at each age. The policy induced increase in marriage rates leads to a moderate improvement in academic outcomes (increases by 0.18 standard deviations at age 15), while its effects on health and mental health are comparatively small.

Table 10: Counterfactual effects on child quality

Child age	Academic	Physical health	Mental health
3	0.025	-0.0036	0.0047
5	0.069	-0.0014	0.0088
9	0.13	0.0094	0.011
15	0.18	0.026	0.010

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

To interpret these magnitudes, we translate the counterfactual changes in latent child quality into expected changes in normalized observed measures. This provides a more concrete, real world benchmark relative to the latent quality measures. As marriage rates rise under the counterfactual, average PPVT test scores at age 9 increase by 0.11 standard deviations. In contrast, the effects on health and mental health measures are smaller: self-reported health at age 3 declines slightly by 0.01 standard deviations, while the measure of bullying behavior improves by 0.004 standard deviations. These patterns indicate that the counterfactual primarily affects academic outcomes, and does not seem to have a substantial impact on health and mental health.

However, the policy leaves the outcomes of the majority of households unchanged, as only marginal households alter their relationship status in response to the policy. Restricting attention to those who change decisions between the baseline and the counterfactual, the effects are substantially larger: normalized test scores increase by 0.37 standard deviations, health improves by 0.034 standard deviations, and mental health improves by 0.008 standard deviations. Taken together, these findings suggest that policies aimed at increasing marriage rates can generate meaningful improvements in children's academic achievement, particularly among the subset of families whose behavior is affected. At the same time, the relatively modest effects on health and mental health highlight that promoting marriage alone is unlikely to produce broad based gains across all dimensions of child outcomes.

Heterogeneous impacts Next, we explore the heterogeneity of the impacts of this policy to understand if certain demographic groups have larger changes to their children's outcome. To do this, for each person we compute their individual average treatment effect (ATE) in this counterfactual.[21] Figures 10-12 in Appendix A present the resulting distributions of the ATEs across the sample. For academic quality, we observe considerable dispersion, with estimated effects ranging from close to zero up to approximately 0.4. The distributions for health and mental health exhibit even greater spread, including negative effects for a subset of individuals. These patterns highlight that the policy's effects are far from uniform: while some children experience sizable gains, others see little change or even slight declines in certain dimensions.

Given this heterogeneity, we examine which subgroups of the population benefit most from this counterfactual; see Kitagawa and Tetenov (2018) and Athey and Wager (2021). This is important to learn how to potentially target policies to have the greatest impact. As a first step, we estimate the relationship between the individual ATE Y and a set of Xs, $E[Y|X]$, nonparametrically using the machine learning LightGBM algorithm of Ke et al. (2017). The X here include the demographics from the model plus latent father well-being and initial child quality. The predicted values from this are a high-dimensional function of X, which we then feed into a shallow 3-level Decision Tree model. The output of this model is a tree structure that separates the sample into subgroups, maximizing the difference in the ATE across the groups. This reduces the dimension of the full relationship between Xs and treatment effects to an interpretable object, allowing us to find the dimensions along which the policy effects vary most strongly.

Figures 5-7 present the estimated trees highlighting which subgroups benefit most from the counterfactual. The rightmost leaves report the estimated ATE for each group, defined by the sequence of splits along observable characteristics. Looking at Figure 5, which shows the results for academic quality, we see that children with fathers with poor well-being and low father income type experience a relatively modest ATE of 0.102. In contrast, the largest gains, an ATE of 0.221, are observed for children with high father well-being and high initial academic quality. This pattern suggests that the counterfactual disproportionately benefits children with more favorable initial conditions, both in terms of family environment and their own baseline skills. A similar pattern emerges for health quality in Figure 6. Children with very low initial mental health experience a negative effect of -0.072, indicating that the policy may even be detrimental for some disadvantaged groups. By contrast, children with high initial mental health and fathers with high well-being experience sizable gains, with an ATE of 0.123. The results for mental health, shown in Figure 7, differ somewhat from this pattern. Here, the largest gains accrue to children with poorer initial conditions -

[21] We simulate their choice and outcome paths ten times under the baseline (control) and ten times under the counterfactual (treatment). We then compute the individual level ATE as the average difference between these simulated outcomes.

specifically, those with low initial health and mental health, while children with high initial health and mental health experience the smallest treatment effects. Taken together, these findings highlight that the effects of marriage promoting policies are highly heterogeneous and depend critically on initial conditions.

Figure 5: Increased utility from marriage: Heterogeneous impacts on academic quality

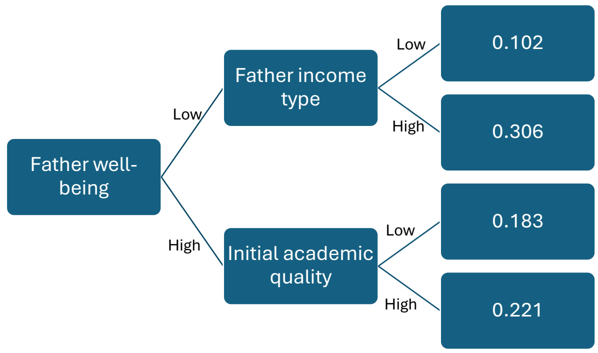

We show the sub-groups most and least affected by this counterfactual in terms of academic quality.

Figure 6: Increased utility from marriage: Heterogeneous impacts on health quality

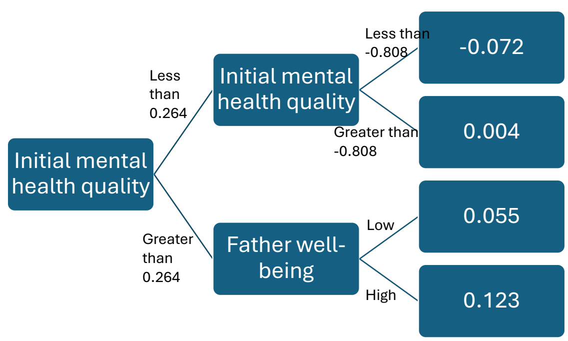

We show the sub-groups most and least affected by this counterfactual in terms of health quality.

Who is most likely to be induced into marriage by this counterfactual? So far, we have learned that when the utility from marriage increases, the impacts on child outcomes are heterogeneous, where on the academic and physical health dimensions children with higher initial conditions improve the most. This is driven by 2 possible mechanisms: (1) certain characteristics make a household more likely to choose marriage in response to the

Figure 7: Increased utility from marriage: Heterogeneous impacts on mental health quality

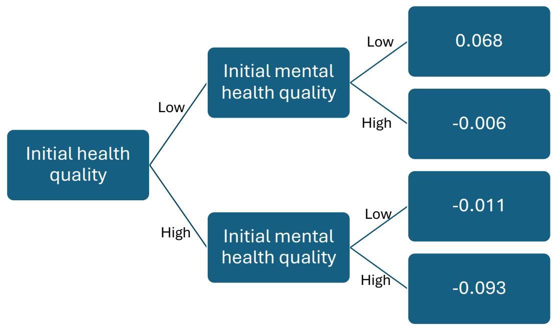

We show the sub-groups most and least affected by this counterfactual in terms of mental health quality.

counterfactual, and (2) characteristics can affect the quality gain in response to marriage. We explore the magnitude of mechanism (1) to understand if it plays a substantial role. For each household, we define an indicator equal to 1 if the mother is married or cohabiting with the child's father and equal to 0 otherwise, and compute the difference in this indicator between the counterfactual and baseline: 0 if there is no change, 1 if the mother is induced from not being with the father to being with him, and -1 if induced the other direction. Figure 8 presents a tree summarizing the characteristics most predictive of transitions into marriage. The results show that when the father has high quality and is a high income type, the household is most likely to transition into marriage or cohabitation in response to the counterfactual. This suggests that the counterfactual primarily induces mothers to form or maintain unions with partners who are already relatively advantaged along key dimensions.

7.2 How does initial child quality affect marriage rates?

In the next counterfactual, we examine how child quality influences family formation and stability. A common concern is that children who struggle may place additional strain on parents, potentially increasing the likelihood of separation. At the same time, if stable families are valued by parents because they help otherwise struggling children, parents may want to stay together. To investigate this, we compare the baseline model simulation to a counterfactual in which each child begins life at the 90th percentile of the initial child quality distribution. Starting from this higher baseline, mothers make optimal relationship and labor market choices, and child quality evolves endogenously in response to these decisions. This exercise allows us to isolate how improvements in initial child ability affect both parental

Figure 8: Increased utility from marriage: Who changes decisions most?

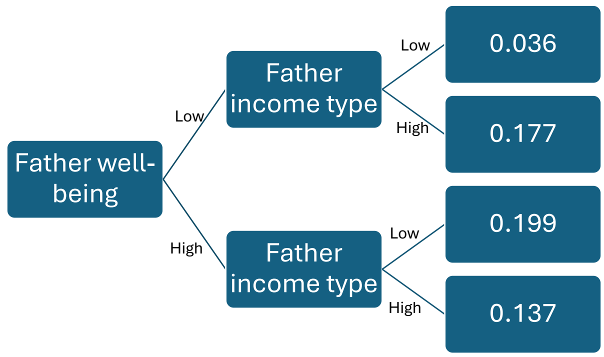

We show the sub-groups most and least likely to choose to marry in response to the increased utility from marriage.

behavior and the subsequent trajectory of family structure.

We define a family as "together" if the child's birth parents are married or cohabiting; in the baseline, this occurs 31% of the time when the child is age 15. When we raise initial child quality, we find modest but meaningful changes in family structure. Increasing initial academic quality leads to a 1 percentage point decrease in the likelihood that parents are together when the child is 15. Similarly, increasing initial health quality reduces this probability by 2 percentage points. In contrast, increasing initial mental health quality results in a 1 percentage point increase in the likelihood that parents remain together. These results suggest that the relationship between child quality and family stability varies across dimensions. Improvements in academic and physical health are associated with slightly lower rates of parental co-residence, while better initial mental health appears to promote family stability.

As before, these average effects mask meaningful heterogeneity across households. For each household in the sample, we compute the probability of being married or cohabiting at age 15 and examine how this probability changes when moving from the baseline to the counterfactual. Figures 13-15 in Appendix A illustrate the distribution of these effects. When initial academic quality is increased, many households exhibit little to no change, or modest increases in co-residence rates, while others experience declines of up to nearly 4 percentage points. In contrast, increasing initial health quality leads to more uniformly negative effects on the mother remaining with the father: for most households, co-residence rates decline, with reductions reaching roughly 4 percentage points at the lower end of the distribution. The effects are most dispersed when initial mental health quality is increased. Some households experience declines in the likelihood of parents remaining together, while others see increases of up to approximately 4 percentage points. This wide distribution indicates the presence of offsetting forces, with improvements in mental health strengthening family stability in some cases while weakening it in others.[22]

7.3 How does cohabitation affect child outcomes?

Cohabitation is common in the Fragile Families sample, which oversamples low-income, urban, unmarried parents. Much of the existing literature on family structure focuses on divorce, but this is not the relevant margin for many families in our sample, for whom the key decision is whether to remain in a cohabiting relationship with the child's father. We therefore examine the extent to which the effects of cohabitation on child outcomes operate through relationship stability.

In the estimated model, cohabiting relationships are substantially less stable than marriages. We interpret this difference as reflecting the lower legal and social costs of separation among cohabiting couples relative to married couples. As a result, cohabiting mothers transition to single parenthood at a considerably higher rate than married mothers. At the same time, the model allows child quality to evolve differently in married and cohabiting households. To isolate the role of family stability, we conduct a counterfactual in which cohabiting relationships are assigned the same separation costs as marriages. This intervention equalizes the persistence of cohabitation and marriage while leaving the direct effects of each relationship state on child quality unchanged. Consequently, any differences in child outcomes can be attributed to the increased stability of cohabiting relationships rather than to changes in the direct effect of cohabitation itself.

In particular, consider the net utility matrix row corresponding to moving from "Married with the Father", the first row of Table 6. The all-else-equal next-period utility cost of moving from married to "New Partner" is 1.787 utils, while the cost of moving from married to "Single" is 1.353. Contrast this with the "Cohabiting with father" row, the second row in Table 6, which gives the net cost of moving from cohabiting to with living with a new partner as 0.557 utils and the net cost of moving from cohabiting with the father to single as -0.286 (that is, a net utility benefit). These estimates show it is much more costly to leave marriage than cohabitation, and in this counterfactual we directly adjust the net costs of moving from cohabitation to new partner and single. We do this by simply decreasing the net utilities of moving to the two states until the costs match the marriage costs: it now

[22] We also construct shallow Decision Trees to identify which subgroups are most affected by this counterfactual, but do not find strong or systematic patterns. This likely reflects the relatively limited variation in treatment effects across observable characteristics. These results are available upon request.

costs 1.787 utils to move from cohabiting to a new partner, and 1.353 from cohabiting to single. (We leave the cost of moving from cohabiting to marriage, and all other entries in the net utility matrix, unchanged.)

These counterfactual results show that relationship paths diverge significantly from the baseline. Marriage rates actually increase from 22% in the baseline to 28% in the counterfactual. This increase in marriage is because the state dependence of the estimated model makes moving from "cohabiting" to "single" to "married to the father" extremely unlikely. In the counterfactual, we reduce the "cohabiting" to "single" part of this transition, thus increasing the rate of eventually moving into marriage, which itself is very stable.

We examine how this counterfactual affects child outcomes. In the terminal period, converting raw quality to test scores, academic test scores increase by 0.04, health scores increase by 0.01, and mental health scores decrease by 0.001. For academic outcomes, these effects are substantial, as they are about 40% of the impact we saw from the counterfactual that increased the utility from marriage. This suggests that the benefit of stability, by decreasing the share of cohabiting households that move to the single state, are substantial.

As before, the average effects reported above include many individuals who never enter into cohabitation and therefore are not directly impacted by the counterfactual. To isolate the impact on those affected by this counterfactual, we restrict the sample to mothers who were cohabiting when the child was born. In this group, we find an increase of 0.10 standard deviations in academic test scores. When we further look only at mothers induced into living with the father in the terminal period (either through cohabitation or marriage), the increase in test scores is much larger, equal to 0.27 standard deviations. These results indicate that the instability of cohabitation, given the higher future probability of separation from the father, has negative impacts on child academic outcomes.

8 Conclusion

This paper studies how family structure shapes children's development across multiple dimensions, including academic achievement, physical health, and mental health. We develop a dynamic model of maternal relationship and labor supply decisions, which we estimate using data on choices and children's outcomes from the Fragile Families dataset. We consider relationship status between birth parents along a broader spectrum than previous work, allowing cohabitation with the father or repartnering as alternatives to marriage or singleness. Additionally, we also use information on the well-being of the child's father, which impacts both relationship choices and the evolution of child outcomes. After estimating the model, we evaluate the impact of policies that alter incentives to form and maintain partnerships, and examine how these changes propagate through both parental behavior and child development.

Our findings highlight that the effects of policies encouraging marriage or cohabitation are fundamentally heterogeneous and operate through multiple channels. While marriage subsidies generate modest improvements in children's academic outcomes on average, they have limited effects on physical and mental health. These average effects mask substantial variation across families: children with more favorable initial conditions benefit primarily through gains in academic and physical outcomes, whereas those with disadvantaged starting points experience larger improvements in mental health. These results emphasize that policies aimed at increasing marriage rates are unlikely to uniformly improve child outcomes, and that targeting specific populations may be more effective.

Additional counterfactual analyses underscore the importance of considering the dynamic effects of policies on choices, since different relationship decisions early in the child's life lead to different relationship paths later in life. First, we find a bidirectional relationship between child quality and family structure. Improvements in children's mental health modestly increase relationship stability, whereas gains in academic achievement or physical health reduce the likelihood that parents remain together. Second, we find that cohabitation can be costly to children relative to marriage not because it is a state that harms them directly, but because its unstable nature means the relationship is more likely to dissolve in the future. Making cohabitation as stable as marriage actually leads to more marriages in the future and gains in child academic 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., Cattan, S., Fitzsimons, E., Meghir, C., and Rubio-Codina, M. (2020a). Estimating the production function for human capital. *The American Economic Review*, 110(1):48–85.

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

Bang, M. and Borghesan, E. (2026). A dynamic evaluation of parental marriage and children's skill development. *Working Paper*.

Bernal, R. (2008). The effect of maternal employment and child care on children's cognitive development. *International 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*.

Jackson, C. K. (2018). What do test scores miss? the importance of teacher effects on non-test score outcomes. *Journal of Political Economy*, 126(5):2072–2107.

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.

A Additional tables and figures

Table 11: 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

We show the mean and standard deviations of academic measures in the data. PPVT and WJ are standardized tests given to survey respondents at different ages. Grades are on a 1-4 scale.

Table 12: 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

We show the mean and standard deviations of self-reported health (1-5 scale).

Table 13: 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	1.2%	5.9%	92.3%	1,598
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

We show the distribution of responses for these health measures. Some are binary Yes/No questions, whereas others also have a Somewhat option.

Table 14: 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. We show the average score for each measure. All measures have a minimum value of 0, and the maximum value as indicated in the table. Higher numbers indicate better mental health outcomes.

Figure 9: Estimated single index values

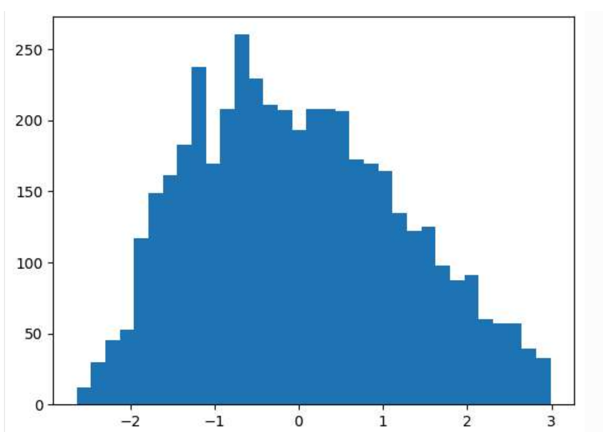

We plot the distribution of the estimated single index values.

Table 15: Child quality transition parameters, married

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	0.109	0.170	0.077	0.107	-0.163	0.238
	(0.062)	(0.082)	(0.066)	(0.058)	(0.086)	(0.080)
Prior academic quality	0.143	-0.056	0.064	0.168	-0.069	0.082
	(0.051)	(0.087)	(0.071)	(0.051)	(0.086)	(0.068)
Prior health quality	0.308	0.125	0.121	0.240	0.131	0.149
	(0.070)	(0.064)	(0.075)	(0.065)	(0.062)	(0.073)
Prior mental health quality	-0.201	-0.187	-0.013	-0.181	-0.230	-0.004
	(0.060)	(0.073)	(0.042)	(0.056)	(0.073)	(0.040)
Single index	-0.044	-0.043	-0.078	0.090	0.082	-0.063
	(0.033)	(0.048)	(0.037)	(0.034)	(0.045)	(0.027)
Separation from father	-0.077	0.075	0.036	-0.077	0.075	0.036
	(0.057)	(0.065)	(0.055)	(0.057)	(0.065)	(0.055)
Father with high well-being	0.050	0.080	0.115	0.115	0.292	0.117
	(0.052)	(0.067)	(0.057)	(0.061)	(0.091)	(0.078)
Family resources	0.00000091	0.11	0.053	0.00000091	0.11	0.053
	(0.022)	(0.033)	(0.029)	(0.022)	(0.033)	(0.029)

Standard errors in parentheses. We report the quality transition parameters for each dimension of quality for working and non working states when the mother is married to the child's father. Separation from father applies to the period where the mother exits marriage or cohabitation with the father, and we assume that this impact is constant across all relationship and labor market states. We assume that the effect of total family resources on child quality does not vary with relationship status or labor market participation.

Table 16: Child quality transition parameters, cohabiting

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	-0.053	-0.030	0.046	-0.050	-0.128	0.245
	(0.055)	(0.076)	(0.055)	(0.057)	(0.094)	(0.075)
Prior academic quality	0.123	-0.116	0.124	0.141	-0.009	0.172
	(0.053)	(0.107)	(0.081)	(0.053)	(0.111)	(0.081)
Prior health quality	0.229	0.177	0.248	0.196	0.182	0.293
	(0.075)	(0.074)	(0.091)	(0.074)	(0.082)	(0.092)
Prior mental health quality	-0.248	-0.339	0.037	-0.257	-0.344	-0.033
	(0.071)	(0.073)	(0.051)	(0.068)	(0.081)	(0.057)
Single index	-0.062	-0.026	-0.060	0.029	0.042	-0.018
	(0.035)	(0.055)	(0.040)	(0.033)	(0.056)	(0.035)
Separation from father	-0.077	0.075	0.036	-0.077	0.075	0.036
	(0.057)	(0.065)	(0.055)	(0.057)	(0.065)	(0.055)
Father with high well-being	0.040	0.203	0.185	0.049	0.238	0.054
	(0.052)	(0.076)	(0.073)	(0.062)	(0.105)	(0.082)
Family resources	0.00000091	0.11	0.053	0.00000091	0.11	0.053
	(0.022)	(0.033)	(0.029)	(0.022)	(0.033)	(0.029)

Standard errors in parentheses. We report the quality transition parameters for each dimension of quality for working and non working states when the mother is cohabiting with the child's father. Separation from father applies to the period where the mother exits marriage or cohabitation with the father, and we assume that this impact is constant across all relationship and labor market states. We assume that the effect of total family resources on child quality does not vary with relationship status or labor market participation.

Table 17: Child quality transition parameters, other partner

	Working			Not working		
	Academic	Health	Mental health	Academic	Mental Health	health
Constant	-0.298	0.378	-0.208	-0.304	0.278	-0.039
	(0.086)	(0.123)	(0.089)	(0.085)	(0.123)	(0.078)
Prior academic quality	0.155	0.120	0.232	0.132	0.005	0.199
	(0.051)	(0.103)	(0.085)	(0.052)	(0.108)	(0.084)
Prior health quality	0.260	0.126	0.152	0.252	0.102	0.163
	(0.075)	(0.078)	(0.080)	(0.076)	(0.085)	(0.090)
Prior mental health quality	-0.188	-0.235	0.023	-0.220	-0.262	0.024
	(0.052)	(0.080)	(0.059)	(0.060)	(0.089)	(0.060)
Single index	-0.044	0.014	-0.073	0.166	0.130	-0.116
	(0.035)	(0.070)	(0.046)	(0.054)	(0.075)	(0.050)
Father with high well-being	-0.112	0.015	-0.033	-0.227	-0.288	0.027
	(0.092)	(0.207)	(0.139)	(0.135)	(0.318)	(0.188)
Family resources	0.00000091	0.11	0.053	0.00000091	0.11	0.053
	(0.022)	(0.033)	(0.029)	(0.022)	(0.033)	(0.029)

Standard errors in parentheses. We report the quality transition parameters for each dimension of quality for working and non working states when the mother is married or cohabiting with a new partner. We assume that the effect of total family resources on child quality does not vary with relationship status or labor market participation.

Table 18: Child quality transition parameters, single

	Working			Not working		
	Academic	Health	Mental health	Academic	Health	Mental health
Constant	-0.043	0.191	0.045	0	0	0
	(0.039)	(0.067)	(0.048)			
Prior academic quality	0.134	-0.032	0.144	0.112	0.008	0.193
	(0.050)	(0.086)	(0.075)	(0.055)	(0.098)	(0.081)
Prior health quality	0.285	0.090	0.144	0.247	0.197	0.213
	(0.073)	(0.064)	(0.078)	(0.075)	(0.076)	(0.084)
Prior mental health quality	-0.170	-0.233	0.034	-0.222	-0.282	0.063
	(0.053)	(0.070)	(0.048)	(0.059)	(0.078)	(0.059)
Single index	-0.070	-0.020	-0.092	0.080	0.072	-0.139
	(0.035)	(0.049)	(0.037)	(0.034)	(0.049)	(0.040)
Father with high well-being	0.096	0.037	0.117	-0.019	0.151	0.234
	(0.050)	(0.055)	(0.048)	(0.058)	(0.093)	(0.074)
Family resources	0.00000091	0.11	0.053	0.00000091	0.11	0.053
	(0.022)	(0.033)	(0.029)	(0.022)	(0.033)	(0.029)

Standard errors in parentheses. We report the quality transition parameters for each dimension of quality for working and non working states when the mother is single. We assume that the effect of total family resources on child quality does not vary with relationship status or labor market participation. As a normalization, the constant terms when not working and single are set to 0.

Table 19: Measurement parameters for child outcomes (academic)

		Constant	Coefficient on quality dimension			Error
Measurement	Age	term	Academic	Physical health	Mental health	term
PPVT score	3	-0.0149	0.4931	-0.0768	0.1057	0.8518
		(0.0304)	(0.0541)	(0.0373)	(0.0449)	(0.0174)
PPVT score	5	0.0250	0.6392	-0.0900	0.0931	0.7294
		(0.0346)	(0.0634)	(0.0439)	(0.0506)	(0.0105)
WJ score	5	-0.0077	0.4808	-0.0248	0.1441	0.8435
		(0.0293)	(0.0562)	(0.0367)	(0.0434)	(0.0106)
PPVT	9	-0.0193	0.6144	0	0	0.5830
		(0.0506)	(0.0747)			(0.0060)
WJ9 score	9	-0.0346	0.5498	0.0291	0.0221	0.6903
		(0.0467)	(0.0691)	(0.0232)	(0.0209)	(0.0079)
WJ10 score	9	-0.0607	0.5103	0.0571	0.0606	0.7174
		(0.0437)	(0.0663)	(0.0230)	(0.0213)	(0.0042)
English grade	15	-0.1369	0.1552	0.0590	0.2623	0.9119
		(0.0355)	(0.0470)	(0.0492)	(0.0408)	(0.0132)
Math grade	15	-0.1297	0.1210	0.0644	0.2532	0.9339
		(0.0339)	(0.0430)	(0.0479)	(0.0405)	(0.0168)
History grade	15	-0.1391	0.2034	0.0784	0.2348	0.8979
		(0.0358)	(0.0491)	(0.0444)	(0.0392)	(0.0190)
Science grade	15	-0.1240	0.1579	0.0462	0.2210	0.9250
		(0.0325)	(0.0432)	(0.0419)	(0.0369)	(0.0141)
Education	22	-0.2702	0.1663	0.1386	0.0982	0.8427
		(0.0611)	(0.0636)	(0.0577)	(0.0541)	(0.0144)
Wages	22	-0.2052	0.0672	0.1941	0.0685	0.9571
		(0.0590)	(0.0394)	(0.0774)	(0.0460)	(0.0319)

Standard errors in parentheses. 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 20: Measurement parameters for child outcomes (health)

Measurement	Age	Constant term	Coefficient on quality dimension: Academic	Physical health	Mental health	Error term
Overall health	1	-0.1228	0	0.3899	0	0.8440
		(0.0604)		(0.0836)		(0.0113)
Physical disabilities	1	-0.0550	0.0206	0.2137	-0.0290	0.9548
		(0.0483)	(0.0175)	(0.0478)	(0.0227)	(0.0492)
Asthma	1	-0.1278	0.0279	0.2664	0.0268	0.9442
		(0.0460)	(0.0200)	(0.0584)	(0.0249)	(0.0173)
Overall health	3	-0.1630	0.0185	0.4203	0.0156	0.8513
		(0.0643)	(0.0201)	(0.0893)	(0.0258)	(0.0132)
Lacks energy	3	0.0607	0.0555	0.1334	0.0384	0.9766
		(0.0336)	(0.0186)	(0.0324)	(0.0231)	(0.0258)
Physical disabilities	3	-0.0599	0.0172	0.1860	-0.0011	0.9653
		(0.0400)	(0.0176)	(0.0429)	(0.0223)	(0.0412)
Overall health	5	-0.1561	0.0153	0.4445	0.0126	0.8103
		(0.0667)	(0.0187)	(0.0956)	(0.0244)	(0.0132)
Asthma	5	-0.1872	-0.1094	0.2937	0.1133	0.9206
		(0.0544)	(0.0229)	(0.0668)	(0.0328)	(0.0230)
Stomach issues	5	-0.0784	-0.0192	0.1192	0.0442	0.9898
		(0.0302)	(0.0168)	(0.0306)	(0.0222)	(0.0473)
Underactive	5	-0.0714	0.0365	0.1515	0.0436	0.9733
		(0.0349)	(0.0192)	(0.0356)	(0.0234)	(0.0215)
Overweight	5	-0.1011	-0.0459	0.0918	0.0813	0.9898
		(0.0370)	(0.0187)	(0.0324)	(0.0294)	(0.0265)
Overall health	9	-0.2740	-0.0653	0.3614	0.1696	0.8644
		(0.0767)	(0.0255)	(0.1026)	(0.0461)	(0.0128)
Allergies	9	-0.1827	-0.1314	0.1224	0.1632	0.9597
		(0.0501)	(0.0262)	(0.0434)	(0.0434)	(0.0088)
Headaches	9	-0.1772	-0.0422	0.1328	0.1254	0.9786
		(0.0410)	(0.0165)	(0.0400)	(0.0341)	(0.0116)
Overweight	9	-0.1054	-0.0388	0.1009	0.0949	0.9864
		(0.0394)	(0.0183)	(0.0353)	(0.0320)	(0.0117)
Underactive	9	-0.1653	-0.0305	0.1244	0.1318	0.9797
		(0.0407)	(0.0174)	(0.0400)	(0.0368)	(0.0342)
Overall health	15	-0.2816	-0.0557	0.3278	0.1952	0.8886
		(0.0739)	(0.0247)	(0.0936)	(0.0504)	(0.0137)
Activity limited	15	-0.2197	-0.0561	0.1831	0.1716	0.9553
		(0.0504)	(0.0199)	(0.0540)	(0.0433)	(0.0077)
Asthma	15	-0.2698	-0.1091	0.2300	0.1966	0.9250
		(0.0645)	(0.0240)	(0.0685)	(0.0501)	(0.0172)

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 21: 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	-0.4379	-0.1009	0.2056	0.4107	0.8413
		(0.0878)	(0.0271)	(0.0662)	(0.0950)	(0.0154)
Social skills	3	-0.3708	-0.0234	0.2308	0.3314	0.8633
		(0.0783)	(0.0252)	(0.0697)	(0.0769)	(0.0221)
Sulks	3	-0.4815	-0.0861	0.2015	0.4171	0.8335
		(0.0966)	(0.0376)	(0.0834)	(0.1378)	(0.0200)
Disobedient	5	-0.5935	-0.1229	0.1775	0.4690	0.8206
		(0.1034)	(0.0474)	(0.0773)	(0.1546)	(0.0114)
Social skills	5	-0.4397	-0.0481	0.1919	0.3240	0.8797
		(0.0795)	(0.0299)	(0.0776)	(0.1087)	(0.0094)
Aggressive	5	-0.4063	-0.0683	0.0993	0.3677	0.8951
		(0.0823)	(0.0311)	(0.0476)	(0.1221)	(0.0227)
Social skills	9	-0.3752	-0.0157	0.0606	0.3138	0.8834
		(0.0734)	(0.0179)	(0.0314)	(0.1066)	(0.0176)
Destructive	9	-0.3522	-0.0183	0.0420	0.3139	0.8896
		(0.0710)	(0.0194)	(0.0276)	(0.1069)	(0.0237)
Argumentative	9	-0.5152	-0.0948	0.0336	0.4892	0.7732
		(0.1028)	(0.0395)	(0.0329)	(0.1657)	(0.0273)
Bullies	15	-0.4272	0	0	0.3827	0.7928
		(0.0887)			(0.1310)	(0.0502)
Disobedient	15	-0.5286	-0.0491	0.0380	0.4735	0.7621
		(0.1029)	(0.0295)	(0.0344)	(0.1596)	(0.0285)
Social	15	-0.2223	-0.0039	0.0932	0.1504	0.9734
		(0.0458)	(0.0168)	(0.0400)	(0.0531)	(0.0140)

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.

Figure 10: Increased utility from marriage: Distribution of impacts on academic quality

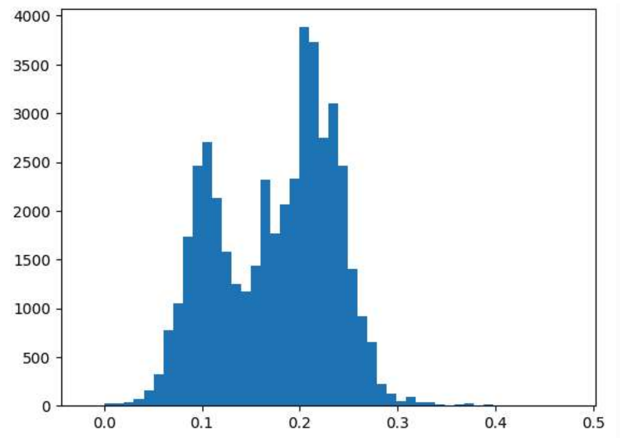

We plot the distribution of the individual ATE for academic quality for this counterfactual.

Figure 11: Increased utility from marriage: Distribution of impacts on health quality

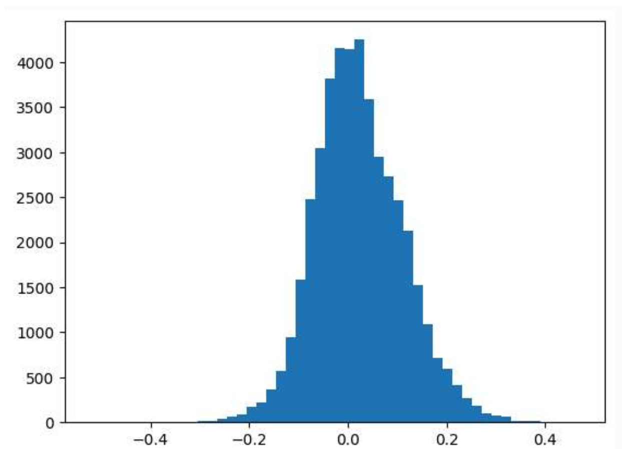

We plot the distribution of the individual ATE for health quality for this counterfactual.

Figure 12: Increased utility from marriage: Distribution of impacts on mental health quality

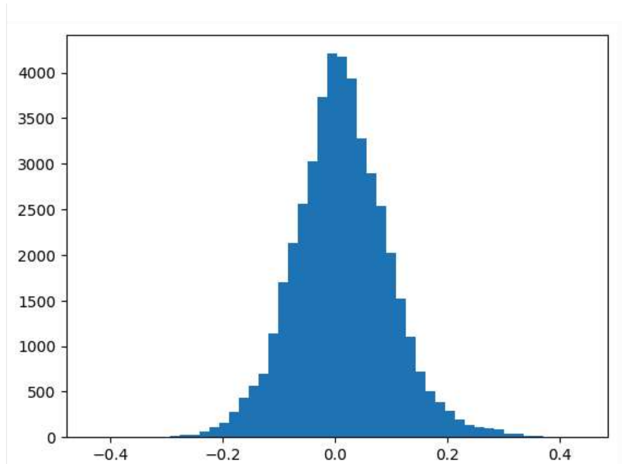

We plot the distribution of the individual ATE for mental quality for this counterfactual.

Figure 13: Increased initial academic child quality: Distribution of impacts on marriage/cohabitation rates

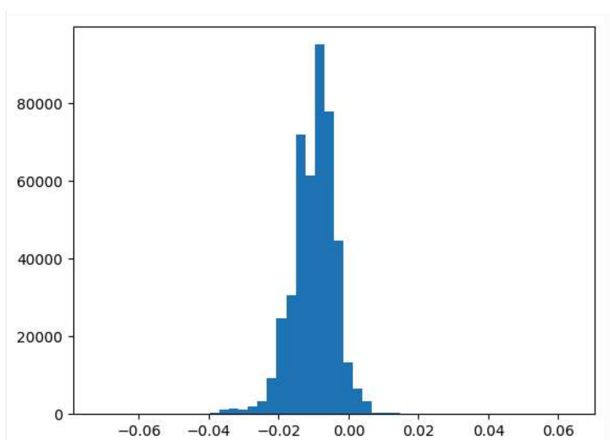

We place all children at the 90th percentile of initial academic quality and look at the impact on marriage/cohabitation rates.

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

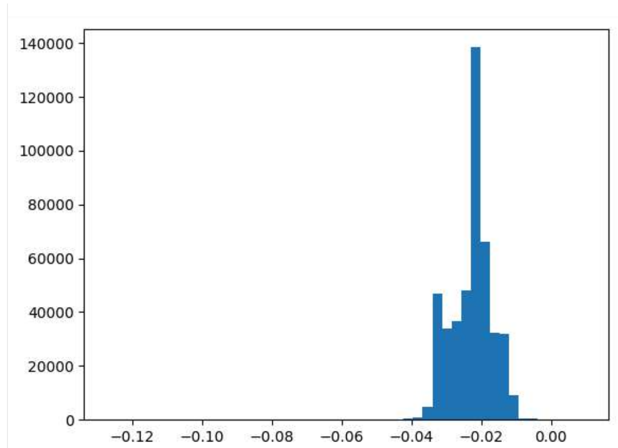

We place all children at the 90th percentile of initial health quality and look at the impact on marriage/cohabitation rates.

Figure 15: Increased initial mental health child quality: Distribution of impacts on marriage/cohabitation rates

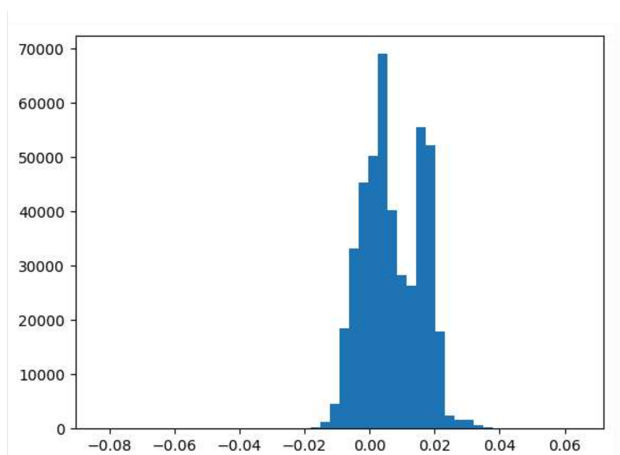

We place all children at the 90th percentile of initial mental health quality and look at the impact on marriage/cohabitation rates.

B 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

$$L\left(\mathbb{S}_i, M_i^c, M_i^f, Y_i|X_i;\Theta\right) = \int \sum g\left(fw,\rho\right) \tilde{L}\left(\mathbb{S}_i, M_i^c, M_i^f, Y_i|q_0^c, X, fw, \rho;\Theta\right) h\left(q_0^c|X_i\right) dq_0^c. \tag{13}$$

Moving the father well-being and the father/new partner income type measures to the outside, and using conditional independence across time and measures, we can write this likelihood for one individual as

$$L_i = E_{fw,\rho|X}[E_{q_0^c|X}[\prod_{t=1}^{T} \Pr(s_t|q_t^c, s_{t-1}, X, fw, \rho) \prod_{v=1}^{\#M} g(m_v^c|q_t^c, X) \prod_{v=1}^{\#Y} g(m_v^y|\rho, X)]], \tag{14}$$

that is, the likelihood is the expected likelihood of the observed choices times the likelihoods of the observed child quality measures and male income 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 child quality 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

$$L_i = C_{\tilde{q}} \cdot E_{fw,\rho}[E_{\tilde{q}_0}[\prod_{t=1}^{T} \Pr(s_t|\tilde{q}_t, s_{t-1}, X, q^f) \prod_{v=1}^{\#Y} g(m_v^y|\rho, X)]], \tag{15}$$

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

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.