Family Law Effects on Divorce, Fertility and Child Investment

Meta Brown
The Ohio State University

Christopher J. Flinn
New York University

Joseph Mullins
University of Minnesota

Abstract

In order to assess the child welfare impact of policies governing divorced parenting, such as child support orders, child custody assignments, and marital dissolution standards, one must consider their influence on not only the divorce rate but also spouses' fertility choices and child investments. We develop a model of fertility, parenting, and divorce, from which we derive estimates of parental preferences and a child cognitive ability production function, using data on parental time allocation, children's cognitive attainment, and realized fertility and divorce. Family policies that reduce divorce are simulated to have significant negative impacts on both fertility and child development.

1 Introduction

Divorced parenting in the U.S. is regulated through a combination of laws regarding marital dissolution, child custody and placement, and the assignment and enforcement of child support obligations. The primary objective of these rules is to increase the well-being of children and parents, and the divorce rate is often regarded as a first-order measure of the success of family law. The rationale for this focus is the preponderance of empirical evidence that suggests that children living in households without both biological parents are more likely to suffer from behavioral problems and have lower levels of a broad range of achievement indicators measured at various points over the life cycle (see, e.g., Haveman and Wolfe 1995). Empirical studies of unilateral divorce laws and child support enforcement have isolated the effects of changes in such legal structures on divorce rates (e.g., Friedberg 1998, Gruber 2004, Wolfers 2006, and Nixon 1997). Recent empirical studies of gender-neutral custody standards and financial incentives for joint custody provide new evidence on the effects of custody policy reforms on children's attainment (Chen and Logan 2020 and Kranz et al. 2021). A complete picture of the influence of family law on family members' welfare should include an understanding of the mechanisms by which family law changes influence fertility, child outcomes, and the distribution of resources within the family, in addition to divorce rates. Our objective is to take a first step in this direction by modeling the interactions of married couples over fertility, child investment, and divorce in the shadow of existing divorce regulations. Our goal is to understand, and quantitatively evaluate, the weight of the law in shaping parental behavior and, ultimately, child outcomes.[1]

Most studies of parental decision-making regarding investments in children and child outcomes take a very limited view of parental interactions. In many cases, the father is

This research was supported in part by grants from the National Science Foundation, Collegio Carlo Alberto and the C.V. Starr Center for Applied Economics at NYU. David Blau, Hanming Fang, Ariel Pakes, Ken Wolpin and seminar participants at the Institute for Research on Poverty, NYU, Virginia, UNC-Greensboro, Wisconsin, Torino, the Society for Economic Dynamics, ESPE, the American Economic Association Meetings, the UNC/Duke Conference on Labor, Health and Aging, the Minnesota Applied Micro Workshop and the Stanford Institute for Theoretical Economics provided valuable comments.

[1]Our more formal analysis builds on the original insights of Mnookin and Kornhauser (1979).

either absent (as in Mullins 2021; 2023) or considered a passive agent in this process, whose role is limited to providing income for the household (as in Bernal 2008, Bernal and Keane 2011, and Liu et al. 2010). In some cases, parents' preferences have been represented by a unitary utility function, as in Del Boca et al. (2014, 2024), which obscures the question of how these unitary preferences are formed from the individual preferences of the parents prior to or after marriage.[2] Much of what we have learned about parents' dynamic decision-making, therefore, has been in the context of a mother's (or mother and father's, assuming a unitary objective) individual dynamic optimization problem. When studying the influence of divorce law on the family, however, the distinct choices of mothers and fathers are paramount. For example, it is virtually impossible to understand the influence of potential child support on fertility, investment, and divorce decisions by studying the mother's perspective in isolation. Hence we model the choices of mothers and fathers as an ongoing, simultaneous-move game. Our model and data begin from the date of marriage, which, while excluding a substantial and non-random segment of parents, has the benefit of granting access to similar information on the mother and father when early fertility, investment and divorce choices are being made.

We draw on an extensive empirical literature on marriage dynamics, including Aiyagari, Greenwood and Guner (2000), Brien, Lillard and Stern (2006), Chiappori, Fortin and Lacroix (2002), Bruze, Svarer, and Weiss (2015), and Voena (2015). Models in these papers emphasize the repeated interaction of a husband and wife in deciding whether to continue a marriage and the allocation of household resources.[3] Within this literature, our contribution is to endogenize fertility and child investment decisions during and after marriage. Given that parents are forward-looking, divorce laws and regulations influence all of these decisions both within intact marriages and when parents are divorced.

We draw methods and insights from a comparatively recent line of research that models

[2]A notable exception is Verriest (2024), who uses the framework of Del Boca et al. (2014) to examine investments in children and child outcomes in intact households. Unlike Del Boca et al., he allows parents to have different preferences over child outcomes and labor supply. Parents' decisions are coordinated through the maximization of a weighted average of their utilities.

[3]In the model of Voena (2015), for example, spouses' endogenous choices are savings and divorce, which are determined under varying divorce access and property distribution standards.

the investment decisions of households with respect to child "quality," which is taken to be cognitive ability. In particular, our framework shares certain modeling choices with those found in Del Boca et al. (2014). These authors examine spouses' time allocation decisions, including investment time with the child and market labor supply, under the assumption of a Cobb-Douglas production technology forming the cognitive ability of the child and a Cobb-Douglas household utility function. Caucutt, Lochner, Mullins, and Park (2020) further expand the set of productive inputs to examine the balance among two parents' time inputs, market goods, and market child care in generating children's attainment. We introduce marriage dynamics and fertility into this framework, with the risk of marital dissolution and the laws governing divorced parenting allowed to influence the choices made by parents prior to and following divorce and even impacting the decision to have a child.

Two prior papers bear noting, as they also model marriage dynamics, fertility, and child investment. Caucutt, Guner and Knowles (2002) model marriage dynamics, fertility, and child expenditures. Their model is more comprehensive than ours on some dimensions, as it includes a marriage market and multiple generations. However, where our approach is one of regular, repeated interactions between parents throughout the fertility and child-rearing process, they employ a three-period overlapping generations framework to study these issues. This suits their object of interest, the life-cycle timing of fertility. In contrast, we model spouses' decisions throughout the fertility and child-rearing process, and our ultimate interest is in child outcomes and parents' welfare and their relationship to family law.

A second related paper is Tartari (2015). Tartari addresses the question of whether a child whose parents divorced would have been better off if the parents had remained married. She also models fertility, time and goods investments in a child, and divorce. Relative to our approach, Tartari elaborates the role of conflict within marriage in spouses' child investment decisions and in shaping divorce and child outcomes, as she models an endogenous mode of interaction for married couples and relies on explicit marital harmony or conflict survey instruments from the National Longitudinal Survey of Youth's 1979 cohort (NLSY79). The extent of conflict in the marriage plays a central role in Tartari's inferences regarding whether a child of divorce would be as well off under the counterfactual continued marriage. While Tartari emphasizes conflict data and choices, we focus on time use, gathered through detailed time diaries collected in the Child Development Survey (CDS) from the Panel Study of Income Dynamics (PSID), in 1997, 2002, and 2007 for both mothers and fathers. Our approach allows us to study parents' explicit time investment in children of all ages, for stable families and for families approaching and experiencing separation and divorce.

In our model, spouses make (simultaneous) choices regarding marriage continuation, fertility, and, where relevant, individual investments in children. A match value of the marriage is drawn from a population distribution and evolves stochastically over time. Fertility choices are influenced by both the expected benefit from the presence of the child and expectations regarding the duration of the marriage, given the state of the marriage quality process. The child progresses as a result of both endogenous parental investment and marital status choices and exogenous productivity factors. Further, child quality is self-productive. In this manner, our child quality production process builds on those of Del Boca et al. (2014) and Cunha and Heckman (2007) by incorporating marriage quality and family structure stability.

Marital dissolution may result from changes in marriage match quality, changes in child presence, and the child reaching maturity. Thus the full history of marriage values and child investments determines current marital status and child investment levels. If the history of child investments and marriage values is poorer for the marginal marriage than it is for the average marriage, then, all else equal, the child welfare gain associated with the continuation of the marginal marriage is smaller than that associated with the continuation of the average marriage. A central objective of our analysis is to study the welfare impacts of variations in family law, which are possible to assess under our assumptions regarding the determination of the utility levels of husbands, wives, and (potential) children. The ability to use the model to infer distinctions between the marginal marriage, which may dissolve in response to changes in prevailing child support, child placement, or divorce standards, and the average marriage is crucial to our ability to recognize distinctions between the divorce rate and family welfare consequences of family policy regimes.

The model is estimated using data from the PSID and the CDS (a subsample of households with young children from the PSID) using the Method of Simulated Moments (MSM). The model we estimate is somewhat parsimonious. The primary benefit of this stylized approach is our ability to infer from our short vector of estimated parameters the primary economic processes that drive any particular observed relationship between family outcomes and policies. The natural drawback to parsimony is often a deficiency in model fit. Nevertheless, we find that the model is able to fit the features of the data used in estimation satisfactorily. The accuracy of fit gives us some confidence in using the estimated model to perform comparative statics exercises and welfare analysis.

The parameter estimates generated by the MSM estimation reveal several new insights regarding the production of children's cognitive ability. Chief among these is the finding that the influence on the growth of child cognitive ability of persisting in a low quality marriage is substantially worse than the influence of realized divorce on the growth of child cognitive ability. In addition, our parameterization of child cognitive development allows us to study the productivity of parents' time investments as children age. We find that mothers' and fathers' time is highly productive in advancing children's cognitive ability in early childhood, but that the productivity of these time investments declines steeply as children age toward independence.[4] Finally, as in Cunha and Heckman (2007) and others, children's current cognitive ability is estimated to be highly productive in generating future cognitive gains.

Our analysis concludes with a series of counterfactual policy experiments in which we manipulate dimensions of family law and simulate the responses of divorce, fertility, and children's cognitive ability to the law changes using the model under the parameter values estimated using our PSID-CDS sample and the prevailing family law regime. An important feature of our model is the incorporation of a fertility decision. In the comparative stat-

[4] This is consistent with the patterns found in Del Boca et al. (2014)

ics exercises we find that family law potentially has an important impact not only on the achievement levels of children from intact and non-intact households, but even more fundamentally on the number of children born and the characteristics of the households having them. The dynamics of selection that determines which households have children and which households divorce turn out to be crucial in determining average skill outcomes in the model.

We first simulate a move from the unilateral divorce standards that have been widely adopted by US states over the past decades, and that characterize our baseline model, to a more restrictive (and outdated) mutual consent divorce standard. Our point estimates indicate that, while these movements toward restricting divorce lead to much lower divorce rates, this comes at the cost of lower levels of both fertility and children's cognitive achievement.[5] A bootstrap estimate of the sampling distribution for the counterfactual suggests that there is considerable uncertainty around these predictions, which is driven by uncertainty in the parameters that govern fertility choices and downstream divorce outcomes by marriage quality.

In additional counterfactual policy experiments, we simulate large changes in the child support transferred from fathers to mothers in the divorce state. This policy change also decreases the rate of divorce, and the decreased rate of divorce comes at the cost of reducing both fertility and child cognitive attainment. In the case of increased child support, the decline in cognitive attainment is noisy, and masks two offsetting effects: on one hand, higher child support frees mothers' time from working for investment in the child (and father's time investment increases as well). On the other hand, the increased child support requirements in divorce change the subgroups of couples who choose to have children during marriage and who choose to enter divorce; the resulting population of parents is characterized by weaker ongoing marriages and lower total factor productivity in the production of child cognitive ability.

The simulated effect of a large change in child support on spouses' welfare and children's

[5] Moreover, evidence from Stevenson and Wolfers (2006) suggests that limiting divorce may bring dire consequences for spouses locked in dangerous marriages.

skill depends critically on the prevailing divorce law. When we simulate doubled child support assuming each state's prevailing divorce laws, we see no clear direction and very wide confidence intervals for the mean changes in wives' welfare and children's skill. However, when we calculate the effects of doubling child support assuming a mutual consent standard for every state, we find that wives' welfare and children's skill increase significantly and substantially. In contrast, when we calculate the effects of doubling child support assuming a unilateral divorce standard for every state, wives' welfare declines modestly (though significantly) and children's mean skill declines meaningfully. Fertility and divorce decrease in all of our child support hike scenarios. These findings underscore the crucial role played by changing the selection among marriages into fertility and divorce for the determination of welfare and child skill outcomes.

A final counterfactual policy simulation experiments with the share of custody, in terms of physical child placement, granted to the father. While divorce, fertility, and mothers' time investment in the child prove relatively unresponsive to expected custody allocations, the father's time investment in the child follows the father's custody share. The child's cognitive attainment is positively (if noisily) associated with the father's share of custody up to the point of equal custody. At point estimates, the difference in skill outcomes between these two regimes is modest, at 5 percent of a standard deviation. As with all of our counterfactuals, this captures the difference in impacts when the policy is known at the time of fertility, which may be quite different from the short-run impacts of unanticipated policy changes.[6]

Overall, our policy simulations indicate that children's cognitive outcomes are not best served either by minimizing divorce or by requiring (implausibly) large resource transfers to the custodial parent in the event of divorce. Despite the extensive evidence that children of divorce, on average, fare worse on several dimensions than children of marriage, a careful analysis of marriage quality heterogeneity demonstrates that the divorce rate among parents

[6]For evidence of this kind, see Chen and Logan (2020) on the role of custody threats in divorce settlement bargaining, a process that lies beyond the dynamic child investment game that is the focus of our analysis. Chen and Logan find that gender-neutral custody standards decrease high school graduation rates, possibly by weakening mothers' bargaining positions.

is a poor policy target when the social objective includes supporting children's cognitive development. Our results emphasize the crucial role played by fertility and divorce decisions when policies are known for the determination of spousal welfare and child skill outcomes. The effect of marriage quality on child outcomes has a particularly formative role in shaping the path from this dynamic selection process to skill outcomes. Such effects are unlikely to fully emerge when studying impacts in the data on children who are already born at the time of a policy reform.

The plan of the paper is as follows. In Section 2, we describe the PSID and CDS data in detail, shows descriptive statistics from our sample, and document some important patterns in the data which we seek to capture in our model. Section 3 develops the details of the model. Section 4 presents our estimation method and discusses the manner in which primitive parameters are identified. Section 5 reports the estimates of the primitive parameters and our assessment of model fit. In Section 6, we describe our various counterfactual policy experiments and interpret and decompose the simulated effects of changes to family law. Section 7 concludes.

2 Data and Motivating Evidence

In this section we document important patterns in the data that shine some light on the dynamics among divorce, the cognitive development of children and parental investment. Although the relationships we discuss here are not necessarily causal, they inform our development of the model and are potentially suggestive of particular causal processes. A key challenge for this paper, and one that will ultimately influence our policy experiments, is decomposing the developmental effect of divorce on children.[7]

[7] This is the key question also posed by Tartari (2015).

2.1 Data

To answer the empirical and policy questions outlined in Section 1, we use data from the Panel Study of Income Dynamics (PSID) and its Child Development Supplement (CDS). The PSID is a dynastic, longitudinal survey conducted annually from 1968 to 1997, and biennially since 1997. It collects information on a range of economic and demographic indicators. The CDS consists of three waves, collected in 1997, 2002 and 2007. Any child in a PSID family between the ages of 0 and 12 at the time of the 1997 survey was considered eligible. These surveys contain a broad array of developmental scores in cognitive and socio-emotional outcomes, as well as information on the home environment of the child. One crucial feature of the survey for our purposes is the availability of time use data, which is collected from the participants using detailed time diaries. We provide further details below.

2.2 Description of Variables and Sample Selection

From the PSID's primary survey we collect data on mothers' and fathers' labor supply, labor income, total family income, and some demographic variables.[8] The PSID additionally provides individual histories of marriage and childbirth that we use to construct our sample.

The CDS is composed of several questionnaires. We use two in particular: the child interview and the primary caregiver (PCG) interview. From the child interview, we draw our measure of cognitive ability based on the Applied-Problems (AP) module of the Woodcock-Johnson Aptitude test. The cognitive ability measure that we incorporate in the model estimation is a record of the number of items on the AP module that the child answers correctly, out of a possible 60.

In addition, the CDS asks each participant child to fill out a time diary. This portion of the survey requires participants to record a detailed, minute by minute timeline of their activities for two days of the week: one random weekday and one random weekend day. Activities are then coded by PSID time diary specialists at a fine level of detail. When necessary, chil-

[8] We use the main survey for all available years between 1975 and 2010.

dren are assisted in completion of the time diary by the PCG. These diaries provide detail on the daily lives of children that is unparalleled among economic data resources.

From the coded time use diaries, we construct a measure of weekly parental time investment by taking a weighted[9] sum of the total hours of time use in which each parent is recorded as actively participating in each diary activity. A key advantage of the PSID-CDS data is this balanced evidence on the time spent with children by both mothers and fathers. Other valuable survey resources containing detailed economic data provide extensive detail on mothers' time investment in children, but they omit or make asymmetric or limited record of the time invested in child development by fathers. The presence of fathers' time in the data is, in some cases, contingent on the parents' continued marriage. In the PSID-CDS, we are able to measure time investment from married, divorced, and single mothers and fathers, and this feature of the PSID-CDS data is crucial to our analysis.

It is worth noting, in this study of married and divorced parenting, that the active time that we measure could conceivably be performed simultaneously by both parents at times in married households. In the model we describe below, the production technology takes each parent's active time as separate inputs without discounting when this active time is spent simultaneously. While this leads to the possibility of "double counting" time investment, it avoids the complication of additionally modeling joint time investment in marriage and divorce.[10]

The focus of this project is on initially married parents and their fertility and divorce decisions, along with their distinct child investment choices. Accordingly we restrict our sample to first marriages that (1) begin between 1975 and 1997; (2) have not ended in widowhood; and (3) involve partners with no recorded childbirths prior to the year of marriage.[11] This

[9]We calculate total weekly time investment from the single measured weekday and weekend day in the time diary as Weekly Time = 5 × Weekday Total + 2 × Weekend Day Total

[10]We thank a referee for pointing out the possibility of simultaneous time investment. In Del Boca et al. (2024), intact households with both parents include individual investment time for each parent alone as well as a third input which is the time investment when the parents are together with the child. We have not included this separate time input here due to the complexity of solving this model with endogenous fertility, divorce, and the wife's labor supply over and above the cognitive ability of the child.

[11]Note that couples that have births very near the time of marriage will be included under this sample

leaves us with a sample of 2,525 marriages to which 1,201 children from the CDS can be linked.[12]

Table 1 reports descriptive statistics for the estimation sample. Wives in the sample are, on average, 2 years younger than husbands at the time of marriage. The attainment rates of four-year college degrees are comparable for husbands and wives at 26 percent and 28 percent, respectively.

We observe a cumulative sample divorce rate of 21 percent by 2018. Figure 1 depicts the timing of these divorces relative to the birth of the first child, with 8 percent of divorces occurring before the birth of any child and the remainder relatively evenly distributed.

2.3 Patterns of Time Investment and Cognitive Skill Outcomes

Central to estimation of the model's parameters is the problem of disentangling the complex dynamic causal processes surrounding divorce and child outcomes. As the model will clarify, divorce is the outcome of a dynamic selection process on unobservable marital quality and other state variables that affect child outcomes both before and after dissolution.

Figures 2 and 3 illustrate two key features of the data that will help pin down causal parameters in the model. Using the panel dimension of the PSID we can separate children into three groups: (1) those whose parents remain married throughout the sample; (2) those whose parents are already divorced; and (3) those whose parents will divorce in the future. A comparison between groups (2) and (3) allows the model to decompose skill differences between divorced and married groups into a component that is due to the reduction in time inputs after divorce and a component that is due to the marriage being of lower quality prior to divorce. This decomposition is critical for determining the skill implications of upcoming counterfactual reforms to divorce law. Current marriages that will soon end are, under the model developed in Section 3, marginal marriages. To the extent that marginal marriages

criterion. Empirically, a meaningful subset of marriages occur during pregnancy. Requiring pregnancy to predate marriage may be an unnecessarily restrictive sample criterion for our purposes.

[12]We retain couples with no children, one child, and multiple children. However, the model in Section 3 only treats the arrival of the first child.

are less productive in the advancement of child skills than average marriages, policies that have the goal and effect of reducing the divorce rate may produce disappointing gains in terms of child attainment.

Figure 2 depicts average time investment of fathers and mothers for each group. We observe a large reduction in fathers' time investment after divorce. Comparison with the "will divorce" group suggests that this difference is not likely to be explained by selection on household-level unobservables. Accordingly, Appendix A.1 supports this conclusion using within-child variation in the marital status of their parents. In contrast, we see only a very small reduction in mothers' time investment following divorce, with no statistically significant difference between the "divorced" and "will divorce" groups, suggesting that any differences could be driven by selection on unobservables. Appendix A.1 shows that marital status plays no role in predicting mothers' time investment after accounting for child or family fixed effects (in contrast to the findings for fathers' time inputs).

The changes in active time investments around divorce among US parents that we report in this section can be compared to changes in overall time with children around divorce among Australian parents reported by Cano and Gracia (2022). Both sets of results indicate a divergence between mothers' and fathers' time with their children as the couple moves into divorce, with fathers' time declining steeply. In addition, both studies reveal a decline in total active or structured time investment in the child from the marriage to the divorce state.[13]

Figure 3 compares average raw test scores for the three marital status groups in four age groups. Although there are small skill gaps between children of divorced parents and children whose parents never divorce, average test scores for the "will divorce" group are consistently the lowest among the three marital status groups.[14] In order to fit this evidence, given a

[13]It may also be worth noting that Cano and Gracia's broader time measure, and perhaps their different cultural context, reveals a large increase in mothers' time with the child from marriage to divorce, which they term a "time penalt(y)" of divorce "associated with gender inequalities".

[14]Note that, although the raw mean of the AP score is lowest for the "will divorce" group in each age range, the "will divorce" mean score is significantly below the never divorced mean score only for the 10 to 13 age range. The analogous difference in means is marginally significant for the zero to five age range.

clear reduction in time inputs after divorce (shown in Figure 2), the model must allow for lower quality marriages to also play some role in determining skill gaps. In Appendix A.1 we replicate this evidence using regressions that alternatively condition for household and child fixed effects, showing that standardized test scores are positively associated with divorce after controlling for time invariant differences across children and households. We also verify in Appendix A.1 that these patterns for skill outcomes are robust to an alternative definition of marital status that uses the timing of separation rather than the timing of divorce.

2.4 Institutional Context

Marital dissolution standards underwent a revolution in the US following the precipitous divorce law reforms in California in 1969. An earlier movement away from fault to no fault divorce was followed by a shift away from mutual consent marital dissolution standards, in which both spouses must consent to a divorce for it to be granted, toward unilateral divorce standards, in which only one spouse must request a divorce for it to be granted. The large majority of states moved to unilateral divorce over the course of the 1970s (Friedberg 1998, Gruber 2004), six states underwent reforms from 1980 to 2014, and all but 12 states maintained unilateral divorce standards by 2014 (Ciacci 2023). Researchers have examined the influence of unilateral divorce standards on divorce rates (Friedberg 1998), children's attainment (Gruber 2004), household savings and women's employment (Voena 2015), and prostitution arrests (Ciacci 2023), among many other family outcomes. By the 1997 to 2007 measurement window available to us in the PSID CDS and its time diaries, most states applied unilateral divorce standards. Our approach in the model and estimation below is to determine divorce choices under the prevailing dissolution standards in the family's state, a

These patterns can be compared with the findings of Tartari (2015), using the National Longitudinal Survey of Youth (NLSY) 1979 cohort. Tartari shows that NLSY children of divorced parents and of married parents who will divorce have similar standardized scores on the PIAT math assessment, and that their scores are significantly and substantially below those of children with parents in stable marriages. Tartari's NLSY data lack time diaries, and so her analysis does not include the time use dynamics around divorce for mothers and fathers that we present in Figure 2. However, Tartari does demonstrate that both divorced NLSY mothers and mothers who will divorce work longer hours than mothers in stable marriages, a result that has influenced our specification of mothers' labor supply and earnings dynamics.

non-negligible minority of which retained mutual consent standards during our estimation window. This implies that the prevailing marital dissolution standard for each sample family contributes to our identification of the parameters governing our baseline estimated model. In later simulations, we investigate the influence of (1) moving all families to mutual consent divorce standards and (2) moving all families to unilateral standards under the estimated model.

We also examine child support standards. Each US state maintains formal child support guidelines. These are functions of the number of children and the incomes of fathers and mothers, with additional factors entering the guidelines of some states. The modal child support rate for a family with one child in our estimation sample is 15 percent. This reflects a modal output from state child support guidelines under the assumptions that the mother retains majority child placement and the couple have one child. However, some state guidelines, under these assumptions, yield a support level as high as 30 percent (Sun 2008, Kranz, Roff, and Sun 2021).[15] Hence, our simulations using the estimated model examine changing outcomes for children, mothers, and fathers when child support levels rise from 15 to 30 percent and, finally, to the likely politically infeasible level of 50 percent.

2.4.1 Auxiliary Data on Custody Allocations

We simplify custody allocations in the model by assigning them randomly at the time of divorce from a known distribution. We use data from the child support module of the Current Population Survey's (CPS) Annual Social and Economic Supplement (Flood et al 2023). We construct from these data an approximate measure of physical custody (measured as the fraction of days per year that the child is with the father) and estimate the distribution of allocations among divorced couples. Appendix C provides additional details.

[15] We thank Hugette Sun for guidance on the implementation of child support and interpretation of child support guidelines.

3 A Model of Fertility, Child Investment and Divorce Decisions

In this section we develop a model to investigate life-cycle patterns of fertility, married female labor supply, divorce, and parental investments in children. Given such a large collection of endogenous dynamic processes, in order for us to empirically investigate the relationships between them some model structure must be imposed. This also enables us to perform counterfactual policy experiments later in the paper.

3.1 Environment, Demographics, and Choices

All individuals begin their life-cycle in a new marriage as either a husband (H) or a wife (W). Time is discrete and indexed by t, with $t = 1$ corresponding to the first period of the marriage. In any given period t, the couple's choice set may contain a one-time fertility decision ($F_t \in \{0, 1\}$), a divorce decision ($D_t \in \{0, 1\}$), time investment decisions in the child's production process ($\tau_{W,t}, \tau_{H,t} \in \mathbb{R}$), and a work decision for wives ($L_t \in \{0, \frac{1}{2}, 1\}$).[16]

The couple faces a finite horizon that corresponds to the date T at which the wife reaches the age of 60, or $A_{W,T} = 60$. Any childless married couple may choose to have a child as long as $A_{W,t} \leq 42$, and we do not model the decision to have subsequent children. We refer to the periods t for which $A_{W,t} \leq 42$ as the "fertile period." Children are characterized by their cognitive skill, k, which is malleable with respect to time investments of the parents until the child reaches age 19. This defines five life-cycle stages. Individuals transition between these stages according to their fertility and divorce decisions, as well as their age. The list below defines these stages and the choices associated with each stage.

- **Stage 1** *(Married without children and with the possibility of having a child)* All married couples begin in stage 1, since at the time of marriage the wife is of child-bearing

[16]These choices correspond to no work, part-time hours, or full-time hours.

age. In this stage couples make fertility decisions, divorce decisions, and the wife makes a work decision.

- **Stage 2** *(Married with a developing child)* In this case we do not consider the possibility of further child-bearing in the current model. Couples in this stage can make time investment decisions in the child and can decide to divorce. The wife additionally makes a work decision.

- **Stage 3** *(Married with a fully developed child or married without the option of fertility)* In this stage, the couple no longer is supplying time to child investment. Their only decision regards divorce and the wife's working hours.

- **Stage 4** *(Divorced with a developing child)* In this stage the couple is making time investment decisions in the child and the wife makes a work decision.

- **Stage 5** *(Divorced with a fully developed child, or divorced without children)* This state is absorbing until the couple reaches their finite horizon. The only choice is the wife's work decision.

Figure 4 summarizes this life cycle structure by depicting the events that precipitate movements in between stages.

3.1.1 Mutual and Unilateral Decisions

In modeling the behavior of married and divorced parents, an important specification choice is the manner in which spouses interact. One may assume that spouses interact cooperatively or noncooperatively.[17] It is unclear that ex spouses are able to interact in a manner that results in efficient outcomes that lie on the Pareto frontier. In a model that moves though married and divorced states, if cooperation is ever attained in marriage it is unclear how spouses' mode of interaction might transition from such cooperation in marriage to the

[17] For examples of the cooperative and non-cooperative approaches, respectively, see Browning and Chiappori (1998), Lundberg and Pollak (1994), and Del Boca and Flinn (2011).

potential cooperation failures of divorce, or how the presence of children might influence interactions in the divorce state. One might assume cooperation throughout, but this skirts important implementation issues.[18] More complex approaches include allowing spouses to choose the current mode of interaction as events progress, following Flinn (2000), or specifying population heterogeneity in spouses' mode of interaction, following Eckstein and Lifshitz (2009), Del Boca and Flinn (2011), and Tartari (2015). Though the latter approaches are appealing, they would add a great deal of complexity to an already intricate model.

For the above reasons, and given evidence in Friedberg (1998) and Gruber (2004) that prevailing unilateral or mutual divorce requirements have impacts on divorce, we have chosen to assume that husbands and wives make their decisions *non-cooperatively.* In particular, each spouse's time investment decisions and the work decisions of the wife are always made in a unilateral fashion.

Nevertheless, the model assumes that all fertility decisions occur by voluntary agreement (each spouse must choose this option independently). Let $F_t \in \{0,1\}$ indicate the final fertility decision in period t, and let $F_{S,t} \in \{0,1\}$ indicate the decision of spouse S. Mutual consent requires that:

$$F_t = F_{W,t} F_{H,t}$$

The nature of the divorce decision ($D_t \in \{0,1\}$) depends on the prevailing policy environment. Let Q indicate the (persistent) prevailing marriage dissolution standard, with $Q = 1$ representing a mutual consent divorce environment and $Q = 0$ representing a unilateral divorce environment. Under our assumption of noncooperative behavior, divorce decisions are not, in general, efficient. In a mutual consent environment, the same rule applies:

$$D_t = D_{W,t} D_{H,t},$$

[18]Cooperative agreements may be sustainable within households because frequent monitoring makes it possible to implement agreements that are not best responses. When the parents of a child do not cohabit, this monitoring is not possible, making the implementation of cooperative agreements relatively implausible.

where $D_{S,t}$ indicates the decision of spouse S. By contrast, in a unilateral or no-fault divorce environment, the decision rule is:

$$D_t = D_{W,t} + (1 - D_{W,t})D_{H,t}.$$

Finally, we assume that the divorce environment is persistent, and therefore spouses need not form expectations regarding the future divorce environment when making their decisions.[19] Hence, all observed historical changes to the divorce standard are implemented in simulation but are not anticipated by agents in the model.

3.1.2 Divorce Law and Divorce Decisions

There are three parameters that define the divorce law environment. First, the dissolution standard determines whether individuals can *unilaterally* dissolve the marriage, or if mutual consent is required. Second, after divorce, the (ex) husband's income is taxed at some rate π, and transferred to the wife. Finally, for couples with children under 18 seeking divorce, the husband is awarded a custody/visitation allocation, $\overline{\tau}$. This allocation is drawn from a distribution $F_{\overline{\tau}}$. While custody and visitation rights are quite distinct legal concepts in reality, the allocation $\overline{\tau}$ will affect the father's access to the child in a way that makes no such distinction.

[19]The unilateral divorce movement in the US occurred largely over the 1970s. Within our PSID CDS sample period, from 1997 to 2007, and in fact between February 1987 and August 2009, only West Virginia (September 2001) and New Jersey (January 2007) moved from a mutual consent to a unilateral divorce standard (see Ciacci 2023). Hence, we allow persistent differences among states in marriage dissolution standards to shape our estimates, but we assume that the marriage dissolution laws faced by a given couple persist from the time of marriage, precluding computationally costly expectations over the evolution of marriage dissolution standards in the couple's state.

3.2 Technology, Endowments, and Constraints

3.2.1 Income

Regardless of marital state, income processes obey:

$$\log(Y_{H,t}) = \gamma_{Y,H,C,1} + \gamma_{Y,H,C,2}A_{H,t} + \gamma_{Y,H,C,3}A_{H,t}^2 + \varepsilon_{Y,H,t} \tag{1}$$

$$\log(Y_{W,t}) = (\gamma_{Y,W,C,1} + \gamma_{Y,W,C,2}\kappa_t + \gamma_{Y,W,C,3}\kappa_t^2)L_t \tag{2}$$

$$\kappa_t = \sum_{s=1}^{t} L_s \tag{3}$$

$$\varepsilon_{Y,H,t} = \rho_\varepsilon \varepsilon_{Y,H,t-1} + \eta_t,\ \ \eta \sim \mathcal{N}(0, \sigma_\eta^2) \tag{4}$$

where $C \in \{0, 1\}$ indicates whether the individual has completed a college education. Earnings for each spouse depend on experience. Husbands are assumed to supply full time labor inelastically, and hence their earnings simply correspond to age $A_{H,t}$. Wives choose among full time work, part time work, and no work at each instance, represented by $L_t \in \{0, \frac{1}{2}, 1\}$, and our model allows the wife's current earnings to depend on accumulated experience κ_t. The addition of the labor supply choice variable and the the experience state variable for the wife helps the model to fit key empirical differences in labor supply among mothers who remain married, mothers who will divorce, and divorced mothers, evidence of which appears in the subsection 2.3 descriptive analysis above and in Tartari (2015).

3.2.2 Time Use

The budget constraint on time use is:

$$l_{H,t} + [1 + \rho(1 - \overline{\tau})]\tau_{H,t} = 72 \tag{5}$$

$$l_{W,t} + \tau_{W,t} = 112 - 40L_t \tag{6}$$

Here we assume that each agent expends 112 waking hours per week, with the husband allocating $72 = 112 - 40$ hours between leisure and child investment and the wife allocating

112 hours among leisure, child investment, and formal labor supply (the latter limited to 0, 20, or 40). In addition, $\overline{\tau}$ is the share of days in the year that the child resides with the father. Thus, during marriage, $\overline{\tau} = 1$, and time is invested without any additional marginal cost. When divorced, ρ summarizes the additional marginal time cost of investing in the child, designed to capture the additional difficulties of investing time when proximity to the child is limited.[20]

3.2.3 Consumption

Consumption is a public good in marriage and equal to total household income (no borrowing or saving):

$$C_t = (Y_{H,t} + L_t Y_{W,t})$$

When divorced, consumption is no longer a public good. With some loss of generality, we define the husband to be the payer of child support in the divorced parenting state. As is the case in many U.S. states, the order amount is determined as a proportion of the father's income: the husband is taxed at a rate π for child support (when the couple have a child of 18 or younger). This reduces the father's income and is transferred as non-labor income to the mother.[21] We assume that each spouse has no non-labor income when married. Thus while divorced with children, consumption obeys:

$$C_{H,t} = (1 - \pi) Y_{H,t} \tag{7}$$

$$C_{W,t} = L_t Y_{W,t} + \pi Y_{H,t} \tag{8}$$

[20]Implicitly, we have set $\rho = 1$ for mothers in the model, based on our findings on the time investment choices of mothers after divorce in Section 2 and Appendix A.1.

[21]We assume that there is perfect compliance with child support order on the part of the father. Del Boca and Flinn (1995) develop and estimate a model in which child support orders and compliance decisions are determined endogenously.

One would expect some relationship between the mother's share of custody and the level of child support she receives from the father. Rather than imposing a functional relationship between $\overline{\tau}$ and π, in the empirical exercise we impose the model custody allocation and base π on predominant child support guidelines and family characteristics. Our model estimation is performed assuming the prevalent child support tax rate, and later policy experiments investigate the effects of reforming to the current maximum among existing state child support guidelines, as well as to a rate far outside of the politically feasible range.

When divorced with no children, we assume a small consumption floor for wives, denoted by b:[22]

$$C_{H,t} = Y_{H,t} \tag{9}$$

$$C_{W,t} = \max\{b, L_t Y_{W,t}\} \tag{10}$$

3.2.4 Fertility and Child Development

Married parents without children have the option to bear them before the wife turns 42 ($A_{W,t} \leq 42$). Children arrive immediately, with initial skill k_0, in the same period in which the decision is made.

Marriages are characterized by a level of quality, ω, that evolves as a first order Markov process given by the transition matrix Π_ω.

The investment problem that we model begins at birth. Our empirical implementation focuses on progress from birth through a set of tests that are completed for most sample children before the age of ten. When married, child skills evolve according to a Cobb-Douglas production function:

$$k_{t+1} = \psi(A, \omega)\tau_{H,t}^{\delta_{H,A}}\tau_{W,t}^{\delta_{W,A}}k_t^{\delta_k}\exp(\xi_t)$$

where A is the age of the child, ω is the current value of the stochastically evolving marriage quality, and ξ_t is a factor-neutral independent and identically distributed (i.i.d.) technology shock that is normally distributed with mean zero and standard deviation σ_ξ.[23] When

[22] In related work, Del Boca et al. (2014) find that most PSID households with children in their sample have little non-labor income.

By assuming that there is no transfer ordered after a divorce if the couple is childless, we are essentially assuming away alimony. Alimony is increasingly uncommon in U.S. divorce cases. According to Case et al. (2003), for example, 4.2 percent of divorced mothers in the 1997 PSID who received a child support or alimony payment (or both) from an ex husband received alimony.

[23] The Cobb-Douglas functional form that we assume for the production function recalls Del Boca et al. (2014). Their approach to modeling two parents' repeated interaction over investment in shared children produces a dynamic solution with various desirable analytic properties. As a result, we, and other researchers, have chosen to inherit their methods. See Del Boca et al. for a more detailed description of the solution benefits of this functional form choice.

parents are divorced, child production is affected through TFP, but is otherwise identical:

$$k_{t+1} = \psi_d(A)\tau_{H,t}^{\delta_{H,A}}\tau_{W,t}^{\delta_{W,A}}k_t^{\delta_k}\exp(\xi_t).$$

The child quality production function described here can be related to leading models of child investment. Cunha and Heckman (2007) and Cunha et al. (2006) argue that a variety of skills that children must develop are subject to "critical periods" early in life, and hence much of intellectual development is accomplished by the time the child reaches school age. Hopkins and Bracht (1975), for example, demonstrate that IQ is stable by the age of 10 or so, suggesting that the critical period for intellectual development occurs by this time. Further, Cunha and Heckman, Cunha et al., and Cunha, Heckman and Schennach (2010) emphasize the importance of both cognitive and non-cognitive skill acquisition to child outcomes, along with the importance of "dynamic complementarity" and "self-productivity" of skill levels in ongoing skill production. Todd and Wolpin (2003, 2007) consider cognitive skill formation, and argue from a different perspective for the importance of both current and lagged inputs to the ongoing production process. They demonstrate the importance of allowing for unobserved endowment effects and the endogeneity of inputs to child skill production.

Like Todd and Wolpin (2003, 2007) and Del Boca et al. (2014), we restrict attention to cognitive skill. In our empirical implementation, we relate k_t to cognitive test scores from the CDS in a way that permits noisy measurement. Age-specific time and ability production parameters allow for critical periods in the production of cognitive ability. Our production function entails both dynamic complementarity and self-productivity of skill. The initial conditions that we specify when estimating the model directly address the need to account for unobserved endowment heterogeneity, and the model permits the endogenous determination of investments throughout the child development process.

3.3 Preferences

Individuals discount the future at an exponential rate, $1-\beta$, and make decisions to maximize the expected discounted stream of utilities, the deterministic component of which can be described as:

$$U_{St} = \begin{cases} \log(C_t) + \alpha_l \log(l_{S,t}) + \nu(\omega_t) + F_t \alpha_F & \text{Stage 1} \\ \log(C_t) + \alpha_l \log(l_{S,t}) + \nu(\omega_t) + \alpha_k \log(k_t) & \text{Stage 2} \\ \log(C_t) + \alpha_l \log(l_{S,t}) + \nu(\omega_t) & \text{Stage 3} \\ \log(C_{S,t}) + \alpha_l \log(l_{S,t}) + \alpha_k \log(k_t) & \text{Stage 4} \\ \log(C_{S,t}) + \alpha_l \log(l_{S,t}) & \text{Stage 5} \end{cases}$$

where

$$\nu(\omega) = \alpha_{\omega,0} + \alpha_{\omega,1}\omega.$$

In order, spouse S derives utility from their consumption, their leisure, their marriage (if married), their decision to bear a child (if childless at the beginning of the period) and the current skill of any potential child.

In addition to this deterministic component of utility, there are three stochastic components. First, when making a work choice, $L_t \in \{0, \frac{1}{2}, 1\}$, the wife receives an additional shock $\epsilon_{L,t}$ to her utility from each choice. Second, when making a fertility decision, $F_t \in \{0, 1\}$, each spouse receives a shock to their utility from that decision, $\epsilon_{F,S,t}$. Third, in every period of marriage, each spouse receives an additional shock $\epsilon_{\omega,S,t}$ to their utility from the marriage. We will assume that each shock is independently and identically drawn from an extreme value distribution of type I with scale parameters $(\sigma_L, \sigma_F, \sigma_\omega)$.

When the child reaches adulthood, each parent receives a terminal payoff equal to $(1-\beta)^{-1}\alpha_k \log(k_t)$.

3.4 Within-Period Timing

In each period decisions are made in a sequence that depends on the life-cycle stage of the couple. The timing of decisions in period t is as follows:

1. First, the stochastic variables $\varepsilon_{Y,H,t}$ and ω_t are realized.
2. If married without children, the shocks $\epsilon_{F,W,t}$ and $\epsilon_{F,H,t}$ are observed *privately* by each spouse, and fertility decisions are made $(F_{W,t}, F_{H,t})$. Fertility occurs $(F_t = 1)$ when there is mutual consent.
3. Next, if married in any state, the shocks $\epsilon_{\omega,W,t}$ and $\epsilon_{\omega,H,t}$ are *privately observed*, and divorce decisions $(D_{W,t}, D_{H,t})$ are made. The outcome (D_t) is determined by the prevailing divorce standard at time t.
4. For any dissolving marriage with children, the custody allocation is drawn from $F_{\bar{\tau}}$.
5. Next, the shock $\epsilon_{L,t}$ is observed, and the wife makes her work decision (L_t).
6. In the final stage, if there are children in the household, time investment decisions are made simultaneously.

3.5 Solution Concept

Spouses in the model behave *non-cooperatively*, using the concept of *Markov Perfect Equilibrium* (MPE).(Pakes and McGuire 2000). The state for each agent in the model is the tuple $(j, t, \tilde{x})$ where j is their life-cycle stage, t is the period of life, and the vector $\tilde{x}$ contains the quality of marriage (ω), the prevailing marriage dissolution standard (Q), the wife's human capital (κ), the age of the child (A), the husband's income shock (ε), and the child's skill level (k):

$$\tilde{x} = \{A_{H,t}, \omega, Q, \kappa, A, \varepsilon, k\}.$$

Naturally, not all of these variables are relevant in all stages of the life-cycle. Let $\tilde{\mathcal{X}}_j$ be the space of relevant state variables in stage j, let $\mathcal{A}_{S,j}$ be the set of actions that spouse S can take given that they are in stage j of the life-cycle, and let $\mathcal{E}_{S,j}$ indicate the space of relevant taste shocks for the decisions made by spouse S in stage j. An MPE is characterized by a pair of *strategy profiles*

$$a^*_{S,t,j} : \tilde{\mathcal{X}}_j \times \mathcal{E}_{S,j} \mapsto \mathcal{A}_{S,j}$$

that solve for each spouse S:

$$a^*_{S,t,j}(\tilde{x}) = \arg\max_{a \in \mathcal{A}_{S,j}} \mathbb{E}_{\epsilon_{\not{S}}} \left\{ u_S(j, \tilde{x}, \epsilon_S, a, a^*_{\not{S}}) + \beta \mathbb{E}_{j', \tilde{x}' | \tilde{x}, a, a^*_{\not{S}}} V_{S,t+1}(j', \tilde{x}') \right\}$$

subject to the constraints, transition rules, and timing defined above. The value function V must adhere to the recursive definition:

$$V_{S,t}(j, \tilde{x}) = \mathbb{E}_{\epsilon} \left[u_S(j, \tilde{x}, \epsilon_S, a^*_S, a^*_{\not{S}}) + \beta \mathbb{E}_{j', \tilde{x}' | j, \tilde{x}, a^*_S, a^*_{\not{S}}} V(j', \tilde{x}') \right]$$

In words, an MPE is a pair of Markov strategy profiles (i.e. profiles that depend only on the state today) that maximize the expected net present value of each spouse's decisions, fixing the strategy of the other spouse and integrating over their privately observed taste shocks.

3.6 Model Solution

The structure of the model allows for important simplifications to its solution, identification, and estimation. Appendix B shows that the value functions in equilibrium are additively separable in child skills and that this leads to (1) closed-form solutions for parental time investments; and (2) an indirect utility function for each spouse that depends only on net household income, labor supply choices, and the age of the child.

Formally, Appendix B shows that the value function for spouse S can be written as:

$$V_{S,t}(j,x) + \alpha_{V,A}\log(k)$$

where $\alpha_{V,A}$ can be expressed recursively as:

$$\alpha_{V,A} = \alpha_k + \delta_{k,A}\beta\alpha_{V,A+1}, \qquad \alpha_{V,19} = (1-\beta)^{-1}\alpha_k \tag{11}$$

Note that these recursive terms are linear in α_k and can therefore be written as:

$$\alpha_{V,A} = \alpha_k \Gamma_A \tag{12}$$

where Γ_A is a function of β and δ_k only.

For a child of age A, the coefficient $\alpha_{V,A+1}$ encodes the net present value return to a log unit increase in skills next period. Discounting this next-period return by β, it follows that a net present value return to a log unit of investment for spouse S is $\Gamma_{S,A} = \beta\delta_{S,A}\alpha_{V,A+1}$. With log preferences in leisure, this results in optimal time investment as

$$\tau^*_{H,t} = \frac{1}{1+\rho(1-\overline{\tau})}\frac{\Gamma_{H,A_t}}{\alpha_l + \Gamma_{H,A_t}} \times 112$$
$$\tau^*_{W,t} = \frac{\Gamma_{W,A_t}}{\alpha_l + \Gamma_{W,A_t}}(112 - L_t 40)$$

Substituting these rules for optimal time investment into utility results in an indirect utility function where the effective weight on non-market time is $\alpha_l + \Gamma_{S,A_t}$. It also results in an additively separable term that encodes how parents value the developmental cost of divorce through reductions in the father's time investment:

$$c(\overline{\tau}, A) = \Gamma_{H,A}\log(1+\rho(1-\overline{\tau})).$$

Since this per-period cost is determined once the custody allocation is drawn, it can be computed in net present value at each age, and is additively separable from the remainder of the value function in divorce.

The coefficient $\alpha_{V,A}$ also determines how parents value the effect that their marital decisions have on their child's development through TFP, expressed as the following components in indirect utility:

$$v(\omega, A) = \beta\alpha_{V,A+1}\log(\psi(A, \omega)) \tag{13}$$

for married couples, and

$$v_d(A) = \beta\alpha_{V,A+1}\log(\psi_d(A)) \tag{14}$$

for divorced couples. These terms act by shifting how parents value marriages of different qualities relative to divorce once they have a child compared to before and after children. In this sense their key behavioral content manifests in differences in divorce dynamics once children are born compared to before and after.

Appendix B collects these terms and explicitly writes indirect utility functions for each spouse and stage of the model. With the indirect utility functions, the remainder of the model can be solved by backward induction across periods and across the within-period decision stages. Appendix B provides further details on the model solution.

3.7 Measurement Assumptions

Two additional assumptions account for avenues of measurement error in the data. First, cognitive skills k are anchored to the probability of getting an item correct on the Applied-Problems (AP) module of the Woodcock-Johnson (WJ) aptitude test. The WJ is designed to avoid ceiling effects by making test items incrementally more difficult. Accordingly, the test score is assumed to be the sum of 60 items:

$$AP = \sum_{s=1}^{60} AP_s$$

where $AP_s = 1$ only if item s is answered correctly. This happens with probability :

$$p(k,s) = \frac{\exp(\log(k) - \gamma_{AP}(-0.5 + s/60))}{1 + \exp(\log(k) - \gamma_{AP}(-0.5 + s/60))}.$$

The parameter γ_{AP} embodies the incremental difficulty of the WJ test items in the Applied Problems Module.[24]

Second, the wife's wages are assumed to be measured with mean zero log-additive measurement error:

$$\log(Y^o_{W,n,t}) = \log(Y_{W,n,t}) + \zeta$$

3.8 Parameterizations

Before moving on to identification and estimation it is necessary to impose some further parameterizations on the model for tractability.

Marriage Quality The matrix Π_ω of transition probabilities for marriage quality is determined by a single parameter π_ω which gives the probability that the marriage quality does not change next period. With equal remaining probabilities $((1 - \pi_\omega)/2)$ the marriage can either move up or down in quality with reflection at the top and bottom of the grid space. To simplify estimation, the model assumes that all marriages begin at the highest quality (and move downs on average quality initially).

Technology Although we allow the factor shares of time inputs to vary arbitrarily by age, for simplicity in estimation we impose the factor share of current skills, δ_k, to be constant.

[24]Modeling the incremental difficulty of the test is necessary to match the age-variance profile of test scores. A simple binomial generates a U-shaped variance profile due to the appearance of ceiling effects, which is not present in the data.

Total factor productivity in marriage is parameterized as

$$\psi(A,\omega) = \exp(\gamma_{\psi,2} + \gamma_{\psi,3}\omega + \gamma_{\psi,4}A)$$

and in divorce:

$$\psi_d(A) = \exp(\gamma_{\psi,1} + \gamma_{\psi,4}A)$$

with an initial condition:

$$k_0 = \exp(\gamma_{\psi,0}).$$

Indirect Utility The terms $v(\omega, A)$ and $v_d(A)$ in (13) and (14) capture the effect that having a child has on periodic payoffs in marriage and divorce. In Section 4 below we argue that these can be identified and estimated directly from fertility and divorce decisions (without imposing the their relationship to underlying production parameters). Accordingly we parameterize an approximate relationship:

$$v(\omega, A) = (\alpha_{v,0,0} + \alpha_{v,0,1}A) + (\alpha_{v,1,0} + \alpha_{v,1,1}A)\omega$$

and

$$v_d(A) = \alpha_{v,d}.$$

While in principle one can estimate these terms non-parametrically, we impose these functional forms to simplify parameter search in estimation.[25] Section 5 will demonstrate that, despite these simplifications, the model can adequately replicate the main empirical patterns in fertility and divorce.

[25]In a fully non-parametric setting, $v_d(A)$ is not separately identified from $v(\omega, A)$. In practice, we found that additional linear terms in age for $v_d(A)$ were only very weakly identified, hence we exclude age from the approximation of $v_d(A)$

4 Identification and Estimation

The model consists of parameters that govern income processes ($\gamma_{Y,W}, \gamma_{Y,H}, \rho_\varepsilon, \sigma^2_\eta$), parameters that govern preferences ($\alpha_l, \alpha_\omega, \alpha_k, \sigma_F, \sigma_\omega, \sigma_L$) and parameters of the technology of skill formation ($\delta_k, \delta_W, \delta_H, \gamma_\psi, \rho, \sigma^2_\xi$). In addition to these primitives, Section 3.5 describes how the ongoing development of children affects parental behavior exclusively through the pair of age-specific preference weights ($\Gamma_{W,A}, \Gamma_{H,A}$), the preference shifters ($v(\omega, A), v_d(A)$) that change the relative utilities in marriage and divorce, and the function $c(\overline{\tau}, A)$ that determines the per-period cost of a particular custody allocation. These terms are functions of combinations of production parameters but can be identified directly from the data which allows for production parameters to be estimated separately from those that govern parental behavior. This section outlines a constructive identification argument along with an estimation routine that follows the logic of identification.

In addition to these, some parameters take externally set values. We set the discount parameter $\beta = 0.95$ and a consumption floor for wives in stage 5 (b) to be \$16.75 per week [26].

4.1 Stage 1: Income Processes

The parameters $\gamma_{Y,H}$ and $\gamma_{Y,W}$ are identified directly by linear regression of log wages on the relevant observables (age for husbands and labor market experience for wives) for each college education type. The persistence of the AR(1) process ε for husbands and the variance of innovations are identified by the relationships:

$$\rho_\varepsilon = \frac{\mathbb{C}(\varepsilon_t, \varepsilon_{t-1})}{\mathbb{V}(\varepsilon_t)}$$

$$\sigma^2_\eta = (1 - \rho^2_\varepsilon)\mathbb{V}(\varepsilon_t)$$

[26] This is the average weekly receipt of food stamps for all participants in the year 2000: `https://www.fns.usda.gov/pd/supplemental-nutrition-assistance-program-snap`

where the latter holds under the assumption that ε is in its stationary distribution. Estimates of these parameters are derived from straightforward sample analogs of these relationships using linear regression and sample variances and covariances of the residuals from these regressions.

4.2 Stage 2: Preference weights

Among married couples, the optimal time use by each spouse with child of age A obeys:

$$\phi_{S,A} = \frac{\tau_{S,t}}{112 - L_{S,t}40} = \frac{\Gamma_{S,A}}{\alpha_l + \Gamma_{S,A}}$$

And hence the preference weight $\Gamma_{S,A}$ is identified up to the scale of α_l for each spouse and age:

$$\Gamma_{S,A}/\alpha_l = \frac{\phi_{S,A}}{1 - \phi_{S,A}} \tag{15}$$

The corresponding estimate of this ratio for each A is derived from age-specific sample means of $\phi_{S,A}$ among married couples for children of age A.

4.3 Stage 3: Time penalty from custody allocations

Fixing the age of the child, optimal time investment implies:

$$\frac{\mathbb{E}[\tau_{H,t}|A_t = a, M_t = 0]}{\mathbb{E}[\tau_{H,t}|A_t = a, M_t = 1]} = \mathbb{E}_{F_{\overline{\tau}}}\left[\frac{1}{1 + \rho(1 - \overline{\tau})}\right]$$

where $M_t \in \{0, 1\}$ indicates the marital status of the husband. Conditional on the population distribution of custody allocations, ρ is identified from the average ratio of time investment between divorced and married fathers (controlling for age). In order to control for unobserved selection into divorce by fathers, this ratio is estimated in a regression framework with individual fixed effects. The population distribution of custody allocations ($F_{\overline{\tau}}$) is identified and estimated from the auxiliary CPS data described in Section 2 and Appendix C.

4.4 Stage 4: Preferences

The remaining parameters that govern fertility, divorce, and work behavior are the coefficients $(\alpha_l, \alpha_\omega)$, the preference shock dispersion parameters $(\sigma_F, \sigma_\omega, \sigma_L)$, the utility shifters with child age $(v(\omega, A), v_d(A))$, and the transition matrix for ω (Π_ω).[27] These parameters are collectively identified by moments describing the dynamic patterns of fertility, divorce, and female labor supply.

In practice, the estimation routine targets (1) the fractions of full and part-time work for married and divorced wives (to identify α_l and σ_L); the fraction of marriages that end in divorce within 5, 10, and 15 years (to identify α_ω, π_ω and σ_ω); the fraction of couples who have their first child within 2, 4, and 6 years of marriage (to identify α_F and σ_F); the fraction of couples who divorce before any child; and the fraction who divorce before their first child reaches age 5, 10, and 15 (to identify the coefficients that approximate $v_d(A)$ and $v(\omega, A)$).

4.5 Stage 5: Child Skill Formation

With behavioral parameters determined, the identification of the technology of skill formation can be considered by exclusively examining child skill outcomes, conditional on parental behavior. Estimation follows suit by fixing a simulated dataset of fertility, divorce, work, and investment decisions and choosing production parameters $(\delta_W, \delta_H, \delta_k, \gamma_\psi, \sigma_\xi, \gamma_{AP})$ to match statistics concerning measured skills in the data.

Given that the model allows for differences in input factor shares by age (A) and by spouse (S), some of the behavioral cross-equation restrictions are useful for pinning down the relative importance of inputs at different ages. To begin, note from the definition of the recursive coefficients in (11) that:

$$\frac{\delta_{H,A}}{\delta_{W,A}} = \frac{\Gamma_{H,A}}{\Gamma_{W,A}}$$

[27]Note that conditional on a choice of α_l, the preference weights $(\Gamma_{H,A}, \Gamma_{W,A})$ and custody time costs $(c(\overline{\tau}, A))$ are already determined from Stages 2 and 3.

and so the values of $\Gamma_{S,A}$ that are identified from time investment patterns in (15) are sufficient to also identify the relative factor shares for each spouse at age A. The behavioral restrictions in (11) and (12) also pin down the relative importance of factor shares at difference ages:

$$\frac{\delta_{S,A}}{\delta_{S,0}} = \frac{\Gamma_{S,A}\Gamma_0}{\Gamma_{S,0}\Gamma_{A+1}}$$

where each Γ_A is a function of δ_k only.

Thus, among the input factor shares, only the self-productivity parameter (δ_k) and the factor share of mother's time at age 0 ($\delta_{W,0}$, which pins down the scale of all time input factor shares) need be identified from test score data. This is achieved by matching in simulated data the coefficients of a regression of test scores on mother's time investment along with lagged test scores:

$$AP_{n,t+5} = \beta_0 + \beta_1 \log(\tau_{W,n,t+5}) + \beta_2 AP_{n,t} + \zeta_{n,t}.$$

where $AP_{n,t}$ is the Applied Problems score for child n from the Woodcock-Johnson battery of test scores. $AP_{n,t+5}$ is the same score measured five years later in the next CDS wave.

According to the model, average differences between children of divorced parents and children of married parents are determined by a combination of time inputs and TFP. These differences may accumulate before divorce due to lower TFP in lower quality marriages, or after divorce due to lesser time inputs and/or lower productivity outside of marriage. Thus, Stage 5 chooses parameters to target the average test score by age for each of three groups: children of parents who never divorce, will divorce, or are already divorced. The variation in exposure to divorce (when paired with estimates of the importance of time inputs) identifies the parameters that determine TFP in each marital state (γ_ψ).

Finally, to identify the standard deviation of the shocks to TFP (σ_ξ), along with the incremental difficulty of the Applied Problems module (γ_{AP}), Stage 5 also targets the age profile of the standard deviation in raw AP test scores.

5 Estimation Results

We present estimates of the model from the five stage estimation process outlined in Section 4. Table 2 reports estimates of income process parameters for both spouses. Unsurprisingly, college-educated wives enjoy higher returns to full-time equivalent experience and college-educated husbands enjoy higher life-cycle growth in earnings. Income shocks for husbands exhibit moderate persistence ($\rho_{\varepsilon} = 0.62$) with substantial time-varying risk ($\sigma_{\eta} = 0.5$).

Table 3 reports the model's fit of the dynamic patterns of fertility, divorce, and labor supply. The model slightly over-predicts wives' rate of part-time work after divorce (26 percent vs 19 percent in the data), and slightly under-predicts the number of marriages that end without the birth of any child (2 percent vs 8 percent in the data). Furthermore, the model seems to exhibit a slightly flatter profile in fertility with respect to the duration of marriage. Having acknowledged these discrepancies, the model does a good job of matching the relevant overall life-cycle patterns of these decisions with relatively few parameters.

Table 4 reports the preference parameter estimates. We estimate that parents place 0.39 times the log consumption preference weight on the log of leisure. Of course, the need to compare leisure hours with consumption dollars implies that this preference weight contains both scaling and relative preference information, and is therefore difficult to interpret alone. Estimates of the parameters governing the utility of marriage indicate that, on average, spouses value the married state more than the divorced state; moreover, the value of the married state is steeply increasing in marriage quality ω. Estimates of the marriage quality transition process and the shock to the utility of marriage show us that, in addition to the couple's current state, the model is relying on both the evolution of the persistent utility of marriage and idiosyncratic shocks to the value of marriage to precipitate divorce, with a sizable reliance on idiosyncratic shocks. Similarly, in addition to the potential welfare gain from parenthood, the fertility decision is guided in meaningful part by the dispersion in the idiosyncratic shock to the value of fertility. The wife's labor supply process, however, is shaped largely by the wage, rather than by idiosyncratic shocks to her taste for work. The preference parameters are, of course, better understood through simulation under existing and counterfactual conditions, which we describe in Table 3, above, and in the following section.

Table 5 reports estimates of the model's production parameters from Stages 3 and 5 of the estimation procedure. This table excludes the factor share of time inputs, which are represented in Figure 8 for simplicity. Several particular estimates deserve discussion. The estimate of $\hat{\rho} = 1.46$ implies a 47 percent decrease in paternal time investment under a 50-50 physical custody split. This mechanically matches the observed reduction in paternal time investment in the data, net of individual fixed effects. The relatively high standard deviation of this parameter reflects the statistical uncertainty in the data regarding the magnitude of this decrease. This is important to note because it will lead to resulting uncertainty in the skill impacts of particular custody arrangement counterfactuals, as well as the contribution of paternal time inputs to observed skill gaps.

Table 5 suggests that TFP is indeed increasing in marital quality ($\gamma_{\psi,3} = 0.11$). Figure 5 plots the implied values of TFP in marriage relative to divorce ($\gamma_{\psi,2} + \gamma_{\psi,3}\omega - \gamma_{\psi,1}$) along with confidence intervals on these quantities to get a better sense of when differences in TFP may appear. Results indicate that TFP in marriage is significantly lower than TFP in divorce for all but the very highest marriage quality. (Recall that all marriages are assumed to begin at the highest marriage quality). Point estimates of TFP decline with worsening marriage quality.

As Section 4 and Figure 3 illustrated, the results on TFP reflect observed differences in children's test scores for each of three groups: (1) Divorced; (2) Never Divorced; and (3) Will Divorce. Figure 6 indicates that these differences are relatively stable over time and that the model fits this pattern well.

Figure 8 depicts the age-patterns in the factor shares of mothers' and fathers' time inputs. Section 4 describes how only the scale of these parameters is pinned down by test scores in Stage 5 of the estimation, while the relative magnitudes are determined by measured patterns in average time investments combined with cross-equation restrictions implied by optimal investment behavior. Consistent with these patterns, Figure 8 shows that both factor shares are large at young ages and decrease as children enter adolescence and approach maturity. Factor shares for the mother, in terms of point estimates, are substantially larger when the child is very young and then decline more steeply as the child ages to reach a similar level to the father's factor shares by the time the child is seven. The overall magnitude of these factor shares implies a major role for time inputs, which will be explored in the next section.

5.1 Decomposing Skill Gaps

A decomposition exercise provides a useful approach to study the quantitative significance of the production parameter estimates. First note that final skills for any child n can be written as a discounted sum of inputs and TFP:

$$\log(\theta_{n,19}) = \sum_{a=0}^{18} \delta_k^{18-a} \left(\delta_{W,a} \log(\tau_{W,n,a}) + \delta_{H,a} \log(\tau_{H,n,a}) + \log(\psi_{n,a}) + \xi_a\right) \quad (16)$$

Due to this additive structure, the difference in final skills for any two groups or under any counterfactual scenario can be decomposed into three sources of interest: (1) Mothers' time inputs; (2) Fathers' time inputs; and (3) TFP.[28]

Table 6 reports such a decomposition for the difference in average cognitive skill between children whose parents remain married for the duration of their skill development and children whose parents divorce before they reach age 18. The distribution of skills is normalized to have a standard deviation of 1. According to the estimated model, skills are 5 percent of a standard deviation higher in the divorced group, but the 90 percent confidence interval is essentially symmetric around zero. This null effect is determined by two countervailing forces: marriages that stay intact are associated with lower TFP relative to divorce (21

[28]It is worth noting at this point that, given the properties of the analytical solution for time investment at each value of the state vector, the time investments of parents in surviving marriages will not be affected by movements in the likelihood of subsequent divorce.

percent of a standard deviation), while the reduction in parental time investments leads to losses in divorce of about 17 percent of a standard deviation.

This decomposition highlights the important role played by time investments in the determination of cognitive skill outcomes. Although the model predicts an increase in maternal labor supply - with a proportional reduction in time investment - as a result of divorce, the impact of this change in investment is small relative to the large reduction in fathers' time. The pattern of the contributions of parents' time to the skill differences between children of marriage and divorce in this decomposition exercise echoes our motivating evidence in Section 2. Figure 2 near the start of the paper depicts the very small decline in mothers' time investment from marriage to divorce, and the large decline in fathers' time from marriage to divorce, that we observe in our PSID sample.

The decomposition evidence should be offered with some careful caveats: the relative magnitudes of the contributions of TFP, mothers' time, and fathers' time derived from the model simulation arise from a combination of measured investment behavior and functional form assumptions regarding the substitutability of these inputs, in combination with children's measured skill progress.

6 Policy Experiments

In this section we study three elements of family law that are particularly relevant to couples' choices of fertility, child investment, and the continuation of a marriage, and that may shape children's cognitive development and family members' welfare. The first is a comparison between a mutual consent divorce standard and a unilateral one. The second policy change we study is an increase in the rate of child support taxation on fathers in the divorced parenting state, π. We increase π from the 15 percent child support tax rate assumed in our estimation first to 30 percent, which is the upper bound of plausible tax rates based on state child support guidelines and relevant family structures, and second to a comparatively extreme, and likely politically infeasible, 50 percent.[29] The third is a comparison between a legal regime in which all mothers receive full physical custody to one in which all custody arrangements are a 50-50 split.

We use the estimated model to simulate outcomes for each couple in the sample under alternative policy arrangements, and we assume that these alternative policies remain in effect from the first year of marriage until their children reach maturity. We consider the ex-ante welfare impacts of each counterfactual for each spouse (as measured by the percentage increase in consumption that would deliver equivalent changes in discounted present values at the beginning of marriage), the effects on the overall divorce and fertility rate, the effects on the wife's (endogenous) wages, and the effects on the cognitive skill of any children who happen to be born under the counterfactual. In the latter case we decompose, as in Section 5, these skill differences into contributions from TFP, maternal time inputs, and paternal time inputs. In the case of TFP, any such differences are driven by differential selection into divorce *and* fertility by marriage quality.

Given the importance of selection into fertility and divorce in the following simulated policy responses, this may be an opportune moment to note one real-world margin of selection that the estimated model omits. Where the model's marriages are formed exogenously before the initial period of the model, in reality marriages form endogenously, possibly in response to anticipated parameters of divorce policy. In one straightforward example, moving from a mutual consent to a unilateral dissolution standard makes divorce easier. Couples may respond to easier divorce by imposing a lower threshold for marriage, and, as a result, couples with lower match quality may select into marriage. Lower quality marriages will have ramifications for all of our simulated outcomes: spouses' welfare, fertility, children's skill, and divorce. Each of our simulated effects of policy must be interpreted with an awareness of the possibility that the policy change would affect marriage formation.

[29] Recall that the 15 percent rate was drawn from the modal rate generated by state child support guidelines for single-child families for the time period and sample that we study.

6.1 Divorce Standards

We begin with a comparison of statistics in which all marriages are moved either to a mutual consent divorce regime (in which both spouses must consent to the divorce) or to a unilateral divorce regime. Table 7 reports the results for each counterfactual. Most striking are the large welfare gains to both spouses from moving from a mutual consent to a unilateral regime. This is intuitive: when both spouses are receiving uncorrelated shocks to the utility they derive from the marriage, there is value to being able to unilaterally dissolve the marriage when a large negative shock is drawn. Estimates in Table 4 indicate that the dispersion of these shocks is large, and so larger welfare effects are not surprising.

As expected, there is an increase in the rate of divorce when moving from a completely mutual consent to a completely unilateral regime (almost 10 percent in total).[30] Accompanying this increase is an increase in fertility, an increase in skill outcomes, and an increase in wives' wages (who work more both after and in anticipation of divorce). Based on the reported 90 percent confidence intervals, only the latter of these increases could be viewed as conventionally statistically significant.

Although it may be surprising that unilateral divorce does not lead to a statistically significant increase in divorce, we emphasise that these are the effects of the policy when it is known to all individuals *ex ante*, not the immediate effect of a regime change as is most commonly studied empirically (Friedberg 1998, Wolfers 2006). Importantly, the estimated effects here also incorporate changes in fertility decisions which leads to changes in divorce decisions downstream. In Appendix A.2 we examine individual bootstrap trials in the counterfactual to show that decreases in divorce are indeed associated with decreases in fertility, and that this is in turn associated with statistical uncertainty in the dispersion of taste shocks for marriage (σ_ω).

[30]Friedberg (1998), for example, estimates a six percent increase in divorce rates for the full married US population as a result of states' gradual (and partial) move to the unilateral standard.

6.2 Child Support

Following the logic of the previous section, we may pursue other policy initiatives that increase the ease with which couples can harmoniously divorce. If mothers prefer to stay in low quality marriages because of their low outside option when living alone, then this may act as a disadvantage to the child. It is reasonable to ask if increasing financial support to mothers after divorce may enable low quality marriages to dissolve and improve child outcomes. Furthermore, more generous child support has a negative income effect on mothers' labor supply which would lead to an increase in mothers' time inputs after divorce, leading to *ceterus paribus* skill improvements.

To evaluate this, we consider two large increases in the rate of child support, π, paid by the father in our model. In particular, we consider taxation rates of 30 percent and 50 percent, which are much higher than the 15 percent used in our baseline analysis. Table 8 shows the aggregate results of these policy changes. Contrary to the intuition that motivated this policy, we find a decrease in divorce rates that is downstream from a decrease in the overall rate of fertility. Overall effects on welfare and child skills appear statistically ambiguous.

Once again, we use the decomposition of equation (16) to look for answers to this puzzle. While we do see increases in mothers' time inputs as expected, the contributions of TFP to average skill outcomes are negative. This is due to a *decrease* in the divorce rate driven by husbands who prefer to remain married given the increased child support burden. Thus, increasing child support while holding other components of the legal framework fixed increases disadvantageous selection by marriage quality and leads to skill losses.

To further understand this result, and emphasize the importance of prevailing dissolution standards on the effect of these child support rates, we recalculate the effect of an increase in the child support rate from 15 to 30 percent when all households face either a mutual consent or a unilateral dissolution standard. The results in Table 9 indicate stark differences in the effect of child support changes under these two regimes. Although fertility and divorce rates decrease in both cases, welfare impacts for wives and skill impacts move in opposite directions. Child support makes divorce more appealing for wives and less appealing for husbands, making them less likely to want to have children in the first place (leading to small reductions in the fertility rate). Since having children makes divorce more likely, this leads to reductions in the divorce rate also.

The mutual consent standard provides a source of commitment that drives differences in the welfare impact of this policy change. Since divorce can only occur when the husband consents, child support leads to welfare gains for the mother only in model states where the husband is also more likely to also want divorce. Although there are reductions in fertility, reductions in the value of this choice are more highly correlated across spouses under the mutual consent regime. Thus, increasing the value of divorce for women can still lead to ex-ante welfare gains. In contrast, under unilateral divorce, wives cannot commit to staying in the marriage in states where the husband would prefer this outcome. Hence, reductions in the value of fertility for the husband are less well correlated with those for the wife, and this results in ex ante welfare losses.

Similarly, the decomposition of skill outcomes shows that most of the difference in impacts is driven by selection on marital TFP, with unilateral divorce leading to a large increase in *disadvantageous selection*. This is also related to differences in the composition of marriage quality that emerge when there are differences in the ability of wives to commit.

More than the specific outcomes of these counterfactuals, the results here emphasise the importance of dynamic selection in determining the response to changes in marital law, well as the important role played by commitment under different legal frameworks for the effects of other policy changes. In particular, it is clear that changes in who decides to have children can be just as important (if not more important) for average skill outcomes and welfare as ex-post changes in behavior after fertility.

6.3 Custody Allocations

Finally, we examine the role of the third family law policy dimension of interest, custody allocation. By custody allocation in this empirical implementation, we refer to the share of physical placement with the child. This can be contrasted with legal custody, which confers decision-making authority to the parent but does not necessarily reflect the parent's share of time with the child. Recall that we have parameterized the child cognitive ability production function so that a key characteristic of custody is the ease with which the father can invest time in the child's skill development. The baseline custody allocation in the estimated and simulated model from Section 5 involves a distribution of custody allocations.

We simulate the response of mean family outcomes to a move from the baseline distribution to sole maternal custody, with results reported in Table 10. On average, we estimate that wives and husbands would each require a two percent increase in consumption to be made indifferent between the baseline and sole maternal custody. However, the dispersion in welfare changes is quite broad for both spouses. Fertility, divorce, child skill, TFP, and wives' wages decline, though insignificantly, with the move to sole maternal custody. Mothers' time is little changed by either large custody reform in Table 9. However, the father's time with the child decreases substantially from the baseline to sole maternal custody.

Perfectly equal custody for both mother and father, however, appears to improve or leave unchanged all relevant family outcome dimensions. Parent welfare is increased substantially at the mean, but with a wide dispersion. Fertility and divorce remain approximately unchanged. Mean child skill improves, driven largely by an increase in the father's time investment, but also by a change in TFP, reflecting a change in the distribution of marriages that produce a child and that enter divorce. In both custody experiments, the simulated impact on child skill is noisy. This is related to the imprecision with which we are able to estimate the cost of physical custody in time investment (ρ).[31]

[31] Chen and Logan (2020) estimate a negative effect of gender-neutral custody laws on children's educational attainment. They emphasize the role of the custody threat in bargaining over financial settlements, a mechanism not captured by our model.

7 Conclusion

We have developed and estimated a model that allows for strategic behavior between parents in making fertility, child quality investment, and divorce decisions. An important component of the behavioral model is the family law environment, which has a large impact on the rewards attached to the marital states and, in turn, the returns to investment in child quality. We use data from the PSID and the PSID-CDS to estimate model parameters using a Method of Simulated Moments estimation procedure. We find that the parameter estimates are roughly in accord with our priors, and that the correspondence between simulated and sample moments is generally close (though in some subsets of the moments measuring the profile of fertility the correspondence is merely adequate).

The most important contribution of our work is to the understanding of the dynamic relationship between divorce decisions and the evolution of fertility and child quality, and the dependence of this process on family law parameters. While there is a well-established empirical relationship between children's outcomes and the characteristics of the households in which they live, we have attempted to disentangle the simultaneous relationships among divorce, fertility, and child development using a behavioral model of these decisions. To our knowledge, this is one of the first studies to link the family law environment to the fertility decisions of intact families, and, in some instances, we find the link to be substantial. [32] While our estimated model is based on a number of restrictive and, ultimately, untestable assumptions, our view is that this type of framework is the only way to begin to understand the complex dynamics present within married households.

Our model estimates are useful in their own right, as they allow us to test predictions including the relative child cognitive development productivity of family settings featuring strong marriages, ex post weak marriages, and divorce. Our estimates indicate that the child cognitive ability total factor productivity of low match quality marriages lies below that of divorced families. This result aligns with descriptive evidence demonstrating the lowest

[32] The other that we are aware of is Aizer and McLanahan (2006).

test score averages for children of households that will eventually divorce. In addition, our parameter estimates are able to show us the levels and age dependence of the child cognitive development productivity of their mother's time inputs, their father's time inputs, and their own current cognitive ability. We find that productivity of mothers' and fathers' time inputs to child development is initially high but declines steeply as the child ages toward independence. The child's own cognitive ability is estimated to have a particularly high level of productivity in generating future cognitive ability gains, and this productivity is substantially more persistent as the child ages than are the various parental inputs.

We have conducted investigations of how substantial changes in the parameters characterizing the family law environment - those reflecting the ease with which marriages may be dissolved, the child support transfers between parents in the divorce state, and fathers' time with the child in divorce - impact fertility, child outcomes, and the distribution of parental welfare. In line with the evidence from the parameter estimates above regarding the low productivity of a weak marriage in producing child cognitive ability gains, counterfactual policy simulations in which we make divorce more difficult to obtain by implementing a mutual consent divorce standard indicate a meaningful decline in child cognitive attainment under legal regimes that restrict access to divorce. Further, the mutual consent standard decreases average fertility in marriage. Our estimated model, in various ways, describes a damaging effect on children of remaining in a low match quality marriage that is approaching or barred from divorce.

In addition, we investigate the effects of two substantial increases in the child support paid by fathers to mothers in the event of divorce on divorce, fertility, and children's cognitive achievements. An increase from a 15 percent child support tax rate to a 30 or 50 percent child support tax rate decreases both the divorce rate and the share of marriages that experience a birth. It also changes the distribution of families that choose to have children, and weakens the quality of ongoing marriages in the resulting children's households. On net, we see a decline in the mean child skill total factor productivity of households with children. The child support increases do, however, enhance the resources available to mothers who engage in divorced parenting, as well as both parents' time investment in children. At the mean, our simulation indicates that an increase in child support transfers decreases the average cognitive achievement of the children who are born, but this simulation result masks large offsetting changes in marriage quality and parental time inputs to child development, and is particularly noisy.

Finally, counterfactual policy simulations in which we shift child custody (placement) standards first to sole maternal custody and next to 50-50 custody predict limited fertility and divorce responses, almost no change in mothers' time investment, substantial shifts in fathers' time investment (increasing with custody), and a noisy but somewhat positive association between paternal custody and child cognitive attainment.

In combination, the evidence generated by our manipulation of divorce standards, child placement, and divorce-state transfers does generate one overarching insight regarding the regulation of divorced parenting: while prior descriptive evidence indicates that children of divorce fare worse on several dimensions that children of marriage, it is not the case that designing family law to minimize divorced parenting unambiguously benefits children. More difficult divorce standards, such as mutual consent divorce requirements, may substantially decrease divorce, but they also lead to a decrease in both total fertility and the average cognitive attainment of the realized population of children. Similarly, (much) higher transfers to the custodial parent in divorce are projected to decrease divorce, but they also lead to declines in fertility and children's cognitive attainment. Our model allows the marginal marriage affected by family law to differ meaningfully from the average marriage. In addition, and importantly, our approach allows stable and unstable marriages to differ in their consequences for children's cognitive development. This approach of formalizing marriage heterogeneity in terms of both marriage stability and child ability production allows us to separate the social goals of raising children in stable, high-investment households and protecting children from the developmental damage that appears to arise from prolonging low match quality marriages.

8 Data Availability

Code replicating the tables and figures in this article can be found in Brown et al. (2024) in the Harvard Dataverse, `https://doi.org/10.7910/DVN/12YTIN`.

References

Aiyagari, S. Rao, Jeremy Greenwood, and Nezih Guner. 2000. On the State of the Union. *Journal of Political Economy* 108, no. 2:213-244.

Aizer, A. and S. McLanahan (2006), "The Impact of Child Support Enforcement on Fertility, Parental Investments, and Child Well-Being," *Journal of Human Resources* 41, 28-45.

Bernal, Raquel. 2008. The Effect of Maternal Employment and Child Care on Children's Cognitive Development. *International Economic Review* 49, no. 4:1173-1209.

Bernal, Raquel, and Michael Keane. 2011. Child Care Choices and Children's Cognitive Achievement: The Care of Single Mothers. *Journal of Labor Economics* 29, no. 3:459-512.

Brown, M., Flinn, C., and Mullins, J. 2024. "Replication Data for: Family Law Effects on Divorce, Fertility and Child Investment'," Harvard Dataverse, `https://doi.org/10.7910/DVN/12YTIN`.

Brien, Michael., Lee Lillard, and Steven Stern. 2006. Cohabitation, Marriage and Divorce in a Model of Match Quality. *International Economic Review* 47, no. 2:451-494.

Browning, Martin, and Pierre-Andre Chiappori. 1998. Efficient Intra-Household Allocations: a General Characterization and Empirical Tests. *Econometrica* 66, no. 6:1241-1278.

Bruze, Gustaf, Michael Svarer, and Yoram Weiss. 2015. The Dynamics of Marriage and Divorce. *Journal of Labor Economics* 33, no.1:123-170.

Cano, Tomás, and Pablo Gracia. 2022. The Gendered Effects of Divorce on Mothers' and Fathers' Time with Children and Children's Development Activities: A Longitudinal Study. *European Journal of Population* 38, no. 5: 1277-1313.

Case, Anne, I-Fen Lin, and Sara McLanahan. 2003. Explaining Trends in Child Support: Economic, Demographic, and Policy Effects. *Demography* 40, no. 1:171-189.

Caucutt, Elizabeth, Nezih Guner, and John Knowles. 2002. Why do Women Wait? Matching, Wage Inequality, and the Incentives for Fertility Delay. *Review of Economic Dynamics* 5, no. 4:815-855.

Caucutt, Elizabeth M., Lance Lochner, Joseph Mullins, and Youngmin Park. 2020. Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments. NBER Working Paper No. 27838, NBER, Cambridge, MA.

Chen, Yang, and Trevon Logan. 2020. Is the Best Interest of the Child Best for Children? Educational Attainment and Child Custody Assignment. *Southern Economic Journal* 86, no. 3:1041-1080.

Chiappori, Pierre-Andre, Bernard Fortin, and Guy Lacroix. 2002. Marriage Market, Divorce Legislation and Household Labor Supply. *Journal of Political Economy* 110, no. 1:37-72.

Ciacci, Riccardo. 2023. On the economic determinants of prostitution: marriage compensation and unilateral divorce in U.S. states. *Review of the Economics of the Household* 21, no. 3:941–1017.

Cunha, Flavio, and James Heckman. 2007. The Technology of Skill Formation. *American Economic Review* 97, no. 2:31-47.

Cunha, Flavio, James Heckman, Lance Lochner, and Dimitriy Masterov. 2006. Interpreting the Evidence on Life Cycle Skill Formation. *Handbook of the Economics of Education*, edited by Eric A. Hanushek and Frank Welch. Amsterdam: North-Holland, 697-812.

Cunha, Flavio, James Heckman, and Susan Schennach. 2010. Estimating the technology of Cognitive and Non Cognitive Skill Formation. *Econometrica* 78, no.3:883-931.

Del Boca, Daniela, and Christopher Flinn. 1995. Rationalizing Child-Support Decisions. *American Economic Review* 85, no. 5:1241-1262.

Del Boca, Daniela, and Christopher Flinn. 2011. Endogenous Household Interaction. *Journal of Econometrics* 166, no. 1:49-65.

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2014. Household Choices and Child Development. *Review of Economic Studies* 81, no. 1:137-185.

Del Boca, Daniela, Christopher Flinn, Ewout Verriest, and Matthew Wiswall. 2023. Parenting with Patience: Parental Incentives and Child Development. CEPR Discussion Paper No. 18424, CEPR Press, Paris and London.

Eckstein, Zvi, and Osnat Lifshitz. 2011. Dynamic Female Labor Supply. *Econometrica* 79, no. 6:1675-1726.

Flinn, Christopher (2000). Modes of Interaction Between Divorced Parents. *International Economic Review* 41, no. 3:545-578.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. 2023. *IPUMS CPS: Version 11.0 [dataset]* Minneapolis, MN: IPUMS, 2023. https://doi.org/10.18128/D030.V11.0

Friedberg, Leora. 1998. Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data. *American Economic Review* 88, no. 3:608-627.

Gruber, Jonathan. 2004. Is Making Divorce Easier Bad for Children? The Long Run Implications of Unilateral Divorce. *Journal of Labor Economics* 22, no. 4:799-833.

Haveman, Robert., and Barbara Wolfe. 1995. The Determinants of Children's Attainments: A Review of Methods and Findings. *Journal of Economic Literature* 33, no. 4:1829-1878.

Hopkins, Kenneth, and Glenn Bracht. 1975. Ten-Year Stability of Verbal and Nonverbal IQ Scores. *American Educational Research Journal* 12, no. 4:469–477.

Kranz, Daniel Fernández, Jennifer Roff, and Hugette Sun. 2021. Can economic incentives for joint custody harm children of divorced parents? Evidence from state variation in child support laws. *Journal of Economic Behavior and Organization* 189:1-27.

Liu, Haiyong, Thomas Mroz, and Wilbert van der Klaauw. 2010. Maternal Employment, Migration and Child Development," *Journal of Econometrics* 156, no. 1:212-228.

Lundberg, Shelly, and Robert Pollak. 1994. Non-cooperative Bargaining Models of Marriage. *American Economic Review* 84, no. 2:132-137.

Mnookin, Robert, and Lewis Kornhauser. 1979. Bargaining in the Shadow of Law: The Case of Divorce. *Yale Law Journal* 88, no. 5:950-997.

Mullins, Joseph. 2021. Designing Cash Transfers in the Presence of Children's Human Capital Formation. *Working Paper*

Mullins, Joseph. 2023. A Structural Meta-analysis of Welfare Reform Experiments and Their Impacts on Children *Working Paper*

Nixon, Lucia. 1997. The Effect of Child Support Enforcement on Marital Dissolution. *Journal of Human Resources* 32, no. 1:159-181.

Pakes, Ariel, and Paul McGuire. 2001. Stochastic Approximation for Dynamic Analysis: Markov Perfect Equilibrium and the 'Curse' of Dimensionality. *Econometrica* 69, no. 5:1261-1281.

Stevenson, Betsey, and Justin Wolfers. 2006. Bargaining in the Shadow of the Law: Divorce Laws and Family Distress. *The Quarterly Journal of Economics* 121, no. 1:267–288.

Sun, Hugette. 2008. Child Support Policy and Divorce. Manuscript, Bureau of Labor Statistics.

Tartari, Melissa. 2015. Divorce and the Cognitive Achievement of Children. *International Economic Review* 56, no. 2:597-645.

Todd, Petra, and Kenneth Wolpin. 2003. On the Specification and Estimation of the Production Function for Cognitive Achievement. *Economic Journal* 113, no. 485:F3-F33.

Todd, Petra, and Kenneth Wolpin. 2007. The production of cognitive achievement in children: Home, school and racial test score gaps. *Journal of Human Capital* 1, no. 1:91-136.

Verriest, Ewout. 2024. Household Labor Supply, Child Development and Childcare Policies in the United States. Manuscript, Pennsylvania State University.

Voena, Alessandra. 2015. Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples? *American Economic Review* 105, no. 8:2295-2332.

Wolfers, Justin. 2006. Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results. *American Economic Review* 96, no. 5:1802-1820.

Table 1: Descriptive Statistics

Divorced by 2018	0.21
Have Child	0.86
College: Wife	0.28
College: Husband	0.26
Age at Marriage: Wife	22.47
Age at Marriage: Husband	24.33
Years to First Birth	2.43
Mutual Consent law in year of marriage	0.42
Average Wage: Wife ($/hour)	14.98
Average Wage: Husband ($/hour)	22.46
Sample Size	2525

NOTE.—This table reports descriptive statistics on Husbands and Wives from the sample of PSID Marriages between 1975 and 1997. Wages are deflated to year 2000 USD. College degree attainment is defined as four years or more of post-secondary education.

Table 2: Estimates of Income Process Parameters

	Wife		Husband	
	Non-College	College	Non-College	College
Const.	1.91	2.34	1.10	-0.08
	(0.02)	(0.03)	(0.11)	(0.23)
FT Exp (κ)	0.057	0.069		
	(0.004)	(0.005)		
FT Exp ^{2} (κ^2)	-0.001	-0.001		
	(0.000)	(0.000)		
Age			0.072	0.148
			(0.007)	(0.012)
Age ^{2}			-0.001	-0.002
			(0.000)	(0.000)
ρ_ϵ			0.62	0.62
			(0.02)	(0.02)
σ_η			0.50	0.50
			(0.01)	(0.01)

NOTE.—This table reports estimates of the parameters governing the income processes for Wives and Husbands. Parentheses report standard errors from a nonparametric bootstrap routine with 50 trials.

Table 3: Model Fit from Stage 4

Moment	Data	Model
FT - Married	0.50	0.47
PT - Married	0.29	0.31
FT - Divorced	0.66	0.72
PT - Divorced	0.19	0.25
Divorce < 5 years	0.03	0.02
Divorce < 10 years	0.09	0.07
Divorce < 15 years	0.14	0.13
Birth < 2 years	0.38	0.46
Birth < 4 years	0.65	0.63
Birth < 6 years	0.78	0.72
Time to Divorce - Time to Birth < 0 years	0.08	0.06
Time to Divorce - Time to Birth < 5 years	0.30	0.30
Time to Divorce - Time to Birth < 10 years	0.61	0.61
Time to Divorce - Time to Birth < 15 years	0.83	0.84

NOTE.—This table reports the model's fit of moments from stage 4 of the estimation procedure, which estimates behavioral parameters governing work, fertility, and divorce.

Table 4: Preference Parameter Estimates

Parameter	Estimate	(Std Error)
α_l	0.40	(0.02)
σ_L	0.54	(0.03)
$\alpha_{\omega,0}$	0.75	(0.06)
$\alpha_{\omega,1}$	4.10	(0.20)
σ_ω	7.97	(0.67)
π_ω	0.91	(0.01)
$\alpha_{\nu,0,0}$	3.10	(0.10)
$\alpha_{\nu,0,1}$	-0.38	(0.01)
$\alpha_{\nu,1,0}$	0.60	(0.07)
$\alpha_{\nu,1,1}$	0.34	(0.02)
$\alpha_{\nu,D}$	-1.80	(0.11)
σ_F	3.76	(0.10)

NOTE.—This table reports estimates of preference parameters from Stage 4 of the estimation procedure. Parentheses report standard errors from a nonparametric bootstrap routine with 50 trials.

Table 5: Production Parameter Estimates

Parameter	Estimate	(Std Error)
$\gamma_{\psi,0}$	-23.30	(2.75)
$\gamma_{\psi,1}$	1.18	(0.15)
$\gamma_{\psi,2}$	0.99	(0.18)
$\gamma_{\psi,3}$	0.10	(0.03)
$\gamma_{\psi,4}$	-0.08	(0.02)
δ_k	0.95	(0.03)
σ_ξ	0.53	(0.12)
γ_{AP}	16.71	(2.12)
ρ	1.17	(0.58)

NOTE.—This table reports estimates of the model's production parameters from Stages 3 and 4 of the estimation procedure. Parentheses report standard errors from a nonparametric bootstrap routine with 50 trials.

Table 6: Decomposition of Skill Differences by Parental Divorce Status

Input	Estimate	Conf. Interval
Total	-0.04	[-0.19 0.09]
TFP	-0.22	[-0.58 -0.04]
Mother's Time	0.04	[0.00 0.15]
Father's Time	0.14	[0.02 0.43]

NOTE.—This table reports the decomposition of average cognitive skill differences (in fractions of a standard deviation) between children whose parents remain married throughout their childhood and children whose parents divorce before age 18. The 90 percent confidence intervals from a non-parametric bootstrap with 50 samples are also reported.

Table 7: Divorce Standard Counterfactual Results

	Mutual Consent	Unilateral
% CEV Husbands	-21.96	27.78
	[-28.54 -20.79]	[17.52 30.83]
% CEV Wives	-5.30	10.80
	[-11.77 -3.03]	[0.61 9.91]
Δ Wife log-wage ($\times$ 100)	-1.12	0.49
	[-1.29 -0.41]	[0.13 0.75]
Δ Fertility (%)	-0.44	0.31
	[-2.49 2.74]	[-1.93 1.95]
Δ Divorce (%)	-5.06	2.07
	[-9.13 0.09]	[-0.11 5.49]
Δ Skill (% sd)	-0.28	-1.67
	[-0.96 7.56]	[-5.35 0.79]
Δ TFP	-2.92	0.15
	[-8.53 5.49]	[-3.43 5.18]
Δ Mother's Time	1.67	-0.95
	[-0.08 7.29]	[-4.40 0.15]
Δ Father's Time	0.96	-0.87
	[0.01 5.83]	[-3.77 0.02]

NOTE.—This table reports the predicted changes in mean family outcomes under two counterfactual scenarios: the first column shows the change in means when all unilateral divorce standard states are moved to mutual consent standards (and all mutual consent states remain unchanged). The second column shows the change in means when all mutual consent states are moved to unilateral divorce standards (and unilateral states remain unchanged). Ninety percent confidence intervals for the change in each mean with the move from the status quo to the uniform law are reported in brackets.

Table 8: Child Support Counterfactual Results

	$\pi = 0.3$	$\pi = 0.5$
% CEV Husbands	-3.04	-5.37
	[-9.38 0.50]	[-13.03 -4.25]
% CEV Wives	-0.61	2.80
	[-4.51 3.05]	[-4.83 1.53]
Δ Wife log-wage (× 100)	-0.38	-0.57
	[-0.60 -0.17]	[-0.88 -0.30]
Δ Fertility (%)	-5.31	-8.81
	[-6.03 -3.35]	[-10.78 -7.93]
Δ Divorce (%)	-4.56	-7.02
	[-5.13 -2.07]	[-7.92 -4.41]
Δ Skill (% sd)	-3.05	-6.91
	[-4.08 1.80]	[-11.58 -1.64]
Δ TFP	-4.13	-8.73
	[-4.85 0.52]	[-12.46 -0.49]
Δ Mother's Time	0.43	1.25
	[-0.03 1.36]	[-2.35 1.84]
Δ Father's Time	0.65	0.57
	[0.08 2.08]	[-2.90 1.75]

NOTE.—This table reports the predicted changes in mean family outcomes under two counterfactual scenarios: the first column shows the change in means when the child support rate is increased from 15 to 30 percent of the father's income. The second column shows the change in means when the child support rate is increased from 15 to 50 percent of the father's income. Ninety percent confidence intervals for the change in each mean with the move from the status quo to the new child support rate are reported in brackets.

Table 9: Comparison of Child Support Counterfactual Results under Different Divorce Standards

	Mutual Consent	Unilateral
% CEV Husbands	-9.67	-1.12
	[-22.30 -5.77]	[-2.19 5.07]
% CEV Wives	4.36	-8.71
	[-0.65 11.17]	[-9.03 -1.19]
Δ Wife log-wage (× 100)	-0.14	-0.50
	[-0.84 0.30]	[-0.70 -0.31]
Δ Fertility (%)	-2.65	-5.72
	[-8.23 -2.13]	[-5.61 -3.43]
Δ Divorce (%)	-1.72	-2.31
	[-9.15 -1.39]	[-4.10 -1.24]
Δ Skill (% sd)	4.04	-1.98
	[-5.07 5.16]	[-4.95 1.70]
Δ TFP	3.33	-3.94
	[-5.82 3.82]	[-6.12 -1.19]
Δ Mother's Time	-0.12	1.36
	[-3.11 1.48]	[0.14 2.05]
Δ Father's Time	0.84	0.60
	[0.07 3.81]	[-0.27 2.04]

NOTE.—This table reports the predicted changes in mean family outcomes under two counterfactual scenarios: the first column shows the change in means when the child support rate is increased from 15 to 30 percent of the father's income and all households are subject to a mutual consent divorce standard. The second column shows the change in means for the same change in child support rate when all households are subject to a unilateral divorce standard. Ninety percent confidence intervals for the change in each mean with the move from the status quo to the new child support rate are reported in brackets.

Table 10: Custody Allocation Counterfactual

	Sole Maternal Custody	50-50 Custody
% CEV Husbands	3.23	-0.18
	[-4.16 2.42]	[-3.68 2.81]
% CEV Wives	3.51	0.07
	[-4.16 2.43]	[-3.55 2.68]
Δ Wife log-wage (× 100)	-0.08	0.03
	[-0.11 0.10]	[-0.12 0.09]
Δ Fertility (%)	-0.62	0.19
	[-0.57 0.08]	[-0.24 0.41]
Δ Divorce (%)	-1.08	-0.02
	[-0.81 0.50]	[-0.54 0.82]
Δ Skill (% sd)	-1.65	-1.16
	[-3.47 0.24]	[-0.93 2.52]
Δ TFP	-0.48	-1.47
	[-1.78 1.08]	[-1.03 1.03]
Δ Mother's Time	-0.11	-0.01
	[-0.24 0.32]	[-0.30 0.49]
Δ Father's Time	-1.06	0.33
	[-2.73 -0.18]	[0.11 2.01]

NOTE.—This table reports the predicted changes in mean family outcomes under two counterfactual scenarios: the first column shows the change in means when all families are subject to sole maternal custody in the event of divorce, compared to the baseline simulation in which families realize a distribution of majority maternal custody allocations. The second column shows the change in means with a move from baseline to 50-50 custody. Ninety percent confidence intervals for the change in each mean with the move from the status quo to the new custody allocation are reported in brackets.

Figure 1

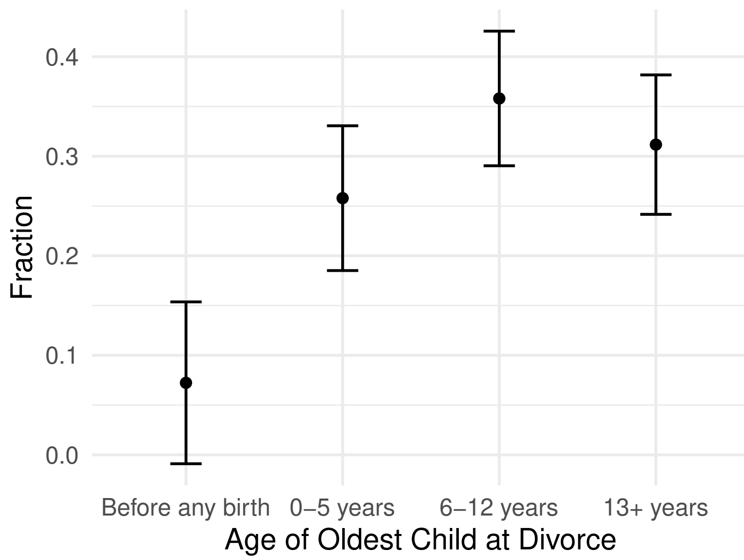

FIG. 1.—Timing of divorce. This figure depicts the timing of divorces on average relative to the birth of the first child. Error bars depict 95% confidence intervals.

Figure 2

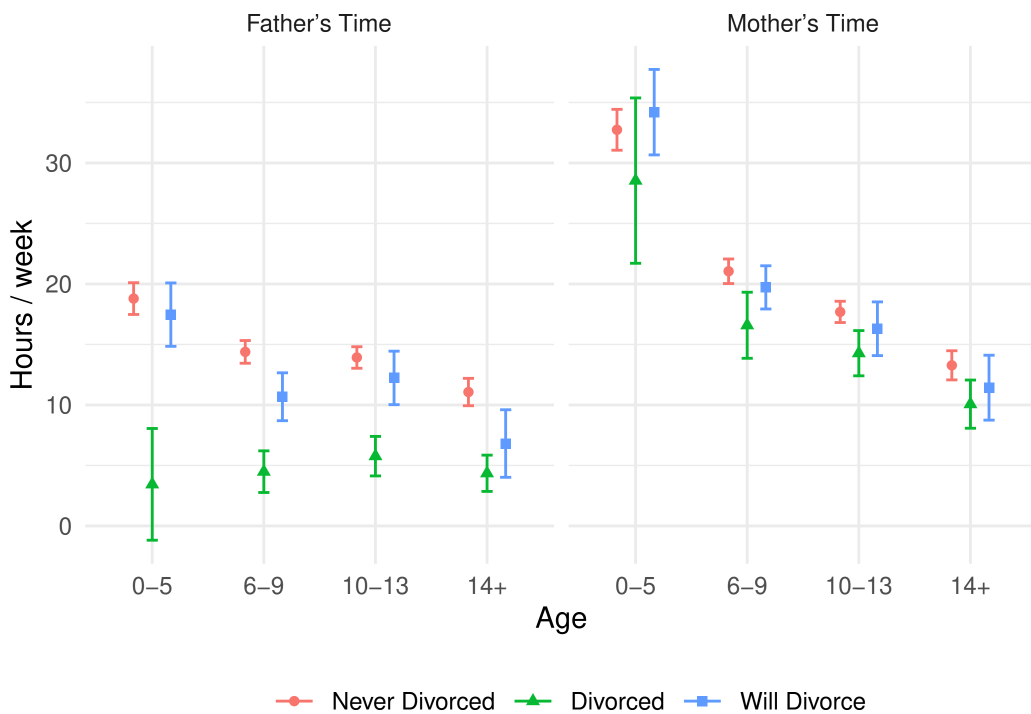

FIG. 2.—Patterns of time investment by marital status. This figure shows average time investment for each parent by the age of the child and by marital status ("Never Divorced", "Divorced", or "Will Divorce". See Section 2 for details on the construction of the time investment variable. Error bars depict 95 percent confidence intervals.

Figure 3

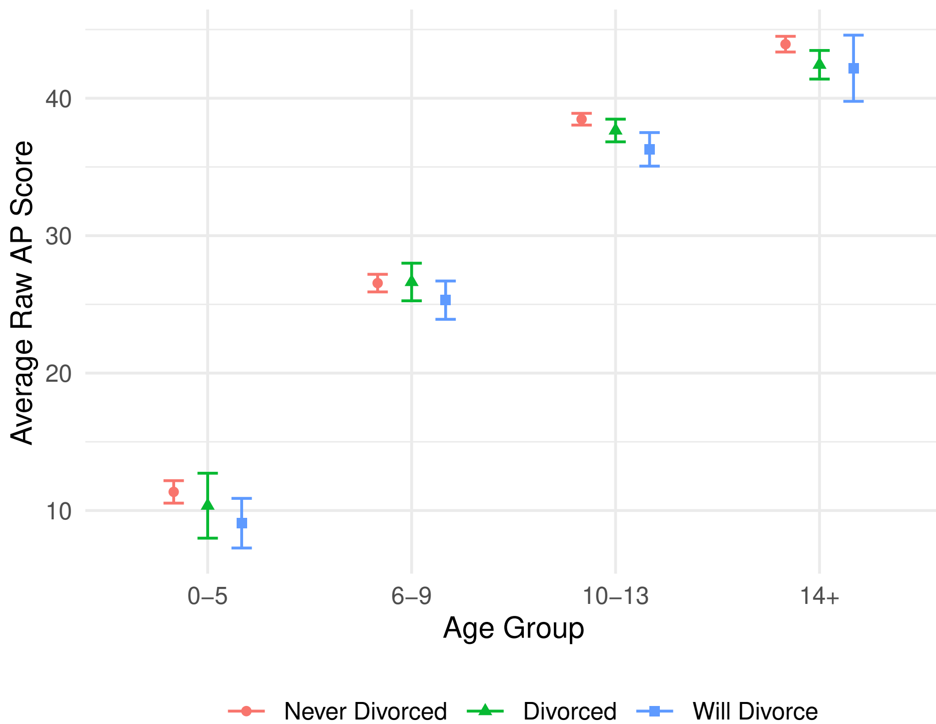

FIG. 3.—Patterns of cognitive skill outcomes by marital status. This figure shows average scores on the Applied Problems (AP) module of the Woodcock-Johnson aptitude test by the age of the child and by marital status ("Never Divorced", "Divorced", or "Will Divorce".) Error bars depict 95 percent confidence intervals.

Figure 4

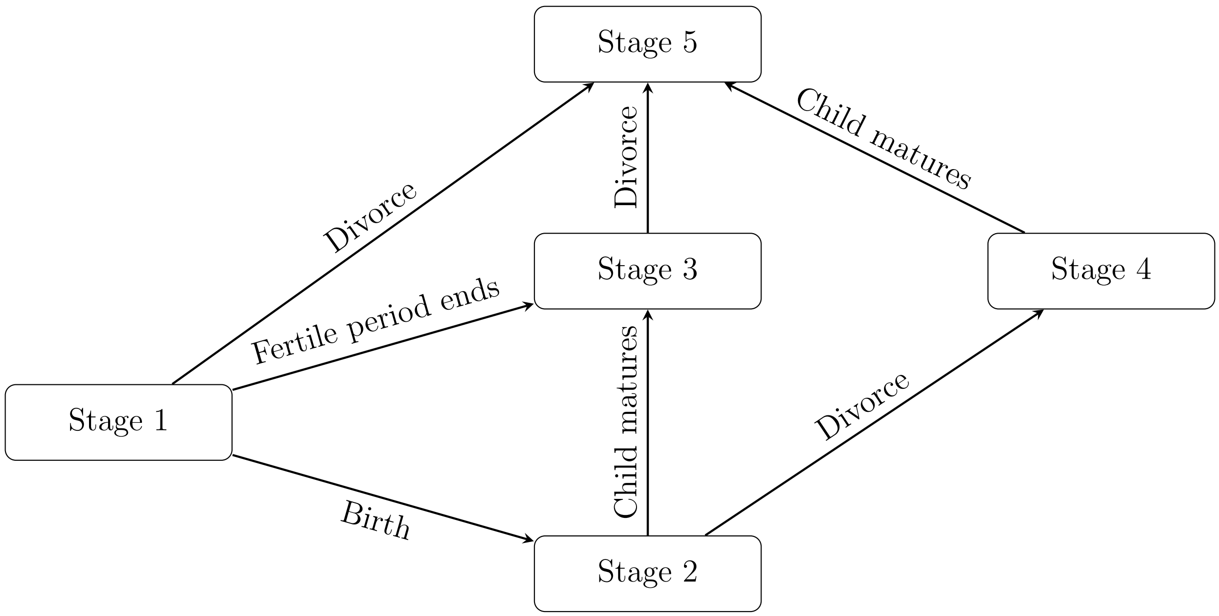

FIG. 4.—The life-cycle stages of the model. This figure depicts the events that precipiate the movement of couples between stage 1 (married and fertile), stage 2 (married with a developing child), stage 3 (married with no children or fertility option), stage 4 (divorced with a developing child), and stage 5 (divorced with no children).

Figure 5

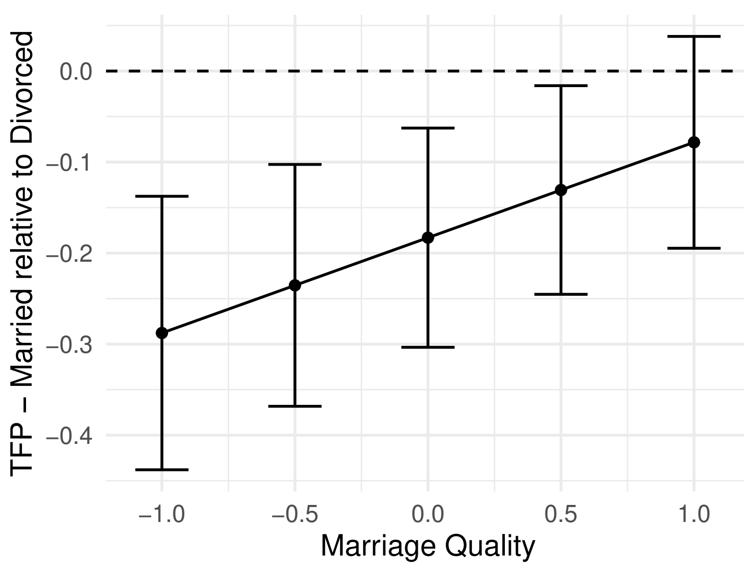

FIG. 5.—TFP in marriage relative to divorce. This figure depicts the implied values of TFP in marriage relative to divorce, $\gamma_{\psi,2}+\gamma_{\psi,3}\omega-\gamma_{\psi,1}$, along with estimated 95 percent confidence intervals.

Figure 6

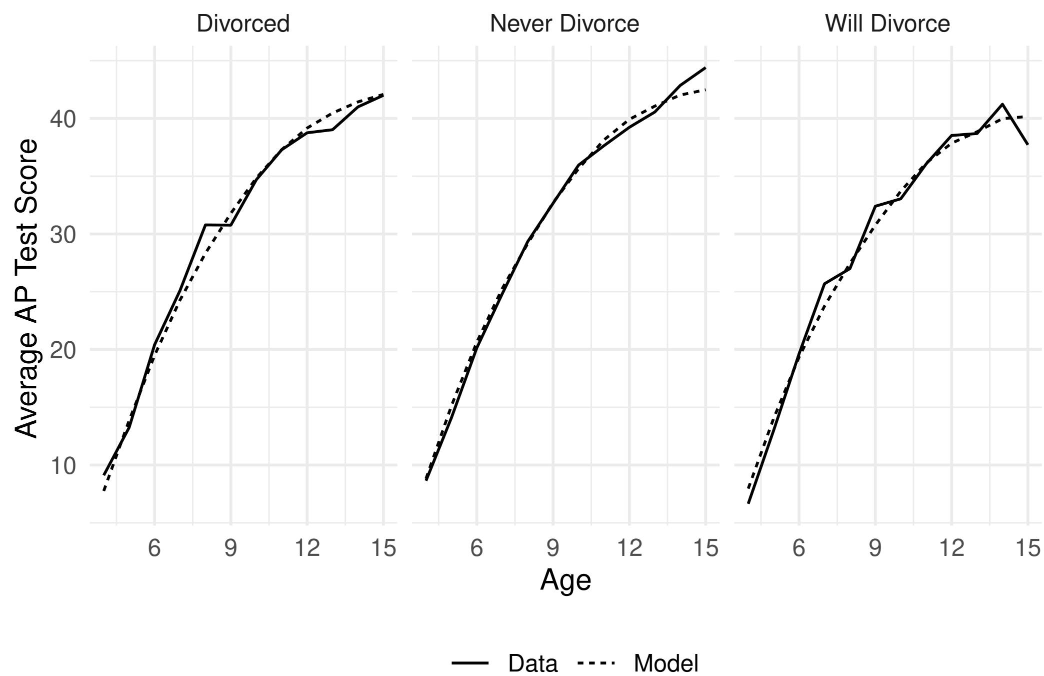

FIG. 6.—Data and simulation test score profiles by age. This figure depicts mean AP test scores by child age in the PSID sample and model simulation.

Figure 7

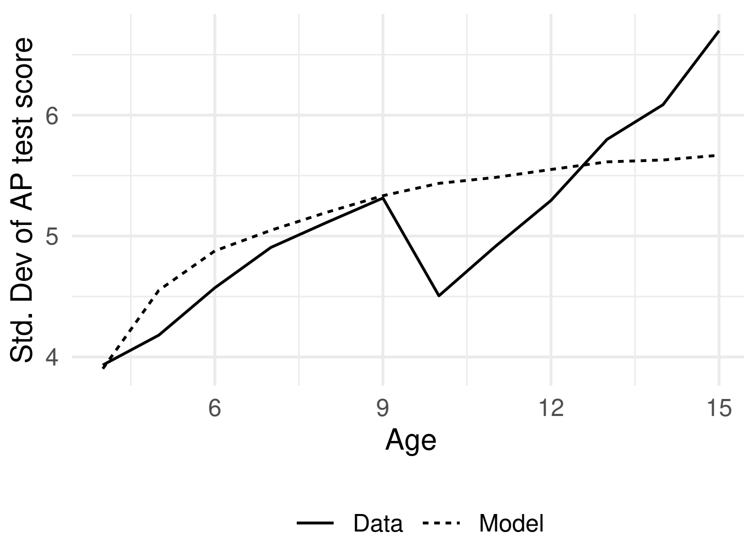

FIG. 7.—Data and simulation test score standard deviation profiles by age. This figure depicts the PSID data and simulated standard deviation of AP test scores by child age.

Figure 8

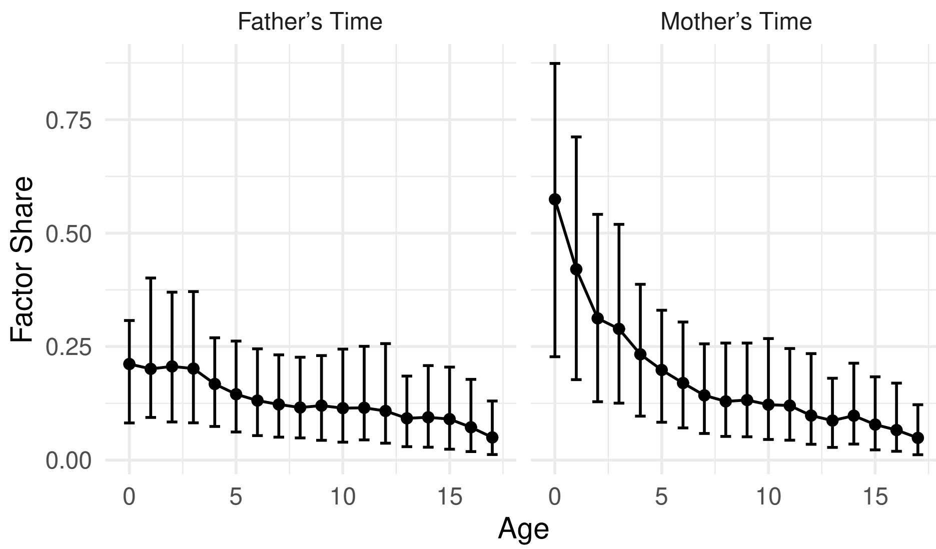

FIG. 8.—Factor shares of mother's and father's time in the production of child ability. This figure depicts the factor shares of mother's and father's time in the production of child cognitive ability for each child age in the model and data, along with estimated 95 percent confidence intervals.