

Investment in Child Quality Over Marital States¹

Meta Brown
Department of Economics
University of Wisconsin, Madison
1180 Observatory Drive
Madison, WI 53706
e-mail: mbrown@ssc.wisc.edu

Christopher J. Flinn
Department of Economics,
New York University
269 Mercer Street
New York, NY 10003
e-mail: christopher.flinn@nyu.edu

July 2006

¹This research is supported in part by a grants from the National Science Foundation and the C.V. Starr Center for Applied Economics at NYU. David Blau, Hanming Fang, John Kennan, Ariel Pakes, Ken Wolpin and seminar participants at the Institute for Research on Poverty, NYU, Virginia, Wisconsin, Torino, the Society for Economic Dynamics, ESPE, the American Economic Association 2002 Meetings, the UNC/Duke Conference on Labor, Health and Aging and the Minnesota Applied Micro Workshop provided valuable comments. Hugette Sun provided extensive child support guideline data. We retain responsibility for all errors, omissions, and interpretations.

Abstract

Policies governing divorce and parenting, such as child support orders and enforcement, child custody regulations, and marital dissolution requirements, can have a large impact on the welfare of parents and children. Recent research has produced evidence on the responses of divorce rates to unilateral divorce laws and child support enforcement. We argue that in order to assess the child welfare impact of family policies, one must consider their influence on parents' investments in their children as well as the stability of the marginal marriage. Further, we expect that changes in the regulatory environment induce changes in the distribution of resources within both intact and divided families. We develop a continuous time model of parents' marital status choices and investments in children, with the main goal being the determination of how policies toward divorce influence outcomes for children. Estimates are derived for model parameters of interest using the method of simulated moments, and simulations based on the model explore the effects of changes in custody allocations and child support standards on outcomes for children of married and divorced parents. We find that, while small changes in children's academic attainment are induced by significant shifts in custody and support, the major effects of these policies in both intact and divided households are on the distribution of welfare between parents. In addition, children's attainments are not necessarily best served by the divorce-minimizing policy.

JEL codes: J12, J13, J18

1 Introduction

Divorced parenting in the U.S. is regulated through a combination of laws controlling marital dissolution, child custody and placement, and the assignment and enforcement of child support obligations. The primary objective of these activities 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 plethora 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 in the life cycle (see, e.g., Haveman and Wolfe 1995). Recent empirical studies of unilateral divorce laws and child support enforcement have had some success in isolating the effects of changes in such legal structures on divorce rates (e.g., Friedberg 1998 and Gruber 2004). We suggest that in developing a complete picture of the influence of divorce regulations on the welfare of family members, particularly that of children, it would be productive to identify the manner in which changes in the legal climate affect child outcomes and the distribution of resources within the family. For example, how might a change in projected child custody allocations influence the probability of divorce? How might it affect each parent's interest in the quality of child outcomes? Taking each of these relationships into account, what is the net effect of the custody change on child welfare? Finally, if marriages vary in quality, what is the child welfare benefit of a policy that succeeds in stabilizing the marginal marriage?

Following the framework developed in Weiss and Willis (1985), Del Boca and Flinn (1995) and Flinn (2000) take as their starting point the problem of expenditure on children faced by divorced parents. The latter two papers model the role of institutions and the agents representing them in determining the welfare of divorced parents and children and take the models to data, but condition on the divorce event. In this research we remedy this potentially important omission by formulating and estimating a dynamic model of divorce and investment decisions. This extends the contribution of Weiss and Willis by looking at the joint evolution of children's human capital and parents' marital status.

We draw on two recent strands of the literature on marriage and childrearing. Analysis of family structure dynamics by Aiyagari, Greenwood and Guner (2000), Brien, Lillard and Stern (forthcoming), Chiappori, Fortin and Lacroix (2002) and others emphasizes the repeated interaction of a husband and wife over marital status and the allocation of household resources. The dynamic individual decision problem of a mother, or married parents with a common objective, is the focus of such child investment studies as Bernal (2006), Bernal and Keane (2006), and Liu, Mroz and van der Klaauw (2003). Our aim is to understand the endogenous growth of children's human capital where family structure and investments result from the distinct choices of mothers and fathers. Such an approach permits the study of divorce regulations with differing effects on the welfare and family attachment of mothers and fathers.

We develop a continuous time model of parents' marital status and child investment decisions. The value of marriage to parents is drawn from a population distribution of match values and

evolves stochastically over time. Parents enjoy utility gains from marriage that result from the exogenously determined match value of the marriage and the output of their child investment decisions. The structure bears similarities to models of turnover and firm-specific human capital investment (e.g., Jovanovic 1979) in that parents invest in a project that produces greater returns while they remain attached to the family, and they have imperfect information on the future values of family attachment to each party. It differs from models of turnover in that parents' returns to investment in children may outlast the marriage match. The theoretical structure allows us to consider the influence of a change in the cost of divorce or each parent's access to the child in the divorce state on married parents' investment in children and their decision to continue in the marriage. Marital dissolution occurs as a result of changes in the state of the marriage, characterized by match quality, child quality, and whether the child quality investment process is on-going. While the evolution of match quality is purely exogenous within an intact marriage, whether a match quality level is sufficiently high to sustain a marriage depends on current levels of child quality, which is a partially endogenous stochastic process that depends on previous realizations of match quality, previous investment decisions of the parents, and other exogenous events. Earlier investment activities contribute to each parent's current benefit from remaining married and enjoying full access to the child. Thus the full history of marriage values and child investments determines current marital status and child investment. If the history of child investments and marriage values is poorer for the marginal marriage than it is for the representative 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 representative marriage.

Regression analysis using our estimation sample of National Longitudinal Survey of Youth (NLSY) only-child families, described in more detail below, demonstrates a negative association between divorce and young children's academic achievement that is consistent with the catalog of such findings in Haveman and Wolfe. Table 1 contains estimates from the ordinary least squares regression of children's age-normed math test scores and test score changes on child and parent characteristics. Children are aged (roughly) six and eight at the first and second test dates. We find that, as compared with having married parents, having divorced parents at both test dates is associated with a decrease of 4.6 percentile points in children's first test scores.¹ Children whose parents divorce between the first and second test dates experience test score changes 14 percentile points lower than those of all other children. A ten thousand dollar increase in family income is associated with a percentile point gain in the first test score and a half of a percentile point gain in the test score change. Mothers earning greater shares of family income have children with substantially higher test scores: an increase of 0.25 in the mother's share of family income is associated with a 3.2 percentile point increase in her child's test score, on average. The effects of the state divorce laws and child support guidelines to which families are subject on child outcomes are difficult to predict without a theory of family behavior; the estimated negative association between unilateral divorce laws and child test scores and positive association between family-specific child

¹The t-statistic on the divorce coefficient in the first test score regression is 1.5.

support guideline prescriptions and child test scores are relatively uninformative absent more careful analysis.

Given the clear problems with this type of treatment of the evolution of family structure and children's academic performance, we abandon the linear regression approach in favor of the richer model of family behavior sketched above. We use the Method of Simulated Moments (MSM) and a sample of NLSY families to estimate the parameters of the model. The estimates indicate a diminishing return to child investments under our specification of the stochastic child quality production process. The contribution of child quality to parents' welfare is considerably less than that of parents' consumption at the sample average income, child quality and investments and the estimated vector of parameters. Where parents remain married at the first test date, the covariances of children's test scores with fathers' and mothers' incomes are 40.71 and 52.14, respectively. Among divorced parents they are 16.84 and 70.93. One question posed by the estimation is whether the model can fit such wide disparities in the covariances between parents' individual incomes and children's test scores without relying on differences in mothers' and fathers' tastes. We find that the model is in fact able to fit the observed patterns in first test scores based on the existing family law's treatment of mothers and fathers, while maintaining the requirement that mothers and fathers place equal preference weight on child quality. It is somewhat less successful at fitting the interaction between marital status, parents' incomes, and score changes between tests.

The repeated interaction between independent decision makers modeled here allows us to perform otherwise impossible policy experiments that address the redistributive aspects of family law changes and their distinct influences on mothers' and fathers' family attachments. Under the assumption that state child support guidelines and custody standards represent optimal decisions on the part of state policymakers, we use characteristics of families residing in three populous states, California, Texas, and New York, along with state policies and the point estimates to infer the relative weights that policymakers in each state place on the welfare of mothers, fathers, and children. Policy parameters assumed available to state policymakers in this exercise are the child support rate paid by the non-custodial to the custodial parent in the event of divorce and the share of child physical placement awarded to each parent. We find that policymakers' objectives favor fathers' welfare in each state, a result not unlike the preference for fathers among Wisconsin family court judges demonstrated by Del Boca and Flinn (1995) for the pre-child support guideline era. Further, while New York policymakers place a weight of roughly 0.30 on child welfare, as compared with 0.22 on mothers' welfare, both California and Texas policymakers place the overwhelming majority of the weight in their objectives on the welfare of parents. The independent decisions of fathers and mothers in the estimated model also allow us to approximate parents' individual welfare, and to simulate child and divorce outcomes, over a grid of child support and custody regimes. We find that the most preferred policies for mothers and fathers differ in the extreme. While mothers prefer high child support rates and low paternal custody shares, fathers prefer the reverse. These results serve to emphasize the obvious point that a major feature of divorce policy is its redistributiveness. The estimated effects of our policy levers on children's attainments fail to moderate the opposition

of parents’ interests in our value function approximations. Finally, we find that children are not necessarily best off under the divorce minimizing policy.

While these experiments demonstrate some of the tensions among the welfare of mothers, fathers and children confronted by policymakers, they must be interpreted with caution given the simplifying assumptions and sample restrictions we have imposed. Future research will address the fertility decision that we have omitted in this study, and we hope that this will lead to more comprehensive policy analysis.

The paper proceeds as follows. In Section 2 we present a model of the dynamic marriage and child investment decisions of two parents involved in rearing a child, along with illustrative simulations. The objectives and choices of state policymakers are discussed in section 3. Section 4 describes methods used in estimating the model parameters given data on parents’ incomes, marital status and marriage quality and children’s ages and attainments. Section 5 introduces data on parents and children from the National Longitudinal Survey of Youth (NLSY), and the results of the empirical analysis are reported in section 6. Section 6 also includes analysis of the social welfare consequences of a range of support and custody policies using the model structure and parameter estimates. This is followed by a section of concluding comments.

2 A Model of Child Investment and Divorce Decisions

There exist three agents in our model, two parents and one child. The welfare of the child at a point in time is summarized by its “quality,” which is a scalar nonnegative index denoted by k . The model is set in continuous time, and the instantaneous utility function of parent p is given by

$$u_p(c_p, k, \theta, b, d) = \alpha_p \ln(c_p) + (1 - \alpha_p) \tau_p(d) \ln(k) + (1 - d)\theta + db, \quad p = 1, 2; \quad (1)$$

where c_p is the consumption of a private good by parent p , d is an indicator variable that takes the value 1 if the parents are divorced, θ is a marriage-specific match value, the value of which can change over time, $\tau_p(d)$ is the amount of contact that the parent has with the child given the divorce status of the parents, b is a parameter that measures the direct disutility of divorce to each parent p , and $\alpha_p \in (0, 1)$ is the preference weight on private consumption. We assume throughout that the price of private consumption is fixed at 1 for both parents.

The utility derived from current child quality by each parent is modified according to the amount of contact the parent has with the child in each marriage state. We assume that when married the parents enjoy complete and concurrent access to the child’s time; without loss of generality, $\tau_1(0) = \tau_2(0) = 1$. Though their intrinsic valuation of the child remains the same in the divorce state, the fact that the child becomes an “excludable” good after divorce reduces the utility flows that parents receive from any given level of child quality.² We assume that parents share time

²This, in fact, is the only way in which our model reflects losses of economies of scale after divorce. Other quasi-fixed costs, such as housing, utilities, etc., which are significant sources of scale economies when three individuals live under one roof, are not included in the current modeling set up.

with the child in divorce, implying $\tau_1(1) + \tau_2(1) = 1$ and $\tau_p(1) \geq 0$, $p = 1, 2$, and that physical custody and visitation allocations are fully anticipated and set exogenously with respect to parental behaviors.

Parents receive incomes of $y_p(d)$, $p = 1, 2$ and $d = 0, 1$. We assume that the income difference across marital states is generated by a child support transfer of $\pi y_1(0)$ from parent 1 to parent 2 in case of divorce. Thus the set of exogenous policy parameters in the model is $(\tau_1(1), \pi)$.

The dynamics of the model are as follows. Parents begin life with a child of quality level k and a marriage-specific match value of θ . At any moment in time, at most one of five possible events may occur. First, the child quality may change value. Child quality assumes one of a finite number of values, $k \in K = \{k_1, \dots, k_T\}$, where $k_1 < k_2 < \dots < k_T$. The current child quality will be interpreted in the analysis that follows as a measure of the child's achievements relative to her or his age cohort. The empirical analog to k that we consider is an age-normed measure of academic performance or behavioral traits. Gains and losses in child quality follow a process that we partition into exogenous and endogenous components. Costly investments in child quality made by the parents increase the rate at which improvements in child quality arrive. The child quality improvement rate is described by the function $\delta(i_1, i_2, \theta)$, where i_p denotes the child quality investment of parent p , $\frac{\partial \delta(i_1, i_2, \theta)}{\partial i_p} > 0$ for $p = 1, 2$, $\frac{\partial \delta(i_1, i_2, \theta)}{\partial \theta} \geq 0$, $\frac{\partial^2 \delta(i_1, i_2, \theta)}{\partial i_p \partial \theta} \leq 0$, $p = 1, 2$. The presence of marriage quality in the child quality production function is meant to capture the impact of the home environment on the effectiveness of a given level of parental investments. We assume a transition function for child quality in which arriving improvements increase current quality from k_t to k_{t+1} with certainty whenever $1 \leq t < T$.

We assume that divorced and married parents share the same child quality production function, and that when in the divorce state marriage quality is equal to 0 in terms of its "productive" value. Since the mean of the (symmetric) marriage quality distribution is normalized to zero, this implies that parents in intact marriages with marriage quality less than 0 are at a productive disadvantage with respect to when they are divorced (for fixed values of i_1 and i_2 , of course), while those with positive match quality values are in a comparatively advantageous position.

Second, the child may experience a setback. Setbacks occur at exogenous rate $\tilde{\sigma}$, and lead to a decline in child quality from k_t to k_{t-1} whenever $1 < t \leq T$. For notational convenience, we define

$$\sigma(k_t) = \begin{cases} \tilde{\sigma} & \text{where } 1 < t \leq T \\ 0 & \text{where } t = 1. \end{cases}$$

The third and fourth possible events are an increase or a decrease in the match value θ . The value of the marriage match is included to permit parents' investments to respond to current information on the stability of the marriage. Like child quality, match quality assumes one of a finite number of values, $\theta \in \Theta = \{\theta_1, \dots, \theta_M\}$, where $\theta_1 < \theta_2 < \dots < \theta_M$. Match quality increases arrive at rate $\tilde{\gamma}^+$. Given a current match quality of θ_m , the arrival of a match quality increase leads with certainty to a new match quality of θ_{m+1} whenever $1 \leq m < M$. Symmetrically, match quality decreases arrive at rate $\tilde{\gamma}^-$, and the arrival of a decrease in match quality leads to a drop from θ_m to θ_{m-1} whenever

$1 < m \leq M$. As with the child quality setback rates, for convenience of notation we define

$$\gamma^+(\theta_m) = \begin{cases} \tilde{\gamma}^+ & \text{where } 1 \leq m < M \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \gamma^-(\theta_m) = \begin{cases} \tilde{\gamma}^- & \text{where } 1 < m \leq M \\ 0 & \text{otherwise.} \end{cases}$$

The values of $\gamma^+(\theta_m)$ and $\gamma^-(\theta_m)$ determine the degree of persistence in marriage quality.

Finally, the child may attain functional independence at the current age-normed child quality, in which case the child quality improvement process ends. The parents enjoy a terminal value that increases with the current child quality level and continues to depend on the parents' marital status.³ Termination of the investment process occurs at exogenous rate η ; state variable $s \in \{0, 1\}$ indicates the current investment condition, and equals 1 when the investment process has been terminated.

Each parent is assumed to have a fixed income flow y_p . While conceptually it is straightforward to augment the model with an exogenous income process for both parents, the computational cost of doing so is prohibitive. Moreover, we face the standard problem of limited data on the income trajectories of divorced fathers. To keep the model tractable, we abstract from the phenomenon of remarriage by assuming that once parents exit the marriage state they never reenter it.

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. Under the cooperative specification, spouses make decisions that place the welfare of family members on the Pareto frontier, and some sharing rule is chosen for the division of the surplus from cooperation.⁴ In the noncooperative case, spouses make decisions representing the equilibrium of a Nash or Stackelberg game, and the family will not, in general, achieve the Pareto frontier.⁵ We look to the empirical literature on divorce regulation for guidance in constructing our model. A testable implication of the hypothesis that married spouses behave cooperatively is that only divorces that are efficient for the family occur. Since laws governing the consent to divorce do not change total family resources, but rather shift property rights within the marriage, a change from bilateral to unilateral divorce laws should have no effect on the decision to divorce when married partners behave cooperatively. As discussed above, in recent studies on the subject both Friedberg and Gruber find evidence of significant effects of unilateral divorce laws on rates of marital dissolution in the U.S., indicating noncooperative interaction in married households. Below we assume that parents behave noncooperatively no matter what their marital state, and that investment strategies constitute a Markov Perfect Equilibrium.⁶ In our discussion of the theoretical results we dedicate some attention to the effects of this modeling choice.

³An alternative approach to finalizing the child investment process would be to impose a fixed time horizon of 18 or 21 years, after which children achieve independence. The drawback to this approach is that it generates strategic manipulations by parents approaching the date of independence that we find unrealistic.

⁴See, for example, Browning and Chiappori (1998).

⁵See, for example, Lundberg and Pollak (1994) and Udry (1996).

⁶See, for example, Pakes and McGuire (2000).

2.1 Divorced Parents

Under our assumptions, divorce is an absorbing state in each parent's marital "life cycle." When the parents are divorced and the child quality improvement process is terminated, a parent p whose child is of quality k_t enjoys terminal value

$$V_p(k_t, b, d = 1, s = 1) = \frac{\alpha_p \ln(y_p(1)) + (1 - \alpha_p)\tau_p(1) \ln(k_t) + b}{\rho},$$

where ρ is the instantaneous discount rate. In the case of divorce with an ongoing child quality improvement process, each parent's only decision is how much to invest in the child. We therefore look for an equilibrium in parental investments, which is determined by the state of child quality and the parental income distribution. To find the equilibrium, we first solve for the reaction function of parent p ; this is the decision rule used by parent p in determining his or her investment level conditional on the investment level of the other parent. The conditional value of the future to divorced parent p is given by

$$\begin{aligned} V_p(k_t, b, d = 1, s = 0 | i_{p'}) = \max_{i_p} & (\rho + \delta(i_p, i_{p'}, 0) + \sigma(k_t) + \eta)^{-1} \{ \alpha_p \ln(y_p(1) - i_p) + (1 - \alpha_p)\tau_p(1) \ln(k_t) + b \\ & + \delta(i_p, i_{p'}, 0) V_p(k_{t+1}, b, d = 1, s = 0) \\ & + \sigma(k_t) V_p(k_{t-1}, b, d = 1, s = 0) + \eta V_p(k_t, b, d = 1, s = 1) \}. \end{aligned} \quad (2)$$

To find the equilibrium investment levels we solve the dynamic reaction functions. Let the function $i_p^*(i_{p'}, k_t, d = 1)$ ⁷ denote the optimal level of investment by divorced parent p given current child quality level k_t and investment by the other parent of $i_{p'}$.⁸ This function is the argument i_p that maximizes the right hand side of [2]. Given the reaction functions $i_1^*(i_2, k_t, 1)$ and $i_2^*(i_1, k_t, 1)$, an equilibrium is a pair of investment values $(\hat{i}_1, \hat{i}_2)(k_t, d = 1)$ such that

$$\begin{aligned} \hat{i}_1 &= i_1^*(\hat{i}_2, k_t, 1) \\ \hat{i}_2 &= i_2^*(\hat{i}_1, k_t, 1). \end{aligned} \quad (3)$$

The properties of this reaction function depend critically on the properties of the improvement rate function δ . Along with $\frac{\partial \delta(i_1, i_2, 0)}{\partial i_p} > 0$, $p = 1, 2$, we assume that δ is twice continuously differentiable and concave, and add to these the restriction that i_1 and i_2 behave as (weak) substitutes. Under these assumptions, $\frac{di_p^*(i_{p'}, k_t, d=1)}{di_{p'}} < 0$ and the reaction function is negatively sloped for each parent p and for all values of $k_t < k_T$.

The expressions in [3] do not fully characterize the equilibrium of the model, since the reaction functions themselves depend upon the equilibrium values $V_p(k_{t'}, b, d = 1, s = 0)$, $\forall t' \neq t$. Equilibrium in the divorce state for a family with an active child investment process is there-

⁷We do not include the state variable s in the investment reaction functions or the equilibrium investment functions since $s = 0$ is a precondition for any investment to occur. Thus the value $s = 0$ is implicit in all of these functions.

⁸Since the direct disutility of divorce is the same across all states for divorced parents, it does not appear as an argument in any function that describes parental investments in this case.

fore determined over the $2T$ parent- and child quality-specific values as well as the $2T$ parent- and child quality-specific investments. The solution is obtained numerically, and the numerical technique employed is simplified by restrictions on the relationships among equilibrium values arising from the theory and the use of the $2T$ values of terminal child qualities. Given the ordering of child qualities and the possibility of setbacks when the investment process is active, we know that $V_p(k_T, b, d = 1, s = 1)$ dominates the divorce-state values of (a) all terminal child qualities k_t such that $t < T$ and (b) *all* non-terminal child qualities. Additionally, $V_p(k_t, b, d = 1, s)$ increases monotonically with k_t for both $s = 0$ and 1 . The numerical solution produces equilibrium investment levels $\{\hat{i}_1(k_t, d = 1), \hat{i}_2(k_t, d = 1)\}_{t=1}^T$ and value functions $\{V_1(k_t, b, d = 1, s = 0), V_2(k_t, b, d = 1, s = 0)\}_{t=1}^T$.

2.2 Married Parents

The experiences they will have if they enter the divorce state can meaningfully affect the investment decisions of forward-looking married parents. In particular, currently married parents who believe that divorce is likely in the near future will make investment decisions that look more like those made by divorced parents than will couples who believe that divorce is a remote possibility. In our model, the likelihood of divorce is (partially) endogenous. We posit the existence of a match value of the marriage θ that evolves according to an exogenous stochastic process. We structure the problem so that when this value of marriage is monotone increasing in θ , so that marriage continuation results only when the current match quality value exceeds a critical value, which is a function of exogenous characteristics, such as the parental income distribution and the family law environment, and on the (endogenous) history of investments in child quality.

We must also specify the manner in which divorce decisions are made. Under our assumption of noncooperative behavior, these decisions are not, in general, efficient. The nature of the decisions depends critically on legal statutes. In our preliminary analysis, we have considered two different cases: one in which it is enough for one of the parents to ask for a divorce for the couple to enter the divorce state and the second in which both parents must agree to the divorce for it to occur. These cases are commonly termed unilateral and bilateral divorce regimes. After much experimentation, we have settled on the assumption that parents interact in a unilateral divorce environment.⁹ With knowledge of the process by which divorce decisions are made, we define $Q_p(k_t, \theta, b, s)$ as the value to parent p of the marital status chosen in equilibrium by both parents in state (k_t, θ, b, s) .

⁹Solutions of the model under the assumption of bilateral or unilateral divorce laws, in the presence and absence of side payments, generate only very small differences in predicted behavior between the two policies. Preliminary estimates using model solutions consistent with the prevailing divorce laws in each family's state showed little response to the added information. This appears to be in large part the result of the narrow range of child and marriage quality levels at which parents disagree over the divorce decision for reasonable parameterizations of the model. Since little is gained by varying the solution technique with state divorce laws, while the computational burden of estimation is doubled by this measure, we estimate with the assumption that all families operate under the (narrowly) more common unilateral divorce law. A remaining question is how theory might be able to generate the non-negligible effects on divorce rates of unilateral divorce laws observed in the U.S. if the extreme assumptions of non-cooperative interactions and no side payments generate little difference in behavior between bilateral and unilateral divorce regimes in this type of model.

The derivation of the marriage state equilibrium is similar to that of the divorce state equilibrium, with one major difference being the search for an equilibrium in divorce decisions as well as investments and values. As before, we begin with the value of a terminated child investment process at k_t for parent p :

$$\begin{aligned}
V_p(k_t, \theta_m, b, d = 0, s = 1) &= (\rho + \gamma^+(\theta_m) + \gamma^-(\theta_m))^{-1} \{ \alpha_p \ln(y_p(0)) \\
&\quad + (1 - \alpha_p) \ln(k_t) + \theta_m + \gamma^+(\theta_m) Q_p(k_t, \theta_{m+1}, b, s = 1) \\
&\quad + \gamma^-(\theta_m) Q_p(k_t, \theta_{m-1}, b, s = 1) \}
\end{aligned} \tag{4}$$

In this case, the only state variable that can change is the marriage match value, θ . Since both parents' welfare levels are increasing in child and marriage quality, an increase in marriage quality (which occurs at rate $\gamma^+(\theta_m)$) cannot lead to a divorce, so that $Q_p(k_t, \theta_{m+1}, b, s = 1)$ corresponds to the value of marriage at those state variables. However, a decrease in marriage quality (which occurs at rate $\gamma^-(\theta_m)$), may lead to a divorce or continuation in the marriage state. It is possible, given values of y_1, y_2 , and k , that marriage will dominate divorce even at the lowest marriage quality value, θ_1 . In this case, marriage is an absorbing state for households characterized by (y_1, y_2, k) .

Next, given the current child quality level and match value, we solve for the equilibrium investment levels and associated values for each parent conditional on the continuation of the marriage. As in the divorce case, using the reaction functions we can define a pair of equilibrium investment levels and parent-specific state values associated with marriage that are given by

$$(\hat{i}_1, \hat{i}_2)(k, \theta, b, d = 0, s = 0); (V_1, V_2)(k, \theta, b, d = 0, s = 0). \tag{5}$$

The investment equilibrium depends on the current marriage quality both through its direct influence on the productivity of child investment and through its effect on the anticipated duration of the parents' marriage, which partially determines the expected gain associated with an increase in child quality.

With the parents' equilibrium investments in the child found as in [5], the value to parent p of marriage, an ongoing child improvement process, and child quality k_t is

$$\begin{aligned}
V_p(k_t, \theta_m, b, 0, 0) &= (\rho + \gamma^+(\theta_m) + \gamma^-(\theta_m) + \delta(\hat{i}_p, \hat{i}_{p'}) + \sigma(k_t) + \eta)^{-1} \{ \alpha_p \ln(y_p(0) - \hat{i}_p) \\
&\quad + (1 - \alpha_p) \tau_p(0) \ln(k_t) + \theta_m + \gamma^+(\theta_m) Q_p(k_t, \theta_{m+1}, b, 0) \\
&\quad + \gamma^-(\theta_m) Q_p(k_t, \theta_{m-1}, b, 0) + \delta(\hat{i}_p, \hat{i}_{p'}) V_p(k_{t+1}, \theta_m, b, 0, 0) \\
&\quad + \sigma(k_t) Q_p(k_{t-1}, \theta_m, b, 0) + \eta Q_p(k_t, \theta_m, b, 1) \}.
\end{aligned}$$

To find equilibrium investments, values, and divorce decisions over the marriage quality distribution and for all child quality levels, we again make use of the restrictions on the relative values of the possible child and marriage quality states implied by the theory. The solution is obtained numerically, but in this case equilibrium occurs over all $2T$ parent- and child quality-specific values

and investments across all M possible values of θ and for all values of b . Computation of the equilibrium is simplified by the presence of the terminal values represented in [4]. Having followed the above steps, we have the complete solution for the marriage state,

$$\left\{ \{(\hat{i}_1, \hat{i}_2)(k_t, \theta_m, b, 0, s), (V_1, V_2)(k_t, \theta_m, b, 0, s)\}_{t=1}^T \right\}_{m=1}^M, s = 0, 1,$$

along with divorce decisions in every state.

2.3 Characterizing the Equilibrium of the Model

Given the relatively large number of state variables, strategic interactions between parents, and the complicated exogenous and endogenous dynamics of the child quality and marital status processes, it is no easy task to characterize the equilibria of the model and conduct comparative statics exercises. The recovery of primitive parameters gives us the capacity to perform comparative statics exercises evaluated at the primitive parameter estimates. In this section we focus on the more limited objective of graphically representing the nature of parental decision rules for some values of the state variables, including those (partially) characterizing the family law environment. By presenting and discussing a few figures, we hope to give the reader a feel for some of the more important characteristics of the equilibria of the model. This will aid in interpreting the parameter estimates and in understanding the outcomes of the welfare exercises reported below.

We begin by presenting parental monetary investments in improving child quality in the divorce and marital states, and for three different values of the child support transfer parameter π . The graphs are plotted for the sample median parental incomes, $y_1 = 22,920$ and $y_2 = 12,360$. We assume that in the divorce state, the father's time allocation, $\tau_1(1)$, is equal to 0.2. The child support transfer rates considered are $\pi = 0.25, 0.17$, and 0. In the marital state, it is assumed that the level of marriage match quality is θ_4 , which is the second highest level that can be realized. In the divorce state, the marriage quality level that enters in the child quality production function is essentially normalized to 0 (which is equivalent to θ_3). In this sense, the graphs of the investment decisions under divorce and marriage are not comparable, since θ , as a productive input in the child quality production function, is not held constant.¹⁰ We plot the decision rules together just so that we can compare their qualitative characteristics.

On the horizontal axes in Figure 1 are the nine levels of k at which investment could possibly take place (recall that no investment occurs when $k = k_{10}$, the highest child quality level attainable). The vertical axes contain the investment levels, in thousands of dollars. The top panels plot the mother's investment as a function of k for each of the three values of π considered. The corresponding investments of the father are plotted in Figures 1.d-1.f.

From all of the figures we see that investments in child quality are nonincreasing in k , which is to be expected given that income and θ are held constant throughout the exercise. We also note

¹⁰ A reasonable alternative would be to plot investments at $\theta = 0$ for the married *and* divorced. This is undesirable as it involves the representation of many conditional marriage-state investments for marriages that would in fact be terminated in equilibrium.

that the mother’s investment is typically far greater than the father’s in the divorce state, and is far less in the marriage state (while these investment levels are for different values of θ , the same pattern is observed if we set $\theta = \theta_3$ in the marriage state as well). This result follows from the shift in the preferences that occurs with divorce, and is reinforced by the income transfer made from the father to the mother under divorce. In Figures 1.c and 1.f, there is no transfer from the father to the mother in the divorce state. However, the father’s valuation of child quality declines by 80 percent when switching to divorce, and thus his contribution to investment falls as well. Nonetheless, he does contribute positive amounts of resources to investment at most child quality levels (Figure 1.f). In the divorce state, the mother’s contributions are greater, however (Figure 1.c).

We see that the differences between the mother’s and father’s investments in the marriage and divorce states are exacerbated as we implement increasingly substantial income transfers. Even at $\pi = 0.17$, the father’s divorce state investment is zero or negligible at all child quality levels (Figures 1.d and 1.e). The mother picks up the slack, and devotes larger resources to investment in child quality as her income (via π) increases.

We note that there are some slight differences in the parental investments, even when married, across the three values of π . While some of these differences arise from approximation error, there can exist real effects of the divorce law on within marriage behavior since the agents are forward-looking. Given that the investment functions are roughly similar across π -regimes, we might speculate that these expectational effects are not exceedingly large for agents with the characteristics assumed in generating these figures.

In Figure 2 we plot the mother’s and father’s investment functions, in marriage, for three different values of k , where we now let the marriage quality take each of its five possible values. The reader should bear in mind that some of these values may not be consistent with equilibrium outcomes, in the sense that for our given value of (y_1, y_2) , a $k = k_1$ and $\theta = \theta_1$ may result in divorce. These investment values are associated with the conditional value function of the marriage state. We assume that $\pi = 0.17$ in this exercise. Figures 2.a-2.c plot the mother’s investment level as a function of θ , while Figures 2.d-2.f contain plots of the father’s investment level.

Before discussing the specific features of the graphs, we consider for a moment the relationship between θ and investment decisions, at any level of k . First, and most directly, as θ increases parental investments in child quality become more productive. The impact of this change on parental investments is ambiguous, due to income and substitution effects. At certain (k, θ) levels the efficiency gains in investment may dominate so that even more resources are devoted to investment, and at other levels the income effect may dominate, resulting in lower investments in children in a nominal sense, even if not in terms of “effective” investment levels.

A second, more subtle, impact is on the long-term value of investment in child quality. Other things equal, for any k level, a higher value of θ indicates a lower likelihood of divorce over any given horizon, given the persistence in marriage match values. Since the payoff from child quality is far lower in the divorce state, a high value of θ increases the long run return to investment in child quality. Thus a high value of θ both impacts the short-run productive value of investment in

child quality, and its expected long-run return.

Given that marginal utility is decreasing in child quality, it comes as no surprise that investment levels decrease in k across all values of θ for both the mother and the father. Figures 2.a and 2.d plot the investment levels of the mother and father, respectively, at the lowest child quality level. We see that for both parents, investment levels are strongly increasing up to θ_4 , but fall at θ_5 . Presumably, the fall at θ_5 mainly reflects the increased productivity of investment associated with the highest match quality level, some of which is recouped in the form of personal consumption. At the child quality level k_4 , qualitative features of the mother's investment function are the same as at k_1 . Due to decreasing marginal returns in child quality, total investment is less at all θ values. We do note that the father's investment function is now monotone increasing. The fact that he invests more at θ_5 than θ_4 may reflect differences in the likelihood of remaining married at the two values. Since the father suffers disproportionately under divorce, in terms of time spent with the child, this factor may be particularly relevant to his investment decision. When child quality is high, at the level k_9 , neither parent makes large investments. The mother's investment level appears to be monotone in θ in this case (Figure 2.c), while the father's investment level is quite low and exhibits some nonmonotonicity (Figure 2.f).

3 Optimal Family Law Environments

In addition to estimating the parameters required to characterize the investment and divorce decisions of parents, a focus of our analysis is the comparison of alternative family law environments. With a common welfare metric that allows us to aggregate the utility of all population members in our modeling environment, we will be able to use our estimates to determine the implicit weights attached to the welfare of various agents that appear in the model. We will also be able to determine the family law configurations that are most beneficial to mothers, fathers, and children.

3.1 Social Welfare

In attempting to assess the implications for the distribution of welfare of a given family law environment, we will need to be able to aggregate welfare levels of the various agents who appear explicitly in our model. Within our context of one-child families, these are the two parents and the child. We will utilize a standard Benthamite social welfare function of the form

$$W(w_1, w_2, w_c) = a_1 w_1 + a_2 w_2 + a_c w_c,$$

$$\text{where } 1 = a_1 + a_2 + a_c, \quad a_i \geq 0, \quad i = 1, 2, c,$$

and where the w_j is the welfare measure associated with agents of type j .

Let a family law environment be denoted by F . The value of the social welfare function depends

on F through the relationship

$$W(w_1, w_2, w_c; F) = a_1 w_1(F) + a_2 w_2(F) + a_c w_c(F). \quad (6)$$

Since the a_j are essentially preference parameters, we do not allow them to be functions of the policy variable F . However, we do make the strong assumption that the same weight is attached to all members of a particular class of agents. For the moment, we continue to remain deliberately vague concerning the exact definitions of the individual welfare measures w_j , particularly since their definition is not an uncontroversial issue. However, because the family law environment is established at an aggregate level (typically the state, in the U.S.), any individual welfare measure we propose for agents of type j will involve aggregating welfare outcomes of individual members of the class.

For members of a given class of agents, that is, mothers, fathers, and children, family i welfare outcomes under a given policy F will depend on household characteristics, X_i . Then the welfare outcomes of an agent of type j in household i are given by $w_j(X_i; F)$. Our population aggregate consists of N households, say, with the observable types given by the sequence $X = \{X_i\}_{i=1}^N$. Using the same additivity assumptions as were used to generate (6), we have

$$w_j(F) = \sum_{i=1}^N w_j(X_i; F).$$

This expression also serves to stress the obvious point that a given family law environment can have very different implications for the distribution of intrafamily welfare for various $X_i \in X$.

3.1.1 The Determination of Welfare Weights

Given a family law system F and consistent estimates of all of the primitive parameters of the model, we will define $\hat{w}_j(X_i; F)$ as the consistent estimator of the welfare outcome for individual agent j in household with characteristics X_i under F . It then immediately follows that

$$\hat{w}_j(F) \equiv \sum_{i=1}^N \hat{w}_j(X_i; F)$$

is a consistent estimator of $w_j(F)$ given continuity of the w function with respect to the primitive parameters.

We assume that a family law environment F belongs to a set of all “possible” family law environments, Ω_F . For simplicity, we will assume that the set Ω_F is compact and continuous, and that $W(F; X, a)$ is continuously differentiable with respect to F on Ω_F for all a and a given X .

Definition 1 A vector of welfare weights a^l are admissible if and only if

$$\frac{\partial W(F; X, a^l)}{\partial F} = 0$$

$$\frac{\partial^2 W(F; X, a^l)}{\partial F \partial F'} \text{ negative definite.}$$

In other words, F must maximize the welfare function as characterized by a^l . Without further assumptions, it is not possible to rule out the possibility of no admissible a for a given problem, and there is also the distinct possibility of there existing a multiplicity of admissible values of the weighting vector. In the latter case, the “social preferences” captured by a could be any of the admissible values, with no obvious refinement to employ to select among them. Whether this is the case in practice depends on the particular empirical application.¹¹

In the application we perform below, we limit attention to two policy choices for the “social planner,” regarding the division of time with the child under divorce, $\tau_1(1)$, and the child support tax rate π . Now these parameters, particularly the child support tax rate, vary significantly across states, so we will consider this analysis to apply strictly at the state level. Let the characteristics of households in state s be given by X^s , and let the welfare weights for state s be given by the vector a^s . The family law choices of state s are the pair $F^s = (\tau_1^s(1) \pi^s)$. Since they are chosen optimally,

$$0 = a_1^s \frac{\partial \sum_{i \in s} \hat{w}_1(X_i; F^s)}{\partial \tau_1(1)} + a_2^s \frac{\partial \sum_{i \in s} \hat{w}_2(X_i; F^s)}{\partial \tau_1(1)} + (1 - a_1^s - a_2^s) \frac{\partial \sum_{i \in s} \hat{w}_3(X_i; F^s)}{\partial \tau_1(1)},$$

$$0 = a_1^s \frac{\partial \sum_{i \in s} \hat{w}_1(X_i; F^s)}{\partial \pi} + a_2^s \frac{\partial \sum_{i \in s} \hat{w}_2(X_i; F^s)}{\partial \pi} + (1 - a_1^s - a_2^s) \frac{\partial \sum_{i \in s} \hat{w}_3(X_i; F^s)}{\partial \pi}.$$

All of the partial derivatives appearing in these two equations are computable using sample characteristics in state s and consistent estimates of the primitive parameters. Then we have a system of two linear equations in two unknowns, which we use to determine unique estimates of a^s .

While we can determine weights uniquely in this manner, there are two further issues that must be considered. The first, and most obvious, is whether the weights satisfy the restriction of nonnegativity. While it may be possible to rationalize negative weights attached classes of agents, we would interpret negative weights as indicating a serious misspecification of the social welfare function.

The second, more subtle, issue, concerns whether the vector F^s actually *maximizes* the welfare function. To address this issue, we determine the estimated Hessian associated with the optimization problem and determine whether it is negative definite. We will only investigate whether the Hessian is negative definite at F^s , that is, we will content ourselves with showing that the objective

¹¹In the application of Del Boca and Flinn (1995), welfare weights were assumed to be specific to judicial actors, and each sample observation was assumed to have been decided by a unique judge. The only policy choice in that example was the amount of the child support order, and the authors showed that each sample observation could be used to determine the weight assigned to the father and the sums of the weights attached to the mother and child. The current application is more complex in that there are multiple policy instruments and hence a much more complicated inversion problem.

function is locally concave in the neighborhood of the chosen values.

3.1.2 Agent-Specific Welfare Measures

In static models, the definition of agent-specific welfare is considerably less involved than in dynamic framework, where one must immediately face the problem of choosing the time at which welfare is to be evaluated. Of course, more sophisticated valuations of agent-specific welfare may be defined by integrating over the instantaneous utility values associated with each point in an agent’s lifetime, but even in this case the choice of weighting convention may influence welfare valuations. While the payoff measures we use seem to be natural ones, the reader should bear in mind that many other reasonable ones could be proposed as well.

For a household characterized by the (initial) state vector X_i , stochastic events and the family law environment F produce a continuous sample path of utility outcomes for the parents and quality outcomes for the child. While the parents already have been assigned utility functions, the child has not. We will assume that the child’s utility at any moment in time is simply equal to $\ln(k)$. Since parental flow utilities include the log of child quality, this seems to be relatively consistent.¹²

Let $U_j(t|X_i, F, \Lambda)$, where Λ denotes the sequence of random events that occur to the household. Then discounted lifetime welfare (beginning from the birth of the child) is given by

$$\int_0^{\infty} \exp(-\rho t) U_j(t|X_i, F, \Lambda) dt$$

for a given sequence of random events Λ . Let the cumulative distribution function of all sequences be given by $B(\Lambda)$. Then the expected welfare of agent j in household i under family law environment F is

$$\int \int_0^{\infty} \exp(-\rho t) U_j(t|X_i, F, \Lambda) dt dB(\Lambda).$$

While the exact form of B is unknown, this integral can be approximated through the use of Monte Carlo integration, which is performed by constructing a large number of sample paths $\Lambda_l, l = 1, \dots, L$, using the stochastic structure of the model, and then forming

$$L^{-1} \sum_{l=1}^L \int_0^{\infty} \exp(-\rho s) U_j(s|X_i, F, \Lambda_l) ds. \tag{7}$$

This is the manner in which we compute w_j under the lifetime welfare measure.

¹²A more elaborate version of our model would embed behavior in an overlapping generations framework. In this case, the children of “today” will be the adults of “tomorrow,” and the child quality level with which they end the investment period would serve as an important initial condition with which they begin their adult life. Without providing a direct link between k and the initial conditions of the parents in our model, which include their schooling and income levels, we are not in a position to close the model in this way. However, the reader should bear in mind that the terminal k value is only an indicator of the child’s lifetime welfare.

3.1.3 Specification of the Optimal Family Law Environment

By construction, under our assumptions regarding the agent class welfare payoffs under the family law environment, $w_j(F)$, and our empirical determination of the weights a_j , the family law system F is optimal. Policy experiments in our framework amount to changes in the welfare weights, while keeping constant the policy invariant payoff functions $w_j(\cdot)$. In estimating the welfare weights, we essentially have determined a set of values that are necessary to rationalize the existing family law system given an assumption of optimal behavior by the institutional agent(s). We are free to consider what type of family law system other social planners, with different welfare weights, would have designed.

Of course, we could consider a continuum of changes in the weighting function, but it seems most instructive to limit our attention to some obvious cases. The first case to consider is one in which all the weight is given to the child's welfare, or $a_c = 1$. While there may exist normative arguments supporting the use of this weighting scheme,¹³ our main interest is in seeing how the optimal family law system would differ from the status quo. Two other very straightforward experiments are those in which all the weight is placed on the father's welfare, implying $a_1 = 1$, or the mother's welfare, implying $a_2 = 1$. Comparison of the results of these experiments will contribute to our understanding of the inherent degree of conflict or harmony among the objectives of mothers, fathers, and children as they are influenced by family law.

4 Estimation Method

We estimate the model using a method of simulated moments (MSM) estimator. In this section we provide the details of the estimation procedure.

While the general estimation strategy we outline below can be used with any number of functional form assumptions on the investment process that satisfy our conditions for uniqueness of the Nash equilibrium investment choices, in the results reported below we assume that

$$\delta(i_1, i_2, \theta) = \delta_0(\theta)[i_1 + i_2]^\nu,$$

where $\nu \in (0, 1)$ and δ_0 is a known function that is increasing in θ and takes values on the nonnegative real line. This form of the δ function satisfies the requirement that $\frac{\partial^2 \delta(i_p^*, i_{p'}, \theta)}{\partial i_p \partial i_{p'}} \leq 0$, for all θ . The specific functional form of $\delta_0(\theta)$ used in the estimation is $\delta_0(\theta) = \tilde{\delta}\Phi(\theta)$, where $\tilde{\delta}$ is a scalar to be estimated.

In specifying the model, we assumed that individuals were characterized by a flow cost of divorce, given by b . To improve the fit of the model, and to introduce some realistic sources of heterogeneity, we allow b to assume two values in the population. Then define the random variable

¹³Given the specification of the payoff functions for each of the agents, in most states of the world the child will have the lowest utility level at each moment in time. A Rawlsian argument would equate household welfare with the welfare level of its worst-off member. But, as was discussed in the previous footnote argued, if k is only an indicator of lifetime welfare, this argument no longer holds.

\tilde{b} , which takes the value b (< 0) with probability $\exp(\beta_r rc)/(1 + \exp(\beta_r rc))$, and takes the value 0 with the complement of that probability, where rc indicates that the mother is Roman Catholic.¹⁴ The normalization implicit in our specification of the probability function is that marriages in which the mother is not a Roman Catholic have probability 0.5 of assuming either value of \tilde{b} , while households in which the mother is Roman Catholic have a higher probability of being a high divorce cost type when $\beta_r > 0$.

The endogenous variables utilized in the estimation procedure consist of a youth's score on a mathematics examination administered as part of the NLSY Child survey (some details on the nature of this examination are provided in the following section) at G_j points in time, $G_j \geq 1$, and whether the child's parents are divorced at the time the test(s) were taken. We denote child j 's score on test number g by $O(j, g)$, and to assist in identification we will assume that there is a deterministic mapping that exists between this score and the child's human capital level. In particular, we assume that

$$k(j, g) = k_t \Leftrightarrow o_{t-1} < O(j, g) \leq o_t, \quad t = 1, \dots, T,$$

where $o_0 = 0$, $o_T = 100$, and $o_0 < o_1 < \dots < o_T$. We denote child j 's age at the time the g^{th} test was taken by $a(j, g)$, and the binary variable that indicates that the parents were divorced at this time is given by $d(j, g)$. The parents' incomes are assumed constant for purposes of this analysis and are denoted by $y_p(j)$, $j = 1, 2$.

Conditional on the parental income observations and the age of the child at the survey, the endogenous variables are $k(j, g)$ and $d(j, g)$. As the model clearly demonstrates, these variables are functions of realizations of exogenous and endogenous stochastic processes. The exogenous stochastic processes are those that describe the termination date of the "window" for child quality improvement and the trajectory of the marriage quality characteristic θ . The endogenous stochastic process is the one determining the timing of improvements in child quality; this process is endogenous since parental behavior determines the (average) rate of change.

For (household) sampling unit j , the dependent variables (jointly) take one of $T^{G_j} \times (G_j + 1)$ possible values.¹⁵ Because the stochastic process generating these outcomes is rather complicated due to the endogeneity of the investment in child quality improvements, we utilize the method of simulated moments to estimate the model. To implement this procedure requires access to a large number of simulated sample paths for each sample household j , all of which terminate at age $a(j, G_j)$, and which produce realizations of $\{k^r(j, g), d^r(j, g)\}_{g=1}^{G_j}$. For the moment, condition on the states of marriage and child quality at the time of the birth of the child ($\theta(1)$ and $k(1)$). Given $\theta(1), k(1), y_1(j), y_2(j)$, and that the parents are married at the time of the birth, we first

¹⁴One could allow the flow costs of divorce to differ across spouses, and this would probably be preferable if the size of the state space of the problem was not an issue and if the appropriate data were available. Since the NLSY is an individual-based survey, detailed information on religious affiliation is only available for the wife.

¹⁵The G_j child quality observations assume T^{G_j} values. Since divorce is assumed to be an absorbing state, the number of distinct marital status sequences in G_j observations are $G_j + 1$.

solve for the equilibrium investment rate in child quality, $(\hat{i}_1, \hat{i}_2)(k(1), \theta(1), d = 0, s = 0)$.¹⁶ The rate of child quality improvement immediately following the birth of the child is given by $\delta(1) = \delta_0(\theta(1))[\hat{i}_1(k(1), \theta(1), d = 0, s = 0) + \hat{i}_2(k(1), \theta(1), d = 0, s = 0)]^\nu$. The rate of arrival of a negative shock to the child's quality level, one that results in a decrease of one level, is given by $\tilde{\sigma} > 0$, for $t = 2, \dots, T$. The rate of arrival of an increase in marriage quality is given by $\gamma^+(\theta(1))$, and the rate of arrival of a decrease in marriage quality is given by $\gamma^-(\theta(1))$. The rate of arrival of the (exogenous) termination of the child quality process is fixed once and for all at η ; recall that this is a one time event.

We have access to a random sample of J one-child families. For each observation we perform R replications for each possible set of initial conditions $(k(1), \theta(1))$. The “base draws” for the random number generation are kept constant across iterations of the estimation algorithm to facilitate the convergence process. For any given individual, we draw a total of $R \times S$ values from a uniform pseudo-random number generator for use in generating the timing of changes in the child quality improvement process (these values are denoted $u^{(1)}$), and another $R \times S$ uniform random numbers for use in the generation of the timing of decreases in the level of child quality ($u^{(2)}$). The timing of increases to marriage quality are generated from the draws contained in the $R \times S$ matrix $u^{(3)}$, while the similarly-dimensioned $u^{(4)}$ contains the draws that generate the timing of decreases in the marriage quality level. Finally, we draw an $R \times 1$ vector for generating the duration of the “window” for child quality improvement ($u^{(5)}$).

Consider the generation of the first event for a sample member with endogenous arrival rate parameter $\delta(1)$ and exogenous rate parameters $\sigma(k(1))$, $\gamma^+(\theta(1))$, $\gamma^-(\theta(1))$, and η in replication r . We define the length of time until the improvement in child quality by

$$\hat{q}_1(r, 1) = -\frac{\ln(1 - u^{(1)}(r, 1))}{\delta(1)}.$$

The length of time until a decrease in the quality level of the child is given by

$$\hat{q}_2(r, 1) = -\frac{\ln(1 - u^{(2)}(r, 1))}{\sigma(k(1))}.$$

Note that if $k(1) = k_1$ then $\sigma(k(1)) = 0$, so that $\hat{q}_2(r, 1)$ is indefinitely large and the first event cannot be a decrease in child quality. The length of time until an increase in marriage quality is given by

$$\hat{q}_3(r, 1) = -\frac{\ln(1 - u^{(3)}(r, 1))}{\gamma^+(\theta(1))},$$

while the timing to a decrease in marriage quality is given by

$$\hat{q}_4(r, 1) = -\frac{\ln(1 - u^{(4)}(r, 1))}{\gamma^-(\theta(1))}.$$

¹⁶We have suppressed the parental income arguments, as well as the divorce cost type (b), in the equilibrium investment functions to simplify the notation.

Finally, the length of time until the exogenous termination of the child quality process is given by

$$\widehat{q}_5(r, 1) = -\frac{\ln(1 - u^{(5)}(r))}{\eta}.$$

Which event is actually observed is determined using a competing risks framework, namely, cause φ is observed if

$$\widehat{q}_\varphi(r, 1) = \min(\widehat{q}_1(r, 1), \dots, \widehat{q}_5(r, 1)).$$

Before the second event is generated the state variables are updated as follows. We define the update function for a random variable x from epoch l to epoch $l + 1$ by

$$\Delta(x(l), n) = \sum_{b=1}^B x_{b+n} \chi[x(l) = x_b],$$

where x takes values (x_1, \dots, x_B) and n takes values $\{-1, 1\}$. For example, if the original child quality ($k(1)$) is equal to k_2 , then if the first event is an improvement in child quality, $k(2) = \Delta(k(1), 1) = k_3$. Under this scenario, $\theta(2) = \theta(1)$. At this new quality level, the parents resolve their investment problems and a new equilibrium rate of arrival of quality changes is found,

$$\delta(2) = \delta_0(\theta(2))[\widehat{i}_1(k(2), \theta(2), d = 0, s = 0) + \widehat{i}_2(k(2), \theta(2), d = 0, s = 0)]^\nu.$$

Recall that the value of marriage is enhanced by a higher level of child quality, so that by definition the arrival of a positive shock to k cannot bring about divorce.

If the first event to occur is a decrease in child quality (so that $k(2) = \Delta(k(1), -1)$ and $\theta(2) = \theta(1)$), then the parents will change their quality investment decisions, and, more importantly, may choose to divorce. If they choose to divorce, then the state variable d switches from 0 to 1 and the value of θ is set to 0 in the instantaneous payoff functions for the spouses. Thus the improvement in child quality parameter will either become

$$\delta(2) = \delta_0(\theta(2))[\widehat{i}_1(k(2), \theta(2), d = 0, s = 0) + \widehat{i}_2(k(2), \theta(2), d = 0, s = 0)]^\nu \quad (8)$$

or

$$\delta(2) = \delta_0(0)[\widehat{i}_1(k(2), 0, d = 1, s = 0) + \widehat{i}_2(k(2), 0, d = 1, s = 0)]^\nu. \quad (9)$$

Other parameter changes may occur as well. For example, if the child quality level was initially the second, i.e., $k(1) = k_2$, then a reduction in the child quality level to k_1 implies that the rate of decrease in child quality now becomes 0; if $k(1) > k_2$, then the rate of decrease remains at $\tilde{\sigma}$.

If the first event to occur is an increase in marriage quality, then $\theta(2) = \Delta(\theta(1), 1)$ and $k(2) = k(1)$; such an event cannot result in a divorce. The new hazard rate for an improvement in child quality is given by (8). If, instead, the first event is a decrease in marriage quality, this can clearly result in a divorce. After the appropriate updating, the hazard of an increase will be given by one of the two equations, (8) or (9), depending on whether a divorce occurs.

Finally, if the first event is an exogenous termination of the child quality improvement process, then the state variable s is reset from 0 to 1. In this case, parental investment will remain at the value of 0 for the remainder of this sample path, and the child's state will remain at $k(1)$ forever. When this event occurs, it is possible that the parents could find that the divorce option dominates the marriage option, so that the divorce decision must be analyzed. If the divorce state dominates the marriage state, then the state variable d is set to 1.

The second event that takes place is then determined as follows. Using the new value of δ , $\delta(2)$, the current values of k and θ , $k(2)$ and $\theta(2)$, and the new values of d and s , define the latent time to the next increase in child quality by

$$\widehat{q}_1(r, 2) = -\frac{\ln(1 - u^{(1)}(r, 2))}{\delta(2)},$$

the latent time until the next decrease in child quality by

$$\widehat{q}_2(r, 2) = -\frac{\ln(1 - u^{(2)}(r, 2))}{\sigma(k(2))},$$

the latent time until the next increase in marriage quality by

$$\widehat{q}_3(r, 2) = -\frac{\ln(1 - u^{(3)}(r, 2))}{\chi(d = 0)\gamma^+(\theta(2))},$$

the latent time until the next decrease in marriage quality by

$$\widehat{q}_4(r, 2) = -\frac{\ln(1 - u^{(4)}(r, 2))}{\chi(d = 0)\gamma^-(\theta(2))},$$

and the remaining time until the end of the possibilities of changes in child quality by

$$\widehat{q}_5(r, 2) = \widehat{q}_5(r, 1) - a_1(r),$$

where $a_1(r)$ is the age of the child when the first event occurred in replication r . Note that the denominators of the latent times to a marriage quality change include the indicator function $\chi(d = 0)$ which takes the value of 0 when the parents have divorced, thus allowing no future change in marriage quality. The time of the second event in the r^{th} simulation is given by

$$a_2(r) = \min(\widehat{q}_1(r, 2), \dots, \widehat{q}_5(r, 2)).$$

Further events are generated in a similar manner.

Let us formally define the sample path associated with the r^{th} replication given parameter values Γ , income values y_1 and y_2 , and initial conditions $k(1)$ and $\theta(1)$ by $\varsigma_r(y_1, y_2, b, k(1), \theta(1); \Gamma)$. Let the state of the sample path r , defined in terms of whether the parents of child j are divorced

and the child quality level at ages $a(j, 1), \dots, a(j, G_j)$, be denoted

$$\{k^r(a(j, g)), d^r(a(j, g))\}_{g=1}^{G_j}(k(1), \theta(1), Z_j),$$

where Z_j includes all of the exogenous, time-invariant household characteristics. We generate R sample paths for each household, where each simulation draw r itself consists of $T \times M$ replications, one for each of the T values of $k(1)$ and M values of $\theta(1)$.

As mentioned above, we utilize the method of simulated moments to obtain estimates of the primitive parameters. Let the dimension of Z_j be L . Then for a given initial condition, we can define the conditional expectation function h by

$$E(f_h(Z_j, d(a(j, 1), \dots, d(a(j, G_j), k(a(j, 1), \dots, k(a(j, G_j)|Z_j, k(1), \theta(1); \Gamma), \quad (10)$$

where the conditioning is with respect to household characteristics Z_j and the initial conditions $k(1)$ and $\theta(1)$. We define H conditional expectations functions, where $H \geq NP$, the dimension of the parameter vector Γ .

Given the complexity of the model there exists no closed form expression for (10) in general; we approximate the value of a conditional moment through the use of simulation. Given the R sample paths associated with a given value of the initial conditions, for household j we have

$$\begin{aligned} \frac{1}{R} \sum_{i=1}^R f_h(Z_j, d^r(a(j, 1), \dots, d^r(a(j, G_j), k^r(a(j, 1), \dots, k^r(a(j, G_j)|Z_j, k(1), \theta(1); \Gamma) \\ \equiv A_h(Z_j, a_j, k(1), \theta(1); \Gamma), \end{aligned}$$

where $a_j = (a(j, 1), \dots, a(j, G_j))$.

Now let the probability distribution over the initial values for household x be defined by

$$\omega(k(1), \theta(1)|Z, \Gamma). \quad (11)$$

Then we have

$$A_h(Z_j, a_j; \Gamma) = \sum_{m=1}^M \sum_{t=1}^T A_h(Z_j, a, k_t, \theta_m; \Gamma) \omega(k_t, \theta_m|Z_j, \Gamma). \quad (12)$$

Finally, unconditioning on Z_j , we have

$$A_h(\Gamma) = J^{-1} \sum_{j=1}^J A_h(Z_j, a_j; \Gamma). \quad (13)$$

Note that some moments computed in this way will only be defined for a subset of the sample. For example, one moment may be the difference in test scores for those who took the test twice. In this case, only the subset of observations for which two test scores are available could be included in the computation of this moment. This subsampling does not affect our interpretation of all moments

as representing the population, since we are assuming, and have reason to believe, that the number of test measurements available is exogenously determined.¹⁷

An important component of this specification of the unconditional moments is the distribution over the initial conditions. We assume that $k(1)$ and $\theta(1)$ are independently distributed, and in keeping with our assumption that the state space is finite, assume that both random variables are discrete. We assume that the set of child quality values is $\{1, 2, \dots, T\}$. Let $Z_k \subset Z_j$ be a set of exogenous household-specific variables that influence the distribution of $k(1)$. Then let

$$\omega_k(k(1) = t|Z_k) = \begin{cases} .5[\Phi(\frac{1-Z_k\beta_k}{\sigma_k}) + \Phi(\frac{2-Z_k\beta_k}{\sigma_k})] & t = 1 \\ .5[\Phi(\frac{(t+1)-Z_k\beta_k}{\sigma_k}) - \Phi(\frac{(t-1)-Z_k\beta_k}{\sigma_k})] & t = 2, \dots, T-1 \\ 1 - .5[\Phi(\frac{(T-1)-Z_k\beta_k}{\sigma_k}) + \Phi(\frac{T-Z_k\beta_k}{\sigma_k})] & t = T \end{cases} \quad (14)$$

where Φ is the standard normal c.d.f. Then the probability distribution of the initial condition is parametric, and determined by β_k and σ_k .

The distribution of match values is similarly determined. Each spousal pair draws a match value from a common support of $\{\theta_1, \dots, \theta_M\}$, with $\theta_1 < 0$ and $\theta_M > 0$. Marriage quality values $\{\theta_1, \dots, \theta_M\}$ are located so that $\{\Phi(\theta_1), \dots, \Phi(\theta_M)\} = \{0.1, 0.3, 0.5, 0.7, 0.9\}$. Thus the range of possible marriage qualities is centered at zero. However, the mass of the initial marriage quality distribution need not be centered at zero and is free to favor either positive or negative marriage values. Define $Z_\theta \subset Z_j$ as a set of household characteristics that affect the match value distribution, and define

$$\omega_\theta(\theta(1) = \theta_m|Z_\theta) = \begin{cases} .5[\Phi(\frac{\theta_1-Z_\theta\beta_\theta}{\sigma_\theta}) + \Phi(\frac{\theta_2-Z_\theta\beta_\theta}{\sigma_\theta})] & m = 1 \\ .5[\Phi(\frac{\theta_{m+1}-Z_\theta\beta_\theta}{\sigma_\theta}) - \Phi(\frac{\theta_{m-1}-Z_\theta\beta_\theta}{\sigma_\theta})] & m = 2, \dots, M-1 \\ 1 - .5[\Phi(\frac{\theta_{M-1}-Z_\theta\beta_\theta}{\sigma_\theta}) + \Phi(\frac{\theta_M-Z_\theta\beta_\theta}{\sigma_\theta})] & m = M \end{cases} \quad (15)$$

Then the probability distribution of the marriage quality value is completely determined by $\{\theta_1, \dots, \theta_M\}$, β_θ , and σ_θ . The distributions of $k(1)$ and $\theta(1)$ are chosen to approximate continuous distributions centered at $Z_k\beta_k$ and $Z_\theta\beta_\theta$ and with variances σ_k and σ_θ , respectively.

Calculation of the decision rules used by agents with current state variables $s \in S$ is an extremely time-intensive task, and to compute the moments from the simulated histories requires access to these rules. We have developed a relatively efficient estimation technique for doing so, a discussion of which is contained in Appendix A. In brief, our method involves solving the model and estimating it at the same time, effectively reducing the computational burden of a dynamic model to that of a static model. We adopt a strategy to speed the convergence process which was inspired by the insightful work of Jain, Imai and Ching (2003). They recognized the wastefulness of recomputing decision rules “from scratch” at each new set of trial parameter values as one works through the iterative process to find the parameter estimates. The idea, as implemented here, is to

¹⁷This implies that the distribution of Z should be invariant among subpopulations defined in terms of the number of times the test has been taken. This can be checked using nonparametric methods and sample estimates of the subsample distribution functions of Z .

compute some “exact” solutions to the household’s investment and divorce problem at a fixed set of parameter values, and to approximate the household investment rule as a convex combination of these parameter values, where the weights attached to the rules are a function of the relative distance between the current parameter guesses and the reference parameter vectors. Using the approximate investment rules and the current guesses of the parameters $\tilde{\Gamma}$, we generate simulated moments. We iterate over $\tilde{\Gamma}$ until we adequately approximate the observed sample moments, and call this estimator $\tilde{\Gamma}_1^*$. We find investments over all states s at this value of the parameter vector, and compare these with the investments predicted from the approximation. If the divergence is sufficiently great for any $s \in S$, we add $\tilde{\Gamma}_1^*$ to our collection of parameter vectors with “exact” investment solutions, and restart the iteration process using as starting value $\tilde{\Gamma}_1^*$. We repeat the process until the exact and approximate investment rules at our estimator are sufficiently close over all $s \in S$. We find that this approach performs well in practical terms. It has many desirable properties, including that the precision of the approximated solution increases most over the course of the estimation procedure in the region of the parameter space in which the estimation algorithm searches most intensely.

4.1 Identification Issues

As always, with such a complex model it is difficult to give precise identification conditions for the various model parameters. However, we will attempt to provide a partially heuristic, partially rigorous discussion of the central issues regarding identification in continuous time, point process models of this type.

In a discrete time modeling framework, with multiple observations per individual, it is natural to look at the transition probability matrix as a leading source of identifying information for underlying model parameters. Say that a random variable that assumes B distinct values is measured at two points in time for N independent realizations of the population stochastic process. Then the transition probability matrix has $B(B - 1)$ independent elements, and as $N \rightarrow \infty$, $\text{plim}_{N \rightarrow \infty} \frac{n_{ij}}{n_i} = \Psi_{ij}$, where i is the origin state, j the destination state, $n_i = \sum_{j=1}^B n_{ij}$, and Ψ_{ij} is the true transition probability. Let the vector of primitive parameters be given by Γ , and write the vectorized transition matrix, after omitting redundant elements, as Ψ . Now define a mapping from the primitive parameters to the (vectorized, non redundant) transition probabilities by $\Psi^*(\Gamma)$, and, for simplicity assume that $\Psi^*(\cdot)$ is everywhere differentiable on the parameter space associated with Γ . As is obvious, for Γ to possibly be identified from knowledge of Ψ requires $NP \leq B(B - 1)$. Then we say that Γ is uniquely identified by knowledge of Ψ if Ψ^* is 1-1, for in this case there exists a unique inverse function $\Gamma = (\Psi^*)^{-1}(\Psi)$. This is a very strict notion of identification. Typically we invoke sample identification criteria in practical applications. Given access to a finite amount of information, we only have access to an estimated value of Ψ , which we will denote by $\hat{\Psi}$. Then we may say that Γ is uniquely identified if an objective function such as $(\hat{\Psi} - \Psi^*(\Gamma))'W(\hat{\Psi} - \Psi^*(\Gamma))$ is globally convex in Γ , where W is some positive definite matrix of the quadratic form (which could

depend on Γ as well). We then say that $\hat{\Gamma}$, the argument that minimizes the value of the quadratic form, is the unique estimate of the primitive parameter vector Γ . In many cases involving complex applied models, one may only be able to establish convexity locally.

The above sketch of an idealized problem corresponds roughly to the one we confront, except our problem is more complex on several fronts. For approximately one-fifth of our sample, we only have access to one measurement on the child's test score (i.e., $G_j = 1$). We begin by considering the transition from the origin states (at the time of birth) $(k(1), \theta(1))$ into the states associated with the first sampling point at time 2. For the moment, assume that this sampling time is the same for all sample members.

The first problem we face, and the most critical, is that the states of the process are imperfectly observed. Denote the state vector at the initial date by S_1 and at the subsequent observation time by S_2 , where the states are the MT ($= B$ in the discussion above) possible values of (k, θ) at each moment in time. There are no measurements available at time 1, creating the usual initial conditions problem. Thus, even if Γ were identified from knowledge of Ψ or $\hat{\Psi}$, it is not possible to estimate Γ from readings only on the (partially observable) state vector S_2 . In order to estimate Ψ consistently, even given full observability of S_2 , it is necessary to assume a prior distribution on values of S_1 , $F_1(S_1|\Gamma, Z)$, which are functions of the primitive parameters Γ and perhaps other covariates Z . Then estimation of the transition matrix parameters Ψ is accomplished by writing the probability of S_2 as

$$p(S_2|\Gamma, Z) = \sum_{S_1} \Psi(S_2|S_1; \Gamma) F_1(S_1|\Gamma, Z).$$

Clearly, even if Γ is identified given observability of S_1 , identification of Γ from $p(S_2|\Gamma, Z)$ depends critically on the functional form of $F_1(S_1|\Gamma, Z)$. Moreover, given the lack of observability of S_1 , identification of Γ from $p(S_2|\Gamma, Z)$ will hinge on untestable assumptions regarding F_1 .

We have two test score measures for some of the children in our sample, and there is no reason to believe that the number of measurements is endogenous.¹⁸ For some random sample of children, then, we observe a second measurement, let us say at some common time S_3 . For these children, if S_2 and S_3 were both perfectly observable, we could eliminate the initial conditions problem introduced by not being able to observe S_1 by looking at the transitions between S_2 and S_3 , and consistently estimating the transition matrix $\Psi(S_3|S_2; \Gamma)$. The problem is that the state vector S_t , $t = 2$ and 3, is only partially observable. For example, at measurement time 2, the state vector S_2 includes the indicator variables d_2 and s_2 indicating whether the parents are divorced and whether the investment process is on-going, as well as the child quality level $k(2)$ and the marriage quality level $\theta(2)$. If the parents are married at time 2, we do not observe the match quality $\theta(2)$, nor do we observe whether the investment process is on-going or not. If the parents are divorced at measurement time 2, then the value of the match is normalized to 0, and hence is known, while the

¹⁸According to the NLSY79 Child and Young Adult Users Guide, all children aged 5-15 and older were intended to undergo PIAT mathematics and reading assessments in each survey wave. However, interviewer error led some children of testing age to be skipped. Other children's tests could not be scored due to interviewer scoring decision errors. Computer-assisted interviewing techniques improved compliance from the 1994 wave on.

investment process indicator $s(2)$ is unknown.

Divorced parents at time 2 have the potential to provide a large amount of information regarding the child quality production process, since the transitions of child quality observed for them are not “contaminated” by changes in the unobservable marriage quality process. For these parents, we observe child quality at times 2 and 3, but do not observe whether the investment process has ended as of time 2. If $s(2) = 1$, so the process is over at time 2, then we know that $k(2) = k(3)$. Conversely, $k(2) \neq k(3) \Rightarrow s(2) = 0$, that is, a change in child quality state between measurement times 2 and 3 implies that the child quality investment process had not ended at time 2. Recalling that the timing of the ending of the investment process is strictly exogenous, we can construct a conditional transition matrix $\pi(k(3)|k(2), d(2) = 1, k(3) \neq k(2), \Gamma)$, that is, in fact, only a function of a subset of primitive parameters, those characterizing the parental preferences and the stochastic production technology, that is, only the parameters $\alpha_1, \alpha_2, \nu, \tilde{\sigma}$, and η appear in the equilibrium investment rules in this case. Given the equilibrium investment rules, and the exogenous processes as defined by $\tilde{\sigma}$ and η , the transition probability function for child quality measurements between times 2 and 3 are determined. Movements between child quality states for this set of individuals provide a large amount of information on a relatively small number of parameters.

Differences in test score transition rates between measurement times 2 and 3 for children of divorced and married parents (as of time 2) are important sources of identification of the impact of the marriage quality process on child outcomes and the relationship between child quality and divorce. This can be appreciated most directly if we condition on the event $k(3) \neq k(2)$, for in this case we know that the child investment process is still on-going at measurement time 2. For the same distribution of pre-transfer parental incomes, parental investment rules (and hence transition rates between child quality states) will differ due to (1) income transfers made in the divorced state (which result in different post-transfer income levels in the two cases), (2) different payoffs to child quality in the two marital states (due to the transformation of child-quality from a public to a semi-private good), and (3) differences in the value of θ between the two states (in the divorced state θ is fixed at 0, and in the married state we know the sample path of θ and k has never resulted in a divorce outcome). Since the divorce law environment is what changes the parental income distribution and parental preferences in the two states, and since the environment is assumed known,¹⁹ the differences in the married and divorced parents’ child quality transition rates mainly serves in identifying the role of the θ process in the child quality transition rates.

Up to this point we have ignored the fact that measurements of child quality are obtained at different ages. This nonconcurrency is essentially impossible to treat in a satisfactory manner using a discrete-time framework (see Flinn and Heckman (1982)). The continuous time framework allows us to accommodate any sampling scheme within the estimation process. Moreover, the fact that measurements are taken at different ages across the children in our sample is an asset in terms of identification. To understand why, consider the example of the previous paragraph, where we now

¹⁹In the estimation we use state-level information on child support percentage-based orders. Heterogeneity in this rate across state jurisdictions further aids in identification of the parameters describing the θ process.

denote the transition matrix between measurement periods 2 and 3 by $\Psi(k(3)|k(2), d(2) = 1, k(3) \neq k(2), \Gamma, a(3) - a(2))$, where the additional conditioning argument $a(3) - a(2)$ is the elapsed time between the first test score measurement and the second.²⁰ Variation in the timing of measurements provides identifying information for Γ because, in general,

$$\Psi(k(3)|k(2), d(2) = 1, k(3) \neq k(2), \Gamma, t) \neq \Psi(k(3)|k(2), d(2) = 1, k(3) \neq k(2), \Gamma)\phi(t),$$

that is, the matrix of transition between states is not invariant with respect to the duration between measurements up to a scale normalization $\phi(t)$. Thus by varying the measurement periods, we actually gain more information regarding Γ than when the measurement times are synchronized.²¹

While we have devoted most of our attention to the case in which two test scores are available for the child, in approximately one-fifth of the families only one observation is available. Given the correctness of our functional form assumptions regarding the initial conditions $((k(1), \theta(1)))$, the same general argument applies regarding the information value of having varying ages of first test measurement. Since measurement time 1 for all children occurs at their birth, and is unobserved, transitions between $k(1)$ and $k(2)$ will be functions of the θ process since all parents are married at measurement time 1. Nonetheless, even if no child in the sample had more than one test score measurement, identification of parameters governing the child quality and marriage quality processes could be distinguished in large part due to the assumption that the marriage quality process is the same within all marriages, while the child quality process, being endogenous, is not. Information on divorce outcomes and test scores (even at only one point in time), assumptions of homogeneity in primitive parameters describing the underlying point processes, and functional form assumptions regarding preferences and child quality production, are sufficient to identify all model parameters.

5 Data and Descriptive Statistics

The estimation employs the Child and Young Adult Data associated with the 1979 cohort of the National Longitudinal Survey of Youth. Our sample consists of families with only children in which the parents were married in the first interview after the date of the child’s birth. The use of single-child families allows us to abstract from issues of investment allocation across children and from problems relating to remarriage and step-siblings.

The child outcome measure employed in the empirical analysis is based on the child’s score on the Peabody Individual Achievement Test (PIAT) in mathematics. The PIAT is administered to all children aged five and older in the NLSY Child sample, and is ceased when children exit the

²⁰Recall that the timing of the measurements is considered to be determined in a strictly exogenous manner, so no selection issues arise from the varying number or timing of measurements across sample families.

²¹Another way to think of this is in terms of the aliasing problem in standard stationary time series analysis. Aliasing occurs when a process defined on a given frequency is sampled at a “coarser” frequency (for example, when a process that changes value every month is observed every quarter). In this case, a sample path defined on the coarser frequency is consistent with a large number of sample paths defined on the true frequency, thus leading to a form of nonidentification. By varying the frequencies of the coarse measurements, we have a way to rule out a potentially large class of candidate “true” processes.

Child sample and enter the Young Adult sample at the age of 14. In order to include children who reach the age of five during the sampling window and are born to mothers with as broad a range of ages as possible, we measure the child's age and test scores and the parents' marital status in the first year in which the child undergoes the PIAT mathematics assessment. Thus, our outcome measures for sample children are collected when the children are between the ages of 5 and 14, but in practice almost all of our sample children are first tested between the ages of 5 and 7. This step adds a third sample selection criterion to the requirements that sample children have no siblings and sample parents were married at the children's births. We require that sample children undergo the PIAT mathematics assessment in at least one interview. A follow-up test score is available for more than four-fifths of the families that meet each of these criteria. The test score measure employed in the estimation and policy experiments is the child's age-specific PIAT mathematics percentile.²²

Parental incomes are measured at the date of birth of the child. A common difficulty faced by empirical studies of married and divorced parents' interactions with their children is tracking divorced fathers who no longer reside with their children. The appeal of the NLSY in this regard is that it allows us to observe families from the date of birth of the child, and therefore we have some information on each sample father no matter how quickly the family dissolves after the birth of the child. We measure each father's income at his child's date of birth in order, since the availability of more later measurements depends on the (endogenous) evolution of marital status. Further, we restrict our measure of the mother's and father's incomes to those observed at the date of birth of the child, and we assume that incomes are constant from that date. This step is also taken to avoid relying on income variation in the data that is only observable for families that remain intact. Each parent's income is determined as the sum of reported incomes in the NLSY that are attributable to the individual parent and not to her or his spouse. Attributable income sources are wage and salary, farm and business income, military income, and unemployment income. Regardless of the date of birth of the child, parents' incomes are inflated to 1998 dollars for the purposes of reporting and estimation.

The first (second) divorce outcome measure used in the estimation is zero if parents remain married and not separated from the first interview after the birth of the child through the interview in which the child's first (second) PIAT assessment takes place. Otherwise, it is one. As discussed in the section on estimation, the probability distribution of initial child quality $k(1)$ is permitted to rely on a vector of characteristics of the parents and child observed at (or, in one case, before) the date of birth of the child. Characteristics entering Z_k are a constant, the mother and father's ages at the date of birth of the child, the mother and father's years of schooling at the date of birth of the child and the mother's score on the Armed Forces Qualifications Test (AFQT) administered to NLSY respondents in 1980. Each is standardized relative to the sample distribution in the actual estimation, so that coefficient estimates β_k can be interpreted as the effect of a standard deviation

²²The age norming performed in the 1979 NLSY Child Data uses an age-based norming sample from 1968. Therefore, NLSY sample children show average scores that exceed those of the norming sample. A useful reference on this point is Dunn and Markwardt (1970).

increase in the parental characteristic on the center of the family's initial child quality distribution.

Like the initial child quality, the probability distribution of the quality of the parents' marriage at the child's birth relies on a vector of characteristics of the parents. Characteristics entering Z_θ pertain to the likely stability of the parents' marriage. The first is the mother's response when asked whether the marriage is "very happy, fairly happy or not too happy" at the interview following the child's birth. An indicator for whether her answer is "very happy" covaries positively with the child's PIAT test score in our sample, and mothers who respond with "very happy" are more often married at the child's first PIAT test. As a second measure of marriage quality, we include the number out of a list of ten topics on which the mother states that she and the father sometimes or often argue. The topics are chores and responsibilities, children, money, showing affection, religion, what to do with leisure time, drinking, other women, her relatives and his relatives. This measure shows a substantial negative covariance with children's test scores, and parents who divorce by the test date argue over more of the ten points on average.

The marriage quality questions described here are fielded with the main surveys of the NLSY79 cohort starting in 1988. This poses a further difficulty in our sample construction. In order to include measures of marriage quality, we must observe families between 1988 and 2002. If we retain children born before 1988 in our sample, then we can only observe the quality of their parents' marriages if the marriages survive until the 1988 interview. As with fathers' incomes, this leads to measurement error in an independent variable whose distribution depends on a dependent variable. While the value of meaningful self-reported marriage quality measures to our approach is clear, we find that restricting the sample to families in which the child is born after the 1987 interview decreases our estimation sample from 426 to 202 families. Our solution is to present estimates based on the 426 family sample that do not use marriage quality measures, and to supplement these with estimates based on the 202 family post-1987 sample that do include marriage quality variables. The choice to include marriage quality measures in the only manner we determine to be valid has one other important effect on the smaller sample estimation. The mothers in this sample are older on average than the set of all NLSY mothers when their children are born. Since respondents were aged 14 to 21 at the time of the 1979 wave, between 1988 and 1998 the first time mothers we study are all aged 23 to 40. In this version of the estimation, we measure marriage quality once at the date of the child's birth and include it in the family's initial conditions distribution. From there we assume that marriage quality improvements and setbacks occur exogenously at rates γ^+ and γ^- .

Since we find that the divorce rate among Roman Catholics in the sample is 8.4 percentage points lower than among non-Catholics in our post-1987 sample (4.0 points in our full sample), we also include an indicator for whether the mother reports her religion as Roman Catholic at the start of the NLSY79 panel. Though a positive coefficient on the Catholic indicator would explain the lower divorce rate among Catholics in the sample as the result of a higher overall welfare while in marriage, it is also possible that this distinction results from a greater cost of divorce among Catholics. For this reason we also permit the family's divorce cost type to depend on the mother's religion.

We have imposed a number of stringent selection conditions in defining our sample, primarily regarding family composition. The benefit of these restrictions is that they allow us to abstract from concerns involving the allocation of parental investments across groups of siblings and step-siblings, and from the complex marital status choices facing unwed parents. The cost is that the use of only-child families in which parents were married at their children's births certainly limits the generalizability of our findings to families with more complex structures. In future research, we plan to investigate the role of evolving family structures, particularly fertility, in parents' ongoing child investment decisions.

Table 2 contains the descriptive characteristics of the variables used in the estimation for the two samples. The first panel of the table describes the 202 family, post-1987 sample, and includes marriage quality measures. The second panel describes the full 426 family sample and includes no marriage quality measures. Full sample mothers and fathers are 25.8 and 28.9 years old at the child's birth on average, while post-1987 mothers and fathers are 29.5 and 32.8 years old at the child's birth on average. Comparison of the means of pertinent estimation variables demonstrates several age-related differences in the two collections of families. Divorce is more common for the larger sample, whose children were born to younger parents. While 17 percent of parents in the post-1987 sample are divorced by the child's first PIAT assessment, 32 percent of parents in the full sample are divorced by the first test. These divorced sample families are crucial for identification of the child quality production parameters, as discussed in the previous section. Divorce rates in the roughly two years between tests are higher for the post-1987 sample, however, which shows a 7 point increase in the percent divorced between tests, as compared with a 5 point increase for the full sample.

In each sample the average child is between 5 and 6 years old at her or his first PIAT assessment, and 7 and 8 at her second PIAT. Ninety-five percent of children complete at least one PIAT test by the age of 7. Eighty-one (72) percent of full (post-1987) sample children complete a second PIAT test, most commonly in the subsequent wave two years later. Not surprisingly, parents who were older at their children's births have higher education levels at that time and higher earnings on average when their roughly 5-6 year old children take their first PIAT tests. While full sample mothers and fathers have average incomes of 14,858 and 28,311 1998 dollars at their children's births, and 13.03 and 12.87 years of education, post-1987 sample mothers and fathers earn \$22,103 and 38,130 on average and have 13.86 and 13.77 years of education.²³ Children's average PIAT score percentiles also show the advantage of families that experienced later first births. While the full sample children score at the 55th percentile on average in both assessments, the post-1987 sample children score at the 58th and 62nd percentiles in the first and second assessment on average. Where observed, more than three quarters of mothers report that their marriages are "very happy", and couples argue over an average of 2.56 of the ten points listed above. The argument measure does show meaningful variation, with a standard deviation of 2.03. A third of each sample is Roman

²³It is worth noting that measurement at the interview following the child's birth likely contributes to the large discrepancy between mothers' and fathers' average incomes. This compromise is not ideal but we accept it in response to the data loss associated with any attempts to follow family incomes past the first post-birth interview.

Catholic.

6 Empirical Results

In order to generate the simulated sample paths used in the MSM estimator, we must specify the family law environment each family faces. We fix the policy parameter $\tau_1(1)$ at a value intended to reflect average outcomes in terms of custody/visitation arrangements. In this exercise, it is set at 0.2, so the mother is assumed to be in contact with the child 80 percent of the time in the divorce state. Custody averages over the period in which we observe our NLSY sample have been studied for eight states. All but California maintain approximately 80-20 custody division averages; California's custody decisions favor fathers substantially more than those of other states, with as much as 40 percent custody going to fathers on average.²⁴

The estimation requires data on the child support payment that will be required of the father in the event of divorce, and works on the assumption that parents are able to predict child support arrangements while still married. The history of child support guidelines suggests that predicting child support from the vantage point of marriage was reasonably feasible for most families in our sample. It also provides a second reason to estimate with post-1987 wave births. Two federal reforms, the Child Support Enforcement Amendments of 1984 and the Family Support Act of 1988, established and extended the requirement that states maintain and use set guidelines for the determination of child support. Some states had guidelines in place before the reforms. Therefore most full sample parents, and all post-1987 sample parents, had access to state child support guidelines from the birth of their children. While state guidelines have changed little since the second reform, some variation in the application and formulas of state guidelines precedes the second reform. We find that current formulas match 1988 formulas for most states. Therefore we determine child support rate π for each family based on the parents' incomes and the current child support guidelines for the state in which the family resided at the birth of the child. Child support formulas are coded by hand, and so where obvious differences exist between current guidelines and the historical guidelines to which particular families would have been subject at the time of the child's birth, appropriate corrections have been made. We apply the guidelines for one child families in which the father enjoys 20 percent of physical placement. Though the resulting assignment of projected child support rates is imperfect, we feel it provides a rich and reasonably accurate picture of the cross-state variation in the child support levels that our sample families would have anticipated. Resulting child support rates average 17 percent, and vary from 0 to 30 percent of the father's income.²⁵

Durations are measured in years. The instantaneous discount rate ρ is fixed at 0.05. We assume $\eta = 0.06$, implying an average age for the termination of investment productivity of between 16 and 17. The number of discrete child quality levels, T , is set to 10 and chosen to reflect deciles of

²⁴See, for example, Cancian (1998) on states' custody averages.

²⁵The authors thank Hugette Sun for sharing her extensive research on state child support guidelines. More details on this can be found in Sun (2005).

the age-specific PIAT math score distribution. We assume that there are $M = 5$ marriage quality levels, and that the probability mass at each level is determined in part by family characteristics and approximates a discrete normal distribution as described above by $\omega(\theta = \theta_m | Z_\theta)$. Parents' incomes are measured in units of \$2500 1998 dollars. Simulated moments are based on $R = 100$ replications per family per $(k(1), \theta(1))$ pair, or 5000 replications per family.

The model relates exogenous household characteristics $X = \{y_1, y_2, a, b, \theta(1), k(1)\}$ to outcomes k and d for a given family. Therefore the moments we choose pertain to the relationship between parents' incomes, children's test ages, determinants Z_θ of initial marriage quality, determinants Z_k of initial child quality, and the child's test score; the relationship between parents' incomes, children's test ages, determinants Z_θ of marriage quality or determinants Z_k of child quality and the parents' marital status at the test; the test and marital status outcome averages for the full sample; and higher-order interactions among the two outcome measures and elements of X and Z . The moments we have selected as the basis for our estimator differ across the two samples due to the fact that self-assessed marriage quality measures are only available in the small sample. While we use 32 moments as a basis for the estimation, four of these involve marriage quality measures and therefore only 28 of the 32 moments are relevant for the full sample. The moments that we attempt to fit are described in Table 4. Note that we choose to measure and simulate unconditional moments in almost every instance, due to the complication associated with simulating and evaluating conditional moments across family-initial condition combinations that are each associated with unique weights. However, the moments chosen contain information equivalent to conditional moments where, for example, we compare moments that condition on marital status to unconditional moments involving products of d and k or elements of X .

The vector of parameters estimated using our MSM procedure govern the parents' preferences, the production of child quality, the relationship of observed marriage quality measures to true marriage quality and the relationship of characteristics of the parents observed at or before the child's birth to the child's initial quality. The complete vector of parameters we estimate is $\Gamma = \{\alpha_1 = \alpha_2, \tilde{\delta}, \nu, \sigma, b, \beta_\theta, \beta_k, \beta_r, \delta_0, \sigma_k, \gamma^+ = \gamma^-\}$. Using the post-1987 sample, with its larger amount of information on marriage quality, Γ contains 18 free parameters, while using the full sample, with less information on marriage quality, Γ contains 16 free parameters.

The parameter estimates are reported in Tables 3a and 3b. We begin by briefly considering the estimates for the smaller sample, which are reported in Table 3a. The point estimate of ν is 0.786, and, given its estimated standard error, we can conclude that the stochastic production function is strongly concave in total parental investments in the child. The precisely estimated $\tilde{\delta}$ estimate of 0.695, together with our estimate of ν , indicates a high degree of efficiency of parental investments in the production of child achievement. Our estimate of σ is 0.098, which indicates a decrease in child quality only every 10 years on average. Thus child quality improvements are relatively irreversible. By the same token, marriage quality shocks are relatively rare; on average, both improvements and setbacks hit θ about every eight years (1/0.13).

The estimated parental preference weights associated with own consumption are extremely large,

at 0.962 for each of the parents (recall that these are restricted to be equal). The size of these weights can be attributed to our lack of direct consumption and investment information. The rate of improvements in child quality are determined by the properties of the stochastic production function and the level of inputs supplied by the parents (in this case restricted to monetary investments). In order to fit a given rate of increase in the sample, either a low level of investments can be very efficient or a higher level of investments could be less efficient. In this case, a high marginal productivity of investment is paired with a low investment level, which is obtained by attributing little interest in child quality on the part of the parents. Utilization of subsidiary data sources on consumption expenditures would enhance identification and produce more reasonable estimates of parental preferences. The estimates of $\alpha_1 = \alpha_2$ are much more reasonable, happily, when the larger sample is used. Nonetheless, identification remains fragile in that case as well.

In terms of the point estimates of β_k , most results are as might be expected. The probability of a high initial quality draw is increasing in mother's and father's age and education, though decreasing in the mother's AFQT score. While we would probably expect that all of these coefficients should be positive, clearly the mothers's AFQT score and educational attainment are strongly related. The parameters characterizing the linear index are all reasonably precisely estimated, which is somewhat unexpected given the fact that the random variable $k(1)$ is never observed. To some degree this is due to the fact that measured test scores are positively associated with all of these characteristics. Given that these characteristics are only allowed to influence outcomes through their effect on the distribution of the initial draw within our model, these strong results are largely due to modeling assumptions.

Our estimates of β_θ indicate that self-reported marital problems decrease the likelihood that the marriage match value at the time of the birth was high. Roman Catholics are likely to have higher initial match values, which is consistent with our result that Roman Catholics are more likely to have high divorce costs. (If this is the case, they should be more selective when choosing a partner and having a child.) Unexpectedly, the number of points on which parents argue is positively associated with having a high initial match value. We don't wish to make too much of this result, but this could indicate that the marriage is sufficiently strong that the partners feel comfortable discussing a wide range of family and personal issues.

The flow cost of divorce for the high divorce cost cases is estimated to be 0.513. A Roman Catholic has a probability of 0.827 of being a high divorce cost type, so that the net difference in probabilities between Catholics and non-Catholics is $0.827 - 0.5 = 0.327$. This, combined with the higher initial θ values, generates the lower divorce rate among Roman Catholics that is observed in the data.

Table 3b estimates, from the sample that is over twice as large as that used to generate the estimates in Table 3a, show some distinct differences. The estimate of the curvature in the production function, $\hat{\nu} = 0.608$, is significantly less than the corresponding estimate in Table 3a (thus indicating more curvature). Further, the coefficient on the child quality update rate, $\tilde{\delta}$, is considerably lower for this sample, at 0.268. The estimate of parental consumption weight, $\hat{\alpha}_p = 0.761$,

is more reasonable than what we observed in the previous table. As mentioned above, even if child improvement rates were roughly the same in the two samples, the estimates in Table 3b imply parents care more about child quality but their investment is less productive.

In terms of the other rate parameters, we see the estimated rate of decrease in child quality has dropped to a point that implies a decrease every 18 years. The rates of shocks to marriage quality have increased markedly, instead, to $\hat{\gamma}^+ = \hat{\gamma}^- = 0.335$, implying an improvement and a setback to marriage quality every three years, on average.

The estimates for the linear index function appearing in the initial conditions distribution for child quality are now all in line with expectations. Mother’s and father’s age and education positively affect the likelihood of a high draw of $k(1)$, as in Table 3a, but now also mother’s AFQT is positively associated with a good outcome as well. In terms of the initial conditions distribution for θ , we can now only estimate a coefficient associated with being a Roman Catholic, which, as before, implies a positive association with the likelihood of having a high $\theta(1)$ draw.

The estimated value of b is very high, at 3.718, for this sample, which is somewhat surprising given that the proportion divorced is much higher in this sample than in the smaller one. Of course, all other parameter estimates are not held constant, so this is not a valid comparison. The probability of being a high divorce type is now much less strongly related to the mother being a Roman Catholic. The net difference in this probability is only 0.067, though the Roman Catholic indicator coefficient in the logit is still estimated to be significantly greater than zero.

The 32 data and simulated moments listed in table 4 give an idea of the fit of the model. Overall, the simulated moments match the patterns in the data reasonably well in each sample. The data and simulation values of the first, third and fourth moments listed in the table, the divorce rate by the date of the child’s test, the average of kd , and the average of $k(1 - d)$ at the test date, together indicate that we have fit the two most obvious targets of the estimation reasonably well. The divorce rate in the post-1987 sample is 17.32 percent, while the simulated divorce rate is 17.98 percent. The data and simulated divorce rates for the full sample are 32.31 and 33.39 percent. Divorces in the model occur as a result of marriage quality shocks, child quality setbacks and productivity terminations, and are of course held off by the endogenous progress of child quality. In each sample the collection of parameters governing the marriage and child quality processes produce a good fit to the observed divorce rates.

The overall average test score in the post-1987 sample is 58.26; we simulate an overall average test score, implied by moments three and four, of 58.19. In the full sample, the overall average test score is 55.89 and the simulated overall average test score is 54.27. Further, the simulated moments replicate the differences in test scores between children with married and divorced parents quite closely. The ability of the model to match these differences provides some encouraging feedback regarding the assumptions on the structure of the child quality production function and its relationship to the marriage state.

Since most simulated moments accurately match their sample analogues, from here we focus on what appear to be the biggest misses. Moments 14 and 15 are the expected differences in fathers’

and mothers' incomes, respectively, from marriage to divorce. While the fit of moments 14 and 15 in the post-1987 sample is quite good, the model underpredicts the large income advantage of married mothers and fathers over divorced mothers and fathers in the full sample. The reason for the difference in the success of the model at predicting the dependence of realized divorce on parents' incomes in the two samples is unclear. It is certainly the case that there is much greater income variation in the full sample, which includes many very young parents, and therefore the task of fitting the income difference from marriage to divorce could well be more difficult in the case of the full sample. The reader should keep this shortcoming in the full sample estimation in mind when considering the point estimates above and the policy analysis that follows.

Other simulated moments that differ from their corresponding data moments by a fairly large proportion of the value of the relevant data moment differ across samples and are generally among moments 22 to 32. Moments 22 to 32 are each based on differences between children's first and second test scores, which are then interacted with other pertinent variables. Since the test score differences may take values from -98 to 98, the range of values assumed in the data and simulations is quite large for each of the moments and always includes zero. In this light, we see that the data and simulated values of moment 27 in the post-1987 sample, for example, imply that the true average test score change among children whose mothers report "very happy" marriages is +5.2, with possible values from -98 to +98. We simulate an average test score change for this group of +3.4, again with possible values from -98 to +98. Though the magnitude of the simulated measure is only about 64 percent of the magnitude of the true measure, with complete understanding of the meaning of each we see that the fit of moment 27 is quite reasonable. Similar arguments regarding the scale and range of the outcome variables that constitute moments 22 to 32 in both samples demonstrate that their simulated values are in fact quite reasonable, with a particular exception. Moments 25 and 31 each describe the relationship between the child's test score gain from test 1 to test 2 and the father's income. Using each sample, a comparison of the data and simulated values of moments 25 and 31 reveals that the model underpredicts the dependence of the child's test score gain on the father's income. It does not, however, underpredict the dependence of the child's first test score on the father's income or the dependence of the child's test score gain on the mother's income.

Recall that the estimates are generated under the restriction that $\alpha_1 = \alpha_2$. One concern is whether the model is able to generate the observed relationships between the incomes of mothers and fathers and children's test scores across marriage and divorce states without relying on a difference in the tastes of mothers and fathers for child quality. Beyond preferences, mothers and fathers differ in the model in their individual incomes and in their treatment by existing policy in the event of divorce. Moments that particularly speak to our success in fitting the described relationship given the preference restriction in question include $E(\text{test score} \times \text{father's income} \times \text{married})$, $E(\text{test score} \times \text{mother's income} \times \text{married})$, $E(\text{test score} \times \text{father's income} \times \text{divorced})$, $E(\text{test score} \times \text{mother's income} \times \text{divorced})$, and the four analogous moments in which the test score is replaced by the test score change. Among these eight relevant moments, the reader will

observe that seven are fit fairly well by the model simulations for each sample. The eighth moment, moment 31, discussed above, reflects the under-prediction of the connection between children’s test score gains and fathers’ incomes in marriage. While we feel that we do a surprisingly good job of fitting the income-test relationships in the data under the assumption that parents have identical tastes for child quality, relaxing the $\alpha_1 = \alpha_2$ requirement might allow us to match moment 31, and therefore better understand the role of the father’s investments in marriage. We are, nonetheless, reasonably satisfied with the ability of the differences in parents’ incomes and legal treatment, absent different preferences, to fit the patterns in the data.

7 Custody and Child Support Experiments

7.1 Determination of State Welfare Weights

Recent law changes and social movements in U.S. states and in western Europe have advocated shared custody and placement or more moderate increases in fathers’ access to their children in divorce. Over the past several years, fathers’ groups in the UK and US such as Fathers 4 Justice, the American Coalition for Fathers and Children, Dads Against Discrimination and the Alliance for Noncustodial Parents’ Rights have agitated for shared physical custody.²⁶ Major law changes from 2000 to 2004 in Iowa, Maine, Wisconsin and Austria, among others, encourage judges to grant joint physical custody, or to divide the child’s time between the two parents as close to equally as possible. In Wisconsin court record data, Cook and Brown (2005) show that the 2000 Wisconsin law change was followed by a continuing upward trend in the rate of joint placement. Early results from the new and ongoing custody research of Atteneder, Boheim, Buchegger and Halla (2005) suggest that, along with increasing the number of joint custody arrangements, the 2001 Austrian custody law reform has had an effect in practical terms on the time that children spend with their non-resident parents. Though placement shares of mothers and fathers averaged 80 and 20 percent, respectively, in most U.S. states over the period covered by our estimation sample, we look forward to the availability of data on children’s academic or behavioral performance following relevant custody law changes.²⁷ For the time being, the structure of the model allows us to perform some initial analysis of the role of placement by including $\tau_1(1)$ in the set of policy parameters that may be manipulated by state legislators.

Child support guidelines required by U.S. federal statute to be in effect by 1988 at the latest in all 50 states allow us to use π as a second choice variable for policymakers in the calculation of states’ social welfare weights, as described in section 3. There is extensive variation in the structure of state guidelines, including but not limited to the existence and level of personal allowances, whether support orders depend on the income of the custodial parent, the shape of the support level as a

²⁶ A description of their activities can be found in Dominus (2005), among other places.

²⁷ The NLSY’s 2002 wave follows the Maine and Wisconsin custody law reforms by a year and two years, respectively. However, NLSY79 cohort children are fairly old by the 2002 wave on average, and in order to study the effects of the custody law change directly we would require a longer (or larger) post-reform panel in order to observe a sufficient number of post-reform divorces.

function of the non-custodial parent’s income, and the dependence of the support level on the number of children.²⁸ Given the simplicity of the policy choices we consider in section 3, we limit our discussion of states’ social welfare weights to those states represented by large numbers of families in our sample and in which the child support guidelines imply a roughly linear dependence of support on the father’s income. Of our 426 larger sample families, 53 reside in California at the time of the child’s birth, 48 in Texas, 29 in Florida, and 24 in New York. Published child support guidelines for California, Texas, and New York indicate approximately flat child support rates of .25, .20, and .17, respectively, for our sample families.²⁹ Florida’s guidelines, on the other hand, produce a clearly nonlinear dependence of child support on father’s income, with rates ranging from 11 to 20 percent in our data. For these reasons, we consider social welfare weights a_1 , a_2 , and a_c in California, New York, and Texas.

We define $\hat{w}_j(X_i; F)$ for parent $j = 1, 2$ to be parent j ’s expected value of married parenting at the birth of the child. This value is approximated numerically using the model structure and the MSM estimates for the full sample, in which no marriage quality measures are used. The expected value integrates over the initial marriage and child quality distributions using the initial conditions parameter estimates and relevant family characteristics contained in X_i . Our choice of welfare measure for the child is somewhat more complicated. We consider the child’s terminal quality to be her relevant welfare measure although, as stated in section 3, many reasonable child welfare measures might be considered. Our reasoning is that children’s achievements are valued by children and parents to the extent that they presage adult success. As was argued in footnote 12, the child’s final “quality” level can be viewed as an important initial condition in determining the utility attained during their adult life. Thus, the terminal value of k should be positively related to their adult welfare, though the exact nature of the mapping is unknown. For this reason, the comparison of parental and child welfare levels and social welfare weights is on shakier ground than is the comparison of welfare outcomes and welfare weights between mothers and fathers.

In order to scale child welfare in a manner consistently with how it enters the parents’ utility functions, we use the log of the child’s quality, and assume that the fixed terminal quality, and then compute the present value of this quantity by dividing by ρ . Formally,

$$\hat{w}_c(X_i; F) = \sum_{m=1}^M \sum_{t=1}^T L^{-1} \sum_{l=1}^L \rho^{-1} \ln(k(\tau(X_i, F, \Lambda_{m,t,l}))),$$

where $\Lambda_{m,t,l}$ represents sample path l originating from initial condition pair (θ_m, k_t) and τ is the time at which the child quality improvement process ends. The expected child welfare used in the calculation of the following social welfare weights is based on $L = 50$ simulated terminal child qualities for each initial conditions pair, or $LMT = 2500$ replications per family.

Table 5 reports the calculated weights on mothers’, fathers’, and children’s welfare levels in

²⁸See, for example, Aizer and McLanahan (2006) and Sun (2005).

²⁹This is true despite a particularly complicated formula for California that includes the incomes of both parents and the determination of "disposable income".

the social welfare functions of policymakers in California, Texas, and New York. Child support and custody allocations clearly have a strong redistributive character, and this is evident in our approximated state welfare weights. While California’s child support standard is relatively generous to the mothers in our sample, New York’s standard is relatively stingy.³⁰ The Texas rate lies in between. Though the calculated welfare weights depend not only on state child support rates but also on the characteristics of sample families residing in the states, we do find that inferred state weights on mothers’ welfare increase with state child support rates from a weight of 0.221 on mothers in New York up to a weight of 0.509 in California. Fathers’ welfare weights are never substantially below mothers’ weights, and are significantly higher in places; the father’s social welfare weight in New York is 0.482, as compared with mothers’ 0.221. The consistently high weight that states place on fathers’ welfare is in line with the findings of Flinn and Del Boca (1995) regarding Wisconsin family court judges’ apparent preference for fathers. New York policymakers are estimated to place a weight of 0.297 on the outcomes of children, while the preference weights on child outcomes in both Texas and California are approximated to be very close to zero. As we have argued above, one rationale for this finding is that the actual welfare outcome for children is substantially undervalued by using only the terminal value of k , which is only an indicator of their lifetime welfare outcome. More cynically, but perhaps realistically, it may be the case that agents setting the family law environment primarily represent tax-payers and voters.

In a second series of policy experiments, we determine welfare $\hat{w}_j(X_i; F)$ for $j = 1, 2, c$ using all $i = 1, \dots, 426$ full sample families over a 42 point grid on policy parameters π and $\tau_1(1)$. Our grid, $\pi \in \{0, .05, .1, .15, .2, .25, .3\}$ by $\tau_1(1) \in \{0, .1, .2, .3, .4, .5\}$, is intended to represent the empirically plausible region of the policy space. Child support orders for fathers of only children based on state guidelines range from roughly 0 to 30 percent of fathers’ incomes in our data, and even the most extreme of the fathers’ rights advocates mentioned above argue for no more than 50-50 child placement. Table 6 shows the policy combination that maximizes the welfare of each category of agents: mothers, fathers, and children. It also reports the percentage difference between the approximated welfare at the particular agent’s optimum and the average approximated welfare for the agent over all 42 policy grid points.

Not surprisingly, the welfare of mothers and fathers is optimized at very different policy combinations, suggesting that child custody and support are largely redistributive in nature. While fathers prefer 50-50 custody and no child support, mothers prefer 80-20 custody and 30 percent child support. Note that where fathers are best off at the extremes, preferring the highest paternal custody and the lowest child support included in the grid, mothers are best off at an interior paternal custody level. Clearly mothers benefit from fathers’ investment in children, which is increased by positive levels of paternal custody in divorce. The welfare redistribution effected by child support and custody changes over what we determine to be the plausible range is large. The welfare of both fathers and mothers is about 7 percent higher under their preferred policies than its average

³⁰It may be worth noting here that many states in our sample maintain less generous child support guidelines than those of New York; Idaho is one easy example. However, few sample families reside in each of these states and therefore we are unable to perform comparable welfare calculations for them.

value across all 42 points of the grid.

We simulate both terminal child quality and divorce rates on the date at which the terminal quality is reached in order to study the effects of our 42 policy combinations on children’s welfare and divorce rates. Children are best off where fathers are best off in our simulations, under 50-50 custody and no child support. The variation in children’s welfare over the 42 grid points is considerably smaller than the variation in mothers’ and fathers’ welfare. Children’s welfare under the optimal policy is only about 3 percent above its average across the full grid. Further, some policy combinations that include both substantial child support and substantial paternal custody generate child welfare levels very close to that of the optimum. Divorce rates are minimized at 20 percent child support and 30 percent paternal custody, a policy arrangement that does not exactly match any party’s preferred policy but is closest to the mother’s preferred policy. Simulated divorce rates show less variation across the policy grid than any of the three agents’ welfare levels, with the minimum divorce rate being less than one percent below the average of the divorce rates over the grid. Despite the current emphasis on divorce rates as a measure of the effectiveness of family law, our model and estimates predict that children are not best off under the divorce-minimizing policy.

In general, we find that the variation in children’s expected welfare with changes in policy parameters π and $\tau_1(1)$ is quite small in both absolute and percentage terms. The variation in parents’ welfare with changes in π and $\tau_1(1)$, however, can be quite large in percentage terms. Observed welfare changes for parents 1 and 2 are opposite and approximately offsetting in percentage terms for every shift in π and $\tau_1(1)$ that we have considered. In sum, the policy experiments generate a substantial redistribution from mothers to fathers as paternal custody shares increase and child support obligations decrease. The magnitudes of the proportional changes in child quality and parental welfare suggest that we consider effects on the distribution of resources between parents to be a key concern of divorce policy. Further, the experiments demonstrate that under the best-fitting parameterization of the model children’s attainments are not necessarily greatest where the divorce rate is minimized.

8 Conclusion

We have developed and estimated a continuous time model that allows for strategic behavior between parents in making child quality investment choices 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 Mother-Child subsample of the NLSY to estimate model parameters using a relatively involved 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 adequate to good.

The most important contribution of our work is to the understanding of the dynamic relationship

between divorce decisions and the evolution of child quality, and the dependence of this process on family law parameters. We have conducted some initial investigations of how substantial changes in these parameters - those reflecting contact time between divorced parents and the child and the child support transfers between parents - impact the parental welfare distribution and the child quality outcomes. To date, our experiments suggest relatively small, but noticeable, impacts of changing the family law environment on the average value of child quality in the population. Instead, the concurrent impact on the welfare distribution of parents is substantially greater. Such a result may suggest a rationale for why changes in family law tend to occur very gradually over time. While “better” family law environments may favorably impact the child outcome distribution, the gains are slight compared to the shifts in the parental welfare distribution. It follows that it may be difficult to attain the wide-spread support from both mothers and fathers that radical changes in family law require.

Though complex, the model is very stylized and it seems important to generalize it along several dimensions in order to bolster the credibility of our policy experiments. As we have mentioned in passing, it would be highly desirable to allow for endogenous fertility decisions, and this is one of our current research directions. We are less convinced that allowing for cooperative behavior on the part of parents will have substantive impacts on our results, but for theoretical reasons believe that it would be a reasonable modeling choice. Perhaps the most daunting task we face is to develop a reasonable set of measures of child quality stretching over the period from birth to young adulthood. The test scores we use are clearly a simplistic measure of child “quality”. We must be able to splice together various measures of child performance over the development period if we are to adequately characterize the long run relationships between the household environment and the growth process.

References

- [1] Aiyagari, S. R., J. Greenwood and N. Guner (2000), "On the State of the Union," *Journal of Political Economy* 108, 213-244.
- [2] 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.
- [3] Atteneder, C., R. Boheim, R. Buchegger and M. Halla (2005), "The Reform of the Austrian Custody Law and its Influence on the Custody Decision After Divorce," manuscript, Johannes Kepler University.
- [4] Bernal, R. (2006), "The Effect of Maternal Employment and Child Care on Children's Cognitive Development," manuscript, Northwestern University.
- [5] Bernal, R. and M. P. Keane (2006), "Child Care Choices and Children's Cognitive Achievement: The Care of Single Mothers," manuscript, Northwestern University.
- [6] Brien, M. J., L. A. Lillard and S. Stern (forthcoming), "Cohabitation, Marriage and Divorce in a Model of Match Quality," *International Economic Review*.
- [7] Browning, M. and P.-A. Chiappori (1998), "Efficient Intra-Household Allocations: a General Characterization and Empirical Tests," *Econometrica* 66, 1241-1278.
- [8] Cancian, M. and D. Meyer (1998), "Who Gets Custody?" *Demography* 36, 147-157.
- [9] Center for Human Resource Research (2000), *1998 Child & Young Adult Data Users Guide*, Columbus, Ohio.
- [10] Chiappori, P.-A., B. Fortin and G. Lacroix (2002), "Marriage Market, Divorce Legislation and Household Labor Supply," *Journal of Political Economy* 110, 37-72.
- [11] Cook, S. T. and P. Brown (2005), "Recent Trends in Children's Placement Arrangements in Wisconsin," Report to the Wisconsin Department of Workforce Development, University of Wisconsin.
- [12] Del Boca, D. and C. J. Flinn (1995), "Rationalizing Child Support Decisions," *American Economic Review* 85, 1241-1262.
- [13] Dominus, S. (2005), "The Fathers' Crusade," *New York Times*, May 8, 2005.
- [14] Dunn, L. M. and F. C. Markwardt, Jr, (1970), *Peabody Individual Achievement Test Manual*, Circle Pines, Minnesota: American Guidance Service, Inc.
- [15] Flinn, C. (2000), "Modes of Interaction Between Divorced Parents," *International Economic Review* 41, 545-578.
- [16] Flinn, C. and J. Heckman (1982), "Models for the Analysis of Labor Force Dynamics," in *Advances in Econometrics, Volume 1* (R. Basman and G. Rhodes, eds.). Greenwich, CT: JAI Press.
- [17] Friedberg, L. (1998), "Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data," *American Economic Review* 88, 608-627.

- [18] Gruber, J. (2004), "Is Making Divorce Easier Bad for Children? The Long Run Implications of Unilateral Divorce," *Journal of Labor Economics* 22, 799-833.
- [19] Haveman, R. and B. Wolfe (1995), "The Determinants of Children's Attainments: A Review of Methods and Findings," *Journal of Economic Literature* 33, 1829-1878.
- [20] Imai, S., N. Jain, and A. Ching (2003), "Bayesian Estimation of Dynamic Discrete Choice Models," Northern Illinois University Research Papers in Economics #2003-19.
- [21] Jovanovic, B. (1979), "Job Matching and the Theory of Turnover," *Journal of Political Economy* 87, 972-990.
- [22] Liu, H., T. Mroz and W. van der Klaauw (2003), "Maternal Employment, Migration and Child Development," manuscript, University of North Carolina.
- [23] Lundberg, S. and R. A. Pollak (1994), "Non-cooperative Bargaining Models of Marriage," *American Economic Review Papers and Proceedings* 84, 132-137.
- [24] Pakes, A. and P. McGuire (2000), "Stochastic Approximation for Dynamic Analysis: Markov Perfect Equilibrium and the 'Curse' of Dimensionality," manuscript, Harvard University.
- [25] Sun, H. (2005), "The Effect of Child Support Guidelines on the Probability of Divorce," manuscript, University of Wisconsin.
- [26] Udry, C. (1996), "Gender, Agricultural Production, and the Theory of the Household," *Journal of Political Economy* 104, 1010-1046.
- [27] Weiss, Y. and R. J. Willis (1985), "Children as Collective Goods and Divorce Settlements," *Journal of Labor Economics* 3, 268-292.

A Estimation Algorithm

Let the number of parameter vectors at which exact solutions are computed be given by H , and let the collection of these parameter vectors be given by $\Lambda = \{\Gamma_1, \Gamma_2, \dots, \Gamma_H\}$, where each $\Gamma_h \in \Omega_\Gamma$, the parameter space associated with Γ . The Nash equilibrium investment rules for the household are given by $i^*(s; \Gamma_h) = i_1^*(s; \Gamma_h) + i_2^*(s; \Gamma_h)$ at the parameter vector Γ_h . Let the true value of the parameter vector be given by Γ_0 . Both Γ_0 and $\tilde{\Gamma}$ are interior points in the K -dimensional parameter space Ω_Γ . Estimation proceeds as follows.

1. Begin by selecting H distinct points in the parameter space Ω_Γ , which we denote by Γ_h , $h = 1, \dots, H$, with the collection of these points defined as Λ . For these H values of the parameter vector we solve for the investment rules for all values s in the finite state space S .
2. Given any current guess of the values of the parameters $\tilde{\Gamma}$, compute the weights

$$w_{\tilde{\Gamma}}(h) = \frac{[D(\tilde{\Gamma}, \Gamma_h)]^{-1}}{\sum_{h=1}^H [D(\tilde{\Gamma}, \Gamma_j)]^{-1}}, \quad (16)$$

where $D(x, y)$ is a distance function so that $D(x, y) = D(y, x)$, $D(x, y) > 0$ for all $x \neq y$, and $D(x, x) = 0$. As a result, $w_{\tilde{\Gamma}}(h) \in [0, 1]$, $\forall h$, and

$$\sum_{h=1}^H w_{\tilde{\Gamma}}(h) = 1, \quad \forall \tilde{\Gamma} \in \Omega_\Gamma. \quad (17)$$

3. Form the approximate decision rules for every value of s ,

$$\hat{i}^*(s; \tilde{\Gamma}) = \sum_{h=1}^H w_{\tilde{\Gamma}}(h) i^*(s; \Gamma_h). \quad (18)$$

4. Generate the simulated moments at the parameter vector $\tilde{\Gamma}$ using the approximate decision rules $\hat{i}^*(s; \tilde{\Gamma})$.
5. Define the distance function

$$L_1(\tilde{\Gamma}; A^N) = (A^N - \hat{A}(\tilde{\Gamma}))' W (C^N - \hat{A}(\tilde{\Gamma})), \quad (19)$$

where A^N are the sample moments, $\hat{A}(\tilde{\Gamma})$ are the analogous moments computed from the simulated sample at the parameter vector $\tilde{\Gamma}$, and W is a positive-definite weighting matrix.

6. Using the Nelder-Mead simplex algorithm, repeat steps (2)-(5) until

$$L_1(\tilde{\Gamma}; A^N) < \varepsilon_N, \quad (20)$$

where ε_N is a small positive number.

7. Denote the value of $\tilde{\Gamma}$ that satisfies (20) by $\tilde{\Gamma}_1^*$, where the subscript ‘1’ suggests that this is an estimator that has passed the first convergence criterion.
8. Compute the optimal investments at $\tilde{\Gamma}_1^*$ for each $s \in S$. Define

$$L_2(\tilde{\Gamma}_1^*) = \max_{s \in S} \{|i^*(s; \tilde{\Gamma}_1^*) - \hat{i}^*(s; \tilde{\Gamma}_1^*)|\}_{s=1}^S. \quad (21)$$

9. If $L_2(\tilde{\Gamma}_1^*) < \zeta_N$, where ζ_N is a small positive number, then we say that the final estimator of Γ is

$$\tilde{\Gamma}_2^* = \tilde{\Gamma}_1^*. \quad (22)$$

If not, then add the point $\tilde{\Gamma}_1^*$ to the set Λ (or $\Lambda' = \Lambda \cup \tilde{\Gamma}_1^*$) so that the cardinality of this set increases to $H + 1$. Then repeat all steps beginning with (2), keeping the current guess of the parameter vector fixed at $\tilde{\Gamma}_1^*$.

In practice we have had good success with this estimation method. At this point we cannot supply a formal proof of consistency of this estimator, but we turn to a sketch its elements.

First consider the approximation of the investment rule as a function of the parameter vector Γ . Given our H element set Λ , for a given value of $\Gamma \in \Omega_\Gamma$, we compute

$$\hat{i}^*(s; \Gamma) = \sum_{h=1}^H w_\Gamma(h) i^*(s; \Gamma_h).$$

If

$$\max |i^*(s; \Gamma) - \hat{i}^*(s; \Gamma)| \geq \zeta_N,$$

then we add the point Γ to the set Λ as element $H + 1$ of the set. If not, we say that we have adequately approximated the decision rule.

If the convergence criterion is not satisfied, we return to recompute the weights attached to the “exact” investment rules associated with the new set of points $\Lambda' = \Lambda \cup \Gamma$. The weight attached to any arbitrary evaluation point h can be expressed as

$$\begin{aligned} w_\Gamma(h) &= \frac{[D(\Gamma, \Gamma_h)]^{-1}}{\sum_{j=1}^{H+1} [D(\Gamma, \Gamma_j)]^{-1}}, \quad h = 1, \dots, H + 1 \\ &= \frac{\frac{1}{D_h(\Gamma)}}{\frac{1}{D_1(\Gamma)} + \frac{1}{D_2(\Gamma)} + \dots + \frac{1}{D_{H+1}(\Gamma)}} \\ &= \frac{\frac{1}{D_h(\Gamma)}}{\frac{D_{-1}(\Gamma) + D_{-2}(\Gamma) + \dots + D_{-(H+1)}(\Gamma)}{D_1(\Gamma) D_2(\Gamma) \dots D_{H+1}(\Gamma)}} \\ &= \frac{D_{-h}(\Gamma)}{D_{-1}(\Gamma) + D_{-2}(\Gamma) + \dots + D_{-(H+1)}(\Gamma)}, \end{aligned}$$

where $D_j(\Gamma)$ is shorthand for $D(\Gamma, \Gamma_j)$ and

$$D_{-j}(\Gamma) = D_1(\Gamma) \cdots D_{j-1}(\Gamma) D_{j+1}(\Gamma) \cdots D_{H+1}(\Gamma).$$

But note that in this case $\Gamma = \Gamma_{H+1}$, so $D_{H+1}(\Gamma) = 0$ and $D_j(\Gamma) > 0, \forall j \neq H+1$, since all points of evaluation are distinct. Then $D_{-(H+1)}(\Gamma) > 0$, while $D_{-j}(\Gamma) = 0, \forall j \neq H+1$. Thus $w_\Gamma(H+1) = 1$, and the new “approximate” decision rule is the “exact” one computed at the point Γ , or

$$\hat{i}^*(s; \Gamma) = i^*(s; \Gamma), \quad \forall s.$$

This completes the discussion of the ability of the investment rule approximation method to fit the actual investment rule solution for every value of s . While it is always capable of providing a perfect fit, we will not want to enforce this in practice since this would imply an indefinite number of iterations over steps (2)-(5). For consistency of the entire estimator, we will only require the critical value used for convergence to get arbitrarily small as sample size grows.

Now we need to consider the convergence of the stage one estimator, $\tilde{\Gamma}_1^*$, which is computed on the basis of a fixed collection of decision rules. Since the weights attached to the exact investment rules used in forming the approximation are functions of the current parameter guess $\tilde{\Gamma}$, the approximation is as well. As long as the distance function is a continuous function of $\tilde{\Gamma}$, then the weights are as well, which implies that the approximation is continuous in $\tilde{\Gamma}$.

Given that certain events involved in the moment computation are discrete (such as divorce), it is not possible to claim that the functions $\hat{A}(\tilde{\Gamma})$ are continuous. However, continuity is not required for consistency, as is made clear in Pakes and Pollard (1989). We have not explicitly noted dependence of \hat{A} on R , but for now write $\hat{A}_R(\tilde{\psi})$. Then we need uniform convergence of $\hat{A}_R(\tilde{\Gamma})$, so that there exists a value \bar{R} and $\kappa > 0$ such that

$$|\hat{A}_R(\tilde{\Gamma}) - A(\tilde{\Gamma})| < \kappa \tag{23}$$

for all $R \geq \bar{R}$ and $\tilde{\Gamma} \in \Omega_\Gamma$. Standard Law of Large Numbers results yield $\text{plim}_{N \rightarrow \infty} A^N = A$. Then the key elements required for $\text{plim}(\tilde{\Gamma}_2) = \Gamma_0$ are:

1. $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$
2. $\zeta_N \rightarrow 0$ as $N \rightarrow \infty$
3. $R \rightarrow \infty$ as $N \rightarrow \infty$
4. $A(\Gamma)$ continuous function of Γ
5. Uniform convergence of $\hat{A}_R(\Gamma)$.

We do not attempt to characterize the requirements for deriving a well-defined limiting distribution for the estimator $\tilde{\Gamma}_2$. Although computation of the estimator is demanding, it is still feasible to construct bootstrap estimates of its sampling distribution.

Table 1: Ordinary Least Squares Regression of PIAT Math Percentile Scores on Family Characteristics

Independent Variable	Dependent Variable	
	PIAT percentile score	PIAT percentile score change
Constant	43.661 [†] (11.771)	-6.842 (11.055)
Married at 1st test, divorced at 2nd	6.414 (8.309)	-14.494* (8.153)
Divorced at both tests	-4.637 (3.145)	0.335 (2.960)
Total income x 10,000 ⁻¹	0.951 [†] (0.350)	0.520* (0.329)
Mother's income share	12.810** (5.617)	-2.159 (5.289)
Age of child at test	1.483 (2.017)	1.046 (1.894)
Divorce law (1 = unilateral)	-8.118 [†] (2.903)	5.007* (2.732)
State child support guideline amount / y_f	6.950 (7.944)	-9.334 (7.463)

$N = 342$; $R^2 = 0.0698$ and 0.0302 , respectively. Income is measured at the child's date of birth (DOB) and is reported in 1998 dollars. * represents significance at the ten percent, ** at the five percent and [†] at the one percent level.

Table 2a: Estimation Sample Descriptive Statistics

Variable	Mean	Standard Deviation	Minimum	Maximum
1 st PIAT percentile	58.26	27.94	1.00	99.00
2 nd PIAT percentile	61.94	25.22	1.00	99.00
1 st Marital status	0.1733	0.3794	0.0000	1.0000
2 nd Marital status	0.2414	0.4294	0.0000	1.0000
Child's age at 1 st test	5.629	1.077	4.000	13.000
Child's age at 2 nd test	7.476	0.7273	6.000	10.000
Mother's income	22,102.74	18,307.84	0.00	146,869.84
Father's income	38,129.54	49,009.40	0.00	595,689.67
$I(\text{marriage very happy})$	0.7673	0.4236	0.00	1.00
Points of argument	2.564	2.029	0.000	10.000
Mother's AFQT score	72.57	19.13	0.00	102.00
Mother's education	13.86	2.45	4.00	20.00
Mother's age	29.53	3.81	19.00	39.00
Father's education	13.77	2.91	0.00	20.00
Father's age	32.79	5.67	22.00	55.00
$I(\text{Roman Catholic})$	0.3366	0.4737	0.0000	1.0000

$N = 202$ at the first test, 145 at the second. The mother and father's age, education and income are each measured at the first interview following the child's birth. Income is reported in 1998 dollars.

Table 2b: Estimation Sample Descriptive Statistics

Variable	Mean	Standard Deviation	Minimum	Maximum
1 st PIAT percentile	55.94	27.29	1.00	99.00
2 nd PIAT percentile	56.44	25.45	1.00	99.00
1 st Marital status	0.3239	0.4685	0.0000	1.0000
2 nd Marital status	0.3750	0.4848	0.0000	1.0000
Child's age at 1 st test	5.601	0.913	4.000	13.000
Child's age at 2 nd test	7.535	0.736	6.000	11.000
Mother's income	14,857.52	16,134.35	0.00	146,869.84
Father's income	28,310.72	36,612.48	0.00	595,689.67
<i>I</i> (marriage very happy)	-	-	-	-
Points of argument	-	-	-	-
Mother's AFQT score	68.41	19.11	0.00	103.00
Mother's education	13.03	2.38	0.00	20.00
Mother's age	25.80	4.98	16.00	39.00
Father's education	12.87	2.20	0.00	20.00
Father's age	28.86	5.85	19.00	55.00
<i>I</i> (Roman Catholic)	0.3498	0.4774	0.0000	1.0000

$N = 426$ at the first test, 342 at the second. The mother and father's age, education and income are each measured at the first interview following the child's birth. Income is reported in 1998 dollars.

Table 3a: Parameter Estimates

Parameter	Estimate (Standard Error)	Parameter	Estimate (Standard Error)
$\tilde{\delta}$	0.6946 (0.1969)	β_{k_4} on father's age	0.7037 (0.2223)
ν	0.7858 (0.0624)	β_{k_5} on father's education	1.1619 (0.2425)
$\tilde{\sigma}$	0.0981 (0.0322)	β_{θ_0} on constant	0.4830 (0.3114)
$\alpha_1 = \alpha_2$	0.9618 (0.0057)	β_{θ_1} on marriage fairly/not too happy	-1.6007 (0.2047)
β_{k_0} on constant	5.2043 (0.8442)	β_{θ_2} on # argument points	0.9687 (0.3503)
β_{k_1} on AFQT	-0.8118 (0.3645)	β_{θ_3} on Roman Catholic	1.9784 (0.2595)
β_{k_2} on mother's age	1.3946 (0.2523)	σ_k	3.2403 (0.8335)
β_{k_3} on mother's education	1.4630 (0.2893)	$\tilde{\gamma}^+ = \tilde{\gamma}^-$	0.1304 (0.0213)
b	0.5130 (0.1822)	β_r	1.5631 (0.3968)

$N = 202$. All elements of X are standardized to have a zero mean and unit variance in the sample. Parents' incomes are scaled to units of 2500 1998 dollars. σ_θ set equal to 1. Standard errors are based on 50 bootstrapped samples.

Table 3b: Parameter Estimates

Parameter	Estimate (Standard Error)	Parameter	Estimate (Standard Error)
$\tilde{\delta}$	0.2681 (0.0766)	β_{k_4} on father's age	0.0267 (0.1725)
ν	0.6077 (0.0215)	β_{k_5} on father's education	1.5128 (0.0946)
$\tilde{\sigma}$	0.0558 (0.0140)	β_{θ_0} on constant	1.2488 (0.0695)
$\alpha_1 = \alpha_2$	0.7612 (0.0169)	β_{θ_3} on Roman Catholic	1.9676 (0.1390)
β_{k_0} on constant	6.8617 (0.1405)	σ_k	2.8512 (0.2897)
β_{k_1} on AFQT	0.8823 (0.1065)	$\tilde{\gamma}^+ = \tilde{\gamma}^-$	0.3348 (0.0655)
β_{k_2} on mother's age	0.6946 (0.1354)	β_r	0.2715 (0.0799)
β_{k_3} on mother's education	1.5092 (0.1125)		
b	3.7175 (0.1194)		

$N = 426$. All elements of X are standardized to have a zero mean and unit variance in the sample. Parents' incomes are scaled to units of 2500 1998 dollars. σ_θ set equal to 1. β_{θ_1} (on marriage fairly/not too happy) and β_{θ_1} (on # argument points) are both set to 0. Standard errors are based on 50 bootstrapped samples.

Table 4: Data and Simulated Moments

Moment	Data	Simulation	Data	Simulation
	$N = 202$		$N = 426$	
[1] Proportion divorced by test	0.1732	0.1798	0.3231	0.3339
[2] $E(\text{child's age at test} \times I(\text{divorced at test}))$	1.0248	1.0755	1.8443	1.9393
[3] $E(\text{test score} \times I(\text{married at test}))$	49.713	49.859	38.991	39.653
[4] $E(\text{test score} \times I(\text{divorced at test}))$	8.545	8.335	16.896	14.617
[5] $E(\text{test score} \times I(\text{marriage very happy}))$	45.124	44.182	–	–
[6] $E(\text{test score} \times \# \text{ argument points})$	142.67	143.81	–	–
[7] $E(\text{test score} \times \text{mother's education})$	826.01	833.88	741.83	750.77
[8] $E(\text{test score} \times \text{mother's AFQT score})$	4405.19	4314.62	4011.24	4009.46
[9] $E(\text{test score} \times \text{mother's age})$	1745.79	1765.22	1464.85	1479.54
[10] $E(\text{test score} \times \text{father's education})$	829.36	830.44	734.82	732.42
[11] $E(\text{test score} \times \text{father's age})$	1941.16	1965.00	1639.20	1642.72
[12] $E(\text{test score} \times \text{father's income} \times \text{married})$	820.89	846.96	537.91	559.90
[13] $E(\text{test score} \times \text{mother's income} \times \text{married})$	501.33	488.05	303.22	317.17
[14] $E(\bar{y}_1 d=0) - E(\bar{y}_1 d=1)$	2.3638	2.5352	4.4575	1.1344
[15] $E(\bar{y}_2 d=0) - E(\bar{y}_2 d=1)$	0.7042	0.6572	1.4625	0.7040
[16] $E(\bar{d} \text{not Catholic}) - E(\bar{d} \text{Catholic})$	0.0839	0.0879	0.0397	0.0432
[17] $E(\text{test score} \times \text{father's income} \times \text{divorced})$	117.98	119.88	150.25	184.52
[18] $E(\text{test score} \times \text{mother's income} \times \text{divorced})$	84.72	76.56	102.25	107.46
[19] $E(\text{test score}^2)$	4170.52	4149.02	3869.37	3879.90

The simulations are based on $R = 5000$ replications per family (100 per initial conditions pair).

Table 4(Cont'd): Data and Simulated Moments

Moment	Data	Simulation	Data	Simulation
	$N = 202$		$N = 426$	
[20] $E(\text{test score} \times \text{father's income} \times \text{mother's income})$	9870.04	9902.60	5959.66	6661.50
[21] $E(\text{test score} \times \text{child's age} \times \text{divorced})$	51.17	49.49	97.26	84.93
[22] $E((\text{test2} - \text{test1}) \times \text{d1} \times \text{d2})$	1.0966	0.9762	0.4486	1.0092
[23] $E((\text{test2} - \text{test1}) \times (1-\text{d1}) \times (1-\text{d2}))$	3.7103	2.2445	1.0409	0.1158
[24] $E((\text{test2} - \text{test1}) \times (1-\text{d1}) \times \text{d2})$	-0.3862	0.1359	-0.3567	0.0012
[25] $E((\text{test2} - \text{test1}) \times \text{y1})$	105.24	49.54	42.81	8.68
[26] $E((\text{test2} - \text{test1}) \times \text{y2})$	25.8971	26.9538	8.8380	4.1528
[27] $E((\text{test2} - \text{test1}) \times \text{I}(\text{marriage very happy}))$	4.0069	2.5792	-	-
[28] $E((\text{test2} - \text{test1}) \times \text{the argue \#})$	8.6345	10.6858	-	-
[29] $E((\text{test2} - \text{test1}) \times \text{y1} \times \text{d1})$	20.3113	13.0441	8.6813	9.2821
[30] $E((\text{test2} - \text{test1}) \times \text{y2} \times \text{d1})$	5.1013	8.0649	1.5148	4.8948
[31] $E((\text{test2} - \text{test1}) \times \text{y1} \times (1-\text{d1}))$	84.9290	36.4944	34.1320	-0.6018
[32] $E((\text{test2} - \text{test1}) \times \text{y2} \times (1-\text{d1}))$	20.7958	18.8890	7.3232	-0.7421

The simulations are based on $R = 5000$ replications per family (100 per initial conditions pair).

Table 5: States' Social Welfare Weights on the Welfare of Mothers, Fathers, and Children

State	a_1 , weight on father's welfare	a_2 , weight on mother's welfare	a_c , weight on child's welfare
New York	0.482	0.221	0.297
Texas	0.516	0.461	0.023
California	0.508	0.509	-0.017

$N = 24, 48, \& 53$ in NY, TX, and CA, respectively. Simulated outcomes are based on $LMT = 2500$ replications for each family.

Table 6: Welfare Maximizing Policies for Mothers, Fathers, and Children

	Optimal π	Optimal $\tau_1(1)$	$\frac{\text{Optimum-Average}}{\text{Average}} \times 100$
\hat{w}_1	0	0.5	6.762
\hat{w}_2	0.3	0.2	7.425
\hat{w}_c	0	0.5	2.846
d	0.2	0.3	0.785

$N = 426$. Simulated outcomes are based on $LMT = 2500$ replications for each family.

Figure 1.a
Mother's Investment
 $\pi = .25$




Figure 1.b
Mother's Investment
 $\pi = .17$

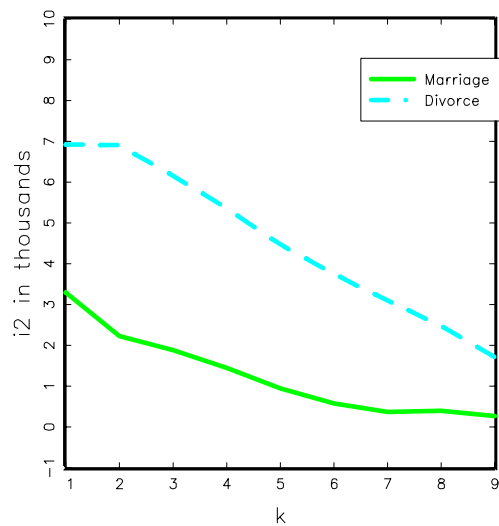


Figure 1.c
Mother's Investment
 $\pi = 0$

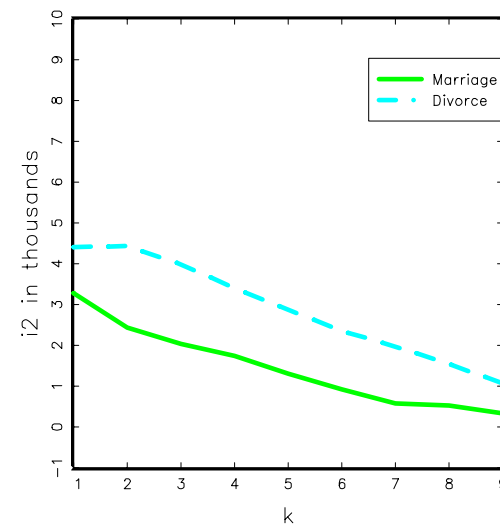


Figure 1.d
Father's Investment
 $\rho = .25$




Figure 1.e
Father's Investment
 $\pi = .17$

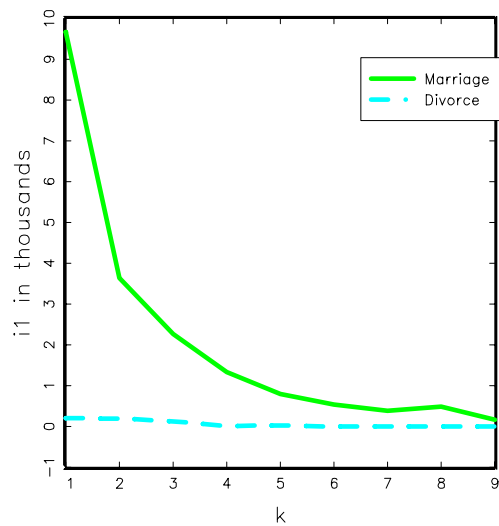


Figure 1.f
Father's Investment
 $\pi = 0$

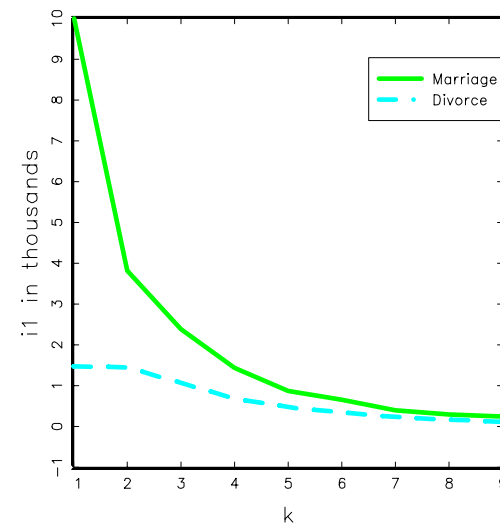


Figure 2.a
Mother's Investment
Married and $k=1$




Figure 2.b
Mother's Investment
Married and $k=4$




Figure 2.c
Mother's Investment
Married and $k=9$




Figure 2.d
Father's Investment
Married and $k=1$




Figure 2.e
Father's Investment
Married and $k=4$




Figure 2.f
Father's Investment
Married and $k=9$

