Children, Time Allocation, and Consumption Insurance

Richard Blundell

University College London and Institute for Fiscal Studies

Luigi Pistaferri

Stanford University

Itay Saporta-Eksten

The Eitan Berglas School of Economics and University College London

We study choices of households deciding on consumption and allocation of spouses' time to work, leisure, and child care. With uncertainty, the allocation of goods and time over the life cycle also serves the purpose of smoothing marginal utility in response to shocks. Combining data on consumption, wages, hours of work, and time spent with children, we compute the sensitivity of consumption and time allocation to transitory and permanent wage shocks. We find that family labor supply responses depend on three counteracting forces: complementarity of leisure time, substitutability of time in the production of child services, and added worker effects.

I. Introduction

In a series of seminal contributions, Becker (1965) and Ghez and Becker (1975) emphasized the importance of understanding the effect of chil-

dren on how households allocate goods and time over the life cycle. To quote from Ghez and Becker, "The parent's utility function is assumed to depend not only on [consumption of goods], but also on commodities measuring child services. . . . The raising of children requires time, especially wife's time, and goods. Thus, time and goods must be allocated between child services and other commodities" (35). One conclusion of that research was that children change the optimal allocation of parents' time to work in the paid labor market, leisure, and child care. Responses to wage and income changes differ across households depending on the importance of "child services" in the parental utility function.

In an environment with uncertainty, the allocation of goods and time over the life cycle also has the goal of smoothing the marginal utility of wealth in the face of shocks to household resources. The key parameters governing the ability to smooth marginal utility are the Frisch and Marshallian (own- and cross-) consumption and labor supply elasticities. For example, a low consumption Frisch elasticity (a low willingness to accept intertemporal fluctuations in consumption) implies a greater desire to smooth consumption relative to, for example, a case in which the elasticity is high. A permanent wage shock faced by the primary earner can be insured through added worker effects (i.e., increased work of secondary earners). This stabilization effect, reflected in the Marshallian elasticity, is stronger the higher the Frisch elasticity of the secondary earner. But since an increase in hours of work reduces the amount of time that can be allocated to the production of child care services, it is clear that there are important trade-offs between producing child care services and insuring consumption against shocks. Combining these two issues (the importance of children in shaping preferences and the demand for insurance) in a unified framework is the goal of this paper. To our knowledge, this is the first paper that attempts to put together these two distinct literatures. Some of the decomposition exercises we discuss below derive precisely from the goal of isolating the distinct forces affecting the choice of how to allocate time in response to different types of events.

To address these goals, we adopt a structural life cycle approach in which a married couple decides how much to consume and how to allocate available time to work activities, leisure, and children (if present). Husband and wife choose time to devote to children as inputs of a home production function for child care. This framework departs significantly from our previous work (Blundell, Pistaferri, and Saporta-Eksten

2016) both because it explicitly takes into account the influence of children on consumption and time allocation decisions of husband and wife and because it is based on a fully structural approach rather than on approximations to the first-order conditions and the lifetime budget constraints. This allows us to fully account for liquidity constraints and extensive margin decisions on time use, which are likely extremely important for younger households with children.

The relation between the hours of husband and hours of wife (as well as the relation between their hours and the spending on goods) can be generated by home production à la Becker (leisure times of the two spouses, as well as market goods, are inputs in the production of "commodities") or by formal nonseparabilities in the utility function (the marginal utility of leisure of one spouse depends on the leisure time of the other; i.e., spouses may enjoy leisure more when they are together). The relationship may also be affected by joint progressive taxation of earnings or correlation of wage shocks (both of which we model explicitly).

In general, we might expect spouses to want to spend time together. Indeed, it is very likely that the complementarity of time together provides a key incentive for relationships to form.[1] We do not directly model family formation, but in order to capture the fact that people have a preference for spending time together or with their children (a public good), we consider the possibility that preferences are nonseparable and complementary over leisure times. In Blundell et al. (2016) we found evidence of Frisch complementarity for time within couples, although we did not relate this to children or measures of time use. However, we found that family labor supply provides insurance against persistent wage shocks and that family leisure times were Marshallian substitutes. When it comes to the care of children, however, there is potentially more room for specialization. For example, it is possible that more efficient technologies require people to separate their time between various activities, especially in the presence of multiple children or children of different ages. In the absence of time use data, it is hard to verify whether covariation in the hours of work of husband and wife descends from explicit nonseparability in utility or from the effect of home production. As Browning, Chiappori, and Weiss (2014) write, "the production function . . . cannot be estimated independently of the utility function unless the home-produced commodities are independently observable" (85).

---

[1] Becker (1976) attributed the existence of marriage to "the desire to raise own children" (211), which requires complementarity of the spouses' leisure time. As he noticed, "Sexual gratification, cleaning, feeding, and other services can be purchased, but not *own* children: both the man and the woman are required to produce their own children and perhaps to raise them" (210), and later, that "complementarity between men and women is the major source of the gain from marriage" (211).

To address these issues we derive structural marginal rate of substitution relations between leisure time of the two partners, the time they allocate to child care, consumption, and prices (hourly wages). We estimate a subset of the structural parameters of the model using data from three data sets: the Panel Study of Income Dynamics (PSID), the American Time Use Survey (ATUS), and the Consumer Expenditure Survey (CEX). We use the latter to impute consumption to all household units in the ATUS (where no consumption information is available). We then use numerical simulation of the model to recover the remaining preference parameters.

The PSID and the ATUS are rich in some aspects but also have important drawbacks. The PSID contains panel data on consumption, assets, hours of work, and hourly wages for both earners, but it does not have consistent information on child care time. During our sample period a special module (the Child Development Survey, CDS) collects time use data for children (including time they spend with their parents). Unfortunately, the CDS is available only occasionally, parental time use has to be inferred from children time use, and the module does not cover all children in the household (in particular, it does not cover young children in later waves). For this reason, we use the PSID to recover parameters when child care time is zero by definition (as in families with no young children). From an identification point of view, households with no young children are key because for them we expect only complementarity of leisure to be important, since no production of child services occurs. Unlike the PSID, the ATUS contains detailed information on time use (including child care) and on hourly wages for the main respondent but lacks information on time use for the partner and has no information on household consumption. We then impute this information using cohort-education-year average values of parental time and hourly wages (from the ATUS male respondents) and household consumption (from the CEX). Time use data bring very useful additional information that allows us to separately identify key aspects of the production of child care services.

Our structural estimates can be used to provide a comprehensive picture of the ability of households to smooth marginal utility in response to shocks. In addition, information on hours of work and hours spent on child care allows us to decompose overall Frisch cross-responses into two components: one reflecting the potential degree of complementarity between husband's and wife's leisure (demand for "companionship" or "love") and another reflecting the degree of substitutability of their child care time in the production of child care services.

Using simulations from our estimated model, we find that this decomposition is important. While overall cross-Frisch responses of hours are small, the responses of the two subutility margins are of opposite signs,

highlighting the tension between substitutability and complementarity in driving the hours response. In particular, the hours of work of young mothers appear to respond little to a temporary increase in the male's wage (which induces an increase in his hours) because the force that pushes her to reduce her leisure time (or work longer) due to the lower leisure time of her companion is counteracted by the force that pushes her to increase child care time because the husband is now allocating fewer hours to child raising and her time is a substitute for his time in the household production of child care services. Similarly, own (especially female) responses are large because they include both an intertemporal substitution component and a home production component: when wages are temporarily higher, the opportunity cost of an extra hour of leisure and an extra hour of child care increases. Ignoring the interaction between children and time use when modeling family labor supply thus misses important components of the adjustments in consumption and time use that families undertake in response to shocks to their resources.

The model can be used to simulate behavior in counterfactual scenarios. We compare two revenue-neutral policies that have direct effects on the welfare of households with children. In one case we provide households with young children with an unconditional subsidy. In another, we compensate households with young children for the child-related fixed costs of work they face. We find that both policies relieve households, in the prechildren period, from saving in anticipation of the decline in family earnings induced by the wife reallocating time from market to child care when children arrive. We also find that the unconditional subsidy policy has greater welfare value than the policy that compensates for the fixed costs of work. This is partly due to the "no-strings-attached" nature of the first policy and partly due to the dynamics induced by credit market frictions.

This paper is related to a large literature that studies family labor supply. Specifically, there are two strands of research: papers that consider the impact of children on female labor supply and those that examine the role of labor supply as an insurance against shocks—in particular, wage shocks. In the first line of work there are papers that establish a clear relationship between the presence of children and female labor supply, such as Angrist and Evans (1998). Attanasio, Low, and Sanchez-Marcos (2008) study the role of decreasing child care prices in explaining changes in female participation rates across cohorts. Other papers use time diary data to examine the precise decomposition of available time between children and other activities (Aguiar and Hurst 2007; Ramey and Ramey 2010; Aguiar, Hurst, and Karabarbounis 2013).[2] In the second

---

[2] Guryan, Hurst, and Kearney (2008) and Ramey and Ramey (2010) also study heterogeneity in child care time by parents' skills.

line of work, Hyslop (2001), Attanasio, Low, and Sanchez-Marcos (2005), and Blundell et al. (2016) are some representative contributions.[3]

Our paper makes three key contributions. The first contribution is methodological. We incorporate consumption and time use activities of husband and wife in a dynamic life cycle model with uncertainty about their underlying wages, liquidity constraints, and heterogeneity of preferences. We then show how to derive marginal rate of substitution relations between leisure of the two partners and consumption. Controlling for the latter is an implicit way to control for the unobservable marginal utility of wealth (as done by Altonji [1986] in a single-earner labor supply context). Similarly, we can derive a structural marginal rate of substitution relation linking child care time of the two partners and consumption. With appropriate selection corrections and instruments, we show how to identify the parameters characterizing utility and the child care production function. Model simulations can then be used to recover the remaining parameters (mostly preference shifters).

Our second contribution is in the context of the family labor supply literature. We demonstrate, both theoretically and empirically, how the presence of young children in the household gives rise to heterogeneity in the (own- and cross-) elasticities of labor supply to wage shocks. Using our data, we provide new evidence on the decomposition of labor supply elasticities to their leisure and child care time sources, especially highlighting the tension between spouses' leisure complementarity and child care time substitutability. Our third contribution is to the consumption smoothing literature. Using simulations from our model, we show that the heterogeneity in labor supply elasticities that is driven by the presence of young kids translates into heterogeneity in the use of market hours to smooth wage shocks.

The paper proceeds as follows. Section II describes the life cycle problem of the household. Section III discusses the derivation of the moment conditions we use in estimation, as well as the decomposition of labor supply elasticities. Section IV describes the PSID, ATUS, and CEX data, while Section V discusses estimation and parameterization. The results, the discussion about the relation between child care and wages, as well as the implications of the results for the degree of insurance are in Section VI. In this section we also discuss the role of credit constraints in shaping responses. In Section VII we discuss the results of the two policy experiments described above (a subsidy to child care cost and a compensation for child-related fixed costs of work). Section VIII presents conclusions.

---

[3] The "added worker effect" literature studies in particular how the wife's employment responds to the husband's unemployment events. See, e.g., Lundberg (1985), Stephens (2002), and Juhn and Potter (2007).

II. A Life Cycle Model with Child Care Production Function

In this section we discuss a life cycle problem faced by a couple. We consider a unitary framework in which the household draws utility from the leisure times of the spouses, from the consumption of non-child-care-related goods, and from "child commodities."[4] The latter are produced using as inputs the hours that husband and wife devote to their children (allowing for observed and unobserved heterogeneity).[5] For simplicity, we assume that fertility decisions are exogenous. Thus, the household's problem is to choose consumption and the allocation of the two members' total hours to three activities: work, child care time, and leisure.[6]

We assume throughout that the hourly wage process is exogenous; we also allow wage shocks to be potentially correlated across spouses. As we highlight, the model shows that the response of family labor supply to shocks is rather complex,[7] as it depends on the spouses' elasticities for (responsiveness of) leisure and child care time to wage changes, the degree of complementarity of their leisure times in utility, the degree of substitutability of child care time inputs in the production of child care, and the degree of progressivity of the tax system (which makes the secondary earner's marginal tax rate depend on the primary earner choice of hours even in the case in which leisure times are additively separable).

A. Wage Process

Similarly to Blundell et al. (2016), we assume that each earner's wage process contains both a permanent and a transitory component. The permanent component evolves as a random walk process, while the transitory component is serially uncorrelated. The distinction between transitory shocks and permanent shocks is important, as in a frictionless world with interior solutions one can interpret transitory shocks as having negligible or no wealth effects on labor choices. Hence in this special case, the response of hours to transitory wage shocks identifies Frisch (or λ-constant)

---

[4] See Chiappori and Mazzocco (2017) for a recent survey of collective models of family labor supply.

[5] Such heterogeneity may capture the influence of other inputs, such as expenditure on toys, books, etc., that we do not model explicitly.

[6] We consider only three possible uses of time (child care, leisure, and work) for the two earners in the household. This is of course a simplifying assumption but is consistent with our goal to focus on and highlight in this paper the interaction between children and labor supply. In practice, what we call "leisure" may subsume multiple adjustment margins in response to changes in resources (housework being the most obvious example), besides "pure leisure."

[7] As Heckman (1988) notes, in a model with multiple time uses of nonmarket time, restrictions on the wife's time in one activity need not imply wage- and income-inelastic labor supply.

elasticities, while the response to permanent wage shocks identifies Marshallian (or uncompensated) elasticities. However, since labor supply behavior of individuals (especially women) is characterized by corner solutions and since households may face frictions in credit markets, it is important to focus more generally on labor supply elasticities in response to wage changes of a different nature.

We assume that the log of real wage of earner $j = \{1, 2\}$ at age t can be written as

$$\log W_{j,t} = x'_{j,t}\beta^j_W + F_{j,t} + u_{j,t}, \tag{1}$$

$$F_{j,t} = F_{j,t-1} + v_{j,t}, \tag{2}$$

where $x_{j,t}$ are observed characteristics affecting wages and are known to the household; $u_{j,t}$ and $v_{j,t}$ are transitory and permanent shocks, respectively. We assume that transitory and permanent innovations are uncorrelated within persons. However, they may be correlated across spouses.[8] The structure of markets is such that shocks are not formally insurable, households have no advance information about them, and they are observed (separately) at time t.[9] We make the simplifying assumption that the variances and covariances of the shocks are constant over the life cycle.[10]

Given the specification of the wage process (1)–(2), the growth in (residual) log wages can be written as

$$\Delta w_{j,t} = \Delta u_{j,t} + v_{j,t}, \tag{3}$$

where Δ is a first difference operator and $\Delta w_{j,t} = \Delta \ln W_{j,t} - \Delta x'_{j,t}\beta^j_W$ (the log change in hourly wages net of observables).

---

[8] This is potentially important since several empirical papers find evidence of a correlation between the labor market outcomes of married couples. See, e.g., Hyslop (2001) and Juhn and Potter (2007).

[9] This is a key assumption in the context of empirical analysis on consumption insurance. See Meghir and Pistaferri (2011) for a discussion about the interpretation of insurance coefficients when this assumption is violated.

[10] We assume that both husband's and wife's offered wages change over the life cycle to capture, in a reduced-form fashion, the effect of human capital accumulation. We do not allow, however, for wages to increase with labor market experience (i.e., cumulated periods of work) or decline with periods out of work, which would induce a feedback between labor supply decisions and future wages. See Attanasio et al. (2008) for examples with female labor supply. The assumption of exogenous wages may potentially affect our estimates of added worker effects, since in a world with endogenous experience capital secondary earners may have greater incentives to respond, since "waiting" to respond to shocks reduces future wages and hence the stabilizing power of the added worker's earnings.

B. Household Decisions

Households solve the following problem over the life cycle:

$$\max \mathbb{E}_t \sum_{s=0}^{T-t} u_{t+s}(C_{t+s}, L_{1,t+s}, L_{2,t+s}, T_{1,t+s}, T_{2,t+s}; z_{t+s}; \varepsilon_{t+s})$$

$$\text{s.t. } A_{t+1} = (1+r)[A_t + \mathcal{T}(z_t, H_{1,t}W_{1,t} + H_{2,t}W_{i,2,t}) - C_t],$$

$$L_{1,t} + H_{1,t} + T_{1,t} = \bar{L}, \tag{4}$$

$$L_{2,t} + H_{2,t} + T_{2,t} = \bar{L},$$

$$A_{t+1} \geq 0, H_{1,t} \geq 0, H_{2,t} \geq 0 \quad \forall\ t,$$

where C represents consumption; L_j, T_j, and H_j are the leisure time, the child care time, and the hours worked by earner j ($j = 1, 2$), respectively; $\bar{L}$ is the maximum time available; and A are assets that pay a nonstochastic interest rate denoted by r (note that $A_{t+1} \geq 0$ introduces an explicit, exogenous liquidity constraint). Finally, the notation $u_{t+s}(\cdot)$ assumes discounting of utility at future dates, that is, $u_{t+s}(\cdot) = \beta^{t+s} u(\cdot)$.

We introduce two sets of demographic conditioning variables in the utility function. The first, z, contains the number of kids and their age composition. We model child care time, leisure, and labor supply decisions as explicit functions of z. Since these family composition variables affect consumption in ways not captured by our model (i.e., not only through child care), we let the effect of these variables on consumption work through an (estimated) adult equivalence scale (see below for actual functional forms). The second set of conditioning variables, ε, includes unobservable taste shifters.

Households face a progressive tax system that mirrors the US schedule. The expression $\mathcal{T}(\cdot)$ is a tax function that maps before-tax into after-tax household earnings. We extend Blundell et al. (2016) by allowing for unearned taxable income and model joint taxation as

$$\mathcal{T}(z_t, H_{1,t}W_{1,t} + H_{2,t}W_{2,t}) \approx \chi_t (b(z_t) + H_{1,t}W_{1,t} + H_{2,t}W_{2,t})^{1-\mu_t}. \tag{5}$$

The way to interpret $b(z)$ is that it represents a consumption floor that is available to households even in cases in which household earnings fall to zero (e.g., means-tested benefits such as the Supplemental Nutrition Assistance Program [SNAP] or Temporary Assistance for Needy Families [TANF] in the United States). In the policy experiment section, we evaluate the welfare consequences of subsidizing child care by changing the value of $b(z)$, that is, targeting transfers to households with young children (similar to an increase in the child tax credit).

III.   The Dynamics of Time Use, Hours, and Goods

A.   Restrictions on Preference and Moment Conditions

For our empirical application, we choose a specific functional form for
the period utility in (4):

$$
u(\cdot) = \exp\big(\tilde{\phi}_C(z,\varepsilon)\big)\frac{\tilde{C}^{1-1/\eta}}{1-1/\eta}
$$

$$
-\frac{1}{1-\rho_L}\left[\exp\big(\tilde{\phi}_{L_1}(z,\varepsilon)\big)L_1^{1-1/\varphi_{L_1}} + \exp\big(\tilde{\phi}_{L_2}(z,\varepsilon)\big)L_2^{1-1/\varphi_{L_2}}\right]^{1-\rho_L} \quad (6)
$$

$$
-\frac{1}{1-\rho_T}\left[\exp\big(\tilde{\phi}_{T_1}(z,\varepsilon)\big)T_1^{1-1/\varphi_{T_1}} + \exp\big(\tilde{\phi}_{T_2}(z,\varepsilon)\big)T_2^{1-1/\varphi_{T_2}}\right]^{1-\rho_T},
$$

where $\tilde{C} = C - \gamma(z)E_2$. Here $E_2 = \mathbf{1}\{H_2 > 0\}$ is a dummy for female em-
ployment.[11] This expression allows us to capture a simple form of non-
separability between consumption and women's employment; see Heck-
man (1974). One way to interpret the parameter $\gamma(z)$ is as a form of
utility fixed cost of work (although this is equivalent to a monetary fixed
cost of work). In the policy experiment section we discuss the welfare ef-
fects of reducing the fixed cost of work for individuals with young chil-
dren by changing the value of $\gamma(z)$. We assume that the taste shifters de-
pend on observable and unobserved heterogeneity, $\tilde{\phi}_x = f_x(z,\varepsilon)$. Finally,
the following restrictions on the parameters yield a well-behaved utility
function: $0 < \varphi_x < 1$ (for $x = \{L_1, L_2, T_1, T_2\}$), $\eta > 0$, and $\rho_x < 1$ (for $x = \{L, T\}$).

The last term in (6) can be interpreted as a home production function
for "child commodities" (or child services), which uses parental time as
inputs. Note also that the sign of ρ_L is informative about whether the lei-
sure times of the partners are complements ($\rho_L > 0$) or substitutes ($\rho_L < 0$)
in utility. A similar interpretation applies, symmetrically, to ρ_T. In particu-
lar, if $\rho_T < 0$ ($\rho_T > 0$), the times that partners devote to the production
of child care services are substitute (complement) inputs.

While this formulation imposes some restrictions on the interactions
between goods in the utility function, it allows for various forms of non-
separability that are crucial for our analysis. First, it allows for nonsepara-
bility between leisure times as well as between child care time of the two
earners in the family. Second, it allows for nonseparability of female labor
supply and consumption through the introduction of a fixed cost of work

---

[11] Note that the consumption equivalence scale is captured by $\tilde{\phi}_C(z,\varepsilon)$.

(which is typically observed in formal estimation; see Cogan [1981] for an early contribution).[12]

We assume a log-linear relationship for the preference shifters, allowing them to change with the age of children z (in practice, we use a dummy for whether the household has a young child—aged 10 or less—which we assume is the age range where production of child care services occurs) and a preference shock $\varepsilon_{x,t}$, which is independently and identically distributed with mean zero and variance $\mathrm{var}(\varepsilon_{x,t}) = \sigma^2_{\varepsilon_x}$:

$$\tilde{\phi}_x(z_t, \varepsilon_{x,t}) = \phi_x^{nk} + \phi_x^{k} z_t + \varepsilon_{x,t}.$$

The first-order conditions of the household's problem lead to a series of marginal rate of substitution (MRS) relations in which the ratio of marginal utilities with respect to two "goods" is equal to the ratio of their prices (at least when solutions are interior). Moreover, there is an Euler equation linking the evolution of the marginal utility of wealth over time. In the appendix (available online) we show that, given our functional form assumptions, a manipulation of the MRS equations leads to the following moment conditions (assuming interior solutions for both partners):

$$E\left(l_{2,t} - K_0 - \varphi_{L_2}(w_{1,t} - w_{2,t}) - \frac{\varphi_{L_2}}{\varphi_{L_1}} l_{1,t} \,\Big|\, I_t \right) = 0, \qquad (7)$$

$$E\Bigg(l_{2,t} - K_1 + \varphi_{L_2} w_{2,t} - \mu\varphi_{L_2} y - \frac{\varphi_{L_2}}{\eta} c_t - \frac{\varphi_{L_2}}{\varphi_{L_1}} \rho_L (1 - \varphi_{L_1}) l_{1,t}$$
$$+ \varphi_{L_2} \rho_L \frac{\varphi_{L_2}(1 - \varphi_{L_1})}{\varphi_{L_1}(1 - \varphi_{L_2})} \frac{W_{2,t} L_{2,t}}{W_{1,t} L_{1,t}} \,\Big|\, I_t \Bigg) = 0, \qquad (8)$$

where lowercase letters indicate logs of uppercase variables (i.e., $l_{2,t} = \log L_{2,t}$) and I_t is a set of exogenous variables known at time t (instruments).[13] Moment condition (7) is the MRS between the wife's and the husband's leisure. It simply shows that the leisure time of the husband may differ from the leisure time of the wife both because of heterogeneity in relative preferences for leisure and because of differences in offered wages. Moment condition (8) is the MRS between the wife's leisure and

---

[12] In principle, the model could be identified even with other forms of nonseparabilities (e.g., nonseparability between consumption and child care time and between leisure and child care time). However, the fact that we have to impute consumption and we have time use for only one respondent in the ATUS sample limits the extent to which we can precisely recover more flexible forms of nonseparability in the data.

[13] Note that y is the log of taxable family income. It appears in the log-linearized MRS because of progressive taxation ($\mu > 0$ in our case).

consumption. The last two terms in (8) originate from the assumption of nonseparability between wife and husband leisure (and hence they drop out if $\rho_L = 0$). Since the error terms include preference heterogeneity shocks (and measurement errors), the set I_t cannot include information on leisure, consumption, or wages. We discuss the choice of instruments we use in the empirical section. The parameters in the moment conditions (7) and (8) are estimated using the PSID. Since the PSID does not contain information on child care time, we focus on a sample in which this is zero by definition (households with no young children—aged 10 or less), and hence leisure time is just total time minus hours of work. Since wages are observed only for workers, we estimate these equations using only working couples and correct for sample selection using standard Heckman correction terms (Heckman 1979). From an identification point of view, the focus on households with older children in estimating (7) and (8) is useful because for those households we expect only complementarity of leisure to be an important determinant of (compensated) cross-responses, since by assumption no production of child services occurs.

How to recover the parameters of the parental time subutility? We use the following moment conditions:

$$E\left(t_{2,t} - K_2 - \varphi_{T_2}(w_{1,t} - w_{2,t}) - \frac{\varphi_{T_2}}{\varphi_{T_1}} t_{1,t} \,\Big|\, I_t \right) = 0, \qquad (9)$$

$$E\left(t_{2,t} - K_3 + \varphi_{T_2} w_{2,t} - \mu\varphi_{T_2}y - \frac{\varphi_{T_2}}{\eta} c_t - \frac{\varphi_{T_2}}{\varphi_{T_1}} \rho_T (1 - \varphi_{T_1}) t_{1,t} \right.$$
$$\left. + \varphi_{T_2}\rho_T \frac{\varphi_{T_2}(1 - \varphi_{T_1})}{\varphi_{T_1}(1 - \varphi_{T_2})} \frac{W_{2,t} T_{2,t}}{W_{1,t} T_{1,t}} \,\Big|\, I_t \right) = 0, \qquad (10)$$

which are the symmetric equivalents of (7) and (8) for parental time. In particular, (7) is the MRS between child care time of the two spouses, while (8) is the MRS between the child care time of the wife and household consumption. An increase in the relative cost of her time would induce a decline in her child care time relative to his child care time. Since the PSID lacks information on parental time, we estimate these two moment conditions using the ATUS. However, we face two missing data problems. First, ATUS collects detailed time use information only for one respondent in the household. Hence, if we focus on female respondents and observe $t_{2,t}$, we have no information on the husband's parental time, $t_{1,t}$. We hence impute the latter from the sample of ATUS male respondents (conditioning on the wife's demographic characteristics).

The second missing data problem is that, unlike the PSID, there is no consumption information in the ATUS. We thus impute consumption using CEX data.[14] Similarly to the case discussed above, the structural estimation of (9) and (10) requires instrumental variables and a selection correction for working couples.[15] We discuss the choice of instruments in the empirical section.

Note that (8) and (10) use ideas developed by Altonji (1986) in a single-earner labor supply context. In general, in a life cycle context one obtains Frisch demand functions for consumption, leisure, time use, and so forth, where the choice variables are a function of observable prices and the unobservable marginal utility of wealth. First-differencing the data (i.e., using the Euler equation) eliminates the unobservable marginal utility of wealth and allows unbiased estimation of intertemporal substitution parameters. A different estimation strategy is to proxy the unobservable marginal utility of wealth with consumption. In this case, one is effectively differencing across economic decisions (given time) instead of differencing across time (given economic choices). Of course, one still needs to instrument the endogenous consumption level, but this is a more conventional problem than modeling directly the structure of the marginal utility of wealth (which requires assumptions about the information set etc.). This strategy is particularly useful in contexts, like ours in the ATUS case, in which there are no panel data allowing us to follow the dynamics of economic decisions across periods.

Direct estimation of (7)–(10) gives consistent estimates of a subset of the parameters of interest: η, φ_{L_1}, φ_{L_2}, ρ_L, φ_{T_1}, φ_{T_2}, and ρ_T. However, the distribution of the preference shift parameters $\tilde{\phi}_x(z, \varepsilon)$, as well as the fixed cost (γ) parameters remains unidentified. To complete estimation, we solve the life cycle problem given the parameters estimated in the first step (i.e., φ_{L_1}, φ_{L_2}, ρ_L, φ_{T_1}, φ_{T_2}, and ρ_T) and identify the rest of the parameters by matching data and simulated moments. The procedure targets the moments for the two earners' distribution of hours of work, time spent with children, employment, and the change in consumption when children are born.

---

[14] We impute using the CEX instead of the PSID for several reasons. First, the CEX sample is larger; second, it is conducted at the same annual frequency of ATUS (while the PSID is biannual); finally, the consumption measure in the CEX is more comprehensive than the PSID measure. Separately for each year, we impute to each ATUS household the average level of CEX consumption observed for households with the same wife's birth cohort and level of education.

[15] The estimation of (10) also requires data on total household taxable income. While in the PSID this measure is readily available, in the ATUS only a proxy can be constructed. Given that the PSID estimates are very close including or omitting the tax-related term, we neglect it for the part of the analysis conducted on the ATUS.

B.   Labor Supply Elasticities

Labor supply elasticities are an obvious object of interest for policy analysis and the focus of a vast literature (Keane 2011). In our context, they are complicated functions of the preference parameters described above. In this section we discuss how they can be obtained from model-based responses.

We first notice that elasticities for hours of work, leisure, and parental time are linked through the time budget constraint. In particular,

$$\eta_{H_i,W_j}\big|_a \equiv -\eta_{L_i,W_j}\big|_a \frac{L_i}{H_i} - \eta_{T_i,W_j}\big|_a \frac{T_i}{H_i} \tag{11}$$

for $i, j = \{1, 2\}$, with

$$\eta_{X_i,W_j}\big|_a = \frac{\partial X_i}{\partial W_j}\frac{W_j}{X_i}\bigg|_a \quad (X = \{H, L, T\}).$$

Here the notation $\big|_a$ refers to the margin that is kept fixed or the type of wage change considered. For example, one could keep constant marginal utility (to get Frisch elasticities) or nothing at all (Marshall or uncompensated elasticities); or one could consider transitory or permanent wage changes. In practice, the mapping between our preference parameters and these various elasticities can be obtained by simulating the response of hours to wage changes of a different nature.

Two special (and policy-relevant) types of elasticities (Frisch and Marshall) obtain when the spouses are choosing interior hours solution and they face no progressive taxes and no liquidity constraints; moreover, in this setting one can approximate the Frisch elasticity as the response to transitory wage changes and the Marshall elasticity as the response to permanent wage changes (see Blundell et al. 2016). In the analysis below we use our model to simulate the response of time uses to transitory and permanent shocks; this also offers an alternative to the approximation approach used by Blundell et al. (2016).[16]

The intuition underlying (11) is simple. Consider the effect of a temporary wage increase faced by the wife ($j = 2$) in a setting with interior hours solution and no credit market frictions. Since the wage change is temporary and the consumers are unconstrained, wealth effects are (approximately) zero and there are only intertemporal substitution effects. The effect of a temporary positive shock on the wife's labor supply ($\eta_{H_2,W_2}\big|_{\Delta W_2 = \text{trans.}} \cong \eta_{H_2,W_2}\big|_{\partial\lambda=0}$, with λ being the marginal utility of wealth) can be decomposed into two distinct forces. First, a positive temporary deviation from the normal wage induces workers to reallocate their work time

---

[16] Blundell, Pistaferri, and Saporta-Eksten (2017) compare the two approaches more formally.

intertemporally to periods in which the wage is temporarily higher and their leisure to periods with lower wages. This effect is captured by the first term, and its strength is measured by the leisure elasticity $\eta_{L_2,W_2}\big|_{\Delta W_2=\text{trans.}} < 0$ (leisure L is a "good" that costs W). When the wage is temporarily high, moreover, it becomes more costly to devote time to children, and this induces a further increase in hours (or, more correctly, a decline in child care hours that is partly reallocated to work), an effect whose strength is measured by $\eta_{T_2,W_2}\big|_{\Delta W_2=\text{trans.}} < 0$. The two elasticities are weighted by L_2/H_2 and T_2/H_2, respectively, because their strength depends on "time allocation shares" (i.e., the second effect is absent for earners who devote no time to children to start with or for households with no young children).

A somewhat more interesting effect is how the wife responds to temporary changes in the husband's wage ($\eta_{H_2,W_1}\big|_{\Delta W_1=\text{trans.}}$), that is, the cross-responses. A temporary increase in the husband's wage induces him to work longer hours (for reasons explained above). If husband and wife enjoy leisure together, $\eta_{L_2,W_1}\big|_{\Delta W_1=\text{trans.}} < 0$, and, in the absence of child care time ($T_2 = 0$), we should see an increase in the wife's hours as well. When husband and wife hours can be used to produce child care, however, the effect is less obvious. If their child care hours are substitutes in production, an increase in hours worked by the husband implies he can now allocate fewer hours to child care production, and hence the wife will allocate more hours to child care production (and fewer hours to market time, i.e., work less). This implies $\eta_{T_2,W_1}\big|_{\Delta W_1=\text{trans.}} > 0$. The magnitude of the effect will depend on the degree of input substitutability in child care production. It follows that the two forces may well counteract each other.

C.   Added Worker Effects

Besides Frisch elasticities, a complete characterization of labor supply behavior requires understanding how Marshallian elasticities vary across households with or without children. Without taxes, the Frisch responses coincide with the goods responses to transitory price shocks. Permanent shocks, however, change the marginal utility of wealth as well. The overall response to permanent shocks to prices is captured by the Marshallian elasticities. Generally, Marshallian responses to (permanent) wage shocks are more involved than Frisch responses because of wealth effects acting alongside substitution effects.[17] In the case of cross-responses, for instance, the effect of a permanent decline in the husband's wage on the wife's labor supply in our model mixes three effects: complementarity of leisure time, substitutability of time inputs in the production of child care ser-

---

[17] Blundell et al. (2016) provide an overview of the forces present in the family labor supply model.

vices, and added worker effects, that is, the fact that women have an incentive to replace some of the (permanently) lost earnings of the husband by working more.[18] Analytically,

$$\eta_{H_i,W_j} = \underbrace{- \eta_{L_i,W_j}|_{\partial\lambda=0} \frac{L_i}{H_i} - \eta_{T_i,W_j}|_{\partial\lambda=0} \frac{T_i}{H_i}}_{\eta_{H_i,W_j}|_{\partial\lambda=0}} + \underbrace{\frac{\partial H_i}{\partial\lambda}\left(\frac{\partial\lambda}{\partial W_j}\frac{W_j}{\lambda}\right)}_{\text{added work effect}}.$$

The uncompensated elasticity (η_{H_i,W_j}) can be decomposed into the three terms described above. The last term, the added work effect, leads to a change in wealth, which in turn induces a change in hours (and of child care time and leisure time correspondingly).

In the empirical application we demonstrate the different roles of our parameters both in shaping the Marshallian labor supply responses and in shaping the Marshallian consumption response, which is closely related to the degree of consumption insurance.

IV.   Data and Estimation Methodology

For our empirical analysis we use three data sets: the Panel Study of Income Dynamics (PSID), the American Time Use Survey (ATUS), and the Consumer Expenditure Survey (CEX). We describe their salient features in turn.

A.   The PSID Interview Data

The PSID started in 1968 collecting information on a sample of about 5,000 households. The initial sample included a subsample that was representative of the US population (roughly 3,000 households, the "core sample") and a subsample of low-income households (the Census Bureau's "Survey of Economic Opportunity sample," about 2,000 households). From then on, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. The PSID interview data were collected annually until 1996 and biennially starting in 1997. Since 1999, in addition to income data and demographics, the PSID collects data about detailed asset holdings and consumption expenditures. Since we need both consumption and assets data (the latter to construct instruments), we focus on interview data from the 1999–2015

---

[18] Adoption of a collective framework of behavior would introduce a further reason for observing a relationship between husband and wife leisure. For example, under limited commitment, negative wage shocks (especially permanent ones) faced by the husband reduce his Pareto weight, implying that the partner gains in terms of consumption of all goods, including the consumption of leisure. This effect runs opposite to the added worker effect that we find dominates empirically. We leave the analysis of this further channel to future work.

sample period. These data therefore provide us with longitudinal information on all the variables required for estimation of (7) and (8) with the exception of the time that each earner spends on child care. We focus on households with married couples in which the wife is aged 25–65. For the estimation of (7) and (8) we further restrict our sample to couples with no young children (defined as age 10 or less). In this sample parental time is zero by assumption, and hence leisure time can be obtained as the difference between total time available and hours of work. We further focus on households with nonmissing data on consumption and education of both earners. Finally, to reduce the impact of measurement error in the estimation of (7) and (8), we drop observations in which one of the earners is working more than 60 hours a week for the entire year, implying very low leisure values.

The construction of the consumption data is similar to that in Blundell et al. (2016). We use only consumption categories that were consistently collected starting in 1999. These include food (at home and away), health expenditures, utilities, gasoline, car maintenance, transportation, education, child care, home insurance, and rent (or rent equivalent for homeowners).[19] We treat subcategories with missing values as zeros. Finally, we winsorize log of hourly wages, leisure, consumption, nonlabor income, and the ratio in (8) at the top and bottom 1 percent to avoid outliers driving the results.

B.   The ATUS Data

The American Time Use Survey (ATUS) collects information on time use for a representative sample of the US population. The sample is drawn from those who have completed their eighth and final month of interviews for the Current Population Survey (CPS). One respondent per household (either the husband or wife in couple households) reports detailed information on how he or she spent his or her time on the previous day. The matching with the CPS can be used to recover demographic information for the respondent and the spouse (including his or her hourly wage), although some information is updated at the time of the ATUS interview. In its modern format, the ATUS is available for the 2003–15 period (older waves are also available but differ substantially in format and content from the more recent waves; we do not use them). The broad type of time use activities covered by the ATUS includes personal care, eating and drinking, household activities, purchasing goods and services, caring for and helping household members, caring for and helping non–

---

[19] In the 2015 wave, because of an error in the PSID data collection procedure, the data on car maintenance are missing whenever a car insurance value is reported. We impute these values using a within-household projection of the car maintenance budget share on all other transportation-related budget shares as well as the food budget share.

household members, working and work-related activities, educational activities, organizational, civic, and religious activities, leisure and sports, telephone calls, mail, and email. Our definition of child care time includes caring for and helping household children (activity codes 030101–030199), activities related to household children's education (activity codes 030201–030299), activities related to household children's health (activity codes 030301–030399), and travel related to caring for children (180381). Separately for each year, we impute child care time of the husbands of wife respondents by computing the average child care time of all husband respondents (averaged over their wives' birth cohort and education, which are always observed).

Our sample selection focuses on households in which the youngest child is 10 or younger. Given that we are interested in the substitution between time spent with the children and hours of work, we use only the weekday diaries (Monday to Friday) and drop national holidays.

We use ATUS data to estimate equations (9) and (10). Unfortunately, the ATUS does not contain any information about consumption. We hence impute consumption using data from the CEX, described next. Separately for each year, we impute to each ATUS household the average level of consumption observed in the CEX for households with the same wife's birth cohort and level of education.

C.   The CEX Data

The Consumer Expenditure Survey (CEX) is the only US data set that has comprehensive and detailed information on household expenditure and its various components. Available on a continuous basis since 1980, it is used by the Bureau of Labor Statistics to form weights that go in the computation of the Consumer Price Index (CPI) (and for other minor matters as well). The CEX is composed of two distinct surveys: the interview survey (where spending information is collected by 3-month recall) and the diary survey (where spending is collected by filling in a 2-week diary). Respondents in the interview survey are sampled every 3 months (for a total of five times, although data for the first interview are not released because they are merely preparatory), while those in the diary survey are sampled only once. The two surveys cover different consumption items, with some overlap. We focus on the interview sample. We construct a measure of nondurable consumption defined as the sum of spending on food at home, food away from home, alcohol and tobacco, utilities, maintenance and repairs, education, housing services, financial services, clothing, health, entertainment, cash contributions, transportation, and other nondurables. We then deflate using the CPI. Our sample excludes rural households, those with incomplete income responses, and those with zero consumption.

D.   Descriptive Statistics

Descriptive statistics on time use from the ATUS (panel A of table 1) show that among working married couples with a child aged 10 or less, about 31 percent of husbands and 8 percent of wives report zero hours spent with their kids.[20] Partly, this difference is due to the diary nature of the data. Suppose that on the day time use data were collected the husband left early in the morning and returned after the child went to sleep. It would be clearly misleading to assume that the husband spent zero annual child care time based on observing zero child care hours on the diary day. Similarly, it is possible that some of the zeros may come from the respondent being out of town on the previous day. For these reasons, estimation of moment conditions (9) and (10) uses actual wife's child care time. Given the way time use data are collected in the ATUS, when child care times are observed for the wife, they are missing for the husband. Hence, we impute them as described above.[21]

Panel A of table 1 also shows the distribution of annual child care hours for husbands and wives (with and without zeros). Among husbands, median hours are 331 and the mean is very similar. Among wives who devote some time to their kids, median hours are 542 and the mean is 656. Consistent with estimates from Guryan et al. (2008) and Ramey and Ramey (2010), we thus find that wife's child care annual hours are almost twice as much as the husband's hours.

Descriptive statistics on PSID variables (consumption, leisure, hours, and hourly wages) for the sample that is used in the estimation of (7) and (8) are reported in panel B of table 1. The survey collects data on annual labor earnings and on annual hours of work. To construct the hourly wage measure we divide annual earnings by annual hours of work. Hence, we have a measure of the average hourly wage. As expected, men work longer hours and have higher hourly wages than women.

V.   Identification and Estimation

There are three sets of parameters in the model: wage parameters, preference parameters, and tax parameters. Some of these parameters are set outside the model (either matching moments in the data or relying on external estimates), while the rest of the parameters need to be estimated. For the rest of this section we provide details about how each model parameter is either set or estimated.

---

[20] The share of husbands with positive child care time is calculated from the sample of ATUS male respondents.

[21] The ATUS number of observations in table 1 includes also the observations with zero child care for the wife, and hence it is not identical to the number of observations in table 3 below.

TABLE 1
DESCRIPTIVE STATISTICS—PSID AND ATUS DATA

	Mean (1)	p25 (2)	Median (3)	p75 (4)
	A. ATUS Descriptive Statistics			
Nonzero child care time (head)	.69			
Nonzero child care time (wife)	.92			
(Imputed) child care annual hours (head)	330	299	331	360
Child care annual hours (wife) including 0's	605	251	503	841
Child care annual hours (wife) excluding 0's	656	303	542	867
Observations	3,151			
	B. PSID Descriptive Statistics			
Total consumption (excluding durables)	42,716	27,608	37,217	51,076
Leisure of husband	2,086	1,730	2,080	2,240
Leisure of wife	2,410	2,080	2,240	2,672
Hours of husband	2,074	1,920	2,080	2,430
Hours of wife	1,750	1,488	1,920	2,080
Hourly wage of husband	30.09	15.89	23.66	36.13
Hourly wage of wife	21.65	12.15	17.89	26.52
Observations	11,195			

NOTE.—ATUS data are from 2003–14 (sample of working married couples in which the wife is aged 25–65 and the youngest child is aged 10 or less). PSID data are from the 1999–2015 (biannual) waves. The sample includes working married couples in which the wife is aged 25–65 and there are no children aged 10 or less. Consumption and wages are in 2010 prices. Leisure is calculated assuming total annual hours are 4,160 ($5 \times 16 \times 52$) (see text for details on further sample restrictions).

A. Parameters Set Outside the Model

Table 2 summarizes the model parameters that are set externally without solving the model or estimating first-order conditions. First we need to set the wage process parameters. We use the estimates from Blundell et al. (2016) for the parameterization of the variances and covariances of transitory and permanent wage shocks. Blundell et al. estimate the stochastic process in (1)–(2), correcting for measurement error (which is crucial for the estimation of the variance of transitory shocks). We rely on their specification with age-constant variances over the life cycle in setting the variance of permanent (transitory) shocks to 0.0303 (0.0275) for husbands and 0.0382 (0.0125) for wives. They report very small (and only marginally significant) positive covariances for the shocks between earners within the family. In contrast to them, who require only estimates of the variance of income shocks, a full solution of our model also requires a characterization of the deterministic component of the life cycle profile of wages, as well as specifying the variance of the initial draw of the permanent component of wages. We assume that husband and wife enter the labor market at 25 and that they retire at age 65. We approximate the life cycle profile of wages using a constant and a quadratic polynomial

TABLE 2
PARAMETERS SET OUTSIDE THE MODEL

Parameter	Value	Description and Source
Wage process:		
$\sigma^2_{v_1}$	.0303	Variance of permanent shocks to husband's wages (Blundell et al. 2016)
$\sigma^2_{v_2}$	.0382	Variance of permanent shocks to wife's wages (Blundell et al. 2016)
$\sigma^2_{u_1}$	.0275	Variance of transitory shocks to husband's wages (Blundell et al. 2016)
$\sigma^2_{u_2}$	.0125	Variance of transitory shocks to wife's wages (Blundell et al. 2016)
$\sigma_{v_1 v_2}$	.0058	Covariance of permanent shocks between spouses (Blundell et al. 2016)
$\sigma_{u_1 u_2}$	.0027	Covariance of transitory shocks between spouses (Blundell et al. 2016)
σ^2_{0,w_1}	.256	Variance of log(wage) at time 0, husband (PSID)
σ^2_{0,w_2}	.285	Variance of log(wage) at time 0, wife (PSID)
$d_{0,w_1}, d_{1,w_1}, d_{2,w_1}$	[1.54, .069, −.0007]	Deterministic age profile parameters, husband (PSID)
$d_{0,w_2}, d_{1,w_2}, d_{2,w_2}$	[1.77, .045, −.0004]	Deterministic age profile parameters, wife (PSID)
Initial assets:		
$\sigma^2_{\log A_0 \mid A_0 > 0}$	2.58	Variance of A_0 when initial assets are positive (PSID)
$P(A_0 > 0)$	.23	Proportion with zero initial assets (PSID)
Tax function:		
χ	2.2	Scale of tax function (Blundell et al. 2016)
μ	.1	Exponent of tax function (Blundell et al. 2016)
b^k	1,000	Monthly taxable income floor, households with kids, $ (TANF and SNAP, 2010)
b^{nk}	367	Monthly taxable income floor, households without kids, $ (SNAP, 2010)
Other parameters:		
γ^k	2,900	Fixed costs of work for women with young children, $ (PSID)
r	.05	Annual return, risk-free assets
β	$1/(1.05)$	Discount factor

NOTE.—See text for details on parameter choices.

in age. To avoid confounding cohort and age effects, we also control for cohort dummies (in five birth year categories). Finally, we include controls for education levels. We use the coefficients from these regressions to generate the average age profile of wages over the life cycle in the model, separately for husbands and wives.[22] To estimate the variance at the time of entry in the labor market, we use the variance of the regression residuals for husbands and wives aged 25–30. The actual estimates are reported in table 2. Both age profiles are concave, and perhaps because of

[22] The coefficients from these regressions are reported as d_{0,w_1}, d_{1,w_1}, and d_{2,w_1} in table 2. To match the mean of log wages over the life cycle, we restrict the sum of the cohort effects to equal zero.

the incidence of early retirement, the male-female wage differential declines after age 50. There is more initial wage dispersion for women than for men.

To discipline the initial assets distribution, we use the distribution of net worth (including housing) in the PSID for families in which the wife is 25–30 years old. In the data, about 23 percent of households start with zero or negative assets. We assume that there is a similar share of zero-assets households in the model. For the nonzero-assets households, we approximate the initial distribution of assets with a lognormal distribution, taking the variance of this distribution from the PSID.

We follow Blundell et al. (2016) when setting the scale and exponent parameters of the tax function (5). However, we extend their work in two directions. First, we allow for a taxable income floor ($b(z)$ in eq. [5]) of $1,000 per month for families with kids (b^k) and of $367 for families without kids (b^{nk}), which is consistent with the TANF and food stamps (SNAP) average allowances for households with no earnings (as of 2010). Second, when the wife is employed, we impose a fixed cost of work—equivalent to a flat tax conditional on employment. For families with kids, this fixed cost equals the average child care cost faced by families with young kids and an employed wife aged 30–35, which we estimate to be around $2,900 (using PSID data on child care spending). For families without kids, we estimate the fixed cost parameter structurally to match wives' participation in the data (see more details in Sec. V.C).

Finally, we simplify family composition in the model by assuming that families either have a young child or do not. Childbirth occurs when the wife is 28, and the child remains "young" for the next 10 years.[23] The last two parameters that we set externally are the interest rate and the discount rate; the actual values chosen are in the last two rows of table 2 and are relatively uncontroversial.

B.   Parameters Estimated Using MRS Equations

As explained above, we break estimation of preference parameters into two steps. In the first step we identify the parameters η, φ_{L_1}, φ_{L_2}, ρ_L, φ_{T_1}, φ_{T_2}, and ρ_T using first-order conditions: the MRS between husband and wife leisure, (7), the MRS between the wife's leisure and household consumption, (8), the MRS between husband and wife child care time, (9), and the MRS between the wife's child care time and household consumption, (10). In the second step (detailed in the next subsection) we estimate the remaining parameters (preference shifters) using estimation by simulation—in particular, we minimize the distance between moments in the

---

[23] We also experimented by introducing some randomness in the age of child arrival (between 26 and 29) but found that none of our conclusions were affected.

data and identical moments simulated in our life cycle model—conditional on the estimates of the wage parameters and the preference parameters that can be identified from estimation of equations (7)–(10).

Estimation of moment conditions (7) and (8) uses PSID data on married couples with children older than 10 and imposes theoretical restrictions across the two equations. Estimation of moment conditions (9) and (10) uses ATUS data on married couples with children aged 10 or younger and also imposes theoretical restrictions across the two equations except for η (the consumption elasticity of intertemporal substitution), which we take from the estimation of (7)–(8) above.[24] In both cases, we assume interior solutions and hence focus on samples of working couples (adding Heckman correction terms to account for sample selectivity). Note that we do not need to impose such selection when estimating preference shifters, nor do we need them when evaluating overall responses of consumption and time use to wage shocks (as we use model-based responses, allowing for corners in time use or assets).

Identification of parameters in equations (7)–(8) requires finding instruments for the endogenous variables (leisure, consumption, and wages).[25] Our instruments include education category dummies for both earners, the wife's actual years of schooling, dummies for the wife's birth cohort and their interactions with the wife's education category, unearned income, and a dummy for whether the household has financial assets that are less than 2 months of household annual income (a traditional indicator for liquidity constraints; Zeldes 1989).[26] Both equations also include a standard (inverse Mills ratio) Heckman selection correction from a probit regression for a working couple, with the husband's education and actual years of schooling used as additional exclusion restrictions.

Turning to the identification of equations (9) and (10), we have similar endogeneity problems and again use a set of instruments for the endogenous variables child care time, consumption, and wages. However, since the ATUS sample is much smaller and we do not have information on assets as in the PSID, we use a smaller instrument set, which excludes unearned income, the liquidity constraints indicator, the dummies for the wife's birth cohort, and their interaction with wife's education category dummies. The inverse Mills ratio is obtained from a similar probit regression that adds the husband's actual years of schooling as an additional exclusion restriction.

---

[24] We do so because in ATUS we have only an imputed consumption measure.

[25] In the theoretical model wages are exogenous. But the possibility of measurement error implies that they should be instrumented too.

[26] We assume that these variables shift wages, time use decisions, and consumption choices but are orthogonal to preference shocks or measurement error. While these may be strong assumptions, our data sets do not offer quasi-experimental variation. When we discuss the results, we show that the results are qualitatively similar if we apply a dynamic panel instruments approach to PSID data.

C.   Parameters Estimated by Matching Moments

The final set of parameters to estimate includes the preference shifters, the fixed costs of female employment for families without children (since we set externally the fixed costs of female employment for families with children), and the equivalence scale for consumption. As explained above, in the model we allow families to either have a young child or not. This implies that we need to estimate the mean of the leisure preference shifter for families with kids ($\phi_{L_1}^k$, $\phi_{L_2}^k$) and without them ($\phi_{L_1}^{nk}$, $\phi_{L_2}^{nk}$). Moreover, for families with kids we also need to estimate the mean of parental time preference shifters ($\phi_{T_1}^k$, $\phi_{T_2}^k$). This is of course not relevant for families without kids since by construction their utility function does not include the parental time subaggregate. We allow for unobserved heterogeneity in wife's leisure preferences, implying that we need to estimate $\sigma_{\varepsilon_{L_2}}^2 = \text{var}(\varepsilon_{L_2,j,t})$ for families with and without kids.[27] The last set of parameters to consider are the preference shifters for consumption ϕ_C. Given that the level of utility is arbitrary, we can normalize one parameter. In particular, we set the $\phi_C^{nk} = 1$ for families without kids. We then estimate the level of this shifter for families with children (ϕ_C^k).

Overall this leaves us with 10 parameters to estimate: preference shifters in husband and wife leisure as a function of having or not having young kids ($\phi_{L_1}^k$, $\phi_{L_1}^{nk}$, $\phi_{L_2}^k$, $\phi_{L_2}^{nk}$), parental time preference shifters ($\phi_{T_1}^k$, $\phi_{T_2}^k$), unobserved heterogeneity in female preferences for leisure, separately by family type ($\sigma_{\varepsilon_{L_2}}^{2,k}$, $\sigma_{\varepsilon_{L_2}}^{2,nk}$), the fixed cost of female employment for families with no young kids (γ^{nk}), and the preference shifter for consumption for families with young kids ϕ_C^k. We choose 11 moments in the data that should be informative about these parameters: male and female hours of work for families with and without young children, male and female parental time hours when young kids are present, the interquartile range of hours and the employment rates of women with and without young kids, and the change in household consumption when a child is born. We calculate these moments when the wife is aged 30–35. To separately match the moments for families with and without kids at these ages, we solve and simulate the model twice—once for families with young kids and once for families without. While there is no one-to-one match between the structural parameters and the targeted moments, some parameters are naturally linked to some of the chosen moments. For example, cross-sectional variation in hours (captured by the interquartile range) should be informative about the variance of unobserved heterogeneity in preferences for leisure; similarly, the arrival of a child leads to reallocation in con-

---

[27] Unobserved heterogeneity is assumed to be fixed over time. We approximate it with a discrete two-state distribution, implying that each family draws either a high or a low value for wife's leisure preferences.

sumption across periods that should be reflected in economies of scale; the levels of hours should be informative about preference shifters in leisure, while the level of time devoted to kids should help pinning down preference shifters for child care time among households with young kids. Finally, employment rates of women with kids (and how they differ from those of women without kids) should reflect the fixed cost of work.

Our two-step estimation approach delivers consistent estimates of the parameters of interest under standard regularity conditions. However, the standard errors of the simulated method of moments (SMM) estimates need to be corrected for the use of first-stage estimates. Following Gourinchas and Parker (2002), we use the gradient of the second-stage objective, along with the first-stage variances, to correct the second-stage covariance matrix of the SMM estimates.

VI.   Results

A.   Parameter Estimates

Panel A of table 3 reports the results of the estimation of preference parameters from equations (7)–(10). Both columns condition on hourly wages for both earners (PSID also conditions on having at least 2 weeks of annual work); they also include an inverse Mills ratio to correct for selection. PSID standard errors are clustered at the household level because of the longitudinal nature of the data set (this is not needed in the ATUS since it is a pure cross section).

The first thing to note is that the estimated consumption Frisch elasticity η is within a plausible range. The estimate is $\eta = 0.903$ with a tight standard error of 0.049—very much in the ballpark of typical literature estimates. Turning to time use parameters, the estimate of ρ_T is negative (-0.197 with a standard error of 0.123), implying that, plausibly, husband's and wife's child care times are substitutes in the home production function for child services. We also find sizable estimates of the utility parameters φ_{T_1} and φ_{T_2}, which are strictly related to the responsiveness of child care time to wage changes. The estimate of φ_{T_2}, in particular, implies a relatively large responsiveness of the wife on the child care time margin. While this may appear surprising at first, note that our sample includes both mothers of preschool children and mothers of older children (aged between 5 and 10), where there is much more flexibility in shifting time allocated to child care (and also "more room to go" given constraints on work schedules higher up in the distribution of hours). Finally, on the leisure time front, the estimate of ρ_L is positive and significant, implying complementarity of leisure time (partners enjoy spending their leisure together). As for leisure responsiveness estimates (φ_{L_1} and φ_{L_2}), they are quite small and not dramatically different across gender.

TABLE 3
PARAMETER ESTIMATES

| | A. MRS ESTIMATES | | |
	Leisure and Consumption (1)		Parental Time (2)
φ_{L_1}	.211 (.037)	φ_{T_1}	.115 (.081)
φ_{L_2}	.162 (.025)	φ_{T_2}	.503 (.201)
ρ_L	.535 (.099)	ρ_T	−.197 (.123)
η	.903 (.049)		
Observations	11,195		2,901
	B. PREFERENCE SHIFTERS		
	With Children		Without Children
ϕ_{L_1}	−8.925 (1.108)		−7.680 (1.013)
ϕ_{L_2}	−9.397 (1.036)		−8.816 (1.024)
ϕ_{T_1}	−23.993 (10.245)		N/A
ϕ_{T_2}	−3.957 (1.201)		N/A
$\sigma^2_{\varepsilon_{L_2}}$	1.476 (.174)		.700 (.087)
γ	(see table 2)		4,794 (438)
ϕ_C	.132 (.024)		Normalized to 0

NOTE.—In panel A the parameters are estimated by generalized method of moments. Standard errors clustered by household are in parentheses. Parameter estimates reported in col. 1 use PSID data; those reported in col. 2 use ATUS and CEX data. In panel B we report parameters estimated by the simulated method of moments (see table 4, panel A, for the list of targeted moments and the main text for data description). All parameters in this panel are for the utility in which all quantities are scaled by total available hours of 4,160. Standard errors are calculated applying a correction for two-step estimation as in Gourinchas and Parker (2002).

Appendix table 1 reports the estimates for the leisure and consumption elasticities obtained using different estimation approaches. In particular, rather than using education, cohort, and unearned income as instruments, we apply a dynamic panel instruments approach, instrumenting the endogenous variables with lagged wages, as well as lagged log difference of leisure of both earners and consumption. While the sample is much smaller (because of the dynamic instruments and the small number of waves), the estimates are qualitatively very similar to the ones reported in panel A of table 3.

Panel B reports the preference shifters and other parameters recovered by matching moments via simulation of the structural model.[28] We match moments for families with and without kids when the wife is 30–35 (using PSID data). While the magnitudes are not directly interpretable, there are a few points that are worth noting. First, given the MRS estimates and the fixed cost of employment for the case with children (which is not estimated but set equal to the average child care costs observed in the data), the estimation implies smaller variance for the wife's leisure preference shifters for families without young children. Second, turning to fixed costs of work, for women with children, fixed cost of work is parameterized to match child care cost ($2,900; see table 2), on top of the child care time cost. The estimated fixed cost of work for families without children is about $4,800, quite close to the parameterized value of families with young children.

How good is the fit of the model? Table 4 shows the data moments along with the simulated moments (i.e., the internal fit of the model). Panel A includes the targeted moments (when the wife's age is 30–35). We estimate 10 parameters targeting 11 moments, and hence we have an overidentified case in which the data moments are not matched mechanically. Nevertheless, for all moments the model does an excellent job.

Because we keep all parameters constant over the life cycle conditional on family composition, a tough test of the model consists of evaluating the match of moments at later stages of the life cycle (which we do not target explicitly). Panel B shows the key labor supply moments when the wife is 50–55 (while, to reiterate, the moments we target in panel A refer to women aged 30–35). The match is qualitatively good. Both in the data and in the model, labor supply of husbands is much lower compared to the case of men with and without children married to younger women (a 200–300 annual hours difference). For wives, labor supply is higher relative to younger women with kids and lower relative to younger women without kids (and the opposite for the dispersion of hours). All labor supply magnitudes at later stages of the life cycle are somewhat higher in the model than in the data. Some of the level difference are likely attributable to the fact that we ignore changes in cohort and education composition when calculating the moments at different ages in the data.

B.   Model Responses

The most relevant use of the structural model is that we can now use our estimates to simulate the response of consumption and time use to tran-

---

[28] Standard errors are calculated applying a correction for two-step estimation similarly to Gourinchas and Parker (2002).

TABLE 4
List of Moments Used in Estimation

	Data	Model
	A. Targeted Moment (Wife 30–35)	
Hours of work: wife with young kids	1,251	1,248
Hours of work: wife without young kids	1,814	1,816
Hours of work: husband with young kids	2,218	2,225
Hours of work: husband without young kids	2,126	2,121
Hours of parental time: wife with young kids	784	778
Hours of parental time: husband with young kids	346	337
p75–p50 hours: wife with young kids	504	434
p75–p50 hours: wife without young kids	212	230
Employment probability of wife with young kids	.77	.76
Employment probability of wife without young kids	.90	.90
Change in consumption when kid is born	.075	.073
	B. Nontargeted Moment (Wife 50–55, No Kids)	
Hours of work: wife	1,411	1,665
Hours of work: husband	1,900	2,038
Employment probability of wife	.78	.83
p75–p50 hours: wife	283	242

Note.—All targeted moments (panel A) are calculated in the data and in the model for households in which the wife is aged 30–35 (with or without young children). All nontargeted moments (panel B) are calculated for households in which the wife is 50–55. All data moments except hours of parental time are from the PSID. Hours of parental time moments are from the ATUS.

sitory and permanent wage shocks over the life cycle. As explained above, response to transitory and permanent shocks will also approximately coincide with, respectively, policy-relevant Frisch and Marshallian elasticities under some special circumstances (interior hours solution, no progressive taxation, and no borrowing constraints). The literature that has focused on exercises of this type typically neglects family labor supply and ignores allocation of time use when children are present. In our context, a key dimension of heterogeneity in the elasticities is indeed the presence and age of children. For this reason, we separately report responses for households with and without young children (aged 10 or less). We obtain the two types of responses by solving and simulating the model twice: first, under the assumption that there are young children in the households and, second, under the assumption that there are no children.

We focus on consumption and time use responses at age 30, when all households in our simulated data have a young child (recall that the assumption is that childbirth occurs when the woman is 28). Our exercise consists of simulating the impact of a 10 percent shock to a given earner's wage. We consider transitory and permanent shocks and look at the impact of shocks received by husband and wife separately. Table 5 reports the consumption and hours responses (hours of work on the intensive

TABLE 5
CONSUMPTION AND LABOR SUPPLY RESPONSES TO TRANSITORY AND PERMANENT SHOCKS

	TOTAL RESPONSE						EXTENSIVE VS. INTENSIVE MARGIN			
	C		H_1		H_2		E_2		$H_2 \mid$ Employed	
	With Kids	Without Kids	With Kids	Without Kids	With Kids	Without Kids	With Kids	Without Kids	With Kids	Without Kids
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Transitory:										
Husband	.119	.123	.180	.222	−.076	.001	−.051	.005	−.041	.006
Wife	.130	.135	.000	−.006	.703	.394	.574	.280	.329	.167
Permanent:										
Husband	.393	.410	.105	.116	−.296	−.140	−.193	−.065	−.170	−.088
Wife	.353	.375	−.070	−.106	.531	.304	.491	.266	.208	.086

NOTE.—Model-simulated responses for transitory and permanent shocks.

and extensive margins), while table 6 reports time use responses (leisure and time devoted to children) corresponding to the various cases. Temporary shocks can be interpreted as one-time changes in the tax code (e.g., a temporary increase/reduction in the payroll tax rate), while permanent shocks are more likely to reflect structural tax reforms. The tables report response elasticities "on impact" (we comment below on longer-run responses).

Starting with consumption, we obtain the traditional result that consumption responds much more to permanent shocks than to transitory shocks. The response to male permanent shocks is larger than to female permanent shocks because of the larger role played by men on family earnings. The nonzero response to transitory shock is due to nonseparability through female employment and the presence of (for some households) binding liquidity constraints (which we study more in detail below).

Consider now the effect of a positive transitory shock to the woman's wage—a wage change that has mostly substitution effects and very modest wealth effects. A woman without kids increases her hours of work and reduces time devoted to leisure: the hours elasticity is 0.39 (col. 6 of table 5). The impact of the same change on the labor supply of a woman with kids is larger (an elasticity of 0.7) because she has another margin of adjustment, child care time (col. 5). For men (cols. 3 and 4), the responses to own transitory shocks in the two cases (with and without children) are similar because men's child care time is small to begin with and because of their lower preferences for spending time with the kids. This confirms the typical result in the literature of larger elasticity for women than for men but highlights a component of heterogeneity (time use in the presence of children) that is often ignored.

The hours cross-responses are also interesting. When the husband's wage increases temporarily, so that he works longer hours, wives in households without kids respond very little. This reflects pure leisure nonseparability: spouses enjoy spending nonworking time together, along with small wealth effects of the transitory shock, giving rise together to virtually no response of leisure or hours. However, hours are reduced (−0.08) in households with kids because wives now have to confront the decline in child care time of the husband, for which her child care hours are substitutes. While these effects are small, they are consistent with the idea that the presence of young children can affect cross–labor supply responses.

Our model allows us also to recover the elasticities of each earner's time use with respect to permanent shifts in wages. For both husband and wife the own elasticities with respect to temporary wage changes are larger than the elasticity with respect to long-run wage changes (independently of children presence), implying a small to moderate wealth

TABLE 6
LEISURE AND PARENTAL TIME RESPONSES TO TRANSITORY AND PERMANENT SHOCKS

	L_1		L_2		T_1	T_2
	With Kids (1)	Without Kids (2)	With Kids (3)	Without Kids (4)	With Kids (5)	With Kids (6)
Transitory:						
Husband	−.230	−.231	−.003	−.001	−.095	.131
Wife	−.007	.006	−.217	−.309	.033	−.538
Permanent:						
Husband	−.131	−.120	.078	.110	−.067	.261
Wife	.085	.110	−.151	−.238	.058	−.443

NOTE.—Model-simulated responses for transitory and permanent shocks.

effect. The fact that the own permanent wage elasticities are positive implies that the substitution effect dominates. Hence, permanent declines in wages induce declines in labor supply for both partners (albeit larger for women, a traditional result in the labor supply literature).

Also of interest is the fact that the cross-Marshallian elasticities are much larger for women than for men. An important observation about these results is the size of these effects for households with and without young children. While for men there is little difference in the elasticities, the responses are much larger in absolute value for women with young kids than for those without. For women in couples without young kids, a permanent decline in the husband's wage means he works less and earns less. She works more because of an added worker effect. In couples with children, husbands who work less allocate some of their leisure time to children, implying that the wife can work more because of substitutability effects in home production of child services.

In the last four columns of table 5 we decompose hours elasticities into extensive and intensive margin effects. The reason why this is a useful exercise is that while young men display very high employment rates, employment rates for young women with children are much lower. This suggests that when analyzing the response of total hours of work to shocks, extensive margin responses are potentially important. In columns 7–8 of table 5 we report extensive margin responses, while in columns 9–10 we report intensive margin responses, conditioning on the wife's being employed in both scenarios (i.e., with and without the shock).[29] We find

---

[29] The intensive and extensive margin elasticities, as we define them, need not sum up to the total hours elasticity. The reason is that the total hours elasticity includes another component, the choice of how many hours to work when employment changes because of the impact of the shock. For example, if the shock implies an increase in employment, it is typically the case that newly employed workers tend to work fewer hours, on average, than those employed both before and after the shock (pushing down the total hours elasticity).
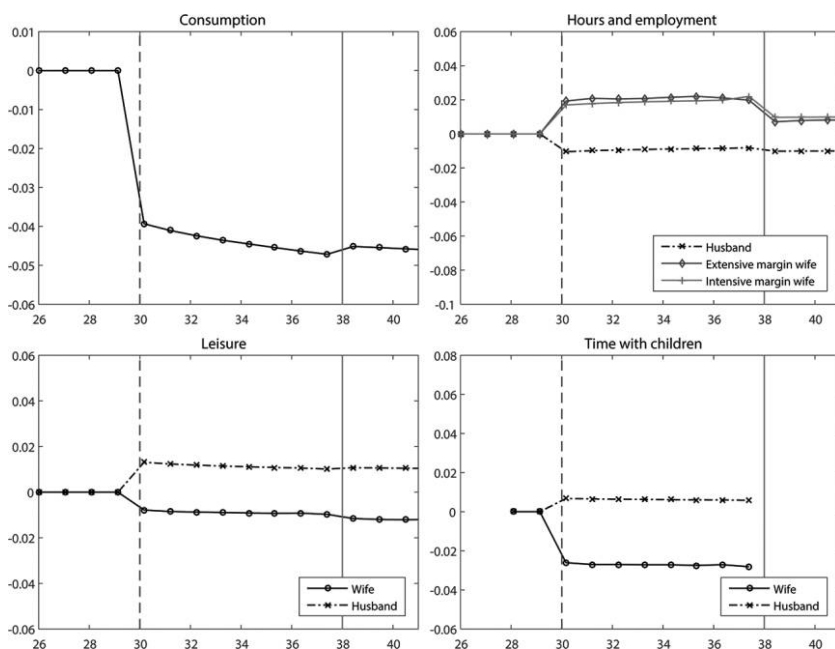
FIG. 1.—Impulse responses to a permanent 10 percent decline to husband's wages. Color version available as an online enhancement.

that both extensive and intensive margin responses are important, although extensive margin responses appear somewhat larger in most cases.

So far we have focused on the responses upon impact of the shock. While for transitory shocks responses have no long-run effects by design, for permanent shocks this is not the case. In fact, for permanent shocks the "on impact" and "long-run" responses coincide in the case of stationary preferences, but not when preferences depend on the presence of kids—as in our context. Figures 1 and 2 show the impulse responses to a 10 percent negative permanent shock to the husband's and wife's wages, respectively. We assume that the shock hits at age 30 (the dashed vertical line is drawn to the left of the shock). The four plots show the impulse responses of consumption, labor supply, leisure, and time devoted to kids. Younger children are present in the household only to the left of the solid vertical line. The changes in the variables of interest at age 30 reproduce the "on-impact" elasticities reported in tables 5 and 6. As is clear from these figures, the long-run labor supply responses to the permanent shocks are attenuated when children no longer require parental time,
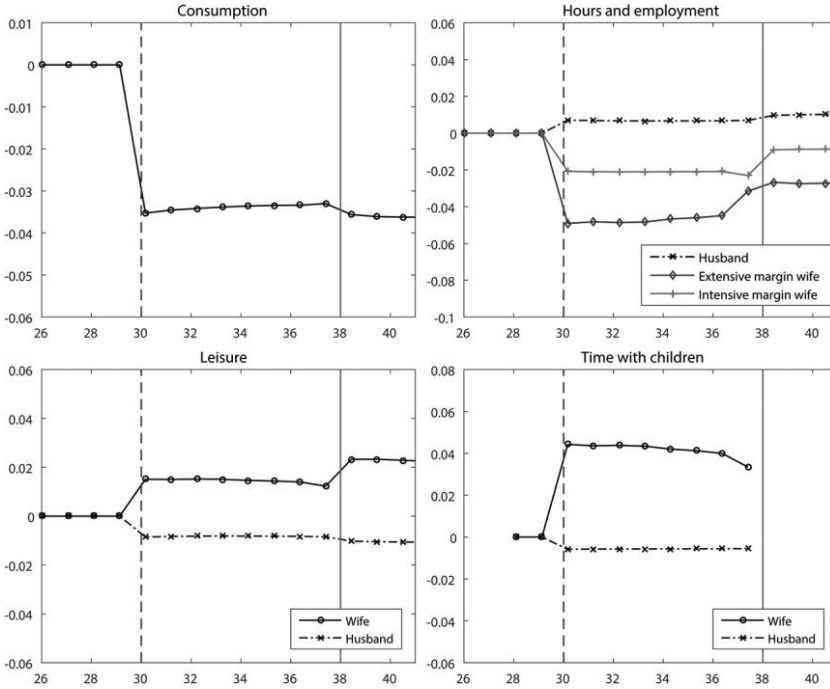
F<small>IG</small>. 2.—Impulse responses to a permanent 10 percent decline to wife's wages. Color version available as an online enhancement.

consistent with the finding of smaller elasticities for the case without young children reported in table 5. An important implication of these findings is that life cycle labor supply elasticities can be dramatically different from short-term labor supply elasticities and either overstate or understate them depending on the presence of young children in the household—a finding consistent with Ghez and Becker's quote in the introduction.

C. Decomposition of Hours Responses

We now use equation (11), reproduced here for convenience, to decompose the labor supply elasticities into leisure and parental time responses:

$$\eta_{H_i, W_j}\big|_a \equiv -\eta_{L_i, W_j}\big|_a \frac{L_i}{H_i} - \eta_{T_i, W_j}\big|_a \frac{T_i}{H_i} \tag{12}$$

for $i, j = \{1, 2\}$ and for two types of wage changes, transitory and permanent.

Table 5 reports labor supply responses (the left-hand side of [12]), while table 6 decomposes these responses into leisure and child care time responses, separately for households with and without young children.

From table 5, husbands have low transitory wage elasticities ($\eta_{H_i,W_i}\big|_{\Delta W_i = \text{trans.}}$ is between 0.18 and 0.22), which are essentially independent of the presence of children. For households with no kids, husband's leisure-to-hours of work ratio is close to one, and hence leisure elasticity is about the same as the negative of the hours elasticity. For households with children, the dominant force is leisure, as the husband's Frisch elasticity for child care time is small and the child care time-to-hours of work ratio is small to begin with. In any case, his reduction in child care time is compensated by an increase in child care time of his wife. In contrast, labor supply elasticities in response to transitory shocks for women when young children are present are much larger (0.7). Here, however, child care time responses are a dominant force: she reduces both her leisure and her time with kids. He consumes less leisure (because of leisure complementarity) and devotes more time to kids (because of substitutability). Her response to transitory shocks to his wages reflects the exact same mechanisms, only amplified. When he receives a positive transitory wage shock, she consumes less leisure (complementarity again) and more child care time (substitutability again).

An increase in the husband's permanent wage increases his hours and reduces both leisure and child care time; the wife steps up by spending less time working in the market and more time working at home (with children). In contrast, if her permanent wage increases, he reduces hours and increases leisure and (by only a small amount) child care time, while her child care time declines considerably. Hence, total time allocated to child services production declines significantly.

The lower two plots in figures 1 and 2 show that the increase in labor supply in response to permanent negative shocks when the children no longer require parental time is driven by the latter. For example, in figure 2, even though wife's leisure response is more positive after the children are no longer in need of care (to the right of the solid vertical line), the total hours response is less negative because most of the hours response was driven by the response of child care time.

D. Consumption Insurance

In Blundell et al. (2016) the focus is to understand the effect of permanent wage shocks on the household's ability to smooth consumption through family labor supply, the tax/transfer system, and savings. In our model, the labor supply response of the partners depends on three margins: leisure complementarity, specialization in child care production, and wealth (added worker) effect. We illustrate these channels by de-

composing the average response to a permanent shock faced by the primary earner (the male, who earns, on average, about 66 percent of household pretax labor earnings) at age 30.

From table 7, a 10 percent permanent decline in the male's hourly wage results in a 3.9 percent decline in consumption. Savings and taxes attenuate the response of consumption relative to after-tax and transfers household earnings (a 5 percent decline) and relative to before-tax (after-transfers) household earnings (a 5.6 percent decline), respectively.[30]

In the rest of table 7 we decompose the earnings response of the husband (a 10.7 percent decline in before-tax earnings) and the earnings response of the wife (a 2.0 percent rise in before-tax earnings) into the three time use margins. The total before-tax and transfers earnings response is 6.4 percent (the average of the husband and wife responses, weighted by their earnings shares).

On the husband's side, hours decline 1.0 percent (as the substitution effect dominates the wealth effect of a permanent decline in wages), while leisure and child care time increase by about 1.3 percent and 0.7 percent, respectively. On the wife's side, while the total pretax earnings increase is 2 percent, the total hours increase is 3 percent. This discrepancy, in the absence of changes in her wages, may seem puzzling at first. The issue is that this is an aggregate response. If assortative mating in wages is relatively small (as it is in the data), low-wage women need to work more to make up for a similar decline (in levels) in the husband's earnings. This induces a postshock decline in the cross-sectional hour-wage covariance, which, in turn, generates a smaller aggregate response in earnings than in hours.

Looking at female time use responses, the bulk of the adjustment in hours is at the expense of child care time (a 2.6 percent decline, over three times larger than the decline in leisure hours). A natural next step would be to study the implications for children outcomes of such shock: do children's test scores or other indicators of school performance decline when the mother is pushed into work by shocks faced by the primary earner? The CDS, a module available in the PSID for some years, contains information on such outcomes, and it is a natural source for questions of this sort.

E.   Borrowing Constraints

In our model, households are subject to an exogenous borrowing constraint, requiring assets to be nonnegative in each period. We discuss two questions here. First, are there dramatic differences of behavior be-

---

[30] Note that the overall consumption response is somewhat larger than the one estimated in Blundell et al. (2016). This is not surprising given that we are focusing on young households in the early stage of their life cycle.

TABLE 7
INSURANCE EFFECTS

Consumption		−3.9%
After-tax and transfers household earnings		−5.0%
Before-tax (after-transfers) household earnings		−5.6%

	Husband	Wife
Earner's average share of before-tax earnings	.66	.34
Earner's before-tax and transfers earnings response:	−10.7%	+2.0%
Hours	−1.0%	+3.0%
Leisure	+1.3%	−.8%
Parental time	+.7%	−2.6%

NOTE.—Insurance decomposition calculations based on model-simulated responses to a 10 percent permanent decline in the husband's wage.

tween households for which the constraint is binding and those for which it is not? Second, are there alternative ways of modeling borrowing restrictions in our context?

To evaluate the sensitivity to liquidity constraints in our model, we use model simulations to study the extent of response heterogeneity in two groups of households (both with young children): those in the bottom quintile of the wealth distribution (and hence more likely to experience binding liquidity constraints) and those in the top quintile (and hence much less likely to be subject to constraints). We report consumption, labor supply, and time use goods responses to transitory and permanent shocks in tables 8 and 9 (comparable to tables 5 and 6, which instead refer to the whole sample). All responses are measured at age 30.[31] As expected, consumption response to transitory shocks is much larger for the low-assets households, with very small responses for high-assets households. Consumption response to permanent shocks is slightly larger for low-assets households as well, but given that assets are not key in insuring against permanent shocks, the differences are much smaller compared to transitory shocks. For time use and labor supply responses to transitory shocks the wealth effect kicks in stronger in offsetting the typical substitution effect. The time use and labor supply responses to transitory shocks of the liquidity-constrained households therefore tend to be similar to their responses to permanent shocks. For example, when a household with young kids is hit by a positive transitory shock to the husband's wage, complementarity between leisures of the two earners implies a decline of leisure for the wife, and substitution of time use with kids implies an increase in her time with kids. However, the wealth effect associated with the transitory shock implies that low-assets wives now want to increase their leisure,

[31] Splitting the sample on the basis of assets early on in the life cycle is better than doing it at a later stage, since the distribution of initial period assets (at age 25) is a random draw from a realistic asset distribution.

TABLE 8
CONSUMPTION AND LABOR SUPPLY RESPONSES: LOWEST VS. HIGHEST QUINTILE OF ASSETS AT THE BEGINNING OF THE LIFE CYCLE

| | TOTAL RESPONSE | | | | | | EXTENSIVE VS. INTENSIVE MARGIN | | | |
| | C | | H_1 | | H_2 | | E_2 | | $H_2 \vert$ Employed | |
	Low (1)	High (2)	Low (3)	High (4)	Low (5)	High (6)	Low (7)	High (8)	Low (9)	High (10)
Transitory:										
Husband	.175	.030	.152	.251	−.120	.024	−.074	.006	−.070	.022
Wife	.190	.036	−.020	.041	.606	.975	.493	.803	.282	.452
Permanent:										
Husband	.415	.344	.092	.144	−.297	−.299	−.188	−.210	−.173	−.161
Wife	.361	.322	−.070	−.070	.493	.671	.445	.649	.198	.248

NOTE.—Model-simulated responses for transitory and permanent shocks. Asset quintile is defined according to simulated initial assets.

as well as increase even further their child care time. The overall impact is a reduction in wife's labor supply, which is stronger than what we see on average and is in contrast with the positive labor supply response for the high-assets households.

While our assumption of an exogenous borrowing constraint is a frequent characterization of borrowing limits in life cycle models of consumption and labor supply, it is not uncontroversial. Several papers consider instead "natural borrowing constraints" à la Aiyagari (1994), in which (because of an Inada condition on preferences) households never choose to have a level of assets that could induce, with positive probability, a future state of the world in which consumption is exactly zero. In this sense, the borrowing limit is the maximum amount that an individual can repay with certainty. As shown by Hai and Heckman (2017), this no-

TABLE 9
LEISURE AND PARENTAL TIME RESPONSES: LOWEST VS. HIGHEST QUINTILE OF ASSETS AT THE BEGINNING OF THE LIFE CYCLE

| | L_1 | | L_2 | | T_1 | | T_2 | |
	Low (1)	High (2)	Low (3)	High (4)	Low (5)	High (6)	Low (7)	High (8)
Transitory:								
Husband	−.205	−.272	.017	−.038	−.088	−.106	.165	.073
Wife	−.020	.051	−.203	−.230	.040	.020	−.528	−.536
Permanent:								
Husband	−.121	−.154	.090	.051	−.065	−.073	.284	.211
Wife	.090	.071	−.153	−.144	.059	.054	−.454	−.406

NOTE.—Model-simulated responses for transitory and permanent shocks. Asset quintile is defined according to simulated initial assets.

tion of natural borrowing constraints needs rethinking in a context with endogenous labor supply or human capital investments, since lenders cannot force individuals to work or to invest in human capital to repay their debt. As a consequence, the borrowing limits imposed by lenders must take into account the incentive compatibility constraints of the borrowers. While these are important theoretically based versions of borrowing limits, endogenizing the nature of liquidity constraints is beyond the scope of our paper.

F.   *Approximation versus Simulation*

Both Blundell, Pistaferri, and Preston (2008) and Blundell et al. (2016) consider a life cycle framework similar to ours but make two key simplifications. In Blundell et al. (2008) there is no hours margin and no secondary earner. In Blundell et al. (2016) there is no time use choice and hence no child care services to produce. Both papers obtain analytical expressions for the response of consumption (and hours) to income and wage shocks by considering approximations of the first-order conditions of the problem and approximations of the lifetime budget constraint. An important limitation of approximation-based approaches is that they cannot handle corner solutions in hours or assets. Moreover, the procedure generates approximation errors that are complex to characterize and difficult to quantify.

Kaplan and Violante (2010) implicitly evaluate the importance of these approximation errors, in a model similar to Blundell et al. (2008), by comparing moments obtained from the approximation-based solution (i.e., the estimation of partial insurance coefficients, which use covariance terms for consumption and income dynamics) against the same moments obtained from simulating the actual life cycle model (where no errors are present). They show that the discrepancies are quite small and are important only for the estimation of the partial insurance coefficient for permanent shocks in a scenario with exogenously imposed liquidity constraints (but not when liquidity constraints are of the "natural," Aiyagari type), and only early on in the life cycle. This shows the importance of corner solutions in assets.

In a companion paper (Blundell et al. 2017), we consider the approximation error bias in the more general case with two earners and endogenous labor supply, similar to the setting considered by Blundell et al. (2016). This allows us to consider the role of corner solutions in hours. We conduct Monte Carlo simulations to evaluate the performance of approximation-based estimators on these simulated data. We show that for most Frisch elasticities the bias is small and that liquidity constraints have relatively small effects on the parameter estimates. We further evaluate the sensitivity of the bias to the choice of functional form and the pos-

sibility of using approximations to guide the choice of moments for indirect inference in a formal structural estimation setting. While these Monte Carlo exercises do not explicitly take into account time use, in a previous version of our paper (Blundell, Pistaferri, and Saporta-Eksten 2015) we extended the approximation-based estimation strategy and estimated Frisch elasticities for the model with time use (applying a somewhat different identification approach and data). Reassuringly, the results were broadly similar to those we report here.

VII. Counterfactual Policy Experiments

In this section we discuss the results of two policy experiments. In the first experiment we provide households with young children with an unconditional subsidy. This is achieved by changing the value of the function $b(z)$ in the tax function:

$$T(z, H_1 W_1 + H_2 W_2) \approx \chi(b(z) + H_1 W_1 + H_2 W_2)^{1 - \mu_t}.$$

In the second experiment we provide households with young children with some compensation for child-related fixed costs of work. This is achieved by changing the value of the parameter $\gamma(z)$ in the adjusted consumption value $\tilde{C}$:

$$\tilde{C} = \frac{C - \gamma(z) E_2}{a(z)}.$$

To make the two policies comparable (in a revenue-neutral sense), we set the experiments such that they give rise to the same deficit, expressed in dollar terms relative to the status quo. Specifically, both subsidies change average discounted after-tax income by 2 percent.[32] We evaluate the welfare impact of the program using the equivalent proportional increase in consumption over the life cycle ("willingness to pay"), that is, the level of consumption that would make an average household indifferent between the pre- and postpolicy regimes. We report changes in consumption and time use over three segments of the life cycle: the years before the birth of children, the period when the household has young children (below the age of 10), and the period that follows it.

As the results in table 10 show, the two policy changes produce very different behavioral and welfare responses. The unconditional subsidy boosts consumption especially in periods when the household has young chil-

---

[32] The value of the subsidy is calculated for the period when the household has young children. Since changes to policy parameters generate behavioral responses, we use an iterative procedure to find the changes in policy parameters that generate the required 2 percent change in after-tax income.

TABLE 10
POLICY EXPERIMENTS

	C (1)	H_1 (2)	H_2 (3)	E_2 (4)	L_1 (5)	L_2 (6)	T_1 (7)	T_2 (8)
	A. Experiment 1: Unconditional Subsidy for Families with Young Children							
Total	.6%	−.4%	−.7%	−.4%	.4%	.3%		
Before young children	.9%	−.4%	−.5%	−.2%	.4%	.4%		
With young children	1.3%	−.6%	−1.8%	−1.0%	.8%	.7%	.2%	1.0%
After young children	.1%	−.1%	−.1%	−.1%	.1%	.1%		
Consumption equivalent utility value	.95%							
	B. Experiment 2: Employment Subsidy for Wives with Young Children							
Total	.1%	−.2%	1.9%	4.6%	.2%	−.5%		
Before young children	.9%	−.4%	−.5%	−.1%	.4%	.4%		
With young children	−.3%	−.3%	6.5%	13.1%	.3%	−1.7%	.3%	−5.6%
After young children	.1%	−.1%	−.1%	~0%	.1%	.1%		
Consumption equivalent utility value	.17%							

NOTE.—See text for detailed description.

dren; these are the periods when the subsidy is received and are also the periods when the price of consumption is higher. Interestingly, while there is little increase in consumption in the post–young children period, there is a significant increase in consumption in the prechildren period. This is a period when households tend to be liquidity constrained, at least partly accumulating savings for the period when they will have young children. Increasing income in the period when children are present (through the subsidy) allows them to reduce savings and increase consumption in the prechildren period. Consistent with the direct wealth effect of the unconditional subsidy, households also increase leisure and time spent with kids at the expenses of hours of work.

Unlike the unconditional subsidy, the employment subsidy has a large impact on labor supply, especially of women. Because of the large decline in the fixed cost of employment, the subsidy increases participation for women with young children by 13 percent (from a status quo value of 0.75) and increases hours worked by 6.5 percent. The employment subsidy also reduces leisure and especially time spent with kids for the wife, along with a slight reduction in labor supply and a slight increase in time spent with the kids for the husband. The latter is far too small to offset the decline in child care hours of the wife, resulting in an overall decline in family child care hours. Note that our model does not explicitly account for how changes in parental employment affect the welfare of children (see Del Boca, Flinn, and Wiswall [2014] for some evidence). As in the first policy, consumption increases in the prechildren period for similar

reasons (the large decline in the fixed cost of work reduces the need to save in anticipation of the arrival of children); in contrast to the first policy, however, the change in consumption during the period when households have younger children is slightly negative because the policy reduces consumption expenditure on the fixed cost of work.

We also report, separately for each policy experiment, the "willingness to pay" (expressed in terms of lifetime consumption). Households appear to gain much more from the unconditional subsidy than from receiving compensation for the fixed cost of work induced by the presence of children. This result is partly intuitive: households have greater leverage in deciding how to allocate the unconditional subsidy to the various margins of behavior. However, this does not necessarily mean that the two policies would be equivalent if all households worked. In fact, borrowing constraints make an extra unconditional dollar much more valuable, in welfare terms, than a conditional subsidy even conditioning on work.

VIII.   Conclusions

In this paper, we have considered a life cycle model in which households enjoy utility from the leisure times of husband and wife and from consuming both non–child care goods and child care services. The latter, in the spirit of one of Becker's seminal contributions, is home produced using as inputs the time allocated by parents to the care of their children and child care–related goods. The wage rate is the opportunity cost of time use alternative to the market (leisure or child care).

We argue that in an environment with uncertainty, hours devoted to the various activities (work, leisure, child care) serve two roles. First, they can be used to produce "commodities," such as child care (as substitute inputs of a production function) or companionship (as complements in utility). Second, they can be changed to provide insurance against wage shocks (either in response to own shocks if wealth effects dominate substitution effects or as added worker effects in response to shocks faced by the other earner).

We find that the response of a mother's labor supply to a transitory wage shock faced by her spouse (cross-Frisch elasticity) is negative. The reason is that the large complementarity desire to spend time together is more than counteracted by the input substitutability in the production function of child care. Similarly, we find that own-Frisch elasticities are large because usual intertemporal substitution effects are compounded by reducing time spent with children because of the higher opportunity cost induced by the wage change.

Our results suggest that, although some of the female labor supply responses are derived from an adjustment in female leisure hours, there is a large response in the time she spends with her children. For an adverse

permanent shock to the husband's wage this implies a significant reduction in her time with her children. The estimates suggest this is not made up with a similar increase in his time with the child. Calculating the resulting impact on the child's well-being is beyond the scope of this paper, but the results suggest that the reduction of the mother's child care time should be an important part of any analysis of the consequences of policies or external shocks that incentivize mothers with young children into work.

While the particular values we have estimated for these trade-offs will depend on our assumption, which separates utility into two aggregates, one comprising consumption and leisure times of the parents and the other comprising the production of child care services, they are present also under less restrictive assumptions. We have also assumed that households are not subject to hours constraints and that they are not subject to employment shocks (such as job destruction or slow arrival rate of offers while out of work). We plan to explore the consequences of more general utility and home production specifications and of labor demand issues in future work.

References

Aguiar, Mark, and Erik Hurst. 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *Q.J.E.* 122 (3): 969–1006.
Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use during the Great Recession." *A.E.R.* 103 (5): 1664–96.
Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Q.J.E.* 109 (3): 659–84.
Altonji, Joseph G. 1986. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *J.P.E.* 94 (3, pt. 2): S176–S215.
Angrist, Joshua D., and William N. Evans. 1998. "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size." *A.E.R.* 88 (3): 450–77.
Attanasio, Orazio, Hamish Low, and Virginia Sanchez-Marcos. 2005. "Female Labor Supply as Insurance Against Idiosyncratic Risk." *J. European Econ. Assoc.* 3 (2–3): 755–64.
———. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *A.E.R.* 98 (4): 1517–52.
Becker, Gary S. 1965. "A Theory of the Allocation of Time." *Econ. J.* 75 (299): 493–517
———. 1976. *The Economic Approach to Human Behavior.* Chicago: Univ. Chicago Press.
Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *A.E.R.* 98 (5): 1887–1921.
Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2015. "Children, Time Allocation and Consumption Insurance." Working Paper no. W15/13, Inst. Fiscal Studies, London.
———. 2016. "Consumption Inequality and Family Labor Supply." *A.E.R.* 106 (2): 387–435.

———. 2017. "Comparing Estimated Life Cycle Responses Using Approximation and Full Scale Solution Approaches." Manuscript, Univ. Coll. London.

Browning, Martin, Pierre-André Chiappori, and Yoram Weiss. 2014. *Economics of the Family*. Cambridge: Cambridge Univ. Press.

Chiappori, Pierre-André, and Maurizio Mazzocco. 2017. "Static and Intertemporal Household Decisions." *J. Econ. Literature* 55 (3): 985–1045.

Cogan, John. 1981. "Fixed Costs and Labor Supply." *Econometrica* 49:945–64.

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2014. "Household Choices and Child Development." *Rev. Econ. Studies* 81:137–85.

Ghez, Gilbert, and Gary S. Becker. 1975. *The Allocation of Time and Goods over the Life Cycle*. New York: Columbia Univ. Press (for NBER).

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70 (1): 47–89.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney. 2008. "Parental Education and Parental Time with Children." *J. Econ. Perspectives* 22 (3): 23–46.

Hai, Rong, and James J. Heckman. 2017. "Inequality in Human Capital and Endogenous Credit Constraints." *Rev. Econ. Dynamics* 25:4–36.

Heckman, James J. 1974. "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption over the Life Cycle." *A.E.R.* 64 (1): 188–94.

———. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153–61.

———. 1988. "Time Constraints and Household Demand Functions." In *Research in Population Economics*, edited by T. Paul Schultz. Greenwich, CT: JAI.

Hyslop, Dean R. 2001. "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings." *A.E.R.* 91 (4): 755–77.

Juhn, Chinhui, and Simon Potter. 2007. "Is There Still an Added Worker Effect?" Staff Report no. 310, Fed. Reserve Bank New York.

Kaplan, Greg, and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Econ. J.: Macroeconomics* 2 (4): 53–87.

Keane, Michael P. 2011. "Labor Supply and Taxes: A Survey." *J. Econ. Literature* 49 (4): 961–1075.

Lundberg, Shelly. 1985. "The Added Worker Effect." *J. Labor Econ.* 3 (1): 11–37.

Meghir, Costas, and Luigi Pistaferri. 2011. "Earnings, Consumption and Life-cycle Choices." In *Handbook of Labor Economics*, vol. 4B, edited by Orley Ashenfelter and David Card, 773–854. Amsterdam: Elsevier.

Ramey, Garey, and Valerie A. Ramey. 2010. "The Rug Rat Race." *Brookings Papers Econ. Activity* (1): 129–76.

Stephens, Mel. 2002. "Worker Displacement and the Added Worker Effect." *J. Labor Econ.* 20 (3): 504–37.

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *J.P.E.* 97 (2): 305–46.