CHILDREN, TIME ALLOCATION AND CONSUMPTION INSURANCE*

Richard Blundell, Luigi Pistaferri and Itay Saporta-Eksten†

March 2015

Abstract

We consider the life-cycle problem of a household that in each period decides how much to consume and how to allocate spouses' time to work, leisure, and childcare. In an environment with uncertainty, the allocation of goods and time over the life cycle plays the further role of providing insurance against shocks. We use longitudinal data on consumption, and husband and wife separate information on hourly wages, hours of work, and time spent with children to estimate structural parameters measuring the sensitivity of consumption and time allocation choices to transitory and permanent wage shocks. These structural parameters provide a full picture regarding the ability of household to smooth marginal utility in response to shocks. In addition, information on hours of work and hours spent on childcare allows to decompose overall Frisch response into two components, one reflecting the degree of complementarity between husband's and wife's leisure ("companionship" or "love") and another reflecting the degree of substitutability of their childcare time in the production of childcare services.

*We thank participants at the Becker Memorial Conference at the University of Chicago for comments. The authors gratefully acknowledge financial support from the UK Economic and Social Research Council (Blundell), and the ERC starting grant 284024 (Pistaferri). All errors are ours.

†Blundell: Department of Economics, University College London, and Institute for Fiscal Studies, Gower Street, London WC1E 6BT, UK (email: r.blundell@ucl.ac.uk); Pistaferri: Department of Economics, Stanford University, Stanford, CA 94305 (email: pista@stanford.edu); Saporta-Eksten: The Eitan Berglas School of Economics, Tel Aviv, Israel, and Department of Economics, University College London (email: itaysap@post.tau.ac.il).

1 Introduction

In a series of seminal contributions, Becker (1965) and Ghez and Becker (1975), emphasized the importance of understanding the effect of children on how households allocate goods and time over the life cycle. To quote from the 1975 volume, "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." One conclusion of that research was that the presence and age of children will change the optimal allocation of parents' time to work in the paid labor market, leisure, and childcare. Responses to wage and income changes will 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 plays the further important role of providing insurance against shocks to household resources. The key parameters governing the ability to insure 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, e.g., 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 to be allocated to the production of childcare services, it is clear that the presence of children impacts also the insurance opportunities of the household. 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 address this goal, we extend our previous work (Blundell, Pistaferri and Saporta-Eksten, 2014) to explicitly take into account the influence of children (both their presence and their ages) on consumption and time allocation decisions of husband and wife, and hence on the household's ability to insure themselves against shocks. Children place additional demands on time and goods, so that in our context it is marginal utility rather than consumption itself that is smoothed.

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 time of the two spouses, as well as market goods, are inputs in the production of "commodities") or by formal non-separabilities in the utility function (the marginal utility of leisure of one spouse depends on the leisure time of the other).

In general, we might expect families 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

[1] Becker (1976) attributed the existence of marriage to "the desire to raise own children" (p. 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" (p.

model family formation, but in order to capture the fact that people have preference for spending time together or with their children (a public good), we consider the possibility that preferences are non-separable and complementary over leisure times. In Blundell, Pistaferri and Saporta-Eksten (2014) we found evidence of Frisch complementarity for time within couples, although we did not relate this to children or use any measure of time use. Even then, we found that family labor supply provides insurance against long-run 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. 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 non-separability 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".

To address the issues of interest we merge data from the Panel Study of Income Dynamics (PSID) and from the Child Development Supplement (CDS), a supplement to the PSID which represents a relatively unexplored source of information on parents' childcare time.[2] The CDS provides a comprehensive accounting of parental (or caregiver) time inputs to children as well as other aspects of the way children spend their time. The data structure allows us to recover how much time each parent spent with the surveyed child over a 24-hours diary period. There are a few issues to take into account when using the time diary data in our analysis, which we discuss at length. However, we argue that time use data bring very useful additional information that allows us to separately identify key aspects of the production of 'childcare services'.

We use longitudinal data from the PSID on consumption of non-childcare goods (available in a much richer form since 1999), and husband and wife separate data on hourly wages, hours of work, and time spent with children to estimate structural parameters measuring the sensitivity of consumption and time allocation choices to transitory and permanent shocks. These structural parameters include the Frisch and Marshallian elasticities, as well as partial insurance parameters ("savings") and human capital shares. They 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 childcare allows to decompose overall Frisch responses into two components, one reflecting the potential degree of complementarity between husband's and wife's leisure ("companionship" or "love") and another reflecting the degree of substitutability of their childcare time in the production of childcare services. We obtain this clean decomposition by assuming that utility is weakly separable over two broad sub-utility aggregates, one comprising consumption of non-childcare related goods and the leisure times of the two spouses, and the other defined over childcare services (which we assume to be produced using childcare related goods and childcare time of the two parents).

We find that this decomposition is important. While overall cross-Frisch responses are small, the responses on the two sub-utility margins are (in absolute value) larger. In particular, female hours appear to respond

210), and later, that " complementarity between men and women is the major source of the gain from marriage" (p. 211).

[2] Sandberg and Hofferth (2001, 2005) use the CDS to estimate trends in trends in childcare time.

little to an increase in the male's temporary 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 decrease her labor supply and increase childcare 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 childcare 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 childcare increases.

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 to shocks - in particular, to 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 prices of childcare in explaining changes in female participation rates over cohorts. There are also papers that access time use data to examine the precise decomposition of available time between children and other activities as in Aguiar and Hurst (2007, 2013), and Ramey and Ramey (2010).[3] In the second line of work, there are papers that emphasize the insurance role of labor supply, as in Hyslop (2001), Attanasio, Low and Sanchez-Marcos (2005), and our own recent study using the PSID (Blundell, Pistaferri and Saporta-Eksten, 2014).[4]

A key contribution of our paper is the extension of the theoretical framework used in previous literature. Building on the family labor supply framework of Blundell et. al. (2014), we demonstrate how to incorporate multiple goods and time use activities in a dynamic life-cycle model with uncertainty about underlying wages. We then show how to apply approximations of the first order conditions and the intertemporal budget constraint of the problem to derive analytical expressions for the sensitivity of consumption and childcare expenditure, as well as each earner's hours of work and childcare time, to wage shocks. These derivations are not only useful in empirical implementation, but they also facilitate the analytical decomposition of labor supply responses to their underlying sources, namely leisure and childcare time.

A second contribution of the paper is in deriving an identification strategy that combines longitudinal data on consumption, wages, earnings and assets with cross-sectional data on time use. This allows us to fully identify the parameters of the model. Implementing this strategy, we estimate the model and provide new evidence on the decomposition of labor supply elasticities to their leisure and childcare time sources, especially highlighting the tension between spouses' leisure complementarity and childcare time substitutability.

The paper proceeds as follows. Section 2 describes the life cycle problem of a household that in each period has to choose spending on childcare related and non-childcare related goods, and how to allocate

[3] Ramey and Ramey (2010), as well as Guryan, Hurst, and Kearney (2008) also study skill heterogeneity in childcare time.

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

spouses' times to work, childcare, and leisure. Section 3 describes the PSID/CDS data, while Section 4 discusses various empirical issues. The results are in Section 5. We further discuss the relation between childcare and wages, as well as the implications of the results to insurance in Section 6. Section 7 concludes.

2 A Life Cycle Model with childcare 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-childcare related goods, and from child services. The latter are produced using as inputs the hours that husband and wife devote to their children and the amount spent on childcare-related goods (such as formal baby-sitting, day care tuition, books, etc.). Hence, the household's problem is to choose consumption and the allocation of the two members' total hours to three activities: work, childcare time, and leisure.[5]

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's goal is to show that the response of family labor supply to shocks is rather complex, as it depends on the spouses' Frisch elasticities for leisure, childcare time, the degree of complementarity of their leisure times in utility, the degree of substitutability of childcare time inputs in the production of childcare, 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).

2.1 Wage Process

We follow Blundell et al. (2014) and assume that each earner's wage process contains both a permanent and a transitory component. The permanent component evolves as a unit root process, while the transitory component is serially uncorrelated. The distinction between transitory shocks and permanent shocks is important from an identification point of view, as we will interpret transitory shocks as having negligible or no wealth effects. Hence, the response of hours to transitory wage shocks will identify Frisch (or λ-constant) elasticities, while the response to permanent wage shocks will identify Marshallian (or uncompensated) elasticities.

Suppose that the log of real wage of earner $j = \{1, 2\}$ in household i at age t can be written as:

$$\log W_{i,j,t} = x'_{i,j,t}\beta_W^j + F_{i,j,t} + u_{i,j,t} \quad (1)$$

$$F_{i,j,t} = F_{i,j,t-1} + v_{i,j,t} \quad (2)$$

where $x_{i,j,t}$ are observed characteristics affecting wages and known to the household. $u_{i,j,t}$ and $v_{i,j,t}$ are transitory shocks and permanent shocks, respectively. We assume that transitory and permanent innova-

[5] We explicitly model only childcare time as non-leisure non-work use of time for the two earners in the household. This is of course a simplifying assumption, but consistent with our goal to focus and highlight in this paper the interaction between kids and labor supply.

tions are uncorrelated within person. However, we assume that they may be potentially correlated across spouses.[6] The structure of markets is such that shocks are not formally insurable, household have no advance information about them, and they are observed (separately) at time t.[7] We let the variances and covariances of the shocks vary by age. Since age-specific cells are quite small given the size of our data set, however, we restrict the age-variation to stages of the life cycle (30-37, 38-42, 43-47, 48-52, 53-57).

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

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

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

2.2 Household Problem

Households solve the following problem over the life cycle:

$$\begin{aligned}
& \max \mathbb{E}_t \sum_{s=0}^{T-t} u_{t+s}\left(\frac{C_{t+s}}{a\left(z_{t+s}\right)}, L_{1,t+s}, L_{2,t+s}, g\left(T_{1,t}, T_{2,t}, K_{,t}; z_{t+s}\right); x_{t+s}\right) \\
\text{s.t. } A_{t+1} &= (1+r)\left(A_t + \mathcal{T}\left(H_{1,t}W_{1,t} + H_{2,t}W_{i,2,t}\right) - C_t - P_{k,t}K_t\right) \\
L_{1,t} + H_{1,t} + T_{1,t} &= \bar{L} \\
L_{2,t} + H_{2,t} + T_{2,t} &= \bar{L} \\
K_t &\geq 0
\end{aligned}$$

where C represents consumption of non-childcare goods (from now on, just "consumption" for simplicity), L_j is the leisure time of earner j ($j = 1, 2$), P_k the price of childcare goods K, and the childcare consumption sub-utility $g(.)$ (from now on, just "childcare" for simplicity) is a function of childcare goods, and the childcare hours T_1 and T_2 of the two spouses. H_j are hours worked by earner j, $\bar{L}$ is maximum time available for allocation, and A are assets that pay a non-stochastic interest rate denoted by r.

We introduce two sets of demographic conditioning variables. The first, z, contains the number of kids and their age composition. We model time use, 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 childcare), we let the effect of these variables on consumption to work through an (estimated) adult equivalence scale $a(z)$. The second set of conditioning variables, x, includes other demographics (e.g. age, education, and so on). $\mathcal{T}(.)$ is a tax function that maps before-tax into after-tax household earnings.

[6] This is potentially important given the empirical findings for the correlation of labor market outcomes of married couples. See for example Juhn and Potter (2007) and Hyslop (2001).

[7] 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.

We follow Blundell et al. (2014) and model joint taxation as:

$$\mathcal{T}\left(H_{i,1,t}W_{i,1,t} + H_{i,2,t}W_{i,2,t}\right) \approx \left(1 - \chi_{i,t}\right)\left(H_{i,1,t}W_{i,1,t} + H_{i,2,t}W_{i,2,t}\right)^{1-\mu_{i,t}}. \tag{4}$$

Each earner allocates the total time endowment to leisure, hours of work (H), and time with children. While this formulation is fairly general, it already imposes the restriction that childcare spending (K) as well as time spent with kids (T_1 and T_2) interact with other goods only through $g\left(.\right)$. This is an identifying assumption similar to the one discussed in Browning, Chiappori and Weiss (2014, ch. 3). The non-negativity constraint on K implicitly imposes that market childcare can be substituted completely by parents' time spent with the kids.[8] This formulation admits corner solutions in childcare expenditure, which is an event frequently observed in the data.

2.3 The Dynamics of Time use, Hours and Goods

2.3.1 Restrictions on preference

For our empirical application, we add some additional restrictions on preferences such that the household's objective function can be rewritten as:[9]

$$\max\mathbb{E}_t \sum_{s=0}^{T-t}\left(\frac{1}{1+\delta}\right)^s \left[\left(1-\gamma\left(z_{t+s}\right)\right)\Upsilon\left(C_{t+s}, \bar{L} - H_{1,t+s} - T_{1,t+s}, \bar{L} - H_{2,t+s} - T_{2,t+s}\right) + \gamma\left(z_{t+s}\right) g\left(T_{1,t+s}, T_{2,t+s}, K_{t+s}\right)\right] \tag{5}$$

where $\Upsilon_m > 0$, $\Upsilon_{mm} < 0$ for $m \in \{C, L_1, L_2\}$, $\Upsilon\left(0, L_1, L_2\right) \to -\infty$, $g_m > 0$, $g_{mm} < 0$ for $m \in \{T_1, T_2, K\}$, and δ is a discount factor. This formulation implies that we have two utility sub-aggregates: (1) consumption, husband leisure, and wife leisure; and (2) husband's childcare time, wife's childcare time, and childcare goods. Note that in this specification, the weight $\gamma\left(.\right)$ is a direct function of the demographic conditioning variables in z. For example, a family with young kids might have a higher $\gamma\left(.\right)$ (i.e., weight the utility associated to childcare more) than a family with older kids. We show below how γ affects labor supply responses even conditioning on the strong separability assumption.

It is worth stressing that we assume separability between sub-utility aggregates, not between the goods themselves. For example, this formulation implies that preferences over non-children consumption and the leisure time of husband and wife are (potentially) non-separable. On the other hand, we assume that spending on non-childcare related goods is separable from spending on childcare related goods.

Our goal is to decompose the response of time allocation (i.e., how household members vary their labor supply, leisure time, and time spent on childcare) to wage shocks. To do that we need to link the growth rates of consumption, childcare goods, hours of work, and childcare time to the wage shocks experienced

[8] Alternatively, it can be interpreted as imposing that a minimum of non-parent childcare is available for free (for example childcare help provided by unmodeled family members or neighbors).

[9] To avoid cluttered notation we omit the normalization of C by the equivalence scale $a\left(z\right)$, as well as all conditioning variables x from all equations.

by the household members. We achieve this in two steps. First, we use a Taylor approximation of the first order conditions of the problem. This yields expressions for the growth rate of consumption and of childcare, the growth rate of hours of work, and the growth rate of childcare time in terms of changes in wages and an additional expectation error term (the innovation in the marginal utility of wealth).[10] This is a standard log-linearization approach, yielding the matrix of Frisch responses of goods to prices. Second, we take a log-linearization of the intertemporal budget constraint. This allows us to map the (unobservable) expectation errors into wage shocks (the only sources of uncertainty of the model).[11]

2.3.2 The Frisch Responses

In Appendix 1 we show that log-linear approximation of the Euler equations yields the following expression:

$$\begin{pmatrix} \Delta c_{\tau+1} \\ \Delta l_{1,\tau+1} \\ \Delta l_{2,\tau+1} \\ \Delta t_{1,\tau+1} \\ \Delta t_{2,\tau+1} \\ \Delta k_{\tau+1} \end{pmatrix} \cong \mathcal{F}^L \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & -\mu \\ 1 & 0 & 1 & 0 & -\mu \\ 1 & 1 & 0 & 0 & -\mu \\ 1 & 0 & 1 & 0 & -\mu \\ 1 & 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \Delta \ln \lambda_{A,\tau+1} \\ \Delta \ln W_{1,\tau+1} \\ \Delta \ln W_{2,\tau+1} \\ \Delta \ln P_{k,\tau+1} \\ \Delta \ln Y_{\tau+1} \end{pmatrix} \tag{6}$$

where the $\mathcal{F}^L$ matrix is:

$$\mathcal{F}^L = \begin{pmatrix} -\eta_{c,p} & \eta_{c,w_1} & \eta_{c,w_2} & 0 & 0 & 0 \\ \eta_{l_1,p} & \eta_{l_1,w_1} & \eta_{l_1,w_2} & 0 & 0 & 0 \\ \eta_{l_2,p} & \eta_{l_2,w_1} & \eta_{l_2,w_2} & 0 & 0 & 0 \\ 0 & 0 & 0 & \varpi_{t_1,w_1} & \varpi_{t_1,w_2} & \varpi_{t_1,p_k} \\ 0 & 0 & 0 & \varpi_{t_2,w_1} & \varpi_{t_2,w_2} & \varpi_{t_2,p_k} \\ 0 & 0 & 0 & \varpi_{k,w_1} & \varpi_{k,w_2} & \varpi_{k,p_k} \end{pmatrix}$$

Note that the zeroes in this matrix come directly from the separability assumption between $\Upsilon(.)$ and $g(.)$. From now on lower-case letters indicate logged variables net of predictable taste shifters. Hence, $c_{i,t}$, for example, is log consumption net of predictable taste shifters. $Y_t = H_{1,t}W_{1,t} + H_{2,t}W_{i,2,t}$ is total household taxable income. Finally, $\lambda_{A,t}$ is the marginal utility of wealth (the Lagrange multiplier on the sequential budget constraint). The parameters $\eta_{q,m}$ represent the Frisch (or λ-constant) elasticities of variable q with respect to changes in the price m which are associated with the Υ sub-utility (p is the "price" of a unit of current consumption relative to future consumption).[12] These parameters capture the response of leisure

[10] In the case of weak separability, this approximation will include another term - the change in the marginal rate of substitution between Υ and g.

[11] Our general specification could give rise to corner solutions in childcare inputs, in which case the log-linearization is inappropriate. Our derivations in the next sections will assume internal solutions for all goods. We address the problem of zeros when we describe the data in Section 3, and when we describe our empirical implementation in Section 4.

[12] Note that in this notation leisure elasticities w.r.t to own price η_{l_1,w_1} and η_{l_2,w_2} are negative.

(or consumption) to small changes in wages, holding the marginal utility of wealth constant.[13] Finally, the parameters $\varpi_{q,m}$ are the Frisch elasticities associated with the g sub-utility, and therefore capture the elasticity of variable q (childcare time of earner j or childcare goods) to a change in price m in the childcare production function, again holding λ constant. Since the wage measures the opportunity cost of time, we expect $\varpi_{t_j,w_j} < 0$.

Our next step is to write (6) in terms of hours (which we observe) rather than leisure (which we do not). This can be done using the following relation between leisure, time use, and hours elasticities:

$$\eta_{h_i,w_j} = -\frac{L_{i,\tau}}{H_{i,\tau}}\eta_{l_i,w_j} - \frac{T_{i,\tau}}{H_{i,\tau}}\varpi_{t_i,w_j}$$

Finally, note that when $\gamma = 0$ and preferences are defined only over the sub-utility Υ, we can define $\eta^{\gamma_0}_{h_i,w_j} = -\frac{L_{i,\tau}}{H_{i,\tau}}\eta_{l_i,w_j}$.[14] This allows us to write the Frisch responses matrix $\mathcal{F}$ for the dynamics of hours rather than leisure:[15]

$$\mathcal{F} = \begin{pmatrix} -\eta^{\gamma_0}_{c,p} & \eta^{\gamma_0}_{c,w_1} & \eta^{\gamma_0}_{c,w_2} & 0 & 0 & 0 \\ \eta^{\gamma_0}_{h_1,p} & \eta^{\gamma_0}_{h_1,w_1} & \eta^{\gamma_0}_{h_1,w_2} & -\frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,w_1} & -\frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,w_2} & -\frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,p_k} \\ \eta^{\gamma_0}_{h_2,p} & \eta^{\gamma_0}_{h_2,w_1} & \eta^{\gamma_0}_{h_2,w_2} & -\frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_1} & -\frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_2} & -\frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,p_k} \\ 0 & 0 & 0 & \varpi_{t_1,w_1} & \varpi_{t_1,w_2} & \varpi_{t_1,p_k} \\ 0 & 0 & 0 & \varpi_{t_2,w_1} & \varpi_{t_2,w_2} & \varpi_{t_2,p_k} \\ 0 & 0 & 0 & \varpi_{k,w_1} & \varpi_{k,w_2} & \varpi_{k,p_k} \end{pmatrix} \tag{7}$$

To illustrate a specific case, consider how hours growth of the wife depends on shocks to the marginal utility of wealth and changes in prices (including wages) and income:

$$\begin{aligned} \Delta h_{2,\tau+1} \quad = \quad & \left(\eta^{\gamma_0}_{h_2,p} + \eta^{\gamma_0}_{h_2,w_1} + \eta^{\gamma_0}_{h_2,w_2} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_1} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_2} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,p_k}\right)\Delta \ln \lambda_{A,\tau+1} \\ & + \left(\eta^{\gamma_0}_{h_2,w_1} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_1}\right)\Delta \ln W_{1,\tau+1} + \left(\eta^{\gamma_0}_{h_2,w_2} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_2}\right)\Delta \ln W_{2,\tau+1} \\ & -\frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,p_k}\Delta \ln P_{k,\tau+1} \\ & -\mu\left(\eta^{\gamma_0}_{h_2,w_1} + \eta^{\gamma_0}_{h_2,w_2} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_1} - \frac{T_{2,\tau}}{H_{2,t}}\varpi_{t_2,w_2}\right)\Delta \ln Y_{\tau+1} \end{aligned}$$

[13] Note that these are after-tax Frisch elasticities. Writting $\Delta Y_{\tau+1}$ as a function of the wages and hours and solving for the dynamics of all goods in terms of wages and shocks to the marginal utility of wealth λ would deliver the before-tax Frisch elasticities. We show this in appendix 1. For more details on before- vs. after-tax Frisch elasticities, see MaCurdy 1983, and Blundel et al. 2014.

[14] This formulation is attractive as we observe hours and not leisure. Moreover, it is usually the elasticities of hours with respect to the wage that researchers care about from a policy point of view. However, to interpret the hours elasticities $\eta^{\gamma_0}_{h_j,w_j}$, we are requires to assume that even though the model is written in term of leisure preferences, the hours elasticities $\eta^{\gamma_0}_{h_j,w_j}$ are (at least approximately) constant.

[15] It is intuitive to use the upper-right block of the matrix to capture the part of the Frisch response that is related to childcare time, highlighting the two components of the hours response. Note however that because of the linear dependence in rows 2-3 and 4-5 in the middle matrix in equation (6), we could also write (7) keeping the 0s in the 4th and 5th column of the upper-right block, simply including the time use response in the upper-left block. For example the 2,2 element of the upper-left block would have become $\eta^{\gamma_0}_{h_1,w_1} - \frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,w_1}$.

Consider the effect of a temporary wage change faced by the wife. Since the wage change is temporary, $\Delta \ln \lambda_{A,\tau+1} = 0$. The effect of the shock on the wife's labor supply is captured by the term $\left(\eta^{\gamma_0}_{h_2,w_2} - \frac{T_{2,\tau}}{H_{2,t}} \varpi_{t_2,w_2}\right)$, and hence can be decomposed into two distinct forces. First, there is a traditional intertemporal substitution effect: a temporary positive deviation from the normal wage induces worker to reallocate their work time 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 Frisch elasticity $\eta^{\gamma_0}_{h_2,w_2}$. 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 childcare hours that is partly reallocated to work), an effect whose strength is measured by ϖ_{t_2,w_2}. This second effect is weighted by $\frac{T_2}{H_2}$ because it depends on initial conditions (i.e., the effect is absent for earners who devote no time to children to start with).

The dashed black line in the upper panel of Figure 1 demonstrates how the overall wife's labor supply Frisch response changes as $\frac{T_{2,\tau}}{H_{2,t}}$ is changing. For illustration purposes, the Figure holds $\eta^{\gamma_0}_{h_2,w_2}$ and ϖ_{t_2,w_2} at their point estimates from Table 3. At very low levels of $\frac{T_{2,\tau}}{H_{2,t}}$, the total Frisch response η_{h_2,w_2} is below 0.5, while it is about 2 as $\frac{T_{2,\tau}}{H_{2,t}}$ approaches 1. The bottom right panel shows how η_{h_2,w_2} is decreasing in ϖ_{t_2,w_2}, this time holding $\frac{T_{2,\tau}}{H_{2,t}}$ fixed at its sample mean value.

A somewhat more interesting effect is how the wife responds to temporary changes in the husband's wage, which is measured by the term $\left(\eta^{\gamma_0}_{h_2,w_1} - \frac{T_{2,\tau}}{H_{2,t}} \varpi_{t_2,w_1}\right)$. The temporary increase in the husband's wage induces him to work longer hours. If husband and wife enjoy leisure together, $\eta^{\gamma_0}_{h_2,w_1} > 0$ and, in the absence of childcare time ($T_{2,\tau} = 0$), we should see an increase in the wife's hours as well. When husband and wife hours can be used to produce childcare, however, the effect is less obvious. If their childcare hours are substitutes in production, an increase in hours worked by the husband implies he can now allocate fewer hours to childcare production, and hence the wife will allocate more hours to childcare production (and fewer hours to market time, i.e., work less). This implies $\varpi_{t_2,w_1} > 0$. The magnitude of the effect will depend on the degree of input substitutability in childcare production. It follows that the two forces may counteract each other. The red solid line in Figure 1 illustrates the various forces at work holding $\eta^{\gamma_0}_{h_2,w_1}$ at its point estimate level from Table 3. When either $\frac{T_{2,\tau}}{H_{2,t}}$ or ϖ_{t_2,w_1} are close to zero, the complementarity of leisures of the two earners is the dominating effect. On the other hand, when the relative time spent with kids or the cross elasticity of childcare time increase, the substitutability of childcare time takes over and the overall effect turns negative.

2.3.3 The Marshallian Responses

So far we have shown how to write the dynamics of consumption, childcare, hours of work and childcare time as functions of the wage shocks and the changes in marginal utility of wealth. Without taxes, the Frisch responses given in $\mathcal{F}$ (7) coincide with the goods responses to transitory price shocks. Permanent shocks, however, change the marginal utility of wealth λ as well. The overall response to permanent shock to prices is captured by the Marshallian elasticities, which incorporate both wealth and substitution effects. To estimate these elasticities, we need to write changes in marginal utility of wealth λ as a function of permanent shocks to prices. Before we proceed, we need to take a stand on the stochastic process for P_k. While long run changes in the cost of childcare have shown to have an effect on overall labor supply (see Blau and Robins, 1988; Attanasio, Low and Sanchez-Marcos, 2008), idiosyncratic changes in the cost of childcare are both less likely, and are unobserved in our data set. For these reasons, we choose to index the price of childcare to the price of overall consumption and hence assume $P_k = 1$. This eliminates ΔP_k from our system of dynamic equations, and hence we are left only with the responses of goods to the idiosyncratic changes in wages.

We proceed in two steps. First, we decompose the growth of the marginal utility of wealth $\Delta \ln \lambda_{A,t+\tau}$ into two components. The first component, ω_t, is a function of the interest rate r, the discount factor δ, and the variance of the change of marginal utility. This component captures the effect of aggregate variables on the consumption slope. Assuming that the only source of uncertainty in our model is the idiosyncratic wage shocks, ω_t is fixed in the cross-section. The second component, $\varepsilon_{i,t}$, captures the revisions in the growth of the marginal utility of wealth. Second, to map innovations in the marginal utility of wealth into innovations in the wage process faced by the two earners, we log-linearize the intertemporal budget constraint:

$$E_t \sum_{s=0}^{T-t} \sum_{s=0}^{T-t} \frac{C_{t+s} + K_{t+s}}{(1+r)^s} = A_t + E_t \sum_{s=0}^{T-t} \frac{\mathcal{T}\left(W_{i,1,t+s}H_{i,1,t+s} + W_{i,2,t+s}H_{i,2,t+s}\right)}{(1+r)^s}$$

and take the difference in expectations between period t and $t-1$ to obtain equations that link consumption and hours growth of the two earners to the wage shocks they face (see Appendix 1 for the exact derivation).

$$\begin{pmatrix} \Delta c_{i,\tau} \\ \Delta h_{i,1,\tau} \\ \Delta h_{i,2,\tau} \\ \Delta t_{i,1,\tau} \\ \Delta t_{i,2,\tau} \\ \Delta k_{i,2,\tau} \end{pmatrix} \simeq \begin{pmatrix} \kappa_{c,u_1} & \kappa_{c,u_2} & \kappa_{c,v_1} & \kappa_{c,v_2} \\ \kappa_{h_1,u_1} & \kappa_{h_1,u_2} & \kappa_{h_1,v_1} & \kappa_{h_1,v_2} \\ \kappa_{h_2,u_1} & \kappa_{h_2,u_2} & \kappa_{h_2,v_1} & \kappa_{h_2,v_2} \\ \kappa_{t_1,u_1} & \kappa_{t_1,u_2} & \kappa_{t_1,v_1} & \kappa_{t_1,v_2} \\ \kappa_{t_2,u_1} & \kappa_{t_2,u_2} & \kappa_{t_2,v_1} & \kappa_{t_2,v_2} \\ \kappa_{k,u_1} & \kappa_{k,u_2} & \kappa_{k,v_1} & \kappa_{k,v_2} \end{pmatrix} \begin{pmatrix} \Delta u_{i,1,t} \\ \Delta u_{i,2,t} \\ v_{i,1,t} \\ v_{i,2,t} \end{pmatrix} \tag{8}$$

Note that in general, the loading factors κ vary over time and across households (i.e., we should write $\kappa^{i,t}_{c,v_1}$, etc.). To avoid unnecessary cluttering we will leave this individual and time-dependence implicit. The response of a good/time m to a permanent wage shock faced by earner j is a general function:

$$\kappa_{m,v_j} = \kappa^{m,v_j}\left(\pi_{i,t}, s_{i,t}, \boldsymbol{\eta}, \boldsymbol{\varpi}, \mu_{i,t}, \chi_{i,t}, \frac{T_j}{H_j}\right)$$

where $\pi_{i,t} \approx \frac{\text{Assets}_{i,t}}{\text{Assets}_{i,t}+\text{Human Wealth}_{i,t}}$ is a "partial insurance" coefficient (the higher $\pi_{i,t}$ the lower the sensitivity of consumption to shocks), $s_{i,t} \approx \frac{\text{Human Wealth}_{i,1,t}}{\text{Human Wealth}_{i,t}}$ is the share of earner 1's human wealth over family human wealth, $\boldsymbol{\eta}$ and $\boldsymbol{\varpi}$ are the vectors of all Frisch elasticities, and $\mu_{i,t}$ and $\chi_{i,t}$ are the tax parameters defined above.[16]

Generally, Marshallian responses to (permanent) wage shocks are more involved than Frisch responses due to 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 childcare services, and added worker effects, i.e., the fact that women have an incentive to replace some of the (permanently) lost earnings of the husband by working more. In the empirical application we demonstrate the different roles of η and ϖ both in shaping the Marshallian labor supply responses, and in shaping the Marshallian consumption response, which is closely related to insurance.

2.3.4 Relating Frisch and Marshallian Responses to Demographic Characteristics

As shown in Figure 1, in our model the ratio of childcare time to work time of each earner (T_j/H_j) is a key statistic as it reflects the weight on the childcare time elasticities ($\boldsymbol{\varpi}$) in the total Frisch labor supply response of each earner. We show now that with the strong separability assumption, this ratio also fully captures the effect of demographic shifters z on labor supply.

Under the strong separability assumption as in (5), $\gamma(z)$ does not show up in the dynamic equations (8). $\gamma(z)$ however, plays an implicit role in shaping $\frac{T_j}{H_j}$. To see that, consider the first order conditions for childcare time and leisure of earner j:

$$\begin{aligned} T_{j,\tau} \quad &: \quad g_{T_j}(\tau) = \frac{1}{\gamma} W_{j,t} \lambda_{A,\tau} \\ L_{j,\tau} \quad &: \quad \Upsilon_{L_j}(\tau) = \frac{1}{(1-\gamma)} W_{j,t} \lambda_{A,\tau} \end{aligned}$$

Given our assumption that $g_{T_j T_j}(\tau) < 0$, $g_{T_j}(\tau)$ is monotonically decreasing in T_j, which implies that for a given level of marginal utility of wealth $\lambda_{A,\tau}$ and wage, T_j is monotonically increasing in γ. By the same argument, leisure L_j is monotonically decreasing in γ. While it is unclear what the total effect on hours is (because $H_j = \bar{L} - L_j - T_j$), it is clear that when z shifts γ, this effect will be reflected in the ratios $\frac{T_j}{H_j}$. To account for this issue, in the estimation we explicitly allow $\frac{T_j}{H_j}$ to depend on the demographic shifters z, including the number of kids and their ages. In the empirical application we explicitly make $\frac{T_j}{H_j}$ a function of the age of the youngest kid in the family. We use the median of $\frac{T_j}{H_j}$ for three separate groups - families with kids 6 or younger, 7 to 11 and 12 to 13.

After explicitly accounting for kids presence and ages through $\frac{T_j}{H_j}$, the dynamic equations in (8) can be estimated using panel data on assets, wages, consumption, childcare, and each earner's hours of work and childcare time. In what follows we describe our data, as well as our identification strategy, where we demonstrate that the model can be estimated even without observing panel data on time use, although it does require panel data on wages.

[16] Human Wealth$_{i,t}$ (Human Wealth$_{i,j,t}$) is the expected discounted flow of lifetime earnings of the household (earner j) at the beginning of period t. The exact expression for Human Wealth is given in Appendix 1.

[17] Blundell et. al. (2014) provide an overview of the forces present in the family labor supply model.

3 Data

For our empirical analysis we use two datasets collected as part of the Panel Study of Income Dynamics (PSID). The first data set is the PSID main interview sample. The second is the Child Development Supplement (CDS), and in particular its diary segment, which was collected alongside the 1997, 2001 and 2007 main interview PSID waves.

3.1 The PSID Interview Data

The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Thereafter, 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. Starting in 1999, in addition to income data and demographics, the PSID collects data about detailed assets holdings and consumption expenditures. Since we need both consumption and assets data, we focus on interview data from the 1999-2009 sample period. These data therefore provide us with longitudinal information on all the variables required for estimation other than the time each earner spends on childcare.

We follow Blundell et al. (2014) in constructing a sample that is suitable for studying family labor supply decisions (see their paper for more details). In particular, we focus on non-SEO households with participating and married male household heads aged between 30 and 57. Whenever there is a change in family composition we drop the year of the change and treat the household unit as a new family starting with the observation following the change. In addition, because our main interest is in the relation between labor supply elasticities and time spent with children, our baseline sample consists of families where the youngest child is 13 or less at the time of the interview.

Finally, we drop observations with missing values for state, education, race, labor earnings, hours, total consumption and total assets, and observations with hourly wages that are lower than half the minimum wage in the state where the household resides. We drop observations with extreme negative values of assets, as well as observations with total transfers (calculated as explained in Section 4.2) more than twice the size of total household earnings. To mitigate measurement error issues, we also drop observations for which consumption, wages or earnings of one of the earners show extreme "jumps", most likely due to extreme measurement error. A "jump" is defined as an extremely positive (negative) change from $t-2$ to t, followed by an extreme negative (positive) change from t to $t+2$. Formally, for each variable (say x), we construct the biennial log difference $\Delta^2 \log(x_t)$, and drop the relevant variables for observation in the bottom 0.25 percent of the product $\Delta^2 \log(x_t) \Delta^2 \log(x_{t-2})$.

3.2 The CDS Data

An important source of information required for estimation of the model is the time spent by each earner on childcare. For this purpose, we use the CDS, a relatively unexplored source of information on children time use. For our sample period, the two relevant CDS surveys come from the 2001 and 2007 waves. The CDS started collecting detailed information on kids 12 year old or younger as a supplement to the 1997 PSID main interview. in 1998, the CDS interviewed 2,394 families, with information on 3,563 children. In 2002-2003, CDS re-interviewed the 1998 CDS families that were still part of the PSID panel (2,021 families and 2,907 children). CDS collects data on a broad range of children outcomes, such as health, intellectual achievement and social relationships. Most importantly for our purposes, the CDS includes a time diary that can help us in recovering how much time each caregiver (parent) spends with the child.

The time diary collects all activities performed by the child over a 24-hours period beginning at midnight for one weekday and one weekend day (randomly chosen). Besides a detailed log for the start and the end time for each activity, the diary provides information about who was doing the activity with the child, and who (else) was present but not directly involved in the activity. This data structure allows us to recover how much time each parent spent with the surveyed child over the 24-hours period. In our empirical analysis we assume that a parent's childcare time includes both the activities he actively participated in and those he was present at. Given that we are interested in the substitution between time spent with the children and hours of work, we only use the weekday diaries (Monday to Friday), and focus on activities between 6am and 7pm.

There are a few issues to take into account when using the time diary data in our analysis. First, while the time diaries we use are supplements to the 2001 and 2007 PSID waves, they are in fact conducted in late 2002/early 2003, and in late 2007/early 2008, respectively. This implies that in terms of timing they need to be matched to the 2003 and 2009 waves interview data (since we assume that the income and consumption data collected in a given wave refer to the previous calendar year). For more details about the timing of the data, see Appendix Figure 1.

Second, the CDS is a child-, not a parent-based time use survey. Moreover, it surveys only one or two children per household in 1998, and re-interviews the same children in 2002 and 2008. It follows that the CDS gives a partial picture of time spent with kids for parents with kids not surveyed, including for parents with more than two kids, or for those who had children between 1998 and 2008. We address these issues in two ways. First, as we do with the other variables, we first regress the time use data on various control variables, and then use the residuals in estimation. For example the fact that the average age of the surveyed children is increasing across the three CDS waves, potentially affecting the level of time spent on childcare, will be absorbed by year effects and demographic controls. Second, even if the level is biased due to not observing all children in the family, our identification strategy relies heavily on logs, therefore it is less sensitive to overall bias in the *level* of time spent on childcare. Nevertheless, the level could still affect the weight $\frac{T_j}{H_j}$. In one of our robustness checks, we show estimates from specifications that are based on alternative time use measures implying different weights $\frac{T_j}{H_j}$.

3.3 Descriptive Statistics

The construction of the consumption data is similar to Blundell et al. (2014) . We use only consumption categories which were consistently collected starting in 1999. These include food (at home and away), health expenditures, utilities, gasoline, car maintenance, transportation, education and childcare. We treat sub-categories with missing values as zeros. Descriptive statistics on consumption, its breakdown for durables and nondurables and for childcare expenditure (nominal averages over the panel waves) are reported in the upper part of Table 1.[18] Data on household's assets holdings is required for the construction of $\pi_{i,t}$, the share of assets out of total wealth. We use data on detailed holdings of cash, bonds, stocks, business, pensions, car value, housing, other real estate holdings as well as household debt to construct a measure of household net wealth for all waves from 1999 to 2009.

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. In the lower part of Table 1 we provide summary statistics on labor supply and earnings for the two earners. We compare our baseline sample (households where the youngest child is 13 or less) with less restrictive samples, where we do not condition on children's ages (Column 2) and where we do not condition on working males (Column 3). It is worth noting that even though female participation rates are slightly lower for families with young children, they are still fairly high (around 75% vs. 80%). On average, females earn about half of what males earn, partly reflecting lower hours of work (conditional on working), and partly reflecting other factors, both explained and unexplained.

The fact that only 75% of women in our sample work raises standard sample selectivity issues. However, Blundell, Pistaferri and Saporta-Eksten (2014) add Heckman-style sample selection correction to their GMM moments and show that the procedure induces only small changes in their estimates, most likely due to the high participation rate (80%) in their female sample. Given this evidence, obtained with a sample that is very close to the one we use here, in the rest of the paper we focus on working women neglecting the sample selectivity correction.[19]

Table 2 summarizes data from the CDS regarding childcare time spend by each parent. Childcare time is annualized as daily hours times 260 days. First, note that only 5% of the wives in our sample report zero childcare time. The proportion of zero childcare time is larger for husbands (22%). Partly, this large difference is due to the diary nature of the data. Suppose that the day the child 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* childcare time based on observing zero childcare hours on the diary day. Because our moments are in logs, zero childcare hours imply a missing observation.

[18] Blundell et. al. (2014) also show that the PSID implied consumption aggregates match well with NIPA estimates.

[19] Of course, it is possible that the correction may matter more in our sample due to the fact that women with young children have lower participation rates than the average women married to a continuously employed men.

In the baseline specification we assume that missing childcare time for husbands refer to men who generally spend less time with their kids. We therefore replace the missing value with a low value (the 5th percentile of the childcare time distribution for husbands who report positive childcare hours). In the robustness section, we show the results assuming instead that the husband's childcare time is missing from the diary at random, implying that no correction is needed.

Consistent with estimates from Guryan, Hurst, and Kearney (2008) and Ramey and Ramey (2010), in our sample wife's childcare annual hours are about twice as much as the husband's hours (conditional on observing a value for the husband).[20] Similarly, the ratio of childcare hours to working hours is much higher for the wife. Interestingly, however, the implied *cost* of childcare time $T_j W_j$ is much closer for the two earners, due to the higher wages earned by the husband, compensating for the lower time he spends with the children. Finally, the cost of market childcare (K) is tiny compared to the implied cost of childcare time. This observation will guide some of our choices in the empirical application below.

4 Empirical Issues

Our goal is to estimate the parameters of the wage process (1) and (2) and the preference parameters in (8). The wage process parameters include the variances and covariances of the transitory and permanent wage shocks $(\sigma_{u_j}, \sigma_{v_j}, j \in \{1, 2\}, \sigma_{u_i u_j}, \sigma_{v_i v_j})$. The preference parameters include the vectors of Frisch elasticities $\boldsymbol{\eta}$ and $\boldsymbol{\varpi}$.

One important problem is that 60% of households in our sample (despite conditioning on having children 13 or younger) report zero expenditure on childcare goods, implying that the log-linear approximation is problematic. In terms of total consumption expenditure, we calculate that the average childcare budget share is only about 3.5%. Moreover, as shown in the last row of Table 2, and as we discussed earlier, the cost of market childcare is negligible compared to the implicit cost of time spent in childcare. For these reasons, we will neglect childcare spending (K) in the structural part of estimation and drop the last row of (8) from the analysis. One justification is that as long as the magnitude of the Frisch elasticities of childcare goods are not too large, neglecting childcare spending in estimation is unlikely to generate misleading estimates of the other parameters. Indeed, we run reduced form regressions aimed at recovering estimates of the elasticities of childcare expenditure in Section 6.1 and show that their magnitudes are rather small.

In the rest of this section we discuss some important estimation issues. First, how measurement error is introduced in the econometric model. Second, how we recover the parameters of the tax function (μ and χ) as well as the human wealth share in constructing π and s. Finally, we discuss the formal identification of the parameter of the model.

[20] Our implied weekly hours for the mother (16) and father (just below 8) are somewhat higher compared to those reported in Table 1 of Guryan, Hurst, and Kearney (2008). This is for two reasons. First, our sample conditions on having at least one child 13 or younger in the household. Second, our measure includes all activities where the parent is present and not participating, therefore might include activities not considered as childcare in their paper.

4.1 Measurement error

Given that we rely on survey data, consumption, wages, earnings and childcare time are all presumably measured with error. In our context, there are three problems one need to confront when adding measurement errors. First, as discussed among others in Blundell et al. (2008), adding measurement errors to models that include a permanent/transitory decomposition (as in our wage process) creates an identification problem, since the distribution of the measurement error is indistinguishable from the distribution of the economically relevant transitory shock. Second, our wage measure is constructed as annual earnings divided by annual hours of work, and therefore the measurement errors of earnings and wages are correlated (the so-called "division bias"). Third, measurement errors are hard to distinguish from stochastic changes in preferences or shocks to higher moments of the distribution of wages in terms of effects on consumption or labor supply choices. We make no attempt to resolve this distinction, and hence identify an aggregate of these various forces, some statistical and some economic.

Ignoring the variance of measurement error in wages or earnings is problematic since it has a direct effect on the estimates of the structural parameters. We thus follow Meghir and Pistaferri (2004) and use findings from validation studies to set *a priori* the amount of wage variability that can be attributed to error. We use the estimates of Bound et al. (1994), who estimate the share of variance associated with measurement error using a validation study for the PSID (which is the data set we are using). We also use the first order autocovariance as an upper bound on the measurement error in consumption (see Blundell et al., 2014, for more technical details).

Finally, childcare time is also measured with error. We assume that this error is uncorrelated with the other sources of measurement error, which is not implausible given that the CDS diary is conducted separately from the PSID main interview. Since we use only terms involving the covariance of the *level* of childcare time (rather than its variance), these assumptions imply that measurement error to childcare time does not pose a serious identification threat.

4.2 Recovering Tax Parameters and Human Wealth Shares

It is important to ensure that the tax function that we use in our model fits well the progressivity of the tax system in the U.S., as well as the generosity of means-tested transfers programs (which tends to be higher for families with children). From equation (4), the mapping from before-tax to after-tax income requires estimation of the tax parameters $\chi_{i,t}$ and $\mu_{i,t}$. We estimate them by regressing after-tax household income on a constant and before-tax household earnings (as reported in the survey), allowing for the regression coefficients to change by year and household characteristics. In particular, we first compute after-tax income as:

$$Y_{i,t} = \sum_{j=1}^{2} W_{i,j,t}H_{i,j,t} - \tau\left(\sum_{j=1}^{2} W_{i,j,t}H_{i,j,t}, z_{i,t}\right) + EITC\left(\sum_{j=1}^{2} W_{i,j,t}H_{i,j,t}, z_{i,t}\right) + FS\left(\sum_{j=1}^{2} W_{i,j,t}H_{i,j,t}, z_{i,t}\right)$$

where $\tau(.)$, $EITC(.)$ and $FS(.)$ are functions that compute taxes and eligible amounts of EITC and Food Stamps benefits using program information for the various years (allowing benefits to vary by demographic, such as number and age of children, etc.). Since using $Y_{i,t}$ directly is infeasible in our log-linear approximation procedure, we approximate the relationship between after-tax income and before-tax household earnings using:

$$Y_{i,t} \approx \left(1-\chi_{i,t}\right)\left(\sum_{j=1}^{2} W_{i,j,t} H_{i,j,t}\right)^{\left(1-\mu_{i,t}\right)}$$

i.e., equation (4).[21] Overall, the approximation appears accurate. In fact, the R^2 of a regression of predicted on actual average tax rates is 0.90.

The calculations of $\pi_{i,t} \approx \frac{\text{Assets}_{i,t}}{\text{Assets}_{i,t}+\text{Human Wealth}_{i,t}}$ and $s_{i,1,t} = \frac{\text{Human Wealth}_{i,1,t}}{\text{Human Wealth}_{i,t}}$ require the knowledge of assets, which we take directly from the data, the expected after-tax human wealth at time t, and the share of human wealth held by each earner. We calculate after-tax human wealth as:

$$\text{Human Wealth}_{i,t} = Y_{i,t} + \frac{\mathbb{E}_t\left(Y_{i,t+1}\right)}{1+r} + ...$$

where $Y_{i,t}$ is total household after-tax income (assuming no changes in tax policy).[22]

Note that the measure of assets we use is defined "beginning-of-period" (i.e., before any consumption decisions are taken), so no endogeneity issues arise.[23] The major difficulty is to form estimates of expected future earnings. We start by applying our tax approximation to pooled household earnings for all years and ages. We then regress after-tax earnings on characteristics ($\varkappa^a$ below) that either do not change over time (such as race and education) or characteristics ($\varkappa^b$) that change in a perfectly forecastable way (such as a polynomial in age, and interaction of race and education with an age polynomial). That is, we regress:

$$Y_{i,t} = \varkappa_i^a \gamma_1 + \varkappa_{i,t}^b \gamma_2 + e_{i,t}$$

To obtain an estimate of expected earnings at $t+s$ given information at t (i.e., $\mathbb{E}_t\left(Y_{i,t+s}\right)$) we simply use $\widehat{Y}_{i,t+s} = \varkappa_i^a \widehat{\gamma}_1 + \varkappa_{it+s}^b \widehat{\gamma}_2$. We assume that agents are working until the age of 65 and that the discount rate is the same as the interest rate, and set the annual interest rate to 2%.

The same idea is applied to calculate expected human wealth for the each earner. To control for participation in the prediction of earnings, we use a probit specification with education, race, polynomial in

[21] To capture the very low and sometimes negative tax rates for low earnings households, we estimate the tax parameters separately for households that are eligible for Food Stamps and for those who are not. Furthermore, for the Food Stamps eligible group we allow the function to change by number of kids, and for the rest of the sample we allow the parameters to change by year in order to capture changes in the progressivity of taxes over time.

[22] The share of human wealth by earner j at time t is calculated similarly, but allowing also for some correlation between the share of earnings in a particular period and the share of each earner in this period. The exact calculation is reported in Appendix 1.

[23] In practice we use assets reported in the previous $(t-2)$ wave.

age and interactions to predict the probability of participation for the secondary earner at each age. The expected earnings for the wife at age $t + j$ are then the product of the predicted offered wages in period $t + j$ and the probability of being employed in that period.

This procedure allows us to (pre-)estimate $\pi_{i,t}$ using asset and human capital data. The average value of these estimates for our sample of families with children 13 or younger is $E(\pi_{i,t}) = 0.097$, but there is quite a large amount of heterogeneity across households (see Blundell et al., 2014). The life cycle path of $\pi_{i,t}$ conforms to expectations. The degree of self-insurance warranted by asset accumulation is negligible at the beginning of the life cycle (all permanent shocks pass more or less through consumption), but the combination of asset accumulation due to precautionary and life cycle motives and the decline of expected human capital due to the shortening of the time horizon imply an increase in $\pi_{i,t}$ as time goes by, and hence the household's ability to smooth permanent wage shocks also increases over time. (These results are all available on request from the authors.)

Our estimates of $s_{i,1,t}$ (the ratio of the husband's human wealth to total household human wealth) can be interpreted as the life cycle evolution of the distribution of earnings power within the household. On average, the husband commands about 70% of total household human wealth. His weight rises initially due to fertility choices made by his wife, and declines at the end of the life cycle due to early retirement choices coupled with age differences within the household.

4.3 Identification

There are four sets of parameters that we are interested in estimating: wage parameters, smoothing parameters, preference parameters, and measurement error variances. Blundell et al. (2014) show how to identify the parameters of the wage function and the measurement error parameters using exclusively wage data and outside information from validation studies (see also details in Section 4.1 above). In what follows, we demonstrate how the childcare production function elasticities can be recovered using cross-sectional data on parents' childcare time matched with panel data on wage changes. This reliance on cross-sectional (rather than longitudinal) moments is imposed on us by data limitations: the CDS data are collected only for two waves, 2001 and 2007, and are sufficiently spaced apart that any interesting dynamics may become hard to detect if long differences were used.

The intuition for why cross-sectional moments may still be useful for identification is that Frisch elasticities can be recovered from the response of goods (or time) to transitory shocks to prices. Using only cross-sectional data on childcare time, we rely on the observation that a negative change in wages from τ to $\tau + 1$ (as visible from (3)) is associated with either a positive transitory shock at τ, or a negative transitory/permanent shock at $\tau + 1$. Given that the *level* of time spent with the kids at τ is independent of the shock at $\tau + 1$, a *negative* correlation between $\Delta w_{i,\tau+1}$ and the level of $t_{j,\tau}$ implies a *positive* Frisch elasticity of j to w_i. Formally, for the case without taxation, the following moments can be used in estimation:[24]

$$\mathbb{E}\left(t_{i,\tau}\Delta w_{i,\tau+1}\right) = -\left(\varpi_{t_i,w_i}\sigma^2_{u_1} + \varpi_{t_i,w_i}\sigma_{u_iu_j}\right) \tag{9}$$

$$\mathbb{E}\left(t_{i,\tau}\Delta w_{j,\tau+2}\right) = -\left(\varpi_{t_i,w_i}\sigma_{u_iu_j} + \varpi_{t_i,w_j}\sigma^2_{u_j}\right). \tag{10}$$

The intuition for the identification argument is similar also with progressive taxation, with the only difference that the feedback effect from taxes to labor supply is taken explicitly into account in estimation. Given that the parameters of the wage process can be identified from longitudinal wage moments as in Blundell et al. (2014), the moments (9) and (10) can be used, alone, to pin down ϖ_{t_i,w_i} and ϖ_{t_i,w_i}.

The total Frisch responses of hours are identified using the response of hours to transitory shocks, and using moments that capture the mean reverting part of wages (similar in idea to 9 and 10). For example (again ignoring taxes for simplicity):[25]

$$\begin{aligned}
E\left(\Delta w_{i,\tau}\Delta h_{i,\tau+1}\right) &= -\left(\eta_{h_i,w_i}\sigma^2_{u_i} + \eta_{h_i,w_j}\sigma_{u_i,u_j}\right)\\
E\left(\Delta w_{j,\tau}\Delta h_{i,\tau+1}\right) &= -\left(\eta_{h_i,w_i}\sigma_{u_i,u_j} + \eta_{h_i,w_j}\sigma^2_{u_j}\right)\\
\eta_{h_i,w_i} &= \left(\eta^{\gamma_0}_{h_i,w_i} - \frac{T_i}{H_i}\varpi_{t_i,w_i}\right)\\
\eta_{h_i,w_j} &= \left(\eta^{\gamma_0}_{h_i,w_j} - \frac{T_i}{H_i}\varpi_{t_i,w_j}\right)
\end{aligned}$$

As is clear from these moments, once the parameters of the childcare time production function ($\boldsymbol{\varpi}$) are known, the use of information on the ratios $\frac{T_j}{H_j}$, allows to recover the parameters η^{γ_0}.

4.4 Empirical Strategy

The following steps summarize our empirical strategy:

1. Regress the log of $C_{i,t}, Y_{i,1,t}, W_{i,1,t}, Y_{i,2,t}$ and $W_{i,2,t}$ onto observable characteristics and construct the first-differenced residuals $\Delta c_{i,t}, \Delta y_{i,1,t}$, $\Delta w_{i,1,t}, \Delta y_{i,2,t}$, $\Delta w_{i,2,t}$. The observable characteristics in the wage equation include year, year of birth, education, race, state and large city dummies as well as education-year, race-year and large city-year interactions. For consumption and earnings we also add, income recipient other than head or wife in the household and whether the couple has children not residing in the household. Note that the wage and earnings regressions use only workers. To account for the fact that consumption level per capita is also different when kids are present, we add dummies for family composition in the consumption regressions. We do not do the same for hours of work, because our empirical strategy explicitly account for the effect of kids on the changes in levels of hours

[24] We omit the τ from the parameters in the RHS of these moments to reduce notation. The variances are all functions of age.

[25] While we demonstrate the identification with hours data, we use earnings data in the estimation. Unlike the moments 9 and 10, the moments that involve hours (or earnings) do potentially suffer from measurement error because wages are calculated as $\frac{earnings}{hours}$. We fully account for these issues in estimation, as described in Section 4.1.

and time use. Finally, we residualize T_1 and T_2 similarly to the way earnings are residualized, but keep the residuals in levels ($t_{1,t}$, $t_{2,t}$) rather than taking first difference as we do with earnings.

2. Estimate the wage variances and covariances using the second order moments of $\Delta w_{i,1,t}$ and $\Delta w_{i,2,t}$;

3. Use (4) to estimate the tax parameters $\chi_{i,t}$ and $\mu_{i,t}$ by regressions of log after-tax joint earnings on a constant term and log before-tax joint earnings allowing for the regression coefficients to change by year and household characteristics;

4. Estimate the smoothing parameters $\pi_{i,t}$ and $s_{i,1,t}$ using asset and (current and projected) earnings data;

5. Estimate the preference parameters using the second order moments for $\Delta y_{i,1,\tau}$, $\Delta y_{i,2,\tau}$, $\Delta c_{i,\tau}$, $t_{i,\tau}$ and $t_{2,\tau}$ conditioning on results (wage variances, covariances, and smoothing parameters) obtained in steps 2, 3 and 4. We impose symmetry of the Frisch substitution matrix for each of the sub-utilities (See Appendix 2 for details).

Steps 1 through 4 are conducted on the entire sample (without conditioning on the age of the kids), as our goal is to obtain consistent estimates of the variance parameters as well as human capital shares by age. The sample restriction is only applied at step 5 for the estimation of the preference parameters, where we want to focus on parents who are likely to spend (at least some of their) time on childcare.

5 Results

5.1 Main Results

Our estimates of the wage parameters by age are similar to the estimates in Blundell et al. (2014) who discuss them extensively and to which interested readers are referred to. The estimates are reported in Appendix Table 1.

Table 3 reports the results for the estimation of the parameters of the model, decomposing the hours response according to the $\mathcal{F}$ matrix as in equation (7). The first column reports the responses of consumption and hours of work implied by the sub-utility $\Upsilon(.)$. For consumption the response coincides with the Frisch elasticity. For hours of work, they can be interpret as a (negative) transformation of the leisure elasticities, or as the Frisch elasticities of hours when the weight on the childcare sub-utility is zero (i.e., $\gamma = 0$). The second column reports the Frisch elasticities of the childcare sub-utility ($g(.)$). These are interpretable as childcare time Frisch elasticities (i.e., the response of childcare time to a change in its price - the hourly wage - keeping the marginal utility of wealth constant). Column 3 reports the weight on the time use elasticity ϖ in the hours response to shocks. This is the median value of $\frac{T_1}{H_1}$ ($\frac{T_2}{H_2}$) for the husband (wife) responses.[26]

[26] As we explain in section 2.3.4, the median changes with demographic characteristics. These are calculated using the sub-sample that has both CDS diary and PSID interview data. The table reports the weighted average of these medians.

Finally, column 4 reports the *total* Frisch response (which accounts for the fact that, given that time has multiple uses, an increase in the price of time keeping wealth constant changes both leisure and childcare time, and hence hours as per the time budget constraint). For example, for the first row (documenting the response of H_1 to W_1) the total response is given by: $\eta^{\gamma_0}_{h_1,w_1} - \frac{T_{1,\tau}}{H_{1,t}} \varpi_{t_1,w_2}$. At the bottom of the table we report the average share of assets in total wealth in our sample (π), as well as the average human wealth of the primary earner (s). The share of human wealth π is directly related to the amount of consumption insurance obtained through savings, and is slightly lower than the 0.15 reported in Blundell et al. (2014) because our sample is composed of younger households (due to the sample restriction of looking at families with young kids).

Before discussing the decomposition of the labor supply response, it is worth noting that the estimated consumption Frisch elasticities are all within plausible ranges. The own price elasticity for consumption of $\eta_{c,p} = 0.64$. Similarly to the estimates in Blundell et al. (2014), we find slightly negative Frisch responses of consumption to wages (η_{c,w_1} and η_{c,w_2}). Note that these negative coefficients do not imply a negative response of consumption to an unexpected increase in wages as, in general, permanent shocks induce an increase in consumption through the wealth effect. We discuss these Marshallian effects further in the next section.

The estimates of the labor supply own total Frisch response, 0.50 for men and 1.43 for women with children, are in line with other estimates in the literature (see Keane, 2011, for a survey). Our main interest is in decomposing the effects of wage changes on hours in the part coming from the change in leisure and the part coming from the change in childcare time.

There are two interesting points worth noting. First, the decomposition of the own Frisch responses is different for the two earners. While for men about 13% of the hours response is associated with changes in time allocated to childcare, for women 75% of the hours response is associated with changes in childcare time. The first reason behind this difference is the husband's low ratio of childcare hours to working hours (on average $\frac{T_1}{H_1} = 0.14$) compared to the high ratio for the wife (on average $\frac{T_2}{H_2} = 0.61$). The second reason is the much larger estimated own price elasticity ($\varpi_{t_2,w_2} = -1.774$) that we find for women. We note however that this elasticity is imprecisely estimated.

Second, childcare time plays an important role in shaping the complementarity of hours of the two earners. The total response of the husband's hours to a change in the wife's wage is small and positive, and the response of the wife's hours to a change in the husband's wage rate is very close to zero.[27] Can this evidence be interpreted as suggesting that the two earners make leisure time choices independently, or with only small complementarities? Looking at the decomposition between the two channels we have

[27] While the estimate for the total Frisch response of H_2 to W_1 is insignificant, the reversal of signs compared to the total response H_1 to W_2 seems puzzling at first, given that we impose the Frisch symmetry. However, note that the Frisch symmetry can be imposed either on the total Frisch response or on the η^{γ_0} paraneters (i.e. $\eta^{\gamma_0}_{h_1,w_2}$ and $\eta^{\gamma_0}_{h_2,w_1}$). While in theory one implies the other, in practice we only impose the latter. Appendix 2 provides more details about the Frisch symmetry condition we impose.

discussed above reveals very clearly that this is a misleading conclusion. The component of the elasticity of husband's hours to changes in the wife's wage that is due to the leisure response is positive, much larger, and statistically significant ($\eta^{\gamma_0}_{h_1,w_2} = 0.39$), while the response of childcare time is positive and significant as well ($\varpi_{t_1,w_2} = 1.55$), implying that the second term is overall negative. The calculation for the wife's elasticity to the husband's change in wages follows a very similar pattern. These estimates imply that while the husband and wife like spending leisure time together, the substitutability of their childcare time in the childcare production function works to weaken or even overturn this effect.

To be more specific, consider a positive transitory shock to the husband's wage. Due to intertemporal substitution effect, this induces the husband to increase his labor supply. This increase triggers two responses from the wife's side. On the one hand, leisure complementarity implies that she would like to increase her labor supply as well (because he has less leisure time available). On the other hand, the husband can now spend less time with the kids, therefore the wife substitutes market time with more of her time spent with the children. For our sample of families with kids under 13, these two effects almost cancel each other up, making it look as if there is almost separability over the hours margin of the two earners.

5.2 Marshallian Elasticities

While Frisch elasticities give an important picture of the responses of consumption, time use and overall labor supply to transitory changes in wages, Marshallian elasticities that reflect the impact of a permanent change of wages (inducing both a substitution and a wealth effect) are also of key importance from a policy point of view. The upper panel of Table 4 summarizes the average Marshallian elasticities in our sample, which correspond to the average κ_{m,v_j} from equation (8).[28] Column 1 shows the elasticity with respect to a shock faced by the male, and column 2 reports the Marshallian elasticities associated to a shock to the wife's wage.

The average labor supply Marshallian elasticity for the husband is close to zero (0.09). It is much larger (0.64) for the wife. Both elasticities are smaller compared to their Frisch counterparts in the "Total" column in Table 3 (0.49 and 1.43 respectively).

Our model allows us to recover also the Marshallian elasticities of each earner's time use with respect to permanent shifts in the other earner's wage. Interestingly, for the wife the Frisch and Marshallian elasticities almost coincide, implying very little wealth effect operating in the context of the response of wife's childcare time to wage shocks. For husbands, we find that both the own and the cross Marshallian elasticities are somewhat smaller than their Frisch counterparts, implying a positive wealth effect associated with a decline in childcare time.

To have a better gauge of the magnitudes of these responses, in the bottom panel of Table 4 we report the response of each "good" to a 10% positive permanent wage shock. The responses are evaluated at the

[28] The κ's from equation (8) represent Marshallian responses to before-tax wage changes. See Blundell et. al. 2014 for an extensive discussion of before- vs. after-tax elasticities.

sample means. Starting from the husband, we calculate that a 10% permanent shock to his own wages is associate with an overall increase of only about 20 annual hours. However, this small increase masks a much larger decrease in childcare time, implying an overall slight increase in husband's leisure. The wife's overall labor supply response to a 10% permanent shock to her wages is much larger (100 annual hours). Similarly to the husband, she experiences a decline in childcare time and some increase in leisure. The wife's response to a 10% permanent shock to her husband's wage is negative and large (152 annual hours) and is almost entirely accounted for by an increase in childcare time. The pattern for the husband is very similar, although much smaller quantitatively.

It is also interesting to look at the total effect of a permanent change to wages on time spent with the child, which is again different for shocks to the husband's and the wife's wages. For a 10% permanent positive shock to the husband's wage (evaluated at the mean) there is overall an increase in the total time the two parents spend with the child (91 hours a year), while for a 10% shock to the wife's wage the overall effect is a 92 hours decline. However, it is worth stressing that children may spend time with both parents, and hence an increase or decline in the sum of hours of the two parents does not necessarily translate one to one into hours that children spend without their parents. In our sample, the mean (median) share of total childcare time in which the two parents either participate actively or are present is 0.32 (0.25).

Finally, as noted in Blundell et al. (2014) and discussed above, even though the consumption Frisch elasticities with respect to wages are negative (implying that consumption and hours are substitutes in the short run), the consumption Marshallian elasticities are always positive (implying complementarity between consumption and hours in the long run, or in response to permanent wage shocks).

5.3 Robustness Checks

Measuring childcare time inputs by the two parents using the time use diary from the CDS is an exercise that is potentially subject to errors. In this section we discuss two robustness tests in which we calculate childcare time differently. First, recall that in our sample 22% of males record zero time spent with the child for the day in which the diary was filled. Since we use an annual measure of childcare time, it is plausible that the infrequency of data collection may bias childcare time towards zero. In our baseline specification we assumed that husbands who report zero hours come from the lower part of the childcare time distribution and assigned them the childcare time value of the 5th percentile of the non-zero distribution of time for husbands. This type of treatment is consistent with the notion that the likelihood of observing positive childcare time dairy entry by the husband is decreasing in the time that he usually spends with the child. While this assumption seems realistic, the choice of which value to assign to those who report zero childcare time is somewhat arbitrary. To check the sensitivity of the results we repeat our estimation assuming that husbands who report zero childcare hours do so at random, which implies that we can use only the observations for husbands who did not report zero time use in estimation. The results are reported in panel A of Table 5. The consumption elasticities estimates, as well as the total labor supply responses, are very similar to our baseline estimates in Table 3. In the decomposition of labor supply to leisure and childcare time, the role of childcare time slightly increases compared to the estimates in Table 3 for both earners, though the difference is not statistically significant. For both earners this is because the childcare time elasticities are estimated to be slightly larger. For the husband it is also because of the increase in the weight of childcare time (which is mechanical due to the fact that we drop the zeros instead of replacing them with a point from the bottom of the distribution).

We assumed that childcare time included both the time that the caregiver is actively engaged with the child (playing, etc.) as well as time where the parent is merely present (supervising, etc.). One concern is that the second component could be interpreted as housework or leisure rather than childcare time. For example, if the child is watching TV, the parent might use the time for housework, or for reading a book. To address this issue, we use the distinction in the CDS diary between a person who *participates* in an activity with the child, and a person who is just *present*. We re-construct the childcare time for each parent only using activities for which the diary marks that the husband or wife participated in. The estimates are reported in panel B of Table 5. While the consumption elasticities and the total labor supply responses are again similar to the previous estimates, the decomposition of leisure and childcare time effects is somewhat different. First, the weight on ϖ ($\frac{T_1}{H_1}$ for the husband and $\frac{T_2}{H_2}$ for the wife) is much smaller. This is by construction, since the childcare time measure is constructed with only a subset of the activities used in the previous specifications. The median participation time is 33% of the total time for the husband and 45% of the time for the wife. These small weights drive down the importance of the childcare time channel. On the other hand, the childcare time elasticity for the wife is larger and somewhat more precisely measured. Combined with the smaller weight, the estimated leisure-related elasticity $\eta^{\gamma_0}_{h_2,w_2} = 0.67$ is now larger and more statistically significant compared to the baseline specification. Finally, the tension between the two parents' leisure complementarity and childcare time substitutability in the production function is dampened as both η_{h_1,w_2} and ϖ_{t_1,w_2} are smaller in absolute value (and similar effects emerge for the wife). Nevertheless, it is worth stressing that the qualitative direction of the effects is mostly confirmed.

6 Discussion

6.1 Childcare Spending

Following the discussion in Section 4, we have so far neglected childcare expenditure in our empirical analysis. In what follows, we look at the relation between childcare expenditure, wages, childcare time, and hours of work for families who do have non-zero childcare expenditure. We show that the substitution patterns between childcare expenditure and time are in line with economic intuition, but also show that these are unlikely to have a large effect on our estimates from Section 5.

To focus on households where childcare spending is likely to be more important in substituting for parents' time, we restrict our sample to households where the youngest child is 5 or less, and hence less likely to be in the public school system. Moreover, noting that 30% of parents with children 5 or less *and no older kids* report non-zero tuition, we add tuition to childcare. Hence, we assume that tuition spending is essentially reflecting pre-school tuition costs. To avoid the possibility that tuition spending refers to parents' rather than children's education, we drop parents who report to be full- or part-time students.

Table 6 provides reduced form evidence for the relation between childcare spending and wages. In Column 1 we regress the log change in childcare expenditure on the log change in wages of the two earners, controlling for the level and the change in the age of the youngest kid as well as the level and the change of number of kids overall and the number of kids in different age groups.[29] Because we are interested in the Frisch elasticities, we control for the change in consumption, which would capture (at least approximately) changes in permanent income (or lifetime wealth). The results show that a 10% change in the wife's wages is associated with about 1.4% change in childcare expenditures. There is no change in childcare expenditure associated with changes in the husband's wage. In the second column we also control flexibly for husband's and wife's education and ages, as well as for year effects. The results are similar. Given that the own Frisch response of childcare time for wages is negative for both earners, the positive Frisch elasticity of childcare expenditure is consistent with substitutability at least between wife's childcare time and market time. The high coefficient attached to the consumption growth variable suggests that childcare spending is a luxury.

In column 3 we take a semi-structural approach. Instead of controlling for wealth changes using change in consumption, we use the model to derive panel data instruments which will allow us to recover the response to transitory shocks. Suppose that we can write

$$\Delta k_{i,\tau} = \theta_{11}\Delta u_{i,1,\tau} + \theta_{12} v_{i,1,\tau} + \theta_{21}\Delta u_{i,2,\tau} + \theta_{22} v_{i,2,\tau} + e_{i,\tau},$$

Assuming $\mathbb{E}\left[e_{i,\tau}|u_{i,1,\tau}, u_{i,2,\tau}, v_{i,1,\tau}, v_{i,2,\tau}\right] = 0$, we can identify the coefficients θ_{11} and θ_{21} by running a regression of $\Delta k_{i,\tau}$ onto $\Delta w_{i,1,t}$ and $\Delta w_{i,2,t}$ and instrumenting the wage growth variables using $w_{i,1,\tau-1}$ and $w_{i,2,\tau-1}$ as instruments.[30] Column 3 shows the results of this exercise. While insignificant, the signs of the coefficients are similar, and so does the magnitude of the elasticity at least for the wife. Finally, Appendix Table 2 repeats this analysis for wives with at least some college attainment, finding that the association between childcare and wife's wages is stronger for the high educated.

While we find (weak) evidence for the substitutability of childcare expenditure and childcare time of the wife, this substitutability is unlikely to affect our structural estimates. To see why this may be the case, suppose that the Frisch elasticity of childcare expenditure with respect to the wife's wages ϖ_{k,w_2} is 0.15 as in column 2 of Table 6. Using the Frisch demand symmetry we can recover the Frisch elasticity of childcare

[29] The age groups are 6-7, 8-9, 10-12, 13-15, 16-18, 19-22, 23-26, $\geq$27. We control separately for these groups to make sure that tuition is not driven by transitions between elementary and junior high, junior high and high school etc.

[30] If wages of husband and wife are uncorrelated, the IV estimator of θ_{11} is $\frac{cov(\Delta k_t, \Delta w_{t-1})}{cov(\Delta w_t, \Delta w_{t-1})}$. The numerator converges asymptotically to $-\theta_{11} var\left(u_{1,t-1}\right)$, while the denominator converges asymptotically to $-var\left(u_{1,t-1}\right)$, proving that this IV procedure is consistent for θ_{11}. A similar argument applies to θ_{21}.

time of the wife with respect to the price of childcare:

$$\varpi_{t_2,p_k} = \varpi_{k,w_2} \frac{P_k K}{T_2 W_2}$$

From Table 2, the sample mean of $\frac{P_k K}{T_2 W_2}$ is 0.15 (with a median of 0), implying $\varpi_{t_2,p_k} = 0.023$, which is very small compared to our estimates of ϖ_{t_2,w_1} and ϖ_{t_2,w_2} in Table 3. Our inference is that neglecting childcare in the structural estimation is unlikely to affect our estimates and conclusions in Section 5.

Even though these results provide justification for neglecting childcare in our empirical application, it could be that childcare does play a greater role for families with very young kids (under 5 for example) and in the context of female participation decisions, which are absent from our framework. We see these issues as a fertile ground for future research on the subject.

6.2 Insurance

In Blundell et al. (2014) 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 this section we show that the role of family labor supply adjustment can be further decomposed into a leisure-related component and a childcare time-related component. To demonstrate the role of these adjustment in smoothing fluctuations in household welfare, we report both the effect they play in smoothing consumption, as well as the effect they have in smoothing the marginal utility of wealth ($\Delta \log \lambda$).[31]

To illustrate how we conduct this exercise, consider the expression for the innovation in the marginal utility of wealth. For simplicity, we focus on the simple case of a single earner who consumes, saves, and allocates time to work, leisure and childcare. If we also assume a linear tax system, the innovation in the growth of marginal utility can be written as:

$$\varepsilon_{i,t} = \frac{\left[(1-\pi_t)\left(1+\eta^{\gamma_0}_{h_1,w_1} - \frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,w_1}\right) - \eta_{c,w_1}\right]}{\eta_{c,w_1} - \eta_{c,p} - (1-\pi_t)\left(\eta_{h_1,p} + \eta_{h_1,w_1} - \frac{T_{1,\tau}}{H_{1,t}}\varpi_{t_1,w_1}\right)} v_{1,t}$$

which is clearly a function of savings (through π), of labor supply and time use (through $\eta^{\gamma_0}_{h_1,w_1}$, $\frac{T_{1,\tau}}{H_{1,t}}$ and ϖ_{t_1,w_1}), and of the consumption elasticities $\eta_{c,p}$ and η_{c,w_1}. Shutting down any of these channels would imply a different transmission of permanent wage shocks to the marginal utility of wealth (and hence to consumption). While too complicated to analyze analytically, our framework allows us to calculate how the marginal utility of wealth and consumption change in scenarios in which we shut down the different insurance channels for the full model as well.

The first row of Table 7 reports the average Marshallian consumption response to a 10% positive shock to the husband's wage rate. This is the consumption response that incorporates all the leisure, time use and labor supply adjustments. Note that while consumption increases by 3.7%, the decline in the marginal

[31] Note that even though the two are closely related, because of non-separability between consumption and leisure, changes in consumption are not proportional to changes in the marginal utility of wealth.

utility of wealth is higher in absolute value (4.9%). Next, we decompose this response into its various parts. In the second row we shut down the wife's time use response (i.e., assign $\varpi_{t_2,w_1} = \varpi_{t_2,w_2} = 0$), hold all the other parameters (including π and s) constant, and re-calculate the transmission coefficients from wage shocks to consumption (the κ's in equation (8)). The implied consumption response in this case is 4.9% (and the decline in marginal utility is 6.3%).

In the third row we shut down the time use responses of the husband as well ($\varpi_{i,j} = 0, \forall i, j$). As expected, this implies a *reduction* in the response of consumption to the permanent wage shock. This is because the husband's Marshallian childcare time elasticity is negative. Holding leisure constant, this implies that when husband time use responses are allowed, a permanent increase in his wages drives an increase in his hours as well (at the expense of childcare time). Finally, shutting down all labor supply responses (i.e., imposing that all elasticities other than $\eta_{c,p}$ are zero), consumption increases on average by 5.5% in response to a 10% permanent shock to the husband's wage, while marginal utility declines by 8.6%.

While this decomposition cannot be considered a formal counterfactual exercise (when shutting down elasticities, we do not allow for adjustment of savings for example), it does illustrate the importance of adjusting childcare time, especially by the wife, in buffering (or amplifying) permanent shocks to wages.

7 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-childcare goods and childcare 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 childcare related goods. The wage rate is the opportunity cost of time use alternative to the market (leisure or childcare).

We argue that in an environment with uncertainty hours devoted to the various activities (work, leisure, childcare) serve two roles. First, they can be used to produce "commodities", such as childcare (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 females to a transitory wage shock faced by the male (cross-Frisch elasticity) is small. However, this is because the large complementarity induced by wanting to spend time together is almost exactly counteracted by the input substitutability in the production function of childcare. Similarly, we find that own-Frisch elasticity are large because usual intertemporal substitution effects are compounded by reducing time spent with children due to the higher opportunity cost induced by the wage change.

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 childcare services, they are present also under less restrictive assumptions. We plan to explore the consequences of more general utility specifications in future work.

References

[1] Mark Aguiar and Erik Hurst. (2007). "Measuring Trends in Leisure: The Allocation of Time Over Five Decades", *The Quarterly Journal of Economics* 122 (3)

[2] Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. (2013). "Time use during the great recession." *The American Economic Review* 103, no. 5: 1664-1696.

[3] Angrist, Joshua D., and William N. Evans. (1998). "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size." *The American Economic Review* 88, no. 3: 450-477.

[4] Attanasio, Orazio, Hamish Low, and Virginia Sanchez-Marcos. (2005). "Female labor supply as insurance against idiosyncratic risk." *Journal of the European Economic Association* 3, no. 2-3: 755-764.

[5] Attanasio, Orazio, Hamish Low, and Virginia Sanchez-Marcos. (2008). "Explaining changes in female labor supply in a life-cycle model." *The American Economic Review*. 98, no. 4: 1517-1552.

[6] Becker Gary S. (1965). "A Theory of the Allocation of Time". *The Economic Journal,* Vol. 75, No. 299: 493-517

[7] Becker, Gary S. (1976), *The Economic Approach to Human Behavior*, University of Chicago Press.

[8] Blau, David M., and Philip K. Robins. (1988). "Child-care costs and family labor supply." T*he Review of Economics and Statistics*: 374-381.

[9] Blundell Richard, Luigi Pistaferri and Ian Preston (2008), "Consumption inequality and partial insurance."*American Economic Review*, 98(5), 1887-1921.

[10] Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. (2014). "Consumption Inequality and Family Labor Supply".

[11] Bound John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers (1994), "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics*, 12(3), 345-368.

[12] Browning, Martin, Pierre-Andre Chiappori, and Yoram Weiss. (2014). "Economics of the Family". Cambridge University Press.

[13] Ghez, Gilbert, and Gary S. Becker. (1975). "The allocation of time and goods over the life cycle." *NBER Books*.

[14] Goux, Dominique, Eric Maurin, and Barbara Petrongolo. (2014). "Worktime Regulations and Spousal Labor Supply." *The American Economic Review* 104, no. 1: 252-76.

[15] Guryan, Jonathan, Erik Hurst, and Melissa Kearney. (2008). "Parental Education and Parental Time with Children." *Journal of Economic Perspectives*, 22(3): 23-46.

[16] Hyslop, Dean R. (2001). “Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings.”*American Economic Review*, 91(4), 755-77.

[17] Juhn, Chinhui and Simon Potter (2007). "Is There Still an Added Worker Effect?". Federal Reserve Bank of New York. Staff Report no. 310.

[18] Keane, Michael P. (2011) "Labor supply and taxes: A survey." *Journal of Economic Literature* 49, no. 4: 961-1075.

[19] Lundberg, Shelly (1985), “The Added Worker Effect”, *Journal of Labor Economics*, 3(1), 11-37.

[20] MaCurdy, Thomas E. (1983), "A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty", *International Economic Review*, 24(2), 265-289.

[21] Meghir, Costas and Luigi Pistaferri (2004), “Income variance dynamics and heterogeneity”, *Econometrica* 72(1), 1-32.

[22] Meghir, Costas and Luigi Pistaferri. (2011). “Earnings, consumption and lifecycle choices”, *Handbook of Labor Economics* , 4(B), 773-854.

[23] Ramey, Garey, and Valerie A. Ramey. (2010). "The Rug Rat Race." *Brookings Papers on Economic Activity 2010*, no. 1: 129-176.

[24] Sandberg, John F., and Sandra L. Hofferth. (2001). "Changes in children’s time with parents: United States, 1981–1997." *Demography* 38, no. 3: 423-436.

[25] Sandberg, John F., and Sandra L. Hofferth. (2005). "Changes in children’s time with parents: A correction." *Demography* 42, no. 2: 391-395.

[26] Stephens, Mel (2002), "Worker Displacement and the Added Worker Effect", *Journal of Labor Economics*, 20(3), 504-537.

Table 1: **Descriptive Statistics - PSID Data**

	(1)	(2)	(3)
Sample	At least one kid under 13	All Married Couples with working males	All Married Couples
Consumption and Assets			
Total Consumption	39,654	39,215	38,715
Nondurable Consumption	8,946	8,509	8,433
Services	30,709	30,707	30,282
childcare	1,403	717	694
Total assets	316,185	354,786	380,476
First earner (head)			
Earnings	66,525	67,048	63,495*
Hours worked	2,314	2,302	2,202*
Second earner (wife)			
Participation rate	0.75	0.80	0.79
Earnings|Work	31,553	32,998	33,522
Hours worked|Work	1,554	1,688	1,698
Observations	5,282	10,476	12,265

Notes: PSID data from 1999-2009 PSID waves. Column 1 refers to the main sample: married couples with working male aged 30-57 and kids 13 years or younger. SEO sample excluded. The table reports averages for all variables. In column 2 we remove the sample restriction of kids aged 13 or younger. In column 3 we also remove the sample restriction for working males. For this final column head's earnings and hours are reported conditional on work (marked with a *). Missing values in consumption and assets subcategories are treated as zeros.

Table 2: **Descriptive Statistics - CDS Data**

	(1)	(2)	(3)	(4)
	mean	p25	median	p75
Non-zero childcare time (head)	0.78			
Non-zero childcare time (wife)	0.95			
childcare annual hours (head)|non zero	414	43	282	608
childcare annual hours (wife)|non zero	832	421	782	1,108
childcare hours/hours worked (wife)|non zero	4.63	0.23	0.49	0.99
childcare hours/hours worked (head)|non zero	0.20	0.02	0.13	0.28
childcare cost: (head time)/(wife time)|non zero	1.41	0.12	0.60	1.52
childcare cost: (market childcare)/(wife time)|non zero	0.15	0.00	0.00	0.05
Observations	686			

Notes: CDS data from 2001, and 2007 waves. Sample of married couples with working male aged 30-57 and kids 13 years or younger with matched CDS diary data. SEO sample excluded. Childcare time includes all daily activities in which the parent was either present or actively participated.

Table 3: **Main Results: Frisch Responses**

	(1)	(2)	(3)	(4)
	η^{γ_0}	ϖ	weight on ϖ	Total
Panel A: Husband Time Use and Labor Supply				
H1 to W1	0.430^{***} (0.163)	-0.483 (1.164)	0.139	0.497^{***} (0.054)
H1 to W2	0.393^{***} (0.120)	1.549^{***} (0.547)	0.139	0.178^{***} (0.052)
H1 to P	0.055^{*} (0.031)	-.-	-.-	0.055^{*} (0.031)
Panel B: Wife Time Use and Labor Supply				
H2 to W1	0.749^{***} (0.235)	1.445^{***} (0.510)	0.605	-0.125 (0.093)
H2 to W2	0.361 (0.926)	-1.774 (1.548)	0.605	1.434^{***} (0.155)
H2 to P	0.242^{*} (0.090)	-.-	-.-	0.242^{*} (0.090)
Panel C: Consumption				
C to W1	-0.063^{*} (0.036)	-.-	-.-	-0.063^{*} (0.036)
C to W2	-0.135^{***} (0.051)	-.-	-.-	-0.135^{***} (0.051)
C to P	0.642^{***} (0.120)	-.-	-.-	0.642^{***} (0.120)
Smoothing Parameters				
$\mathbb{E}(\pi)$	0.097			
$\mathbb{E}(s)$	0.698			
Observations	$3,453$			

Notes: Parameters estimated by GMM. PSID data from 1999-2009 PSID waves. Sample includes married couples with working male aged 30-57 and kids 13 years or younger. SEO sample excluded. Childcare time includes all daily activities in which the parent was either present or actively participated. *, **, *** = Significant at 10%, 5%, and 1%. GMM standard errors clustered by household in parenthesis.

Table 4: **Marshallian Elasticities**

	Husband wage	Wife Wage
Elasticities		
Consumption	0.366	0.274
Husband labor supply	0.087	-0.192
Wife labor supply	-0.976	0.641
Husband childcare time	-1.086	1.005
Wife childcare time	1.631	-1.606
Derivative w.r.t. 10% wage shock		
Consumption ($)	1,451	1,087
Husband labor supply (Annual hours)	20	-44
Wife labor supply (Annual hours)	-152	100
Husband childcare time (Annual hours)	-45	42
Wife childcare time (Annual hours)	136	-134

Notes: Parameters estimated by GMM. PSID data from 1999-2009 PSID waves. Sample includes married couples with working male aged 30-57 and kids 13 years or younger. SEO sample excluded. Childcare time includes all daily activities in which the parent was either present or actively participated.

Table 5: **Robustness Checks**

	A: No Replacement of 0 Time for Husband				B: Exclude "Time Present" from Time Calculation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	η^{γ_0}	ϖ	weight on ϖ	Total	η^{γ_0}	ϖ	weight on ϖ	Total
Panel A: Husband Time Use and Labor Supply								
H1 to W1	0.364^{**} (0.163)	-0.750 (0.918)	0.182	0.501^{***} (0.054)	0.403^{***} (0.084)	-1.866 (1.208)	0.060	0.516^{***} (0.056)
H1 to W2	0.480^{***} (0.182)	1.791^{***} (0.623)	0.182	0.154^{***} (0.045)	0.139^{***} (0.075)	0.789 (0.606)	0.060	0.092^{**} (0.042)
H1 to P	0.060^{*} (0.031)	-.-	-.-	0.060^{*} (0.031)	0.070^{**} (0.032)	-.-	-.-	0.070^{**} (0.032)
Panel B: Wife Time Use and Labor Supply								
H2 to W1	0.913^{***} (0.290)	1.671^{***} (0.581)	0.605	-0.097 (0.082)	0.265^{*} (0.142)	0.944 (0.726)	0.272	0.008 (0.00)
H2 to W2	0.124 (0.926)	-2.149 (1.484)	0.605	1.424^{***} (0.153)	0.665^{*} (0.404)	-2.660^{*} (1.453)	0.272	1.389^{***} (0.151)
H2 to P	0.227^{**} (0.089)	-.-	-.-	0.227^{**} (0.089)	0.174^{**} (0.087)	-.-	-.-	0.174^{**} (0.087)
Panel C: Consumption								
C to W1	-0.069^{*} (0.036)	-.-	-.-	-0.069^{*} (0.036)	-0.080^{**} (0.036)	-.-	-.-	-0.080^{**} (0.036)
C to W2	-0.127^{**} (0.050)	-.-	-.-	-0.127^{**} (0.050)	-0.097^{**} (0.049)	-.-	-.-	-0.097^{**} (0.049)
C to P	0.636^{***} (0.120)	-.-	-.-	0.636^{***} (0.120)	0.596^{***} (0.113)	-.-	-.-	0.596^{***} (0.113)
Smoothing Parameters								
$\mathbb{E}(\pi)$	0.097							
$\mathbb{E}(s)$	0.698							
Obs.	$3,453$							

Notes: Parameters estimated by GMM. PSID data from 1999-2009 PSID waves. Sample includes married couples with working male aged 30-57 and kids 13 years or younger. SEO sample excluded. In Panel A, childcare time includes all daily activities in which the parent was either present or actively participated. Whenever the husband's childcare time was recorded zero, the observation is not used in estimation of the time use related moments. In Panel B, childcare time includes only activities in which the parent actively participated. *, **, *** = Significant at 10%, 5%, and 1%. GMM standard errors clustered by household in parenthesis.

Table 6: **Childcare Wage Elasticities**

	(1)	(2)	(3)
Dependent variable: $\Delta \log(\text{childcare})$			
Δw_2 (wife wage)	0.143^{*}	0.150^{*}	0.109
	(0.078)	(0.080)	(0.095)
Δw_1 (husband wage)	-0.032	0.018	-0.125
	(0.117)	(0.124)	(0.253)
Δc (consumption)	1.331^{***}	1.509^{***}	1.498^{***}
	(0.207)	(0.229)	(0.221)
observations	581	581	581
Estimation	OLS	OLS	IV
Age, educ, year controls	No	Yes	Yes
F-test (P-value)			
Δw_2 (wife wage)			27.1
			(<0.001)
Δw_1 (husband wage)			461.2
			(<0.001)

Notes: PSID data from 1999-2009 PSID waves. Sample includes married couples with working male aged 30-57 and kids 5 years or younger. SEO sample excluded. Column 1 reports the results from a regression of change in log childcare on the change in hourly wages of both earners and the change in log consumption, controlling also for the level and the change in the age of the youngest kid as well as the level and the change of number of kids overall and the number of kids in different age groups. Column 2 repeats the regression controlling also for year effects, age and education of the two earners. Column 3 instruments for changes in log wages using lagged log wages. *, **, *** = Significant at 10%, 5%, and 1%. Standard errors clustered by household in parenthesis.

Table 7: **Decomposing Welfare Response to a 10 percent Positive Permanent Shock**

	Consumption Response (%)	Marginal Utility (λ) response (%)
Observed consumption response	3.7	-4.9
Shutting down wife's time use response	4.9	-6.3
Shutting down all time use responses	4.4	-5.8
Shutting down all labor supply responses	5.5	-8.6

Notes: Insurance decomposition calculations based on the estimates from Table 3. See Section 6.2 in the text for details.

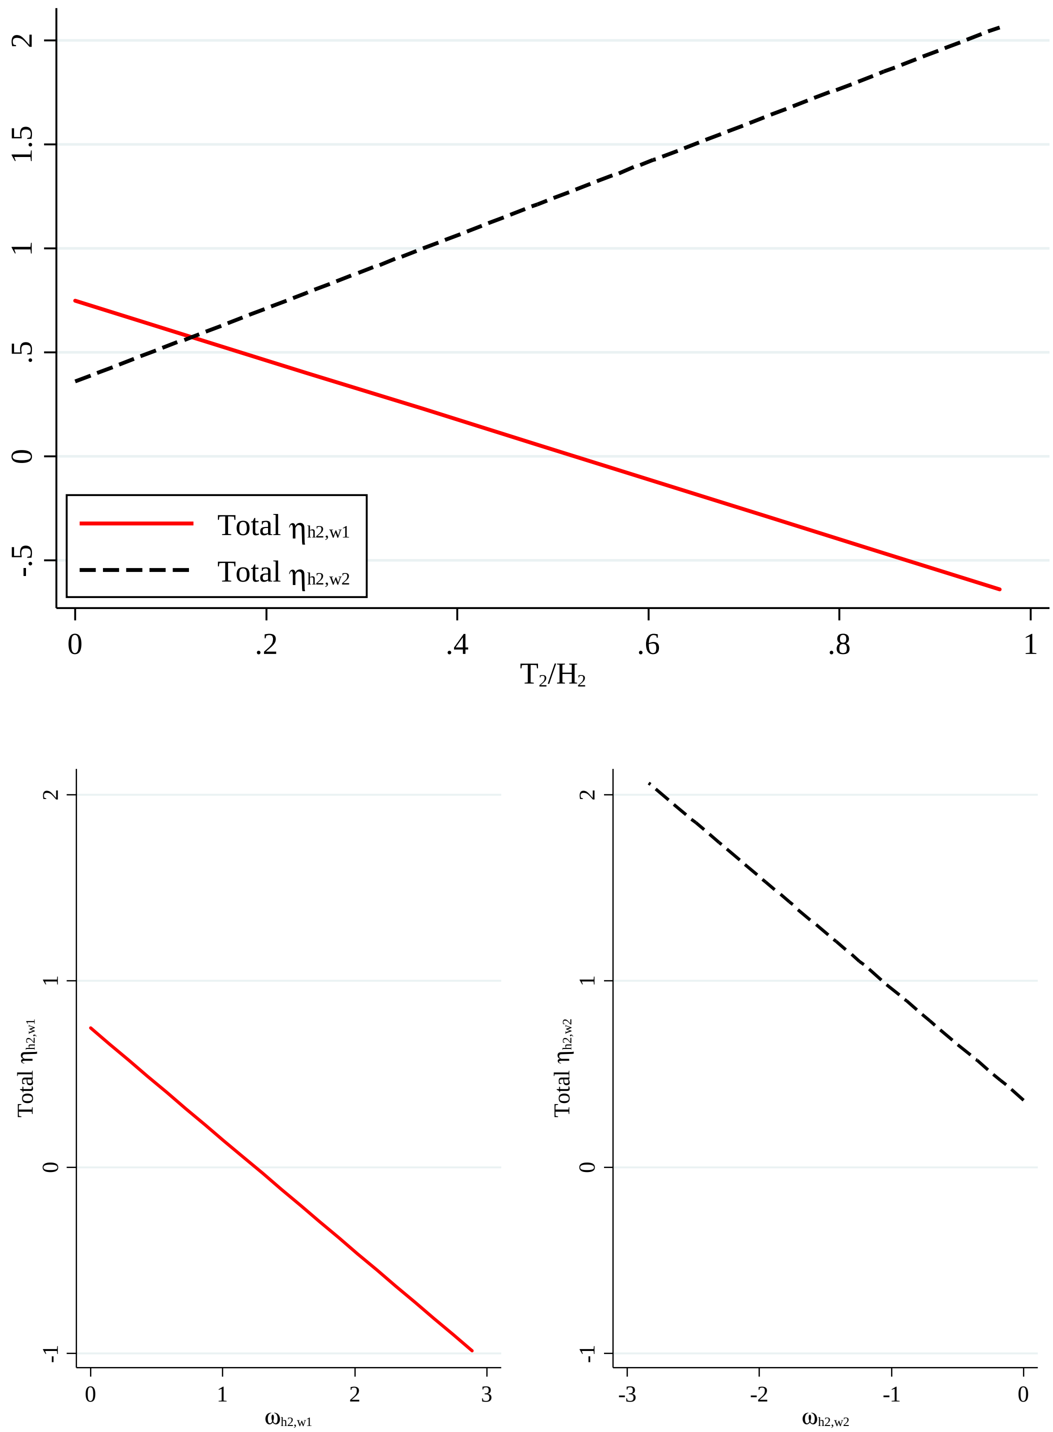

Figure 1: Sensitivity of Total Hours Frisch Elasticity to T_2/H_2 and ϖ