

NBER WORKING PAPER SERIES

AGGREGATING ELASTICITIES:
INTENSIVE AND EXTENSIVE MARGINS OF FEMALE LABOUR SUPPLY

Orazio Attanasio
Peter Levell
Hamish Low
Virginia Sánchez-Marcos

Working Paper 21315
<http://www.nber.org/papers/w21315>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2015

We are grateful for a number of useful conversations with Richard Blundell, Guy Laroque, Costas Meghir, Richard Rogerson and Tom Sargent. We received several useful comments by different seminar audiences and during presentation at the NBER Summer Institute and the Society for Economic Dynamic Conference. Attanasio's research was partially funded by an ESRC Professorial Fellowship and by the ESRC Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies. Sánchez-Marcos thanks Spanish MCYT for Grant ECO2009-09614. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Orazio Attanasio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Aggregating Elasticities: Intensive and Extensive Margins of Female Labour Supply
Orazio Attanasio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos
NBER Working Paper No. 21315
July 2015
JEL No. D12,J22

ABSTRACT

There is a renewed interest in the size of labour supply elasticities and the discrepancy between micro and macro estimates. Recent contributions have stressed the distinction between changes in labour supply at the extensive and the intensive margin. In this paper, we stress the importance of individual heterogeneity and aggregation problems. At the intensive margins, simple specifications that seem to fit the data give rise to non linear expressions that do not aggregate in a simple fashion. At the extensive margin, aggregate changes in participation are likely to depend on the cross sectional distribution of state variables when a shock hits and, therefore, are likely to be history dependent. We tackle these aggregation issues directly by specifying a life cycle model to explain female labour supply in the US and estimate its various components. We estimate the parameters of different component of the model. Our results indicate that (i) at the intensive margin, Marshallian and Hicksian elasticities are very heterogeneous and, on average, relatively large; (ii) Frisch elasticities are, as implied by the theory, even larger; (iii) aggregate labour supply elasticities seem to vary over the business cycle, being larger during recessions.

Orazio Attanasio
Department of Economics
University College London
Gower Street
London WC1E 6BT
UNITED KINGDOM
and NBER
o.attanasio@ucl.ac.uk

Peter Levell
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
peter_l@ifs.org.uk

Hamish Low
Faculty of Economics
University of Cambridge
Cambridge CB3 9DD ENGLAND
hamish.low@econ.cam.ac.uk

Virginia Sánchez-Marcos
Department of Economics
Universidad de Cantabria
Santander
virginia.sanchez@unican.es

1 Introduction

How much do hours of work and participation rates respond to changes in wages? For a long time, there has been a tension between labour economists, who estimated labour supply elasticities from individual level data at relatively low levels, especially for men, and macroeconomists, who, from business cycle fluctuations of wages and hours, have argued that labour supply elasticities are relatively large. Ljungqvist and Sargent (2014) have recently summarized this debate citing Carneiro and Heckman (2003) and Prescott (2002)¹. The controversy has stimulated a number of recent papers, such as those published in the AER papers and proceedings in 2011², Rogerson and Keane (2012) and Chetty et al. (2013), as well as many others. This debate is important because of the implications it has for the effect of changes in the structure of labour income taxes on labour supply and to interpret variations in employment and hours of work over the business cycle.

As argued by Blundell and MaCurdy (1999) and, more recently by Keane (2009), the term ‘wage elasticity’ may refer to different quantities depending on the type of variation in wages one is considering. On the one hand, one can consider the effect of changes in the entire wage structure, as induced, for instance, by a permanent changes in labour income taxation (or in the comparison between different countries). On the other, one can consider short term variations in wages, such as those one observe over the business cycle, akin to what Blundell and MaCurdy (1999) and MaCurdy (1985) define ‘evolutionary’ wage changes and that might be of particular interest to macroeconomists.³ Different type of variations in wages can be mapped in different theoretical concepts. The size of changes in labour supply induced by evolutionary wage changes is related to the size of the Frisch (or marginal utility of wealth constant) elasticity, while the size of changes induced by permanent shifts to the wage structure are determined by the size of Hicksian or Marshallian elasticities, depending on whether the changes in wages are compensated or not.⁴

In each of these cases, the labour supply response can be thought of in terms of the intensive (hours) or the extensive (participation) margin. At the individual level, an elasticity is easily defined when thinking of the intensive margin, while the same concept is a bit vaguer when thinking of the extensive margin, especially when thinking of the Frisch elasticity that is supposed to keep the marginal utility

¹Carneiro and Heckman (2003, p. 196): “In a modern society, in which human capital is a larger component of wealth than is land, a proportional tax on human capital is like a nondistorting Henry George tax as long as labor supply responses are negligible. Estimated intertemporal labor supply elasticities are small, and welfare effects from labor supply adjustment are negligible.”

Prescott (2002, pp. 13, 1): “The differences in the consumption and labor tax rates in France and the United States account for virtually all of the 30-percent difference in the labor input per working-age person. . . . if France modified its intratemporal tax wedge so that its value was the same as the U.S. value, French welfare in consumption equivalents would increase by 19 percent.”

²See Blundell et al. (2011), Chang et al. (2011), Chetty et al. (2011), Ljungqvist and Sargent (2011)

³A similar distinction is made by Chetty (2012) and Chetty et al. (2013).

⁴Blundell and MaCurdy (1999) and Keane (2009) discuss clearly how the concepts of Marshallian and Hicksian elasticities, which are typically derived within the framework of a static model, can be put within the framework of a dynamic life cycle model through the machinery of two-stage budgeting, as developed by Gorman (1959) and applied to labour supply by MaCurdy (1981, 1983) and Blundell and Walker (1986).

of wealth constant. For a macroeconomist, the next step is to think of how these individual responses are reflected in changes in employment and hours of work. Indeed, in the case of the extensive margin, one can think of the impact that a change in wages has on the fraction of individuals that change their participation status, given the distribution of state variables. In this sense, the consideration of the extensive margin brings to the forefront aggregation issues that have not figured prominently in the discussion of labour supply elasticities. Aggregate participation responses to an aggregate shock are bound to depend on the distribution of state variables in the cross section. As we discuss below, aggregation issues can also be relevant for the intensive margin.

The extent of disagreement over the values of the labour supply elasticities depends on which elasticity is being considered. Chetty (2012) finds that the estimates of the Hicksian elasticity from micro data are consistent with macroeconomic estimates once we allow for small optimization frictions such as adjustment costs or inattention of the order of 1%. By contrast, he finds that estimates of the Frisch elasticities are inconsistent: estimates of the higher values of the Frisch elasticity from a macroeconomic perspective such as Rogerson and Wallenius (2009) appear at odds with the microeconomic estimates that some papers identify from temporary tax reforms or other natural experiments. Many other recent contributions to understanding the disagreement over labour supply elasticities have focused on the extensive margin, as discussed by Keane and Rogerson (2012) and Chetty (2012). Rogerson and Wallenius (2009) argue that indivisible labour explains discrepancies between the micro and macro elasticities. They develop a macro model in which elasticities at the extensive and intensive margins are effectively unrelated. The explanation for this is that if there is fixed cost of entry into the labour market the aggregate employment rate depends on the distribution of reservation wages.

In this paper, we step back from the concept of an *elasticity* as a single parameter. Instead, our focus is on the determinants of different elasticities and how they relate to the quantities that are discussed in the literature. The key feature of our approach is that we consider an integrated model of intratemporal and intertemporal labour supply choices at both the intensive and the extensive margins. We estimate the parameters of this model using rich data that include information on consumption. We can then study how these parameters translate into individual elasticities of labour supply, both in terms of hours of work and in terms of participation in the labour force, and to show how these elasticities vary across individuals, and with characteristics such as age, number of children, and the extent of uncertainty in the economy.⁵ Finally, we can study how aggregate labour supply responses arise from individual behaviour.

The explicit consideration of even relatively simple preference specifications makes it apparent that labour supply elasticities might be very heterogeneous in the population and over time. Aggregation issues undermine the very concept of an aggregate labour supply elasticity. The concept of labour

⁵The distinction between estimating preference parameters and calculating elasticities in different economic environments is stressed by Keane and Rogerson (2011) and Domeji and Floden (2006).

supply elasticity as a structural parameter is particularly elusive in the case of the extensive margin, where by the very nature of the problem, responses might be time varying and aggregate differently over the business cycle.

While specifying a utility function is unavoidably restrictive as it imposes on the data a substantial amount of structure, it makes clear what aspects of the data generate certain levels of elasticities. We use relatively flexible specifications that allow for different degrees of substitutability between consumption and leisure, intertemporal substitutability, different utility costs of changes to labour supply at the intensive and the extensive margin, a rich role for demographic and other variables.

To estimate preference parameters, we use a variety of approaches. In particular, we use different sets of equilibrium conditions, and therefore different sources of variability to estimate different components of preferences. In this respect, in the estimation of each set of parameters, we try to minimize the assumptions needed for the identification of a specific set of parameters. We show that intra-temporal first order conditions can be used to identify a set of preference parameters that determine Marshallian and Hicksian labour supply elasticities. In order to get estimates of these parameters one can in principle use cross-sectional variation in prices. It is important, however, to use variability in wages that is plausibly exogenous and unrelated to preference heterogeneity. For this reason, information from different labour markets, possibly over time, can be useful.

As discussed above, the consideration of intra-temporal first order conditions is useful in itself as the elasticities that can be identified from such framework can be appropriate to judge the extent of labour supply responses to changes in the entire structure of wages. Moreover, the Hicks elasticity provides a lower bound on the Frisch elasticity. However, the intratemporal first-order condition is uninformative about the separability of consumption and leisure and about how much larger the Frisch elasticity is. To get a grip on these issues and estimate the parameters that allow the computation of Frisch elasticities it is necessary to bring a new set of equilibrium conditions to bear on the data. In particular, we use intertemporal Euler equations to identify these parameters. The data requirements that are necessary for the identification of these parameters are obviously larger than those required to identify the determinants of Marshallian or Hicksian elasticities. In particular, to avoid making strong and unrealistic assumptions about the completeness of markets, we need a long time series of data to identify the parameters of the Euler equation.

Finally, whilst in estimating the Euler equation we allow for the possibility of corner solutions in hours (that is the possibility of the extensive margin), we do not model the extensive margin explicitly. Therefore, Euler equation estimation cannot be used to estimate all the parameters of the utility function and learn about the relevance of the extensive margin. The big advantage of the Euler equation is that focusing on an equilibrium condition on a specific margin avoids the necessity of solving the model explicitly to derive policy functions. It also avoids the necessity of specifying all

the details of the dynamic problem solved by the individuals considered. However, by its very nature, to study the extensive margin it is necessary to get such policy functions and, therefore, specify all the details of the life cycle problem. In principle, the result one obtains on the extensive margin depends on every single detail of the life cycle problem considered, from the nature of the income process to pension arrangements, to the type of markets agents have access to. It should be stressed however, that some of these channels have only a marginal effect and that we can perform a number of robustness exercises, in addition to the standard matching of certain moments of the data. It should also be stressed that the set of parameters that is identified from the equilibrium conditions discussed above (the intratemporal ones and the Euler equation) are robust to the specific details of the life cycle problem considered.

The second crucial step in our approach is going from the characterization of individual preferences to the determination of ‘aggregate’ elasticities or elasticities defined at the macro level. In what follows, we stress the difficulty of this exercise. In the case of the intensive margin, a number of important non-linearities generate a substantial level of heterogeneity that makes aggregation very difficult. And matters are considerably more complicated at the extensive margin. The presence of non-convexities (such as fixed costs to go to work) induces some level of inertial behaviour (such as that studied in Chetty (2012)) and clustering around kinks and corners of the budget constraint. The relevance of this clustering for aggregate fluctuations depends on the size of shocks to wages and, crucially, on how thick these clusters are. The extent to which individuals are spread around kinks and corners of individual budget constraints is bound to depend on the history of individual and aggregate shocks. Therefore the aggregate ‘extensive margin elasticity’ will be time varying and bound to be cyclical. Responses are likely to be higher after a sequence of shocks with the same size than after a period of relative calm.

Armed with our empirical estimates and the flexible labour life cycle model, we study female labour supply in the US. The results we obtain are somewhat surprising. First, even when considering the elasticity of labour supply at the intensive margin, we find a substantial amount of heterogeneity in the size of elasticity. The elasticities vary by age, family composition, and the level of consumption. There is no sense in which we can talk about an aggregate labour supply elasticity, even as an approximation. Second, the size of these elasticities is considerably larger (in absolute value) than many of the estimates reported in the literature. The Marshallian median elasticity for females about 0.70 and, as theory predicts, the figures for the Hicksian (1.08) and Frisch (1.35) elasticities are higher. We believe that the higher values for our estimates of the elasticities is linked to the explicit use of consumption data we make. Interestingly, our results are consistent with recent evidence presented by Blundell et al. (2015), who use a completely different methodology from the one we employ and data from the PSID. They estimate a 0.40 Marshallian elasticity and a Frisch elasticity of 1. While

the method is different, they also use explicitly data on non-durable consumption. In regard to the extensive margin Frisch elasticity we find that it is decreasing in age, being 0.8 at the age of 36. Third, we find that consumption and hours are complements consistent with findings in Ziliak and Kniesner (2005) for male labor supply.⁶ Finally, the conclusion of our aggregation exercise is that the emphasis of the literature on ‘the’ elasticity of labour supply to wages is misplaced. Not only does aggregation fail even for relatively simple specification of preferences, but it fails in fundamental and economically relevant ways in a variety of dimensions. Particularly important is the elasticity of participation to wages: by the very nature of the decision, such elasticity is likely to be dependent not only on cross sectional heterogeneity but to be time varying, with different values in different parts of the business cycle. We show that estimated elasticities do vary over the business cycle by a substantial amount. To the best of our knowledge, ours is the first systematic evidence of such a fact.

Our exercise is not without important caveats. In much of our analysis, we do not consider the effect of tenure and experience on wages. Such effects can obviously be important, as labour supply choices will change future wages and, therefore, future labour supply behaviour. Imai and Keane (2004) argue that assuming wages are exogenous may introduce a downward bias in the estimates of the Frisch elasticity. Indeed, they present estimates of such a parameter as high as 3.8 in a model that accounts for returns to labor market experience.⁷ We notice, however, that if tenure effects happen only through participation (rather than hours of work), the analysis we present of the intensive margin goes through and our estimates of the Marshallian, Hicksian and Frisch elasticities for the number of hours (conditional on working) are unbiased. What does change, in such a case, is the analysis of the extensive margin. In section 7, we discuss the implications of introducing returns to tenure on the extensive margin. It should be noted, however, that if the return to tenure operate on the number of hours (rather than only on participation), we would need to change our analysis substantially. We leave that for future work.

When estimating the Euler equation for consumption we also ignore the possibility of liquidity constraints that might prevent households from being at the relevant intertemporal margin. As discussed by Domeij and Floden (2006), omitting credit constraints may lead to underestimates of the Frisch elasticity, and as shown by Low (2005) uncertainty over future wages may reduce individuals’ willingness to exploit inter-temporal substitution opportunities.

The rest of the paper is organized as follows. In section 2, we describe the life cycle model we use as a framework for our analysis. We provide details of our preference specification and show

⁶Ziliak and Kniesner (2005) estimate the incentive effects of income taxation in a life-cycle model of consumption and male labor supply that allows for non-separability between consumption and labour supply. They are able to identify both within-period preference parameters and inter-temporal preference parameters. While their exercise focuses on male labour supply and uses a different data source (the PSID), consistently with what we find, their result indicate that consumption and labour supply are complement. Their estimates of compensated labour supply elasticities (Hicksian) are also considerably larger than those previously reported in the literature.

⁷However, as discussed in Wallenius (2011), Imai and Keane (2004) base their identification on the early periods of the life-cycle. The model does a less good job of accounting for the life-cycle profile at later ages using these estimates.

how preference parameters can be mapped into static and intertemporal elasticities. In section 3 we explain the various components of our empirical strategy to identify the preference parameters, that is, using intraperiod first order conditions, intertemporal first order conditions and full structural estimation. Section 4 describes the data and provides some descriptive statistics. Section 5 presents and discusses the parameter estimates. In section 6 we report the implications of our estimates for labour supply elasticities, distinguishing between Marshallian, Hicksian and Frisch elasticities. We also discuss aggregate responses on the extensive margin and, more generally, the aggregation issues that are central to our argument. Section 7 extends the analysis to include returns to experience and section 8 concludes.

2 A life cycle model of female labour supply

To study the elasticity of female labour supply to wages, we use a rich model of female labour supply choices embedded in a life cycle framework. A unitary household makes choices about consumption and female labour supply, given exogenous processes for male earnings and female wages and an intertemporal budget constraint. Both the intensive and extensive margins are meaningful because of the presence of fixed costs of going to work (possibly related to family composition) and/or because of the presence of preference costs specifically related to participation.

We assume that couples are expected utility maximisers and choose consumption, saving and female labour supply to solve the following dynamic problem under uncertainty:

$$\max_{c,t} E_t \sum_{j=0}^T \beta^j u(c_{t+j}, l_{t+j}, P_{t+j}; z_{t+j}, \zeta_{t+j}, \chi_{t+j}) \quad (1)$$

subject to an intertemporal budget constraint:

$$A_{t+1} = R_{t+1} \left(A_t + \left(w_t^f (H - l_t) - F(a_t) \right) P_t + w_t^m \bar{h} - c_t \right) \quad (2)$$

where c_t is consumption, l_t female labour supply, A_t are beginning of period assets, R_t is the interest rate, F the fixed cost of work which depends on a_t , the age of the youngest child. P_t is an indicator of labour force participation. z_{t+j} is a vector of observable variables (such as family composition) and χ_{t+j} and ζ_{t+j} represent unobservable taste shifters. Female wages are given by w_t^f , and husband wages are given by w_t^m , with fixed husband hours of \bar{h} . In any period, households are able to borrow against the minimum income they can guarantee for the rest of their lives.

We denote the child care units needed by a family whose youngest child is age a_t by $G(a_t)$ and the price of each unit of child care by p . Therefore, the total child care cost faced by a household when women participate in the labor market is given by

$$F(a_t) = pG(a_t) + \bar{F} \quad (3)$$

We estimate the function $G(a_t)$ from expenditure data of households with children of the relevant ages. The presence of fixed costs of going to work and discrete utility costs introduces the possibility that some women will decide not to work at all, especially at low levels of productivity. By the same token, it will be unlikely that women who do choose to work will work only very few hours.

We assume men always work. Male earnings are given by

$$\ln w_t^m = \ln w_0^m + \alpha_1^m t + \alpha_2^m t^2 + v_t^m \quad (4)$$

v_t^m is a random process that we describe below.

In this baseline specification, female wages are given by

$$\ln w_t^f = \ln w_0^f + \ln h_t^f + v_t^f \quad (5)$$

where h_t^f is the level of female human capital at the start of the period and v_t^f is a permanent productivity shock. There is an initial distribution of wages, w_0^f .

In our baseline specification we assume that human capital does not depend on the history of labour supply and evolves exogenously:

$$\ln h_t^f = \alpha_1^f t + \alpha_2^f t^2 \quad (6)$$

We relax the assumption that there are no returns to experience in section 7. We distinguish the cases where returns to experience depend on participation and where returns depend on hours worked. Much of our estimation steps will go through if returns to experience operate through the participation margin rather than through the hours of work margin.

Both female and male wages, w_t^f and w_t^m , in the household are subject to permanent shocks, v_t^f and v_t^m , that are positively correlated. In particular we assume

$$v_t = v_{t-1} + \xi_t \quad (7)$$

$$\xi_t = (\xi_t^f, \xi_t^m) \sim N(\mu_\xi, \sigma_\xi^2) \quad (8)$$

$$\mu_\xi = \left(-\frac{\sigma_{\xi^f}^2}{2}, -\frac{\sigma_{\xi^m}^2}{2}\right) \text{ and } \sigma_\xi^2 = \begin{pmatrix} \sigma_{\xi^f}^2 & \rho_{\xi^f, \xi^m} \\ \rho_{\xi^f, \xi^m} & \sigma_{\xi^m}^2 \end{pmatrix}$$

In this framework, innovations to wages and to interest rates constitute the uncertainty that households face. They could also face uncertainty over fertility and child care costs. We assume that they know they will remain married. When we proceed to step 3 of our estimation through solving numerically the model, we will impose additional restrictions, namely that the interest rate is constant and fertility is known. Further, from the point of view of the consumer, current taste shocks are observed. From the point of view of the econometricians, there are several sources of unobserved

variation: the innovations to wages and earnings, innovations to interest rates and the unobserved heterogeneity terms.

So far we have described the process faced by an individual household. This household takes the stochastic processes that generate female wages, male earnings and possibly interest rates as given. In making predictions about future factor prices (wages and interest rates), the household will consider the current level of the stochastic variables and make the best use of this information. We assume that households are subject to both idiosyncratic and aggregate shocks and so the shocks that affect individual households at a given point in time are correlated. However, from the household's perspective, they do not distinguish aggregate and idiosyncratic shocks and condition their future expectations only on their own observed wage realisations. Households have no insurance markets to smooth aggregate or idiosyncratic shocks and must rely on self-insurance. We assume there are no explicit borrowing constraints.

Our framework is not a general equilibrium one: we do not construct the equilibrium level of wages (and interest rates). However, we study aggregate female labour supply and its elasticity to wages. We do so by simulating a large number of households and aggregating explicitly their behaviour.

2.1 Preference Specification

We need to specify the functional form for the direct utility function for our estimation. Although this parametric specification is necessary, we keep it as general and flexible as possible, allowing for example, for non-separability between consumption and leisure both at the intensive and extensive margin, and for the effect of demographic variables and unobserved taste shocks to affect utility.

We start by defining the aggregator:

$$M_t = \left(\frac{\alpha_t(z_t, \chi_t)(c_t^{1-\phi} - 1)}{1 - \phi} + \frac{(1 - \alpha_t(z_t, \chi_t))(l_t^{1-\theta} - 1)}{1 - \theta} \right)$$

where z_t is a vector of observable demographic variables and the term χ_t represents ‘taste shocks’ or ‘unobserved heterogeneity’ in within period preferences. The function α_t is specified so that it is always between 0 and 1:

$$\alpha_t = \frac{1}{1 + \exp(\psi z_t + \chi_t)} \quad (9)$$

We assume that the utility function is of the form:

$$u(c_t, l_t) = \frac{M_t^{1-\gamma}}{1-\gamma} \exp(\pi z_t + \varphi P_t + \zeta_t) \quad (10)$$

where the vector of observable variables z_t appears again and ζ_t is another taste shock which affects intertemporal preferences; this is different from but not necessarily uncorrelated with χ_t . Notice that the observable variables that appear in equations (9) and (10) need not be the same. These terms

(and the two different taste shocks χ_t and ζ_t) play different roles as they operate at the intratemporal and intertemporal margins respectively.

We require that the MRS between consumption and leisure is decreasing in leisure and increasing in consumption. After estimating the relevant parameters, these conditions can be verified empirically.

2.2 Marginal Rate of Substitution and Marshallian and Hicksian Elasticities.

In a dynamic context, a Marshallian elasticity describes how hours of work within a period change holding constant the full income available within the period (defined as the value of consumption plus the value of leisure), whereas a Hicksian response conditions on utility within the period. As suggested by Keane (2009), an alternative representation of the Hicks elasticity can be given considering a tax change with a lump-sum transfer, keeping life-cycle wealth constant.⁸ In such a situation, the Marshallian elasticity would describe the change in labour supply if the tax change is not compensated. Therefore, if one wants to think about the implications for labor supply of changes in taxes, the Marshallian and Hicksian elasticities are the relevant concepts. Following the change in the structure of wages (possibly induced by changes in taxes), resources may be reallocated over time through changes to the time path of hours of work changing or through changes to the time path of the marginal utility of wealth changing, or both. The Frisch elasticity captures the change over time in hours worked in response to the anticipated evolution of wages, with the marginal utility of wealth unchanged because the wage change conveys no new information.⁹ This is then the right concept if one wants to think about the implications of changes in wages over the business cycle.

Standard two-stage budgeting imply that we can first consider the problem of allocating resources between consumption and female leisure within each period. If the optimum implies a strictly positive number of hours, the first order condition for within period optimality implies that the ratio of the marginal utility of leisure to that of consumption, that is the Marginal Rate of Substitution, equals the after tax real wage. For our specification of preferences, for $l_t < H$, this equation will be:

$$w_t = \frac{u_{l_t}}{u_{c_t}} = \frac{1 - \alpha_t}{\alpha_t} \frac{l_t^{-\theta}}{c_t^{-\phi}} \quad (11)$$

This equation is useful for computing static labour supply elasticities. Differentiating the MRS equation (11) and the budget constraint with respect to wages we obtain an expression for Marshallian elasticities for consumption and female leisure:

⁸This concept of a Hicks elasticity is used in Chetty (2012) and Keane and Rogerson (2011). It is equivalent to the static concept under the assumption that resources are freely transferable between periods and preferences are separable between consumption and leisure. Alternatively, it is equivalent if preferences are quasilinear, in which case the Marshallian, Hicksian and Frisch elasticities coincide.

⁹When a wage changes stochastically, the response of hours worked will partly be due to the Frisch intertemporal substitution motive, but will also be affected by the change in the marginal utility of wealth due to the particular wage realisation.

$$\varepsilon^M = \begin{bmatrix} \frac{\partial \ln c}{\partial \ln w} \\ \frac{\partial \ln l}{\partial \ln w} \end{bmatrix} = \begin{bmatrix} 1 & \frac{wl}{c} \\ \phi & -\theta \end{bmatrix}^{-1} \begin{bmatrix} \frac{w(H-l)}{c} \\ 1 \end{bmatrix}$$

By using the Slutsky equation, we can obtain Hicksian elasticities by adding to the Marshallian elasticities the expressions for the income elasticities

$$\begin{aligned} \varepsilon_c^H &= \varepsilon_c^M + \frac{\partial \ln c}{\partial \ln(c+wl)} \frac{wl}{(c+wl)} \\ \varepsilon_l^H &= \varepsilon_l^M - \frac{\partial \ln l}{\partial \ln(c+wl)} \frac{w(H-l)}{(c+wl)} \end{aligned}$$

where the expressions for the income elasticities can be obtained by differentiating the MRS equation and the budget constraint with respect to income:

$$\begin{bmatrix} \frac{\partial \ln c}{\partial \ln y} \\ \frac{\partial \ln l}{\partial \ln y} \end{bmatrix} = \begin{bmatrix} \phi & -\theta \\ 1 & \frac{wl}{c} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \frac{y}{c} \end{bmatrix}$$

Several facts are worth noting. First, despite their simplicity, these equations result in non-linear expressions for the elasticities that have the potential of varying greatly across consumers and do not aggregate in a straightforward way. Second, for the specification we have used, the Marshallian and Hicksian elasticities depend only on ϕ and θ (and on the values of earnings and consumption). In particular, they do not depend on the inter-temporal parameters or on whether consumption and leisure are separable in the utility function. Third, by log-linearizing equation (11), we can derive an expression that can be used to estimate the parameters needed to identify the Marshallian and Hicksian elasticities. Taking logs of the Marginal Rate of Substitution equation (11), and noticing that $\log\left(\frac{1-\alpha_t}{\alpha_t}\right) = \psi z_t + \chi_t$, we obtain:

$$\ln w_t = \psi z_t - \theta \ln l_t + \phi \ln c_t + \chi_t \quad (12)$$

As we discuss below, the first stage of our estimation process estimates this equation to identify the parameters that enter α_t (that is the vector ψ), as well as ϕ and θ . This pins down the within period elasticities. In addition, economic theory requires that Frisch intertemporal elasticities must be at least as great as Hicks elasticities. Thus, our estimates of static elasticities provide a bound on the intertemporal elasticity. This is particularly useful if there is limited data or complications in estimating Frisch elasticities directly.¹⁰

2.3 Euler equations

Having considered the intratemporal margin conditional on participation (MRS), we now characterize the intertemporal equilibrium conditions for the optimization problem in equations (1) and (2), which

¹⁰In the context of quasilinear utility as used by Chetty (2012), the Frisch elasticity collapses to equal the Hicks elasticity (and the Marshallian) because there are no wealth effects on hours of work.

are given by a set of Euler equations. While in principle we could consider either the Euler equation for hours or that for consumption, only one is relevant, when coupled with the intratemporal condition. To avoid considering interior points (and the selection problems they involve) at different points in time, which would be relevant for the Euler equation for labour supply, we focus on the Euler equation for consumption. Assuming that the household is not at a corner solution for savings, and so they are not in a situation where they cannot consume as much as they would like today because of binding borrowing restrictions, the following intertemporal condition will hold:

$$\begin{aligned} \beta(1 + R_{t+1})u_{c_{t+1}}(\cdot) &= u_{c_t}(\cdot)\varepsilon_{t+1} \\ \text{where } E[\varepsilon_{t+1}|I_t] &= 1 \end{aligned} \tag{13}$$

where I_t denotes the information available to the household at time t . The first line of (13) defines ε_{t+1} , while the second line characterizes the optimality conditions. ε_{t+1} represents the innovation to the discounted marginal utility of consumption and will incorporate innovations about present and future expected wages, male earnings and interest rates as well as the taste shifters $z_{t+1}, \chi_{t+1}, \zeta_{t+1}$. We assume that the marginal utility of consumption and the discount factor are always strictly positive, and that the real interest rate R_{t+1} is bounded away from -1, so that the support of ε_{t+1} is \mathbb{R}^+ . We can then take the log of equation (13). Taking the log of the marginal utility of consumption (and adding the superscript h to the relevant variables to denote household we have:

$$\ln u_{c_t^h} = -\gamma \ln M_t^h + \ln \alpha_t^h - \phi \ln c_t^h + \varphi P_t^h + \pi z_t^h + \zeta_t^h$$

Log-linearizing the Euler equation and rearranging we therefore get:

$$\begin{aligned} \eta_{t+1}^h &= \kappa_t^h + \ln \beta + \ln(1 + R_{t+1}) - \phi \Delta \ln c_{t+1}^h + \Delta \ln \alpha_{t+1}^h \\ &\quad - \gamma \Delta \ln(M_{t+1}^h) + \varphi \Delta P_{t+1}^h + \pi \Delta z_{t+1}^h \end{aligned} \tag{14}$$

where $\eta_{t+1}^h \equiv \ln \varepsilon_{t+1}^h - E[\ln \varepsilon_{t+1}^h | I_t^h] + \Delta \zeta_{t+1}^h$ and $\kappa_t^h \equiv E[\ln \varepsilon_{t+1}^h | I_t^h]$. This error term combines the expectation error and the taste shifters that are unobserved to the econometrician. Notice that $E[\eta_{t+1}^h | I_t^h] = 0$ by construction. We discuss the identification and estimation of the parameters of this equation in section 3.2 below. Frisch elasticities on the intensive margin can be calculated directly from the Euler equations and are given by the following expressions (the derivation is in Appendix C):

$$\varepsilon_c^F = \frac{-u_c u_{cl} w_t}{c_t(u_{ll} u_{cc} - u_{cl}^2)} = \frac{w_t \gamma \alpha_t c_t^{-\phi} l_t}{\left\{ \gamma \phi (1 - \alpha_t) l_t^{1-\theta} + \theta \gamma \alpha_t c_t^{1-\phi} + M_t \phi \theta \right\}} \quad (15)$$

$$\varepsilon_l^F = \frac{u_{cc} u_c w_t}{l_t(u_{ll} u_{cc} - u_{cl}^2)} = \frac{-\left(\gamma \alpha_t c_t^{1-\phi} + M_t \phi \right)}{\left\{ \gamma \phi (1 - \alpha_t) l_t^{1-\theta} + \theta \gamma \alpha_t c_t^{1-\phi} + M_t \phi \theta \right\}} \quad (16)$$

2.4 Returns to Experience

In our baseline specification and in the estimation of the parameters identified by the Euler equation and by the Marginal Rate of Substitution, we neglect returns to tenure and experience and assume that female wages are given by an exogenous process, as specified in equations (6), (7) and (8). If, instead, the evolution of human capital, and therefore wages, is not exogenous as in (6) but depends on past labour supply histories, rational individuals will take this into account when making their current labour supply choices. This issue has been argued to be important, for instance by Imai and Keane (2004).

If returns to experience operate only through the participation decision, rather than hours, then the use of the first order condition for hours of work (which conditions on participation) and the Euler equation for consumption (which also conditions on optimal participation) is still valid. Therefore, under this assumption, the estimation strategy that we discuss below will be valid, regardless of whether returns to experience are operational or not. If, however, the returns to experience depend on hours of work, rather than (or in addition to) the participation decision, then the MRS conditions will no longer be valid, as individuals will choose hours taking into account not only the current wage, but also the effect that current hours have on future wages. The Euler equation analysis will not be affected, except by the fact that some of the quantities we use come from the estimates of the MRS.

In section 7, we explore the possible role of the returns to experience, but only when these operate through the extensive margin. This implies that we will not need to change our empirical strategy for the analysis of the MRS and of the Euler equation. However, we will need to change our analysis of the extensive margin to take the possibility of returns to experience into consideration.

3 Empirical Strategy

Given the model we have sketched in the previous section, we use US household level data on consumption, labour supply, earnings and wages (as well as a variety of demographic variables) to estimate the model parameters. We use a variety of methods and exploit different restrictions imposed by the model on different sets of moments to estimate different sets of parameters. In this section, we discuss our empirical approach and the identification assumptions we make. We divide our discussion

into three sections, corresponding to the three sets of equilibrium conditions that we use to identify different parts of the model.

We start with a discussion of the Marginal Rate of Substitution conditions and of what parameters they can identify. We then move on to discuss intertemporal conditions and their use to estimate the parameters that determine the intertemporal elasticity of substitution. For these two steps, it is not necessary to solve the model explicitly and derive policy functions that determine consumption and leisure choices as a function of state variables. Instead, we use equilibrium conditions and some assumptions about the nature of the random variables that enter the problem (that can be either representing uncertainty faced by individuals or unobserved (by the econometrician) components of preferences).

As we discuss below, however, these conditions are not sufficient to identify all components of preferences or to characterise fully the decision rules implied by our model. To complete our exercise, therefore, we need to solve the full model. By matching certain moments of the data with the corresponding theoretical moments, we identify the parameters that could not be identified by the MRS and the Euler Equation. With the complete set of parameters we can then characterise the properties of the decision rules for all endogenous variables, including participation and hours of work.

3.1 Intratemporal margins

As mentioned in Section 2.2, standard two-stage budgeting considerations imply that, for households not at a corner, that is where the wife works, the relevant intra-temporal equilibrium condition is given by equation (12). Notice the importance of the unobserved heterogeneity term χ_t in that equation: in its absence we would have an equation with perfect fit that would obviously be rejected by the data and would imply the ad-hoc consideration of measurement error in the relevant variables.

The econometric estimation of the MRS equation poses two problems. First, the subset of households for whom the wife works and the MRS condition holds as an equality is not a random subset. This would therefore imply that the unobserved heterogeneity term χ_t would not average out to zero and would be correlated with the variables that enter equation (12). Second, even in the absence of participation issues and corner solutions, it is likely that individual wages (and consumption and leisure) will be correlated with the unobserved heterogeneity term, so that the use of OLS to estimate such an equation would result in biased estimates of the structural parameters ϕ and θ . We discuss these two issues in turn.

For participation, we use a two step procedure. We specify first a reduced form equation for the extensive margin. Having estimated such a participation equation, we use an Heckman (1979)-type selection correction approach to estimate the MRS equation (12) only on the households where the wife

works and augmenting it with a polynomial in the estimated residuals of the participation equation. Non-parametric identification requires that some variables that enter the participation equation do not enter the specification for the MRS: consistently with the model we assume that these variables are given by male earnings and male employment status.

Whilst the participation equation is consistent with our structural dynamic model, in that we model participation as a function of the state variables of the dynamic problem in equation (1), we do not solve it explicitly at this stage. Beside its simplicity, this approach has the advantage of delivering consistent estimates of the parameters of the MRS equation even when some of the details of our model are mis-specified, such as the specification of the innovation process.

The second issue in the estimation of equation (12) is that our measures of wages, which is obtained by dividing earnings by hours, might be correlated with the residual term χ_t . This could be due either to measurement error in hours or earnings or to the possible correlation between taste and productivity heterogeneity. To avoid these problems, we use an instrumental variable approach and exploit only part of the observed variability in wages to identify the parameters of interest. In particular, we use as instruments fully interacted regional, time and education groups dummies. This is equivalent to taking averages within cells defined by time periods (in quarters), region and education groups and so we exploit only the variability across these groups, rather than the individual variability. While this does mean that we use the differences between wages at different levels of education, the vector of taste shifter variables z includes education dummies, which effectively absorbs average differences in the wages of individuals with different levels of education, differences in their taste for work and taste consumption. Within each education group, the variability that we exploit is that over time and across regions.

Finally, notice that if $\gamma = 0$, then the utility function collapses to the additively separable form and the elasticity of intertemporal substitution of consumption would equal ϕ and could be estimated from the within period MRS condition alone. However, it should be stressed that we cannot test non-separability from the within period MRS alone.

3.2 Euler Equation Estimation

A natural approach to the estimation of the Euler equation (13) is GMM. However, given the nature of the data we have, all that is possible to bring to data is its log-linearized version, as in equation (14). Moreover, as discussed in Attanasio and Low (2004), the small sample properties of non linear GMM estimators can be poor when applied to Euler equations similar to that we are studying. We therefore focus on the estimation of equation (14).

The identification and estimation of the parameter of this equation depends, obviously, on the nature of the ‘residual’ term η_{t+1}^h on its right-hand-side. As noted above, η_{t+1}^h contains expectations

errors (ε_{t+1}^h) and taste shifters unobservable to the econometrician (ζ_{t+1}^h). As for the former, the rational expectations assumption that is typically invoked, implies that any variable known to the household at time t is a valid instrument. On the other hand, to achieve consistency using such an argument, it will be necessary to exploit explicitly the time series variation and, therefore, as discussed in Attanasio and Low (2004), a long time series is required to achieve consistency.¹¹

If we can use a sample that covers a large number of time periods, we then need to assume that the lagged variables that are used as instruments are uncorrelated with the innovations to the taste shifters $\Delta\zeta_{t+1}^h$. This is trivially true if individual taste shifters are constant over time or if they are random walks. In what follows we will be making this assumption, which can be in part be tested by considering over-identifying restriction tests.

The nature of the data we use, the Consumer Expenditure Survey (CEX), which we describe in Section 4, poses some additional challenges to the identification and estimation of equation (14). In particular, although the CEX covers now a substantial time period (from 1980 to 2010) over which we can consider quarterly data, as in many other household surveys, each household is only observed for a few time periods (in our case 4 quarters). Therefore, it is not possible to observe the same households over an extended time period.

For this reason, we follow a well-established tradition in the literature on the estimation of life cycle models of consumption and labour supply and use a synthetic cohort approach (see Browning, Deaton and Irish, 1985; Attanasio and Weber, 1993, 1995; Browning, Blundell and Meghir, 1994). An equation such as (14) can be aggregated over certain groups and we follow the average behaviour of the variables of interest (or their non-linear transformation) for a group of households with constant membership. A time series of cross sections can be used to construct consistent estimates of these aggregates and, in this fashion, use a long time period to estimate the parameters of the Euler equation and test its validity.

We define groups by year of birth. The assumption of constant membership of these groups might be questioned at the beginning and at the end of the life cycle for a variety of reasons, including differential rates of family formation, differential mortality and so on. To avoid these and other issues, we limit our sample to households whose husband is aged between 25 and 67 and where wives are aged between 25 and 60.¹²

Having identified groups and denoting them with the superscript g , we define as X_t^g the (population) average for group g of the variable X_t^h . We then aggregate equation (14) across households

¹¹The reason for the need of a long time series is that, even under rational expectations, expectations errors do not necessarily average out to zero (or are uncorrelated with available information) in the cross section, but only in the time series: expectation errors may be correlated with available information in the cross section in the presence of aggregate shocks. See the discussion in Hayashi (1987), Miller and Sieg (1997), Attanasio (1999), or Attanasio and Weber (2010).

¹²If credit constraints are binding, the Euler equation will not be holding as an equality. Very young consumers are excluded because they are more likely to be affected by this issue. For older consumers, in addition to changes in labour force participation and family composition, health status also changes in complex ways that maybe difficult to capture with the taste shifters that we have been considering.

belonging to group g to get:

$$\begin{aligned}\eta_{t+1}^g &= \kappa_t^g + \ln\beta + \ln(1 + R_{t+1}) - \phi\Delta \ln c_{t+1}^g + \\ &\quad \Delta \ln \alpha_{t+1}^g - \gamma\Delta \ln(M_{t+1}^g) + \varphi\Delta P_{t+1}^g + \pi\Delta z_{t+1}^g\end{aligned}\tag{17}$$

For this approach to work, however, it is necessary that the relationship one studies is linear in parameters. If M_t^h and α_t^h were observable, this would be the case for equation (17). However, both M_t^h and α_t^h are non linear functions of data *and* unobserved parameters, so that, in principle they cannot not be aggregated within groups to obtain M_t^g and α_t^g .

A solution to this problem uses the fact that the parameters that determine M_t^h and α_t^h can be consistently estimated, as discussed in Section 3.1, using the MRS conditions. Given these consistent estimates of the parameters that enter M_t^h and α_t^h , one can construct consistent estimates of these variables and, effectively, treat them as data. This is the procedure we use in what follows.

Finally, we need to consider the fact that the quantities that enter equation (17) are population means of the relevant variables and, as such, are not directly observable. However, we can obtain consistent estimates of these quantities from the time series of cross section that we have. We can therefore substitute these observable quantities and obtain:

$$\begin{aligned}\tilde{\eta}_{t+1}^g &= \bar{\kappa} + \ln\beta + \ln(1 + R_{t+1}) - \phi\Delta \overline{\ln c_{t+1}^g} + \\ &\quad \Delta \overline{\ln \hat{\alpha}_{t+1}^g} - \gamma\Delta \ln(\overline{\hat{M}_{t+1}^g}) + \varphi\Delta \overline{P_{t+1}^g} + \pi\Delta \overline{z_{t+1}^g}\end{aligned}\tag{18}$$

The residual term $\tilde{\eta}_{t+1}^g$ now includes, in addition to the average of the expectation errors and of the changes in taste shifters, several other terms. In particular, it includes: (i) a linear combination of the difference between the population and sample averages at time t and $t + 1$ for all the relevant variables (induced by the fact that we are considering sample means rather than population means for group g); (ii) the difference between the (consistently) estimated M_t^g and α_t^g and their actual value (induced by estimation error in the parameters of the MRS); (iii) the difference between the innovation over time to the average value of κ_t^g , which we have denoted with the constant $\bar{\kappa}$.

All the variables on the right hand side of equation (18) are observable. We can therefore use this equation to estimate the parameters of interest. However, care has to be taken to choose the instruments so that they are plausibly uncorrelated with $\tilde{\eta}_{t+1}^g$.¹³

¹³As noted by Deaton (1985) and discussed extensively in the context of the CEX by Attanasio and Weber (1995), the use of sample rather than population averages for all the ‘group’ variables induces an MA(1) in the residuals, induced by the sampling variation in the rotating panel structure. We need to assume that the instruments are not correlated with the (average) estimation error of the M_t^h ’s and α_t^h or with the innovations to the higher moments of the expectation errors ($\kappa_t^g - \bar{\kappa}$). This last assumption is discussed in Attanasio and Low (2004). In the Appendix, we discuss some of the sample selection choices to avoid some of the problems caused by the CEX.

While the assumptions we make guarantee that the appropriate choice of instruments yield consistent estimates, the covariance structure of the $\tilde{\eta}_{t+1}^g$ is quite complex. The contemporaneous covariance of $\tilde{\eta}_{t+1}^{g_i}$ and $\tilde{\eta}_{t+1}^{g_j}$ is not, in general zero, as aggregate shocks will have effects that correlate across the various groups. We should take this structure into account when computing the variance covariance matrix of the estimates, if not to improve their efficiency. Whilst it is in principle possible, given our assumptions, to construct an estimate of the variance covariance of $\tilde{\eta}_{t+1}^g$ from the estimated parameters, in practice this turns out to be cumbersome, as there is no guarantee that, in small sample, these estimates are positive definite. Given these difficulties, we decided to follow a different and, as far as we know, novel approach based on bootstrapping our sample, with a structure consistent with the basic assumption of our model. We describe the bootstrapping procedure in detail in Appendix B.

3.3 Extensive margins

One of the main goals of this paper is to characterise the labour supply reaction to wage shocks at the extensive as well as the intensive margin. And one could argue that the extensive margin is particularly important, as in the presence of fixed costs of participation and other non-convexities, it might generate a considerable amount of action and, therefore, be particularly salient for evaluating the size of ‘aggregate’ labour supply elasticities (as argued by Rogerson and Keane (2012) among others).

However, when considering the extensive margin, it will be necessary to solve explicitly the dynamic problem we have been considering. This involves specifying completely the economic environment the individual households live in, including both present and future conditions (at least as perceived by the household). Moreover, often it will be impossible to obtain a closed form solution for the policy rules or general results for the elasticities of interest. It will therefore be necessary to solve the model numerically and estimate or calibrate its parameters using the properties of the solutions so obtained.

In what follows, we use a number of life cycle facts and match similar moments computed by simulating our model to obtain the missing parameters of the model. Not all the parameters of the model will be calibrated, however. First, we will use the estimates of the other parameters that we obtained from the MRS and the Euler Equation. Second, some parameters will be taken from other sources: either the literature or auxiliary regressions. Armed with these parameters we will be simulating the model we have constructed for a large number of individuals to study the properties of individual and ‘aggregate’ labour supply.

To obtain the calibrated parameters of our model, we target the labour supply behaviour of women born in the 1950s. We assume one model period is one quarter. We assume individuals leave for 50 years, the last 10 in retirement. We assume 15% of women are childless (see OECD Family Database).

We assume there are two different groups of mothers, young mothers who have their first child at the age of 23 and old mothers who have their first child at the age of 28. In order to target the US average age at first child arrival of 25, we assume there are 60% of mothers in the first group and 40% in the second. Second child arrives 2 years after the first. We describe in detail the moments we target and the results we get in our result section.

Goodness of fit Having obtained all our estimates, we simulate the model and check whether it is able to fit several features of the data, over and above those that have been used to derive the parameter estimates (either by econometric methods or by calibration). In particular we explore: participation and hours life-cycle profiles, participation rates conditioning on several characteristics such as motherhood and the distribution of hours worked.

4 Data and Descriptive Statistics

We use data from the Consumer Expenditure Survey (CEX) for the years 1980-2003.¹⁴ The CEX includes detailed recall questions on household expenditures as well as some information on the assets, demographics, incomes and labour supply of household members. Households can be followed for up to four quarters.

Our definition of consumption covers nondurable goods excluding medical and education spending. While we are able to tell whether an individual earns an income in the current quarter or not, labour supply and income questions in the CEX typically cover the previous 12 months. To obtain quarterly hours worked we therefore divide by four the product of average hours worked per week when working over the past year and the number of weeks worked over the past year. Hours of leisure are then given by 1250 minus quarterly hours. Net wages are calculated by dividing annual salary income by annual hours (and dividing by four), and then subtracting marginal federal income tax rates generated using the NBER TAXSIM model (Feenberg and Coutts, 1993).¹⁵ We deflate all expenditures, wages and incomes using the Consumer Price Index for the appropriate period.

Our sample consists of couples where the female is aged between 25 and 60 and males are aged between 25 and 67. Labour supply and income data are only collected in the first and final interviews of the CEX unless a member of the household changes their employment status. We therefore restrict our sample to households interviewed in the first interview for our estimation. We use information from these households' fifth interviews to calculate growth rates and transitions for our calibration. We drop those in rural areas and those in the top and bottom 2.5% of the distribution of hours

¹⁴We stop at 2003, as income imputation was introduced to the data from 2004 onwards (and the original non-imputed variables were only reintroduced in 2006).

¹⁵We are grateful to Lorenz Kueng for making his mapping of the CEX to TAXSIM publically available.

Table 1: Descriptive statistics

	Mean	Standard Deviation		Mean	Standard Deviation
Consumption	8510	3662.15	Husband's age	42.6	10.52
% Wives employed	0.69	0.46	Wife's net wage	15.6	7.12
% Husbands employed	0.90	0.30	Husband's net wage	20.6	8.63
Wife's hours	432	152.05	Number of children	1.1	1.22
Husband's hours	546	131.13	Number of adults	2.4	0.87
Wife's age	40.2	9.75			

Note: Monetary values expressed in 2011 \$

(conditional on participation) or consumption. We also drop the top 2.5% of wages and those who are seen to earn less than 3 quarters of the national minimum wage in any given year. This leaves us with a sample of just over 27,500 households. Interest rates are for 3 month Treasury Bills and are taken from the Federal Reserve Bank of St Louis. Table 1 presents some summary statistics including on hours, consumption and wages in our data.

5 Parameter Estimation and Calibration.

In this section, we report the estimates we obtain using the econometric techniques discussed in Section 3. In the first two sub-sections we report the estimation results obtained used the MRS conditions and the Euler equation, respectively, while in the third, we discuss our calibration results. In the last subsection, we show how well the model matches additional statistics.

5.1 MRS estimates

In Table 2, we report the estimates of key parameters for the MRS equation:¹⁶

$$\ln w_t = \psi z_t - \theta \ln l_t + \phi \ln c_t + v_t$$

We estimate a value of ϕ of 0.43 and of θ of 0.87. Both are less than one and so satisfy the concavity requirements of the utility function. A standard CES specification imposes $\phi = \theta$, which is rejected by our estimates. A Cobb-Douglas specification imposes that $\phi = \theta = 1$ which is also rejected.

Table 2 also shows the coefficients attached to variables included in z_t , reflecting the impact of some demographic variables on the MRS. We report the coefficient on the number of children of various ages and on family size. A positive coefficient on one of these variables implies that women will supply less hours of work in the market, for a given level of consumption and wages, when this

¹⁶The results for the probit model for participation are reported in Appendix A.

Table 2: Estimation of MRS equation

Parameter	Estimate	(Standard Error)
θ	0.87***	(0.120)
ϕ	0.43***	(0.035)
Ψ		
$\ln(famsize)$	-0.234***	(0.018)
kids 0-2	0.138***	(0.014)
kids 3-15	0.023**	(0.007)
kids 16-17	-0.001	(0.010)
e_1	0.077	(0.049)
e_2	0.0934*	(0.044)
e_3	0.0357	(0.029)

Note: N = 17,852, standard errors in parentheses

*p<0.05, ** p<0.01, *** p<0.001

Additional controls for elderly person in the household,
a cubic in age, race, education, region and season

variable increases.¹⁷ The coefficient on young children aged 0-2 is positive and highly significant. The variable measuring children aged 3-15 also attracts a positive (albeit smaller) coefficient, while the coefficient for older children is not statistically different from zero. The log of family size attracts a negative and significant coefficient.

Finally, Table 2 reports the coefficients on the estimates of the first three moments of the residuals conditional on positive participation, calculated using the estimates of the probit for participation. These coefficients are jointly significant (p-value = 0.037), indicating that it is important to take into account selection in obtaining the estimates of the MRS coefficients.

5.2 Euler Equation estimates

As discussed above, we use the parameters obtained from estimating the MRS condition to calculate cohort average values of the logs of M_t and α_t for different time periods, and this gives the variables we need to estimate our Euler equation. We calculate α for each individual by evaluating $1/(1 + \exp(\psi_z z_{i,t} + \chi_i))$, where χ_i is the residual from the MRS equation.¹⁸

¹⁷A positive coefficient means that the marginal utility of leisure must be lower, and this in turn means hours of leisure must be higher.

¹⁸This must also be calculated for non-participants for whom we do not have estimates of the MRS residuals. we do this by imputing wages to those out of work using a regression of wages on family characteristics and region dummies, calculating a lower bound on what this would imply for their residuals given our MRS coefficients and their non-participation, and then adjusting these residuals such that for all participants and non-participants $E[v_i] = 0$. Once we have obtained α_i the calculation of M_i is straightforward.

Table 3 shows results we obtain from estimating the Euler equation (18). We estimate γ at 2.64, a value that given the precision of our estimates, is significantly different from zero, implying that preferences are non-separable and that consumption and leisure are substitutes.¹⁹ The coefficients on the control variables included in the vector z_t are imprecisely estimated and are not significantly different from zero. In what follows, we impose that the parameter ξ and the coefficient on having children aged 16-17 are both zero. A Hansen J test of overidentifying restrictions (Hansen, 1982) fails to reject the null hypothesis at the 5% significance level with a p-value of 0.083.

Table 3: Estimation of Euler equation

Parameter	Estimate	Standard Error	[95% Confidence Interval]
γ	2.638**	0.774	[1.505, 4.671]
$\bar{\kappa} + \ln(\beta)$	-0.038	0.504	[-1.506, 0.635]
π			
$\ln(famsize)$	-0.004	0.681	[-1.813, 0.740]
kids 0-2	0.353	0.340	[-0.396, 0.988]
kids 3-15	-0.175	0.246	[-0.208, 0.770]

Hansen J-statistic P value = 0.083

N = 835, standard errors in parentheses, *p<0.05, ** p<0.01, *** p<0.001

Note: Additional controls for season dummies. Instruments are second, third and fourth lags of the logs of consumption and M_t , and first, second, third and fourth lags of the logs of the interest rate, leisure, and α_t

5.3 Calibration of the remaining parameters.

As discussed in Section 3.3, to estimate the responsiveness on the extensive margin, we need to specify all the details of the model and quantify each element of the model. This need to specify the full model in order to identify the extensive margin is in contrast to the lower informational requirement needed to identify intensive parameters. There are three sets of parameters in the calibration: those estimated via the MRS conditions and the Euler equation, those coming from external sources and those that we calibrate using the full model.

External Parameters. Table 4 reports the estimated and external parameters used in the

¹⁹A value of 0 for γ would imply additive separability in preferences over consumption and leisure. For cases when θ , $\phi > 0$ (as we have here), a value of γ less than zero would imply that leisure and consumption are complements, and a value of γ greater than zero would imply that consumption and leisure are substitutes.

calibration. The first panel reports the estimated parameters from Tables 2 and 3 above. The second panel reports parameters which come from external sources.

Table 4: External Parameters

Estimated Parameters (from first-order conditions)		
Curvature on leisure	θ	0.87
Curvature on consumption	ϕ	0.43
Curvature on utility	γ	2.64
Exogenous Parameters		
Discount Factor (annual)	β	0.98
Interest Rate (annual)	R	1.015
Regression Log Wage on Age and Age ² (Men)	α_1^m, α_2^m	0.05991, -0.00064
Husband and Wife Wage Correlation	ρ	0.25
Standard Deviation of Permanent Shock (Men)	σ_{ξ^m}	0.077
Standard Deviation of Permanent Shock (Women)	σ_{ξ^f}	0.063
Standard Deviation of Initial Wage (Men)	$\sigma_{\xi^m,0}$	0.447
Standard Deviation of Initial Wage (Women)	$\sigma_{\xi^f,0}$	0.387
Length of Life (in years)	T	50
Length of Working Life (in years)	r	40

We fix the annual interest rate to equal the average real return on three monthly T-bill at 0.015, and set an annual discount factor equal to 0.98. This implies a discount rate slightly higher than the interest rate. The deterministic component of the male earnings process is estimated from the CEX: we take the two parameters of a regression of husband log earnings on age and age squared. Both the innovations to male earnings and those to female wages are assumed to have a unit root, consistent with the evidence on men produced by MaCurdy (1983) and Abowd and Card (1989). The standard deviation of the innovation for husband's earnings is assumed to be 0.077, consistent with Hugget, Ventura and Yaron (2011) and Meghir and Pistaferri (2004). Furthermore, we assume an initial standard deviation of husband earnings of 0.447 as measured in the CEX. There is limited evidence on the variability of female wages and/or earnings. In contrast with men, this statistic is highly affected by non-random self-selection into the labour market. We set the initial standard deviation of wages from the CEX equal to 0.387. We set the standard deviation of female wages innovations to 0.063, which is consistent with the increase in the variance of wages over the life-cycle for women born in the 1950s. We assume that the correlation coefficient between the two shocks (for husband and wife) is equal to 0.25 as estimated by Hyslop (2001).

As in Attanasio et al. (2008), there are two components to child care costs: the function $G(a_t)$ and

the price p . We estimate the function $G(a_t)$ directly from data. In particular, for households where the mother is working, we regress total childcare expenditure on the age of the youngest child, the age of the oldest child, the number of children and a dummy that equals one if the youngest child is 0. The shape $G(a_t)$ can be derived from the coefficients of this regression function, considering that in our model all women with children have two of them and at the same interval between children of two years.²⁰ This implies that the child care cost can be expressed as a function of the age of the oldest child. Finally, we assume that the household receives a pension equal to 70% of the husband's earnings in the final working period.

Table 5: Baseline economy: Calibrated Parameters and Targets

Parameters		Value
Childcare Cost	p	51
Fixed Cost of Working	\bar{F}	17
Offered Wage Gender Gap	y_0^f/y_0^m	0.72
Constant term weight of consumption	ψ_0	2.91
Exogenous wage growth	α_1	0.052
Exogenous wage growth	α_2	-0.0004

Targets	Data	Model
Participation Rate	0.714	0.704
Participation Rate of Mothers	0.550	0.549
Observed Wage Gender Gap	0.770	0.757
Wage Growth (if younger than 40)	0.0192	0.0203
Wage Growth (if older than 40)	0.0142	0.0107
Hours worked	432	436

Statistics for individuals aged 25 to 52. Wage growth is over 3 quarters.

Calibrated parameters. There are six parameters that we calibrate within our decision model and that relate to the participation decision: the fixed cost of working, \bar{F} ; the price of child care, p ; the offered wage gender gap, y_0^f/y_0^m ; two parameters that determine exogenous wage growth; and the ‘constant term’ of the $\alpha(\cdot)$ function in the CES utility which determines, together with a set of demographics, the weight of consumption in the utility function, parameter ψ_0 . In order to identify these parameters we target statistics from the cohort born in the 1950s: the female participation rate, the participation rate of mothers, average hours worked, the observed wage gender gap, and observed

²⁰Our estimate of $G(a_t)$ combines the cost of the first born child along with any subsequent costs associated with additional children who are born later. In this way, any economies of scale in child costs will be captured by $G(a_t)$, but we do not identify separately the marginal cost of extra children.

wage growth at two different stages of the life-cycle.

In Table 5, we report the value of the parameters we obtain in this calibration exercise as well as the value of the targeted moments in the data and in the simulated data. Both the monetary fixed cost of working and the monetary fixed childcare cost are small compared to household earnings.

5.4 Goodness of fit

Our next step is to show to what extent the model can account for observed female labor supply behaviour that was not targeted in the calibration. The calibration was focused on averages taken over the life-cycle. Our focus here is on life-cycle paths. Figure 1 shows life-cycle profiles in the simulations and in the data and Table 6 reports additional moments showing the heterogeneity and distribution of behaviour.

Figure 1: Life-Cycle Profiles: Baseline Model (solid black line) versus Data (dashed red line)



The life-cycle path of female labor supply both at the extensive and intensive margin is similar in the model and in the data, except at the early part of the life-cycle when the model underestimates hours and participation. This underestimation is driven by women without children who work more in the data than in the simulations. Observed female wages and the variance of wages are increasing with age in our simulations, consistent with what we observe in the data. The evolution of the wage gender gap over the life-cycle is not reported but is stable over the life-cycle in both the simulations and the data. The profiles shown are shaped not only by our assumptions on the wage process, but

also by the selection of women into the labor market.

The distribution of hours worked is close to the data, except that the model implies higher hours worked for those women at the 90th percentile of the hours distribution. Furthermore, the fraction of women working 520 hours (an average of 40 hours a week) is higher in the data. The distribution of observed wages in the simulations is similar in the model and in the data.

Table 6: Statistics on Heterogeneity

	Data	Model
Participation Rate Young Mothers	0.487	0.526
Participation Rate Old Mothers	0.618	0.559
Participation Rate Mothers with Children Aged 3-18	0.714	0.726
Participation Rate Childless Women	0.840	0.710
Average Hours Worked 10th Percentile	189	168
Average Hours Worked 25th Percentile	330	281
Average Hours Worked 50th Percentile	520	441
Average Hours Worked 75th Percentile	520	593
Average Hours Worked 90th Percentile	585	713
Median Duration of Spells (years)		1.8
Wage 10th Percentile	7.96	8.97
Wage 50th Percentile	14.18	14.73
Wage 90th Percentile	25.42	27.24

Finally, we perform two additional exercises to compare the correlations in the data to those observed in the data. First, we use a simulated sample to reestimate the MRS equation, employing the same procedure used in getting our estimates from the data and described in section 2.2. Second, we estimate a probit model for female labour force participation as a function of husband earnings and demographics, both on simulated and actual data

The estimates of the MRS parameters θ and ϕ that we obtained from actual data (and that were used to generate the simulated data) are almost identical to those we recover from the simulated data. Given the complexity of the model that includes discrete choices over the life cycle, it is an important validation of our strategy that we are able to recover the MRS parameters from the simulated data.

In Table 7, we report the marginal effects of the probit model for participation decision obtained from actual data and simulated data. Although the correlation between wife's employment and husband earnings is higher in the model (the marginal effect being -0.14) than in the data (the marginal effect being -0.06), it should be noted that our specification of preferences helps to produce a much closer correlation to the data than the one implied by standard preferences, such as those in

Table 7: Probit of the Employment Decision

	Data	Model
$\ln(\text{husband earnings})$	-0.0591 (0.0085)	-0.1442 (0.0017)
kids 0-2	-0.2006 (0.0102)	-0.1397 (0.0022)
kids 3-15	-0.0868 (0.0046)	0.0065 (0.0011)
kids 15-17	0.0032 (0.0157)	0.0919 (0.0028)

Coefficients are marginal effects. Standard errors in brackets.

6 Labour Supply Elasticities

In this section, we use the estimates of the model to discuss implications for various wage elasticities. We start our discussion with the Marshallian and Hicksian elasticities that can be obtained from the MRS parameters. We then move on to the Frisch elasticities at the intensive margin. We then simulate the model to obtain elasticities at the extensive margin. In the final subsection, we look at aggregation issues and discuss what are the implications of our estimates for aggregate labour supply elasticities.

6.1 Marshallian and Hicksian Hours Elasticities

The first two panels in Tables 8 and 9 show how the MRS parameters translate into within-period Marshallian and Hicksian wage elasticities for hours of work, for leisure and for consumption. These elasticities vary according to family characteristics and the levels of consumption and leisure. Table 8 reports elasticities at different percentiles of the distribution of Marshallian elasticities to highlight the heterogeneity in the elasticities, while Table 9 shows them at different percentiles of the distribution of consumption to indicate how elasticities differ for households with different levels of welfare.

The median Marshallian hours elasticity is 0.70. As theory predicts, Hicksian elasticities are always greater than Marshallian elasticities: for the household with the median Marshallian elasticity, the Hicksian hours elasticity is around 50% larger at 1.08. We also notice that Marshallian elasticities are positive across the distribution, implying an upwardly sloping labour supply function with no evidence of a backward-bending supply curve.

²¹If we add as an additional regressor lagged employment (3 quarters) the marginal effect of husband earnings decreases both in the simulated sample (to -0.05) and in the data sample (to -0.02).

These estimates of the elasticities, and especially the Hicksian elasticities, are larger than found elsewhere in the literature. There are several reasons for this finding: first, the specific functional form for the utility function which allows for nonseparabilities between consumption and hours of work; second, the specific variability which is used to identify the parameters of the Marginal Rate of Substitution; and third, the explicit use of consumption data in estimating the elasticities through the Marginal Rates of Substitution. All these features of our exercise result in relatively large elasticities, which is similar to that found by Ziliak and Kniesner (2005). More recently, Blundell et al. (2015) also report elasticities that are similar in size to what we obtain.²² The importance of using consumption data is that this imposes consistency in the data such that as wages change, either hours of work change, consumption changes or savings change. Estimation strategies that do not impose this consistency can result in estimates of the elasticity of hours of work which are very low because the implied adjustment of consumption is unconstrained by the data.

There is a large variation in elasticities in the cross section. The interquartile range of the Marshallian hour elasticity is 0.66 (from 0.45 to 1.11) and Hicksian elasticities increase with Marshallian elasticities. Finally we notice a considerable amount of heterogeneity in the size of elasticities as the level of non durable consumption varies. Those with the highest levels of consumption make labour supply decisions that are the most responsive to wage changes, and make consumption decisions that are the least responsive to wage changes.

Table 8: Wage Elasticities at Percentiles of Marshallian distribution

Percentiles:	Marshallian			Hicksian			Frisch		
	25th	50th	75th	25th	50th	75th	25th	50th	75th
Hours Worked	0.45 (0.113)	0.70 (0.158)	1.11 (0.274)	0.90 (0.143)	1.08 (0.194)	1.34 (0.313)	1.22 (0.136)	1.35 (0.171)	1.50 (0.245)
Leisure	-0.58 (0.143)	-0.44 (0.102)	-0.30 (0.0660)	-0.83 (0.178)	-0.71 (0.128)	-0.55 (0.0775)	-0.95 (0.144)	-0.90 (0.113)	-0.84 (0.083)
Consumption	1.14 (0.112)	1.42 (0.125)	1.70 (0.145)	1.64 (0.0913)	1.96 (0.107)	2.19 (0.138)	0.40 (0.124)	0.51 (0.158)	0.63 (0.21)

Note: Standard errors calculated for individuals at quantiles (as opposed to quantiles themselves)

6.2 Frisch hours elasticity

We use the estimates we obtain from the the Euler equation reported in Section 5.2 to estimate the Frisch elasticities with respect to wages (at the intensive margin). Notice that these elasticities can

²²For female, using PSID data and a completely different approach from ours, they report a Marshallian elasticity of 0.64 and a Frisch elasticity of 1.43

Table 9: Wage Elasticities at Percentiles of Consumption distribution

Percentiles:	Marshallian			Hicksian			Frisch		
	25th	50th	75th	25th	50th	75th	25th	50th	75th
Hours Worked	0.70 (0.159)	0.66 (0.147)	1.22 (0.291)	1.05 (0.197)	1.07 (0.181)	1.47 (0.331)	1.28 (0.169)	1.40 (0.165)	1.66 (0.263)
Leisure	-0.50 (0.113)	-0.42 (0.093)	-0.71 (0.169)	-0.75 (0.140)	-0.68 (0.115)	-0.86 (0.193)	-0.91 (0.121)	-0.88 (0.105)	-0.97 (0.153)
Consumption	1.32 (0.123)	1.47 (0.130)	0.89 (0.117)	1.82 (0.104)	1.99 (0.114)	1.19 (0.102)	0.48 (0.144)	0.54 (0.164)	0.37 (0.110)

Note: Standard errors calculated for individuals at quantiles (as opposed to quantiles themselves)

be obtained directly from the Euler equation using equations (15) and (16). These are shown in the right hand panels in Tables 8 and 9. Table 8 reports estimates of the Frisch elasticities with respect to the wage rate for hours of work, leisure and consumption at different points in the distribution of the Marshallian elasticity, and Table 9 reports estimates at different points in the distribution of consumption.

The Frisch elasticity for hours of work is larger than the Hicksian elasticity, as theory would predict. The elasticity varies with the Marshallian elasticity and it is larger for those with higher values of consumption, rising from 1.28 at the 25th consumption percentile to 1.66 at the 75th percentile. These estimates of the elasticities are larger than those found elsewhere in the literature.

For consumption, the elasticity of consumption with respect to wages varies in a non-monotonic fashion along the distribution of consumption, being smaller at the 25th and 75th percentile than at the median. At the median consumption level it takes the value of 0.54.

Part of the heterogeneity we observe in the Frisch elasticities is due to differences across the life-cycle, but much of the heterogeneity is due to differences in the level of hours of work. The most responsive individuals are those who are working relatively few hours. This greater responsiveness is observed for young and older women who are working few hours.

6.3 The Extensive Margin and Aggregate Elasticities

The focus of the previous two subsections was on how responsive individuals' decisions over hours worked are to wage changes. This subsection reports on how responsive individuals' decisions about whether or not to participate are to wage changes. We approach this question by asking how much the percentage of women who work changes as the wage changes. We then combine this extensive measure with the intensive measure to show how total labour supply (ie. total hours worked) changes

as the wage changes. This is what we call the “macro elasticity”.

Table 10 summarises these responses for women of different ages. Each of the three columns in the table corresponds to the elasticities at a different age (26,36 and 46). The first row in the Table refers to the ‘extensive margin’ elasticity and represents the percentage of women who are shifted from non working to working as a consequence of an anticipated 5% increase in real wages that persists for 1 year. Rows 2 to 4 report the hours (intensive margin) elasticity for the same change in wages, with the different rows showing the distribution of the elasticity. Finally, in the last row, we aggregate explicitly the responses of all women to report what we label the ‘macro’ elasticity: this is the change in the total number of hours worked, considering both intensive and extensive margins.

Table 10: Labor supply elasticities in baseline economy

	age 26	age 36	age 46
Extensive	0.85	0.80	0.67
Intensive Margin 25 th percentile	1.30	0.85	0.81
Intensive Margin 50 th percentile	2.08	1.37	1.29
Intensive Margin 75 th percentile	3.87	2.58	2.45
Macro Elasticity	2.55	1.72	1.51

Elasticities are calculated by comparing labour supply in two economies where the difference between them is that wages in one year in one of the economies are 5% higher than in the other. This difference generates differences in participation rates, differences in hours worked for participants and differences in total labour supply. These differences are converted into elasticities and reported in the table. The different percentiles are percentiles of the distribution of elasticities defined by age.

Younger women are more elastic to wage increases at the intensive and extensive margins and in aggregate. One reason may be that they face more uncertainty and have less assets than older women. The degree of heterogeneity is considerable: for instance in the case of hours, the elasticity at the median goes from 2.08 for the 26 years old to 1.29 for the 46 years old; for participation, it goes from 0.85 to 0.67. This heterogeneity means the reporting of a single “elasticity” does not make sense: the effect of wage changes will depend crucially on whose wages are changing. The other notable feature of this table is that the extensive margin elasticity is smaller than the intensive margin elasticity.²³

Finally, we explore the elasticities at the age of 26, dividing women according to their maternity type, Table 11. We find that those who are *young mothers* are more elastic at the extensive margin than childless women or *older mothers*. Young mothers are the group of women that at the age of 26 are less attached to the labor market because of the fixed childcare cost they face and so are most

²³We consider whether this conclusion is robust to a scenario where the amount of idiosyncratic uncertainty is considerably higher. The extensive elasticity is somewhat lower in this economy, but the intensive elasticity is fairly similar. Results available on request.

responsive to wage changes. As in Table 10, one key message here is that there is no single elasticity which captures behaviour.

Table 11: Labor supply elasticities across maternity groups

	All	Childless	Young Mother	Older Mother
Extensive	0.85	0.73	0.97	0.80
Intensive Margin <i>25th</i> percentile	1.30	1.29	1.29	1.37
Intensive Margin <i>50th</i> percentile	2.08	2.15	1.99	2.10
Intensive Margin <i>75th</i> percentile	3.87	3.87	3.59	3.95
Macro Elasticity	2.55	2.46	2.71	2.48

Elasticities are calculated by comparing labour supply in two economies where the difference between them is that wages in one year in one of the economies are 5% higher than in the other. This difference generates differences in participation rates, differences in hours worked for participants and differences in total labour supply. These differences are converted into elasticities and reported in the table. The different percentiles are percentiles of the distribution of elasticities defined by maternity type.

6.4 Aggregate shocks: elasticities in recessions and booms

In the previous tables, we have shown that there is a substantial amount of heterogeneity in elasticities in the cross section. This heterogeneity is driven both by the assumptions we have made for the utility function and heterogeneity in variables that determine heterogeneous responses (such as the level of consumption or age). This also highlights that differences in the economic environment will lead to differences in the estimated elasticity for the same underlying preference parameters, as also discussed by Rogerson and Keane (2012). This issue is likely to be relevant particularly for the extensive margin, which is driven by non-convexities in the dynamic problem, such as fixed costs of going to work. If these non-convexities are important, it is likely that a certain sequence of aggregate shocks will tend to bunch (or disperse more) households around the kinks that determine the extensive margin response. As a consequence, different distributions of the state variables will trigger different responses in the aggregate. In particular, whether an economy is in a boom or a recession may well affect labour supply elasticities.

In Table 12, we analyse the labour supply responses of women aged 26 and 36 to deterministic changes in wages at different points of the business cycle to highlight how the state of the economy affects Frisch labour supply responses. In the simulation used to derive these tables, we define a recession as a situation in which all men and women receive an unexpected negative earnings shock for four consecutive quarters. Analogously, an expansion is a situation in which all men and women receive an unexpected positive earnings shock during four consecutive quarters. These wage changes are to

the permanent wage and will affect the marginal utility of wealth as well as changing intertemporal incentives. In the context of these different stochastic realisations of the aggregate state, we compare the labour supply response to a deterministic change in the wage and report the responses in Table 12. In other words, we report how individuals respond to intertemporal incentives in booms compared to recessions.

The key finding in table 12 is that elasticities are substantially higher in recessions than in the baseline and slightly higher in the baseline than in booms. There are differences across ages: at age 36, a boom decreases the extensive margin elasticity much more than at age 26.

Table 12: Alternative economies: Labor supply elasticities over the Business Cycle

Elasticities at age 26			
	Recession	Baseline	Boom
Extensive	1.03	0.85	0.84
Intensive Margin 25th percentile	1.35	1.30	1.25
Intensive Margin 50th percentile	2.18	2.08	2.03
Intensive Margin 75th percentile	3.86	3.87	3.59
Macro Elasticity	2.81	2.55	2.40

Elasticities at age 36			
	Recession	Baseline	Boom
Extensive	0.91	0.80	0.63
Intensive Margin 25th percentile	0.91	0.85	0.81
Intensive Margin 50th percentile	1.44	1.37	1.37
Intensive Margin 75th percentile	2.72	2.58	2.52
Macro Elasticity	1.86	1.72	1.61

Elasticities are calculated by comparing labour supply in two economies where the difference between them is that wages in one year in one of the economies are 5% higher than in the other. This difference generates differences in participation rates, differences in hours worked for participants and differences in total labour supply. These differences are converted into elasticities and reported in the table. The different percentiles are percentiles of the distribution of elasticities defined by age.

7 Returns to Experience

In this section, we consider an alternative framework in which returns to experience operate at the extensive margin; that is, we assume that the returns to experience are not affected by the number of hours worked but only by the decision to participate. In particular, we assume that human capital accumulates according to the following process:

$$\ln h_t^f = \ln h_{t-1}^f + (\eta_0 + \eta_1 h_{t-1}) I(P_{t-1} = 1) - \delta I(P_{t-1} = 0)$$

As we mentioned above, in this case the estimates of the MRS and Euler equations remain valid. However, we need to change the solution for the discrete choices.

We begin by recalibrating the parameter values that were chosen in the baseline economy to fit some of the features of participation: the fixed cost of working, \bar{F} , child care price, p , the offered wage gender gap and ψ_0 . In addition to these parameters, we also need to calibrate the two parameters that characterize human capital accumulation function and its depreciation rate. In order to identify all these parameters we target the female participation rate, the participation rate of mothers, the average hours worked, the observed wage gender gap, the observed wage growth at two different stages of life, and the observed depreciation of wages during non-participation. Note that the value of the statistics on wages are shaped by selection so we need to identify the underlying parameters by solving the model. We report the implied parameters in Table 13.

In order to assess the ability of the model to account for female labor supply behaviour we provide several statistics beyond the targets of the calibration. First, analogously to Figure 1, Figure 2 shows life-cycle profiles in the simulations and in the data. Second, Table 14 reports some additional statistics.

There are two main differences between the model with returns to experience and the baseline we considered above. First, with returns to experience, 8.8% of workers are at the corner solution, working the minimum hours possible per quarter, and yet obtaining the return to experience. Second, the median duration of spells out of the labour force is much longer: those who do exit, exit for long periods or do not return. This can be seen in the declining participation profiles at ages beyond 35. These patterns are not observed either in the data or in the baseline model. Further, very few women change their participation decisions. For example, the fraction of women who worked in all previous periods at the age of 52 is 50.2%, which compares to 42.8% in the economy without returns to experience.

In table 15, we report the aggregate labour supply response in the economy with returns to experience. The key finding is that, in contrast to the economy without returns to experience, the extensive margin elasticity is essentially zero. In this economy, there is a strong incentive to participate to obtain the return to experience. The calibrated fixed cost of participation is therefore larger in

Table 13: Returns to experience: Calibrated Parameters and Targets

Parameters		Value
Childcare Cost	p	2,050
Fixed Cost of Working	\bar{F}	130
Offered Wage Gender Gap	y_0^f/y_0^m	0.69
Female Human Capital Tech	η_0	0.03
Female Human Capital Tech	η_1	-0.018
Depreciation rate	δ	0.014
Constant term weight of consumption	ψ_0	2.89

Targets	Data	Model
Participation Rate	0.714	0.707
Participation Rate of Mothers	0.550	0.547
Observed Wage Gender Gap	0.767	0.770
Wage Growth (if younger than 40)	0.0192	0.0196
Wage Growth (if older than 40)	0.0142	0.006
Observed Depreciation Rate	-0.050	-0.050
Hours worked	432	450

Statistics for individuals aged 25 to 52. Wage growth is over 3 quarters and depreciation is annual.

Table 14: Returns to Experience: Other Statistics

	Data	Model
Participation Rate Young Mothers	0.487	0.510
Participation Rate Old Mothers	0.618	0.564
Participation Rate Mothers with Children Aged 3-18	0.714	0.699
Participation Rate Childless Women	0.840	0.743
Average Hours Worked 10th Percentile	195	25
Average Hours Worked 25th Percentile	339	291
Average Hours Worked 50th Percentile	520	483
Average Hours Worked 75th Percentile	520	630
Average Hours Worked 90th Percentile	585	734
Median Duration of Spells (years)		8
Wage 10th Percentile	7.99	7.91
Wage 50th Percentile	14.18	15.14
Wage 90th Percentile	25.47	27.69

Figure 2: Life-Cycle Profiles, Ret to Exp Model (blue) versus Data (red)



this model than in the baseline in order to match observed participation rates. This large fixed cost alongside the strong incentive to participate implies that changes in the current wage make little difference to the incentive to participate.

Table 15: Returns to Experience: labor supply elasticities

	age 26	age 36	age 46
Extensive	0.03	0.02	0.03
Intensive Margin <i>25th</i> percentile	0.97	0.61	0.71
Intensive Margin <i>50th</i> percentile	1.58	0.92	1.14
Intensive Margin <i>75th</i> percentile	2.88	1.54	2.16
Macro Elasticity	1.72	1.25	1.29

It may well be that the small response of the extensive margin labor supply that we find is related to the simple model of return to experience we have considered. Whether returns to experience operate in a more subtle manner through intensive margins and the number of hours is a question we leave for future research. If that is the case, we would need to change substantially the estimation methods we used in the first part of the paper.

One possibility, of course, is that returns to tenure are important for some occupations and/or skill levels and not for others. In such a case, it would be necessary to introduce an additional dimension of heterogeneity that would make the aggregation issues we have stressed repeatedly even more salient.

8 Conclusions

In this paper, we have proposed an integrated approach to evaluate the aggregate and micro response of labour supply to changes in wages. To frame these issues, we start from a comprehensive specification of preferences in a life cycle model of consumption and labour supply decisions. Different parameter values have different implications for labour supply elasticities. Our first and somewhat negative result is that aggregation issues are important enough to prevent us from talking meaningfully about the elasticity of labour supply to wages as a single number. This is particularly true for the extensive margin, or participation decisions: in this case the aggregation problems arise naturally from the discreteness of the decision involved and for the non-convexities that make such a decision discrete. We stress that in such a case, aggregate elasticities are likely to vary over time and the business cycle.

We use our framework to study female labour supply in the US, using a long time series of cross sectional data which contains information on both household consumption, labour supply and wages.

Our comprehensive approach yields a number of estimates that are characterized by different degrees of robustness: we obtain the parameters for Marshallian and Hicksian elasticities from intratemporal first order conditions that are relatively robust, the parameters to estimate the Frisch elasticities from Euler equations and the parameters relevant for computing the extensive margin elasticities from the calibration of the full life cycle model which we fit to some life cycle moments. The results of this estimation exercise yield elasticities that are, on the one hand, very heterogeneous in the cross section and, on the other, considerably larger than those estimated in many labour supply papers. We believe that these differences are driven by our explicit use of consumption and the explicit consideration of the marginal rate of substitution between consumption and leisure.

Finally, we show that aggregate responses of female to labour supply do vary both in the cross section and over time. This result is important because it shows that the aggregation issues that are central to our argument have a practical relevance and cannot be ignored. In particular, we find that female labour supply is considerably more responsive to changes in wages during recession than booms.

Our two key points in understanding the controversy over micro and macro estimates of elasticities are first, that previous micro estimates were too low and instead using our consistent and integrated estimation strategy yields much larger estimates of the elasticities; and second, that there is no behavioural content in talking about an aggregate elasticity.

The research we present also poses a number of unanswered questions. In particular, whilst we present some discussion of the effects of return to tenure when it operates through participation, we have not analyzed in any depth the issue of return to tenure in terms of number of hours. Whether returns to experience operate through the intensive or extensive margin is an empirical question and one on which we have not presented much evidence. Should the evidence point to important returns on the intensive margin, our analysis should be changed substantially.

References

- [1] Abowd, J., and Card, D. (1989) “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, vol. 57(2), pages 411-455.
- [2] Attanasio, O. (1999) “Consumption,” Handbook of Macroeconomics, in: J.B Taylor and M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 11, pages 741-812 Elsevier.
- [3] Attanasio, O. and Low, H. (2004) “Estimating Euler Equations,” *Review of Economic Dynamics*, Elsevier for the Society for Economic Dynamics, vol. 7(2), pages 405-435, April.
- [4] Attanasio, O., Low, H. and Sanchez-Marcos, V. (2008) “Explaining Changes in Female Labor Supply in a Life-Cycle Model,” *American Economic Review*, American Economic Association, vol. 98(4), pages 1517-52, September.
- [5] Attanasio, O. and Weber, G. (1993) “Consumption Growth, the Interest Rate and Aggregation,” *Review of Economic Studies*, Wiley Blackwell, vol. 60(3), pages 631-49, July.
- [6] Attanasio, O. and Weber, G. (1995) “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy*, University of Chicago Press, vol. 103(6), pages 1121-57, December.
- [7] Attanasio, O. and Weber, G. (2010) “Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy,” *Journal of Economic Literature*, American Economic Association, vol. 48(3), pages 693-751, September.
- [8] Blundell R, Bozio, A. and G. Laroque, (2011). “Labor Supply and the Extensive Margin,” *American Economic Review*, vol. 101(3), pages 482-86.
- [9] Blundell, R., Browning, M. and Meghir, C. (1994) “Consumer Demand and the Life-Cycle Allocation of Household Expenditures,” *Review of Economic Studies*, Wiley Blackwell, vol. 61(1), pages 57-80, January.
- [10] Blundell, R. and MaCurdy, T. (1999) Labor supply: A review of alternative approaches in *Handbook of Labor Economics*.
- [11] Blundell, R., Meghir, C., and Neves, P. (1993) “Labour supply and intertemporal substitution,” *Journal of Econometrics*, Elsevier, vol. 59(1-2), pages 137-160, September.
- [12] Blundell, R., Pistaferri L. and I. Saporta-Eksten (2015) “Children, Time Allocation and Consumption Insurance”, IFS Working Paper , http://www.ucl.ac.uk/uftp39a/BPS_Becker_Draft_Paper_8_March_2015.pdf .

- [13] Browning, M., Deaton, A. and Irish, M. (1985) "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, Econometric Society, vol. 53(3), pages 503-43, May.
- [14] Carneiro, P. and J.J. Heckman (2003) "Human Capital Policy." In B. Friedman (ed.) *Inequality in America: What Role for Human Capital Policies?*, pp. 77239. Cambridge, Mass.: MIT Press.
- [15] Chang Y, Kim, S., Kwon, K. and Rogerson R. (2011) "Interpreting Labor Supply Regressions in a Model of Full- and Part-Time Work ", *American Economic Review*, vol. 101(3), pages 476-81.
- [16] Chetty, R. Guren, A. Manoli, D. and Weber, A. (2011) "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review*, American Economic Association, vol. 101(3), pages 471-75, May.
- [17] Chetty, R. (2012). "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, Econometric Society, vol. 80(3), pages 969-1018.
- [18] Chetty R., Guren A., Manol, D. and Weber A. (2013). "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," NBER Macroeconomics Annual, University of Chicago Press, vol. 27(1), pages 1 - 56.
- [19] Deaton, A. (1985) "Panel data from time series of cross-sections," *Journal of Econometrics*, Elsevier, vol. 30(1-2), pages 109-126.
- [20] Domeij and Floden (2006). "The labor-supply elasticity and borrowing constraints: Why estimates are biased," *Review of Economic Dynamics*
- [21] Feenberg, D. and Coutts, E. (1993) "An introduction to the TAXSIM model," *Journal of Policy Analysis and Management*, vol. 12(1), pages 189-194.
- [22] Hansen, L. P. (1982) "Large sample properties of generalized method of moments estimators, " *Econometrica*, vol. 50, pages 1029-1054.
- [23] Hayashi, F. (1987) "Tests for Liquidity Constraints: A Critical Survey and Some New Results, " In *Advances in Econometric Fifth World Congress*, vol. 2, edited by Truman F. Bewley. Amsterdam: North-Holland (for Econometric Soc.).
- [24] Heckman, J. (1979) "Sample Selection Bias as a Specification Error," *Econometrica*, Econometric Society, vol. 47(1), pages 153-61, January.
- [25] Hyslop, D (2001) "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings ", *American Economic Review*, vol. 91(4), pages 755-777.

- [26] Huggett, M., Ventura, G. and Yaron, A. (2011) "Sources of Lifetime Inequality," *American Economic Review*, American Economic Association, vol. 101(7), pages 2923-54, December.
- [27] Imai, S. and Keane, M. (2004) "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review*, vol. 45(2), pages 602-631
- [28] Keane, M.(2009) "Labor Supply and Taxes: A Survey ". UTS: Business Working Paper 160
- [29] Keane, M. and Rogerson, R. (2011) "Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective," NBER Working Papers 17430, National Bureau of Economic Research, Inc.
- [30] Keane, M., and Rogerson, R. (2012) "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom ", *Journal of Economic Literature*, vol. 50(2), pages 464-76.
- [31] Ljungqvist, L. and T. Sargent (2011) "A Labor Supply Elasticity Accord? ", *American Economic Review*, vol. 101(3), pages 487-91.
- [32] Ljungqvist, L. and T. Sargent (2014) "An Open Letter to Professors Heckman and Prescott ", http://www.tomsargent.com/research/openletter_H&P.pdf
- [33] Low, H. (2005) "Self-Insurance in a Life-Cycle Model of Labor Supply and Savings," *Review of Economic Dynamics*, Elsevier for the Society for Economic Dynamics, vol. 8(4), pages 945-975, October.
- [34] MaCurdy, T. (1983) "A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty ", *International Economic Review* , vol. 24, pages 265-289.
- [35] Meghir and Pistaferri (2004) "Income Variance Dynamics and Heterogeneity," *Econometrica*, Econometric Society, vol. 72(1), pages 1-32, 01.
- [36] Miller, R. and Sieg, H. (1997) "A Microeconomic Comparison of Household Behavior between Countries," *Journal of Business Economics and Statistics*, vol. 15(2), pages 237-53.
- [37] Prescott, E. C. (2002) "Prosperity and Depression", *American Economic Review*, Vol. 92(1) , pp. 1-15
- [38] Rogerson, R. and Wallenius, J. (2009) "Micro and macro elasticities in a life cycle model with taxes," *Journal of Economic Theory*, Elsevier, vol. 144(6), pages 2277-2292, November.
- [39] Wallenius, J. (2011) "Human Capital Accumulation and the Intertemporal Elasticity of Substitution of Labor: How Large is the Bias?," *Review of Economic Dynamics*, Elsevier for the Society for Economic Dynamics, vol. 14(4), pages 577-591, October.

- [40] Ziliak, J. and Kniesner, T. (2005) “The Effect of Income Taxation on Consumption and Labor Supply,” *Journal of Labor Economics*, University of Chicago Press, vol. 23(4), pages 769-796, October.

Appendix A: Probit results

Probit for wife's labour force participation

Log earnings of husband	-0.258***	(0.016)
husband employed	-2.762***	(0.148)
Elderly in HH	0.004	(0.041)
Log family size	-0.167***	(0.038)
Age	-0.073	(0.054)
Age ²	0.003*	(0.001)
Age ³ /100	-0.003**	(0.001)
Child 0-2	-0.597***	(0.024)
Child 3-15	-0.208***	(0.013)
Child 16-17	0.069*	(0.028)
Wife: White	0.014	(0.051)
Husband: White	-0.146**	(0.052)
Wife: Less than high school	-1.001***	(0.139)
Wife: High school	-0.538***	(0.098)
Wife: College	-0.320***	(0.097)
Husband: Less than high school	0.074*	(0.034)
Husband: High school	0.137***	(0.025)
Husband: College	0.128***	(0.025)
<i>N</i>	27672	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Additional controls education-region-year interactions

Omitted education group is "university or higher"

Appendix B: Bootstrap procedure

We bootstrap standard errors and confidence intervals for both our MRS and Euler equations. The two step Heckman-selection procedure for estimating the MRS coefficients is bootstrapped in the standard way.

Bootstrapping results for our Euler equation requires a more novel approach. This is because we aggregate our data into cohort groups and then implement an IV procedure. Taking Z_t as a vector of exogenous variables, and X_t and Y_t as endogenous variables (with Y_t as our dependent variable) we can reformulate our approach as estimating the equations

$$\begin{aligned} X_t &= \Pi Z_t + v_t \\ Y_t &= X_t \beta + u_t \end{aligned}$$

where v_t is a vector of errors in our first stage. These can be thought of as economic shocks which may have a complicated structure. For instance they may be correlated across time for a given cohort, or may have an aggregate component which is correlated across cohorts for a given time period. Errors may also be correlated across the equations for different exogenous variables Z_t . We

will wish to preserve these correlations when we implement our bootstrap procedure. In order to do this, we attempt to construct the variance-covariance matrix of the residuals v . Rather than filling in all possible cross-correlations in this matrix, we calculate the following moments for each cohort c , and equation i

$$\begin{aligned} & var(v^{i,c}) \\ & cov(v_t^{i,c}, v_{t-1}^{i,c}) \\ & cov(v_t^{i,c}, v_t^{j,c}) \\ & cov(v_t^{i,c}, v_t^{i,k}) \end{aligned}$$

Setting all other correlations to zero. Thus we impose for instance that there is zero correlation between $v_t^{i,c}$ and $v_{t-1}^{i,k}$. Unfortunately, there is no guarantee that this matrix will be positive definite. In our procedure we therefore apply weights to the non-zero elements of our ‘off-diagonal’ matrices - which give the covariances across different cohorts for the same equation - and to our 1st autocovariances for residuals for the same cohort and same equation. The weights we apply to these are the maximum that ensure the resulting matrix is positive definite: in our case 0.28 and 0.22 respectively.

Once we have this matrix we can Cholesky decompose it to obtain a vector of orthogonalised residuals

$$\Omega = vv' = \epsilon C C' \epsilon'$$

We then draw from the orthogonalised residuals, premultiply them by C and then add them to ΠZ_t to reconstruct the endogenous variables (including Y). We then reestimate our reduced form equation to obtain a new set of estimates of β .

The values of Z_t in our case will depend on the results we obtain from our MRS equation, so in each iteration of our bootstrap we resample with replacement from our disaggregated data, re-run the MRS equation, reaggregate to obtain the cohort averages which make up Z_t and then make a draw from our residuals.

Appendix C: Derivatives

In this section we provide the formulae for the first and second derivatives that are used to calculate the different elasticities. We define $D_t = \exp(\pi z_t + \xi P_t + \zeta_t)$. Then it is easy to show that:

$$u_c(c_t, l_t) = D_t M_t^{-\gamma} \alpha_t c_t^{-\phi} \tag{19}$$

$$u_l(c_t, l_t) = D_t M_t^{-\gamma} (1 - \alpha_t) l_t^{-\theta} \tag{20}$$

$$u_{cl}(c_t, l_t) = (-\gamma) D_t M_t^{-\gamma-1} \alpha_t (1 - \alpha_t) c_t^{-\phi} l_t^{-\theta} \quad (21)$$

$$u_{ll}(c_t, l_t) = (-\gamma) \frac{u_l(c_t, l_t)}{M_t} (1 - \alpha_t) l_t^{-\theta} - u_l(c_t, l_t) \theta l_t^{-1} \quad (22)$$

$$u_{cc}(c_t, l_t) = (-\gamma) \frac{u_c(c_t, l_t)}{M_t} \alpha_t c_t^{-\phi} - u_c(c_t, l_t) \phi c_t^{-1} \quad (23)$$

Finally, note that:

$$u_{cl}(c_t, l_t) = (-\gamma) u_c(c_t, l_t) l_t^{-\theta} \frac{(1 - \alpha_t)}{M_t} = (-\gamma) u_l(c_t, l_t) c_t^{-\phi} \frac{\alpha_t}{M_t} \quad (24)$$

These expressions can be used to calculate the Frisch elasticities in the paper. The formula for the Frisch can be derived as follows:

$$\begin{aligned} \begin{bmatrix} u_{cc} & u_{cl} \\ u_{cl} & u_{ll} \end{bmatrix} \begin{bmatrix} \frac{\partial c}{\partial w} \\ \frac{\partial l}{\partial w} \end{bmatrix} &= \begin{bmatrix} 0 \\ \lambda \end{bmatrix} \\ \begin{bmatrix} \frac{\partial c}{\partial w} \\ \frac{\partial l}{\partial w} \end{bmatrix} &= \begin{bmatrix} u_{cc} & u_{cl} \\ u_{cl} & u_{ll} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \lambda \end{bmatrix} \\ \begin{bmatrix} \frac{\partial c}{\partial w} \\ \frac{\partial l}{\partial w} \end{bmatrix} &= \frac{1}{u_{cc}u_{ll} - u_{cl}^2} \begin{bmatrix} u_{ll} & -u_{cl} \\ -u_{cl} & u_{cc} \end{bmatrix} \begin{bmatrix} 0 \\ \lambda \end{bmatrix} \\ \varepsilon_c &= \frac{w}{c} \frac{\partial c}{\partial w} = -\frac{u_c u_{cl}}{u_{cc}u_{ll} - u_{cl}^2} \frac{w}{c} \\ \varepsilon_l &= \frac{w}{l} \frac{\partial l}{\partial w} = \frac{u_c u_{cc}}{u_{cc}u_{ll} - u_{cl}^2} \frac{w}{l} \\ \varepsilon_h &= \frac{w}{h} \frac{\partial h}{\partial l} \frac{\partial l}{\partial w} = -\frac{u_c u_{cc}}{u_{cc}u_{ll} - u_{cl}^2} \frac{w}{h} \varepsilon_h = -\varepsilon_l \frac{l}{h} \end{aligned}$$