

Blundell, Richard; Pistaferri, Luigi; Saporta-Eksten, Itay

Working Paper

Consumption inequality and family labor supply

ECB Working Paper, No. 1656

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Blundell, Richard; Pistaferri, Luigi; Saporta-Eksten, Itay (2014) : Consumption inequality and family labor supply, ECB Working Paper, No. 1656, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/154089>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM



In 2014 all ECB publications feature a motif taken from the €20 banknote.



WORKING PAPER SERIES

NO 1656 / MARCH 2014

CONSUMPTION INEQUALITY AND FAMILY LABOR SUPPLY

Richard Blundell, Luigi Pistaferri
and Itay Saporta-Eksten

HOUSEHOLD FINANCE AND
CONSUMPTION NETWORK

NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Household Finance and Consumption Network

This paper contains research conducted within the Household Finance and Consumption Network (HFCN). The HFCN consists of survey specialists, statisticians and economists from the ECB, the national central banks of the Eurosystem and a number of national statistical institutes.

The HFCN is chaired by Gabriel Fagan (ECB) and Carlos Sánchez Muñoz (ECB). Michael Haliassos (Goethe University Frankfurt), Tullio Jappelli (University of Naples Federico II), Arthur Kennickell (Federal Reserve Board) and Peter Tufano (University of Oxford) act as external consultants, and Sébastien Pérez Duarte (ECB) and Jiri Slacalek (ECB) as Secretaries.

The HFCN collects household-level data on households' finances and consumption in the euro area through a harmonised survey. The HFCN aims at studying in depth the micro-level structural information on euro area households' assets and liabilities. The objectives of the network are:

- 1) understanding economic behaviour of individual households, developments in aggregate variables and the interactions between the two;
- 2) evaluating the impact of shocks, policies and institutional changes on household portfolios and other variables;
- 3) understanding the implications of heterogeneity for aggregate variables;
- 4) estimating choices of different households and their reaction to economic shocks;
- 5) building and calibrating realistic economic models incorporating heterogeneous agents;
- 6) gaining insights into issues such as monetary policy transmission and financial stability.

The refereeing process of this paper has been co-ordinated by a team composed of Gabriel Fagan (ECB), Pirmin Fessler (Oesterreichische Nationalbank), Michalis Haliassos (Goethe University Frankfurt), Tullio Jappelli (University of Naples Federico II), Sébastien Pérez Duarte (ECB), Jiri Slacalek (ECB), Federica Teppa (De Nederlandsche Bank), Peter Tufano (Oxford University) and Philip Vermeulen (ECB).

The paper is released in order to make the results of HFCN research generally available, in preliminary form, to encourage comments and suggestions prior to final publication. The views expressed in the paper are the author's own and do not necessarily reflect those of the ESCB.

Acknowledgements

We thank Mark Aguiar, Nick Bloom, Olympia Bover, Martin Browning, Chris Carroll, Jon Levin, Gianluca Violante and seminar participants at the 2012 ESEM, 2012 Royal Economic Society, 2012 UCL Economics Phd Alumni Conference, 2012 UCL Conference on Labour Markets and the Welfare State, 2011 SED, the 2011 Workshop on Stabilization Policies at the University of Copenhagen, the Bank of Italy, Venice, Bologna, Mannheim, and Cattolica Milan for comments. Thanks to Kerwin Charles for initial help with the new consumption data from the PSID. The authors gratefully acknowledge financial support from the UK Economic and Social Research Council (Blundell), and the ERC starting grant 284024 (Pistaferri). All errors are ours.

Richard Blundell

Department of Economics, University College London, and Institute for Fiscal Studies; e-mail: r.blundell@ucl.ac.uk

Luigi Pistaferri

Department of Economics, Stanford University; e-mail: pista@stanford.edu

Itay Saporta-Eksten

Graduate School of Business, Stanford University, and National Bureau of Economic Research; e-mail: isaporta@stanford.edu

© European Central Bank, 2014

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu
Fax	+49 69 1344 6000

All rights reserved.

ISSN	1725-2806 (online)
EU Catalogue No	QB-AR-14-030-EN-N (online)

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=2161237.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

Abstract

In this paper we examine the link between wage inequality and consumption inequality using a life cycle model that incorporates household consumption and family labor supply decisions. We derive analytical expressions based on approximations for the dynamics of consumption, hours, and earnings of two earners in the presence of correlated wage shocks, non-separability and asset accumulation decisions. We show how the model can be estimated and identified using panel data for hours, earnings, assets and consumption. We focus on the importance of family labour supply as an insurance mechanism to wage shocks and find strong evidence of smoothing of male's and female's permanent shocks to wages. Once family labor supply, assets and taxes are properly accounted for there is little evidence of additional insurance.

Key words: Consumption, Labor Supply, Earnings, Inequality.

JEL codes: D11, D12, D31, D91

Non-technical Summary

The link between household consumption inequality and idiosyncratic income changes has been the focus of a large body of recent economic research. This literature usually relates movements in consumption to predicted and unpredictable income changes as well as persistent and non-persistent shocks to economic resources. One remarkable and consistent empirical finding in most of this recent work is that household consumption appears significantly smoothed, even with respect to highly persistent shocks. But what are the mechanisms behind such smoothing? This is the question we attempt to answer in this paper.

To do so, we set up a life cycle model that allows for three potential sources of smoothing. The first, a traditional one in the literature, is self-insurance through credit markets. The second source is family labor supply, i.e., the fact that hours of work can be adjusted along with, or alternatively to, spending on goods in response to shocks to economic resources. While this is not a new channel, the focus on family labor supply has not received much attention. As we shall see, our empirical analysis suggests that this is a key insurance channel available to families, and hence its omission is particularly glaring if the goal is to have an accurate view of how households respond to changes in their economic fortunes. Finally, households may have access to external sources of insurance, ranging from help received by networks of relatives and friends, to social insurance such as unemployment benefits and food stamps, to formal market insurance. It is hard to model in a credible way the myriad of external insurance channels potentially available to households. We hence choose to subsume these mechanisms into a single parameter, measuring all consumption insurance that remains after accounting for the two “self-insurance” sources discussed above. We use our estimates to measure how much of the consumption smoothing we find in the data can be explained by these various forces in different stages of the life cycle.

From a modeling point of view, our paper has three distinctive features. First, the labor supply of each earner within a household is endogenous (hours are chosen to reflect preferences for work and the dynamics of market wages), heterogeneous (spouses respond differently to wage changes), and potentially non-separable with respect to consumption and also with respect to each other (e.g., partners may enjoy spending time together). The focus on endogenous labor supply makes market wages the primitive source of uncertainty faced by households; the focus on heterogeneity and non-separability agrees with most influential work on labor supply. Second, we model the stochastic component of the wage process as being the sum of transitory and permanent components - these components are allowed to be freely correlated across spouses, reflecting for example assortative mating or risk sharing arrangements. Finally, since our goal is to understand the transmission mechanisms from wage shocks to consumption and labor supply, we obtain analytical expressions for consumption and labor supply as a function of wage shocks using approximations of the first order conditions of the problem and of the lifetime budget constraint. The usefulness of our approach is that it gives a very intuitive and transparent view of how the various structural parameters are identified using panel data on individual wages and earnings (or hours), and household consumption and assets.

One of our contributions is to enrich and extend the theoretical framework used in previous literature. In particular, we consider a life cycle setup in which two individuals within the family (husband and wife) make unitary decisions about household consumption and their individual labor

supply, subject to uncertainty about offered market wages. We allow for partial insurance of wage shocks through asset accumulation; heterogeneous Frisch elasticities for husband and wife; non-separability; and differences between the extensive and intensive margins of labor supply. These extensions are not merely formal, but substantial. Estimating a single earner model when two earners are present potentially yields biased estimates for the level of self-insurance and for the elasticity of intertemporal substitution of consumption, a key parameter for understanding business cycle fluctuations. Studies of the added worker effect that disregard self-insurance through savings may find little evidence for an added worker effect if couples have plenty of accumulated assets to run down in case of negative shocks to resources. Ignoring non-separability could yield biased estimates for the response of consumption to permanent wage shocks (and also distorts the measurement of the welfare effect of risk). The direction of the bias is ambiguous and depends on the substitutability or complementarity of consumption and leisure. If consumption and hours are complements, the response of consumption to permanent shocks is over estimated. If consumption and hours are substitutes the result is reversed. A similar bias emerges in estimating the elasticity of intertemporal substitution in labor supply. With fixed consumption costs of work, differences will naturally appear between elasticities at the extensive and the intensive margin.

From the empirical side, we highlight the separate identification of Marshallian and Frisch elasticities, obtained by looking at the response of labor supply to permanent and transitory wage shocks, respectively. Given the life cycle focus, we allow for age-varying impact of shocks onto consumption and we also consider the possibility that wage shocks are drawn from age-varying distributions. In this framework, the distinction between permanent and transitory shocks is important, although in a finite horizon model the effect of a permanent shock is attenuated by the horizon of the consumer.

Our work has important policy implications. First, most families (i.e., poor or young families) do not have the assets that would allow them to smooth consumption effectively. Without the labor supply channel one could conclude that they have little in the way of maintaining living standards when shocks hit. For a correct design of public and social insurance policies, it is important to know whether households can use labor supply as an alternative insurance mechanism and to what extent they do so. Much depends on whether labor supply is frictionlessly changeable, which can be modelled at the cost of some simplifications, and how strong preferences for leisure are. Moreover, studying how well families smooth income shocks, how this changes over the life and over the business cycle in response to changes in the economic environment confronted, and how different household types differ in their smoothing opportunities, is an important complement to understanding the effect of redistributive policies and anti-poverty strategies.

1 Introduction

The link between household consumption inequality and idiosyncratic income changes has been the focus of a large body of recent economic research (Blundell et al., 2008; Heathcote et al., 2009).¹ This literature usually relates movements in consumption to predicted and unpredictable income changes as well as persistent and non-persistent shocks to economic resources. One remarkable and consistent empirical finding in most of this recent work is that household consumption appears significantly smoothed, even with respect to highly persistent shocks.² But what are the mechanisms behind such smoothing? This is the question we attempt to answer in this paper.

To do so, we set up a life cycle model that allows for three potential sources of smoothing. The first, a traditional one in the literature, is self-insurance through credit markets. The second source is family labor supply, i.e., the fact that hours of work can be adjusted along with, or alternatively to, spending on goods in response to shocks to economic resources. While this is not a new channel (see Heckman, 1974; Low, 2005), the focus on *family* labor supply has not received much attention. As we shall see, our empirical analysis suggests that this is a key insurance channel available to families, and hence its omission is particularly glaring if the goal is to have an accurate view of how households respond to changes in their economic fortunes. Finally, households may have access to external sources of insurance, ranging from help received by networks of relatives and friends, to social insurance such as unemployment benefits and food stamps, to formal market insurance. It is hard to model in a credible way the myriad of external insurance channels potentially available to households. We hence choose to subsume these mechanisms into a single parameter, measuring all consumption insurance that remains after accounting for the two "self-insurance" sources discussed above. We use our estimates to measure how much of the consumption smoothing we find in the data can be explained by these various forces in different stages of the life cycle.

From a modeling point of view, our paper has three distinctive features. First, the labor supply of each earner within a household is endogenous (hours are chosen to reflect preferences for work and the dynamics of

¹Meghir and Pistaferri (2011) and Jappelli and Pistaferri (2010) review the relevant theoretical and empirical literature.

²See also Krueger and Perri (2006), Primiceri and van-Rens (2009), Kaufman and Pistaferri (2009), Kaplan and Violante (2010) and Hryshko (2011). See moreover Guvenen (2007) and Guvenen and Smith (2010) for an alternative view about the nature of the income process and its implications for the consumption-income nexus.

market wages), heterogeneous (spouses respond differently to wage changes), and potentially non-separable with respect to consumption and also with respect to each other (e.g., partners may enjoy spending time together). The focus on endogenous labor supply makes market wages the primitive source of uncertainty faced by households; the focus on heterogeneity and non-separability agrees with most influential work on labor supply (see Blundell and MaCurdy, 1999, for a survey). Second, we model the stochastic component of the wage process as being the sum of transitory and permanent components - these components are allowed to be freely correlated across spouses, reflecting for example assortative mating or risk sharing arrangements. Finally, since our goal is to understand the transmission mechanisms from wage shocks to consumption and labor supply, we obtain analytical expressions for consumption and labor supply as a function of wage shocks using approximations of the first order conditions of the problem and of the lifetime budget constraint (as illustrated in Blundell and Preston, 1998; Blundell et al., 2008). A similar goal is pursued in Heathcote et al. (2009), but it differs from ours because the authors focus on one-earner labor supply models, assume that preferences are separable, and decompose permanent shocks into two components (measuring the fraction of permanent shocks which is insurable). The usefulness of our approach is that it gives a very intuitive and transparent view of how the various structural parameters are identified using panel data on individual wages and earnings (or hours), and household consumption and assets.

But where do we find such rich data? In the US there are two sources of data that have been extensively used, the CEX and the PSID. The CEX has complete consumption data, but lacks a long panel component and the quality of its income, asset and consumption data has recently raised some worries. The PSID has traditionally been used to address the type of questions we are concerned with in this paper, but until recently had incomplete consumption data, which has meant that authors have either used just food data (Hall and Mishkin, 1982), or resorted to data imputation strategies (Blundell et al., 2008). In this paper we make use of new consumption data that as far as we know are untapped for the type of questions asked here. Starting in 1999 the PSID was drastically redesigned. In particular, it enriched the consumption information available to researchers, which now covers over 70% of all consumption items available in the CEX. On the other hand, as part of its redesign, data are now available only every other year. However, this can be easily accounted for in our framework.

Our paper is related to several literatures in macroeconomics and labor economics. A large literature in macroeconomics is devoted to understanding the response of consumption to income changes, both anticipated changes and economic shocks. A good understanding of how consumers respond to income changes is of course crucial when evaluating policy changes that impact households' resources (such as tax and labor market reforms), as well as for the design of stabilization, social insurance, and income maintenance policies. Recent contributions that assume exogenous labor supply include Krueger and Perri (2006), and Blundell et al. (2008). In contrast, Attanasio et al. (2008), Blundell and Preston (2004), and Heathcote et al. (2009), relax the exogeneity of labor supply but either focus on a single earner, aggregate hours across spouses, or impose restrictions on the nature and type of insurance available to consumers.³ Most of these papers find a significant degree of consumption smoothing against income shocks, including very persistent ones.

A related literature in labor economics asks to what extent a secondary earner's labor supply (typically, the wife's) increases in response to negative wage shocks faced by the primary earner (Lundberg, 1985). This literature, also known as the "added worker effect" literature, investigates the role of marriage as a risk sharing device focusing mostly on the wives' propensity to become employed when their husbands exit employment. Saving choices are typically not modeled.⁴

A somewhat distinct, but equally large and influential literature estimates the responsiveness of individual labor supply to wage changes using micro data (see Keane, 2011, for a recent review of this literature). Most of the papers in this literature do not consider the joint consumption-labor supply choice (with some exceptions, Altonji 1986) and focus on the single earner case. We show how the labor elasticities of intertemporal substitution can be identified allowing for non-separability with respect to consumption and the labor supply

³More in detail, Attanasio et al. (2008) introduce a model with a two earners' household. They do not explicitly model the labor supply decision of the household, but rather use Markov process for the evolution of the participation of the second earner. Blundell and Preston (2004) develop a lifecycle framework with two earners modeling the simultaneous decision on consumption and hours worked for each earner. As in Blundell and Preston (1998) they assume that permanent shocks to earnings are fully transmitted to consumption. Finally, Heathcote et al. (2009) develop an analytical framework for the estimation of the response of consumption to insurable and uninsurable shocks to wages in a single earner setup.

⁴The most relevant paper for our purposes is Hyslop (2001). He uses a life cycle model to look directly at the response of hours worked by one earner to the other earners' wage shocks, decomposing it as the response to transitory and permanent components. He finds that the permanent shocks to wages are correlated for first and second earner, and that the relatively large labor supply elasticity for wives can explain about 20% of the rise in household earnings inequality in the early 1980's. A recent paper by Juhn and Potter (2007) finds that the value of marriage as a risk sharing device has diminished due to an increase in correlation of employment among couples. See also Stephens (2002).

of the partner. As we shall see, allowing for non-separability is important, as previously found in micro data (Browning and Meghir, 1991). In a recent macro literature, the degree of complementarity between consumption and hours plays an important role for explaining multiplier effects (see for example Christiano et al., 2011). Adding consumption information besides labor supply information increases efficiency of estimates and imposes on the model the tougher requirement of fitting not just labor supply moments, but also consumption moments.

One of our contributions is to enrich and extend the theoretical framework used in previous literature. In particular, we consider a life cycle setup in which two individuals within the family (husband and wife) make unitary decisions about household consumption and their individual labor supply, subject to uncertainty about offered market wages. We allow for partial insurance of wage shocks through asset accumulation; heterogeneous Frisch elasticities for husband and wife; non-separability; and differences between the extensive and intensive margins of labor supply. These extensions are not merely formal, but substantial. Estimating a single earner model when two earners are present potentially yields biased estimates for the level of self insurance and for the elasticity of intertemporal substitution of consumption, a key parameter for understanding business cycle fluctuations. Studies of the “added worker effect” that disregard self-insurance through savings may find little evidence for an added worker effect if couples have plenty of accumulated assets to run down in case of negative shocks to resources. Ignoring nonseparability could yield biased estimates for the response of consumption to permanent wage shocks (and also distorts the measurement of the welfare effect of risk). The direction of the bias is ambiguous and depends on the substitutability or complementarity of consumption and leisure. If consumption and hours are complements, the response of consumption to permanent shocks is over estimated. If consumption and hours are substitutes the result is reversed. A similar bias emerges in estimating the elasticity of intertemporal substitution in labor supply. With fixed consumption costs of work, differences will naturally appear between elasticities at the extensive and the intensive margin.

From the empirical side, we highlight the separate identification of Marshallian and Frisch elasticities, obtained by looking at the response of labor supply to permanent and transitory wage shocks, respectively. Given the life cycle focus, we allow for age-varying impact of shocks onto consumption and we also consider

the possibility that wage shocks are drawn from age-varying distributions. In this framework, the distinction between permanent and transitory shocks is important, although in a finite horizon model the effect of a permanent shock is attenuated by the horizon of the consumer.

Our work has important policy implications. First, most families (i.e., poor or young families) do not have the assets that would allow them to smooth consumption effectively. Without the labor supply channel one could conclude that they have little in the way of maintaining living standards when shocks hit. For a correct design of public and social insurance policies, it is important to know whether households can use labor supply as an alternative insurance mechanism and to what extent they do so. Much depends on whether labor supply is frictionlessly changeable, which can be modelled at the cost of some simplifications, and how strong preferences for leisure are. Moreover, studying how well families smooth income shocks, how this changes over the life and over the business cycle in response to changes in the economic environment confronted, and how different household types differ in their smoothing opportunities, is an important complement to understanding the effect of redistributive policies and anti-poverty strategies.

The rest of the paper is organized as follows. Section 2 describes the life cycle model we use and develops the two cases of interest of additive separability and non-separability; we also discuss identification and how we estimate the parameters of interest. In Section 3 we describe the data, discuss the empirical strategy, and the estimation problems we face. Section 4 discusses the main results (including robustness checks), while Section 5 includes a discussion of intensive vs. extensive margin, labor supply elasticities, and a quantification of the degree and importance of the various insurance channels. It also examines the effect of introducing non-linear taxation. Taxes change the interpretation of our results. In particular, we find that the Frisch elasticities are typically larger when we explicitly account for non-linear taxes. However, the overall results on consumption smoothing and the insurance value of family labor supply are largely unaffected. Section 6 concludes.

2 Two Earners Life-Cycle Model

In this section we develop the link between wage shocks, labor supply and consumption in a life cycle model of a two earners' household drawing utility from consumption and disutility from work. The household

chooses consumption and hours of the first and second earner to optimize expected life time utility. We assume throughout that the hourly wage process is exogenous. For the time being we assume that the utility function is separable in consumption and both earners' hours. We relax this later. We maintain the assumption of separability over time throughout the paper. We also assume decisions are made by the two household members within a unitary framework. The difficulty with relaxing this is that identification becomes particularly cumbersome in the dynamics case (see Chiappori, 1988, for a static approach).

2.1 Wage Process

For each earner within the household we adopt a permanent-transitory type wage process, assuming that the permanent component evolves as a unit root process. We confine our analysis to the case of households with two potential earners, husband and wife. Suppose that the log of real wage of individual $j = \{1, 2\}$ of household i at time t can be written as

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

$$F_{i,j,t} = F_{i,j,t-1} + v_{i,j,t}$$

where $x_{i,j,t}$ are observed characteristics affecting wages and known to the household. $u_{i,j,t}$ and $v_{i,j,t}$ are transitory shocks (such as short illnesses that may affect productivity on the job) and permanent shocks (such as technological shocks that make one's marketable skills less or more valuable), respectively. We make the following assumptions regarding correlation of shocks over time and within household:

$$\mathbb{E}(u_{i,j,t} u_{i,k,t-s}) = \begin{cases} \sigma_{u_j}^2 & \text{if } j = k \text{ and } s = 0 \\ \sigma_{u_j u_k} & \text{if } j \neq k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

$$\mathbb{E}(v_{i,j,t} v_{i,k,t-s}) = \begin{cases} \sigma_{v_j}^2 & \text{if } j = k \text{ and } s = 0 \\ \sigma_{v_j v_k} & \text{if } j \neq k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

and $\mathbb{E}(u_{i,j,t} v_{i,k,t-s}) = 0$ for all $j, k = \{1, 2\}$ and all s . The shocks are not formally insurable. In one of the robustness checks we conduct, we let the variances of the shocks vary over stages of the life cycle. This is

done to capture the possibility that there is more dispersion in shocks for older workers due, for example, to worsening of health conditions.

Assumptions (2) and (3) imply that the process for each shock does not vary with time and it is serially uncorrelated. Our data do not span a long time period (six waves, covering eleven years) and hence these assumptions are less strong than they appear at first (the variance of wages were rather flat over the 1999-2009 period covered by our data). We also assume that contemporary shocks (transitory or permanent) can be correlated across spouses.⁵ This correlation is theoretically ambiguous. If spouses were to adopt perfect risk sharing mechanisms, they would select jobs where shocks are negatively correlated. Alternatively, assortative mating or other forms of sorting can imply that spouses work in similar jobs, similar industries, and sometimes in the same firm - hence their shocks may be potentially highly positively correlated. Finally, we assume that transitory and permanent shocks are uncorrelated within and between persons.⁶

While the stochastic wage structure embedded in (1) is widely used in models of the type we are considering here, it is far from being uncontroversial. Some authors have stressed the role of superior information issues (Primiceri and van Rens, 2009); other researchers have emphasized the importance of allowing for growth heterogeneity (Guisen and Smith, 2010). Nevertheless, we will show that (1) fits wage data rather well. We also assume that the household has no advance information about the shocks and that the shocks are observed (separately) at time t .⁷ We provide a test of no superior information in Section 5.1.

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

$$\Delta w_{i,j,t} = \Delta u_{i,j,t} + v_{i,j,t} \quad (4)$$

where Δ is a first difference operator and $\Delta w_{i,j,t} = \Delta \ln W_{i,j,t} - \Delta x'_{i,j,t} \beta_W^j$ (the log change in wages net of observables). We discuss measurement error issues in Section 3.2.1.

⁵This is potentially important given the empirical findings for the correlation of labor market outcomes of married couples. See for example Juhn and Potter (2007) and Hyslop (2001).

⁶Hryshko et al. (2011) considers the consequences of relaxing this assumption for partial insurance models.

⁷This is a key assumption in the context of empirical analysis on consumption insurance. See Meghir and Pistaferri (2011) for a discussion about the interpretation of insurance coefficients when this assumption is violated.

2.2 Household Maximization Problem

Given the exogenous wage processes described above, we assume that the household's maximization problem is given by:

$$\max \mathbb{E}_t \sum_{s=0}^{T-t} u_{t+s} (C_{i,t+s}, H_{i,1,t+s}, H_{i,2,t+s}; z_{i,t+s}, z_{i,1,t+s}, z_{i,2,t+s}) \quad (5)$$

subject to the intertemporal budget constraint

$$A_{i,t+1} = (1+r)(A_{i,t} + H_{i,1,t}W_{i,1,t} + H_{i,2,t}W_{i,2,t} - C_{i,t}) \quad (6)$$

The time subscript on the utility function $u_{t+s}(\cdot)$ captures intertemporal discounting. The primary arguments of the utility function are household consumption $C_{i,t}$, and the hours chosen by the two earners, respectively $H_{i,1,t}$ and $H_{i,2,t}$. The utility function also includes preference shifters specific to the household, such as number of children ($z_{i,t}$), or specific to the earner, such as his or her age ($z_{i,1,t}$ and $z_{i,2,t}$). These preference shifters can potentially include stochastic components as well. Note that we can (and will) specialize $u_{t+s}(C_{i,t+s}, H_{i,1,t+s}, H_{i,2,t+s}; z_{i,t+s}, z_{i,1,t+s}, z_{i,2,t+s})$ to cover the case of additive separability and the non-separability case. We assume that $u_{t+s}(\cdot)$ is twice differentiable in all its primary arguments with $u_C > 0$, $u_{CC} < 0$, $u_{H_j} < 0$, $u_{H_j H_j} > 0$ for $j \in \{1, 2\}$ and $u(0, H_1, H_2) \rightarrow -\infty$. Finally, $A_{i,t}$ denotes the assets at the beginning of period t and r is the fixed interest rate (i.e., this is a Bewley-type model in which consumers have access to a single risk-free bond).

There are only a few special cases for which the problem (5)-(6) can be analytically solved. One is the case of quadratic utility and additive separability (Hall, 1978) which predicts that consumption evolves as a random walk. Unfortunately, a quadratic utility model does not generate precautionary savings and is therefore unrealistic. The exponential utility specification is another case for which analytical solutions exist (Caballero, 1990). A caveat of exponential utility is that it implies constant absolute risk aversion.

While analytical solutions are based on strong counterfactual assumptions regarding preferences, approximations for the evolution of consumption and hours can be found in the literature for more realistic assumptions about preferences. In the following subsection we apply a two-step approximation procedure similar to the one used in Blundell and Preston (2004), Blundell et al. (2008), and Attanasio et al. (2002). The overall accuracy of this approximation under a variety of preference and income specifications is assessed

in detail in Blundell et al. (2011b).

2.3 The Dynamics of Consumption, Hours and Earnings

Our goal is to link the growth rates of consumption and hours to the wage shocks experienced by the household. We achieve this in two steps. First, we use a Taylor approximation to the first order conditions of the problem. This yields expressions for the growth rate of consumption and the growth rate of hours in terms of changes in wages and an additional expectation error term (the innovation in the marginal utility of wealth). This is a standard log-linearization approach. Second, we take a log-linearization of the intertemporal budget constraint. This allows us to map the (unobservable) expectation error in the consumption and hours growth equations into wage shocks. We discuss the two empirically relevant cases, the additive separability case first and the non-separable case next.

2.3.1 The Additive Separability Case

In the additive separability case, we write the utility function in (5) as:

$$u_{t+s}(\cdot) = (1 + \delta)^{-s} [u(C_{i,t+s}; z_{i,t+s}) - g^1(H_{i,1,t+s}; z_{i,1,t+s}) - g^2(H_{i,2,t+s}; z_{i,2,t+s})]$$

Assuming that the solution for hours is always interior,⁸ we approximate the first order conditions to yield the following growth equations for household i 's consumption and for earner j 's earnings (See Appendix 1 for a proof):⁹

$$\begin{aligned} \Delta c_{i,t} &\simeq -\eta_{c,p} \Delta \ln \lambda_{i,t} \\ &= -\eta_{c,p} (\omega_t + \varepsilon_{i,t}) \end{aligned} \tag{7}$$

$$\begin{aligned} \Delta y_{i,j,t} &\simeq (1 + \eta_{h_j,w_j}) \Delta \ln w_{i,j,t} + \eta_{h_j,w_j} \Delta \ln \lambda_{i,t} \\ &= (1 + \eta_{h_j,w_j}) (\Delta u_{i,j,t} + v_{i,j,t}) + \eta_{h_j,w_j} (\omega_t + \varepsilon_{i,t}) \end{aligned} \tag{8}$$

⁸By the properties of $u(\cdot)$ the solution for consumption is always interior. Assuming that hours are always positive is a much stronger assumption. However, since the goal of this procedure is to derive an analytical estimation framework one can think of correcting the distribution of observed wages, earnings and consumption for the selection to employment, rather than to explicitly model the participation decision. See section 3.2.3 for further discussion.

⁹Given that definitionally $\Delta \log Y_{i,j,t} = \Delta \log H_{i,j,t} + \Delta \log W_{i,j,t}$, we will find it useful to work with log earnings rather than log hours in what follows.

where $c_{i,t}$ and $y_{i,j,t}$ are log consumption and log earnings of earner j (net of predictable taste shifters). We decompose the growth of the marginal utility of wealth, as captured by the Lagrange multiplier on the sequential budget constraint $\lambda_{i,t}$, into two components. The first component, ω_t , is a function of the interest rate r , the discount factor δ , and the variance of the change of marginal utility. This component captures the intertemporal substitution and precautionary motives for savings. Assuming that the only source of uncertainty in this setup is the idiosyncratic wage shocks, ω_t is fixed in the cross-section. The second component, $\varepsilon_{i,t}$, captures the revisions in the growth of the marginal utility of wealth. The parameter $\eta_{c,p} = -\frac{u_C}{u_{CC}} \frac{1}{C} > 0$ is the elasticity of intertemporal substitution (EIS) for consumption and $\eta_{h_j,w_j} = \frac{g_{H_j}^j}{g_{H_j H_j}^j} \frac{1}{H_j} > 0$ is the EIS for labor supply of earner j , both assumed to be constant.¹⁰

While the characterization (7)-(8) is theoretically appealing, it is empirically not very useful because we do not know how to characterize the marginal utility of wealth and hence its innovations. To make some progress, we follow Blundell et al. (2008), and log-linearize the intertemporal budget constraint

$$E_t \sum_{s=0}^{T-t} \frac{C_{i,t+s}}{(1+r)^s} = A_t + E_t \sum_{s=0}^{T-t} \frac{W_{i,1,t+s} H_{i,1,t+s}}{(1+r)^s} + E_t \sum_{s=0}^{T-t} \frac{W_{i,2,t+s} H_{i,2,t+s}}{(1+r)^s} \quad (9)$$

and then take the difference in expectations between period t and $t-1$ to obtain equations that link consumption and earnings growth of the two earners to the wage shocks they face (see Appendix 2 for the exact derivation). From the second step of the approximation, we can write the shock to the growth in the marginal utility of wealth $\varepsilon_{i,t}$ as a linear function of the change in transitory shocks ($\Delta u_{i,1,t}$ and $\Delta u_{i,2,t}$) and the permanent shocks ($v_{i,1,t}$ and $v_{i,2,t}$) faced by the two earners. From now on, however, we will assume that the transitory wage shocks of either spouse ($\Delta u_{i,1,t}$ and $\Delta u_{i,2,t}$) have no wealth effect (which is likely true when the horizon is sufficiently long). The assumptions we have made yield the following equations for consumption growth and for the growth of earnings of the two earners under additive separability:

$$\begin{pmatrix} \Delta c_{i,t} \\ \Delta y_{i,1,t} \\ \Delta y_{i,2,t} \end{pmatrix} \simeq \begin{pmatrix} 0 & 0 & \kappa_{c,v_1} & \kappa_{c,v_2} \\ \kappa_{y_1,u_1} & 0 & \kappa_{y_1,v_1} & \kappa_{y_1,v_2} \\ 0 & \kappa_{y_2,u_2} & \kappa_{y_2,v_1} & \kappa_{y_2,v_2} \end{pmatrix} \begin{pmatrix} \Delta u_{i,1,t} \\ \Delta u_{i,2,t} \\ v_{i,1,t} \\ v_{i,2,t} \end{pmatrix} \quad (10)$$

¹⁰Heretofore, $\eta_{x,y}$ measures the Frisch (marginal-utility constant) elasticity of x relative to changes in price y .

where

$$\kappa_{c,v_j} = \frac{\eta_{c,p} (1 - \pi_{i,t}) s_{i,j,t} (1 + \eta_{h_j,w_j})}{\eta_{c,p} + (1 - \pi_{i,t}) \bar{\eta}_{h,w}} \quad (11)$$

$$\kappa_{y_j,u_j} = 1 + \eta_{h_j,w_j} \quad (12)$$

$$\kappa_{y_j,v_j} = 1 + \eta_{h_j,w_j} \left(1 - \frac{(1 - \pi_{i,t}) s_{i,j,t} (1 + \eta_{h_j,w_j})}{\eta_{c,p} + (1 - \pi_{i,t}) \bar{\eta}_{h,w}} \right) \quad (13)$$

$$\kappa_{y_j,v_{-j}} = - \frac{\eta_{h_j,w_j} (1 - \pi_{i,t}) s_{i,-j,t} (1 + \eta_{h_{-j},w_{-j}})}{\eta_{c,p} + (1 - \pi_{i,t}) \bar{\eta}_{h,w}}, \quad (14)$$

In the expression above $\pi_{i,t} \approx \frac{\text{Assets}_{i,t}}{\text{Assets}_{i,t} + \text{Human Wealth}_{i,t}}$ is the "partial insurance" coefficient (the higher $\pi_{i,t}$ the lower the sensitivity of consumption to shocks), $s_{i,j,t} \approx \frac{\text{Human Wealth}_{i,j,t}}{\text{Human Wealth}_{i,t}}$ is the share of earner j 's human wealth over family human wealth (with $\sum_{j=1}^2 s_{i,j,t} = 1$), and $\bar{\eta}_{h,w} = \sum_{j=1}^2 s_{i,j,t} \eta_{h_j,w_j}$ is the household's weighted average of the EIS of labor supply of the two earners.¹¹ Note that $\text{Human Wealth}_{i,t}$ is the expected discounted flow of lifetime earnings of the household at the beginning of period t .¹²

2.3.2 The Non-separable Case

Consider now removing the assumption of separability between consumption and leisure, i.e., leave $u_{t+s}(\cdot)$ unrestricted. A direct implication of relaxing the separability assumption is that the marginal utility of consumption now depends on hours. This changes the decision making process of the household in the sense that it has to choose hours considering the effect that this decision may have on the utility from consumption. This implies that while in the separable case the Frisch elasticity with respect to own price and the elasticity of intertemporal substitution coincide, in the non-separable case this is no longer the case (see the online Appendix 3 for definitions).¹³ The signs of the Frisch elasticities η_{c,w_j} and $\eta_{h_j,p}$ determine

¹¹We use the notation " $-j$ " to indicate variables that refer to the other earner. For example, $\kappa_{y_j,v_{-j}}$ measures the response of earner j 's earnings ($j = \{1, 2\}$) to the other earner's permanent shock.

¹²This system of equations assume that the terms related to ω_t are absorbed in the observables. Note also that the κ parameters vary in principle by i, j, t . However, they only do so through $\pi_{i,t}$ and $s_{i,j,t}$, which will be "pre-estimated" using asset and human capital data. Hence, for simplicity we omit the i, j, t subscripts onto the transmission parameters κ .

¹³See for example Browning et al. (1999). They show that the EIS for consumption in the nonseparable case is the sum of the Frisch elasticities of consumption with respect to own price and with respect to wages. In the separable case the latter is zero, therefore the Frisch elasticity and the EIS coincide.

whether consumption and hours of earner j are Frisch complements ($\eta_{c,w_j} > 0, \eta_{h_j,p} < 0$) or Frisch substitutes ($\eta_{c,w_j} < 0, \eta_{h_j,p} > 0$).

We show in Appendix 3 that the approximation to the Euler equations and the log-linearization of the intertemporal budget constraint yield the following dynamics for consumption and earnings of the two earners:

$$\begin{pmatrix} \Delta c_{i,t} \\ \Delta y_{i,1,t} \\ \Delta y_{i,2,t} \end{pmatrix} \simeq \begin{pmatrix} \kappa_{c,u_1} & \kappa_{c,u_2} & \kappa_{c,v_1} & \kappa_{c,v_2} \\ \kappa_{y_1,u_1} & \kappa_{y_1,u_2} & \kappa_{y_1,v_1} & \kappa_{y_1,v_2} \\ \kappa_{y_2,u_1} & \kappa_{y_2,u_2} & \kappa_{y_2,v_1} & \kappa_{y_2,v_2} \end{pmatrix} \begin{pmatrix} \Delta u_{i,1,t} \\ \Delta u_{i,2,t} \\ v_{i,1,t} \\ v_{i,2,t} \end{pmatrix} \quad (15)$$

where, as before, the parameters $\kappa_{m,n}$ measure the response of variable m ($\Delta c_{i,t}$ and $\Delta y_{i,j,t}$) to the wage shock n ($\Delta u_{i,j,t}$ and $v_{i,j,t}$).

Compared to the case of additive separability, in the non-separable case the parameters $\kappa_{c,u_1}, \kappa_{c,u_2}, \kappa_{y_1,u_2}$ and κ_{y_2,u_1} are not restricted to be zero. In particular, one can show that, quite intuitively, $\kappa_{c,u_j} = \eta_{c,w_j}$ and $\kappa_{y_j,u_{-j}} = \eta_{h_j,w_{-j}}$ for $j = \{1, 2\}$ (see Appendix 3). In essence, a test of non-separability between consumption and the leisure of earner j is a test of whether consumption respond to transitory shock of that earner (shocks that do not have, or have only negligible, wealth effects). With non-separability a transitory wage shock induces a change in hours and, through preference shifts, requires an adjustment also of consumption.¹⁴ Similarly, a test of non-separability between the leisures of the two spouses is a test of whether earnings (that is, labor supply) of earner j respond to the (wealth-constant) transitory shock faced by the other earner. When preferences are separable these transitory shocks have no wealth effect in the contexts considered, so no response is expected. But in the non-separable case these shocks shift preferences (for example because spouse enjoy leisure together), so they generate a response that depends on the degree of complementarity/separability between the arguments of the period utility function.

The remaining transmission coefficients $\kappa_{m,n}$ are - as before - complicated functions of the Frisch elasticities (including those measuring the extent and sign of non-separability), partial insurance (and possibly external insurance parameter), as well as the human wealth shares. To save space, we report the relevant

¹⁴Of course, the test can also reject if consumption responds to transitory shocks due to failure of self-insuring against it. As we shall see, in this case the coefficient κ_{c,u_j} should be positive, while in the empirical analysis we find that $\kappa_{c,u_j} < 0$.

expressions in Appendix 3. Given its importance, we report here only the expression for κ_{c,v_j} , the response of consumption to a permanent shocks to earner j 's wage:¹⁵

$$\kappa_{c,v_j} = \eta_{c,w_j} + \frac{\left(\eta_{c,p} - \left(\eta_{c,w_j} + \eta_{c,w_{-j}}\right)\right) \left[(1 - \pi_{i,t}) \left(s_{i,j,t} + \overline{\eta_{h,w_j}}\right) - \eta_{c,w_j}\right]}{\left(\eta_{c,p} - \left(\eta_{c,w_j} + \eta_{c,w_{-j}}\right)\right) + (1 - \pi_{i,t}) \left(\overline{\eta_{h,w_j}} + \overline{\eta_{h,w_{-j}}} + \overline{\eta_{h,p}}\right)} \quad (16)$$

which of course collapses to κ_{c,v_j} of the additive separable case if $\eta_{c,w_j} = \eta_{c,w_{-j}} = \eta_{h_j,p} = \eta_{h_{-j},p} = \eta_{h_j,w_{-j}} = \eta_{h_{-j},w_j} = 0$.

Special cases are easily obtained from the more general formulation (15). If we assume that labor supply is exogenous (which is equivalent to assuming $\eta_{h_j,w_j} = 0$ for $j = \{1, 2\}$), that there is a single earner ($s_{i,j,t} = 1$), and that preferences are separable ($\eta_{c,w_j} = \eta_{c,w_{-j}} = \eta_{h_j,p} = \eta_{h_{-j},p} = \eta_{h_j,w_{-j}} = \eta_{h_{-j},w_j} = 0$), then we obtain the specification of Blundell et al. (2008). The specification of Heathcote et al. (2009) can be obtained further imposing $s_{i,j,t} = (1 - \beta)(1 - \pi_{i,t}) = 1$.

2.3.3 Insurance Above Self-Insurance

Expressions (11)-(14) and (16) are derived under the assumption that there is no insurance over and above self-insurance. However, households may have access to multiple external sources of insurance, ranging from help received by networks of relatives and friends, to social insurance such as unemployment benefits and food stamps, to formal market insurance. It is hard to model in a credible way the myriad of external insurance channels potentially available to households. We hence choose to subsume these mechanisms into a single parameter (β), which factors $\pi_{i,t}$ whenever it appears. For example, the response of consumption to a permanent shock to male wages in the separable case (11) becomes

$$\kappa_{c,v_j} = \frac{\eta_{c,p} (1 - \beta) (1 - \pi_{i,t}) s_{i,j,t} \left(1 + \eta_{h_j,w_j}\right)}{\eta_{c,p} + (1 - \beta) (1 - \pi_{i,t}) \overline{\eta_{h,w}}}$$

(κ_{y_j,v_j} , $\kappa_{y_j,v_{-j}}$, and κ_{c,v_j} in the non-separable case are revised accordingly).

The parameter β measures all consumption insurance that remains after accounting for the "self-insurance" sources represented by asset accumulation (through the risk free bond A) and labor supply of the primary and secondary earner. Here, $\beta = 0$ means that there is no external insurance over and above self-insurance

¹⁵Where the notation $\overline{\eta_{h,n}} = s_{i,t} \eta_{h_1,n} + (1 - s_{i,t}) \eta_{h_2,n}$, and $n = \{w_1, w_2, p\}$

through assets and labor supply, while $\beta > 0$ would imply some external insurance is present. Note that it is also possible that $\beta < 0$ - which may capture the fact that consumption over-respond to shocks, for example because assets are held in illiquid forms and transaction costs exceed the benefit of smoothing (see for a similar argument Kaplan and Violante, 2011).

2.3.4 Interpretation

To aid in the interpretation of the parameters, let us take the case of separable preferences for simplicity (the set of equations (10) and transmission coefficients (11)-(14)). The interpretation in the non-separable case is similar, and we will discuss it at the end of this section.

Let us start with labor supply responses. Because in the separable framework transitory shocks have negligible or no wealth effects, the earnings of a given earner do not respond to the transitory wage shocks faced by the other earner (and *vice versa*) - hence the zero restrictions on $\kappa_{y_j, u_{-j}}$. In contrast, each earner's labor supply respond to his/her own transitory wage shock to an extent that depends on his/her labor supply EIS (and since transitory shock translate one-to-one in wage changes, the coefficient $\kappa_{y_j, u_j} = (1 + \eta_{h_j, w_j})$). This is almost definitional: the Frisch elasticity (which here coincides with the EIS) measures the labor supply response to a wealth-constant wage change, which here is represented by a pure transitory shock.

The response of earner j 's to a permanent shock to his/her own wage is informative about whether labor supply is used as a *consumption smoothing device*, i.e., as a shock absorber. This depends crucially on the traditional tension between the wealth and the substitution effect of a wage change. This response is hence unrestricted by theory, and indeed the response of earner j 's to a permanent shock to his/her own wage is the closest approximation to a *Marshallian* labor supply effect (as opposed to the *Frisch* effect discussed above). For labor supply to be used as a consumption smoothing device, we require $\kappa_{y_j, v_j} < 1$ (implying that *hours* move in the opposite direction as the permanent shock - they rise, or people work longer, when wages decline permanently). This occurs when the wealth effect dominates the substitution effect of a permanent wage change. In particular, to build intuition, assume there is only one earner for simplicity ($s_{i,j,t} = 1$). In this case, the condition that ensures that labor supply is used as a consumption smoothing device is:

$$(1 - \beta)(1 - \pi_{i,t}) - \eta_{c,p} > 0$$

This condition is more likely to be satisfied when consumers have little or no accumulated assets and/or no access to external sources of insurance ($\pi_{i,t} \rightarrow 0$ and/or $\beta \rightarrow 0$), so that labor supply appears as the sole source of consumption smoothing available to consumers, and when consumers are highly reluctant to intertemporal fluctuations in their consumption ($\eta_{c,p} \rightarrow 0$), so that adjustment is delegated to declines in leisure rather than declines in consumption.¹⁶

The response of earner j 's to a permanent shock faced by the other earner is instead informative about the so-called *added worker effect*. Looking at $\kappa_{y_j, v_{-j}}$, it is easy to see that the latter effect is unambiguously negative, i.e., earner j always increases her labor supply when earner i is hit by a permanent negative shock. Why? The reason is that a permanent negative shock faced by earner i has only a wealth effect as far as earner j is concerned, and no substitution effect (the household is permanently poorer when earner i has a permanently lower wage and hence a reduction in all consumptions, including consumption of leisure of earner j is warranted).

What about consumption responses to shocks? The first thing to notice is that in the additive separability case, and if credit markets are assumed to work well, consumption does not respond to transitory shocks ($\kappa_{c, u_j} = 0$ for $j = \{1, 2\}$). This is because (for consumers with a long horizon) transitory shocks have no lifetime wealth effect (they have negligible impact on the revision of the marginal utility of wealth). As for the response to permanent shocks, we know that in traditional analyses with e.g. quadratic utility, consumption respond one-to-one to permanent shocks. Equations (10) shows how misleading this can be when we account for family labor supply and precautionary behavior. This is important because neglecting these two forces may give a misleading view of the response of consumption to, say, tax policies that change permanently after-tax wages.

In our framework, the response of consumption to permanent wage shocks depends on the insurance

¹⁶In the more general case with multiple earners, labor supply of the primary earner is more likely to be used to smooth consumption if the secondary earner counts little in the balance of life time earnings ($s_{i,2,t}$ is low, so the primary earner cannot count on the added worker effect contributing much to the smoothing of family earnings) or if her labor supply is relatively inelastic (η_{h_2, w_2} is small - for similar reasons).

parameters $\pi_{i,t}$ and β , on the human wealth shares $s_{i,j,t}$, the consumption EIS $\eta_{c,p}$, and the labor supply EIS of the two earners, η_{h_1,w_1} and η_{h_2,w_2} . Interpreting the role of $s_{i,j,t}$ is straightforward: when $s_{i,j,t}$ is large, the j -th earner's importance (in terms of his human wealth relative to the household's) is large, and hence consumption responds more to the permanent wage shock faced by this earner. *Ceteris paribus*, the sensitivity of consumption to the first earner's permanent wage shock (κ_{c,v_1}) is decreasing in the labor supply elasticity of the other earner (because in that case the added worker effect is stronger, and hence adjustment is partly done through increasing labor supply of the other earner); and it is decreasing in the own labor supply EIS if the response of hours of this earner to a shock is negative (i.e., if there is smoothing done through own labor supply, as discussed above). The sensitivity of consumption to a permanent shock also increases with $\eta_{c,p}$ because consumers with high values of the consumption EIS are by definition less reluctant to intertemporal fluctuations in their consumption.

Finally, note that the sensitivity of consumption to a permanent shock is higher whenever insurance through savings or other external sources is small ($\pi_{i,t}$ and β are low). The intuition is that the smaller is $\pi_{i,t}$ (β), the less assets (external insurance) the household has to smooth consumption when hit by a permanent shock of either spouse. It is indeed accumulation of these precautionary reserves that make consumption smoother than household earnings.

Let us now consider the interpretation of the coefficients under the non-separable preference assumption. Given that consumption and leisures can be complements or substitutes in utility, it is much more complicated to derive clear-cut comparative statics of the parameters $\kappa_{m,n}$ in the non-separable case (apart from the straightforward cases discussed above). A heuristic interpretation can be offered, though. Consider the approximation of the first order condition for consumption (see Appendix 3 for the derivation):

$$\Delta c_{i,t} \simeq (\eta_{c,w_1} + \eta_{c,w_2} - \eta_{c,p}) \Delta \ln \lambda_{i,t} + \eta_{c,w_1} \Delta w_{i,1,t} + \eta_{c,w_2} \Delta w_{i,2,t} \quad (17)$$

As originally remarked by Heckman (1974), the dynamic response of consumption to wage changes will depend on whether consumption and hours are complements or substitutes in utility. In particular, when C and H are substitutes ($\eta_{c,w_j} < 0$), we may have "Excess Smoothing" of consumption with respect to wage shocks; while complementarity ($\eta_{c,w_j} > 0$) may induce "Excess Sensitivity" (excess response to shocks

relative to the additive separable case). As an illustration, consider the case in which the primary earner faces a negative transitory wage shock (for which wealth changes are neutralized, or $\Delta \ln \lambda_{i,t} = 0$): $\Delta w_{i,1,t} < 0$ and $\Delta w_{i,2,t} = 0$. Under additive separability ($\eta_{c,w_j} = 0$), we would record a minimal decrease in consumption (in fact, in our set-up we imposed this to be zero) and a concurrent decrease in hours. When C and H are substitutes ($\eta_{c,w_1} < 0$), equation (17) shows that the consumption decrease is attenuated (it may even become an increase). Hence, consumption is smoother (there is more "insurance") in the presence of substitutability between consumption and hours.

2.4 Identification

There are four sets of parameters that we are interested in estimating: wage parameters, smoothing parameters, preference parameters, and measurement error variances. We discuss identification of these parameters in Appendix 4, both in terms of what moments we use and what kind of "variability in the data" we exploit to obtain our empirical estimates.

3 Data, Estimation Issues, and Empirical Strategy

3.1 The PSID Data

We use the 1999-2009 Panel Study of Income and Dynamics (PSID) to estimate the model. The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. The PSID data was collected annually until 1996 and biennially starting in 1997. A great advantage of PSID after 1999 is that, in addition to income data and demographics, it collects data about detailed assets holdings and consumption expenditures. To the best of our knowledge this makes the PSID the only representative large scale US panel to include both income, consumption, and assets data. Since we need both consumption and assets data, we focus on the 1997-2009 sample period.

For our baseline specification we focus on non-SEO households with participating and married male

household heads aged between 30 and 65. Whenever there is a change in family composition we drop the year of the change and treat the household unit as a new family starting with the observation following the change. We drop observations with missing values for state, education, race, labor earnings, hours, total consumption and total assets. We drop observations with wages that are lower than half the minimum wage in the state where the household resides. Finally, we drop observations for which consumption, wages or earnings of one of the earners show extreme "jumps" most likely due to measurement error. A "jump" is defined as an extremely positive (negative) change from $t-2$ to t , followed by an extreme negative (positive) change from t to $t+2$. Formally, for each variable (say x), we construct the biennial log difference $\Delta^2 \log(x_t)$, and drop observation in the bottom 0.25 percent of the product $\Delta^2 \log(x_t) \Delta^2 \log(x_{t-2})$.

3.1.1 Descriptive Statistics¹⁷

To estimate our model we need to construct a series of household consumption. Since we do not model the household decision on durables purchase, it is natural to focus on nondurables and services. Before 1999, PSID collected data on very few consumption items, such as food, rent and child care. However, starting in 1999 consumption expenditure data cover many other nondurable and services consumption categories, including health expenditures, utilities, gasoline, car maintenance, transportation, education and child care. A few other consumption categories have been added starting in 2005 (such as clothing). We do not use these categories to keep the consumption series consistent over time. The main items that are missing are clothing, recreation, alcohol and tobacco.

While rent is reported whenever the household rents a house, it is not reported for home owners. To construct a series of housing services for home owners we impute the rent expenditures for home owners using the self reported house price.¹⁸ We then aggregate all nondurable and services consumption categories to get the household consumption series.¹⁹ Descriptive statistics on the various components of aggregate consumption (nominal values) are reported in the upper part of Table 1. A comparison of the main aggregates (total consumption, nondurables, and services) against the NIPA series is offered in Table 2. As shown in

¹⁷For detailed list of consumption categories covered in the PSID in different years refer to http://psidonline.isr.umich.edu/Data/SL/ConsumptionQsinPSID_1968-2009.xls

¹⁸For our baseline measure we approximate the rent equivalent as 6% of the house price.

¹⁹We treat missing values in the consumption (and asset) subcategories as zeros.

Table 2, taking into account that the PSID consumption categories that we use are meant to cover 70% of consumption expenditure, the coverage rate is remarkably good.

Data on household's assets holdings is required for the construction of $\pi_{i,t}$, the share of assets out of total wealth. Starting in 1999, the PSID collects data on assets holding in each wave (between 1984 and 1999, asset data were collected every five years). The data include detailed holdings of cash, bonds, stocks, business, pensions, cars value, house and other real estate holdings. In addition, data is collected on household debt including first and second mortgage and other debt. Since we are interested in the net assets holdings, our measure of assets is constructed as the sum of cash, bonds, stocks the value of any business, the value of pension funds, the value of any house, the value of other real estate, the value of any car, net of any mortgage and other debts.

In addition to consumption and assets, data on wages and earnings of the first and second earner are also required. The survey collects data on annual labor earnings and on annual hours of work. To construct the hourly wage we divide annual earnings by annual hours.

In the lower part of Table 1 we provide summary statistics on asset holdings, and on labor supply and earnings for the two earners. It is worth noting that the female participation rate in this sample is fairly high (around 80%) and that on average they earn about half of what males earn, partly reflecting lower hours of work (conditional on working), and partly reflecting other factors, both explained and unexplained.

3.2 Estimation Issues

From an estimation point of view, we need to take a stand on a number of difficult issues. These include: (1) Allowing for measurement error in consumption, wages, and earnings; (2) Adopting the correct inference for our estimation procedure, and (3) Controlling for the selection into work of the secondary earner. We discuss these problems in the rest of this section.

3.2.1 Measurement Error

Consumption, wages and earnings are most invariably measured with error. In our context, there are three problems one need to confront when adding measurement errors. First, as discussed among others in Blundell

et al. (2008), adding measurement errors to models that include a permanent/transitory decomposition (as in our wage process) creates an identification problem, in that the distribution of the measurement error is indistinguishable from the distribution of the economically relevant transitory shock. Second, our wage measure is constructed as annual earnings divided by annual hours, and therefore the measurement errors of earnings and wages are correlated (the so-called "division bias"). Third, measurement errors are hard to distinguish from stochastic changes in preferences or shocks to higher moments of the distribution of wages in terms of effects on consumption or labor supply choices. We make no attempt to resolve this distinction, and hence identify an aggregate of these various forces, some statistical and some economic.

Ignoring the variance of measurement error in wages or earnings is problematic since it has a direct effect on the estimates of the structural parameters. We thus follow Meghir and Pistaferri (2004) and use findings from validation studies to set *a priori* the amount of wage variability that can be attributed to error. We use the estimates of Bound et al. (1994), who estimate the share of variance associated with measurement error using a validation study for the PSID (which is the data set we are using). Details are in the online Appendix 8.

3.2.2 Inference

We use multiple moments, which we deal with using a GMM strategy and an identity matrix as a weighting matrix. Given the multi-step approach, and the fact that we use longitudinal data, (unless explicitly noted) we compute the standard errors of our estimated parameters using the block bootstrap. In this way we account for serial correlation of arbitrary form, heteroskedasticity, as well as for the fact that we use pre-estimated residuals.²⁰

3.2.3 Selection Into Work by the Second Earner

Above, we have derived the expressions for earnings and hourly wage growth assuming interior solutions for labor supply for both spouses. A major concern when modeling labor supply is endogenous selection into work and therefore the need to distinguish between the intensive and the extensive margin of employment.

²⁰To avoid the standard errors being affected by extreme draws, we apply a normal approximation to the inter-quartile range of the replications.

Male participation is very high (for example in our sample, before conditioning on working, men between the age of 30 and 65 have average participation rates of 93%). This justifies our decision to focus on a sample of always-employed males. As for wives, their participation is 80% on average, and hence it is potentially important to account for their selection into work (see Table 1).

One approach is to explicitly model the decision to participate of the secondary earner. While appealing from a theoretical point of view, it makes the solution of the life cycle problem much more difficult - in fact, it would make our approximation procedure infeasible. We therefore decide to adopt a solution that is more statistical in nature, and in particular derive an empirical correction for the sample selection in the spirit of Low et al. (2010). We use “conditional covariance restrictions” rather than unconditional ones as done in most of the literature. Finding exclusion restrictions is the challenging part of this exercise. We use a set of state-year dummies intended to capture labor market related policy changes at the state level and the presence of first and second mortgage. There is some evidence showing that female participation rises when households move into home ownership (see Del Boca and Lusardi, 2003). Details are in Appendix 9.

3.3 Empirical Strategy

The following steps summarize our empirical strategy:

1. Regress the log difference of $C_{i,t}$, $Y_{i,1,t}$, $W_{i,1,t}$, $Y_{i,2,t}$ and $W_{i,2,t}$ onto observable characteristics and construct the first-differenced residuals $\Delta c_{i,t}$, $\Delta y_{i,1,t}$, $\Delta w_{i,1,t}$, $\Delta y_{i,2,t}$, $\Delta w_{i,2,t}$. The observable characteristics in the wage equation include year, year of birth, education, race, state and large city dummies as well as education-year, race-year and large city-year interactions. For consumption and earnings we also add dummies for number of kids, number of family member, employment status (at the time of interview), income recipient other than head or wife in the household and whether the couple has children not residing in the household. For observables which are not fixed over time we use both the level and the change. Note that the wage and earnings regressions use only workers;
2. Estimate the wage variances and covariances using the second order moments for $\Delta w_{i,1,t}$ and $\Delta w_{i,2,t}$;
3. Estimate the smoothing parameters $\pi_{i,t}$ and $s_{i,1,t}$ using asset and (current and projected) earnings

data;

4. Estimate the preference parameters using the second order moments for $\Delta y_{i,1,t}$, $\Delta y_{i,2,t}$ and $\Delta c_{i,t}$ conditioning on results (wage variances, covariances, and smoothing parameters) obtained in steps 2 and 3.

Our baseline specification uses only workers and does not correct for selection into work. In the robustness section we show that the correction for selection makes little difference. When we apply the sample selection correction described in section 3.2.3, we run the regressions that calculate residual measures for the wife's wages and earnings equations (step 1) controlling for selection into work (which is done by preliminarily running female employment probits and then constructing conventional Mills ratio terms). As said above, the exclusion restrictions we use are a set of state-year dummies intended to capture labor market related policy changes at the state level, and the presence of first and second mortgage.

4 Results

4.1 Estimating $\pi_{i,t}$ and $s_{i,1,t}$

The calculations of $\pi_{i,t} \approx \frac{\text{Assets}_{i,t}}{\text{Assets}_{i,t} + \text{Human Wealth}_{i,t}}$ and $s_{i,j,t} = \frac{\text{Human Wealth}_{i,j,t}}{\text{Human Wealth}_{i,t}}$ require the knowledge of assets, which we take directly from the data, and of expected human wealth at time t for both earners, i.e.:

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

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

$$Y_{i,1,t} = q_{i,1}^a \gamma_1 + q_{i,1,t}^b \gamma_2 + e_{i,1,t}$$

To obtain an estimate of expected earnings at $t + s$ given information at t (i.e., $\mathbb{E}_t(Y_{i,1,t+s})$) we simply use $\hat{Y}_{i,1,t+s} = q_{i,1}\hat{\gamma}_1 + m_{i,1,t+s}\hat{\gamma}_2$. We assume that agents are working until the age of 65 and that discounting rate is the same as the interest rate, and set the annual interest rate to 2%.

The same idea is applied to calculate expected human wealth for the secondary earner. However, since we allow for non-participation of the second earner, we run the earnings regressions controlling for selection using the Heckman correction. Moreover, to control for participation in the prediction of earnings, we use a probit specification with education, race, polynomial in age and interactions to predict the probability of participation for each second earner. The expected earnings for the wife at time $t + j$ are then the product of the predicted offered wages in period $t + j$ and the probability of being employed in that same period.

This procedure allows us to (pre-)estimate $\pi_{i,t}$ using asset and human capital data. The average value of these estimates is $E(\pi_{i,t}) = 0.18$; the age-specific averages are reported in Figure 1 (on the left axis), together with the life-cycle evolution of the household's total assets (on the right axis). These trends remain very similar if we use medians rather than means.

The estimates of $\pi_{i,t}$ conform to expectations. The degree of self-insurance warranted by asset accumulation is negligible at the beginning of the life cycle (all permanent shocks pass more or less through consumption), but the combination of asset accumulation due to precautionary and life cycle motives (visible from the evolution of the right axis variable) and the decline of expected human capital due to the shortening of the time horizon imply an increase in $\pi_{i,t}$ as time goes by, and hence the household's ability to smooth permanent wage shocks also increases over time. As the household head nears retirement, the average value of $\pi_{i,t}$ is as high as 0.5. What needs to be noted, however, is that this estimate - reflecting "actual" saving decisions of households - embeds all forms of insurance (or constraints to them) that households have available. In other words, there is no obvious way to benchmark the pattern shown in Figure 1. The closest equivalent is the hypothetical pattern presented by Kaplan and Violante (2011). We also estimate the pattern of $\pi_{i,t}$ by terciles of the asset distribution and find that the average value of $\pi_{i,t}$ increases with the rank in the wealth distribution, suggesting greater ability to smooth consumption among the wealthier, a result also found by Blundell et al. (2008).

Our estimate of $s_{i,1,t}$ (the ratio of the husband's human wealth to total household human wealth) are

plotted in Figure 2 against the head's age. These estimates can be interpreted as the life cycle evolution of the distribution of earnings power within the household. On average, the husband commands about 2/3 of total household human wealth. His weight rises initially due to fertility choices made by his wife. His weight declines at the end of the life cycle due to early retirement choices coupled with age differences within the household.

4.2 Main results

4.2.1 Wage Variances

The estimates of the wage variances are presented in Table 3. Three things are worth noting. Estimates of "wage instability", typically associated with the variance of the transitory component (see Gottschalk and Moffitt, 2008) are larger for males, perhaps reflecting a larger influence of turnover, etc. The variance of the more structural component (the variance of permanent shocks) is similar for males and females, although slightly higher for females, perhaps reflecting greater dispersion in the returns to unobserved skills etc. Finally, the transitory components of the two spouses are positively correlated, and (to a less extent) the permanent shocks as well, most likely reflecting the fact that (due perhaps to assortative mating) spouses tend to work in sectors, occupations, or even firms that are subject to similar aggregate shocks. These estimates are fairly noisy, however.

4.2.2 Consumption and Labor Supply Parameters

Table 4 reports the estimates of preference parameters. We start with the separable case in column 1. The elasticity of intertemporal substitution in consumption ($\eta_{c,p}$) is estimated to be 0.2, implying a relative risk aversion coefficient of around 5, which is fairly high, although in the ballpark of previous estimates from the literature. The elasticity of intertemporal substitution in male labor supply is smaller than that for female labor supply, again supporting previous evidence and intuition. For males, we estimate a Frisch elasticity $\eta_{h_1,w_1} = 0.4$; for females, we estimate $\eta_{h_2,w_2} = 0.8$. Our estimate of men's Frisch elasticity sits in the upper range of MaCurdy's (1981) estimates (0.1-0.45) and Altonji's (1986) estimates (0.08-0.54), which vary depending on the specification or set of instruments used. Keane (2011) surveys 12 influential studies and

reports an average estimate of 0.83 and a median estimate of 0.17.²¹ For women, Heckman and MaCurdy (1980) report an elasticity of substitution of 1 (p. 65), which is close to our estimate. The literature surveyed in Keane (2011) confirms, with a few exceptions, the finding of high Frisch elasticities for women.

Finally, we estimate $\beta = 0.74$, implying quite a large amount of "external" insurance over and above self-insurance. This estimate implies that there is very little responsiveness of consumption to permanent wage shocks even early on in the life cycle. This can be seen more clearly by looking at the estimated loading factor matrix from (10), which we report in column 1 of Table 5.²² These estimates show that theoretical restrictions are satisfied: the Frisch elasticities are positive and exceed the Marshallian elasticities, and the cross-labor supply elasticities are negative as they reflect only wealth effects. Nevertheless, there are two warning flags suggesting that this model may be misspecified. The first is that the estimate of β implies an "excessive" degree of excess smoothing. The second is that the fit of the model is far from being perfect. While the fit is excellent for hourly wage moments, we find that the variance of consumption growth is understated, the covariance of consumption growth and wages is overstated, and that the zero-restriction test on omitted moments rejects the null with a p-value of 0.014.²³

Next, we relax the assumption of separability and report the results in column 2 of Table 4. Allowing for non-separability introduces six new parameters into the analysis: the (λ -constant) elasticity of substitution of consumption with respect to the price of labor of the two spouses (η_{c,w_1} and η_{c,w_2}), which is informative about the possibility that consumption and leisure are non-separable; the (λ -constant) elasticity of substitution of earner i 's labor supply with respect to the price of j 's labor supply (η_{h_1,w_2} and η_{h_2,w_1}), which is informative about the possibility of non-separability between the leisure of the two spouses; and the (λ -constant) elasticity

²¹A number of other papers have challenged the notion of that the Frisch labor supply elasticity for males is close to zero. See for example Domeji and Floden (2006) and Wallenius (2011).

²²These loading factors, or transmission coefficients, are obtained replacing the estimates of the structural parameters in the relevant theoretical expressions. To give an example, the implied transmission coefficient κ_{c,v_1} reported in column 1 of Table 5 is obtained replacing the estimates of $\eta_{c,p}$, β , η_{h_1,w_1} , η_{h_2,w_2} , $\pi_{i,t}$ and $s_{i,1,t}$ (from Table 4) in the theoretical expression:

$$\kappa_{c,v_1} = \frac{\eta_{c,p} (1 - \beta) (1 - \pi_{i,t}) s_{i,1,t} (1 + \eta_{h_1,w_1})}{\eta_{c,p} + (1 - \beta) (1 - \pi_{i,t}) \bar{\eta}_{h,w}}$$

and then averaging across households and over the life cycle since $\pi_{i,t}$ and $s_{i,1,t}$ vary along these two dimensions.

²³These results remain overall fairly similar when we consider (i) correcting for the wife's non-participation decision as described in Section 3.2.3; when we change (ii) age selection criteria; or when we use (iii) individual rather than average assets when constructing $\pi_{i,t}$, and (iv) different weighting matrices in the GMM estimation procedure. These results are fully shown in Appendix 6.

of substitution of each earner's labor supply with respect to price of consumption ($\eta_{h_1,p}$ and $\eta_{h_2,p}$).

To increase efficiency of our estimates, we impose symmetry of the Frisch substitution matrix (see the online Appendix 7 for details). In Table 4, column 2, there are a number of results that are of some interest. First, additive separability is strongly rejected, as four of the new parameters are individually statistically significant. In particular, we find evidence of complementarity of husband and wife leisure (spouses enjoy spending time together), and we also find that both husband's and wife's leisure are complements with respect to household consumption. Second, the own-price Frisch elasticities are estimated to be larger than in the separable case. In particular, $\eta_{c,p}$ increases to 0.44. The labor supply elasticities also slightly rise. Finally, we no longer find evidence of "excess" outside insurance - the estimate of β is actually negative (-0.12), although statistically insignificant. Column 3 of Table 4, reports the results imposing that $\beta = 0$. The estimated parameters are very similar to the unrestricted case (often with smaller standard errors), stressing the small role played by β once allowing for non-separability.

The implied estimates of the transmission coefficients from (15) are reported in column 2 of Table 5. There are a few things to note. As before, the Frisch elasticities are positive and exceed the Marshallian elasticities and the cross-labor supply elasticities are negative (although they now reflect both wealth effects and non-separability with respect to the spouse's leisure). We also find larger consumption responses to permanent shocks (we offer an interpretation of these numbers in the next Section). Finally, there may be some worry that the response of consumption to transitory shocks reflect liquidity constraints rather than non-separability. However, with liquidity constraints the estimates of κ_{c,u_j} would be positive, not negative as we find. If liquidity constraints explain the behavior of consumption, then the implication is that we are even underestimating the degree of complementarity between consumption and leisure.²⁴

In terms of fit of the model, it improves substantially. The fit of wage moments is of course unchanged as they are estimated "outside" the model. We fit better the consumption variance and the covariance between earnings of the two spouses; however, there is no discernible improvement with respect to consumption covariances with the spouses' hourly wages. The model implies symmetries which are not rejected for the female moment in the data (p-value 0.31) while for the male the p-value is borderline (0.06).

²⁴In Appendix 7 we formally test the hypothesis that preferences are quasi-concave.

Finally, we recover Marshallian elasticities of labor supply for the husband and wife, and compare these to the Frisch elasticities. As mentioned above, in our estimation framework, recovering Marshallian elasticities is direct as these are given by the responses of hours to permanent shocks to wages $(\kappa_{y_j, v_j} - 1)$.²⁵ Allowing for non-separability (as in Table 4, column 2) the average Marshallian elasticity for males is very close to zero (-0.02 with a standard error of 0.13). We find a larger average elasticity of 0.32 for females (with a standard error of 0.09). As expected, these Marshallian elasticities are smaller than the corresponding Frisch elasticities. One advantage of recovering the Marshallian elasticities from the responses of hours to permanent shocks is that we can allow for heterogeneity in the elasticities as a function of household human and financial wealth (as reflected in $\pi_{i,t}$ and $s_{i,t}$). Figure 3 plots the Marshallian elasticities for both the husband and the wife against age. As is clear from the graph, late in the life cycle, as the household accumulate wealth, the role of the income effect is decreasing, driving the Marshallian elasticities up.²⁶

4.2.3 Robustness

We conducted a number of additional empirical exercises with the goal of assessing how robust our results are to some changes in sample selection and specification. The results are reported in Table 6. First, we focus on a sample that excludes older workers, focusing on heads aged 30-55. Second, we restrict our analysis to the high education group. Third, we allow for age-varying variances. Since transitory shocks do not play an important role on consumption, we allow only the distribution of permanent shocks to vary with age. Finally, we apply the selection correction, described in section 3.2.3.

In column (1), the estimates show that the degree of partial insurance accounted for by asset accumulation declines when we focus on a sample of younger workers who have had less time to accumulate assets (the estimate of $\pi_{i,t}$ on average decreases relative to the baseline case). The estimates of the other parameters remain very similar. In column (2), the estimate of $\pi_{i,t}$ increases on average reflecting more asset accumulation among the highly educated. We also find that men's labor supply is slightly more elastic and women's slightly less elastic than the baseline. Allowing for age-varying variances in column (3), does not

²⁵The theoretical relation between Marshallian and Frisch elasticities is well known. See for example Keane (2011), for a derivation in a constant elasticities set-up. The derivation in our setup is very similar, allowing for 2 earners and non-separability.

²⁶As $\pi \rightarrow 1$ the Marshallian elasticities are converging to their Frisch counterparts.

affect our conclusions much. Estimates of the permanent shock variances reveal that dispersion tends to increase with age, perhaps reflecting the growing incidence of health shocks. Finally, the estimates with selection correction, reported in column (4), are very similar to the results in the non-separable case without the correction (column (2) of Table 4).

5 Discussion

In this section we discuss our empirical findings. In particular, we focus on four issues: (a) advance information; (b) extensive and intensive margin in labor supply; (c) measuring how much of the insurance we estimate from the data can be explained by the various channels we allow for in our model, and (d) the role of non-linear taxation.

5.1 Advance Information

Some of the muted response of consumption to wage shocks may be due to wage changes not being shocks at all. In other words, consumers may have some advance information about shocks, and may have therefore adapted their consumption in advance of the shocks themselves. To test whether this is an explanation of our findings, we present a test of "superior or advanced information". We follow the intuition of Cunha et al. (2005) that with advanced information we should find that *future* wage growth predicts *current* consumption growth. We hence compute the covariances $E(\Delta c_{i,t} \Delta w_{i,j,t+\tau})$ for $\tau = \{2, 4, 6\}$ (as our panel is biennial) and test whether they are jointly insignificant (the null of no advanced information). The test does not reject the null of zero correlation with a p-value of 13%. We conclude that superior or advance information do not appear to be responsible for our findings.²⁷

²⁷There are two problems with this test. First, suppose that the true income process is a heterogeneous growth model and that the individual growth rate is known at time 0. In this case the correlation between current consumption growth and future consumption growth is going to be zero. However, this model would predict that also the correlation between current consumption growth and current income growth is zero, something that is clearly violated in our data. Second, the test is weak due to the fact that changes in income may reflect measurement error. It is worth noting, however, that if there is advance information about the permanent shocks, then the test will still be valid. Moreover, we are pre-adjusting our measure of income growth to account for measurement error.

5.2 Extensive vs Intensive Margin

In general, whether consumption and leisure are Frisch complements or substitutes is an empirical question. In this paper we find that there is evidence for complementarity: keeping constant the marginal utility of wealth, consumption and leisure tend to move in parallel. In the literature, evidence for complementarity has been rare (Browning et al., 1985). A more frequent finding is that of substitutability (see e.g., Aguiar and Hurst, 2005). However, it is worth noting that this evidence comes primarily from studying the relationship between changes in consumption and *large* changes in hours, often associated to events like retirement or unemployment, i.e., extensive margin shifts. For example, it is frequently found that consumption falls at retirement (when leisure increases dramatically in an anticipated manner). In this paper, in contrast, we have mainly focused on the relationship between changes in consumption and small changes in hours (i.e., intensive margin shifts).²⁸

Can we reconcile our "intensive margin" complementarity findings with the evidence of "extensive margin" Frisch substitutability in the literature? Consider the following example. Suppose that there are fixed costs associated with employment (i.e., when the extensive margin becomes active). For example a worker needs to buy a suit in order to show up at work; or pay a monthly road transit fee to travel to work. These costs exists independently of the number of hours worked - a suit is needed regardless of whether the worker works 1 or 40 hours, and the road transit fee must be paid regardless of whether the worker commutes 1 or 28 days a month to work. This is an example where consumption is *substitute* with respect to leisure (on the extensive margin). But the consumer's budget may include other goods that are Frisch *complements* with respect to leisure, such as utilities. The use of electricity or gas at home depends on the number of hours the worker actually spends at home or at work. Blundell et al. (2011a) derive a model with both margins.

To test in an informal way whether this story holds up in our data, we estimate "conditional" Euler equations, controlling for growth in hours (the intensive margin) and changes in participation (the extensive

²⁸Home production may induce substitutability also at the intensive margin. For example, individuals who are working a reduced number of hours may have more time to devote to home production and, if time and goods are substitutes, this may induce lower spending on goods needed to produce a given amount of consumption. Our finding that there is Frisch complementarity at the intensive margin suggests that this effect, while possibly present, is dominated by the alternative interpretation we offer here.

margin) - and instrumenting the two appropriately. The results are presented in Table 7. In the first column we use the PSID sample without conditioning on male participation, so we control for the growth in hours and the changes in participation of both spouses. To avoid the issues of zeros in hours for non-participant, we approximate the growth in hours with the expression $\Delta \ln h_t \approx \frac{h_t - h_{t-1}}{(1/2)(h_t + h_{t-1})}$. We use lagged (hours and participation) instruments dated $t - 2$ and $t - 4$ (recall the biennial nature of the PSID). The results seem consistent with the story above. For both males and females there is evidence of Frisch complementarity with leisure on the *intensive* margin (consumption falls when hours grow), confirming the results of the previous sections; however, consistent with most findings in the literature, we also find evidence that consumption and leisure are substitute at the *extensive* margin (consumption rises when participation rises). The estimates are more precise for females and remain so even when we focus on our estimation sample (always-working husbands), as shown in column 2. Reassuringly, the signs of the estimate do not change. In column 3 we tackle the issue of measurement error using only instruments dated $t - 4$. The point estimates are similar but the instruments are less powerful (information on the power of instrument is given in the last three columns of the table) resulting in less precise estimates overall. In conclusion, while some of these estimates are noisy, it appears to be able to reconcile the internal evidence of the previous sections with the external evidence coming from most of the literature once allowance is made for the distinction between intensive and extensive margin, a crucial one as such.

5.3 How Much Insurance?

The response of household earnings to a permanent shock to the male's hourly wage can be decomposed as follows:

$$\frac{\partial \Delta y}{\partial v_1} = s \frac{\partial \Delta y_1}{\partial v_1} + (1 - s) \frac{\partial \Delta y_2}{\partial v_1}$$

Empirically, a 10% permanent decrease in the husband's wage rate ($v_1 = -0.1$) induces a 4.4% decline in household earnings, since $\hat{s} = 0.69$, $\hat{\kappa}_{y_1, v_1} = 0.98$ and $\hat{\kappa}_{y_2, v_1} = -0.81$.

We can decompose this overall change into three effects: (a) earnings distribution within the family (s); (b) behavioral response of the husband ($\frac{\partial \Delta y_1}{\partial v_1}$); and (c) behavioral response of the wife ($\frac{\partial \Delta y_2}{\partial v_1}$).

If the husband was the sole source of income ($s = 1$) and if there were no behavioral response in terms of his labor supply ($\kappa_{y_1, v_1} = 1$), then a 10% permanent decrease in male wages will induce a fall in household earnings of the same amount. Introducing earnings distribution within the family ($s \neq 1$) (but still assuming away behavioral responses, i.e., $\kappa_{y_1, v_1} = 1, \kappa_{y_2, v_1} = 0$) would reduce the effect to 6.9%. Hence the mere presence of an additional earner, albeit supplying labor inelastically, acts as a significant source of income smoothing. Allowing for behavioral responses by the husband ($\kappa_{y_1, v_1} \neq 1$) makes little difference (household earnings decline by 6.8%), because the husband's labor supply to a permanent decline in wages is fairly inelastic (the substitution effect is approximately the same as the wealth effect). Finally, allowing for behavioral responses in female labor supply gives the baseline result (i.e., a 10% permanent decrease in the husband's wage rate induces a 4.4% decline in household earnings). This comes from the fact that the women works more when the male's wage falls permanently, a pure wealth effect. Clearly, this is a significant "insurance" channel in terms of household earnings.

What about the effect on consumption? From the empirical estimates, we know that $\frac{\partial \Delta c}{\partial v_1} = \hat{\kappa}_{c, v_1} = -0.38$ (see column 2 of Table 5), implying that a 10% permanent decrease in the male's wage rate decreases consumption by 3.8%. This insurance is for the most part coming from family labor supply (as we have just shown) and partly from self-insurance through savings and external sources. In fact, we can decompose the response of consumption into three steps: (a) with fixed labor supply and no savings or transfers, household earnings would fall by 6.9% as the male's wage falls permanently by 10%, and the fall in consumption would be of the same magnitude given the absence of self-insurance through savings etc.; (b) with the family labor supply insurance channel active, the fall in household earnings would be only 4.4% and again, so would be the consumption decline; (c) finally, with both insurance channels active, the fall in household earnings is 4.4%, but the fall in consumption is further attenuated, to 3.8%. In other words, of the 31 cents of consumption "insured" against the shock to the male's wage,²⁹ 25 cents (81% of the total insurance effect) come from family labor supply (she works, hence reducing s , and she increases her labor supply when his wage fall permanently) and only 6 cents (19% of the total insurance effect) come from self-insurance.

²⁹The figure 31 cents is derived from the difference between the response of consumption with insurance and labor supply response (a 3.8% decline) and without insurance and without labor supply response (but taking into account that $s_1 = 0.69$), i.e., a 6.9% decline.

In contrast, we find that the husband's labor supply is a relatively poorer insurance channel against shocks to the wife's wages. We can go through the same decomposition exercise, but this time considering a 10% permanent decline in the wife's wage (and focusing on the intensive margin response).

As before, the response of consumption can be decomposed working through three steps. First, with fixed labor supply and no savings or transfers, household earnings would fall by 3.1% as the woman's wage falls permanently by 10% (since her "weight" on total household earnings is $(1 - s) = 0.31$). The fall in consumption would be of the same magnitude given the absence of self-insurance through savings etc. Second, with the family labor supply insurance channel active, the fall in household earnings would be only slightly smaller (2.5%). This is for two reasons. First, the woman's labor supply declines a lot given her larger behavioral response. Second, although the husband's labor supply increases, this is not enough to keep household earnings stable due to his low behavioral responses (and despite his larger share in household earnings). Finally, with both insurance channels active, the fall in household earnings is 2.5%, but the fall in consumption is 2.1% (again see column 2 of Table 5). In other words, of the 10 cents of consumption "insured" against the shock to her wage, about 40% can be attributed to conventional insurance sources (savings and transfers) and 60% to family labor supply effects.

In Figure 4 we offer a graphical representation of the decomposition exercise focusing on the life cycle aspects. We focus on the experiment in which we let the permanent wage of the husband decline by 10%. Early on in the life cycle, essentially all consumption insurance (the "triangles" line) can be explained by labor supply responses (the "circle" line) as households do not have enough assets to smooth consumption through savings. As assets start to cumulate, though (after age 50), some of the insurance is taken up by saving, and the role of labor supply as an insurance device declines in importance.

How do our estimates compare to Blundell et al. (2008)? They consider the impact of (transitory and permanent) shocks to disposable income on consumption. In this paper, instead, we have looked at the impact of (transitory and permanent) shocks to hourly wages on consumption. In Blundell et al. (2008), a 10% permanent shock to disposable income induces a 6.5% decline in consumption. Suppose that disposable income includes only household earnings. We know that a 10% permanent shock to the male's hourly wage reduces household earnings by 4.4% and consumption by 3.8%. It follows that these estimates imply that

a 10% permanent shock to household earnings would reduce consumption by 8.7% with a standard error of 2.35%. One way to get this result is to divide the consumption response, 0.38, by the household earnings response, 0.44 (or $\frac{\partial \Delta c / \partial v_1}{\partial \Delta y / \partial v_1}$). While our estimate is higher compared to the estimate reported by Blundell et al. (2008), their estimate falls well within the 95% confidence interval of our estimate. Note, however, that these estimates are not strictly comparable. A decline in household earnings most likely reduces disposable income by less than 10% (due to taxes and transfers attenuating the fall in earnings). The next section addresses the impact of taxes in more detail.

5.4 Non-linear Taxes

So far, we have neglected any discussion of the impact of taxes on behavior. Introducing a proportional tax system will have little impact on our results, but non-linear progressive taxation may affect our estimates and interpretations because the optimal choice of consumption and hours of work is now done under a convex, non-linear budget set. Hence, the (after-tax) price of leisure changes with the amount of hours worked, inducing feedback effects and dampening the overall hours response to an exogenous shock to (before-tax) wages. The complications are exacerbated when considering the labor supply choice of a married couple that files for taxes jointly. In this case, variation in one earner's labor supply may change the marginal tax rate (and hence the return to work) faced by the other earner even in the presence of separable preferences between spouses' leisures.

Nonlinear progressive tax system induces kinks in the budget constraint. Taking fully into account tax progressivity in our framework is infeasible - as long as the goal is to study the transmission of wage shocks onto choices in a tractable manner (which we achieved through log-linear approximations of first order conditions and lifetime budget constraint). To partly account for tax progressivity, we adopt a further approximation procedure. In particular, we consider approximating the tax schedule faced by individuals using a continuous mapping from before- to after-tax earnings.³⁰ The mapping we use is:

³⁰This mapping is a simple algorithm for calculating the amount of taxes due, and hence after-tax income, starting from information on before-tax income. To provide an example, TAXSIM is another, albeit very sophisticated, mapping. To perform our mapping, we disregard income from sources other than labor (apart from EITC payments). Given this information, we calculate $\tilde{Y}_{it}$ using information on the tax schedule for the various tax years. Hence, what we do is akin to a log-linearization of the tax system.

$$\tilde{Y}_{it} = (1 - \chi_t) (H_{1,t}W_{1,t} + H_{2,t}W_{2,t})^{1-\mu_t} \quad (18)$$

where the parameters χ and μ vary over time to reflect changes in the degree of progressivity of the tax system, and $\tilde{Y}_{it}$ is after-tax family earnings (see Heathcote et al., 2009). In a proportional tax system μ_t will be zero and χ_t will represent the proportional tax rate. Researchers have proposed a number of alternative mappings (see Carroll and Young, 2011, and the references therein). We prefer this mapping because it provides a simple log-linear relationship between after- and before-tax income.³¹

In Appendix 5 we show that the lifetime budget constraint is rewritten with $\tilde{Y}_{it}$ used in place of $(H_{1,t}W_{1,t} + H_{2,t}W_{2,t})$. We can use assumption (18) and thus extend our previous analysis to account for progressive taxation. In particular, the approximated first order conditions of the problem (equivalent to equations (7) and (8) in the non-separable case) are:

$$\Delta c_{t+1} \simeq (-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2}) \Delta \ln \lambda_{t+1} + \eta_{c,w_1} \Delta w_{1,t+1} + \eta_{c,w_2} \Delta w_{2,t+1} \quad (19)$$

$$\begin{aligned} & -\mu_{t+1} (\eta_{c,w_1} + \eta_{c,w_2}) \Delta y_{t+1} \\ \Delta h_{j,t+1} & \simeq (\eta_{h_j,p} + \eta_{h_j,w_j} + \eta_{h_j,w_{-j}}) \Delta \ln \lambda_{t+1} + \eta_{h_j,w_j} \Delta w_{j,t+1} + \eta_{h_j,w_{-j}} \Delta w_{-j,t+1} \\ & -\mu_{t+1} (\eta_{h_j,w_j} + \eta_{h_j,w_{-j}}) \Delta y_{t+1} \end{aligned} \quad (20)$$

for $j = \{1, 2\}$. These equations show very clearly the feedback effect of taxes. To make the discussion more transparent, consider the labor supply equilibrium condition in the single-earner case and assume for simplicity that the tax system is stationary:

$$\Delta h_{j,t+1} \simeq (\eta_{h_j,p} + \eta_{h_j,w_j}) \Delta \ln \lambda_{t+1} + \eta_{h_j,w_j} \Delta w_{j,t+1} - \mu \eta_{h_j,w_j} (\Delta w_{j,t+1} + \Delta h_{j,t+1})$$

Consider first the case $\mu = 0$. In this case there is no distinction between after-tax and before-tax earnings (as it would happen if individual behavior was unaffected by the progressivity of taxes). Hence, the equilibrium conditions in the progressive tax case collapse to those we have been using so far. However, in the more realistic case $\mu > 0$, the response of hours to own wage shocks is no longer given by the Frisch

³¹In Appendix 5 we show this mapping to provide an accurate approximation.

elasticity η_{h_j, w_j} , but by the parameter $\tilde{\eta}_{h_j, w_j} = \frac{\eta_{h_j, w_j}(1-\mu)}{1+\mu\eta_{h_j, w_j}}$ (with $\tilde{\eta}_{h_j, w_j} \leq \eta_{h_j, w_j}$ for $0 \leq \mu \leq 1$), which we call the Frisch elasticity with respect to before-tax wage changes. Clearly, when μ increases, the elasticity of labor supply with respect to before-tax wage changes is dampened relative to the no-tax or flat-tax case, because any labor supply increase induced by an exogenous increase in before-tax wages would be attenuated by a decrease in the return to work as people cross tax brackets (which they do "continuously" in our case).³²

In the consumption case, one can calculate that the response of consumption to a before-tax transitory wage shock is no longer η_{c, w_j} (which was identifying the extent of non-separability between consumption and leisure) but equals $\frac{\eta_{c, w_1}(1-\mu)}{1+\mu\eta_{h_1, w_1}}$, and hence it is also dampened (in absolute value). The reason is that this coefficient captures the extent of consumption co-movement with hours, but in the case with taxes hours move less and this lower sensitivity of hours to wage shocks spills over to a lower sensitivity of consumption to wage shocks induced by preference non-separability.

We re-run our approximation procedure accounting for taxes and obtain the equivalent of (15). See Appendix 5 for full details. The transmission parameters, $\kappa_{k,j}$ in equation (15), change because both $\pi_{i,t}$ and $s_{i,t}$ change as a consequence of introducing taxes into the picture. However, the changes are negligible and are discussed in Appendix 5. The estimates of the Frisch elasticities (which are now capturing the response of labor supply with respect to after-tax wage changes), of the average π , and of β are reported in column 4 of Table 4. The estimates of the Frisch elasticities are typically larger (in absolute values) than in the flat tax case (column 2), because those that we estimated without accounting for taxes were downward biased - the feedback effect of taxes was already there, but we interpreted it instead as a low elasticity of response. Nevertheless, it is worth noting that introducing taxes does not affect our qualitative results and the amount of bias is small.

Researchers interested in the effect of taxes on labor supply may want to distinguish between the Frisch (Marshallian) elasticity of labor supply with respect to before-tax changes in wages $\tilde{\eta}$ and the Frisch (or Marshallian) elasticity of labor supply with respect to after-tax changes in wages η (MaCurdy, 1983). In the simple single-earner case, these objects are linked through the relationship $\frac{\tilde{\eta}}{1-\mu(1+\tilde{\eta})} = \eta$.

³²When the tax system changes over time and we allow for multiple earners, the elasticity to before-tax wage changes $\tilde{\eta}$ will vary to reflect time and individual characteristics (earnings share within the household). In Tables 5 and 8, we report a value of $\tilde{\eta}$ averaged across all periods and individuals.

We compare the before- and after-tax elasticities in Table 8, both for the Frisch and the Marshallian case. None of our conclusions are materially affected by the introduction of progressive taxation. There is, however, an important change of interpretation. Non-linear progressive taxation effectively attenuates the impact of wage shocks and consequently, when non-linear taxation is ignored, the Frisch elasticities are underestimated. Nevertheless, the before-tax elasticities and the smoothing parameters change very little.

6 Conclusions

This paper estimate a life cycle model with two earners making consumption and labor supply decisions. We allow for flexible preferences (possibly non-separable among all arguments of the utility function, consumption and leisure of the two spouses), correlated wage shocks, and use approximations of the first order conditions and the lifetime budget constraint to derive expressions linking changes in consumption and hours to wage shocks. The sensitivity of consumption and hours to shocks depend on the structural parameters of the problem (Frisch elasticities and cross-elasticities), as well as terms that measure the relevance of self-insurance, insurance through external channels, and through family labor supply. We reject separability. We also reject advance information as an explanation for consumption smoothing relative to wage shocks. Once we allow for nonseparable preferences, we find no evidence of additional insurance channels. Most of the consumption smoothing we observe can be explained by decisions that are within the boundaries of the household, i.e., an extended view of self-insurance. We find a particularly important role for family labor supply, and calculate that of the 31 cents of consumption "insured" against the shock to the male's wage, 81% come from family labor supply and only 19% comes from self-insurance. We find a smaller insurance role for the husband's labor supply, calculating that of the 10 cents of consumption "insured" against the shock to her wage, about 40% can be attributed to conventional insurance sources (savings and transfers) and 60% to family labor supply effects.

Our work could be fruitfully extended in a number of directions. Here we suggest a few avenues. First, it is important to understand the role played by liquidity constraints in affecting consumption and labor supply choices. In our framework, consumption responds to transitory shock, but while liquidity constraints predict a positive response to transitory shocks, we find that the response is negative and interpret this as evidence

for complementarity between leisure and household consumption. It is possible that complementarity is even higher and this masks a role for liquidity constraints (perhaps concentrated among low wealth households). Future work should aim at disentangling these two distinct forces. Second, we need to understand the role of nonseparability of consumption and hours separately from the effect of fixed cost of work. Third, intra-family allocation issues have been neglected. This is not because we think they are unimportant, but because identification is extremely challenging and is only now started to being confronted with more appropriate data (i.e., spending on "exclusive" goods) and methodologies. Finally, we have assumed that hours can be freely adjusted in response to wage shocks, but with adjustment costs in hours this is less obvious. Our results, suggesting an important role for family labor supply in self-insuring household consumption against wage shocks would be presumably even more prominent if adjustment costs in labor supply were important.

References

- [1] Aguiar, Mark and Erik Hurst (2005), "Consumption vs. Expenditure", *Journal of Political Economy*, 113(5), 919-948
- [2] Altonji, Joseph (1986), "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy* 94(3 Part 2), S176-S215.
- [3] Attanasio, Orazio, Gabriella Berloffa, Richard Blundell and Ian Preston (2002), "From Earnings Inequality to Consumption Inequality." *Economic Journal*, 112(478), C52-C59.
- [4] Attanasio, Orazio, Hamish Low and Virginia Sanchez-Marcos (2008), "Explaining changes in female labour supply in a life-cycle model", *American Economic Review*, 98(4), 1517-1552.
- [5] Blundell, Richard and Thomas MaCurdy (1999), "Labor supply: A review of alternative approaches", *Handbook of labor economics* 3, O. Ashenfelter and D. Card (Eds.), 1559-1695.
- [6] Blundell, Richard, and Ian Preston (1998), "Consumption Inequality and Income Uncertainty." *Quarterly Journal of Economics*, 113(2), 603-640.

- [7] Blundell, Richard and Ian Preston (2004), "Consumption inequality and income uncertainty: The case with endogenous family labor supply", unpublished manuscript, University College London.
- [8] Blundell Richard, Luigi Pistaferri and Ian Preston (2008), "Consumption inequality and partial insurance." *American Economic Review*, 98(5), 1887-1921.
- [9] Blundell, Richard, Antoine Bozio and Guy Laroque (2011a). "Labor Supply and the Extensive Margin." *American Economic Review*, 101(3), 482-86.
- [10] Blundell, Richard, Hamish Low and Ian Preston (2011b). "Decomposing Changes in Income Risk Using Consumption Data", IZA Working Papers, 6125.
- [11] Bound John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers (1994), "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics*, 12(3), 345-368.
- [12] Browning, Martin, Angus Deaton, and Margaret Irish (1985), "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle", *Econometrica*, 5(3), 503-44.
- [13] Browning, Martin, Lars P. Hansen, and James Heckman (1999), "Micro Data and General Equilibrium Models." in *Handbook of Macroeconomics*, 1A: 543-633.
- [14] Browning, Martin and Costas Meghir (1991), "The Effects of Male and Female Labor Supply on Commodity Demands," *Econometrica*, 59(4), 925-51.
- [15] Caballero, Ricardo (1990), "Consumption Puzzles and Precautionary Savings." *Journal of Monetary Economics*, 25(1), 113-136.
- [16] Carroll, Daniel R. and Eric R. Young (2011), "The long run effects of changes in tax progressivity," *Journal of Economic Dynamics and Control*, 35(9), 1451-1473.
- [17] Chiappori, Pierre-Andre (1988), "Rational household labor supply", *Econometrica*, 56(1), 63-90.
- [18] Christiano, Larry J, Martin Eichenbaum and Sergio Rebelo (2011), "When is the Government Spending Multiplier Large?." *Journal of Political Economy*, 119(1), 78-121.

- [19] Cunha, Flavio, James Heckman and Salvador Navarro (2005), "Separating Uncertainty from Heterogeneity in Life Cycle Earnings", *Oxford Economic Papers*, 57(2), 191-261.
- [20] Del Boca, Daniela and Annamaria Lusardi (2003), "Credit market constraints and labor market decisions." *Labour Economics*, 10(6), 681-703.
- [21] Domeij, David and Martin Floden (2006). "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased", *Review of Economic Dynamics*, 9(2), 242-262.
- [22] Gottschalk, Peter and Robert A. Moffitt (2008), "Trends in the Transitory Variance of Male Earnings in the U.S., 1970-2004". Boston College Working Papers 697.
- [23] Guvenen, Fatih. (2007), "Learning your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97(3), 687-712.
- [24] Guvenen, Fatih and Anthony Smith (2010), "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach." Staff Report 450, Federal Reserve Bank of Minneapolis.
- [25] Hall, Robert E. (1978), "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86(2), 971-87.
- [26] Hall, Robert, and Frederic Mishkin (1982), "The sensitivity of consumption to transitory income: Estimates from panel data of households", *Econometrica*, 50(2), 461-81.
- [27] Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante. (2009). "Consumption and labor supply with partial insurance: An analytical framework". Mimeo, New York University.
- [28] Heckman, James J. (1974), "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption Over the Life Cycle", *The American Economic Review*, 64(1), 188-194.
- [29] Heckman, James J. and Thomas E. Macurdy (1980), "A Life Cycle Model of Female Labour Supply", *The Review of Economic Studies*, 47 (1), 47-74.

- [30] Hyslop, Dean R. (2001). "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings." *American Economic Review*, 91(4), 755-77.
- [31] Hryshko, Dmytro (2011). "Excess Smoothness of Consumption in an Estimated Life Cycle Model." Mimeo, University of Alberta.
- [32] Jappelli, Tullio and Luigi Pistaferri. (2010). "The Consumption Response to Income Changes." *Annual Reviews in Economics*, 2, 479-506.
- [33] Juhn, Chinhui and Simon Potter (2007). "Is There Still an Added Worker Effect?". Federal Reserve Bank of New York. Staff Report no. 310.
- [34] Kaplan, Greg and Giovanni Violante (2010), "How Much Consumption Insurance Beyond Self-Insurance." *American Economic Journal*, 2(4), 53-87.
- [35] Kaplan, Greg and Giovanni Violante (2011), "A Model of the Consumption Response to Fiscal Stimulus Payments," NBER Working Papers 1733.
- [36] Kaufmann, Katja and Luigi Pistaferri (2009). "Disentangling Insurance and Information in Intertemporal Consumption Choices." *American Economic Review*, 99(2), 387-92
- [37] Keane, Michael P. (2011). "Labor Supply and Taxes: A Survey." *Journal of Economic Literature*, 49(4), 961-1075.
- [38] Krueger, Dirk, and Fabrizio Perri. (2006). "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies*, 73(1), 163-193.
- [39] Low, Hamish (2005), "Self-insurance in a life-cycle model of labour supply and savings" , *Review of Economic Dynamics* 8(4), 945-975.
- [40] Low, Hamish, Costas Meghir and Luigi Pistaferri, (2010), "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review*, 100(4), 1432-67.
- [41] Lundberg, Shelly (1985), "The Added Worker Effect", *Journal of Labor Economics*, 3(1), 11-37.

- [42] MaCurdy, Thomas E. (1981), "An Empirical Model of Labor Supply in a Life-Cycle Setting", *The Journal of Political Economy*, 89(6), 1059-1085.
- [43] MaCurdy, Thomas E. (1983), "A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty", *International Economic Review*, 24(2), 265-289.
- [44] Meghir, Costas and Luigi Pistaferri (2004), "Income variance dynamics and heterogeneity", *Econometrica* 72(1), 1-32.
- [45] Meghir, Costas and Luigi Pistaferri. (2011). "Earnings, consumption and lifecycle choices", *Handbook of Labor Economics* , 4(B), 773-854.
- [46] Primiceri , Giorgio E. and Thijs van Rens. (2009) "Heterogeneous life-cycle profiles, income risk and consumption inequality", *Journal of Monetary Economics*. 56(1), 20-39
- [47] Stephens, Mel (2002), "Worker Displacement and the Added Worker Effect", *Journal of Labor Economics*, 20(3), 504-537.
- [48] Wallenius, Johanna, (2011). Human capital accumulation and the intertemporal elasticity of substitution of labor: How large is the bias? *Review of Economic Dynamics*, 14(4), 577–591

Table 1: **Descriptive Statistics**

	1998	2000	2002	2004	2006	2008
Consumption	27,290	31,973	35,277	41,555	45,863	44,006
Non durable Cons.	6,859	7,827	7,827	8,873	9,889	9,246
Food at home	5,471	5,785	5,911	6,272	6,588	6,635
Gasoline	1,387	2,041	1,916	2,601	3,301	2,611
Services	21,319	25,150	28,419	33,755	36,949	35,575
Food out	2,029	2,279	2,382	2,582	2,693	2,492
Health ins.	1,056	1,268	1,461	1,750	1,916	2,188
Health serv.	902	1,134	1,334	1,447	1,615	1,844
Utilities	2,282	2,651	2,702	4,655	5,038	5,600
Trasportation	3,122	3,758	4,474	3,797	3,970	3,759
Education	1,946	2,283	2,390	2,557	2,728	2,584
Child care	601	653	660	689	648	783
Home ins.	430	480	552	629	717	729
Rent (or rent eq.)	8,950	10,465	12,464	15,650	17,623	15,595
Total assets	332,625	352,247	382,600	476,626	555,951	506,823
Housing and RE	159,856	187,969	227,224	283,913	327,719	292,910
Financial assets	173,026	164,657	155,605	192,995	228,805	214,441
Total debt	72,718	82,806	98,580	115,873	131,316	137,348
Mortgage	65,876	74,288	89,583	106,423	120,333	123,324
Other debt	7,021	8,687	9,217	9,744	11,584	14,561
First earner (head)						
Earnings	54,220	61,251	63,674	68,500	72,794	75,588
Hours worked	2,357	2,317	2,309	2,309	2,284	2,140
Second earner						
Participation rate	0.81	0.80	0.81	0.81	0.81	0.80
Earnings Work	26,035	28,611	31,693	33,987	36,185	39,973
Hours worked Work	1,666	1,691	1,697	1,707	1,659	1,648
Observations	1,872	1,951	1,984	2,011	2,115	2,221
Notes: PSID data from 1999-2009 PSID waves. PSID means are given for the main sample: married couples with working male aged 30-65. SEO sample excluded. PSID rent is imputed as 6% of reported house value. Missing values in consumption and assets subcategories are treated as zeros.						

Table 2: **Comparison of PSID data with NIPA**

	1998	2000	2002	2004	2006	2008
PSID Total	3,276	3,769	4,285	5,058	5,926	5,736
NIPA Total	5,139	5,915	6,447	7,224	8,190	9,021
ratio	0.64	0.64	0.66	0.70	0.72	0.64
PSID Nondurables	746	855	887	1,015	1,188	1,146
NIPA Nondurables	1,330	1,543	1,618	1,831	2,089	2,296
ratio	0.56	0.55	0.55	0.55	0.57	0.50
PSID Services	2,530	2,914	3,398	4,043	4,738	4,590
NIPA Services	3,809	4,371	4,829	5,393	6,101	6,725
ratio	0.66	0.67	0.70	0.75	0.78	0.68

Note: We use PSID weights (we have a total of 47,206 obs. for 1999-09). Total consumption is defined as Nondurables + Services. PSID consumption categories include food, gasoline, utilities, health, rent (or rent equivalent), transportation, child care, education and other insurance. NIPA numbers are from NIPA table 2.3.5. All figures are in nominal terms.

Table 3: **Variance Estimates**

Sample	All		
Males	Trans.	$\sigma_{u_1}^2$	0.033 (0.008)
	Perm.	$\sigma_{v_1}^2$	0.032 (0.005)
Females	Trans.	$\sigma_{u_2}^2$	0.012 (0.006)
	Perm.	$\sigma_{v_2}^2$	0.043 (0.005)
Correlation of shocks	Trans.	σ_{u_1,u_2}	0.244 (0.164)
	Perm	σ_{v_1,v_2}	0.113 (0.082)
Observations	8,191		
Notes: Wage process parameters estimated using GMM.			
Block bootstrap standard errors in parenthesis.			

Table 4: Parameters Estimates

	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	Additive separability	Non- separability	Non- separability	Nonlinear taxes	Additive separability	Non- separability	Non- separability	Nonlinear taxes
	$\beta = 0$				$\beta = 0$			
π, β and own elasticities:								
$E(\pi)$	0.181 (0.008)	0.181 (0.008)	0.181 (0.008)	0.193 (0.005)	-.-	-0.141 (0.051)	-0.141 (0.053)	-0.186 (0.067)
β	0.741 (0.165)	-0.120 (0.298)	0	-0.140 (0.351)	-.-	0.082 (0.030)	0.082 (0.031)	0.108 (0.039)
$\eta_{c,p}$	0.201 (0.077)	0.437 (0.124)	0.448 (0.126)	0.495 (0.151)	-.-	-0.138 (0.139)	-0.158 (0.121)	-0.175 (0.146)
η_{h_1,w_1}	0.431 (0.097)	0.514 (0.150)	0.497 (0.150)	0.621 (0.176)	-.-	0.162 (0.166)	0.185 (0.145)	0.206 (0.176)
η_{h_2,w_2}	0.831 (0.133)	1.032 (0.265)	1.041 (0.275)	1.133 (0.289)	-.-	0.128 (0.052)	0.12 (0.064)	0.197 (0.061)
η_{h_2,w_1}					-.-	0.258 (0.103)	0.242 (0.119)	0.397 (0.127)
Observ.	8,191	8,191	8,191	8,191				

Cross-elasticities:

Notes: Parameters estimated using GMM. Column 1 reports the estimates for the additive separable case. Column 2 allows for non-separability of hours of the two earners and for non-separability of hours and consumption. Column 3 is similar to column 2, but restricts insurance over and above savings to be absent ($\beta = 0$). Column 4 is similar to column 2 but allows for non-linear taxes (see section 5.4). Block bootstrap standard errors in parenthesis.

Table 5: **Implied Transmission Coefficients**

	(1)	(2)
κ_{c,v_1}	0.13 (0.060)	0.38 (0.057)
κ_{c,v_2}	0.07 (0.040)	0.21 (0.037)
κ_{c,u_1}	0	-0.14 (0.051)
κ_{c,u_2}	0	-0.14 (0.139)
κ_{y_1,v_1}	1.15 (0.067)	0.98 (0.131)
κ_{y_1,v_2}	-0.16 (0.057)	-0.23 (0.048)
κ_{y_1,u_1}	1.43 (0.097)	1.51 (0.150)
κ_{y_1,u_2}	0	0.13 (0.051)
κ_{y_2,v_1}	-0.54 (0.206)	-0.81 (0.180)
κ_{y_2,v_2}	1.53 (0.101)	1.32 (0.087)
κ_{y_2,u_1}	0	0.26 (0.103)
κ_{y_2,u_2}	1.83 (0.133)	2.03 (0.265)

Note: The estimated transmission coefficients in columns (1) and (2) are obtained using the estimates of the structural parameters reported in columns (1) and (2) of Table 4, respectively. Reported values are averages across households and time periods. Block bootstrap standard errors in parenthesis.

Table 6: **Robustness Checks**

	(1)	(2)	(3)	(4)		(1)	(2)	(3)	(4)
	Age	Some	Flexible	Selection		Age	Some	Flexible	Selection
	30-55	college	variance	Correction		30-55	college	variance	Correction
		or more	by age				or more	by age	
π, β and own elasticities:					Cross-elasticities:				
$E(\pi)$	0.142	0.202	0.181	0.176	η_{c,w_1}	-0.113 (0.018)	-0.162 (0.022)	-0.15 (0.018)	-0.15 (0.017)
β	-0.177 (0.089)	0.117 (0.072)	-0.109 (0.077)	-0.129 (0.076)	$\eta_{h_1,p}$	0.065 (0.01)	0.087 (0.012)	0.087 (0.01)	0.88 (0.01)
$\eta_{c,p}$	0.465 (0.044)	0.368 (0.05)	0.42 (0.037)	0.473 (0.041)	η_{c,w_2}	-0.083 (0.029)	-0.142 (0.032)	-0.11 (0.026)	-0.122 (0.028)
η_{h_1,w_1}	0.467 (0.036)	0.542 (0.045)	0.575 (0.04)	0.509 (0.038)	$\eta_{h_2,p}$	0.097 (0.034)	0.169 (0.038)	0.129 (0.038)	0.143 (0.033)
η_{h_2,w_2}	1.039 (0.099)	0.858 (0.097)	1.005 (0.086)	1.095 (0.092)	η_{h_1,w_2}	0.101 (0.011)	0.115 (0.012)	0.141 (0.011)	0.125 (0.01)
					η_{h_2,w_1}	0.205 (0.022)	0.255 (0.027)	0.285 (0.022)	0.253 (0.021)
Observ.	6,942	5,014	8,191	8,191					

Notes: Parameters estimated using GMM. All columns allow for non-separability of hours of the two earners and for nonseparability of hours and consumption. In column 1 the sample is restricted to households with heads aged 30-55. In column 2 the sample is restricted to households with heads that have at least some college education. In column 3, we use the baseline sample, and allow the variance of the permanent shock to change with age (using five age groups). Column 4 uses the baseline sample and applies the participation correction (see section 3.2.3). Second stage GMM standard errors in parenthesis.

Table 7: **Conditional Euler Equations**

	Regression results			First stage F-stats		
	(1)	(2)	(3)	(1)	(2)	(3)
$\Delta EMP_t(Male)$	0.144 (0.369)			23.4		
$\Delta h_t(Male)$	-0.073 (0.175)	-0.013 (0.021)	-0.014 (0.064)	26.3	135.5	6.0
$\Delta EMP_t(Female)$	0.456 (0.199)	0.362 (0.186)	0.232 (0.892)	98.4	91.2	12.6
$\Delta h_t(Female)$	-0.220 (0.100)	-0.171 (0.094)	-0.128 (0.314)	86.5	77.7	18.5
Sample	All	EMP _t (Male)=1	EMP _t (Male)=1			
Instruments	2 nd ,4 th lags	2 nd ,4 th lags	4 th lag			
Observations	7,247	6,678	6,678			


Notes: The table reports “conditional” Euler equations estimates. Δx_t is defined as $(x_t - x_{t-1}) / [0.5(x_t + x_{t-1})]$. In columns 1 and 2, hours growth and change in employment are instrumented using the second and the fourth lags of hours and employment, as well as the two-year change in average wages by cohort, education and year, and the second lag of this change. In column 3 the two-year lag of hours and participation is omitted. All specifications also control for year effects, change in family size, number of kids, # of kids outside the household, extra earner, and leaving in SMSA. Column 1 does not condition on heads employment, (but does condition on non-missing lagged instruments). Columns 2 and 3 also condition on head’s employment. Standard errors are clustered at the household level.

Table 8: **The effect of taxes on labor supply elasticities**

	(1)	(2)	(3)
	No taxes	Before-tax	After-tax
	accounted for	response	response
η_{h_1, w_1}	0.514 (0.150)	0.621 (0.176)	0.530 (0.147)
η_{h_2, w_2}	1.032 (0.265)	1.133 (0.289)	1.033 (0.260)
η_{h_1, w_2}	0.128 (0.052)	0.197 (0.061)	0.143 (0.055)
η_{h_2, w_1}	0.258 (0.103)	0.397 (0.127)	0.228 (0.100)
η_{c, w_1}	-0.141 (0.051)	-0.186 (0.067)	-0.146 (0.052)
η_{c, w_2}	-0.138 (0.139)	-0.175 (0.146)	-0.152 (0.136)
Marshallian Elasticities (w.r.t. own wage shocks)			
Head	-0.018 (0.131)	-0.011 (0.129)	-0.010 (0.123)
Wife	0.322 (0.087)	0.322 (0.058)	0.370 (0.078)
Observ.	8,191	8,191	8,191


Notes: Parameters estimated using GMM. Column 1 reports the elasticities for the nonseparable case not accounting for taxes (as in Table 4, column 2). Column 2 reports the before-tax elasticities for the nonseparable case (as in Table 4, column 4). Column 3 reports the after-tax elasticities for the nonseparable case.

Figure 1: π by Age of Head of Household




Notes: The figure plots the average value of π and the average Total assets by age categories.

Figure 2: s by Age of Head of Household



Notes: The figure plots the average value of s .

Figure 3: Marshallian Elasticities by Age



Note: The figure plots the average Marshallian labor supply elasticities (κ_{h_1, v_1} and κ_{h_2, v_2}) for the head and for the wife over the life cycle, respectively.

Figure 4: Decomposition of consumption smoothing by age

Note: The response of consumption in the "fixed labor supply and no insurance" case ("diamonds"), is calculated as s times the shock. The response of consumption in the "family labor supply adjustment only" case ("circles") is calculated as the response of household earnings to the shock times the shock. The response of consumption in the "family labor supply adjustment and other insurance" case is calculated as κ_{c, v_1} , times the shock.