

Consumption and Labor Supply with Partial Insurance: An Analytical Framework*

Jonathan Heathcote

Federal Reserve Board of Governors, Georgetown University, and CEPR
jhh9@georgetown.edu

Kjetil Storesletten

University of Oslo, Frisch Centre (Oslo), and CEPR
kjstore@econ.uio.no

Giovanni L. Violante

New York University, NBER and CEPR
glv2@nyu.edu

First draft: October 2006 – This draft: August 2007

Abstract

This paper develops an analytical framework to study consumption and labor supply in heterogeneous-agent economies with partial insurance. The environment allows for insurance beyond non-contingent bonds, for permanent and transitory idiosyncratic productivity shocks, and for permanent preference heterogeneity and idiosyncratic preference shocks. Exact closed-form solutions are obtained for equilibrium allocations and for the first and second moments of the equilibrium joint distribution over wages, hours and consumption. With these expressions in hand, we show that all the structural preference and risk parameters in the model can be identified, even when productivity risk varies over time, given panel data on wages and hours, and cross-sectional data on consumption. We estimate the model on CEX and PSID data for the U.S. economy over the period 1967-1996. We then use the estimated parameter values to decompose inequality in all variables of interest, both over the life-cycle and across time, into cross-sectional variation in preferences, uninsurable wage risk, insurable wage risk, and measurement error.

*We are grateful to Jeffrey Campbell, Luigi Pistaferri and Thomas Sargent for useful comments, and to Greg Kaplan for outstanding research assistance. We thank the Federal Reserve Banks of Chicago and Minneapolis for their hospitality at various stages of this project. Heathcote and Violante's research is supported by a grant of the National Science Foundation (SES 0418029). The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve System.

1 Introduction

How does labor market risk at the individual level transmit to consumption, leisure, and, ultimately, to household welfare? The answer depends on the persistence of labor market shocks, on agents' ability and willingness to adjust savings and hours worked in response to these shocks, and on the roles of a multitude of further potential insurance mechanisms, such as formal and informal credit markets, insurance within the family, long-term wage contracts, institutions such as bankruptcy laws, and public redistributive programs.

In this paper we develop an original theoretical framework that we use to organize a large amount of panel and cross-sectional data on wages, hours and consumption, and to extract from this data information about the nature of risk and insurance. Specifically, we use the framework to address the following important questions at the intersection between macroeconomics and labor economics:

1. What fractions of the variance of individual wages reflect permanent shocks, transitory shocks, and measurement error?
2. How willing are households to substitute consumption and labor supply inter-temporally?
3. What is the relative importance of productivity versus preference variation in accounting for heterogeneity in consumption and hours worked across U.S. households?

The dramatic changes in the U.S. wage structure over the past thirty years are well documented (Katz and Autor, 1999). Although inequality in wages has risen sharply, cross-sectional data show that consumption inequality has gone up less markedly (Cutler and Katz, 1991; Slesnick, 2001; Attanasio et al., 2005; Krueger and Perri, 2006.). Dispersion in hours worked has remained fairly stable (Heathcote, Storesletten and Violante, 2004).

We can use our framework to connect changes in the wage structure to changes in other dimensions of inequality, and thus to measure how risk and insurance have changed over time. In particular, we address the following additional questions:

4. What fraction of the recent rise in wage dispersion was insurable from the standpoint of U.S. households?
5. How can one reconcile the different dynamics for dispersion in wages, hours and consumption?

Our model is designed to study consumption and labor supply decisions in an environment where heterogeneous households have only partial insurance against idiosyncratic labor productivity shocks. The key advantage of our model, relative to the existing literature, is that it is *analytically tractable*: equilibrium allocations can be derived without using numerical methods or analytical approximations. Thus we introduce a degree of transparency in the analysis of equilibrium allocations that is rare in models with idiosyncratic risk and incomplete markets. At the same time, the model is rich enough to include a number of desirable features such as flexible labor supply, standard (CRRA) specifications for preferences, permanent and transitory productivity shocks, permanent and transitory heterogeneity in the taste for leisure, and a flexible financial market structure incorporating a risk free asset, and, potentially, additional opportunities for insuring certain risks.

We manage to maintain tractability in this rich environment by generalizing a result first discovered by Deaton (1991) and further developed by Constantinides and Duffie (1996). Deaton (1991) argued that, in an economy where a risk-free bond is the only financial asset, agents cannot borrow to smooth permanent exogenous earnings shocks without violating the budget constraint. Thus, for a sufficiently low interest rate, they will slowly run down their assets to zero, and then consume their earnings every period. Constantinides and Duffie (1996) proved that this latter scenario can constitute a “no bond-trading” equilibrium: at the right interest rate, a positive precautionary saving motive is exactly offset by a negative intertemporal saving motive. Thus, the bond is not traded in equilibrium, even though it is available, and agents’ optimal wealth-holdings are always zero. This makes the model analytically tractable, but at a cost: since the bond is the only available asset in the Constantinides and Duffie economy, their equilibrium is autarkic and no risk-sharing is achieved.

We show that tractability can be maintained under four important generalizations that better equip the model to confront the data. First, we introduce flexible labor supply, so that earnings dynamics are determined endogenously. Second, we introduce a second source of risk and inequality in the form of permanent and transitory shocks to the taste for leisure relative to consumption. Third, we add a transitory component to wage shocks on top of the permanent component, to generate a more realistic statistical representation for wage dynamics. Fourth, and most importantly, we allow for individuals to trade a larger set of financial assets than just a risk-free bond. This latter feature yields an equilibrium where some financial markets are active, where some risks are insured in equilibrium, and where there is a non-degenerate time-varying distribution of wealth.

We are able to characterize allocations in closed form without keeping track of this distribution, because wealth holdings are known functions of other individual states and parameters defining the process for labor market risk.

Our model economy consists of a set of “islands”, where each island is a group of agents. Some (preference and productivity) shocks that hit agents are purely idiosyncratic, while others are island- (or group-) specific. Agents within each group can fully insure against idiosyncratic shocks, whereas the only asset that can be traded between islands in response to island-specific shocks is a non-contingent bond. Under appropriate conditions, there exists an equilibrium where this bond is not traded across islands. Since the no-bond-trade result holds across groups instead of holding across individual agents (as in Constantinides and Duffie), the economy features partial (within-group) risk-sharing in equilibrium.

Casual observation of actual economies suggests that, at the very minimum, individuals can use their labor supply and self-insurance through a risk-free asset to respond to shocks. These two insurance channels are, indeed, the only ones available when all shocks are island-specific. In practice, however, individuals also have access to a variety of additional insurance mechanisms ranging from more sophisticated financial instruments to family labor supply, and from insurance within firms to government transfers. These are the additional insurance channels that we have in mind when we introduce within-group insurance in the model. Indeed, when all the shocks are individual-specific, our economy delivers perfect risk-sharing. In general, with shocks at both the idiosyncratic and the island level, our market structure occupies an intermediate position between a traditional bond economy model and complete markets, and insurance is neither absent nor perfect, but partial.

The spirit of our exercise is that originally advocated by Deaton (1997): given the complexity and multiplicity of insurance channels potentially available to households, rather than modelling each individual channel explicitly, a useful first step is to quantify the overall degree of insurability of income innovations, remaining somewhat agnostic about the specific sources.¹ This intent is shared, for example, by Blundell, Pistaferri

¹Deaton (1997, pages 371-374) writes: “Saving is only one of the ways people can protect their consumption against fluctuations in their income. An alternative is to rely on other people, to share risk with friends and kin, with neighbors, or with other anonymous participants through private or government insurance schemes, or through participation in financial markets...The problems of implementing such schemes are as obvious as their advantages. As a result these schemes are more likely to exist among groups of people among whom information is good, income difficult to hide, and where behavior can be monitored and influenced by the group...While it is hard to believe that any of these mechanisms could provide complete insurance for poor households, the very multiplicity of existing mechanisms makes it likely that there is at least partial insurance through financial or social institutions, and that such risk sharing adds to [...] intertemporal transfers of money...There has therefore arisen a line of research that examines consumption behavior to find the extent of [overall] insurance...”.

and Preston (2006). Their framework, centered on the permanent-income-hypothesis, is quite different from ours. Moreover, they only look at consumption data to learn about insurability of income shocks. Instead, we will use data on both consumption and labor supply to infer exactly where the actual U.S. economy lies on the bond economy-complete markets spectrum.

A complementary literature has, instead, quantified the risk-sharing value of specific channels within structural equilibrium models. Examples are investigations of the roles of financial markets in the presence of limited commitment (Krueger and Perri, 2006), self-insurance and labor supply (Heathcote, Storesletten and Violante, 2004), bankruptcy laws (Livshits, McGee and Tertilt, 2006), consumer durables (Fernandez-Villaverde and Krueger, 2005), and government redistribution to the poor (Low, Meghir and Pistaferri, 2006). The common denominator in these models is that they require the numerical solution of complex fixed-point problems in order to characterize equilibrium cross-sectional distributions. Even though extensive comparative statics can be performed to try to disentangle the forces at work in these artificial economies, this requirement severely limits the limpidity of the analysis. Moreover, through this approach one can only quantify the role of one channel at the time rather than measuring what fraction of total labor market risk is insurable in the economy.

The present paper makes contributions on two levels. The first set of contributions is “qualitative”: this new model yields exact closed-form solutions for all the first and second cross-sectional moments of the equilibrium joint distribution of wages, hours worked, and consumption as functions of structural model parameters. Thus, the mapping between the evolution of inequality and basic preferences, shocks and market structure is transparent. For example, we show how the theoretical closed-form expressions for second moments by year and cohort can be used to identify all the deep preference and wage risk parameters, even when the latter are allowed to change over time in an unrestricted manner. Identification can be achieved by using panel data information on wages and hours (such as that contained in the Panel Study of Income Dynamics [PSID]) and cross-sectional data on consumption (such as that contained in the Consumer Expenditure Survey [CEX]).

The other set of contributions is “quantitative”. We first compile a large set of empirical moments characterizing the profiles for cross-sectional dispersion over the life-cycle and across time in the United States. It is worth emphasizing that our model permits us to use information contained in both the “macro facts” on the *levels* of consumption, earnings and hours that have been the focus of recent macroeconomic investigations (e.g., Attanasio and Davis, 1996; Krueger and Perri, 2006) and the “micro-facts” on *first-*

differences that have been the target of investigations by labor economists (e.g., Abowd and Card, 1989; Blundell, Pistaferri and Preston, 2006). Next, we use the closed-form model expressions for cross-sectional moments in levels and first-differences to estimate the model with a Minimum Distance Estimator. This exercise delivers quantitative answers to the questions we started with concerning the transmission of individual wage shocks to household consumption, leisure and welfare.

Our quantitative answers to the questions laid out above are as follows. First, out of total cross-sectional wage variance, measurement error accounts for 20%, transitory shocks for 10% and permanent shocks for the rest. Second, the evolution of cross-sectional dispersion in wages, hours and consumption calls for an intertemporal elasticity of substitution between 0.35 and 0.45 and a Frisch labor supply elasticity between 0.15 and 0.50, depending on the sample selection criteria and various parametric restrictions. Third, the key determinants of life-cycle inequality depend on the variable of interest: dispersion in earnings is mainly due to productivity shocks, dispersion in hours is mostly attributable to preference heterogeneity and measurement error, and dispersion in consumption is an intermediate case. Fourth, insurable wage shocks represent around two thirds of the observed rise in cross-sectional wage variation between 1967 and 1996, with uninsurable shocks accounting for the rest. Finally, the predominantly insurable nature of the increase in labor market risk is consistent both with the small observed rise in consumption inequality, and also with stable hours inequality, since the bulk of cross-sectional hours variation is accounted for by preference heterogeneity.

The rest of the paper is organized as follows. Section 2 develops our framework, derives the equilibrium allocations and explains how we preserve tractability. Based on these allocations, in Section 3 we present closed-form expressions for all the equilibrium cross-sectional moments of interest. Section 4 proves how these cross-sectional moments allow us to identify all the structural parameters of the model. Section 5 describes the data and the estimation algorithm. Then we report the estimation results and perform a robustness analysis. Section 6 concludes the paper.

2 Model economy

Preliminaries The economy consists of a continuum of islands (i.e., “groups” of agents) with a continuum of agents on each island. It is convenient to defer to the end of this section the formal definition of an island/group. Individuals are hit by shocks affecting both their preferences and their labor productivity. Some shocks are island-specific (i.e.,

they take the same value for all individuals within an island/group) and independent across islands. Other shocks are purely idiosyncratic and independent across individuals within an island. In this economy a law of large numbers (Uhlig, 1996) can be applied twice so that individual-level shocks “wash out” within an island and island-level shocks induce no aggregate uncertainty in the economy as a whole (see Attanasio and Rios-Rull, 2000, for a similar structure). Production takes place through a constant returns to scale technology with labor as the only input. All assets in the economy are in zero net supply.

Demographics We adopt the perpetual youth, overlapping-generations framework developed by Yaari (1965). Agents are born onto a given island/group at age zero and survive from age a to age $a + 1$ with constant probability $\delta < 1$. A new generation with mass $(1 - \delta)$ enters the economy each period. Thus the measure of agents of age a is $(1 - \delta)\delta^a$ and the total mass of agents is one. Upon birth, each agent is endowed with zero financial wealth.

Preferences Lifetime utility for an agent born (i.e., entering the labor market) at time $t = b$ is given by

$$\mathbb{E}_b \sum_{t=b}^{\infty} (\beta\delta)^{t-b} u(c_t, h_t, \zeta_t, \varphi), \quad (1)$$

where the expectation is taken over sequences of preference and productivity shocks. Period utility is

$$u(c_t, h_t, \zeta_t, \varphi) = \frac{c_t^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi + \sigma\zeta_t) \frac{h_t^{1+\sigma}}{1+\sigma}. \quad (2)$$

Here c_t denotes consumption, and h_t is hours worked at time t . Agents discount the future at rate $\beta\delta$, where $\beta < 1$ is the pure discount factor. The parameter γ represents the coefficient of relative risk-aversion ($1/\gamma$ is the intertemporal elasticity of substitution for consumption). The Frisch elasticity of labor supply is equal to $1/\sigma$.²

The parameters φ and ζ_t are individual-specific preference weights that capture the strength of an individual’s aversion to work relative to his preference for consumption: φ is a permanent island-level effect known at the time of labor market entry, while ζ_t is a transitory idiosyncratic shock drawn at each age which is *i.i.d.* across agents and over

²Separability in preferences between consumption and hours worked is a common assumption in the micro literature that estimates elasticities for consumption and labor supply (for a survey, see Browning, Hansen and Heckman, 1999). A popular alternative that is more common in the macro literature is to assume consumption and leisure are aggregated in a Cobb-Douglas fashion. In Heathcote, Storesletten and Violante (2005b) we show that it is also possible to solve the model in closed form in that case. Here we adopt the separable specification, primarily because it affords valuable flexibility in distinguishing between agents’ willingness to substitute consumption and hours inter-temporally.

time. Let F_ζ be the distribution of the transitory preference shock with variance v_ζ . We specify the distribution of φ below.³

We include preference variation in the model because some authors have argued that a large share of hours variation occurs at fixed wage rates (e.g., Abowd and Card, 1989). The fixed effect φ represents permanent differences in industriousness across agents, while the shock ζ_t captures transitory changes in the relative taste for leisure (e.g., short bouts of illness).

Productivity shocks The process for individual efficiency units of labor (labor productivity) w_t is given (in logs) by the sum of two orthogonal stochastic components:

$$\log w_t = \alpha_t + \varepsilon_t. \quad (3)$$

The difference between these two components is that α_t is an island-level shock, whereas ε_t is an individual-level shock.⁴

The α_t shocks are assumed to be permanent and follow a random walk:

$$\alpha_t = \alpha_{t-1} + \omega_t,$$

where the innovation ω_t is drawn from the distribution $F_{\omega t}$ with variance $v_{\omega t}$ at time t . The individual-level shocks ε_t are the sum of two orthogonal components:

$$\varepsilon_t = \kappa_t + \theta_t,$$

where θ_t is a transitory (independently distributed over time) shock drawn from $F_{\theta t}$ with variance $v_{\theta t}$ at time t , and κ_t is a permanent component that follows a second unit root process:

$$\kappa_t = \kappa_{t-1} + \eta_t,$$

where the innovation η_t is drawn from the distribution $F_{\eta t}$ with variance $v_{\eta t}$ at time t .

Upon entering the labor market at age $a = 0$ agents draw initial realizations for the two components, κ_0 and α_0 . The initial value for κ_0 is drawn from the distribution F_{κ_0} with variance v_{κ_0} . The fixed island-specific effects (φ, α_0) are jointly drawn from a bivariate distribution $F_{\varphi\alpha}$ with variances v_{α_0} and v_φ and covariance $v_{\varphi\alpha}$.⁵

³The fact that in the period utility φ and ζ are multiplied by preference parameters is just an innocuous normalization which will turn out to simplify the expressions for equilibrium allocations.

⁴We could also handle an aggregate shock z_t (affecting all islands in the same way) in addition to the idiosyncratic and island-level shocks, i.e. $\log w_{it} = z_t + \alpha_{it} + \varepsilon_{it}$. The process for z_t would play no role in determining equilibrium cross-sectional second moments, so for the sake of simplicity we abstract from aggregate risk in the exposition of the model.

⁵The initial draw of α_0 can be thought of as innate ability, or initial human capital obtained through schooling. If human capital accumulation were endogenous, agents with low φ would be more prone to accumulate since they would be willing to work more hours and hence face a higher return on their investment. They would therefore start their working life with a higher value for α_0 . Human capital theory, thus, seems to suggest a negative value for $v_{\varphi\alpha}$.

Following the tradition in this literature, the evolution of efficiency units (w_t) is taken as exogenous. However, earnings are endogenous because labor supply is a choice variable, so individuals have some control over labor income.⁶ The statistical process for wages described above is quite rich, and is potentially consistent with both the key features of individual wage dynamics as well as with the broad trends in wage dispersion across the life-cycle and through time. For example, the empirical autocovariance function for individual wages displays a sharp decline at the first lag, indicating the presence of a transitory component in wages. At the same time cross-sectional wage dispersion increases approximately linearly with age, suggesting the presence of permanent shocks.

In the empirical labor literature, the processes for wages or earnings are often specified as combinations of a unit root component and an MA(1) component (e.g., Meghir and Pistaferri, 2004). However, the estimated coefficient on the lagged error component is typically small, and often not significantly different from zero.⁷ In the interest of parsimony, we therefore abstract from a moving average component.⁸

Since the process for preference shocks is time-invariant, all nonstationarity in the model is induced by time-variation in the distributions for wage shocks. All these dynamics in the wage-generating process are modelled as time effects; in our baseline economy we assume no cohort effects. In Heathcote, Storesletten and Violante (2005a) we argue that time effects are required to account for the observed trends in inequality in thirty years of U.S. data, while there is little evidence that cohort effects have been important. As part of a sensitivity analysis, we generalize the model to introduce cohort effects in the form of cross-cohort variation in the variances of the initial fixed effects α_0 and κ_0 .

Information Agents are assumed to know the current and future distributions $\{F_{\varphi\alpha}, F_{\kappa_0}, F_{\zeta}, F_{\omega t}, F_{\eta t}, F_{\theta t}\}$, i.e. they have perfect foresight over future changes in wage dynamics. This assumption is not required for tractability. Alternatively, one could assume that the distributions are Normal, and that the variances $v_{\omega t}$, $v_{\eta t}$ and $v_{\theta t}$ themselves

⁶Huggett, Ventura and Yaron (2006a, 2006b) take a different approach: by building on the prototypical Ben-Porath model, they allow individuals to allocate time between work and investment in a risky human capital accumulation technology. In their model, individuals have inelastic labor supply but, through these investments, they partially control the life-cycle evolution of efficiency units of labor.

⁷For college graduates, Meghir and Pistaferri cannot reject the hypothesis that the number of moving average lags is equal to zero. Similarly, Abowd and Card (1989) find no evidence of a moving average component in bi-annual National Longitudinal Survey of Men 45-59 data.

⁸Recently, Guvenen (2007) and Primiceri and van Rens (2006) have revived a long-standing debate on whether the best statistical representation for the wage process is an ARMA process or one with deterministic heterogeneity in wage profiles across individuals. We have opted for the former (and more common) hypothesis, but it is important to note that we could also allow for within-group heterogeneity in the growth rate of individual wages. The next draft of the paper will include this extension.

follow unit root processes.⁹

Island/group structure An island is a group of agents who share the same age a , the same pair of fixed effects (φ, α_0) , and the same infinite sequence for innovations to the island-level component of wages $\{\omega_t\}$. Within an island, agents are heterogeneous with respect to their sequences $\{\zeta_t\}$, $\{\kappa_t\}$ and $\{\theta_t\}$. Each period, a continuum of new islands is formed by agents who share common (known) values for (φ, α_0) and a common (unknown) sequence for future shocks $\{\omega_t\}$. The size of each island/group shrinks at rate δ as time passes.¹⁰

Market structure The economy-wide goods market and the labor market are perfectly competitive, so individual wages equal individual productivities (units of effective labor per hour worked). Within each island, we assume the existence of complete insurance markets so that agents can pool all idiosyncratic risks associated with the shocks (ζ, κ, θ) .¹¹ These markets are unable to provide insurance against the island-specific shocks (φ, α) . Trade among islands is restricted by assumption: there exists only a market for a risk-free non-contingent bond (which is in zero net supply, like every other asset).

The island/group structure is a convenient device to allow for a tractable description of an environment in which some risks are fully insured while others are not. Instead of imposing restrictions on the set of assets traded *among individuals* (in the Bewley-Imrohoroglu-Aiyagari-Huggett tradition), we restrict the set of assets that can be traded *between groups*, which allowing for perfect risk-sharing within groups. This way, our economy can achieve any degree of insurance between a bond economy and a complete market economy. For $v_\varphi = v_{\alpha t} = 0$, the market structure is complete, while for $v_{\kappa t} = v_{\theta t} = v_\zeta = 0$, it converges to the bond economy. In general, it is an economy offering *partial insurance* against shocks.¹²

In reality, individuals pool risks within different groups through a range of mechanisms. We will focus on a decentralization involving explicit insurance markets, but the same

⁹These processes for the variances of the idiosyncratic shocks could also be functions of an aggregate productivity shock. This would allow for counter-cyclical variance in idiosyncratic risk, an important ingredient for asset pricing in incomplete market models according to Constantinides and Duffie (1996).

¹⁰We segregate individuals across islands by age a and fixed effects (φ, α_0) because this is convenient when it comes to solving for equilibrium allocations: when individuals are segregated this way, within-island allocations can be determined using equal-weighted island-level planning problems. However, given our proposed market structure, we could envision islands populated by individuals with different ages and fixed effects, as long as all individuals on an island experience the same sequence $\{\omega_t\}$.

¹¹We assume markets open before the realization of the initial draw κ_0 .

¹²We borrow this term from Blundell, Pistaferri and Preston (2006) in the expectation that it will become the standard way to define this large class of economies which offer more risk sharing than a bond-economy but still less than full insurance.

allocations could obviously be achieved through non-market mechanisms. As Deaton (1997) noted, identifying all the details of precisely who is sharing risks and precisely how they are being shared is a major challenge. The strength of our approach is that we can and will remain deliberately agnostic about what constitutes an island/group in the data, and about how within-group risk-sharing is achieved. When we outline the models predictions for cross-sectional moments and prove identification, it will become clear that we do not need to specify these details *a priori*. Rather there is sufficient information in the aggregate cross-section to let the data speak as to the fraction of preference and productivity variation that can be insured in equilibrium, without exploiting any of the model's predictions for within-group or between group inequality.

2.1 Digression: two explicit models of the family

The family may be an important source of risk-sharing. Our structure implicitly allows for some within-family insurance (insofar as members of the same family belong to the same island/group), but we can also be more explicit. In what follows we develop two alternative models of the family that are entirely consistent with the individual agent's problem above. We will envision a household as composed of two members (e.g., head and spouse indexed by $i = 1, 2$), but it will be obvious how the logic could be extended to households with additional members.

Model I Suppose a household's period utility function is a simple extension of the one in (2), i.e.,

$$u(c_t, h_{1t}, h_{2t}, \zeta_{1t}, \zeta_{2t}, \varphi) = \frac{c_t^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi) \sum_{i=1,2} \exp(\sigma\zeta_{it}) \frac{h_{it}^{1+\sigma}}{1+\sigma},$$

and assume that head and spouse belong to the same island/group, i.e. they share the same φ and α_0 (as specified in the preferences above) and face the same sequence of ω_t shocks over time. However, spouses will experience orthogonal draws of the preference shocks ζ_t and of the individual specific productivity shocks η_t and θ_t . Thus, hours worked by the two spouses will differ and will be imperfectly correlated. Nonparticipation of one of the spouses should be interpreted as the spouse drawing either a very low value for w_t , or a very high value for the preference shock ζ_t . Given that both members face the same labor productivity process, using data on male wages and hours allows identification of all the structural productivity parameters.

Finally, note that in this model of the family it does not matter whether we view consumption as a private good or as a public good at the household level. The reason is that

all agents within an island (and thus both spouses) enjoy the same level of consumption anyway (see Section 2.3 below).

Model II Under this alternative specification, household preferences display perfect complementarity in leisure, i.e.,

$$\begin{aligned} u(c_t, h_{1t}, h_{2t}, \zeta_{1t}, \zeta_{2t}, \varphi) &= \frac{c_t^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi + \sigma\zeta_{1t}) \frac{(\max\{h_{1t}, \exp(\sigma\zeta_{2t})h_{2t}\})^{1+\sigma}}{1+\sigma} \\ &= \frac{c_t^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi + \sigma\zeta_{1t}) \frac{h_{1t}^{1+\sigma}}{1+\sigma}, \end{aligned}$$

where the second equality is obtained by plugging back into the utility function the optimality condition for hours worked, which requires $h_{1t} = \exp(\sigma\zeta_{2t})h_{2t}$. This is exactly the original preference specification in (2). In this reduced form model, heads' labor supply is not just chosen based on his own wage, but rather on an "effective wage" \hat{w}_{1t} which can be simply measured in the data as family earnings divided by heads' hours, or $\hat{w}_{1t} = [w_{1t}h_{1t} + w_{2t}h_{2t}]/h_{1t}$.

2.2 Agent's problem

Let s^t denote the history of the shocks up to date t for agents on a given island, where s^t has typical element $s_\tau = (\alpha_\tau, \varphi, \kappa_\tau, \theta_\tau, \zeta_\tau)$. Let $B_t^a(s_{t+1}, s^t)$ and $Q_t(s_{t+1}, s^t)$ denote the quantity and the price of a claim purchased at node s^t by an agent of age a that pays off if and only if the state at $t+1$ is s_{t+1} . These assets are only traded within islands. Non-contingent one-period bonds are also traded. These bonds can be traded between any agents in the economy, i.e. also between islands. Let $b_t^a(s^t)$ and $q_t(s^t)$ denote the quantity and price of riskless bonds purchased.

An agent's sequential budget constraint at node s^t is then

$$c_t^a(s^t) + \int Q_t(s_{t+1}, s^t) B_t^a(s_{t+1}, s^t) ds_{t+1} + q_t(s^t) b_t^a(s^t) = w_t(s^t) h_t^a(s^t) + d_t^a(s^t) \delta^{-1} \quad (4)$$

where realized wealth is

$$d_t^a(s^t) = B_{t-1}^{a-1}(s_t, s^{t-1}) + b_{t-1}^{a-1}(s^{t-1}),$$

and s^t is the history s^{t-1} followed by s_t . We assume that agents are initially endowed with zero initial financial wealth.¹³ For what follows, it is also useful to define "net" financial wealth as $\hat{d}_t^a(s^t) = d_t^a(s^t) - [c_t^a(s^t) - w_t(s^t) h_t^a(s^t)]$, i.e. wealth net of the financial income used to finance the gap between current consumption and earnings.

¹³It is straightforward to relax the assumption of zero initial financial wealth. The key requirement for retaining tractability is that *average* initial wealth on each island is zero.

The problem for an agent entering the labor market at date t is to maximize (1) subject to a sequence of budget constraints of the form (4). In addition to these budget constraints, agents face limits on borrowing that rule out Ponzi schemes, and non-negativity constraints on consumption and hours worked.

2.3 Competitive equilibrium

Given a path for the distributions $\{F_{\varphi\alpha}, F_{\kappa_0}, F_{\zeta}, F_{\omega t}, F_{\eta t}, F_{\theta t}\}$, a sequential competitive equilibrium for this economy is a set of allocations

$\{c_t^a(s^t), h_t^a(s^t), d_t^a(s^t), b_t^a(s^t), B_t^a(s_{t+1}, s^t)\}$, prices $\{q_t(s^t), Q_t(s_{t+1}, s^t)\}$ for all histories s^t such that 1) allocations maximize expected lifetime utility, 2) insurance markets clear island by island, and 3) the economy-wide markets for the final good, labor services and for the non-contingent bond clear.

In general, incomplete-market economies do not admit an analytical solution, and numerical methods are required to solve for equilibrium allocations. The key difficulty is the determination of the equilibrium wealth distribution, since wealth is a state variable. We now state and will later prove that, in our environment, the vector s_t together with the age a of the individual contains sufficient information to describe individual choices and prices in equilibrium. Note that s_t does not include individual financial wealth d_t , thus all the entries in s_t are either constant (e.g., permanent preference heterogeneity) or evolve stochastically according to known exogenous processes, while age evolves deterministically.¹⁴ This simplifies enormously the task of characterizing allocations.

We now state an assumption ensuring that the agent's optimization problem is well defined and that allocations are finite.

ASSUMPTION A1: *The term*

$$\mathcal{M}_t^a \equiv M_{\zeta} + M_{\kappa_0} + M_{\theta t} + \sum_{j=t-a+1}^t M_{\eta j} \quad (5)$$

is finite, where

$$\begin{aligned} M_{\zeta} &\equiv \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp(-\zeta) dF_{\zeta} \right) \\ M_{\kappa_0} &\equiv \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp \left(\frac{1 + \sigma}{\sigma} \kappa_0 \right) dF_{\kappa_0} \right) \\ M_{\eta t} &\equiv \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp \left(\frac{1 + \sigma}{\sigma} \eta_t \right) dF_{\eta t} \right) \end{aligned}$$

¹⁴Moreover, we will show that, in equilibrium, individual financial wealth itself d_t^a can be expressed as a function of s_t and age a .

$$M_{\theta t} \equiv \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp \left(\frac{1 + \sigma}{\sigma} \theta_t \right) dF_{\theta t} \right).$$

The time-varying age-specific term \mathcal{M}_t^a will be convenient when characterizing individual consumption and hours worked. We are now ready to state the first main result of the paper, the characterization of equilibrium allocations for consumption, hours and wealth.

PROPOSITION 1 [COMPETITIVE EQUILIBRIUM] *Under Assumption A1, there exists a sequential competitive equilibrium characterized as follows:*

(i) *Consumption and hours are given by*

$$\log c_t^a(s_t) = \log c_t^a(\varphi, \alpha_t) = -\varphi + \frac{1 + \sigma}{\sigma + \gamma} \alpha_t + \mathcal{M}_t^a \quad (6)$$

$$\log h_t^a(s_t) = -\varphi - \zeta_t + \frac{1 - \gamma}{\sigma + \gamma} \alpha_t + \frac{1}{\sigma} (\kappa_t + \theta_t) - \frac{\gamma}{\sigma} \mathcal{M}_t^a, \quad (7)$$

and net financial wealth is given by

$$\hat{d}_t^a(s_t) = c_t^a(s_t) \left[1 - \frac{\exp \left(\frac{1 + \sigma}{\sigma} \kappa_t \right)}{\int \exp \left(\frac{1 + \sigma}{\sigma} \kappa_t \right) dF_{\kappa t}^a} \right] \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j \left(\frac{c_{t+j}^a(s_{t+j})}{c_t^a(s_t)} \right)^{1-\gamma}, \quad (8)$$

where $F_{\kappa t}^a$ denotes the distribution of κ_t across agents of age a at date t .

(ii) *The prices of the non-contingent bond and of state-contingent Arrow securities are given by*

$$q_{t-1} = \beta \exp(-\gamma \Delta \mathcal{M}_t) \int \exp \left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_t \right) dF_{\omega t} \quad (9)$$

$$Q_{t-1}(s_t, s^{t-1}) = \beta \exp(-\gamma \Delta \mathcal{M}_t) \exp \left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_t \right) f_{s_t} \quad (10)$$

where f_{s_t} is the probability density of s_t and $\Delta \mathcal{M}_t = \mathcal{M}_t^a - \mathcal{M}_{t-1}^{a-1} = M_{\eta t} + \Delta M_{\theta t}$.

(iii) *In equilibrium, the bond is not traded across islands.*

PROOF: The logic of the proof is straightforward. We first guess that there will be no bond trade between islands and, hence, zero net savings on each island. Given complete markets on an island, we apply the first welfare theorem and solve for the island-specific allocations via a simple static planner problem with equal weights. With the allocations in hand, we can compute the prices that support these allocations in the decentralized equilibrium and verify the no bond-trade conjecture. See the Appendix for the formal proof.

It should be clear now that in equilibrium the individual-specific shocks are fully insurable, while the island-level shocks are not insurable at all. Hereafter, we refer to $(\kappa_t, \theta_t, \zeta_t)$ as “insurable shocks” and to (α_t, φ) as “uninsurable shocks”.

Consumption, hours and wealth Individual consumption is independent of the realization of the idiosyncratic shocks (κ_t, θ_t) , since they can be fully insured in equilibrium, but it is re-scaled by the effect that the island-level permanent shock α_t and preference for leisure parameter φ have on earnings. This consumption allocation is not what the simplest version of the PIH would imply. Consumption is still a random walk, as in Hall (1978), but some permanent shocks are fully insurable, and thus do not affect consumption. In other words, our consumption allocations exhibit “excess smoothness” (as originally defined by Campbell and Deaton, 1989). It is precisely this feature of the data that has motivated a large amount of recent research aimed at developing insurance models that lie in between the bond economy and complete markets (e.g., Krueger and Perri, 2006; Attanasio and Pavoni, 2006).¹⁵

Hours worked are increasing in the idiosyncratic shocks (κ_t, θ_t) , proportionately to the Frisch elasticity, since these shocks have a substitution effect but no income effect given that they are perfectly insured. Island-level permanent shocks do have an income effect, and hours increase with α_t if and only if $\gamma < 1$. Both φ and ζ_t reduce hours worked through their effect on the marginal utility of leisure.

Wealth is inversely related to the current value of the permanent insurable shock κ_t . In particular, the larger the permanent insurable component of earnings, $\exp\left(\frac{1+\sigma}{\sigma}\kappa_t\right)$ relative to the average on the island, the lower financial wealth. In fact, the purchases of the state-contingent claims at s^t are such that the payout in any future event s_{t+1} is equal to the difference between average consumption on the island and individual earnings given history $s^{t+1} = (s^t, s_{t+1})$. Finally, wealth is proportional to island consumption, since different islands have different (uninsurable) productivities and taste for leisure.

2.3.1 No bond trading and tractability

The reason allocations and prices can be characterized without reference to wealth in our environment is twofold. First, within a particular island, markets are effectively complete, so allocations can be computed using an island-planner abstraction. The planner worries only about its objective and an island-level resource constraint, thus allocations can be

¹⁵In Section 2.3.2 we explain the dependence of the allocations on age.

characterized without reference to the within-island wealth distribution.¹⁶

Second, the inter-island wealth distribution does not show up in allocations because, in equilibrium, this distribution remains degenerate at zero. To understand this result, recall that our individuals have three saving motives: an intertemporal motive (given by the interaction of patience, the interest rate and earnings growth rate), a smoothing motive (that of the strict permanent income hypothesis), and a precautionary motive (due to island-level earnings uncertainty). In general, the presence of this latter motive makes it impossible to obtain exact closed form solutions for equilibrium consumption allocations. In our model, however, the precautionary saving motive is the same on each island, since $v_{\omega t}$ is common across all islands. Also the intertemporal saving motive is the same on each island, because expected island-level earnings growth is constant across all islands. Thus one can construct an equilibrium where the interest rate is low enough so that the (negative) intertemporal saving motive exactly offsets the precautionary motive. In Section 2.3.2, we derive such interest rate explicitly in closed-form. As a result, agents are only left with a saving motive related to consumption-smoothing. However, island-level shocks are permanent, hence also this motive is zero. In conclusion, there is no island-level saving, and island-level wealth remains at zero.¹⁷

Another way to state this result is that, absent inter-island trade, agents in different islands share a common (but time-varying) expected growth rate of marginal utility of consumption at each date. Therefore, even in presence of an economy-wide bond market, the bond is not traded across islands. This property emerges when (i) islands all start out with zero average wealth, (ii) preferences over consumption are in the constant relative risk aversion class (so that high and low income agents have the same attitude towards risk), (iii) island-specific wage shocks are multiplicative and permanent (so that wage growth has the same mean and variance across islands), and (iv) expected island-level earnings growth is equal across islands, which implies restrictions on how labor supply enters preferences.

Our result represents a generalization of Constantinides-Duffie (1996) in several dimensions. First, in our economy some risks are insurable, so that the Constantinides-Duffie no-trade result applies across groups rather than across individuals. This extension is

¹⁶This first part of our result is reminiscent of some existing contributions. With complete markets, CRRA preferences, and heterogeneity in initial wealth endowments, Chatterjee (1994) shows that Gorman-aggregation holds within the neoclassical growth model: individual savings are linear in individual wealth, so average wealth is a sufficient statistics for aggregate dynamics. Maliar and Maliar (2003) generalize this result to the case of fully insurable productivity shocks. In these complete markets economies, the dynamics of the wealth distribution are easy to track.

¹⁷Within an island, though, some individuals will have positive wealth while others will have negative wealth positions; some will save, others will dissave.

important, because in Constantinides-Duffie the equilibrium entails no risk-sharing at all, whereas we demonstrate that the no-bond-trade result (and the tractability that follows) can survive in more interesting environments with partial insurance. Second, our agents supply labor elastically, so that the process for earnings is endogenous. Thus the no-trade result survives when we introduce a second intra-temporal dimension to household optimization. Third, our agents are subject to preference shocks, beyond productivity. An additional source of shocks in the economy permits, potentially, a better fit of the data.

To our knowledge, it is possible to solve for the equilibrium of incomplete-market economies analytically only in two other cases. Caballero (1990) shows that, in a bond economy, under CARA utility the precautionary saving motive is constant across all agents, independently of the earnings process. Wang (2003) builds on this result and proves that one can construct an equilibrium where the risk-free interest rate is low enough that the intertemporal dissaving motive offsets the precautionary motive. In this case, in general, the residual consumption-smoothing motive is non-zero and in equilibrium consumption end up following the permanent income hypothesis. Benabou (2002) outlines an economy where each agent lives in isolation from the rest of the economy: by working, she produces a final good that can be consumed or used to increase her stock of human capital. With CRRA utility over consumption and hours worked, given an assumption of i.i.d. log-normal productivity shocks to the human capital production function, Benabou shows that the cross-sectional distribution of human wealth has a closed form.

2.3.2 Normality of the shocks

To make further progress in characterizing equilibrium allocations and prices, distributional assumptions for the shocks are required. Under Normality of preferences and productivity shocks (implying lognormal wages in levels), the term \mathcal{M}_t^a can be computed analytically. This result is summarized in the following corollary to Proposition 1.

COROLLARY 1 [NORMALITY OF SHOCKS] *Suppose that $(\alpha_0, \omega_t, \kappa_0, \eta_t, \theta_t, \varphi, \zeta_t)$ are Normally distributed, with variance v_x and mean $-v_x/2$ for each random variable x . Then the terms $M_\zeta, M_{\kappa_0}, M_{\theta_t}, M_{\eta_t}$ are given by, respectively,*

$$\begin{aligned} M_\zeta &= \frac{\sigma}{\sigma + \gamma} v_\zeta \\ M_{\kappa_0} &= \left(\frac{1 + \sigma}{\sigma + \gamma} \right) \frac{v_{\kappa_0}}{2\sigma} \end{aligned}$$

$$\begin{aligned}
M_{\eta t} &\equiv \left(\frac{1 + \sigma}{\sigma + \gamma} \right) \frac{v_{\eta t}}{2\sigma} \\
M_{\theta t} &= \left(\frac{1 + \sigma}{\sigma + \gamma} \right) \frac{v_{\theta t}}{2\sigma},
\end{aligned}$$

and the term \mathcal{M}_t^a is given by

$$\mathcal{M}_t^a = \left(\frac{1 + \sigma}{\sigma + \gamma} \right) \left(\frac{v_{\varepsilon t}^a}{2\sigma} + \frac{\sigma}{1 + \sigma} v_{\zeta} \right), \quad (11)$$

where $v_{\varepsilon t}^a = v_{\theta t} + v_{\kappa 0} + \sum_{j=0}^{a-1} v_{\eta, t-j}$ is the variance of the insurable component of the wage on an island of age a at date t . The risk-free bond price becomes

$$q_{t-1} = \beta \exp \left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \left(\frac{v_{\eta t} + \Delta v_{\theta t}}{2\sigma} - \left(\gamma \frac{1 + \sigma}{\sigma + \gamma} + 1 \right) \frac{v_{\omega t}}{2} \right) \right). \quad (12)$$

The Normality assumption is not needed to derive closed-form expressions for the cross-sectional second moments of the joint distribution of wages, hours and consumption, However, it is useful in two ways. First, by solving explicitly for the dummy variables \mathcal{M}_t^a , we see from (11) that consumption is increasing in the variance of the insurable shocks. The logic for this result is that greater insurable wage dispersion increases average labor productivity and thus expected earnings and consumption, an important consideration for welfare analysis (see Heathcote, Storesletten, and Violante, 2005b). Second, through Normality we can get an intuitive closed-form expression for the price of the risk-free bond, which we discuss next.

The risk-free rate Let's begin by assuming that $\Delta v_{\varepsilon t} = v_{\eta t} + \Delta v_{\theta t} = 0$. If $v_{\omega t} = 0$, the equilibrium interest rate would simply guarantee that the intertemporal saving motive is zero, i.e. $q_t = \beta$. In this case, there is no precautionary saving motive since the only uninsurable risks agents face are the fixed effects (φ, α_0) which are drawn before markets open. In general, the bond price is increasing in the variance of ω_t , indicating that the equilibrium interest rate is less than the rate of time preference, or $q_t > \beta$. Intuitively, greater risk increases the precautionary demand for safe assets, which drives up the price of these assets. The larger is risk aversion, the stronger is this effect.¹⁸ The bond price is decreasing in the growth rate of the variance of the insurable risk $\Delta v_{\varepsilon t}$. Since earnings are increasing in the variance of insurable risk, faster growth in this variance - a higher value for $\Delta v_{\varepsilon t}$ - signals faster labor income growth, which in turn lowers the demand for saving and the price of safe bonds.

¹⁸The effect of σ depends on the value of γ . If $\gamma > 1$, then the income effect on labor supply is stronger than the substitution effect and hours respond negatively to permanent shocks (see equation 7). In other words, labor supply is used as an insurance device. In this case, a higher Frisch elasticity of labor supply (lower σ) reduces the precautionary saving motive, since labor supply provides a hedge against risk.

Overall, the equilibrium interest rate is such that the intertemporal dissaving motive is exactly offset by the precautionary saving motive. This feature of the equilibrium is particularly clear in the special case $\sigma \rightarrow \infty$ (inelastic labor supply). In this case, expression (12) simplifies to $(\rho - r_t)/\gamma \simeq (1 + \gamma)v_{\omega t}/2$, where $\rho \simeq \log\left(\frac{1}{\beta}\right)$ and $r_t \simeq \log\left(\frac{1}{q_{t-1}}\right)$. The right hand side captures the precautionary motive for saving, which turns out to be proportional to the variance of the innovation to the uninsurable component of productivity multiplied by the coefficient of relative prudence. The left hand side is a measure of the intertemporal motive to dissave, since the rate of time preference exceeds the equilibrium interest rate. With elastic labor supply the expression is somewhat more involved, but the interpretation remains unchanged.¹⁹

3 Equilibrium cross-sectional moments

The first step towards making the model operational is recognizing that measurement error is pervasive in micro data. We assume that consumption, earnings and hours worked are measured with error, and that this error is classical, i.e., i.i.d. over time and across agents. Let the measurement error in a variable x be denoted μ^x , with mean zero and variance v_x , and let $\hat{x}_t = x_t + \mu_t^x$ denote the empirical observation of x_t including measurement error. While we directly observe consumption, hours, and earnings, we compute hourly wages as earnings divided by hours, hence measurement error in hourly wages reflects errors in both earnings and hours.

If we abstract from the age/time dummies \mathcal{M}_t^a in the allocations (6) and (7) (this same treatment will be applied to individual wages, hours and consumption observations in the data), then the measured individual log allocations at time t are given by

$$\log \hat{w}_t = \alpha_t + \kappa_t + \theta_t + \mu_t^y - \mu_t^h \quad (13)$$

$$\log \hat{c}_t = -\varphi + \frac{1 + \sigma}{\sigma + \gamma} \alpha_t + \mu_t^c \quad (14)$$

$$\log \hat{h}_t = -\varphi - \zeta_t + \frac{1 - \gamma}{\sigma + \gamma} \alpha_t + \frac{1}{\sigma} (\kappa_t + \theta_t) + \mu_t^h \quad (15)$$

Let $\Delta x_t \equiv x_t - x_{t-1}$ denote the change in variable x . From the above expressions, the

¹⁹Using the log-approximation, with $\sigma < \infty$, the equilibrium risk-free rate satisfies

$$\frac{\rho - r_t}{\gamma} + \left(\frac{1 + \sigma}{\sigma + \gamma}\right) \frac{\Delta v_{\varepsilon t}}{2\sigma} \simeq \left(\frac{1 + \sigma}{\sigma + \gamma}\right) \left(\gamma \frac{1 + \sigma}{\sigma + \gamma} + 1\right) \frac{v_{\omega t}}{2}$$

where the left hand side shows a second term capturing the extra dissaving motive due to expected earnings growth.

individual changes in measured log allocations can be expressed as

$$\Delta \log \hat{w}_t = \omega_t + \eta_t + \Delta \theta_t + \Delta \mu_t^y - \Delta \mu_t^h \quad (16)$$

$$\Delta \log \hat{c}_t = \frac{1 + \sigma}{\sigma + \gamma} \omega_t + \Delta \mu_t^c \quad (17)$$

$$\Delta \log \hat{h}_t = \frac{1 - \gamma}{\sigma + \gamma} \omega_t + \frac{\eta_t + \Delta \theta_t}{\sigma} - \Delta \zeta_t + \Delta \mu_t^h. \quad (18)$$

In what follows, we denote by $v_{\alpha t}^a$ and $v_{\kappa t}^a$ (with $v_{\varepsilon t}^a = v_{\kappa t}^a + v_{\theta t}$) the cumulation of the permanent uninsurable and insurable shocks over the life-cycle of a cohort of age a at time t , e.g., $v_{\alpha t}^a = v_{\alpha 0} + \sum_{j=0}^{a-1} v_{\omega, t-j}$. Since we have filtered out differences in mean values for allocations across age groups, the unconditional cross-sectional variance is $v_{\alpha t} = (1 - \delta) \sum_{a=0}^{\infty} \delta^a v_{\alpha t}^a$, and similarly for $v_{\kappa t}$.²⁰

3.1 Time-series moments

Given the allocations (13) – (18), we can easily derive closed-form expressions for the unconditional second moments (variances and covariances) of the equilibrium cross-sectional joint distribution of wages, hours and consumption. These moments are only functions of preference, risk and measurement error parameters.

We start from the moments in levels, which we call the “macro moments” and then move to the variances and covariances of first differences, which we will refer to as the “micro moments”. Both sets of moments are informative in their own way. Macroeconomists working with heterogeneous-agents models have traditionally been more interested in the levels (e.g., Aiyagari 1994; Huggett 1996; Castaneda et al. 2003; Krusell and Smith 1998). Labor economists, by contrast, have typically studied the properties of these variables in first differences (e.g., Abowd and Card 1989; Blundell, Pistaferri and Preston 2006). We view both micro and macro facts as containing valuable information about structural parameters, and will use both to identify and estimate the model.

Macro moments The cross-sectional moments involving wages and hours are, respectively,

$$\text{var}(\log \hat{w}_t) = v_{\alpha t} + v_{\kappa t} + v_{\theta t} + v_{\mu y} + v_{\mu h} \quad (19)$$

$$\text{var}(\log \hat{h}_t) = v_{\varphi} + v_{\zeta} - \frac{2(1 - \gamma)}{\sigma + \gamma} v_{\varphi \alpha} + \left(\frac{1 - \gamma}{\sigma + \gamma} \right)^2 v_{\alpha t} + \frac{1}{\sigma^2} (v_{\kappa t} + v_{\theta t}) + v_{\mu h} \quad (20)$$

$$\text{cov}(\log \hat{h}_t, \log \hat{w}_t) = \frac{1 - \gamma}{\sigma + \gamma} v_{\alpha t} + \frac{1}{\sigma} (v_{\kappa t} + v_{\theta t}) - v_{\mu h} \quad (21)$$

²⁰This follows from the variance decomposition

$$\text{var}(x_t) = \mathbb{E}[\text{var}(x_t|a)] + \text{var}[\mathbb{E}(x_t|a)],$$

where the second term is zero if we abstract from the terms $\mathcal{M}_{a,t}$ in the allocations.

The dispersion of measured wages is simply the sum of all the orthogonal productivity components, plus the variances of measurement error in earnings and hours.

The variance of hours has five components. First, the larger is heterogeneity in the taste for leisure v_φ and in preference shocks v_ζ , the larger is cross-sectional variance in hours. Second, the variance of the uninsurable shock translates into hours dispersion proportionately to the distance between γ and one. As γ approaches one, uninsurable shocks have no effect on hours. Third, the variance of the insurable shocks increases hours dispersion proportionately to the Frisch elasticity (squared). Fourth, the effect of the covariance between taste for leisure and ability depends on the role of labor supply. Recall that if $\gamma > 1$, strong income effects mean that high-ability individuals will work fewer hours. If $v_{\varphi\alpha} > 0$, then high-ability (high α_0) individuals also desire to work less for pure taste reasons. Thus, the stronger this covariance is, in absolute value, the higher the variance of hours will be. Finally, measurement error in hours contributes positively to observed dispersion.

The covariance between wages and hours has three components. Whether the variance of uninsurable shocks increases or decreases the covariance once again depends on the value for risk aversion γ . If $\gamma > 1$, then uninsurable shocks tend to decrease the wage-hours covariance. Insurable shocks, by contrast, always induce positive co-movement between hours and wages. Measurement error in hours reduces the observed covariance between wages (earnings divided by hours) and hours.²¹ For different parameter values, the model can generate both positive and negative covariance between hours and wages.

We now turn to the moments involving consumption:

$$var(\log \hat{c}_t) = v_\varphi - \frac{2(1+\sigma)}{\sigma+\gamma}v_{\varphi\alpha} + \left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\alpha t} + v_{\mu c}. \quad (22)$$

$$cov(\log \hat{h}_t, \log \hat{c}_t) = v_\varphi - \frac{(1+\sigma) + (1-\gamma)}{\sigma+\gamma}v_{\varphi\alpha} + \frac{(1-\gamma)(1+\sigma)}{(\sigma+\gamma)^2}v_{\alpha t} \quad (23)$$

$$cov(\log \hat{c}_t, \log \hat{w}_t) = \frac{1+\sigma}{\sigma+\gamma}v_{\alpha t}. \quad (24)$$

The variance of consumption is increasing in the variance of uninsurable preference heterogeneity and uninsurable wage shocks, as expected. The covariance term $v_{\varphi\alpha}$ reduces consumption dispersion because while a high positive uninsurable shock increases consumption, a high value of φ reduces consumption. Note the role of labor supply: if $\gamma > 1$, then a lower σ (higher Frisch) translates into lower consumption dispersion for given v_α because the impact of uninsurable wage shocks on earnings is dampened.²²

²¹The covariances involving earnings can be obtained easily as linear combinations of the three moments above.

²²By computing the variance of log earnings one would notice that the first three terms in consumption

The covariance between hours and consumption is increasing in the degree of preference heterogeneity since a large value for φ reduces hours, and thus earnings and consumption. The effect of uninsurable risk depends on the value of γ : when $\gamma > 1$, a positive uninsurable shock reduces hours worked but increases consumption. A positive correlation between the ability-component of productivity and taste for leisure ($v_{\varphi\alpha} > 0$) in general decreases the covariance.²³ For different parameter values, the model can generate both positive and negative covariance between consumption and hours.

The covariance between consumption and wages is only affected by uninsurable wage shocks: fluctuations in uninsurable productivity affect both wages and consumption in the same direction. The model predicts that this covariance should be positive. Finally, note that neither of the two covariances involving consumption are affected by the variance of insurable preference and wage shocks or by measurement error. This is a useful property for proving identification.

Micro moments These moments are computed as cross-sectional variances and covariances of individual changes in wages, hours and consumption between $t - 1$ and t :

$$\text{var}(\Delta \log \hat{w}_t) = v_{\omega t} + v_{\eta t} + v_{\theta t} + v_{\theta, t-1} + 2v_{\mu y} + 2v_{\mu h} \quad (25)$$

$$\text{var}(\Delta \log \hat{h}_t) = \left(\frac{1-\gamma}{\sigma+\gamma}\right)^2 v_{\omega t} + \frac{1}{\sigma^2} (v_{\eta t} + v_{\theta t} + v_{\theta, t-1}) + 2v_{\zeta} + 2v_{\mu h} \quad (26)$$

$$\text{cov}(\Delta \log \hat{h}_t, \Delta \log \hat{w}_t) = \frac{1-\gamma}{\sigma+\gamma} v_{\omega t} + \frac{1}{\sigma} (v_{\eta t} + v_{\theta t} + v_{\theta, t-1}) - 2v_{\mu h} \quad (27)$$

$$\text{var}(\Delta \log \hat{c}_t) = \left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\omega t} + 2v_{\mu c}, \quad (28)$$

$$\text{cov}(\Delta \log \hat{h}_t, \Delta \log \hat{c}_t) = \frac{(1-\gamma)(1+\sigma)}{(\sigma+\gamma)^2} v_{\omega t} \quad (29)$$

$$\text{cov}(\Delta \log \hat{c}_t, \Delta \log \hat{w}_t) = \frac{1+\sigma}{\sigma+\gamma} v_{\omega t}. \quad (30)$$

Note that while the first three moments based on wages and hours are quite involved, the last three moments for consumption are rather simple, since they only depend on the innovation to the uninsurable wage shock (plus measurement error in the case of the variance of changes in log consumption). In our data analysis, we only use cross-sectional consumption data from CEX (given the limitations of the panel dimension in the CEX).²⁴

dispersion are precisely the terms in the variance of earnings that depend on the uninsurable components α and φ .

²³A positive covariance term $v_{\varphi\alpha}$ may increase $\text{cov}(\log h, \log c)$ only in the presence of very strong income effects, i.e., $\gamma \gg 1$. In this case, individuals with high α and high φ work less but they also consume less than individuals with low α and low φ . See the consumption allocation in equation (14).

²⁴The panel aspect of CEX is quite weak. It consists of two, generally noisy, observations spaced nine months apart. See Davis (2003) for a discussion.

Thus we will prove identification when the set of micro moments is restricted to those that involve only wages or hours.

3.2 Life-cycle moments

Given the OLG structure of our economy, we can recover all the above moments conditional on a given age a . In particular, it is interesting to analyze the within-cohort change in the various dimensions of inequality as households age. Because productivity shocks are either permanent or i.i.d., the model-implied within-cohort changes in these moments will be *independent* of age.

Macro moments Let $\Delta var(\hat{x}_t^a)$ be the within-cohort change (i.e., between age $a-1$ and age a and between $t-1$ and t) in the variance of the measured variable x . Then, we obtain

$$\Delta var(\log \hat{w}_t^a) = v_{\omega t} + v_{\eta t} + v_{\theta t} - v_{\theta, t-1} \quad (31)$$

$$\Delta var(\log \hat{h}_t^a) = \left(\frac{1-\gamma}{\sigma+\gamma}\right)^2 v_{\omega t} + \frac{1}{\sigma^2} (v_{\eta t} + \Delta v_{\theta t}) \quad (32)$$

$$\Delta cov(\log \hat{h}_t^a, \log \hat{w}_t^a) = \frac{1-\gamma}{\sigma+\gamma} v_{\omega t} + \frac{1}{\sigma} (v_{\eta t} + \Delta v_{\theta t}). \quad (33)$$

The rise in wage inequality over the life-cycle is determined by the variance of the innovations to the permanent insurable and uninsurable components, and by the change in the variance of the transitory insurable component. Wage dispersion will increase over the life-cycle as permanent shocks cumulate. The model suggests that the variance of hours should be increasing over the life cycle for the same reasons as wages, though with different weights on the insurable and uninsurable permanent variances. In the log-utility case, only the former matters for hours. The model can generate a *decline* in the variance of hours over the life-cycle of a cohort only through a fall over time in the variance of the transitory component of productivity ($\Delta v_{\theta t} < 0$).

Whether the covariance between wages and hours rises or falls over the life cycle depends on the value for risk aversion and the relative size of permanent and transitory innovations. When $\gamma > 1$, hours and the permanent component of wages covary negatively, and thus the cumulation of permanent shocks pushes the covariance down as individuals age. The cumulation of the permanent insurable shocks pulls the hour wage covariance up, in proportion to the Frisch elasticity $\frac{1}{\sigma}$.

Turning to the life-cycle moments involving consumption, we obtain:

$$\Delta var(\log \hat{c}_t^a) = \left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\omega t} \quad (34)$$

$$\Delta cov(\log \hat{h}_t^a, \log \hat{c}_t^a) = \frac{(1-\gamma)(1+\sigma)}{(\sigma+\gamma)^2} v_{\omega t} \quad (35)$$

$$\Delta cov(\log \hat{c}_t^a, \log \hat{w}_t^a) = \frac{1+\sigma}{\sigma+\gamma} v_{\omega t}. \quad (36)$$

The change in the variance of consumption over the life-cycle is determined only by the size of the variance of the innovation to the uninsurable component, while the size of insurable shocks (both permanent and transitory) has no impact. The uninsurable-shock multiplier $\left(\frac{1+\sigma}{\sigma+\gamma}\right)^2$ for consumption is exactly one either in the log case or in the case of inelastic labor supply, i.e. $\sigma \rightarrow \infty$.²⁵

Hours and consumption are related only through uninsurable shocks. When $\gamma > 1$, hours move up in response to a negative shock, while consumption moves down, so that the covariance falls with age as permanent shocks cumulate over the life-cycle. The model predicts that the covariance between consumption and wages, in contrast, will increase over the life cycle, in proportion to the variance of uninsurable innovations.

Finally, note that none of the changes in variances and covariances over the life-cycle are affected by permanent or transitory preference heterogeneity or by measurement error, since these variances are all assumed constant. We will exploit this result to achieve identification.

Micro moments It is also straightforward to study how the variances and covariances of growth in wages, hours and consumption evolve over the life-cycle. For example, for wages

$$\Delta var(\Delta \log \hat{w}_t^a) = \Delta v_{\omega t} + \Delta v_{\eta t} + v_{\theta t} - v_{\theta, t-2}.$$

This expression makes clear that in the stationary version of the model all the moments in first-differences at any date t are common across cohorts of different ages.

Finally, note that covariances of levels and changes at lag one can be obtained as a linear combinations of some of the other moments. For example,

$$\begin{aligned} cov(\Delta \log \hat{w}_{t-1}^a, \Delta \log \hat{w}_t^a) &= \frac{1}{2} [\Delta var(\log \hat{w}_t^a) - var(\Delta \log \hat{w}_t^a)] = -(v_{\theta, t-1} + v_{\mu y} + v_{\mu h}), \quad (37) \\ cov(\log \hat{w}_{t-1}^{a-1}, \log \hat{w}_t^a) &= var(\log \hat{w}_{t-1}^{a-1}) - \frac{1}{2} [\Delta var(\log \hat{w}_t^a) - var(\Delta \log \hat{w}_t^a)] = v_{\alpha, t-1}^{a-1} + v_{\kappa, t-1}^{a-1} \end{aligned}$$

²⁵In general, the model predicts that the growth in the variance of earnings over the life-cycle should be larger than that for consumption but, interestingly, it could be lower than the growth in the variance of wages, if γ is sufficiently larger than one, so that wages and hours comove very negatively.

where v_{at}^a is the cumulated variance of the uninsurable shocks for a cohort of age a at time t . Hence, these moments are superfluous since they do not add any new information relative to the ones listed above. Given the specification of the stochastic process for shocks and measurement error, in the model covariances of the individual changes are all zero beyond lag one.

Taking stock The nature of the equilibrium in our heterogeneous-agent economy displays a sharp dichotomy: some shocks (within-group) are perfectly insured, while others (between-group) are not insured at all: when slicing the aggregate economy cross-sectionally, it emerges an equilibrium with partial insurance. One of the strengths of our approach is that it is not necessary to identify explicitly what the groups are in order to identify the parameters and learn about the insurability of shocks—it suffices to use the aggregate cross-section. However, we can go one step further. Our cross-sectional variances and covariances all contain both a within-group and a between-group term. Therefore, one can use the predictions of the model for the evolution of within and between group variances and covariances to try to identify the specific groups of agents where risk-sharing is particularly strong.²⁶

Another key feature of our cross-sectional moments is that they are all composed by separate terms capturing the role of preference heterogeneity/shocks, insurable productivity shocks, uninsurable productivity shocks, and measurement error. Since all these terms enter additively and are orthogonal, (co-)variance decompositions are always unique and immediate to compute, once we have the parameter estimates. In Section 5 we document these decompositions in detail.

We conclude this section by noting that while virtually every paper in the literature exploits only cross-sectional data on consumption and income to assess the degree of persistence and/or insurability of shocks, by using cross-sectional moments involving labor supply as well, we bring to the table new valuable, and previously unexploited, information.

4 Identification

This section exploits the closed-form solution for the cross-sectional second moments to prove identification of the structural model parameters. The conditions for identification depend on data availability and on the parametric restrictions one is willing to make. We

²⁶In the next draft of the paper, we will do an explorative analysis in this direction by looking at groups defined based on education, age, marital status and race.

therefore consider an array of different scenarios. Our baseline scenario (Proposition 2 below) is the model described in Section 2 under the assumption that one has access to an unbalanced panel on wages and hours (e.g., PSID) and a repeated cross-section on consumption (e.g., CEX), both supplying data for the time interval $t = 1, \dots, T$.

Next, we consider several variants on data availability. We will assume that the consumption cross-section begins later than the wage/hours panel (as for the CEX and PSID). We will assume that the consumption cross-section also contains wage and hours data, so one can exploit also the information contained in the covariances between consumption and hours, and between consumption and wages. We will also discuss the identification strategy when only one cohort of data on wages and hours is available (e.g., NLSY).

As for parametric restrictions, we will prove identification when a less rich structure of preference heterogeneity is imposed, i.e., $v_{\varphi\alpha} = 0$ and $v_{\zeta} = 0$. And we will study the case where we allow for cohort effects in $v_{\kappa 0}$ and $v_{\alpha 0}$.

PROPOSITION 2 [IDENTIFICATION] *Suppose one has access to an unbalanced panel on wages and hours and a repeated cross-section on consumption from $t = 1, \dots, T$. Assume that external estimates of both $v_{\mu h}$ (or $v_{\mu y}$) and $v_{\mu c}$ are available. Then, the parameters $\{\sigma, \gamma, v_{\varphi}, v_{\varphi\alpha}, v_{\zeta}, v_{\alpha 0}, v_{\kappa 0}, v_{\mu y}\}$ and the sequences $\{v_{\eta t}, v_{\omega t}, v_{\theta t}\}_{t=1}^T$ are identified.*

PROOF The proof is organized in five sequential steps.

1. Consider the following within-cohort changes (between age $a - 1$ and age a from $t - 1$ to t) in the following macro moments:

$$\begin{aligned} \Delta \text{var}(\log \hat{w}_t^a) &= v_{\omega t} + v_{\eta t} + \Delta v_{\theta t} \\ \Delta \text{var}(\log \hat{c}_t^a) &= \left(\frac{1 + \sigma}{\sigma + \gamma}\right)^2 v_{\omega t} \\ \Delta \text{var}(\log \hat{h}_t^a) &= \left(\frac{1 - \gamma}{\sigma + \gamma}\right)^2 v_{\omega t} + \frac{(v_{\eta t} + \Delta v_{\theta t})}{\sigma^2} \\ \Delta \text{cov}(\log \hat{w}_t^a, \log \hat{h}_t^a) &= \left(\frac{1 - \gamma}{\sigma + \gamma}\right) v_{\omega t} + \frac{(v_{\eta t} + \Delta v_{\theta t})}{\sigma}. \end{aligned}$$

These four nonlinearly independent equations, available from $t = 2, \dots, T$, identify σ, γ and $\{v_{\omega t}, v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$.²⁷ Intuitively, the preference parameters (γ, σ) mediate the extent to which the cumulation of insurable and uninsurable shocks translates into the growth over the life-cycle in the dispersion of wages, hours and consumption.

²⁷Given the nonlinearity in (γ, σ) , the solution for this system of equation may not be unique. We return to this point later.

2. The following expressions, available from $t = 1, \dots, T - 1$,

$$\begin{aligned}\frac{1}{2} \left(\text{var} \left(\Delta \log \hat{w}_{t+1}^{a+1} \right) - \Delta \text{var} \left(\log \hat{w}_{t+1}^{a+1} \right) \right) &= v_{\theta t} + v_{\mu y} + v_{\mu h} \\ \frac{1}{2} \left(\text{var} \left(\Delta \log \hat{h}_{t+1}^{a+1} \right) - \Delta \text{var} \left(\log \hat{h}_{t+1}^{a+1} \right) \right) &= \frac{1}{\sigma^2} v_{\theta t} + v_{\zeta} + v_{\mu h} \\ \frac{1}{2} \left(\text{cov} \left(\Delta \log \hat{h}_{t+1}^{a+1}, \Delta \log \hat{w}_{t+1}^{a+1} \right) - \Delta \text{cov} \left(\log \hat{h}_{t+1}^{a+1}, \log \hat{w}_{t+1}^{a+1} \right) \right) &= \frac{v_{\theta t}}{\sigma} - v_{\mu h},\end{aligned}$$

together with an external estimate of measurement error in hours $v_{\mu h}$, yield $v_{\mu y}$, v_{ζ} and the sequence $\{v_{\theta t}\}_{t=1}^{T-1}$. Using this latter sequence, from step 1 we obtain $\{v_{\eta t}\}_{t=2}^{T-1}$. Substituting the value for $v_{\theta, T-1}$ into $v_{\eta T} + \Delta v_{\theta T}$ from step 1, we obtain an estimate for $v_{\eta T} + v_{\theta T}$. This is a crucial step of the proof: by taking the difference between the dispersion in growth rates and the growth rate of within-cohort dispersion, we eliminate the variances of permanent innovations, which allows us to identify the transitory component.

3. Exploiting the unbalanced panel structure of the data on (w, h) and the assumption of no cohort effects, at year $t = 1$,

$$\begin{aligned}\text{var} \left(\log \hat{w}_1^1 \right) - \text{var} \left(\log \hat{w}_1^0 \right) &= v_{\omega 1} + v_{\eta 1} \\ \text{cov} \left(\log \hat{h}_1^1, \log \hat{w}_1^1 \right) - \text{cov} \left(\log \hat{h}_1^0, \log \hat{w}_1^0 \right) &= \frac{1 - \gamma}{\sigma + \gamma} v_{\omega 1} + \frac{1}{\sigma} v_{\eta 1}\end{aligned}$$

which allows us to identify $v_{\omega 1}$ and $v_{\eta 1}$. And using the same pair of moments for year $t = T$ allows us to identify $v_{\eta T}$ as well as $v_{\theta T}$ from step 2.

4. In any period t , for every age a we can use

$$\begin{aligned}\text{var} \left(\log \hat{w}_t^a \right) &= v_{\alpha}^a + v_{\kappa}^a + v_{\theta t} + v_{\mu y} + v_{\mu h} \\ \text{cov} \left(\log \hat{h}_t^a, \log \hat{w}_t^a \right) &= \frac{1 - \gamma}{\sigma + \gamma} v_{\alpha}^a + \frac{1}{\sigma} (v_{\kappa}^a + v_{\theta t}) - v_{\mu h}\end{aligned}$$

to identify the cumulated variances for age group a for the entire sample period $\{v_{\alpha}^a, v_{\kappa}^a\}_{t=1}^T$. In particular, evaluating these two equations for the new entrants of age $a = 0$, we can identify the variance of the fixed effects $v_{\kappa 0}$ and $v_{\alpha 0}$.²⁸

5. Finally, for any age group a in any period t , the pair of equations

$$\begin{aligned}\text{var} \left(\log \hat{c}_t^a \right) &= v_{\varphi} - \left(\frac{1 + \sigma}{\sigma + \gamma} \right) v_{\alpha \varphi} + \left(\frac{1 + \sigma}{\sigma + \gamma} \right)^2 v_{\alpha t}^a + v_{\mu c} \\ \text{var} \left(\log \hat{h}_t^a \right) &= v_{\varphi} + v_{\zeta} - \frac{2(1 - \gamma)}{\sigma + \gamma} v_{\varphi \alpha} + \left(\frac{1 - \gamma}{\sigma + \gamma} \right)^2 v_{\alpha t}^a + \frac{1}{\sigma^2} (v_{\kappa t}^a + v_{\theta t}) + v_{\mu h}\end{aligned}$$

²⁸At this point, it is easy to see that the unconditional cross-sectional moments $\text{var}(w_t)$ and $\text{cov}(w_t, h_t)$ identify the variances $v_{\alpha t}$ and $v_{\kappa t}$ taken over the entire cross-section (i.e., at across all ages) from $t = 1, \dots, T$.

identify the covariance $v_{\varphi\alpha}$ and the variance of preference heterogeneity v_φ , given an external estimate of the variance of measurement error in consumption $v_{\mu c}$. This concludes the proof.

It should be noted that there are several alternative set of moments which identify the model: we picked a simple and intuitive combination. We now turn to various alternative scenarios that we summarize in the following corollary.

COROLLARY 1 (A) Suppose data on $\text{var}(\log \hat{c}_t)$ are available only from $t = t^*$, with $0 < t^* < T$. Then, the model is identified under the same assumptions as in Proposition 2. (B) Suppose that either $v_{\varphi\alpha} = 0$ or $\text{cov}(\log \hat{h}_t, \log \hat{c}_t)$ is available for at least one period t . Then, the model can be identified without any a priori knowledge of the size of the measurement error in consumption $v_{\mu c}$. (C) Suppose $v_\zeta = 0$. Then, the model is identified without any a priori knowledge of the size of the measurement error in hours $v_{\mu h}$ (or earnings $v_{\mu y}$). (D) Suppose we allow for cohort effects in $v_{\alpha 0}$ and $v_{\kappa 0}$. Then, under the same assumptions of Proposition 2, one can identify the parameters $\{\sigma, \gamma, v_\varphi, v_{\varphi\alpha}, v_\zeta, v_{\mu y}\}$ as well as the sequences $\{v_{\theta t}, v_{\alpha 0t}, v_{\kappa 0t}\}_{t=1}^T$ and $\{v_{\eta t}, v_{\omega t}\}_{t=2}^T$. (E) Suppose wage and hours data are only available for one cohort from $t = 1, \dots, T$. Then, under the same assumptions of Proposition 2, we can identify the parameters $\{\sigma, \gamma, v_\varphi, v_{\varphi\alpha}, v_\zeta, v_{\alpha 0}, v_{\kappa 0}, v_{\mu y}\}$ as well as the sequences $\{v_{\omega t}\}_{t=2}^T$, $\{v_{\eta t}\}_{t=2}^{T-1}$, $\{v_{\theta t}\}_{t=1}^{T-1}$, and $v_{\eta T} + v_{\theta T}$.

PROOF We prove the five parts of the corollary one by one.

(A) The elasticities γ and σ can be identified using the four moments in step 1 in the main proof from $t = t^*, \dots, T$. Given values for (γ, σ) , $\Delta \text{var}(\log \hat{w}_t^a)$ and $\Delta \text{var}(\log \hat{h}_t^a)$ then identify $\{v_{\omega t}, v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$. The proof proceeds exactly as for Proposition 2.

(B) When $v_{\varphi\alpha} = 0$, it is easy to see from step 5 that the two equations for $\text{var}(\log \hat{c}_t^a)$ and $\text{var}(\log \hat{h}_t^a)$ are now linearly independent functions of only two unknowns, v_φ and $v_{\mu c}$, instead of three. When $\text{cov}(\log \hat{h}_t^a, \log \hat{c}_t^a)$ is available, step 5 can be augmented with an additional equation:

$$\text{cov}(\log \hat{h}_t^a, \log \hat{c}_t^a) = v_\varphi - \frac{(1 + \sigma) + (1 - \gamma)}{\sigma + \gamma} v_{\varphi\alpha} + \frac{(1 - \gamma)(1 + \sigma)}{(\sigma + \gamma)^2} v_{\alpha t}^a.$$

Now we have three linearly independent equations in $\{v_\varphi, v_{\alpha\varphi}, v_{\mu c}\}$.

(C) When $v_\zeta = 0$, in step 2, the three equations simplify, so the last two equations permit identification of $v_{\theta t}$ and $v_{\mu h}$, and the first one yields $v_{\mu y}$.

(D) If we allow for cohort effects $\{v_{\kappa 0t}, v_{\alpha 0t}\}$, step 4 computed for period $t = 1, \dots, T$ allows identification of all the cohort effects. However, now the first equation in step 3

becomes

$$\text{var}(\log \hat{w}_1^1) - \text{var}(\log \hat{w}_1^0) = v_{\omega 1} + v_{\eta 1} + (v_{\alpha 0 0} - v_{\alpha 0 1}) + (v_{\kappa 0 0} - v_{\kappa 0 1})$$

and we cannot identify the pair $(v_{\eta 1}, v_{\omega 1})$ any longer since we do not have estimates of $v_{\alpha 0 0}$ and $v_{\kappa 0 0}$.²⁹ Evaluating the equations in step 3 for year $t = T$ shows that we can still identify $v_{\eta T}$ and $v_{\theta T}$.

(E) This statement is a direct consequence of the fact that only step 3 of the main proof uses the unbalanced panel structure of the data, which is unavailable for cohort data by definition. This concludes the proof.

Note that all the equations used for identification are linearly independent in the variances. However, in step 1 of the proof those four equations are nonlinear functions of (γ, σ) . Even though this is unlikely to be a problem given the abundance of moments where these two parameters appear, it may be that multiple admissible pairs of (γ, σ) could generate the same data.

One way to ensure this issue does not hinder identification is to include $\Delta \text{cov}(\log \hat{w}_t^a, \log \hat{c}_t^a)$ among the moments used in step 1. The ratio between $\Delta \text{var}(\log \hat{c}_t^a)$ and $\Delta \text{cov}(\log \hat{w}_t^a, \log \hat{c}_t^a)$ allows us to uniquely identify the term $(1 + \sigma) / (\sigma + \gamma)$, hence $v_{\omega t}$ and $(v_{\eta t} + \Delta v_{\theta t})$. Using then the expression for $\Delta \text{var}(\log \hat{h}_t^a)$ and acknowledging that the coefficient on $v_{\omega t}$ is simply $\left(\frac{1+\sigma}{\sigma+\gamma} - 1\right)^2$ allows us to uniquely identify σ , and thus γ .

4.0.1 A comparison with Blundell-Preston (1998)

Blundell and Preston (1998, BP thereafter) developed a theoretical framework to identify permanent and transitory income innovations from joint household data on consumption c and income y . Their baseline model is that of a household with quadratic-utility, discount rate equal to the interest rate, inelastic labor supply and no liquidity constraint, i.e., the classical Hall (1978) permanent income hypothesis (PIH) where consumption follows a random walk. They assume that labor income is the sum of two components, a permanent shock (ϖ_t) and a transitory shock (τ_t) .³⁰ They show that, in their model, for a cohort of age a at time t ,

$$\begin{aligned} \Delta \text{var}(c_t^a) &= \Delta \text{cov}(c_t^a, y_t^a) \simeq v_{\varpi t} \\ \Delta \text{var}(y_t^a) - \Delta \text{var}(c_t^a) &\simeq \Delta v_{\tau t}. \end{aligned} \tag{38}$$

²⁹One could still identify the pair $(v_{\eta 1}, v_{\omega 1})$ under the additional assumption $v_{\alpha 0 0} = v_{\alpha 0 1}$ and $v_{\kappa 0 0} = v_{\kappa 0 1}$ (e.g., before $t = 1$ the model is in “steady-state”).

³⁰They allow τ_t to be a MA(1) process, but this generalization is not crucial for their results and for the point we make below.

In our model, shocks are to wages and propagate to earnings through labor supply decisions. From equations (31) to (34), in our economy:

$$\begin{aligned} \Delta var(\log c_t^a) &= \Delta cov(c_t^a, y_t^a) = \left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\omega t} \\ \Delta var(\log y_t^a) - \Delta var(\log c_t^a) &= \left(\frac{1+\sigma}{\sigma}\right)^2 (v_{\eta t} + \Delta v_{\theta t}). \end{aligned} \quad (39)$$

A comparison of (38) and (39) reveals immediately the role of endogenous labor supply. The change in the variance of log consumption correctly identifies the variance of uninsurable permanent wage shocks v_{ω} only if either $\gamma = 1$ or $\sigma \rightarrow \infty$. In the first case, income and substitution effect of an uninsurable shock on labor supply cancel out, thus the original permanent innovation to productivity translates one for one into consumption. In the second case, trivially, labor supply has no role. Given our point estimates for (γ, σ) in Table 2 (see below), our study suggests that BP could underestimate the average variance of the pure uninsurable permanent component of labor productivity by over 40%, i.e. the estimated value of the term $\left(\frac{1+\sigma}{\gamma+\sigma}\right)^2$ is 0.56. Its time trend would be consistently estimated.

The second equation in the BP model provides a difference in difference estimator for the change in the variance of the transitory earnings shock. Once again, our model suggests that the more elastic is labor supply (σ low), the larger will be the difference between the left hand side and the true change in the transitory productivity shock $\Delta v_{\theta t}$. Intuitively, when agents' labor supply is very sensitive to transitory insurable shocks, the response of labor income is magnified. Given the point estimate for σ , BP would overestimate this variance by over 45%.

This latter equation highlights another important issue: productivity shocks could contain a permanent component (η_t in our notation) that is insurable, and that it is not passed through consumption. This possibility is not considered explicitly by BP, but it is the core of the analysis by Blundell, Pistaferri and Preston (2006) (see Section 5.3.5).

5 Estimation

This section is organized as follows. We begin by describing how we construct the PSID and CEX samples used for the estimation. Next, we outline the details of the minimum distance estimator. We then report the parameter estimates, discuss the fit in various dimensions, and present variance decompositions over time and over the life-cycle. Finally, we perform a robustness analysis on various aspects of the sample and the model.

5.1 Data

Our data are drawn from two data sets, the *Michigan Panel Study of Income Dynamics* (PSID), and the *Consumer Expenditure Survey* (CEX). From the 1968-1997 waves of the PSID, we construct an unbalanced sample containing information on individual demographic characteristics, individual annual hours and annual earnings. The PSID asks questions about earnings in the previous year, so our data refers to the period 1967-1996.³¹

The CEX data contain detailed information on nondurable and durable household consumption, as well as on demographic characteristics of household members, their individual earnings and hours worked. Consistent data over time are available only since the 1980 survey. Our starting point is the same sample used by Krueger and Perri (2006). This includes all households who are complete income respondents and for whom we observe data from four consecutive quarterly interviews. We use two alternative measures of consumption expenditure, both constructed by Krueger and Perri. One only includes nondurable consumption goods, while the other includes an estimate of the service flow from some durables (vehicles and housing).

Since we use all these data jointly, we try to construct comparable samples by imposing the same selection criteria across the two datasets. In every year, we select all males between ages 25 and 54. We drop observations if earnings are top-coded, if the hourly wage is below half the federal minimum wage in that year, or if the individual works less than 520 hours or more than 5096 hours that year, i.e., less than one full-time week for a quarter or more than 14 hours a day, every day of the week.³² In both datasets, the hourly wage is computed as annual pre-tax labor earnings divided by annual hours worked. Consumption for such individuals is defined as consumption expenditure of their household expressed in per-adult-equivalent units.³³ All monetary variables are deflated using the Consumer Price Index (CPI) and expressed in 1992 dollars.

The final PSID sample contains 51,512 individual-year observations, comprising 4,834 individuals of which 812 are present in the sample for at least 20 of the 30 possible years. The final CEX sample contains 23,220 individual/year observations, i.e., an average of 968 individuals per year. Table 1 shows the number of observations lost at each stage of

³¹The more recent waves of the PSID, after 1997, include questions about income only every second year and are still in their “early release” format, hence we exclude them from our analysis. Since we do not include observations from the Survey of Economic Opportunities (SEO), a sub-sample of low-income households, the sample is representative of the US population and weights are not used in any calculations.

³²This last restriction is imposed to limit the size of measurement error in hours, we relax it later.

³³Our equivalence scale is the same Census equivalence scale used by Krueger and Perri (2006).

Table 1: Sample selection

	Sample Size	
	PSID	CEX
Baseline sample	290,375	69,816
Exclude SEO sample	(136,078)	NA
Keep male heads of households	(81,836)	(19,232)
Drop obs with missing earnings data	(4,452)	(20,740)
Drop top-coded earnings	(63)	NA
Hours restrictions	(1,649)	(880)
Minimum wage restrictions	(897)	(229)
Age restriction	(13,888)	(5,515)
Final Sample	51,512	23,220

the selection process.³⁴

To purge the data from variation due to demographic factors that the model is not designed to address we regress individual log wages, hours and consumption (the former two in both datasets, the latter only in CEX) on year and race (white, black or other) dummies, and on a quartic in age.³⁵ We then use the residuals from these regressions to construct variances and covariances in levels and first-differences for age/year cells constructing by grouping observations into six non-overlapping age classes: 25-29, 30-34, and so on until 50-54.

When constructing the unconditional moments by year, we average across all age groups.³⁶ When constructing moments by age-group to document the typical life-cycle evolution of dispersion, we control for time effects in both model and data.³⁷

5.2 Estimation method

The structural estimation of the model uses the minimum distance estimator (MDE) introduced by Chamberlain (1984) which minimizes a weighted squared sum of the differences between each moment in the model and its data counterpart.

³⁴A comparison across the two datasets in each year of the overlapping sample period 1980-1996 shows that (i) the proportion of college graduates, (ii) average earnings, and (iii) average hours worked are very similar in levels and move closely together over time. Individuals are two years older, on average, in the CEX. The time series for the variance of hours in PSID tracks very closely the one in CEX. The same is true for wages, except for the last year in the sample, when CEX wages become much more dispersed than PSID wages. Overall, we conclude that the two samples are broadly consistent.

³⁵Recall that controlling for age effects in the individual wage, hours and consumption allocations is consistent with the way we constructed equilibrium cross-sectional (co-) variances in Section 3.

³⁶Since the model's period is one year, to construct the model's moments by 5-year age brackets, we assume a constant yearly survival probability $\delta = 1 - 1/(54 - 25)$ and aggregate across the five adjacent ages.

³⁷Effectively, we regress age/year observations for second moments on a full set of time-dummies and plot the residuals by age group, averaging across all cohorts. This is how the lines labelled "Data" are constructed in Figures 2-4. Consistently, the lines labelled "Model" are generated through the closed-form expressions for age a by setting the variances of the shocks to their estimated sample mean.

The structural parameters to be estimated in the model are preference parameters $\{\gamma, \sigma, v_\varphi, v_{\alpha\varphi}, v_\zeta\}$, productivity parameters $\{v_{\alpha 0}, v_{\kappa 0}\}$, and $\{v_{\theta t}, v_{\omega t}, v_{\eta t}\}_{t=1}^T$, where $T = 30$ is the length of the PSID sample, together with variances of measurement error $\{v_{\mu y}, v_{\mu h}, v_{\mu c}\}$. Given our discussion of identification above, we set $v_{\mu h}$ exogenously, and include the covariance between log consumption and log hours and the covariance between log consumption and log wages from the CEX. In the benchmark estimation we use nondurable consumption. In total, we have $3T + 9 = 99$ parameters to be estimated. Denote by Θ this parameter vector.

Let $\mathbf{m}(\Theta)$ be the $(J \times 1)$ vectors of the stacked theoretical covariances with typical element $m(\Theta, j)$ where the index $j = 1, \dots, J$ denotes the position in the vector. Correspondingly, we define $\hat{\mathbf{m}}$ as the vector of empirical covariances with typical element \hat{m}_j . In the baseline estimation, we use 1,149 moment conditions.

Our MDE solves the following minimization problem

$$\min_{\Theta} [\hat{\mathbf{m}} - \mathbf{m}(\Theta)]' \mathcal{W} [\hat{\mathbf{m}} - \mathbf{m}(\Theta)], \quad (40)$$

where \mathcal{W} is a $(J \times J)$ weighting matrix. Standard asymptotic theory implies that the estimator $\hat{\Theta}$ is consistent, asymptotically Normal, and has asymptotic covariance matrix $V = (D' \mathcal{W} D)^{-1} D' \mathcal{W} \Delta \mathcal{W} D (D' \mathcal{W} D)^{-1}$, where the matrix $D \equiv \mathbb{E}[\partial \mathbf{m}(\Theta) / \partial \Theta']$ and the matrix $\Delta \equiv \mathbb{E}[(\hat{\mathbf{m}} - \mathbf{m}(\Theta))(\hat{\mathbf{m}} - \mathbf{m}(\Theta))']$ are estimated via their empirical analogs to compute standard errors.

To implement the estimator, we need a choice for \mathcal{W} . The bulk of the literature follows Altonji and Segal (1996) who argue that in common applications there is a substantial small sample bias, hence using the identity matrix for \mathcal{W} is a strategy superior to the use of the optimal weighting matrix characterized by Chamberlain. With this choice, the solution of (40) reduces to a nonlinear least square problem.³⁸

5.3 Estimation results

We begin by discussing the parameter estimates. We then analyze the fit of the model along the life-cycle dimension and the time-series dimension. In both cases, we perform our variance decomposition to study the determinants of inequality over the life-cycle and over time.

³⁸We deviate only slightly from Altonji and Segal. We normalize the weight of each age group/year moment by one. Thus, we weight the unconditional moments by age with the number of years in the (PSID or CEX) sample, and the unconditional moments by year with the number of age brackets.

5.3.1 Parameter estimates

Our discussion of identification highlighted the need for an external estimate of either $v_{\mu h}$ or $v_{\mu y}$. We set the variance of measurement error in hours to 0.020, which represents 25% of the unconditional variance of hours, and is in line with the findings of Bound et al. (1994) based on the 1986 PSID Validation Study.

Table 2 (column 1) reports parameter estimates and asymptotic standard errors.³⁹ For our two key preference parameters, we estimate $\gamma = 2.97$ and $\sigma = 4.84$. The implied intertemporal elasticity of substitution, 0.34, is within the standard range of existing estimates (for a survey, see Attanasio, 1999). The implied Frisch elasticity of labor supply is 0.21, a value that is consistent with the microeconomic evidence for males (for a survey, see Blundell and MaCurdy, 1999). While γ is tightly estimated, σ has a large standard error. We return to this point in the robustness analysis, where we show that by imposing further parametric restrictions we can get more precise estimates of σ .

The variance of preference heterogeneity is three times larger than the variance of insurable preference shocks. The covariance term $v_{\varphi\alpha}$ is negative, as predicted by human capital theory, albeit statistically insignificant.

The initial variance of the insurable and uninsurable wage components $v_{\alpha 0}$ and $v_{\epsilon 0}$ are comparable in size, so they each explain roughly half of the initial wage variance. The variance of the innovation to the uninsurable shock (Figure 1) declines slightly until the early 1980s. It then increases again, with peaks in the mid 1980s and in the early 1990s. The variance of the permanent insurable shock starts rising before the uninsurable shock in the mid 1970s, and remains high throughout the 1980s. Recall that the 1980s is the period when the bulk of the rise in inequality takes place. The variance of the i.i.d. insurable productivity shock is four times larger than its permanent counterpart and almost twice the variance of i.i.d. insurable preference shocks. It grows steadily throughout the sample and it peaks in the first half of the 1990s, consistently with Moffitt and Gottschalk's (2002) findings from estimated earnings dynamics.⁴⁰

Table 2 also shows that measurement error in earnings and consumption account for, respectively, 5% and 29% of the total cross-sectional variances for these variables. While our estimate of the measurement error in earnings is lower than some of the existing external estimates based on validation studies (e.g., Bound et al., 1994, report measurement errors in earnings to be between 15% and 30%), the estimate of $v_{\mu c}$ confirms the widespread view that measurement error in consumption data is sizeable.

³⁹In the next draft of the paper we will compute bootstrapped standard errors.

⁴⁰The variances in this figure are plotted together with their 90% asymptotic confidence interval.

Table 2: Parameter Estimates

Parameter	(1)	(2)	(3)	(4)	(5)	(6)	(7)
γ	2.968 (0.143)	2.367 (0.066)	2.872 (0.040)	2.931 (0.078)	2.763 (0.143)	3.025 (0.019)	2.796 (0.090)
σ	4.836 (2.507)	5.611 (1.873)	1.987 (0.271)	6.734 (0.022)	2.289 (0.225)	2.857 (0.172)	2.779 (0.080)
v_φ	0.038 (0.001)	0.040 (0.001)	0.056 (0.000)	0.038 (0.000)	0.055 (0.002)	0.043 (0.000)	0 –
$v_{\varphi\alpha}$	-0.010 (0.020)	-0.005 (0.009)	-0.016 (0.010)	0 –	0 –	-0.024 (0.004)	0 –
v_ζ	0.012 (0.001)	0.012 (0.000)	0.024 (0.000)	0.012 (0.001)	0.025 (0.001)	0 –	0.015 (0.000)
$v_{\alpha 0}$	0.079 (0.003)	0.098 (0.003)	0.079 (0.003)	0.072 (0.006)	0.077 (0.003)	0.084 (0.000)	0.077 (0.003)
$v_{\kappa 0}$	0.067 (0.004)	0.049 (0.003)	0.069 (0.006)	0.074 (0.005)	0.069 (0.006)	0.063 (0.000)	0.072 (0.004)
$v_{\mu y}$	0.013 (0.009)	0.011 (0.007)	0.001 (0.008)	0.007 (0.000)	0.001 (0.012)	0.014 (0.001)	0.025 (0.005)
$v_{\mu c}$	0.048 (0.020)	0.074 (0.008)	0.0378 (0.010)	0.057 (0.000)	0.055 (0.001)	0.033 (0.004)	0.110 (0.000)

Note: (1) baseline model; (2) consumption definition includes services from durables; (3) sample w/o hours restriction; (4) $v_{\varphi\alpha} = 0$; (5) $v_{\varphi\alpha} = 0$ + sample w/o hours restriction; (6) $v_\zeta = 0$; (7) $v_\varphi = v_{\varphi\alpha} = 0$.

Figure 1 also shows the evolution of the unconditional variances of insurable and uninsurable shocks. The latter variance declines throughout the 1970s in accordance with the well-known decline in the skill premium. It rises sharply in the first half of the 1980s and keeps rising slowly through the 1990s. The variance of insurable shocks starts at roughly the same level as the uninsurable component in the late 1960s, but it starts rising earlier, in line with the well documented rise in within-group wage dispersion. Its increase continues steadily throughout the 1980s. At the end of the sample, the insurable component accounts for around 60% of the total dispersion, meaning that 2/3 of the observed rise in labor market dispersion was insurable in nature. Of this rise in insurable risk, around half of it was permanent and half transitory.

5.3.2 Life-cycle: fit and decomposition

Fit Figures 2, 3 and 4 compare model and data in the life cycle dimension. In the data, the variance of log wages increases by 14 log points between age 25 and 54, displaying a slightly concave pattern. The variance of log earnings increases by a smaller amount, roughly 11 log points. The variance of log hours is flat, thus what accounts for this smaller rise in the variance of log earnings is the fact that the covariance between hours and wages falls over the life-cycle. The variance of log non-durable consumption grows by only 3.0

log points over the life cycle. As we emphasized in Heathcote, Storesletten, and Violante (2005a), this number represents a much smaller increase than previously reported in the pioneering work of Deaton and Paxson (1994).⁴¹ To emphasize the different dynamics of these series, we have plotted them on the same scale.

The model is able to match the observed large rise in labor productivity dispersion and, simultaneously, the small rise in consumption dispersion thanks to the fact that a large part of the cumulated permanent shocks are insurable. However, it generates an increase in the variance of earnings that is mildly above the one in the data, by 3 log points, in the age-range 25-35. The reason is twofold. First, the variance of hours in the model is slightly increasing, while in the data it is decreasing in that age range. Equation (20) for the variance of hours in the model shows that the model cannot generate a declining pattern for this moment: over the life-cycle both insurable and uninsurable shocks cumulate, leading to more hours dispersion. Since γ is not too far from one, and the Frisch elasticity is low, overall the model predicts a weak rise.⁴² Second, the model predicts an increasing correlation between hours and wages between 25 and 35, while in the data this moment has a small downward trend. From equation (21) one can see the model can predict both positive and negative trends, depending on the relative size of insurable and uninsurable innovations to productivity, as well as on the value for γ . At the point estimates for parameter values, the impact of insurable shocks dominates.

The correlation between consumption and wages grows both in the model and in the data, by roughly the same amount. However, the model generates a declining pattern for the consumption-hours correlation while this correlation is increasing in the data. The cumulation of permanent uninsurable shocks over the life cycle explains these dynamics.

Note that the model can generate, at the same time, a negative wage-hour correlation and a positive consumption-hour correlation. A casual look at expressions (21) and (23) shows that, given $\gamma > 1$ (and ignoring the $v_{\varphi\alpha}$ term which is negligible), this sign combination is possible only thanks to preference heterogeneity. We discuss this point further in Section 5.3.4.

⁴¹Deaton and Paxson (1994, Figure 8) report an increase in the variance of log consumption of almost 20 log points between ages 25 and 55. The reason is that the Deaton and Paxson analysis, by controlling only for cohort-effects, implicitly incorporates rising dispersion over time into the age profiles. Moreover, their study covers the period 1980-1990, precisely when the bulk of the rise in cross-sectional dispersion is concentrated (see also Slesnick and Ulker 2004, for a related discussion), while our sample covers a longer time period. Another important difference is that Deaton and Paxson do not use any equivalence scale.

⁴²Within a self-insurance model, Kaplan (2006) can generate the observed decline in the variance of hours at the beginning of the life-cycle. This is due to (1) age effects in the transitory component of wages which are decreasing at younger ages; (2) a non-degenerate distribution of initial wealth which increases the variance of hours for younger agents.

Finally, Figure 4 documents moments in first-differences over the life cycle. The model has a very stark implication: all these moments should be constant. In the data the variance of changes in log wages and the correlation between changes in log wages and log hours are indeed quite flat. The variance of changes in log hours, in contrast, declines between ages 25 and 35, but is essentially constant thereafter. Note that the average levels of all these moments are matched successfully.

Variance decomposition To decompose the life-cycle moments into preference heterogeneity/shocks, uninsurable/insurable productivity shocks and measurement error, we set productivity shocks to their average value over the sample period and use the theoretical moment expressions to produce artificial life-cycle profiles. Figures 5 and 6 tell an interesting story. The rise in wages and earnings dispersion over the life cycle is explained by the cumulation of permanent productivity shocks that are, for the most part, insurable. By contrast, preference heterogeneity and measurement error are the main determinants of the variance of hours worked, accounting for over 3/4 of its level, though movements in the variance of hours over the life-cycle are entirely explained by productivity shocks. Clearly, the rise in consumption dispersion is entirely due to uninsurable wage shocks, even though measurement error and preference heterogeneity account for quite significant fractions of the level of consumption dispersion, around 1/4 for both factors.

The pattern for the wage-hour covariance is determined by the tension between uninsurable and insurable shocks, with measurement error also playing a sizeable role. Preference heterogeneity is the main determinant of the positive covariance between consumption and hours: workers with a strong taste for leisure work less, earn less and consume less.

We conclude that there is no unique answer to the question: What determines measured life-cycle inequality? The answer depends on the variable of interest: for hours it's preference heterogeneity and measurement error, for wages and earnings it's productivity shocks, while for consumption it's a mix of everything.

5.3.3 Time series: fit and decomposition

Fit The variance of log wages increases from 0.24 to 0.34 between 1967 and 1996. This 10 log-point increase is mostly concentrated in the period 1978-1992. The rise in the variance of earnings is larger by roughly 5 log points. Behind this more rapid increase lies a substantial rise in the wage-hour correlation, whereas there is no noticeable change in the variance of log hours. CEX data on consumption are only available in a consistent way since the 1980 survey. The variance of log consumption has risen since 1980 by a

much smaller amount relative to the variance of earnings: only 4 log points. This fact has been previously documented and discussed by Krueger and Perri (2006), and Attanasio, Battistin and Ichimura (2006).

The model is able to replicate a steep increase in earnings inequality vis-a-vis the relatively small rise in consumption dispersion and the flat profile of hours dispersion (Figure 7). Since most of the observed rise in inequality over the period was insurable, this did not translate into a large rise in consumption dispersion. At the same time, a relatively low Frisch elasticity and large preference heterogeneity are needed to explain why the variance of hours has not increased in the wake of such a large surge in insurable productivity dispersion.

Figure 8 shows that the model also replicates the rise in the wage-hours correlation. This emerges for the same reason that growth in earnings dispersion exceeds growth in consumption dispersion: most of the increase in wage dispersion was insurable in nature, which drives up co-movement between hours and wages. Thus the model teaches us that the message about the nature of the rise in wage dispersion contained in the picture for consumption versus earnings dispersion is the same message contained in the picture for the wage-hours correlation. Data on labor supply seem therefore to validate the patterns emerging from consumption data.

The CEX data also document a remarkable rise in the wage-consumption correlation, by 10 log points, which the model replicates thanks to the increase of the uninsurable component over time. This same force induces a slightly declining pattern in the consumption-hours correlation, whereas the data display a flat pattern since 1980.

Figure 9 documents the moments in first-differences. The variance of changes in log wages shows a slow but continuous rise over the entire sample period, with a peak in the early 1990s. This moment is matched almost perfectly by the model, since it identifies the variance of the transitory insurable shock θ_t (recall step 2 of the identification proof). The variance of changes in log hours is quite stable both in the data and the model: the low Frisch elasticity makes this moment quite unresponsive to the trends in insurable and uninsurable productivity dispersion. The correlation between changes in wages and changes in hours rises substantially (by 15 log points), especially in the 1970s, the period of the sharp rise in insurable wage dispersion. The model replicates this rise successfully.

Variance decomposition The results of the time-series decompositions are in Figures 10 and 11. As mentioned already, two thirds of the rise in wage dispersion between 1967-1996 is attributable to insurable risk. For earnings inequality, the insurable share is even larger. This is because part of the rise in earnings inequality is due to the increase

in the wage-hour covariance, which in turn reflects the finding that insurable shocks become increasingly important. Overall, more than 80% of the measured rise in earnings dispersion is insurable from the point of view of households.

Even though we do not have data on consumption before 1980, we can still identify the variances of the uninsurable shock prior to that date: this information is embedded in the consumption dispersion of the older cohorts. Projecting backward, the model predicts a U-shape in consumption inequality over the whole sample period, with inequality declining in the 1970s. Even though we do not have direct evidence of this phenomenon, because of lack of data, this pattern is potentially consistent with the well known dynamics of the skill premium over the period. The fall in consumption inequality in the 1970s is also broadly consistent with evidence provided by Slesnick (2001, chapter 6) who, notwithstanding data comparability issues, uses CEX data for 1961 and 1973 in order to construct a longer series for U.S. consumption dispersion.

5.3.4 Robustness

The first of our robustness experiments is to use a definition of consumption that includes services from durables (column 2 in Table 2) instead of nondurable consumption in the estimation. Under this new definition, the variance of log consumption is higher in levels, but it rises by roughly the same amount over the sample period (3 points). It is not surprising therefore that the model requires a higher initial variance of the uninsurable component and larger measurement error. All the other parameters estimates are very close to the benchmark estimates, so the precise definition of consumption used does not seem to affect the results.

In column 3, we relax our sample-selection restrictions on hours worked. The cross-sectional dispersion of log hours now displays a stronger increase over the sample period, from 0.10 to 0.20 in the mid 1980s followed by a modest decline in the 1990s. These larger swings in the variance and the rise in the insurable variance suggest a larger Frisch elasticity, thus σ is now estimated to be around 2. Most of the other parameter values are unchanged.⁴³ This exercise suggests that the largest adjustments in hours worked in the face of rising inequality may have occurred at the bottom of the hours distribution, along the extensive margin.

Given that $v_{\varphi\alpha}$ is not significantly different from zero, in column (4) we have restricted this parameter to be zero. When using the baseline sample, the Frisch elasticity decreases

⁴³In this experiment, we double the external estimate of $v_{\mu h}$ to 0.04 so that measurement error in hours still represents around 25% of the total cross-sectional variation.

slightly and is much more precisely estimated. In column (5) we combined this constraint with the sample without an hours restriction and, once again, we estimate a larger Frisch with a lower standard error.

In all these alternative exercises, the fit of the model remains as good as in the baseline case and the estimates of the productivity shocks follow the same general dynamics.

To understand the role of preference shocks and preference heterogeneity, we run two experiments where we shut down these channels, one at the time. When $v_\zeta = 0$ (column 6), the model fits the data very well, except for the variance of changes in hours whose level is somewhat underestimated. The model compensates for the lack of preference shocks through a higher Frisch elasticity, which raises the level of dispersion in hours, but generates a counter-factual rise in hours dispersion over time. Overall, though, transitory preference shocks do not seem critical to understanding the evolution of cross-sectional moments either over the life-cycle or through time.

When we set $v_\varphi = v_{\varphi\alpha} = 0$ (column 7), we reach the opposite conclusion. The major problem with this version of the model in which labor productivity is the only source of heterogeneity is its inability to jointly generate a negative covariance between hours and wages and a positive covariance between consumption and hours. The former requires $\gamma > 1$, but for $\gamma > 1$ and $v_\varphi = 0$, the latter covariance must be negative (recall equations (21) and (23)). Moreover, in the absence of preference heterogeneity, the model can generate a large variance of hours only when σ is very low. This causes other problems, since a high Frisch elasticity means that the variance of hours and the covariance between hours and wages are very sensitive to the variance of insurable shocks. As a result the model overestimates the rise in both moments over this period. The best possible fit of the model without preference heterogeneity is displayed in Figure 12.

Throughout the estimation, we used both micro and macro moments. One natural question is whether focusing on only one set of moments yields a distorted perspective on the evolution of inequality over time and over the life-cycle. To this end, we ran one estimation where we substantially overweighted the micro moments.⁴⁴ The main discrepancy relative to the baseline case is that in this case we estimate a larger contribution of uninsurable shocks, relative to insurable ones. This can be explained by the absence of consumption micro-moments: the consumption data call sharply for a large insurable component.

⁴⁴One cannot completely exclude the moments in levels because identification of some parameters would be lost.

5.3.5 The transmission from wage shocks to consumption

Our model can be used to measure the fraction of labor productivity shocks that transmit to consumption. We define pass-through coefficients as the average percentage change in individual consumption in response to one percent change in the wage. Let π_t^w denote the pass-through coefficients for permanent wage shocks at date t .⁴⁵ There are two potential reasons for incomplete pass-through. First, labor supply mediates the transmission from wage shocks to earnings. Second, if changes in earnings reflect insurable shocks, they will not be reflected in changes in consumption.

In the context of our model, permanent wage changes for an individual are given by $\omega_t + \eta_t$. The elasticity of consumption with respect to an uninsurable permanent innovation ω_t is equal to $\frac{1+\sigma}{\sigma+\gamma}$ (see equation (6)), while consumption does not respond to permanent insurable innovations. Thus the average cross-sectional response of consumption to permanent wage changes at date t is given by

$$\pi_t^w = \left(\frac{1 + \sigma}{\sigma + \gamma} \right) \left(\frac{v_{\omega t}}{v_{\eta t} + v_{\omega t}} \right) \quad (41)$$

where the second term represents the fraction of permanent shocks that are uninsurable.

The higher is the risk-aversion coefficient γ , the smaller is π_t^w . Intuitively, a strong income effect on labor supply dampens the impact of uninsurable wage shocks such that they have a smaller impact on consumption. The impact of the Frisch labor supply elasticity ($1/\sigma$) depends on the value for γ . When $\gamma > 1$, a larger Frisch elasticity further dampens the transmission of uninsurable wage shocks to earnings and, ultimately, to consumption. If $\gamma = 1$ (when income and substitution effects on labor supply cancel out) or if $\sigma \rightarrow \infty$ (the inelastic labor supply case), then π_t^w is simply the share of uninsurable risk in all permanent shocks.

Given our estimates of the structural parameters in Table 2, we can compute the average pass-through coefficient over our sample period:

$$\bar{\pi}^w = \left(\frac{1 + 4.836}{4.836 + 2.968} \right) \left(\frac{0.0036}{0.0060 + 0.0036} \right) = 0.28.$$

To isolate the role of labor supply relative from the role of other insurance channels in the economy, note that as $\sigma \rightarrow \infty$, $\bar{\pi}^w \rightarrow 0.38$. Therefore, flexible labor supply reduces the pass-through from wages to consumption by one fourth.

Comparison with Blundell-Pistaferri-Preston In our model, the primitive source of labor income variation comes from wages, while earnings are endogenous. However, in

⁴⁵Transitory shocks do not transmit to consumption since they are insurable.

order to compare our results to Blundell, Pistaferri and Preston (2006, BPP hereafter) it is useful to repeat our exercise on earnings. Let π_t^y denote the pass-through coefficients from permanent innovations to earnings into consumption. Permanent uninsurable earnings innovations translate one-for-one into consumption, while permanent insurable earnings innovations have no effect. Thus π_t^y is the share of uninsurable earnings innovations in all permanent earnings changes:

$$\pi_t^y = \frac{\left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\omega t}}{\left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 v_{\omega t} + \left(\frac{1+\sigma}{\sigma}\right)^2 v_{\eta t}} = \frac{1}{1 + \left(1 + \frac{\gamma}{\sigma}\right)^2 \frac{v_{\eta t}}{v_{\omega t}}} \quad (42)$$

Given our estimates in Table 2, it is straightforward to compute that, on average in our sample: $\bar{\pi}^y = 0.19$. Thus permanent innovations to earnings pass through to consumption to a much lesser extent than permanent shocks to wages. For higher elasticities of labor supply, the difference between π_t^w and π_t^y would be even larger. Suppose, for example, that $\sigma \simeq 0$. Then $\pi_t^y \simeq 0$ while $\pi_t^w \simeq (1/\gamma) v_{\omega t} / (v_{\omega t} + v_{\eta t})$. In this case, insurable earnings variation, which is not passed through to consumption, becomes infinitely large, while insurable wage variation is unaffected by the value of σ .

In our model the primitive shocks are to wages and earnings are endogenous. We therefore view π_t^w , the pass through coefficient for wages as the most appropriate measure of pass-through, and interpret π_t^y as overstating the true extent to which permanent shocks can be insured.

By manipulating our moments in first differences, an observables-based expression for π_t^y can be derived as follows:

$$\pi_t^y = \frac{\text{cov}(\Delta \log \hat{c}_t^a, \Delta \log \hat{y}_t^a)}{\text{var}(\Delta \log \hat{y}_t^a) + \text{cov}(\Delta \log \hat{y}_t^a, \Delta \log \hat{y}_{t-1}^{a-1}) + \text{cov}(\Delta \log \hat{y}_t^a, \Delta \log \hat{y}_{t+1}^{a+1})}.$$

To operationalize this expression requires panel data on consumption. Based on an empirical approach dating back to Hall and Mishkin (1982), Blundell, Pistaferri and Preston (2006, BPP hereafter) propose to jointly use panel data on consumption and earnings to measure the extent of *permanent earnings variation* which is transmitted to household consumption.⁴⁶ BPP estimate π_t^y using a very similar ratio involving covariances of changes in earnings and in consumption.

The BPP observables-based estimate for $\bar{\pi}^y$ using data on male earnings and a definition of household consumption that excludes durables (like ours) is 0.29. This is somewhat

⁴⁶While Hall and Mishkin used only food consumption from PSID, BPP impute a measure of nondurable consumption to PSID based on the relation between nondurable and food consumption expenditures estimated from CEX.

larger than our model-based estimate, but both approaches lead to the conclusion that the lion's share of permanent innovations to earnings do not transmit to consumption. The BPP approximated model, however, does not yield an alternative expression in terms of structural parameters. Our equation (42) therefore offers a more structural interpretation of the partial insurance coefficient in the BPP analysis.

6 Conclusions

This paper has laid out a new theoretical framework to study the evolution of inequality in labor supply and consumption over the life-cycle and over time. The most distinguishing feature of the framework is that it is an incomplete-markets model that can be solved analytically. The theoretical tool that we exploited to reach this outcome is an extension of the no-bond-trade equilibrium studied by Constantinides and Duffie (1996). Instead of holding among individuals, in our economy the no-trade result holds across groups within which there is full risk-sharing. The model therefore allows for any degree of partial insurance in between a bond economy and complete markets.

The tractability gained from this structure allows us to derive closed forms expressions for the equilibrium cross-sectional moments of the joint distribution of wages, hours and consumption. These analytical expressions make the theoretical characterization very transparent, and the empirical analysis simple to conduct, in spite of the number of parameters to estimate and moment conditions to handle. Therefore, the current set up could be further extended in several directions.

From a theoretical perspective, there are several extensions that we intend to explore in the near future. First, one can derive the same set of moments for non-separable preferences of the Cobb-Douglas class (as done in Heathcote, Storesletten and Violante, 2005b, with a simpler productivity process). Second, one can link the variances of insurable and uninsurable innovations to the aggregate state of the economy. Within this extended model, one can estimate the cyclicity of idiosyncratic risk and quantify its role for asset pricing (as in Constantinides and Duffie, 1996, and Storesletten, Telmer and Yaron, 2006), as well as for the welfare costs of business cycles. With respect to the latter issue, interestingly, fluctuations in the size of insurable and uninsurable individual risk will impact aggregate output through their effect of individual labor supply decisions, even without any shocks at the aggregate level. For example, times of large insurable uncertainty will be "good times" in the cycle. Third, it is possible to introduce taxes and transfers in a way that maintains tractability and allows for the explicit study of the impact of taxation

on the cross-sectional moments of interest. Fourth, under some conditions, we can allow for a discrete participation decision.

From an empirical perspective, one could further exploit the Normality assumption on the shocks. Given the closed-form expressions for individual consumption and labor supply as functions of productivity and preference shocks, it is possible to write down explicitly the likelihood function for individual histories. Then, using the approach of Blundell, Pistaferri and Preston (2005), a measure of total non-durable consumption could be imputed to PSID households drawing from information on the demand for food in the CEX. This augmented PSID panel could then be used for the likelihood estimation.

References

- [1] Abowd, J.M. and D. Card (1989), “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57, 411-445.
- [2] Aiyagari, R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109, 659-684.
- [3] Altonji, J.G. and L.M. Segal (1996), “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business and Economic Statistics*, 14, 353-366.
- [4] Attanasio, O. (1999), “Consumption,” in J.B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, vol. 1., Amsterdam: Elsevier.
- [5] Attanasio, O., E. Battistin, and H. Ichimura (2005), “What Really Happened to Consumption Inequality in the US?” in E. Berndt and C. Hulten (eds.), *Measurement Issues in Economics—The Paths Ahead: Essays in Honor of Zvi Griliches*, (Chicago: University of Chicago Press)
- [6] Attanasio, O. and S.J. Davis (1996), “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104(6), 1227-1262.
- [7] Attanasio, O. and N. Pavoni (2006), “Testing Private Information Models with Asset Accumulation,” *Mimeo*, UCL.
- [8] Attanasio, O. and J.V. Rios-Rull (2000), “Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?,” *European Economic Review*, vol. 44, 1225-1258.
- [9] Benabou, R. (2002), “Tax and Education Policy in a Heterogeneous Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?,” *Econometrica*, 70(2), 481-517.
- [10] Blundell, R. and MaCurdy (1999), “Labour Supply: a Review of Alternative Approaches,” in Ashenfelter, O. and D. Card, *Handbook of Labor Economics*, vol. 3A, Amsterdam: Elsevier.
- [11] Blundell, R., L. Pistaferri and I. Preston (2006), “Consumption Inequality and Partial Insurance,” *Mimeo*, Stanford University.
- [12] Blundell, R., L. Pistaferri and I. Preston (2005), “Imputing Consumption in the PSID Using Food Demand Estimates from the CEX,” *IFS Working Paper*, WP04/27.
- [13] Blundell, R. and I. Preston (1998), “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 113(2), 603-640.
- [14] Bound, J., Brown, C., Duncan, and G. Rodgers (1994), “Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data,” *Journal of Labor Economics*, 12(3), 345-368.
- [15] Browning, M., L.P. Hansen, J.J. Heckman (1999), “Micro Data and General Equilibrium Models,” in *Handbook of Macroeconomics*, vol. 1A, Amsterdam: Elsevier.

- [16] Caballero, R.J. (1990), "Consumption Puzzles and Precautionary Savings," *Journal of Monetary Economics*, 25(1), 113-136.
- [17] Campbell, J., and A. Deaton (1989), "Why is Consumption So Smooth?," *Review of Economic Studies*, 56, 357-374.
- [18] Castaneda, A., J. Diaz-Gimenez and J.V. Rios-Rull (2003), "Accounting for the US Earnings and Wealth Inequality," *Journal of Political Economy*, 111, 818-857.
- [19] Chamberlain, G. (1984), "Panel Data," in Griliches and Intriligator (eds), *Handbook of Econometrics*, vol. 3, New York: North-Holland.
- [20] Chatterjee, S. (2004), "Transitional Dynamics and the Distribution of Wealth in a Neoclassical Growth Model," *Journal of Public Economics*, 54, 97-119.
- [21] Constantinides, G.M. and D. Duffie (1996), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104(2), 219-240.
- [22] Cunha, F. and J. Heckman (2006), "A Framework for the Analysis of Inequality," *Journal of Macroeconomics*, forthcoming.
- [23] Cunha, F., J. Heckman, and S. Navarro (2005), "Separating Uncertainty from Heterogeneity in Life-Cycle Earnings," *Oxford Economic Papers*, 57(2), 191-261.
- [24] Cutler, D., and L. Katz (1991), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1990s," *American Economic Review*, 82, 546-551.
- [25] Davis, S. (2003), Comment to "On the Welfare Consequences of the Increase in Inequality in the United States," by D. Krueger and F. Perri, *Macroeconomics Annual 2003*, 121-132.
- [26] Deaton, A. (1991), "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221-1248.
- [27] Deaton, A. (1997), *The Analysis of Household Surveys*, Baltimore: John, Hopkins University Press.
- [28] Deaton, A. and C. Paxson (1994), "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102, 437-467.
- [29] Fernandez-Villaverde, J. and D. Krueger (2005), "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?," *Mimeo*, University of Pennsylvania.
- [30] Guvenen F. (2007), "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, vol. 97(3), 687-712.
- [31] Hall, R.E. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971-987.
- [32] Hall, R.E. and F. Mishkin (1982), "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, vol. 50(2), 461-81.

- [33] Heathcote, J., K. Storesletten and G.L. Violante (2004), “The Macroeconomic Implications of Rising Wage Inequality in the U.S.,” *Mimeo*, NYU.
- [34] Heathcote, J., K. Storesletten and G.L. Violante (2005a), “Two Views of Inequality Over the Life Cycle,” *Journal of the European Economic Association*, 3(2-3), 765-775.
- [35] Heathcote, J., K. Storesletten and G.L. Violante (2005b), “Insurance and Opportunities: The Welfare Implications of Rising Wage Dispersion,” *Mimeo*, NYU.
- [36] Huggett, M. (1996), “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics*, 38, 469-494.
- [37] Huggett, M., G. Ventura, and A. Yaron (2006a), “Human Capital and Earnings Distribution Dynamics,” *Journal of Monetary Economics*, 53, 265-290.
- [38] Huggett, M., G. Ventura, and A. Yaron (2006b), “Sources of Lifetime Inequality,” *Mimeo*, Georgetown University.
- [39] Kaplan, G. (2006), “The Patterns of Lifecycle Inequality,” *Mimeo*, NYU.
- [40] Katz, L., and D. Autor (1999), “Changes in the Wage Structure and Earnings Inequality,” in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 3A, Amsterdam: Elsevier.
- [41] Krueger, D. and F. Perri (2006), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies*, 73(1), 191-261.
- [42] Krusell, P. and A. Smith (1998), “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867-896.
- [43] Livshits, I., J. MacGee, and M. Tertilt, “Consumer Bankruptcy: A Fresh Start,” (2006), *American Economic Review*, forthcoming.
- [44] Low, H., C. Meghir, and L. Pistaferri (2006), “Wage Risk and Employment Risk Over the Life-cycle,” *Mimeo*, UCL.
- [45] Maliar L. and S. Maliar (2003), “The Representative Consumer in the Neoclassical Growth Model with Idiosyncratic Shocks,” *Review of Economic Dynamics*, 6, 362-380.
- [46] Meghir, C. and L. Pistaferri (2004), “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72, 1-32.
- [47] Moffitt, R.A. and P. Gottschalk (2002), “Trends in the Transitory Variance of Earnings in the United States,” *Economic Journal*, 112(3), C68-C73.
- [48] Primiceri, G. and T. van Rens (2006), “Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality,” CEPR DP 5881.
- [49] Slesnick, D. (2001), *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*, Cambridge: Cambridge University Press.

- [50] Slesnick, D. and A. Ulker (2004), "Inequality and the Life-Cycle: Age, Cohort Effects and Consumption," *Mimeo*, University of Texas - Austin.
- [51] Storesletten, K., C. Telmer. and A. Yaron (2004), "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695-717
- [52] Storesletten, K., C. Telmer, and A. Yaron (2006), "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," *Review of Economic Dynamics*, forthcoming.
- [53] Uhlig, H. (1996), "A Law of Large Numbers for Large Economies," *Journal of Economic Theory*, 8(1), 41-50.
- [54] Wang, N. (2003), "Caballero Meets Bewley: The Permanent-Income Hypothesis in General Equilibrium," *American Economic Review*, 93(3), 927-936.
- [55] Yaari, M.E. (1965), "Uncertain Lifetime, Life Insurance and the Theory of the Consumer," *Review of Economic Studies*, 32, 137-150.

7 Appendix

PROOF OF PROPOSITION 1

Guess no bond trade across islands Suppose the net demand for bonds is zero on every island, for every history s^t , i.e.

$$\int \int \int b_t^a(s^t) dF_{\kappa t}^a dF_{\theta t} dF_{\zeta} = 0$$

Planner problem Given no bond trade between islands and complete markets on each island, the welfare theorems apply, so the competitive equilibrium allocation can be computed as the outcome of an island-specific social planner problem. Since agents on an island are ex-ante identical, the planner weights must be equal for all agents. Moreover, since each island by assumption transfers zero net financial wealth between periods and preferences are time-separable, the island-specific planner problem is static.

The state- s_t planner problem for an island populated by agents of type (a, φ, α_t) is then given by

$$\max_{c_t^a(s_t), h_t^a(s_t)} \int \int \int \left(\frac{c_t^a(s_t)^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi + \sigma\zeta_t) \frac{h_t^a(s_t)^{1+\sigma}}{1+\sigma} \right) dF_{\kappa t}^a dF_{\theta t} dF_{\zeta} \quad (43)$$

subject to the resource constraint

$$\begin{aligned} \int \int \int c_t^a(s_t) dF_{\kappa t}^a dF_{\theta t} dF_{\zeta} &= \int \int \int \exp(\alpha_t + \kappa_t + \theta_t) h_t^a(s_t) dF_{\kappa t}^a dF_{\theta t} dF_{\zeta} \\ &= \exp(\alpha_t) \int \exp(\kappa_t + \theta_t) \left(\int h_t(s_t) dF_{\zeta} \right) dF_{\kappa t}^a dF_{\theta t}. \end{aligned}$$

Solving for island-specific allocations The first-order conditions with respect to $c_t^a(s_t)$ and $h_t^a(s_t)$ are, respectively,

$$\begin{aligned} c_t^a(s_t)^{-\gamma} &= \chi_t \\ \exp((\gamma + \sigma)\varphi + \sigma\zeta) h_t^a(s_t)^\sigma &= \chi_t \exp(\alpha_t) \exp(\kappa_t + \theta_t), \end{aligned}$$

where χ_t is the multiplier on the resource constraint. Combining the two conditions gives

$$h_t^a(s_t) = c_t^a(s_t)^{\frac{-\gamma}{\sigma}} \exp\left(\frac{\alpha_t}{\sigma}\right) \exp\left(\frac{1}{\sigma}(\kappa_t + \theta_t)\right) \exp\left(-\frac{(\gamma + \sigma)}{\sigma}\varphi - \zeta_t\right). \quad (44)$$

Note that from the first-order conditions, $c_t^a(s_t)$ is the same for all agents on the island, and as such it cannot depend on κ_t , θ_t , or ζ_t . Using this fact, and substituting (44) into the planner's island-specific resource constraint, gives

$$\begin{aligned} c_t^a(s_t) &= \exp(\alpha_t) \int \exp(\kappa_t + \theta_t) \left(\int h_t(s_t) dF_{\zeta} \right) dF_{\kappa t}^a dF_{\theta t} \\ &= \exp(\alpha_t) \int \exp(\kappa_t + \theta_t) \left(\int c_t^a(s_t)^{\frac{-\gamma}{\sigma}} \exp\left(\frac{1}{\sigma}(\kappa_t + \theta_t)\right) \exp\left(-\frac{\gamma + \sigma}{\sigma}\varphi - \zeta_t\right) dF_{\zeta} \right) dF_{\kappa t}^a dF_{\theta t} \\ &= c_t^a(s_t)^{\frac{-\gamma}{\sigma}} \exp\left(\frac{1 + \sigma}{\sigma}\alpha_t\right) \exp\left(-\frac{\gamma + \sigma}{\sigma}\varphi\right) \left(\int \exp(-\zeta_t) dF_{\zeta} \right) \left(\int \exp\left(\frac{1 + \sigma}{\sigma}(\kappa_t + \theta_t)\right) dF_{\kappa t}^a dF_{\theta t} \right). \end{aligned}$$

Taking logs:

$$\begin{aligned} \log c_t^a(s_t) &= \frac{1 + \sigma}{\sigma + \gamma} \alpha_t - \varphi + M_{\zeta} + \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp\left(\frac{1 + \sigma}{\sigma}(\kappa_t + \theta_t)\right) dF_{\kappa t}^a dF_{\theta t} \right) \\ &= \frac{1 + \sigma}{\sigma + \gamma} \alpha_t - \varphi + M_{\zeta} + \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp\left(\frac{1 + \sigma}{\sigma}\theta_t\right) dF_{\theta t} \right) \\ &\quad + \frac{\sigma}{\sigma + \gamma} \log \left(\int \exp\left(\frac{1 + \sigma}{\sigma}\kappa_0\right) dF_{\kappa_0} \right) + \frac{\sigma}{\sigma + \gamma} \sum_{j=t-a+1}^t \log \left(\int \exp\left(\frac{1 + \sigma}{\sigma}\eta_j\right) dF_{\eta_j} \right) \\ &= \frac{1 + \sigma}{\sigma + \gamma} \alpha_t - \varphi + \mathcal{M}_t^a, \end{aligned}$$

which is the expression in the text.

From the logarithm of the intra-temporal first order condition (44):

$$\begin{aligned}\log h_t^a(s_t) &= -\frac{\gamma}{\sigma} \log c_t^a(s_t) + \frac{1}{\sigma} (\alpha_t + \kappa_t + \theta_t) - \frac{\gamma + \sigma}{\sigma} \varphi - \zeta_t \\ &= \left(\frac{1}{\sigma} - \frac{\gamma}{\sigma} \frac{1 + \sigma}{\sigma + \gamma} \right) \alpha_t + \left(\frac{\gamma}{\sigma} - \frac{\gamma + \sigma}{\sigma} \right) \varphi - \zeta_t + \frac{1}{\sigma} (\kappa_t + \theta_t) - \frac{\gamma}{\sigma} \mathcal{M}_t^a \\ &= \frac{1 - \gamma}{\sigma + \gamma} \alpha_t - \varphi - \zeta_t + \frac{1}{\sigma} (\kappa_t + \theta_t) - \frac{\gamma}{\sigma} \mathcal{M}_t^a,\end{aligned}$$

which is the expression in the text.

Solving for equilibrium prices— We now compute prices supporting these allocations. The (island-specific) Arrow-Debreu price of state s_t in period t is $P_t^{\tilde{\omega}}(s_t) \equiv \beta^t (c_t^a(s_t))^{-\gamma}$. Note that this price is indexed by $\tilde{\omega}$, indicating that this price applies to claims traded on an island with one particular sequence of length a for $\tilde{\omega} \equiv \{\dots, \omega_{t-1}, \omega_t, \omega_{t+1}, \dots\}$. The island-specific stochastic discount factor is then

$$\begin{aligned}m_t &\equiv \frac{P_t^{\tilde{\omega}}(s_t)}{P_{t-1}^{\tilde{\omega}}(s_{t-1})} = \beta \left(\frac{c_t^a(s_t)}{c_{t-1}^a(s_{t-1})} \right)^{-\gamma} = \beta \frac{\exp\left(-\gamma \frac{1+\sigma}{\sigma+\gamma} \alpha_t + \gamma \frac{1}{\sigma+\gamma} \varphi - \gamma \mathcal{M}_t^a\right)}{\exp\left(-\gamma \frac{1+\sigma}{\sigma+\gamma} \alpha_{t-1} + \gamma \frac{1}{\sigma+\gamma} \varphi - \gamma \mathcal{M}_{t-1}^a\right)} \\ &= \beta \exp(-\gamma (M_{\eta t} + \Delta M_{\theta t})) \exp\left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_t\right).\end{aligned}$$

The price of the one-period non-contingent bond in period $t - 1$ must then be given by

$$q_{t-1} = \mathbb{E}_{t-1}[m_t] = \beta \exp(-\gamma (M_{\eta t} + \Delta M_{\theta t})) \int \exp\left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_t\right) dF_{\omega t}, \quad (45)$$

which is the expression in the text. Moreover, the price in period $t - 1$ of the one-period contingent claims paying one unit of consumption in state s_t is given by

$$Q_{t-1}(s_t, s^{t-1}) = m_t f_{s_t} = \beta \exp(-\gamma (M_{\eta t} + \Delta M_{\theta t})) \exp\left(-\gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_t\right) f_{s_t},$$

where f_{s_t} is the density function of F_{s_t} .

Verifying the no bond trade guess— We must now verify the initial conjecture that islands do not trade. The bond price implied by the no trade guess (45) does not depend on any island-specific shocks, since the expected growth rate of consumption is the same across all islands. Hence, at the proposed bond price q_t , the choice of zero net savings on each island (i.e., $\int \int \int b_t(s_t) dF_{\kappa t}^a dF_{\theta t} dF_{\zeta} = 0$) satisfies the Euler equation on every island, confirming the guess. With zero net demand for non-contingent bonds on each island, the world market for bonds must clear by construction.

Financial wealth— Finally, we compute the implied individual financial wealth. The net transfer from the planner to an agent with state s_t and age a in period t is given by:

$$\begin{aligned}T_t^a(s_t) &= c_t^a(s_t) - w_t(s_t) h^a(s_t) = \left(1 - \frac{w_t(s_t) h^a(s_t)}{c_t^a(s_t)}\right) c_t^a(s_t) \\ &= \left(1 - \exp\left(-\zeta_t + \frac{1 + \sigma}{\sigma} (\kappa_t + \theta_t) - \frac{\sigma + \gamma}{\sigma} \mathcal{M}_t^a\right)\right) \exp\left(-\varphi + \frac{1 + \sigma}{\sigma + \gamma} \alpha_t + \mathcal{M}_t^a\right).\end{aligned}$$

The budget constraints in (4) imply that individual net financial wealth at date t must equal the present discounted value of all future transfers from $t+1$ onward, evaluated at the island-specific stochastic discount factors. Thus we have:

$$\begin{aligned}\hat{d}_t^a(s_t) &= \mathbb{E}_t \sum_{j=1}^{\infty} \frac{P_{t+j}^{\tilde{\omega}}(s_{t+j})}{P_t^{\tilde{\omega}}(s_t)} T_{t+j}^{a+j}(s_{t+j}) \\ &= \mathbb{E}_t \sum_{j=1}^{\infty} \frac{P_{t+j}^{\tilde{\omega}}(s_{t+j})}{P_t^{\tilde{\omega}}(s_t)} c_{t+j}^{a+j}(s_{t+j}) \left\{ 1 - \mathbb{E}_t \left[\exp\left(-\zeta_{t+j} + \frac{1 + \sigma}{\sigma} \left(\kappa_t + \sum_{i=1}^j \eta_{t+i} + \theta_{t+j}\right) - \frac{\sigma + \gamma}{\sigma} \mathcal{M}_{t+j}^{a+j}\right)\right] \right\} \\ &= c_t^a(s_t) \left(1 - \frac{\exp\left(\frac{1+\sigma}{\sigma} \kappa_t\right)}{\int \exp\left(\frac{1+\sigma}{\sigma} \kappa\right) dF_{\kappa t}^a}\right) \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j \left(\frac{c_{t+j}^{a+j}(s_{t+j})}{c_t^a(s_t)}\right)^{1-\gamma},\end{aligned}$$

which is the expression in the text. This concludes the proof.

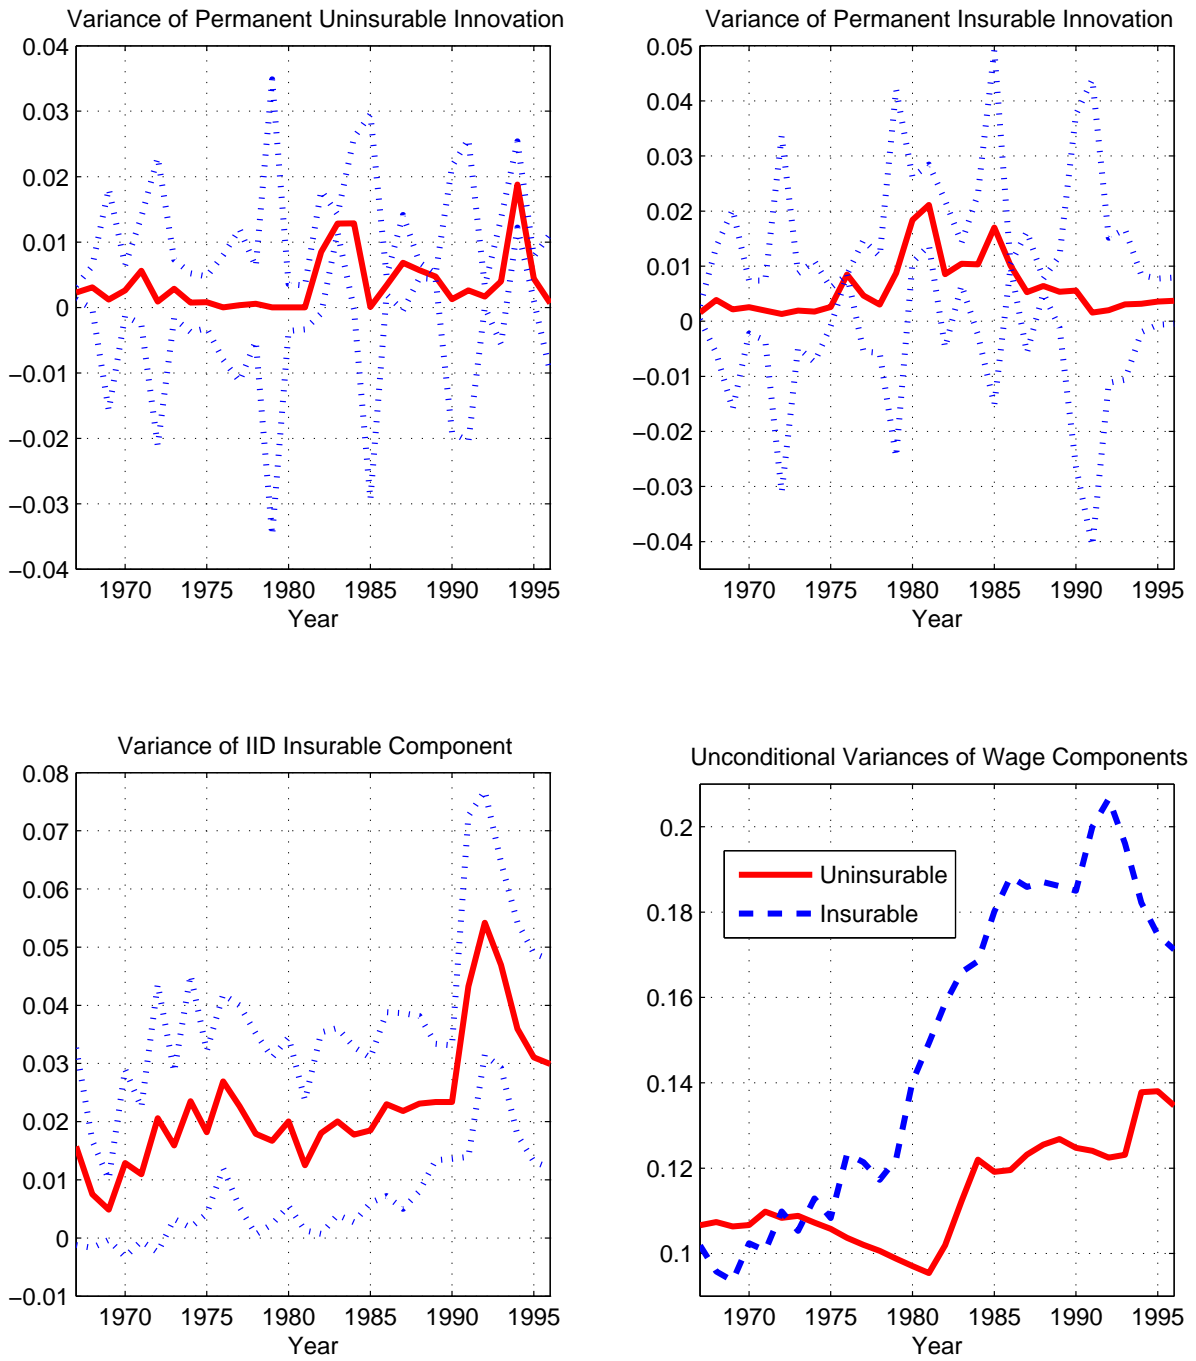


Figure 1: Estimated variances of uninsurable innovation, insurable innovations, transitory insurable shocks and unconditional variances.

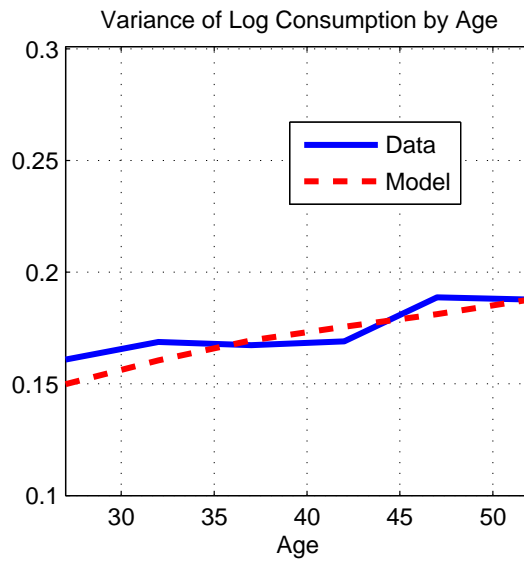
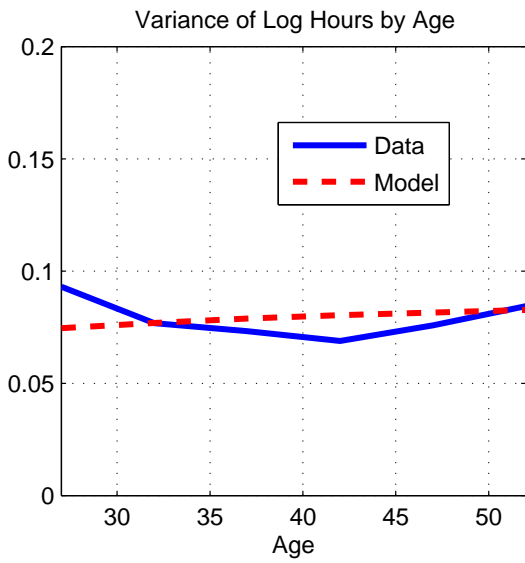
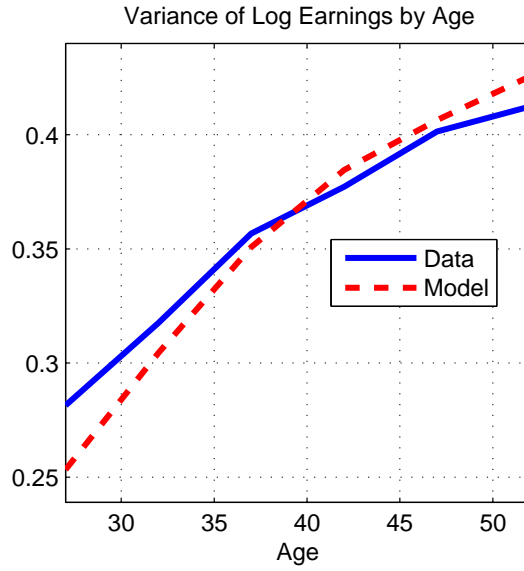
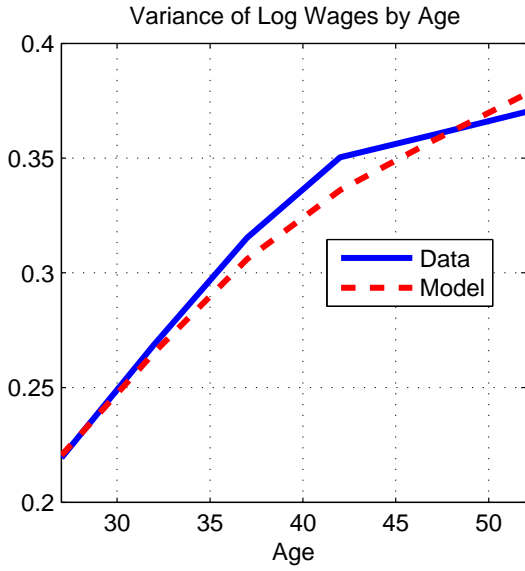


Figure 2: Life-cycle dimension – Model's fit






Figure 3: Life-cycle dimension – Model's fit




Figure 4: Life-cycle dimension – Model's fit



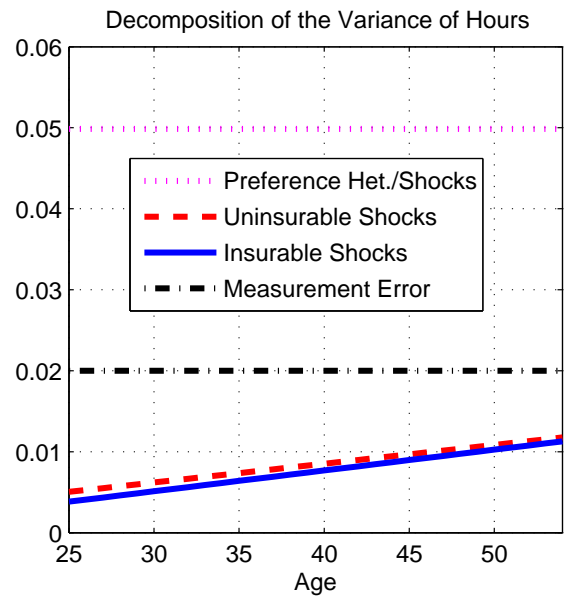
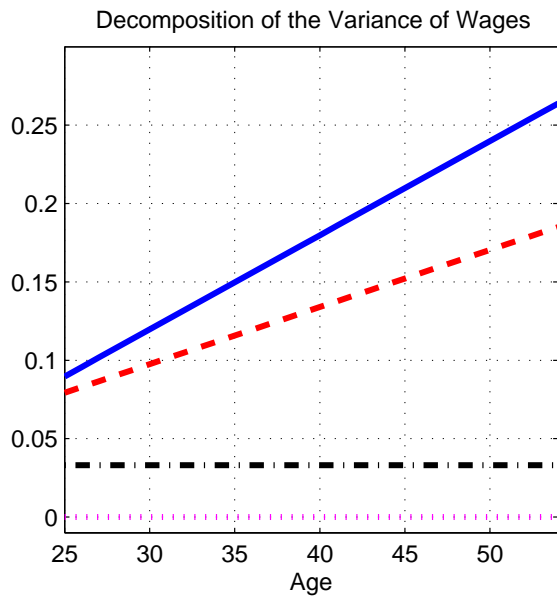


Figure 5: Life-cycle dimension – Decomposition


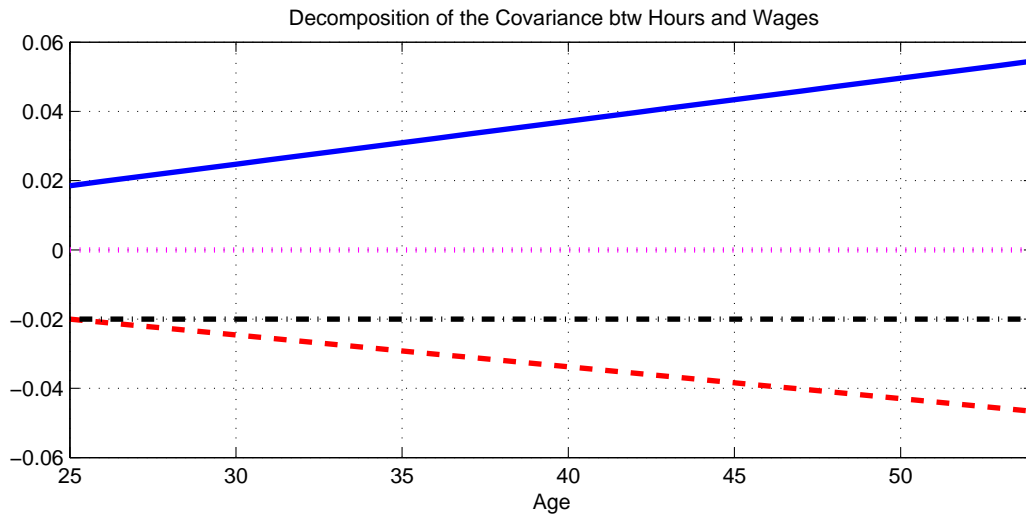


Figure 6: Life-cycle dimension – Decomposition

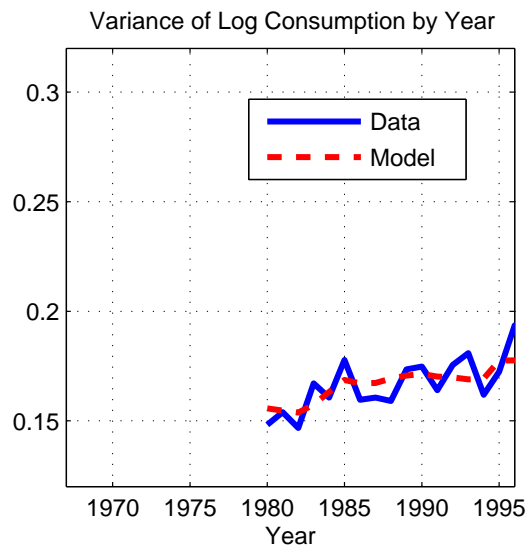
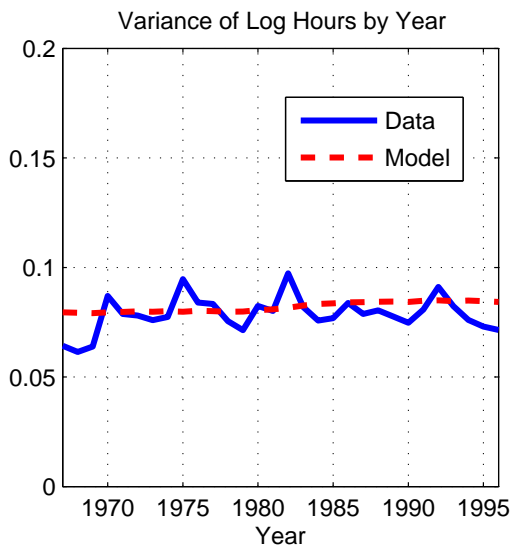
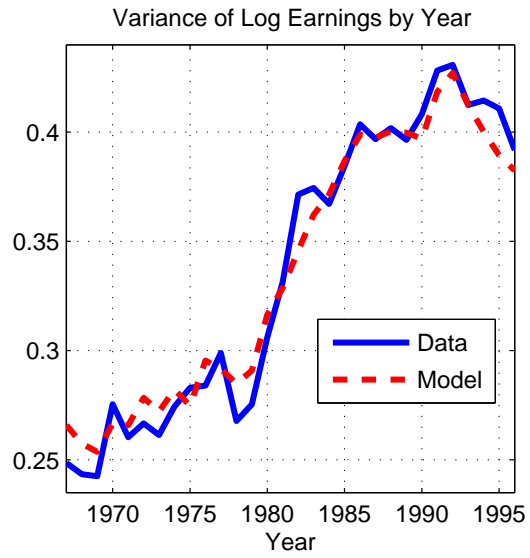
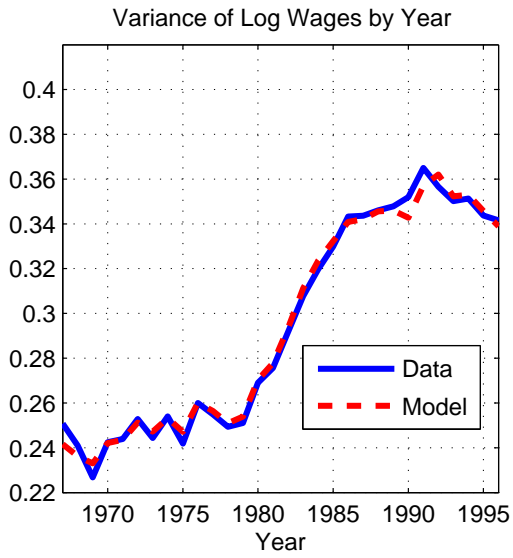


Figure 7: Time series dimension – Model's fit

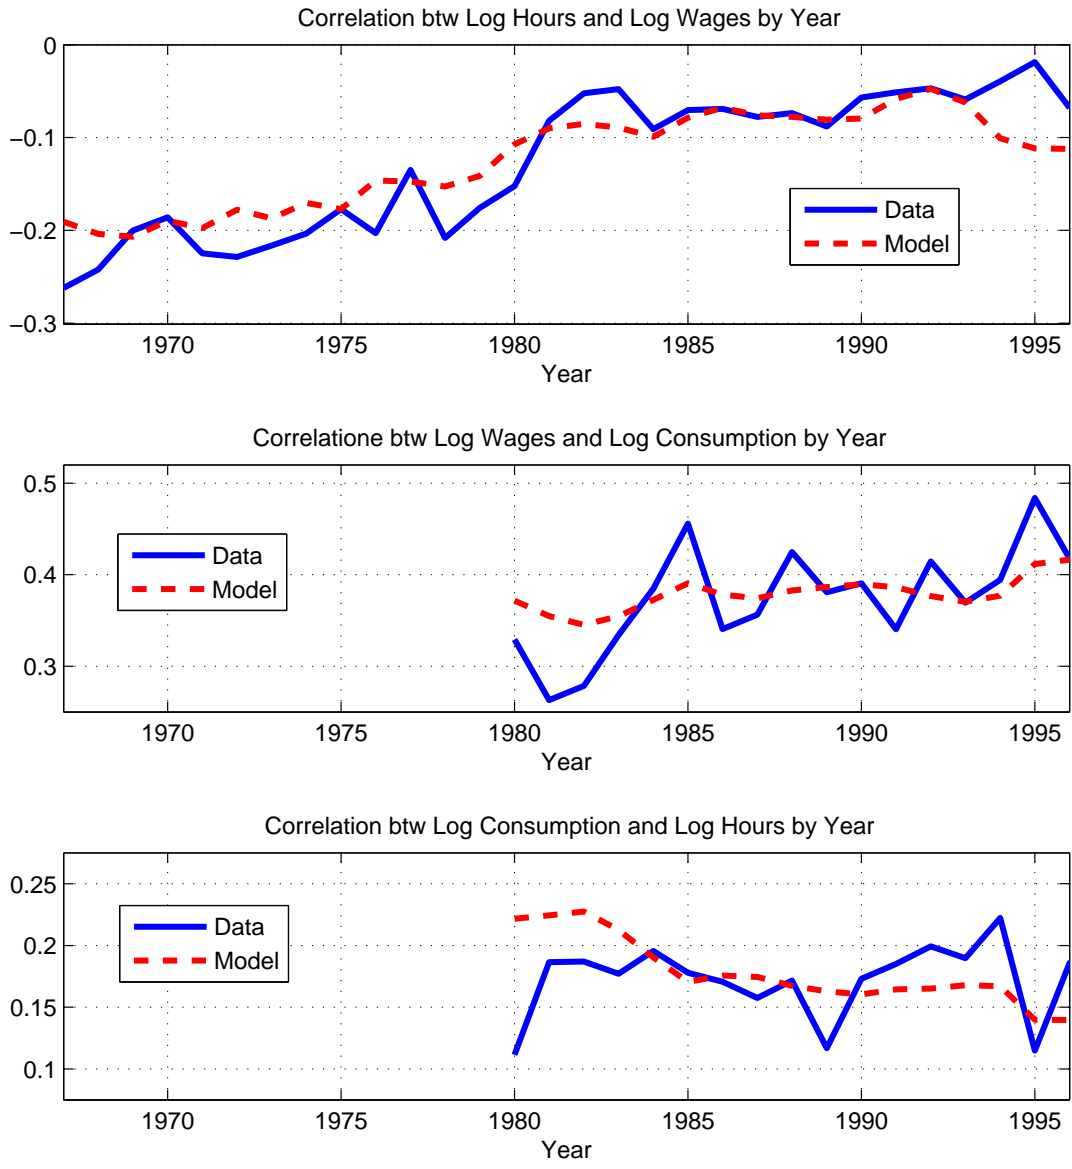


Figure 8: Time series dimension – Model's fit

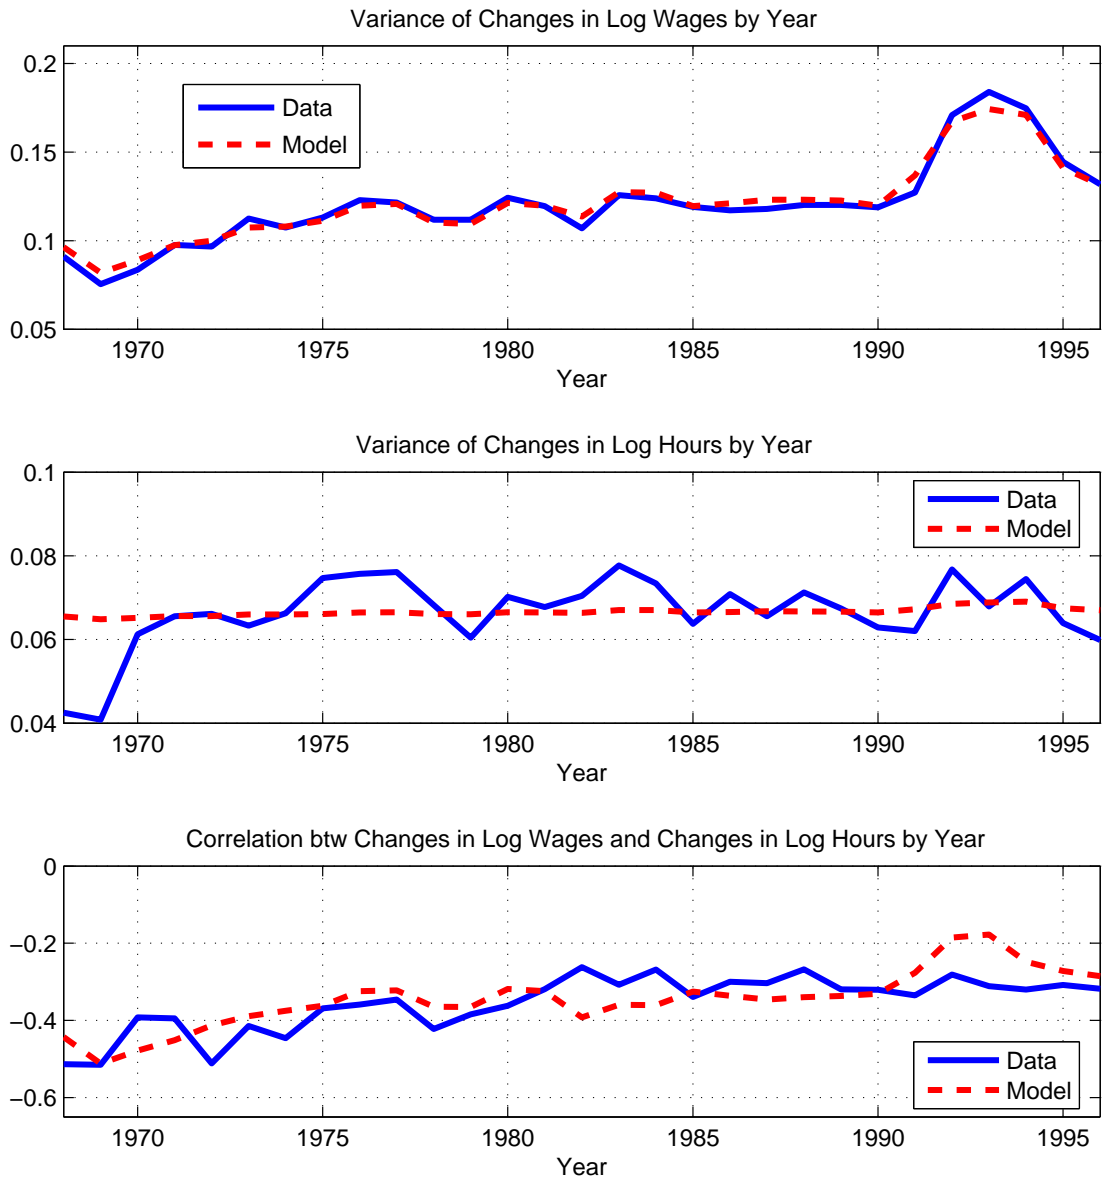


Figure 9: Time series dimension – Model's fit


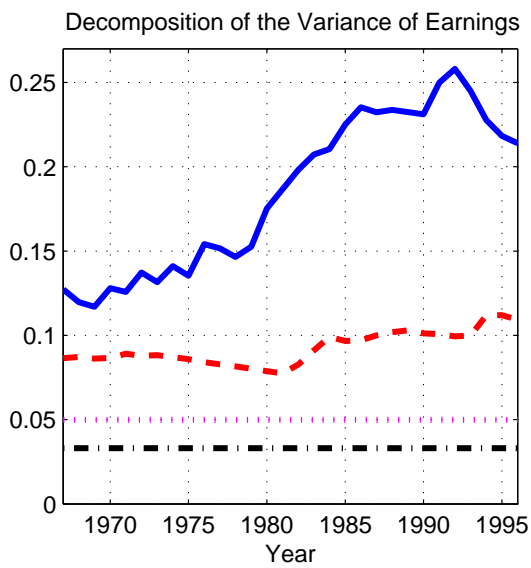
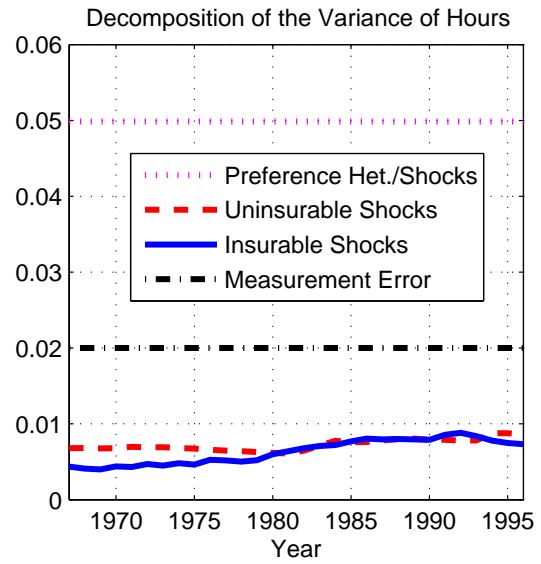
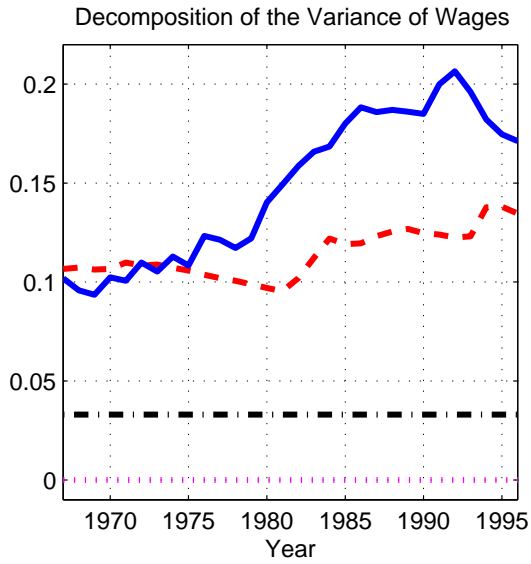


Figure 10: Time series dimension – Decomposition





Figure 11: Time series dimension – Decomposition







Figure 12: Robustness analysis – Fit without preference heterogeneity