

NBER WORKING PAPER SERIES

CONSUMPTION AND LABOR SUPPLY WITH PARTIAL INSURANCE:
AN ANALYTICAL FRAMEWORK

Jonathan Heathcote
Kjetil Storesletten
Giovanni L. Violante

Working Paper 15257
<http://www.nber.org/papers/w15257>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2009

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). Storesletten thanks the Norwegian Research Council for support for this research. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Consumption and Labor Supply with Partial Insurance: An Analytical Framework
Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante
NBER Working Paper No. 15257
August 2009
JEL No. E21,J22,J31

ABSTRACT

This paper studies consumption and labor supply in a model where agents have partial insurance and face risk and initial heterogeneity in wages and preferences. Equilibrium allocations and variances and covariances of wages, hours and consumption are solved for analytically. We prove that all parameters of the structural model are identified given panel data on wages and hours, and cross-sectional data on consumption. The model is estimated on US data. Second moments involving hours and consumption show that the rise in wage dispersion in the 1970s was effectively insured by households, while the rise in the 1980s was not.

Jonathan Heathcote
Federal Reserve Bank of Minneapolis
Research Department
90 Hennepin Ave.
Minneapolis, MN 55401
heathcote@minneapolisfed.org

Giovanni L. Violante
Department of Economics
New York University
19 W. 4th Street
New York, NY 10012-1119
and NBER
glv2@nyu.edu

Kjetil Storesletten
Federal Reserve Bank of Minneapolis
Research Department
90 Hennepin Ave.
Minneapolis, MN 55401
kjetil.storesletten@econ.uio.no

1 Introduction

How large is idiosyncratic risk for individual workers, and how does it translate into cross-sectional dispersion in hours worked, consumption, and welfare? To answer these questions, we develop a new theoretical framework within which it is possible to follow the microeconomic lives of individual workers, and at the same time to understand how these lives aggregate to generate cross-sectional dispersion in the macroeconomy. In our model, agents face two fundamental sources of risk: shocks to wages and to preferences. They choose hours worked and risk-free saving. In addition, they have access to insurance against some but not all shocks. Prices and quantities reflect individual optimization and market clearing.

The first result of the paper is that equilibrium allocations for consumption, earnings and hours can be characterized analytically: they are log-linear in three latent random variables — an idiosyncratic preference parameter, and two components of idiosyncratic labor productivity, only one of which is insured in equilibrium. Given allocations in closed form, first and second cross-sectional moments of the equilibrium joint distribution over wages, hours worked, and consumption can be expressed as simple functions of structural model parameters. These moments serve as a powerful lens through which to interpret the equilibrium mapping between preference parameters and the insurability of idiosyncratic risk on the one hand, and the evolution of cross-sectional inequality on the other.

How do we maintain analytical tractability? In the prototypical bond economy with constant relative risk aversion utility and mean reverting income dynamics (Huggett, 1993; Aiyagari, 1994), agents smooth consumption by adjusting saving inversely with shocks to income. Aggregate dynamics typically depend on the entire equilibrium distribution of wealth, a high dimensional endogenous state variable, which renders the models intractable. When shocks are fully permanent, however, there is an equilibrium in which bonds are not traded (Constantinides and Duffie, 1996). Their result is the starting point for our framework, but we do not want to buy tractability if the price is an economy in which labor market shocks translate one-for-one into consumption. On the contrary, there is evidence of substantial (but not perfect) risk-sharing among households (Attanasio and Davis, 1996; Blundell, Pistaferri, and Preston, 2008).

We therefore extend the Constantinides-Duffie environment in two ways. First, we focus on wage risk and allow agents to choose labor supply, thereby endogenizing earnings and introducing a margin of response to wage shocks through hours worked. Second, we allow agents to perfectly insure a subset of idiosyncratic risks – through trade in bonds,

more sophisticated financial assets, or non-market mechanisms. Intuitively, a model with full insurance against some risks and no insurance against others is tractable for the same reason that models with complete markets or autarky are tractable: in the first case, equilibrium allocations can be characterized without reference to the (nondegenerate) wealth distribution by exploiting the welfare theorems; in the second case, the wealth distribution is degenerate. Because the degree of risk-sharing in our framework lies in between a bond economy and complete markets, the model features “partial insurance.”¹

While labor supply and risk-free savings are perhaps the most obvious margins of adjustment to idiosyncratic shocks, we allow for additional insurance to capture myriad other adjustment mechanisms and institutions that spread risks across individuals or over time. Examples include sophisticated financial instruments, government taxation and transfer schemes, and risk-sharing within a range of networks including families, firms and unions. Fully structural incomplete-markets models have been developed that explicitly incorporate a variety of such insurance mechanisms.² However, simultaneously incorporating multiple adjustment mechanisms poses obvious numerical challenges. Moreover, if certain channels are not explicitly modeled, one risks exaggerating or underestimating the quantitative importance of the mechanisms that are included, depending on whether they are substitutes for or complements to the mechanisms that are missing.

Deaton (1997) suggests an alternative approach: “Although it is possible to examine the mechanisms, the insurance contracts, tithes and transfers, their multiplicity makes it attractive to look directly at the magnitude that is supposed to be smoothed, namely consumption.”³ Our framework combines basic self-insurance in the Bewley tradition (through borrowing/lending and labor supply) with an assumption that insurance—in some form—exists against a fraction of permanent wage risk. In the spirit of Deaton’s

¹Recent examples of models displaying partial insurance include Krueger and Perri (2006), Attanasio and Pavoni (2007), Blundell et al., (2008), and Ales and Maziero (2009).

²Examples are investigations of self-insurance and individual and family labor supply (Heathcote, Storesletten and Violante, 2008a), bankruptcy laws (Livshits, MacGee and Tertilt, 2007), consumer durables (Fernandez-Villaverde and Krueger, 2004), and government redistribution to the poor (Hubbard, Skinner and Zeldes, 1995; Low, Meghir and Pistaferri, 2006). See Heathcote, Storesletten and Violante (2009) for a recent survey of this literature.

³Deaton (1997, pages 372–374) 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, the schemes are more likely to exist among groups of people among whom information is good, where income is 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 the possibilities for autarkic consumption smoothing through intertemporal transfers of money or goods.”

ambition, we then quantify the overall extent of risk-sharing, while remaining somewhat agnostic about the specific sources. One indication that a large fraction of permanent innovations to wages is effectively insured is that the increase in dispersion over the life cycle is much larger for earnings than for consumption (Heathcote et al. 2005).

A second feature of our model that is new to the heterogeneous-agents incomplete-markets literature is that we allow for initial dispersion and life cycle shocks to the relative taste for consumption versus leisure. It is important to incorporate preference heterogeneity because there is strong evidence that some part of cross-sectional dispersion in consumption and hours is unrelated to dispersion in wages. For example, there is substantial variation in hours worked at fixed wage rates (e.g., Abowd and Card, 1989). Moreover, over the past forty years there have been substantial changes in the distribution of household composition and in social norms and attitudes regarding gender roles in the workplace. Some of the effects of these changes on inequality in hours and consumption are better modeled as reflecting changes in the cross-sectional distribution of preferences, rather than the distribution of wages.

The return from tractability is a transparency lacking in standard incomplete-markets models. In particular, we obtain theoretical closed-form expressions for second moments of the joint equilibrium distribution of wages, hours, and consumption. These moments by year and cohort can be used to formally prove that we can identify all structural parameters defining preferences and idiosyncratic risk, even when allowing for classical measurement error and time variation in the risk parameters. Identification requires panel data on wages and hours (as in the Panel Study of Income Dynamics [PSID]) and cross-sectional data on consumption (as in the Consumer Expenditure Survey [CEX]). An important insight is that labor supply contains substantial information about the degree of insurance available to individuals. Thus there is a big payoff from endogenizing labor supply and earnings, rather than focusing exclusively on moments involving consumption to identify the extent of equilibrium insurance, which is the typical approach in the literature (Hall and Mishkin, 1982; Blundell and Preston, 1998; Blundell et al., 2008; Guvenen and Smith, 2008; Primiceri and van Rens, 2009).

Our final contribution is to apply these qualitative insights on the connection between observable second moments and structural parameters by estimating the model using a minimum distance estimator.⁴ It is worth emphasizing that because we solve for the *levels* of consumption, earnings and hours, we can exploit information contained in the “macro

⁴Note that estimation would be enormously challenging in a standard incomplete-markets model given the numerical challenges involved in approximating equilibrium allocations for a single parameterization.

facts” on cross-sectional dispersion in levels that have primarily engaged macroeconomists (e.g. Castañeda, Diaz-Gimenez and Ríos-Rull, 2003; Krueger and Perri, 2006; Heathcote et al., 2008a, Huggett, Ventura and Yaron, 2009) in addition to the “micro facts” on the distribution of individual growth rates which have long been a focus of labor economists (e.g., Abowd and Card, 1989; Blundell et al., 2008). Both micro and macro facts contain valuable information about structural parameters, and we use both to identify and estimate the model.

One key finding from the estimation exercise is that cross-sectional data yield precise estimates for preference parameters that are consistent with recent estimates derived from very different methodologies. Our estimate for the CRRA curvature coefficient on consumption is 2.02, and our baseline estimate for the Frisch elasticity of labor supply is 0.38. A second important finding is that labor supply data tell the same story as consumption data regarding changes over time in the fraction of idiosyncratic productivity risk that is essentially insurable: both sets of moments indicate that the rise in wage dispersion until the early 1980s was largely insurable, while the subsequent rise in wage dispersion reflects larger uninsurable shocks. Over the entire period 1967–1996, around two-thirds of the rise in idiosyncratic wage risk was insurable.

Estimation is the natural way to quantify the extent of latent dispersion in preferences. Preference dispersion is an independent source of variation in hours worked, and those who work relatively long hours enjoy relatively high earnings and consumption. We find that preference dispersion accounts for around a third of cross-sectional dispersion in consumption and hours worked and also helps explain the fact that the empirical correlation between consumption and hours is strongly positive.

The rest of the paper is organized as follows. Section 2 develops our framework, derives the equilibrium allocations, and explains how we preserve tractability. 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, and discusses and interprets the estimation results. Section 6 concludes the paper and discusses a range of further applications and extensions that maintain tractability.

2 Model economy

We first describe the model formally. Next, we discuss in detail the key assumptions.

Demographics We adopt the Yaari perpetual youth model: agents are born 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 at each date t . Thus, the measure of agents of age a is $(1 - \delta)\delta^a$, and the total population size is unity.

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

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

where the expectation is taken over sequences of shocks defined below. 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. Period utility is

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

The parameter γ is the inverse of the intertemporal elasticity of substitution for consumption. The Frisch elasticity of labor supply is equal to $1/\sigma$. The variable φ_t is a time-varying preference weight that captures the strength of an individual's aversion to work relative to his preference for consumption.⁵ We assume that φ_t follows the unit root process

$$\varphi_t = \varphi_{t-1} + \chi_t,$$

where the innovation χ_t is drawn from the distribution $F_{\chi t}$ with variance $v_{\chi t}$ at time t . Agents entering the labor market at age $a = 0$ in year t draw an initial realization φ_t^0 from a distribution $F_{\varphi^0 t}$ with cohort-specific variance $v_{\varphi^0 t}$.

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 α_t component follows 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 second component ε_t is itself the sum of two orthogonal random variables:

$$\varepsilon_t = \kappa_t + \theta_t.$$

⁵In the utility function φ is multiplied by preference parameters. This is an innocuous normalization which simplifies the expressions for equilibrium allocations.

Here θ_t is a transitory (independently distributed over time) shock drawn from $F_{\theta t}$ with variance $v_{\theta t}$, 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 $v_{\eta t}$.

Agents entering the labor market at age $a = 0$ in year t draw initial realizations α_t^0 and κ_t^0 from distributions $F_{\alpha^0 t}$ and $F_{\kappa^0 t}$ with cohort-specific variances $v_{\alpha^0 t}$ and $v_{\kappa^0 t}$. The initial draws φ_t^0 , α_t^0 and κ_t^0 are assumed to be independent.⁶

Production Production of the final consumption good takes place through a constant returns to scale technology with labor as the only input. The economy-wide goods and labor markets are perfectly competitive. By assumption, individual wages equal individual productivities (units of effective labor per hour worked).

Information Agents are assumed to take as given the sequences of distributions $\{F_{\varphi^0 t}, F_{\alpha^0 t}, F_{\kappa^0 t}, F_{\chi t}, F_{\omega t}, F_{\eta t}, F_{\theta t}\}$. Thus they have perfect foresight over future wage dynamics.⁷

Island partitioning The population in the economy is partitioned across islands, or groups. Each period, a continuum of new islands is formed, with each island comprising a continuum of newborn agents who share common sequences $\{\varphi_t, \alpha_t\}_{t=b}^{\infty}$.⁸ Within an island, agents are heterogeneous with respect to the sequences $\{\kappa_t, \theta_t\}_{t=b}^{\infty}$. Note that there is wage dispersion within islands due to the idiosyncratic component (ε_t), and wage dispersion across islands due to the island-level component (α_t). In contrast, there is no preference heterogeneity within islands.⁹

Thus, the economy has a continuum of islands, each inhabited by a continuum of agents whose mass shrinks at rate δ as time passes. A law of large numbers (e.g., Uhlig, 1996) can be applied twice so that idiosyncratic individual-level shocks “wash out” within

⁶The initial draws ($\varphi^0, \alpha^0, \kappa^0$) could in principle be correlated if, for example, wages at labor market entry depend on schooling and schooling depends on the initial preference parameter, φ^0 . In a previous version of the paper we allowed for correlation between α^0 and φ^0 , but the estimated correlation coefficient turned out to be insignificantly different from zero.

⁷This assumption is not required for tractability. Alternatively, one could assume that the variances of these distributions themselves follow some stochastic process.

⁸We segregate individuals across islands by birth year t and fixed effects (φ^0, α^0) because this is convenient when it comes to solving for equilibrium allocations: as shown later, when individuals are segregated this way, within-island allocations can be determined using equal-weighted island-level planning problems. We could alternatively envision groups composed by individuals with different birth years and fixed effects, as long as all individuals in a group experience the same sequence $\{\omega_t, \chi_t\}$. Such an approach would yield exactly the same allocations for consumption and hours.

⁹In Section 5.3 we discuss the implications of introducing within-group preference dispersion.

an island, and island-level shocks induce no aggregate uncertainty in the economy as a whole (see Attanasio and Ríos-Rull, 2000, for a similar economy structure).

Market structure At birth, each agent is endowed with zero financial wealth.¹⁰ All assets in the economy are in zero net supply. Within each island, there are complete insurance markets. In particular, any agent can purchase insurance contracts at date t that deliver consumption at $t + 1$ contingent on the realization of their individual shocks $s_{t+1} \equiv (\omega_{t+1}, \chi_{t+1}, \eta_{t+1}, \theta_{t+1})$. Profit-maximizing insurance providers operate in markets that are competitive, and segmented by island. As a result, in equilibrium agents can pool the idiosyncratic shocks $(\eta_{t+1}, \theta_{t+1})$ but not the island-level shocks $(\omega_{t+1}, \chi_{t+1})$. Asset trading across islands is exogenously restricted to a risk-free non-contingent bond.

2.1 Discussion

With respect to our specification for preferences, some authors have invoked nonseparability between consumption and leisure to explain certain features of individual life cycle data (for a survey, see Browning et al., 1999). In Heathcote, Storesletten and Violante (2008b) we show that it is possible to solve for the equilibrium of the model in closed form when consumption and leisure are aggregated in a Cobb-Douglas fashion. Here we adopt the separable specification because it affords valuable flexibility in distinguishing between agents' willingness to substitute consumption and hours intertemporally.

Preference dispersion in the model is intended to capture various sources of heterogeneity and risk which generate cross-sectional variation in hours worked and consumption that is independent of variation in wages. In addition to actual differences in the relative weighting of consumption and leisure across agents, our specification also encompasses, in reduced form, all those factors that create individual-specific wedges in the intratemporal first-order condition. Examples include changes in household composition (and thus economies of scale) affecting the marginal utility of consumption, shocks to the time endowment available for work (such as disability), and variation in effective marginal tax rates on labor income. In the Conclusions, we explain how one could explicitly model both demographic changes and progressive taxation, while retaining tractability.

The evolution of individual efficiency units/wages (w_t) is taken as exogenous by agents. However, earnings are endogenous because labor supply is a choice variable, so individuals have some control over labor income. The statistical process for wages described above (unit root plus i.i.d. shocks) is quite standard in the literature, and is consistent with the

¹⁰It is straightforward to relax the assumption of zero initial individual financial wealth. The key requirement, as will become clear below, is that *average* initial wealth on each island is zero.

key features of individual wage dynamics as well as with trends in wage dispersion across the life cycle. 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, within-cohort wage dispersion increases approximately linearly with age, suggesting the presence of permanent shocks.

Recently, Guvenen (2007) and Primiceri and van Rens (2009) have revived a long-standing debate on the extent to which individual wages changes are predictable. We have labeled all individual wage changes in our environment as shocks. However, as will become clear, in the context of our model it is difficult to distinguish between insurable shocks to wages, and changes in wages that are foreseen but not explicitly insurable.

Our model allows for two potential exogenous drivers of the evolution of the joint equilibrium distribution of consumption and labor supply across households. The first is the evolution of the wage distribution, driven by a mix of cohort effects and time effects. New cohorts enter the economy with cohort-specific dispersion in initial labor productivity, which generates changes over time in the cross-sectional variance of wages for the economy as a whole. At the same time, all agents are subject to life cycle productivity shocks, whose cohort-independent variance changes over time. The second driver of changes in inequality is changes in the distribution of preference heterogeneity. Symmetrically with wages, cross-sectional dispersion in preferences reflects a mix of cohort effects and time effects. When taken to the data, the model is free to assign a large role to wage factors, non wage factors, or both.

The market structure of our economy embeds exogenous restrictions on the set of assets that can be traded, in keeping with the Bewley-Imrohroglu-Huggett-Aiyagari tradition. However, instead of restricting the set of assets traded *among individuals*, we instead restrict the set of assets that can be traded *between islands*, while allowing for perfect risk-sharing within groups. Our economy can achieve any degree of insurance between a bond economy and a complete markets economy. For $v_{\omega t} = v_{\chi t} = 0$ the market structure is complete (all islands are alike), while for $v_{\eta t} = v_{\theta t} = 0$ it becomes a standard bond economy (with a representative agent on each island). In general, it is an economy offering *partial insurance* against shocks.¹¹

We will show that there is sufficient information in the aggregate cross-sectional joint distribution of wages, hours, and consumption to identify the fraction of productivity

¹¹Microfounding the asset trading friction is beyond the scope of this paper, where we focus on the empirical content of our environment. One could pursue a microfoundation based on asymmetric information about shocks (e.g., better observability of shocks within a group than outside the group) along the lines of Cole and Kocherlakota (2001), and Ales and Maziero (2009).

variation that is effectively insured, without exploiting any of the model's predictions for within-island or between-island inequality. As Deaton (1997) emphasizes, examining precisely who shares risks with whom and how is a major empirical challenge. In this light, the fact that we can quantify the scope of overall insurance without having to identify islands in the data constitutes a strength of our framework.

2.2 Agent's problem

Let $s^t = \{s_b, s_{b+1}, \dots, s_t\}$ denote the individual history of the shocks for an agent from birth year b up to date t , where

$$s_\tau = \begin{cases} (\alpha_\tau^0, \varphi_\tau^0, \kappa_\tau^0, \theta_\tau) & \tau = b \\ (\omega_\tau, \chi_\tau, \eta_\tau, \theta_\tau) & \tau > b \end{cases}$$

with $s_\tau \in S = \mathbb{R}^4$, and $s^t \in S^t$.

Let $b_t^a(s^t)$ and q_t denote the quantity and price of riskless bonds, purchased by an agent of age a with history s^t , that pay one unit of consumption at $t + 1$. Let $B_t^a(A; s^t)$ and $Q_t(A; \omega^t, \chi^t)$ denote the quantity and the price of insurance claims purchased that pay one unit of consumption if and only $s_{t+1} \in A \subseteq S$. Note that since riskless bonds can be traded freely across the whole economy, their price does not depend on the individual state s^t . Since insurance claims are only sold within islands, their price could in principle depend on the island-specific history (ω^t, χ^t) .

An agent's sequential budget constraint is then

$$c_t^a(s^t) + \int_S Q_t(s; \omega^t, \chi^t) B_t^a(s; s^t) ds + q_t b_t^a(s^t) = w_t(s^t) h_t^a(s^t) + d_t^a(s^t) \delta^{-1}, \quad (4)$$

where realized wealth at node $s^t = (s^{t-1}, s_t)$ is given by

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

For what follows, it is also useful to define "net" financial wealth for an individual with history s^t 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 current earnings.

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, non-negativity constraints on consumption and hours worked, and have zero initial financial wealth.

2.3 Competitive equilibrium

Given sequences $\{F_{\varphi^0 t}, F_{\alpha^0 t}, F_{\kappa^0 t}; F_{\chi t}, F_{\omega t}, F_{\eta t}, F_{\theta t}\}$, a sequential competitive equilibrium is a set of allocations $\{c_t^a(s^t), h_t^a(s^t), d_t^a(s^t), b_t^a(s^t), B_t^a(A; s^t)\}$ and prices $\{q_t, Q_t(A; \omega^t, \chi^t)\}$ for

all dates t , all histories $s^t \in S^t$, and all $A \subseteq S$ 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 the non-contingent bond clear.

For what follows, it is convenient to define a set of constants indexed by time and age: $\mathcal{M}_t^a \equiv M_{\theta t} + M_{\kappa t}^a$, where $M_{\theta t} = \sigma / (\sigma + \gamma) \log \left(\int \exp \left(\frac{1+\sigma}{\sigma} \theta_t \right) dF_{\theta t} \right)$, and $M_{\kappa t}^a = \sigma / (\sigma + \gamma) \log \left(\int \exp \left(\frac{1+\sigma}{\sigma} \kappa_t \right) dF_{\kappa t}^a \right)$. Here, $F_{\kappa t}^a$ denotes the distribution of the cumulated value of κ across agents of age a at date t (i.e., for the cohort born at $b = t - a$).

We are now ready to state the first theoretical result of the paper, the characterization of equilibrium allocations for consumption, hours, and wealth.

PROPOSITION 1 [COMPETITIVE EQUILIBRIUM] *There exists a sequential competitive equilibrium characterized as follows:*

(i) *Consumption and hours are given by*

$$\log c_t^a(s^t) = -\varphi_t + \frac{1+\sigma}{\sigma+\gamma} \alpha_t + \mathcal{M}_t^a \quad (5)$$

$$\log h_t^a(s^t) = -\varphi_t + \frac{1-\gamma}{\sigma+\gamma} \alpha_t + \frac{1}{\sigma} (\kappa_t + \theta_t) - \frac{\gamma}{\sigma} \mathcal{M}_t^a, \quad (6)$$

and net financial wealth is given by

$$\hat{d}_t^a(s^t) = c_t^a(s^t) \left[1 - \frac{\exp \left(\frac{1+\sigma}{\sigma} \kappa_t \right)}{\int \exp \left(\frac{1+\sigma}{\sigma} \kappa \right) dF_{\kappa t}^a} \right] \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j \left(\frac{c_{t+j}^{a+j}(s_{t+j})}{c_t^a(s_t)} \right)^{1-\gamma}. \quad (7)$$

(ii) *The prices of the bond and the insurance claims are given by*

$$q_t = \beta \exp(-\gamma \Delta \mathcal{M}_{t+1}) \int_S \exp \left(\gamma \chi_{t+1} - \gamma \frac{1+\sigma}{\sigma+\gamma} \omega_{t+1} \right) dF_{s,t+1} \quad (8)$$

$$Q_t(A; \omega^t, \chi^t) = \beta \exp(-\gamma \Delta \mathcal{M}_{t+1}) \int_A \exp \left(\gamma \chi_{t+1} - \gamma \frac{1+\sigma}{\sigma+\gamma} \omega_{t+1} \right) dF_{s,t+1} \quad (9)$$

where $F_{s,t+1}$ is the joint distribution at $t+1$ over $(\omega, \chi, \eta, \theta)$ and $\Delta \mathcal{M}_{t+1} = \mathcal{M}_{t+1}^{a+1} - \mathcal{M}_t^a = M_{\eta,t+1} + \Delta M_{\theta,t+1}$.

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

We now sketch the strategy for proving Proposition 1 (see the Appendix for details). We first assume that there is no bond trade between islands and, hence, zero net savings on each island. Given complete markets within an island, we then apply the first welfare theorem and show that the allocations that solve a static planner's problem (with equal planner weights corresponding to the assumption of equal initial wealth) take the form

(5) and (6). We then show that these allocations can be supported in a decentralized equilibrium by prices of the form (8) and (9), and verify the no-bond-trade conjecture.¹² We discuss the intuition for this no-trade result in more detail in the next section.

As is clear from Proposition 1, no distributional assumptions are needed to derive closed-form expressions for the cross-sectional second moments of the joint distribution over wages, hours, and consumption. The absence of an explicit solution for terms involving \mathcal{M}_t^a is no obstacle for the empirical analysis, since they can be modeled through age and time dummies in individual consumption and hours observations. However, a distributional assumption allows us to compute these terms in closed form, further illuminating the nature of equilibrium allocations and prices.

COROLLARY [NORMALITY OF SHOCKS] *Suppose that $(\alpha_t^0, \kappa_t^0, \varphi_t^0, \omega_t, \eta_t, \theta_t, \chi_t)$ are independently, normally distributed, with variance v_x and mean $-v_x/2$ for each random variable x . Then, the term \mathcal{M}_t^a is given by $\mathcal{M}_t^a = \left(\frac{1+\sigma}{\sigma+\gamma}\right) \frac{\text{var}_t^a(\varepsilon)}{2\sigma}$, where $\text{var}_t^a(\varepsilon) = v_{\kappa^0, t-a} + \sum_{j=0}^{a-1} v_{\eta, t-j} + v_{\theta t}$ is the variance of the insurable component of the wage within islands of age a in year t . The risk-free bond price is*

$$q_t = \beta \exp \left(\gamma \frac{1+\sigma}{\sigma+\gamma} \left(\left(\gamma \frac{1+\sigma}{\sigma+\gamma} + 1 \right) \frac{v_{\omega, t+1}}{2} - \frac{v_{\eta, t+1} + \Delta v_{\theta, t+1}}{2\sigma} \right) + \gamma(\gamma+1) \frac{v_{\chi, t+1}}{2} \right). \quad (10)$$

This entirely closed-form expression for the equilibrium bond price is useful for understanding why our framework is tractable.

Tractability In general, incomplete-market economies do not admit an analytical solution, and numerical methods are required to solve for equilibrium allocations. In our environment, the vector of cumulated values for the shocks $(\alpha_t, \varphi_t, \kappa_t, \theta_t)$, together with the age a of the individual, contains sufficient information to fully describe equilibrium choices. The power of this result lies in the fact that these are all exogenous states. Crucially, individual wealth is not an independent state variable, in that it can be expressed as a function of $(\alpha_t, \varphi_t, \kappa_t, \theta_t)$ and a .

The reason individual wealth can be excluded from the individual state space is twofold. First, within a particular island, markets are effectively complete. Thus within-island allocations can be computed as the solution to the problem of an island-planner who maximizes island-level welfare subject to an island-level resource constraint. Because agents are endowed with zero initial wealth, the corresponding planner's problem accords all agents equal planner weights. But given these weights, the wealth distribution plays no further role in the planner's welfare maximization problem.

¹²Note that $Q_t(A; \omega^t, \chi^t) = Q_t(A)$. Thus the price of insurance against individual risk is independent of the island-specific history (ω^t, χ^t) . Note also that the equilibrium price of insurance against events $s_{t+1} \in A$ reflects both probabilities and state-contingent marginal rates of substitution.

Second, the inter-island wealth distribution does not show up in allocations because, in equilibrium, this distribution remains degenerate at zero. The result that there is no cross-island heterogeneity in desired saving rates is specific to our environment, and can be explained as follows.

First, recall that our individuals have three saving motives: an intertemporal motive given by the interaction of the degree of patience β and the interest rate, a smoothing motive linked to expected earnings growth, and a precautionary motive reflecting the variance of island-level shocks ω and χ . A key feature of our model is that each of these saving motives applies with equal strength across all islands. The intertemporal motive is the same across islands, because all agents share common discount factor β and face the same economy-wide risk-free interest rate q_t^{-1} . The precautionary motive is also identical, reflecting the assumptions that island-level shocks are multiplicative, permanent and drawn from common distributions $F_{\eta t}$, $F_{\omega t}$, and $F_{\chi t}$, that preferences are in the power utility class with respect to both consumption and hours worked, and that all islands start out with zero average wealth. These same features imply a common rate of island-level earnings growth associated with rising dispersion in the insurable component of wages, which in turn ensures that all islands have the same smoothing motive.¹³

Because the strength of these three saving motives is identical across all islands, there exists an economy-wide interest rate q_t^{-1} at which, in equilibrium, the (negative) intertemporal motive and (negative) smoothing motive exactly offset the (positive) precautionary motive, and no agent wants to either borrow or lend across islands. The closed-form expression for the interest rate derived in the corollary helps us visualize this argument.

Let us begin by taking the limit as $\sigma \rightarrow \infty$ (inelastic labor supply). In the absence of island-level risk ($v_{\omega,t+1} = v_{\chi,t+1} = 0$), the equilibrium interest rate would satisfy the standard condition for a complete-markets endowment economy, $q_t = \beta$. In the presence of island-level risk, expression (10) simplifies to $(\rho - r_{t+1})/\gamma = (1 + \gamma)(v_{\omega,t+1} + v_{\chi,t+1})/2$, where $\rho = -\log(\beta)$ and $r_{t+1} = -\log(q_t)$. The right-hand side captures the precautionary motive for saving, which is proportional to the variance of the island-level productivity and preference shocks with the constant of proportionality defined by the coefficient of relative prudence ($-u'''(c)c/u''(c) = 1 + \gamma$). The left-hand side measures the intertemporal motive to dis-save, which is proportional to the difference between the rate of time preference and the equilibrium interest rate, with the constant of proportionality defined by the elasticity of intertemporal substitution ($1/\gamma$). Here, it is clear that the equilibrium interest rate is

¹³At the individual-level, agents with low draws for the transitory component θ have relatively high expected earnings growth. However, innovations to the idiosyncratic wage component ε are perfectly insured within the island, so they play no role in determining the island-level demand for risk-free bonds.

such that the two saving motives exactly offset each other.

With elastic labor supply, the expression for the interest rate is somewhat more involved, but the interpretation remains similar. An extra term appears involving life cycle growth in the variance of the insurable component of wages, $v_{\eta,t+1} + \Delta v_{\theta,t+1}$. This reflects a motive to borrow given that cumulating insurable productivity shocks imply expected growth in earnings and consumption. In addition, allowing for elastic labor supply changes the constant of proportionality defining the strength of the precautionary motive for saving in the face of uninsurable productivity risk. In particular, if $\gamma > 1$ then labor supply is adjusted inversely to shocks to wages in order to smooth earnings. In this case, the higher is the Frisch elasticity of labor supply, the weaker is the precautionary saving motive, since labor supply provides a hedge against risk (Bodie, Merton and Samuelson, 1992, develop similar intuition).

Relation to the literature Our no-inter-island-trade result represents a generalization of Constantinides and 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 important, because in Constantinides-Duffie the equilibrium entails no risk-sharing at all, whereas our environment allows for insurance within islands/groups. Second, our agents supply labor elastically, so that the process for earnings is endogenous. This is important because, as will shortly become clear, data on hours worked are a rich source of information on the nature of risk and risk-sharing, and endogenizing hours worked allows us to exploit this information alongside information on consumption choices. The third key extension relative to Constantinides-Duffie is that our agents are subject to idiosyncratic preference shocks, in addition to productivity shocks. This is important because we do not want to impose a priori that all cross-sectional dispersion in consumption and hours worked is driven by dispersion in wages.

To our knowledge, it is possible to solve for the equilibrium of incomplete-market economies analytically only in a few other cases. Caballero (1990) shows that with constant absolute risk aversion utility, the precautionary saving motive in a bond economy is constant across all agents, independently of the earnings process. Wang (2003) builds on this result and proves that one can construct an equilibrium in which the risk-free interest rate is low enough that the negative intertemporal saving motive offsets the positive precautionary motive, and all saving is driven by the residual consumption-smoothing motive. Thus equilibrium consumption ends up following the permanent income hypothesis.

There is also a class of tractable models which consider idiosyncratic risk to human or

physical capital. Examples include Benabou (2002), Krebs (2003), and Angeletos (2007).

Consumption and hours worked Because markets are complete within islands, while there is no trade between islands, individual-specific shocks are fully insurable, while island-level shocks are not insurable at all. Thus we will refer to (η_t, θ_t) as “insurable shocks” and to (ω_t, χ_t) as “uninsurable shocks”. This stark dichotomy is the key to preserving tractability, while delivering an environment in which the extent of overall insurance is somewhere in between autarky and complete markets. We now interpret the equilibrium allocations for hours worked and consumption described in (5) and (6).

Hours worked are increasing in the insurable component ε_t , and the response of hours to insurable shocks (η_t, θ_t) is defined by the Frisch elasticity, $1/\sigma$. This reflects the fact that full insurance with respect to ε_t rules out any income effect on hours worked. Uninsurable permanent shocks to α_t do have an income effect which is regulated by γ . If $\gamma > 1$, the income effect dominates the substitution effect and hours worked decline in response to an increase in α_t . If $\gamma < 1$, the relative size of the two forces is reversed, and hours worked increase. The impact of the preference parameter φ_t on hours and consumption is readily interpreted: a stronger relative distaste for work (higher φ_t) reduces labor supply.

Individual consumption is independent of ε_t , since these shocks can be fully insured in equilibrium. Consumption is equal to the uninsurable component of earnings (wages times hours), which is increasing in the uninsurable component of the wage, α_t , and decreasing in the preference parameter φ_t : stronger distaste for work induces a reduction of hours worked which transmits to earnings and consumption.

By assuming normally distributed shocks, we can further refine the expressions for equilibrium allocations. Substituting the explicit solutions for the dummy variables \mathcal{M}_t^a in the Corollary into eqs. (5) and (6), it is immediate that consumption is increasing in the variance of insurable shocks to labor productivity, while hours are decreasing. The logic for this result is that with endogenous labor supply, greater insurable wage dispersion increases average labor productivity per hour worked, an important consideration for welfare analyses.¹⁴ In Heathcote et al. (2008b), we discuss the welfare implications of increased idiosyncratic uncertainty in detail.

Note that the expression for individual consumption is not what the permanent income hypothesis would imply. Consumption is still a random walk, but some permanent shocks (innovations η_t) are fully insurable, and thus do not affect consumption. In other words, our consumption allocation exhibits “excess smoothness” It is precisely this feature of the

¹⁴By construction, the average wage per worker is independent of the variance of shocks to labor productivity.

data that has motivated a large amount of recent research aimed at developing “partial insurance” models that lie in between the bond economy and complete markets (e.g., Krueger and Perri, 2006; Attanasio and Pavoni, 2007).

Insurability versus predictability Two important remarks are in order with respect to insurability. First, our characterization of equilibrium assumes that all within-island risk-sharing arises from explicit markets and state-contingent financial income flows, while in reality agents have access to a large range of different smoothing mechanisms, as we discussed in the Introduction. One could support the same allocations for consumption and hours through a variety of market and non-market insurance mechanisms (e.g., government redistribution, within-family insurance) as long as these mechanisms deliver within-island interpersonal transfers exactly equal, state by state, to the difference between individual earnings and individual consumption.

Second, as is often emphasized in the literature (e.g., Cunha, Heckman and Navarro, 2005; Guvenen and Smith, 2008), it is very challenging, in general, to distinguish between insurability and predictability. This is true in our model as well. Consider the following variation on the environment described above. Suppose that there are no insurance instruments explicitly indexed to η_t or θ_t on the island, just a non-state-contingent bond with a loose borrowing limit. Suppose also that the statistical representation for the wage process is unchanged, but that from the perspective of individual agents (not the econometrician) all future innovations (η_t, θ_t) are perfectly foreseen. In this alternative environment, agents will use the bond to perfectly smooth consumption in the face of forecastable wage changes η_t and θ_t . The elasticities of consumption and labor supply to each of the shocks $(\omega_t, \chi_t, \eta_t, \theta_t)$ will be identical to those in the baseline model. Rather than imposing additional structure to try to differentiate between insurable shocks and uninsurable but foreseen fluctuations, we will simply use the label “insurable shocks” as a catch-all for what in reality is likely to be a mix of both.¹⁵

Indeterminacy of financial wealth These two considerations have important consequences for the determination of the “true” equilibrium distribution of financial wealth. Whether within-island risk sharing is achieved through market or non-market mechanisms has obvious implications for financial wealth: if all risk-sharing were achieved through non-market mechanisms, the wealth distribution would remain degenerate at zero.

¹⁵Blundell et al. (2008) have suggested that the covariance between consumption growth at t and income growth at $t + n$ may be informative about current news about future income changes. This strategy requires longitudinal household survey data on consumption and income, which are not directly available in the United States.

Moreover, switching from the “insurable” interpretation to the “predictable” interpretation would also change wealth allocations. In both cases, individual wealth following history $s^t = \{s_\tau\}_{\tau=b}^t$ is equal to the difference between the expected present value of lifetime consumption and the expected present value of lifetime earnings. The difference is that under the insurable view, these expectations are conditional on s^t (eq. (7)), while under the predictable view, expectations would be conditional on both s^t and the sequences of known future innovations $\{\eta_\tau, \theta_\tau\}_{\tau=t+1}^\infty$.

While equilibrium wealth depends on the details of how agents smooth wage shocks, the efficient responses of consumption and labor supply do not. This is why we focus on the model’s predictions for the joint dynamics of wages, consumption, and hours worked, rather than its predictions for wealth.

3 Equilibrium cross-sectional moments

With allocations in hand, we can derive closed-form expressions for cross-sectional moments of the joint equilibrium distribution of wages, hours, and consumption. These theoretical moment expressions provide a transparent framework for interpreting the dynamics of their empirical counterparts over time and over the life cycle.

Measurement error The first step toward making the model operational is recognizing that measurement error is pervasive in micro data. We allow consumption, earnings and hours worked to be measured with error. Since we compute hourly wages as earnings divided by hours, measurement error in hourly wages contains errors in both earnings and hours. We assume measurement error μ_t^x for variable x is *classical*, i.e., i.i.d. over time and across agents, has mean zero and variance v_{μ_x} . Let $\log \hat{x}_t = \log x_t + \mu_t^x$ denote the log of the measured (with error) value for x .

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

$$\log \hat{w}_t = \alpha_t + \kappa_t + \theta_t + \mu_t^y - \mu_t^h \quad (11)$$

$$\log \hat{c}_t = -\varphi_t + \frac{1 + \sigma}{\sigma + \gamma} \alpha_t + \mu_t^c \quad (12)$$

$$\log \hat{h}_t = -\varphi_t + \frac{1 - \gamma}{\sigma + \gamma} \alpha_t + \frac{1}{\sigma} (\kappa_t + \theta_t) + \mu_t^h. \quad (13)$$

Let $\Delta \log \hat{x}_t \equiv \log \hat{x}_t - \log \hat{x}_{t-1}$ denote the observed individual growth rate for variable x . From the above expressions, individual changes in measured log allocations can be

expressed as

$$\Delta \log \hat{w}_t = \omega_t + \eta_t + \Delta \theta_t + \Delta \mu_t^y - \Delta \mu_t^h \quad (14)$$

$$\Delta \log \hat{c}_t = -\chi_t + \frac{1 + \sigma}{\sigma + \gamma} \omega_t + \Delta \mu_t^c \quad (15)$$

$$\Delta \log \hat{h}_t = -\chi_t + \frac{1 - \gamma}{\sigma + \gamma} \omega_t + \frac{\eta_t + \Delta \theta_t}{\sigma} + \Delta \mu_t^h. \quad (16)$$

Moments in closed form We now use eqs. (11)–(16) to derive expressions for the unconditional second moments (variances and covariances) of the equilibrium joint distribution of wages, hours, and consumption. These expressions are closed-form functions of structural 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.”

Macro moments Let $var_t^a(\alpha)$ denote the variance of cumulated permanent uninsurable shocks over the life cycle for a cohort of age a at time t , with $var_t^a(\alpha) = v_{\alpha^0, t-a} + \sum_{j=0}^{a-1} v_{\omega, t-j}$. Similarly, let $var_t^a(\varphi)$ denote the variance of the permanent preference shocks. The corollary defines the corresponding expression for the variance of insurable shocks, $var_t^a(\varepsilon)$.

The macro moments involving wages and hours are, respectively,

$$var_t^a(\log \hat{w}) = var_t^a(\alpha) + var_t^a(\varepsilon) + v_{\mu y} + v_{\mu h} \quad (17)$$

$$var_t^a(\log \hat{h}) = var_t^a(\varphi) + \left(\frac{1 - \gamma}{\sigma + \gamma}\right)^2 var_t^a(\alpha) + \frac{1}{\sigma^2} var_t^a(\varepsilon) + v_{\mu h} \quad (18)$$

$$cov_t^a(\log \hat{w}, \log \hat{h}) = \frac{1 - \gamma}{\sigma + \gamma} var_t^a(\alpha) + \frac{1}{\sigma} var_t^a(\varepsilon) - v_{\mu h}. \quad (19)$$

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

The variance of hours has four components. First, the larger is heterogeneity in the taste for leisure, $var_t^a(\varphi)$, the larger is cross-sectional variation 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 (the log-consumption case), uninsurable shocks have no effect on hours. Third, the variance of the insurable shocks increases hours dispersion in proportion to the Frisch elasticity (squared). Finally, measurement error in hours contributes positively to observed dispersion.

The covariance between wages and hours has three components. The effect of uninsurable wage shocks on this covariance depends on the value for the curvature parameter, γ . If $\gamma > 1$, then uninsurable shocks decrease the wage-hours covariance, since strong income

effects induce low (uninsured) wage workers to work longer hours. 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.

Covariances involving earnings can be obtained as linear combinations of the three moments above.

We now turn to the moments involving consumption:

$$var_t^a(\log \hat{c}) = var_t^a(\varphi) + \left(\frac{1+\sigma}{\sigma+\gamma}\right)^2 var_t^a(\alpha) + v_{\mu c} \quad (20)$$

$$cov_t^a(\log \hat{h}, \log \hat{c}) = var_t^a(\varphi) + \frac{(1-\gamma)(1+\sigma)}{(\sigma+\gamma)^2} var_t^a(\alpha) \quad (21)$$

$$cov_t^a(\log \hat{w}, \log \hat{c}) = \frac{1+\sigma}{\sigma+\gamma} var_t^a(\alpha). \quad (22)$$

The variance of consumption is increasing in the variance of uninsurable preference heterogeneity and uninsurable wage shocks, as expected. Note the role of labor supply: if $\gamma > 1$, then a lower σ (higher Frisch) translates into lower consumption dispersion for a given $var_t^a(\alpha)$, because labor supply dampens the impact of uninsurable wage shocks on earnings.

The covariance between hours and consumption is increasing in the degree of preference heterogeneity since individuals with higher values for φ work relatively few hours, and thus earn and consume relatively less. The effect of uninsurable wage risk depends on the value of γ : when $\gamma > 1$, a positive uninsurable shock reduces hours worked but increases consumption.

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. Finally, note that neither of the two covariances involving consumption are affected by insurable wage shocks or by measurement error. This is a useful property for proving identification.

Dispersion in the entire cross section Since we have filtered out differences in mean values for allocations across age groups, the expressions for dispersion in the entire cross section are identical to those above, but without the age a superscripts.¹⁶ Thus, for example,

$$var_t(\log \hat{w}) = var_t(\alpha) + var_t(\varepsilon) + v_{\mu y} + v_{\mu h}, \quad (23)$$

¹⁶This follows from the variance decomposition $var_t(x) = \mathbb{E}[var_t^a(x)] + var_t[\mathbb{E}(x|a)]$, where the second term is zero if we abstract from the terms \mathcal{M}_t^a in the allocations.

where $var_t(\alpha) = (1 - \delta) \sum_{a=0}^{\infty} \delta^a var_t^a(\alpha)$ is the unconditional cross-sectional variance of the uninsurable component of log wages, and $var_t(\varepsilon)$ and $var_t(\varphi)$ are the corresponding variances for the insurable component of wages and for preference dispersion.

Dispersion over the life cycle Let $\Delta var_t^a(\log \hat{x}) = var_t^a(\log \hat{x}) - var_{t-1}^{a-1}(\log \hat{x})$ be the within-cohort change (i.e., between age $a - 1$ in year $t - 1$ and age a in year t) in the variance of $\log \hat{x}$:

$$\Delta var_t^a(\log \hat{w}) = v_{\omega t} + v_{\eta t} + \Delta v_{\theta t} \quad (24)$$

$$\Delta var_t^a(\log \hat{h}) = v_{\chi t} + \left(\frac{1 - \gamma}{\sigma + \gamma}\right)^2 v_{\omega t} + \frac{1}{\sigma^2} (v_{\eta t} + \Delta v_{\theta t}) \quad (25)$$

$$\Delta cov_t^a(\log \hat{w}, \log \hat{h}) = \frac{1 - \gamma}{\sigma + \gamma} v_{\omega t} + \frac{1}{\sigma} (v_{\eta t} + \Delta v_{\theta t}) \quad (26)$$

$$\Delta var_t^a(\log \hat{c}) = v_{\chi t} + \left(\frac{1 + \sigma}{\sigma + \gamma}\right)^2 v_{\omega t} \quad (27)$$

$$\Delta cov_t^a(\log \hat{h}, \log \hat{c}) = v_{\chi t} + \frac{(1 - \gamma)(1 + \sigma)}{(\sigma + \gamma)^2} v_{\omega t} \quad (28)$$

$$\Delta cov_t^a(\log \hat{w}, \log \hat{c}) = \frac{1 + \sigma}{\sigma + \gamma} v_{\omega t} \quad (29)$$

Note, first, that none of these moments involve measurement error, reflecting our assumption that the variance of measurement error is independent of age and time. Second, note that because all shocks in our economy are either permanent or i.i.d., all these moments are independent of age.

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-consumption utility case, only the former matters for hours.

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 $\gamma > 1$, the cumulation of permanent uninsurable shocks pushes the covariance down as individuals age, while the cumulation of permanent insurable shocks pulls it up.

The change in the variance of consumption over the life cycle is determined by the variances of uninsurable productivity and preference shocks. The uninsurable-wage-shock multiplier for consumption, $(1 + \sigma) / (\sigma + \gamma)^2$, is exactly one either in the log case or in

the case of inelastic labor supply, i.e., $\sigma \rightarrow \infty$.

When $\gamma > 1$, hours move up in response to a negative uninsurable wage shock, while consumption moves down, driving the consumption-hours covariance down over the life cycle as $var_t^a(\alpha)$ rises with age. Preference shocks work in the opposite direction: as $var_t^a(\varphi)$ rises over the life cycle, the covariance between hours and consumption is driven up. The model predicts that the covariance between consumption and wages will increase over the life cycle, in proportion to $var_t^a(\alpha)$.

Expressions for growth in overall cross-sectional dispersion follow naturally from those above: for variable x ,

$$\Delta var_t(\log \hat{x}) = \delta \Delta var_t^a(\log \hat{x}) + (1 - \delta) (var_t^0(\log \hat{x}) - var_{t-1}(\log \hat{x})).$$

The first term in this expression reflects the change in cross-sectional dispersion within the set of agents that survive from $t - 1$ to t . The second term reflects the change in cross-sectional dispersion as new entrants (with age $a = 0$) replace the mass $1 - \delta$ of agents who die off.

Micro moments These moments are computed as variances and covariances of individual changes in log wages and log hours between $t - 1$ and t .¹⁷ Let $var_t^a(\Delta \log \hat{x})$ denote the cross-sectional variance of $(\log \hat{x}_t^a - \log \hat{x}_{t-1}^{a-1})$ for the set of individuals of age a at date t for whom variable x is observed at both $t - 1$ and t :

$$var_t^a(\Delta \log \hat{w}) = v_{\omega t} + v_{\eta t} + v_{\theta t} + v_{\theta, t-1} + 2v_{\mu y} + 2v_{\mu h} \quad (30)$$

$$var_t^a(\Delta \log \hat{h}) = v_{\chi t} + \left(\frac{1 - \gamma}{\sigma + \gamma}\right)^2 v_{\omega t} + \frac{1}{\sigma^2} (v_{\eta t} + v_{\theta t} + v_{\theta, t-1}) + 2v_{\mu h} \quad (31)$$

$$cov_t^a(\Delta \log \hat{w}, \Delta \log \hat{h}) = \frac{1 - \gamma}{\sigma + \gamma} v_{\omega t} + \frac{1}{\sigma} (v_{\eta t} + v_{\theta t} + v_{\theta, t-1}) - 2v_{\mu h} \quad (32)$$

Note that, once again, none of these moments depend on age a . Thus the model implies that the variances and covariances of individual growth rates should be invariant to age and thus common across cohorts.

4 Identification

This section exploits the closed-form solutions for cross-sectional second moments to describe conditions under which one can prove identification of the structural parameters.

¹⁷Given 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. Moreover, we omit moments involving changes in consumption since we do not use the limited longitudinal dimension of the CEX.

The conditions for identification depend on data availability and on the parametric restrictions one is willing to make. 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 $t = 1, \dots, T$.

PROPOSITION 2 [IDENTIFICATION] *Suppose one has access to an unbalanced panel on wages and hours and a repeated cross section on wages, hours, and consumption from $t = 1, \dots, T$. Then, the parameters $\{\sigma, \gamma, v_{\mu h}, v_{\mu y}, v_{\mu c}\}$ as well as the sequences $\{v_{\varphi^0 t}, v_{\alpha^0 t}\}_{t=1}^T$, $\{v_{\kappa^0 t}, v_{\theta t}\}_{t=1}^{T-1}$, $\{v_{\omega t}, v_{\chi t}\}_{t=2}^T$, and $\{v_{\eta t}\}_{t=2}^{T-1}$ are identified. The values $v_{\eta T} + v_{\theta T}$ and $v_{\kappa^0 T} + v_{\theta T}$ are also identified.*

See the Appendix for the proof.

We now turn to alternative scenarios which reflect the limitations of current survey data for the United States. We first assume that one only has access to cross-sectional data on consumption, hours, and earnings. Next, we assume that the consumption cross section begins later than the wage/hours panel. Finally, we prove identification when only one cohort of panel data on wages and hours is available, along with cross-sectional data on consumption. These three cases would be relevant if one were to estimate the model using, respectively, (i) only CEX data, (ii) the entire available time series for the PSID and the CEX, (iii) the NLSY and the CEX.

COROLLARY [LIMITED DATA AVAILABILITY] *(A) Suppose we only have access to repeated cross-sectional data on consumption, hours, and wages from $t = 1, \dots, T$. Then, the parameters $\{\sigma, \gamma, v_{\mu y}, v_{\mu h}, v_{\mu c}\}$ and the sequences $\{v_{\omega t}, v_{\chi t}, (v_{\eta t} + \Delta v_{\theta t})\}_{t=2}^T$ and $\{v_{\varphi^0 t}, v_{\alpha^0 t}, (v_{\kappa^0 t} + v_{\theta t})\}_{t=1}^T$ are identified. (B) Suppose data on consumption are available only from $t = t^*$, with $0 < t^* < T$. Then, the model is identified under the same assumptions as in Proposition 2. (C) Suppose wage and hours data are only available for one cohort from $t = 1, \dots, T$. Then, $(v_{\varphi^0}, v_{\alpha^0}, v_{\kappa^0})$ are identified only for the available cohort, and all other parameters are identified under the same assumptions as in Proposition 2.*

5 Estimation

We begin by describing 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). Our PSID and CEX samples are those constructed by Heathcote, Perri and Violante (2009). Since these datasets are widely used, we refer to that paper for a detailed description of these two surveys, sampling methodology, and variable definitions.

We use the 1968–1997 waves of the PSID to build an unbalanced panel with information on demographic characteristics, individual annual hours, and annual earnings. The PSID asks questions about earnings in the previous year, so our data refer to the period 1967–1996.¹⁸ Since we only use the SRC sample, we do not use sample weights.

The CEX contains detailed information on nondurable and durable household consumption, as well as on demographic characteristics of household members, individual earnings, and hours worked. We focus on the quarterly Interview Survey, and we only retain households for which we have a complete set of quarterly interviews, so that we can aggregate consumption to an annual frequency in line with data on earnings and hours worked.¹⁹ Our baseline measure of consumption includes expenditures on nondurables, services, small durables, and an estimate of the service flow from vehicles and housing. Consistent and continuous data over time are available only since 1980, hence we restrict attention to the 1980–1996 cross sections and use weights provided in the survey.

Since we use both the PSID and CEX data jointly, it is important that we apply the same sample selection criteria across the two datasets. We first drop seriously incomplete or implausible observations.²⁰ We use an imputation procedure to adjust for top-coding based on the Pareto distribution. We then select households in which the head is between the ages of 25 and 59, and works at least 260 hours in the year. In both datasets, the hourly wage is computed as annual pre-tax labor earnings divided by annual hours worked.²¹ To avoid severe selection issues, we use wages and hours for males only. We

¹⁸PSID surveys became biannual after 1997. Exploring this more recent data would require appropriately modifying the identification scheme.

¹⁹See Attanasio, Battistin and Ichimura (2007) for a study that uses both the Diary and the Interview Surveys. The former is better designed to capture frequently purchased items (such as food, beverages, and personal care items), while the latter aims to provide comprehensive information on up to 95% of the typical household’s consumption expenditures.

²⁰We drop records if 1) there is no information on age for either the head or the spouse, 2) either the head or spouse has positive labor income but zero annual hours, and 3) either the head or spouse has an hourly wage less than half of the corresponding federal minimum wage in that year. In the CEX, we drop households which report implausibly low quarterly consumption expenditures (less than \$100, in 2000 dollars). In order to reduce measurement error, we also exclude CEX households flagged as “incomplete income reporters.”

²¹Labor earnings are defined in both surveys as the sum of all income from wages, salaries, commissions,

do not equalize any variables. All monetary variables are deflated using the Consumer Price Index (CPI).²²

Because our model is not designed to address variation in first moments, we regress individual log wages, hours, and consumption (the former two in both datasets, the latter only in the CEX) on year dummies, and on a quartic in age.²³ We then use the residuals from these regressions to construct variances and covariances in levels and first-differences for all available age/year cells constructed by grouping observations in any given year into 31 age classes: “age” 27 (25–29), “age” 28 (26–30), and so on until “age” 57 (55–59).²⁴

5.2 Estimation method

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

The structural parameters to be estimated in the model are preference parameters $\{\gamma, \sigma\}$, variances of measurement error $\{v_{\mu h}, v_{\mu y}, v_{\mu c}\}$, and the sequences of variances $\{v_{\varphi^0 t}, v_{\alpha^0 t}, v_{\kappa^0 t}, v_{\omega t}, v_{\eta t}, v_{\chi t}, v_{\theta t}\}_{t=1}^T$, where $T = 30$ is the length of the PSID sample, and $T^* = 17$ is the length of the CEX sample.

As discussed in Section 4, certain parameter values are not identified at the end points of our sample period. Thus we make two additional identifying assumptions that allow us to estimate every parameter. First, we assume that prior to year $t = 1$ (1967) all parameter values were constant and equal to their values in 1967. Given this assumption, we can immediately construct theoretical expressions for all moments in levels and first differences for all $t = 1, \dots, T$. Moreover, this assumption also identifies $(v_{\omega 1}, v_{\eta 1}, v_{\chi 1})$, the variances for the permanent shocks in 1967.²⁵ Second, we assume $v_{\theta T} = v_{\theta, T-1}$, which

bonuses, and overtime, and the labor part of self-employment income.

²²See Heathcote et al. (2009), Table 1, for a summary of the dropped observations at each step of sample selection, and their Table 2 for a comparison of basic demographic and labor market statistics across the two datasets. In short, we find that they are broadly consistent, and also line up very closely with analogous statistics computed from the much larger March Current Population Survey.

²³Recall 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.

²⁴The only reason to group observations into five-year age groups is to reduce sampling variation: the number of observations in some single-age, single-year cells is very small.

²⁵For example, under this convenient assumption:

$$\begin{aligned} cov_1^2(\log \hat{w}, \log \hat{h}) - cov_1^1(\log \hat{w}, \log \hat{h}) &= \frac{1-\gamma}{\sigma+\gamma} v_{\omega 1} + \frac{1}{\sigma} v_{\eta 1} \\ var_1^2(\log \hat{w}) - var_1^1(\log \hat{w}) &= v_{\omega 1} + v_{\eta 1} \\ var_1^2(\log \hat{h}) - var_1^1(\log \hat{h}) &= v_{\chi 1} + \left(\frac{1-\gamma}{\sigma+\gamma}\right)^2 v_{\omega 1} + \frac{1}{\sigma^2} v_{\eta 1}. \end{aligned}$$

then identifies $v_{\eta T}$ and $v_{\kappa^0, T}$ (see steps 2 and 3 in the proof of Proposition 2).

One last issue to discuss is that around 1992 (survey year 1993), the first year of computer-assisted telephone interviewing in the PSID, earnings and hours appear unusually volatile. This short-lived surge in volatility has been widely attributed to a temporary increase in measurement error. To account for this, we simply assume that the variance of measurement error in earnings and hours in 1992 is twice as large as in other years.

In the baseline estimation, we use $J = 6,981$ moment conditions, corresponding to 31 age groups over 30 years for the three moments in levels involving wages and hours, 31 age groups over 17 years for the three moments in levels involving consumption, and 30 age groups over 29 years for the three moments in first differences. Let Λ denote the $(N \times 1)$ vector of parameter values to estimate, with $N = 5 + 7T - 1 = 214$. Let $m(\Lambda)$ denote the $(J \times 1)$ vector of theoretical covariances with typical element $m_j(\Lambda)$, where the index $j = 1, \dots, J$ denotes the position in the vector. Correspondingly, we define \hat{m} as the vector of empirical covariances with typical element \hat{m}_j .

Our estimator solves the following minimization problem:

$$\min_{\Theta} [\hat{m} - m(\Lambda)]' \mathcal{W} [\hat{m} - m(\Lambda)], \quad (33)$$

where \mathcal{W} is a $(J \times J)$ weighting matrix. Standard asymptotic theory implies that the estimator $\hat{\Lambda}$ is consistent and asymptotically normal. Due to the small sample size, we make two choices. First, in the baseline estimation we use an identity matrix for \mathcal{W} , but we also report results for the case in which \mathcal{W} is a diagonal matrix with entries corresponding to the number of observations used for each of the J moments.²⁶ Second, instead of asymptotic standard errors, we compute 90–10 confidence intervals through a block-bootstrap procedure based on 600 replications.²⁷

5.3 Results

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

²⁶The bulk of the literature follows this strategy, in light of the Monte Carlo simulations of Altonji and Segal (1996) who argue that in common applications there is a substantial small sample bias when using the optimal weighting matrix characterized by Chamberlain.

²⁷Bootstrap samples are drawn at the household level with each sample containing the same number of observations as the original sample. The implied confidence intervals thus account for arbitrary serial correlation, heteroskedasticity, and estimation error induced by the first-stage regression of individual observations on age and time.

Parameter estimates Table 1 reports parameter estimates and bootstrapped 90–10 confidence intervals. For our two key preference parameters, we estimate $\gamma = 2.02$ and $\sigma = 2.62$. In both cases the confidence intervals are narrow. In the context of the set of moments we use to estimate the model, the key feature of the parameter γ is that it determines the wealth effect on labor supply associated with uninsurable shocks to wages. Thus moments involving the covariance between (changes in) wages and hours are informative about the true value for γ . In a similar spirit, Chetty (2006) argues that existing empirical evidence on the response of hours to permanent shocks to wages can be used to bound estimates for risk aversion. An advantage of our fully structural approach is that we can identify γ in an environment with a mix of uninsurable and insurable permanent wage shocks, even though these are associated with very different labor supply responses (see eq. 16).²⁸ The implied Frisch elasticity of labor supply given our estimate for σ is 0.38, a value that is consistent with the microeconomic evidence for males (see e.g. Browning et al., 1999). Our estimate for the variance of measurement error in log hours is 0.036, while we find virtually no measurement error in earnings. The estimated variance of measurement error in log consumption is 0.033.

Table 1 reports average estimated values for the time-varying parameters.²⁹ The top left panel of Figure 1 shows the cross-sectional variance of the uninsurable component of wages (α_t) implied by our estimates, year by year. The overall variance of the uninsurable component declines in the 1970s, and then rises sharply in the 1980s. Interestingly, this pattern broadly accords with the fall of the skill premium in the late 1960s to mid-1970s, and the subsequent increase in the 1980s and beyond. Under this interpretation, “skill-biased demand shifts” represent the main source of uninsurable wage shocks.³⁰

The corresponding picture (top right panel of Figure 1) for the permanent insurable component of wages (κ_t) is somewhat different. The cross-sectional variance of this component is generally increasing throughout the first two decades, but declines in the last decade of the sample. At the same time, the variance of transitory insurable shocks (bottom left panel) grows steadily throughout the sample and peaks in the first half of the 1990s, consistently with Moffitt and Gottschalk’s (2002) findings from estimated earnings dynamics. Overall, the relative contribution of the insurable component of wage

²⁸The intertemporal elasticity of substitution implied by our estimate for γ , 0.5, is within the standard range of existing estimates.

²⁹See Table 4 in the Appendix for the complete year-by-year set of parameter estimates, and associated confidence intervals.

³⁰This interpretation is consistent with Heathcote et al. (2008a) who, in the context of an augmented version of the standard incomplete-markets model, show that skill-biased demand shifts are the main driver of the rise in consumption inequality.

Table 1: Parameter Estimates

Preferences		Life cycle Shocks			
σ	γ	$\overline{v_{\omega}}$	$\overline{v_{\eta}}$	$\overline{v_{\theta}}$	$\overline{v_{\chi}}$
2.633	2.019	0.004	0.004	0.032	0.001
[2.503,2.830]	[1.942,2.157]	[0.001,0.009]	[0.001,0.009]	[0.026,0.039]	[0.000,0.003]
Cohort Effects			Measurement Error		
$\overline{v_{\varphi^o}}$	$\overline{v_{\kappa^o}}$	$\overline{v_{\varphi^o}}$	$v_{\mu y}$	$v_{\mu h}$	$v_{\mu c}$
0.084	0.081	0.045	0.000	0.036	0.033
[0.070,0.095]	[0.064,0.098]	[0.039,0.051]	[0.000,0.000]	[0.034,0.036]	[0.029,0.037]

Table 2: Variance Decomposition, 1967 – 1996

	Total Variance	Contribution of Variance in			
		Insurable Wages	Uninsurable Wages	Preferences	Measurement Error
$var(\log \hat{w})$	0.340	0.161	0.142	0.000	0.037
$var(\log \hat{h})$	0.118	0.023	0.008	0.050	0.037
$var(\log \hat{y})$	0.446	0.308	0.088	0.050	0.000
$var(\log \hat{c})$	0.169	0.000	0.086	0.050	0.033

dispersion rises in the late 1960s and 1970s, and is broadly stable thereafter.

Figure 1 also shows what part of aggregate dispersion reflects initial heterogeneity at the time of labor market entry (the dotted line labeled “cohort”) with the remainder reflecting the cumulation of subsequent permanent shocks.³¹ The decomposition indicates that the dynamics of both uninsurable and insurable wage dispersion are largely driven by shocks received during the life cycle. Put differently, even though initial heterogeneity in labor productivity is large, differences across adjacent cohorts are small and induce only slow movements in cross-sectional wage dispersion.

Cross-sectional dispersion in preferences (bottom right panel) is small relative to dispersion in wages, but it shows a constant and slow increase. Contrary to wages, rising dispersion in preferences reflects pre-existing heterogeneity rather than subsequent shocks. One potential explanation for the prominent role of cohort effects is the secular trend away from the traditional married household, and toward greater variation in household composition and the gender division of market work.

Variance decomposition One attractive feature of our cross-sectional moment expressions is that they are all additively separable in terms capturing the roles of preference dispersion, insurable wage dispersion, uninsurable wage dispersion, and measurement error. Thus (co-)variance decompositions are unique and easy to compute given parameter

³¹In the earlier part of the sample, the decomposition between the contributions of initial dispersion versus subsequent shocks is somewhat sensitive to the identifying assumptions we make regarding the parameters that define the wage process prior to 1967.

estimates. In Table 2, we report the average contribution of each component across the entire 1967–2006 period: in the next section, we discuss changes in dispersion over time.

Wage dispersion is roughly equal parts insurable and uninsurable in nature, with a small additional role for measurement error. The effect of differential comovement between hours and wages for insurable versus uninsurable shocks shows up clearly in the variance decomposition for earnings: relative to the variance of wages, endogenous labor supply doubles the variance of the insurable component, and reduces the variance of the uninsurable component by a third.

Preference heterogeneity matters for both hours and consumption dispersion: preference heterogeneity accounts for almost half of the variance of hours worked, and one-third of the variance of consumption. Measurement error also plays a large role, accounting for one-third of the observed variance of hours, and one-fifth the variance of consumption.

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

Time series fit Before discussing the model fit in the time dimension, we briefly describe the evolution of cross-sectional inequality in the data.³² Figure 2 indicates that the variance of log male wages increases steadily by around ten log points between the late 1960s and early 1990s, with an acceleration in the 1980s. Over the same period there is a much smaller increase in the variance of male hours—around two log points. The correlation between log wages and log hours rises quite dramatically in the first part of the sample, approaching zero from negative values. These facts combined imply that the increase in the variance of log earnings is larger than the increase in the variance of log wages. The variance of log consumption grows by only about five log points between 1980 and 1996, in line with earlier estimates by Krueger and Perri (2006). The correlation between wages and consumption is positive and increases slightly over time. The correlation between hours and consumption is around 0.2 and relatively stable.

These features of the data underlie the parameter estimates in Table 1 and Figure 1, and given these estimates the model in turn replicates the data very closely.

As is clear from Figure 2, to match the sharp widening of the wage distribution, the model requires a large increase in the variance of either the insurable or the uninsurable

³²When plotting unconditional moments by year in Figure 2, we average across all age groups within each year. In both model and data we weight successive age groups by survival probabilities $\delta = 1 - (1/35)$ to account for mortality. Thus both model and data exhibit the same stationary age distribution. 90–10 confidence intervals based on bootstrapped samples are plotted around the time series.

component of wages. Before 1980, CEX data are not available, and thus it is primarily wage and labor supply data which convey information on structural parameters. Over the first half of the sample, we see a sharp rise in the wage-hours correlation, and a modest increase in the variance of hours. The model fits both data patterns through a sharp growth in the insurable wage component, and a relatively small Frisch elasticity (high σ) which mutes the impact of insurable shocks on hours dispersion.

The model also replicates the observed average negative correlation between hours and wages, even though insurable wage dispersion, dominant over this period, induces a positive correlation. Examining eq. (19) suggests two main factors at work in the model. First and foremost, uninsurable wage variation pushes the correlation below zero when $\gamma > 1$, since income effects dominate substitution effects. This is precisely what the estimation delivers for γ . Second, measurement error in hours drives down the correlation (in both model and data) because wages are measured as earnings divided by hours.

After 1980 consumption data are available and further inform the estimation.³³ Since the early 1980s, both consumption dispersion and the wage-consumption correlation have risen, calling for an increase in uninsurable wage dispersion in the model, and a slowdown in the rise of insurable wage dispersion. This pattern is also consistent with the end of growth in the wage-hours correlation.

The main reason why the increase in the variance of consumption is small relative to the increase in uninsurable wage dispersion is that the estimated value for γ implies that some portion of uninsurable changes in relative wages show up in changes in relative leisure rather than consumption. In particular, given $\gamma = 2.02$ and $\sigma = 2.62$, eq. (20) indicates that each log point increase in the variance of the uninsurable component of wages translates into a 0.6 log point increase in the variance of consumption.

The correlation between consumption and hours is informative about the role of preference dispersion. This correlation is positive in the data, while uninsurable wage dispersion tends to induce a negative correlation when $\gamma > 1$ (see eq. 21). The model replicates the observed positive correlation thanks to preference heterogeneity: households that place relatively more weight on consumption in utility work longer hours and consume more. Larger uninsurable wage shocks tend to drive the consumption-hours correlation down over time. To offset this force and replicate the roughly flat pattern for the correlation in the data, the estimation calls for a rising time profile for preference dispersion.

Finally, the model can be used to infer the path for consumption inequality in the

³³Information about parameter values before 1980 is also embedded in the moments involving consumption for older cohorts present in the early years of CEX.

Table 3: Decomposition of Changes in (Co-)Variances

	Change over Time (1967-1996)				Change over Life Cycle (1980-2006)			
	Total	Contribution of			Total	Contribution of		
		v_ε	v_α	v_φ		v_ε	v_α	v_φ
$var(\log \hat{w})$	0.084	0.059	0.025	0.000	0.208	0.095	0.113	0.000
$var(\log \hat{h})$	0.032	0.009	0.001	0.022	0.027	0.014	0.006	0.007
$cov(\log \hat{w}, \log \hat{h})$	0.018	0.023	-0.005	0.000	0.011	0.036	-0.025	0.000
$var(\log \hat{y})$	0.152	0.114	0.016	0.022	0.257	0.181	0.069	0.007
$var(\log \hat{c})$	0.038	0.000	0.016	0.022	0.076	0.000	0.069	0.007
$cov(\log \hat{h}, \log \hat{c})$	0.018	0.000	-0.004	0.022	-0.012	0.000	-0.019	0.007
$cov(\log \hat{w}, \log \hat{c})$	0.020	0.000	0.020	0.000	0.088	0.000	0.088	0.000

1970s, before the start of our CEX sample. Projecting backward, the model predicts flat or modestly declining consumption inequality pre-1980. As mentioned, this flat then increasing pattern for consumption inequality parallels the well-known dynamics of the skill premium over the period. It is also broadly consistent with evidence provided by Slesnick (2001, chapter 6) and Guvenen and Smith (2008, Figure 4) who, notwithstanding data comparability issues, use CEX data pre-1980 in order to construct a longer series for U.S. consumption dispersion.

Table 3 describes the total predicted changes between 1967 and 2006 for a set of variances and covariances, and decomposes these changes into components attributable to changes in insurable wage dispersion, uninsurable wage dispersion, and preference dispersion. The variance of measurement error is assumed constant over time (except for the 1992 dummy), and thus contributes nothing to changes in cross-sectional dispersion.

The table indicates that two-thirds of the overall increase in cross-sectional wage dispersion was insurable in nature. As explained above, both labor supply data (namely, the strong growth in the wage-hours correlation), and consumption data (namely, the modest rise in the variance of log consumption) point in this direction.

Interestingly, roughly two-thirds of the predicted increases in the variances of hours and consumption are attributed to a rise in preference dispersion, and changes in preference dispersion are also the main driver of predicted changes in the covariance between hours and consumption. Therefore, the model indicates that non-wage factors played a significant role in explaining the dynamics of cross-sectional inequality in the United States over the period under study.

Life cycle fit Figure 3 compares the evolution of model and data in the life cycle dimension.³⁴ In the data, the variance of log wages increases by twenty-one log points,

³⁴These plots are constructed by averaging across cohorts for each age, using only data for the 1980–

approximately linearly, between ages 27 and 57. The variance of log hours has a slight U-shape, while the correlation between hours and wages declines modestly. The variance of log consumption grows by about seven log points over the life cycle. As emphasized in Heathcote et al. (2005), this number represents a much smaller increase than previously reported in the pioneering work of Deaton and Paxson (1994).³⁵

The model is able to match the large life cycle rise in wage dispersion and, simultaneously, the small rise in consumption dispersion thanks to the fact that most permanent shocks are insurable (though even uninsurable shocks do not completely pass through to consumption because they are partially smoothed via labor supply).

Equation (25) shows that inequality in hours worked will typically increase over the life cycle in the model, as shocks cumulate. However, the overall increase turns out to be small, since γ is not too far from one (so increasing uninsurable wage dispersion has little impact on hours), and the Frisch elasticity is low (so the effect of rising insurable dispersion is also small).³⁶

From equation (26) one can see that the model can predict both positive and negative trends for the correlation between hours and wages over the life cycle, depending on the relative size of insurable and uninsurable innovations to productivity, as well as on the value for γ . In the estimation, the two effects roughly offset, so the model predicts a roughly constant correlation, while in the data this moment has a small downward trend.

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 hump-shaped in the data. The cumulation of permanent uninsurable shocks over the life cycle drives this dynamic (see Table 3) with preference shocks playing a modest offsetting role.

These life cycle facts are central to understanding how the estimation assigns cross-sectional dispersion to initial heterogeneity (at age 27) versus subsequent life cycle shocks. Table 3 indicates that the evolution of dispersion over the life cycle is primarily driven by shocks to labor productivity, with preference shocks playing in a minor role. This is

2006 period. This ensures comparability across the life cycle profiles based on CEX data and those from the PSID. The decomposition in Table 3 is based on model estimates corresponding to the same period.

³⁵Deaton and Paxson (1994, Figure 8) report an increase in the variance of log consumption of almost twenty log points between ages 25 and 55. One reason is that their study covers the period 1980–1990, precisely when the bulk of the rise in cross-sectional dispersion is concentrated, while our sample covers a longer time period.

³⁶Within 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, and (2) a non-degenerate distribution of initial wealth which increases the variance of hours for younger agents.

consistent with the estimation results indicating that the bulk of cross-sectional preference dispersion is pre-determined at age 27. Intuitively, permanent life cycle shocks to preferences must be relatively small, since otherwise the model would deliver counterfactually large life cycle increases in dispersion in consumption and hours worked. At the same time, substantial initial heterogeneity in preferences is required to match the levels of inequality in hours, consumption, and the hours-consumption correlation.

Fit in first-differences Figure 4 documents moments in first-differences, both over time (top panels) and over the life cycle (bottom panels). Over time, 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 plays a key role in identifying the variance of the transitory insurable shock θ_t . The variance of growth in hours is quite flat, except for a spike in the early 1990s.

The measurement error dummy for 1992 allows us to match the spike in the variance of wage and hours growth in the early 1990s, and the dip around the same time in the correlation between wage growth and hours growth. By virtue of our assumption that measurement error doubled in the first year of the new methodology, the model mirrors these movements.

The correlation between growth in wages and growth in hours increases in the data between 1967 and 1990, a trend that is broadly replicated by the model. As discussed above, this rising correlation reflects an increase over time in the relative share of wage fluctuations that can be insured.

Over the life cycle, the model predicts that the (co-)variances of all growth rates should be constant across age groups. This prediction is in line with the empirical variance of wage growth and the covariance between wage and hours growth, but not the variance of growth in hours declines until around age forty. The model replicates the average variance for wage growth, but under-predicts the observed levels for the variance of hours growth, and the covariance between hours growth and wage growth. We discuss this shortcoming of the model in the next section.

Robustness Exploring alternative specifications sheds more light on the mapping between model ingredients and parameter values on the one hand, and the large set of targeted empirical cross-sectional moments on the other.³⁷

First, we consider an alternative measure for consumption, which excludes durables. This measure of consumption exhibits slightly higher variance than our baseline mea-

³⁷Table 5 in the Appendix shows how parameter estimates vary across the set of alternative models we consider.

sure, but a lower covariance with wages. The estimation routine then imputes more consumption dispersion to measurement error, and less to initial dispersion in the uninsurable component of wages. Less uninsurable wage dispersion in turn points to a higher estimate for γ (2.32 rather than 2.02) to replicate the observed wage-hour covariance, which improves the fit for the correlation between wages and hours over the life cycle by generating a small decline, as observed in the data.

Second, we consider a version of the model in which the variances of initial dispersion for wages and preferences are assumed time-invariant. This has a negligible impact on any of the preference or measurement error parameter estimates. While this model assumes, counter-factually, that dispersion among labor market entrants is time-invariant, the fit for the evolution of aggregate cross-sectional dispersion over time remains good.

Third, we estimate the model without using any moments in first-differences. Recall from the corollary that without moments in first-differences we cannot decompose the overall variance of insurable wage dispersion into permanent and transitory components, but all other model parameters are still identified. Reassuringly, estimates for these other parameters are little changed relative to the baseline model.

Fourth, we consider a version of the model with transitory i.i.d. preference shocks. These serve as an additional source of variance in hours growth, and thus as a simple way for the estimation routine to replicate the high observed volatility in hours growth without relying on big measurement error in hours or a high Frisch elasticity. The effect of including these shocks is similar to the effect of excluding the moments in first-differences: the estimated variance of measurement error in hours is reduced and the Frisch elasticity falls (from 0.38 to 0.22).

Fifth, we consider more restricted models for preference variation. Assuming no life cycle preference shocks has relatively little impact, which is not surprising given that these shocks are estimated to be small in the baseline model. Assuming no initial preference heterogeneity forces the model to try to replicate the observed cross-sectional dispersion in consumption and hours through a mixture of more measurement error, and a higher Frisch elasticity (0.58). The main failure of this restricted model is that it is unable to generate a positive consumption-hours correlation, a key feature of the data.

Finally, we tried a model with no transitory shocks to wages. Not surprisingly, this model does poorly in terms of replicating moments in first differences (notwithstanding a higher estimate for the variance of measurement error in earnings). However, all other parameters are little changed, and the fit for moments in levels remains very good. This finding reflects the fact that in the context of our model, what is important for cross-

sectional moments is not whether shocks to wages are permanent or transitory, but whether they are insurable or uninsurable.

6 Concluding Remarks

This paper has laid out a new theoretical framework to study consumption and labor supply in the presence of idiosyncratic risk to wages and preferences. A distinguishing feature of the model is that it can be solved analytically. Tractability is achieved by extending the environment of Constantinides and Duffie (1996) to incorporate endogenous labor supply, insurable wage risk, and preference heterogeneity. The model allows for any degree of risk-sharing between a bond economy and complete markets.

The framework delivers closed-form expressions for the equilibrium cross-sectional variances and covariances of the joint distribution of wages, hours, and consumption. These expressions make the link between structural parameters and observed inequality transparent, permit a formal proof of identification, and facilitate estimation of rich dynamic models for the structure of wages and preferences.

The model implies unique decompositions of cross-sectional inequality into components attributable to insurable and uninsurable wage dispersion, preference dispersion, and measurement error. Wage and preference dispersion, in turn, reflect a mix of permanent heterogeneity and life cycle shocks.

One insight from the moment expressions is that labor supply contains useful information on the extent of insurance available to households. This complements evidence from consumption, the typical focus of the literature (e.g., Blundell and Preston, 1998). We find that moments involving both hours and consumption paint a coherent picture: over the 1970s and 1990s, the increase in wage risk was insurable for US households, while in the 1980s it was largely uninsurable. Non-wage factors, which we capture through preference heterogeneity, also contributed to rising inequality in hours and consumption.

The model outlined in Section 2 incorporates the essential ingredients for the main application of the paper: a study of the factors determining the evolution of US inequality over time and over the life cycle. The same model can address other issues. First, assuming lognormal shocks, one can obtain closed-form solutions for the welfare implications of idiosyncratic wage risk. In a greatly simplified version of our setup, Heathcote et al. (2008b) show that rising insurable risk can be welfare improving given flexible labor supply. Second, the allocations can be used to derive closed-form expression for the transmission coefficients of permanent wage and earnings fluctuations to consumption

growth, providing a structural interpretation for the estimates in Blundell et al. (2008).

In the baseline framework, we have chosen to remain agnostic about the mechanisms that deliver partial insurance against wage risk, and about the underlying sources of preference dispersion. The model can be generalized to shed more light on these issues by explicitly modeling families and taxation, without compromising tractability.

It is straightforward to reinterpret equilibrium allocations from the baseline model as arising in the context of a model of families. In particular, if family members belong to the same island, and family utility is separable between public consumption and each member's labor supply, then one of the mechanisms which deliver within-island insurance is risk-sharing within the family (Hayashi, Altonji, and Kotlikoff, 1996). Dispersion and fluctuations in household composition are an obvious source of preference heterogeneity and risk. Modeling these explicitly would require including a measure of household size as an extra state variable for the household, defining a corresponding stochastic process, and introducing an equivalence scale for household consumption.

Our framework can also be extended to incorporate progressive taxation and government redistribution. Following Benabou (2002), we can introduce a balanced-budget progressive tax/transfer system, where after-tax labor earnings y are given by $y = (wh)^{1-\tau} \bar{y}^\tau$, where τ is the average marginal tax rate (weighted by income), and \bar{y} is average labor earnings. Explicitly modeling government redistribution this way has several payoffs. First, we can show that progressive taxation changes the response of labor supply to insurable wage shocks, and thus the interpretation of the estimated Frisch elasticity. Second, by estimating τ we can explicitly quantify the role of progressive taxation as an insurance mechanism. Third, we can pursue a closed-form welfare analysis of progressive taxation.

Finally, even though in the baseline framework we have abstracted from aggregate risk, we can easily introduce economy-wide fluctuations and link idiosyncratic risk to the aggregate state. The framework remains analytically tractable. In particular, let the source of aggregate fluctuations be shocks to productivity (the average wage per efficiency unit of labor). Follow Constantinides and Duffie (1996) by assuming conditional heteroskedasticity of the four shocks modeled as $v_{jt} = \alpha_j + \beta_j Z_t$, where $j \in \{\omega, \eta, \theta, \chi\}$. The framework still yields analytical expressions for aggregate allocations and asset prices. This extended set-up can be used to study the link between idiosyncratic fluctuations and aggregate business cycles, the welfare costs of business cycles, and asset pricing.

References

- [1] Abowd, John, and David Card. 1989. “On the Covariance Structure of Earnings and Hours Changes.” *Econometrica*, 57(2): 411–445.
- [2] Aiyagari, S. R. 1994. “Uninsured Idiosyncratic Risk and Aggregate Saving.” *Quarterly Journal of Economics*, 109(3): 659–684.
- [3] Ales, Laurence, and Pricila Maziero. 2009. “Non-Exclusive Dynamic Contracts, Competition, and the Limits of Insurance.” Mimeo, University of Minnesota.
- [4] Altonji, Joseph, and Lewis M. Segal. 1996. “Small-Sample Bias in GMM Estimation of Covariance Structures.” *Journal of Business and Economic Statistics*, 14(3): 353–366.
- [5] Angeletos, George-Marios. 2007. “Uninsured Idiosyncratic Investment Risk and Aggregate Saving.” *Review of Economic Dynamics*, 10(1): 1–30.
- [6] Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura. 2007. “What Really Happened to Consumption Inequality in the US?” In *Measurement Issues in Economics—The Paths Ahead: Essays in Honor of Zvi Griliches*, ed. Ernst R. Berndt and Charles R. Hulten. Chicago: University of Chicago Press.
- [7] Attanasio, Orazio, and Steven J. Davis. 1996. “Relative Wage Movements and the Distribution of Consumption.” *Journal of Political Economy*, 104(6): 1227–1262.
- [8] Attanasio, Orazio, and Nicola Pavoni. 2007. “Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption.” NBER Working Paper 12994.
- [9] Attanasio, Orazio, and José V. Ríos-Rull. 2000. “Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?” *European Economic Review*, 44(7): 1225–1258.
- [10] Benabou, Roland. 2002. “Tax and Education Policy in a Heterogeneous Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?” *Econometrica*, 70(2): 481–517.
- [11] Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance.” *American Economic Review*, 98(5): 1887–1921.
- [12] Blundell, Richard, and Ian Preston. 1998. “Consumption Inequality and Income Uncertainty.” *Quarterly Journal of Economics*, 113(2): 603–640.
- [13] Bodie, Zvi, Robert C. Merton, and William F. Samuelson. 1992. “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model.” *Journal of Economic Dynamics and Control*, 16(3–4): 427–449.
- [14] Browning, Martin, Lars P. Hansen, and James J. Heckman. 1999. “Micro Data and General Equilibrium Models.” In *Handbook of Macroeconomics*, vol. 1, ed. John B. Taylor and Michael Woodford, 543–633. Amsterdam: Elsevier.
- [15] Caballero, Ricardo. 1990. “Consumption Puzzles and Precautionary Savings.” *Journal of Monetary Economics*, 25(1): 113–136.
- [16] Castañeda, Ana, Javier Diaz-Gimenez, and José V. Ríos-Rull. 2003. “Accounting for the US Earnings and Wealth Inequality.” *Journal of Political Economy*, 111(4): 818–857.
- [17] Chamberlain, Gary. 1984. “Panel Data.” In *Handbook of Econometrics*, vol. 3, ed. Zvi Griliches and Michael Intriligator, 1247–1318. New York: North-Holland.
- [18] Chetty, Raj. 2006. “A New Method of Estimating Risk Aversion.” *American Economic Review*, 96(5): 1821–1834.
- [19] Cole, Harold, and Narayana Kocherlakota. 2001. “Efficient Allocations with Hidden Income and Hidden Storage.” *Review of Economic Studies*, 68(3): 523–542.
- [20] Constantinides, George, and Darrel Duffie. 1996. “Asset Pricing with Heterogeneous Consumers.” *Journal of Political Economy*, 104(2): 219–240.
- [21] Cunha, Flavio, James J. Heckman, and Salvador Navarro. 2005. “Separating Uncertainty from Heterogeneity in Life-Cycle Earnings.” *Oxford Economic Papers*, 57(2): 191–261.

- [22] Deaton, Angus. 1997. *The Analysis of Household Surveys*. Baltimore: John Hopkins University Press.
- [23] Deaton, Angus, and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy*, 102(3): 437–467.
- [24] Fernandez–Villaverde, Jesus, and Dirk Krueger. 2004. "Consumption and Saving over the Life Cycle: How Important Are Consumer Durables?" Society for Economic Dynamics, Meeting Paper 357b.
- [25] Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97(3): 687–712.
- [26] Guvenen, Fatih, and Anthony Smith. 2008. "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach." Mimeo, Yale University.
- [27] Hall, Robert E. and Frederic Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica*, 50(2): 461–481.
- [28] Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff. 1996. "Risk–Sharing between and within Families." *Econometrica*, 64(2): 261–294.
- [29] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2005. "Two Views of Inequality Over the Life Cycle." *Journal of the European Economic Association*, 3(2-3): 765–775.
- [30] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008a. "The Macroeconomic Implications of Rising Wage Inequality in the U.S." Mimeo, New York University.
- [31] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008b. "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk." *Journal of Monetary Economics*, 55(3): 501–525.
- [32] Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante. 2009. "Quantitative Macroeconomics with Heterogeneous Households." NBER Working Paper 14768.
- [33] Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2): 360–399.
- [34] Huggett, Mark. 1993. "The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies." *Journal of Economic Dynamics and Control*, 17(5-6): 953–969.
- [35] Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2009. "Sources of Lifetime Inequality." Mimeo, Georgetown University.
- [36] Kaplan, Greg. 2006. "The Patterns of Lifecycle Inequality." Mimeo, New York University.
- [37] Krebs, Tom. 2003. "Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk." *Review of Economic Dynamics*, 6(4): 846–868.
- [38] Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies*, 73(1): 191–261.
- [39] Livshits, Igor, James MacGee, and Michele Tertilt. 2007. "Consumer Bankruptcy: A Fresh Start." *American Economic Review*, 97(1): 402–418.
- [40] Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2006. "Wage Risk and Employment Risk Over the Life Cycle." NBER Working Paper 14901.
- [41] Moffitt, Robert and Peter Gottschalk. 2002. "Trends in the Transitory Variance of Earnings in the United States." *Economic Journal*, 112(3): C68–C73.
- [42] Primiceri, Giorgio, and Thijs van Rens. 2009. "Heterogeneous Life Cycle Profiles, Income Risk and Consumption Inequality." *Journal of Monetary Economics*, 56(1): 20–39.
- [43] Slesnick, Daniel. 2001. *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*. Cambridge: Cambridge University Press.
- [44] Uhlig, Harald. 1996. "A Law of Large Numbers for Large Economies." *Journal of Economic Theory*, 8(1): 41–50.
- [45] Wang, Neng. 2003. "Caballero Meets Bewley: The Permanent Income Hypothesis in General Equilibrium." *American Economic Review*, 93(3): 927–936.

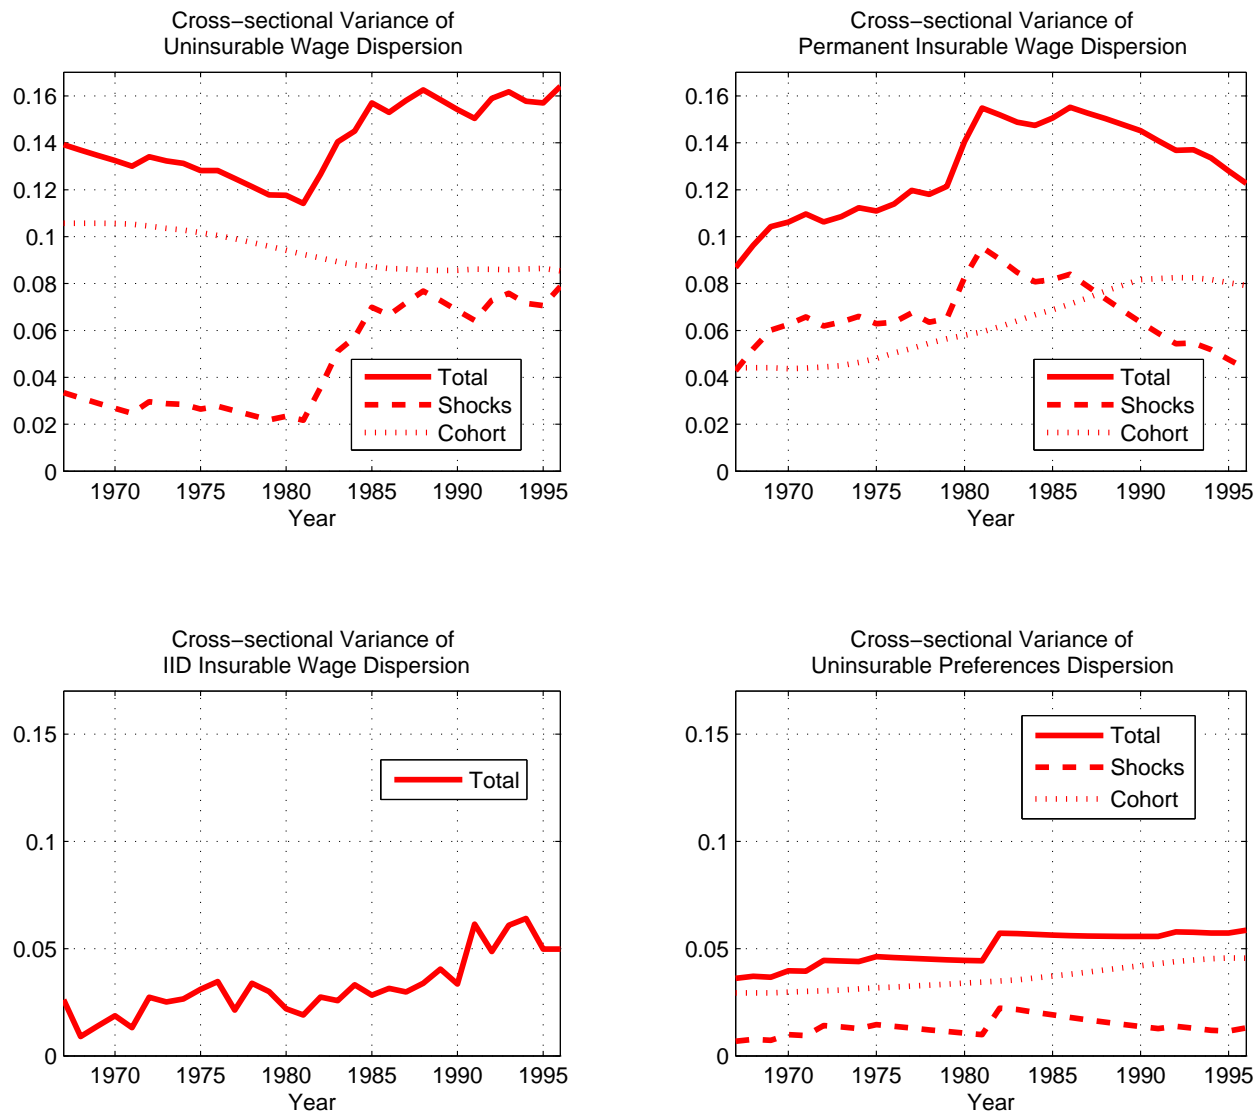


Figure 1: Plot of estimated values of all the time-varying parameters of the model. See Table 4 for the corresponding 90–10 bootstrapped confidence intervals.

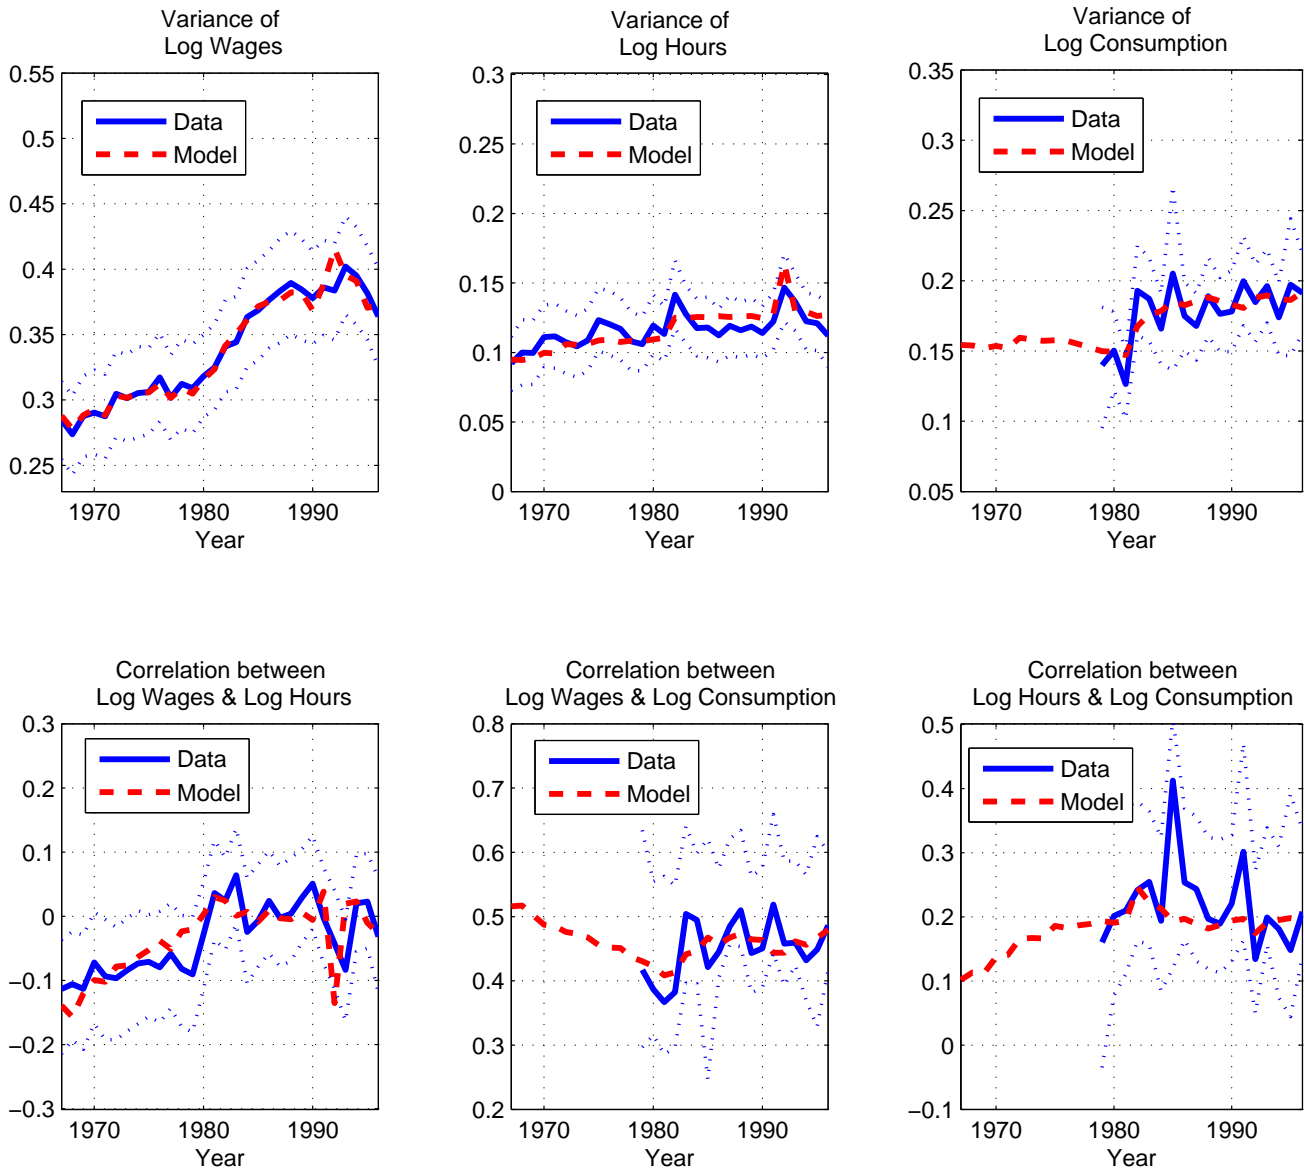


Figure 2: Data and model fit for moments in levels along the time dimension. Dotted lines denote 90–10 bootstrapped confidence intervals for the empirical moments.

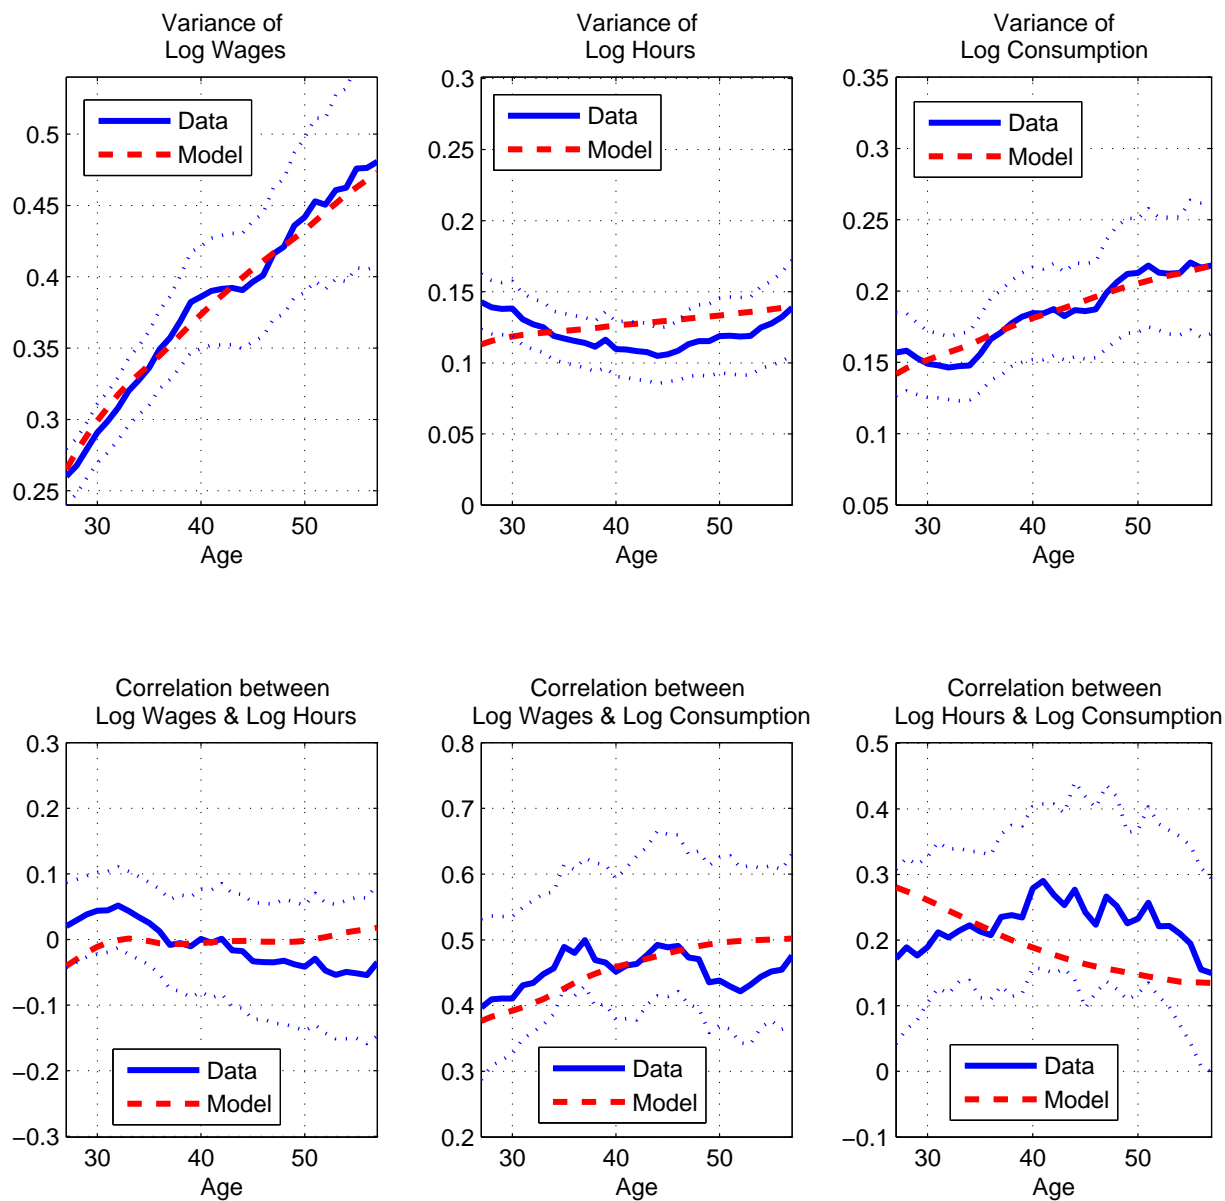


Figure 3: Data and model fit for moments in levels along the age dimension. Dotted lines denote 90–10 bootstrapped confidence intervals for the empirical moments.

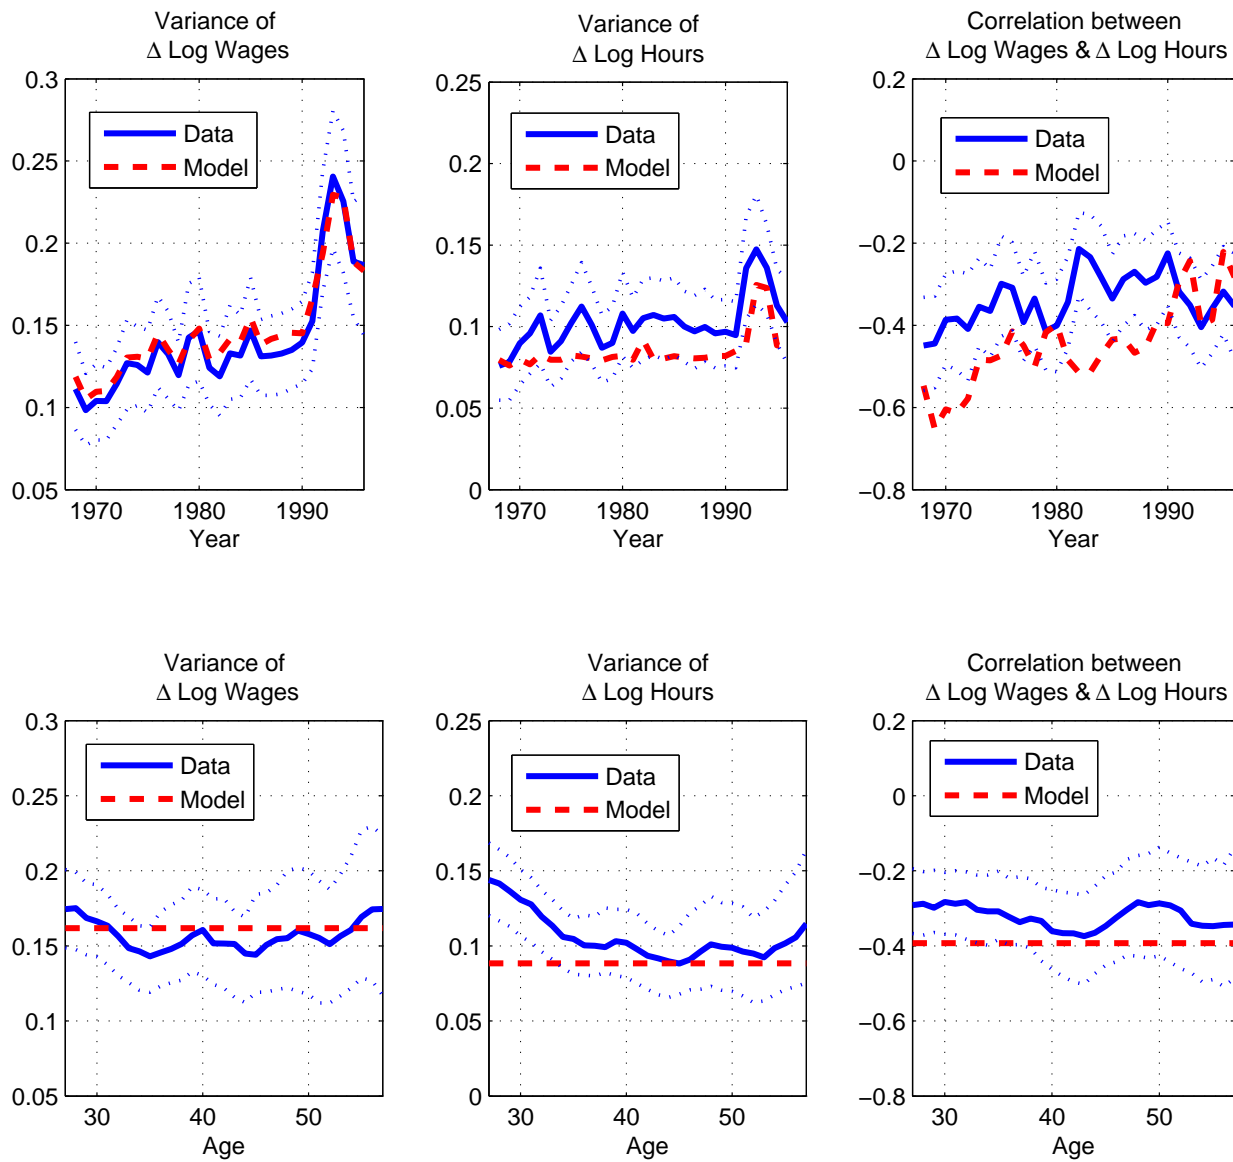


Figure 4: Data and model fit for moments in first-differences along the time dimension (panels in the upper row) and age dimension (panels in the lower row). Dotted lines denote 90–10 bootstrapped confidence intervals for the empirical moments.

Appendix

Not for Publication

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.

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

Planner problem Given no bond trade between islands and complete markets on each island, the first welfare theorem applies, 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 planner problem for an island populated by agents of age a and type (ω^t, χ^t) is then given by

$$\max_{\{c_t^a(s^t), h_t^a(s^t)\}} \int \int \left(\frac{c_t^a(s^t)^{1-\gamma} - 1}{1-\gamma} - \exp((\gamma + \sigma)\varphi_t) \frac{h_t^a(s^t)^{1+\sigma}}{1+\sigma} \right) dF_{\kappa t}^a dF_{\theta t}$$

subject to the resource constraint

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

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,

$$\begin{aligned} c_t^a(s^t)^{-\gamma} &= \lambda_t \\ \exp((\gamma + \sigma)\varphi_t) h_t^a(s^t)^\sigma &= \lambda_t \exp(\alpha_t) \exp(\kappa_t + \theta_t), \end{aligned}$$

where λ_t is the multiplier on the resource constraint at date t . Combining the two conditions gives

$$h_t^a(s^t) = c_t^a(s^t)^{\frac{-\gamma}{\sigma}} \exp\left(\frac{\alpha_t}{\sigma}\right) \exp\left(\frac{1}{\sigma}(\kappa_t + \theta_t)\right) \exp\left(-\frac{(\gamma + \sigma)}{\sigma}\varphi_t\right). \quad (34)$$

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 or θ_t . Using this fact, and substituting (34) into the

planner's island-specific resource constraint, gives

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

Taking logs:

$$\begin{aligned}
\log c_t^a(s^t) &= \frac{1 + \sigma}{\sigma + \gamma}\alpha_t - \varphi_t + \frac{\sigma}{\sigma + \gamma} \log\left(\int \exp\left(\frac{1 + \sigma}{\sigma}(\kappa_t + \theta_t)\right) dF_{\kappa t}^a dF_{\theta t}\right) \\
&= \frac{1 + \sigma}{\sigma + \gamma}\alpha_t - \varphi_t + \mathcal{M}_t^a,
\end{aligned}$$

which is the expression in Proposition 1.

From the logarithm of the intratemporal first-order condition (34):

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

which is the expression in Proposition 1.

Solving for equilibrium prices We now compute prices supporting these allocations. Given the candidate expressions for equilibrium consumption, the marginal rate of substitution between consumption in s^t and in $s^{t+1} = (s^t, (\omega_{t+1}, \chi_{t+1}, \eta_{t+1}, \theta_{t+1}))$ is

$$\begin{aligned}
m_{t+1}^{a+1}(s^t, (\omega_{t+1}, \chi_{t+1}, \eta_{t+1}, \theta_{t+1})) &\equiv \beta \left(\frac{c_{t+1}^{a+1}(s^t, (\omega_{t+1}, \chi_{t+1}, \eta_{t+1}, \theta_{t+1}))}{c_t^a(s^t)} \right)^{-\gamma} \\
&= \beta \frac{\exp\left(-\gamma \frac{1+\sigma}{\sigma+\gamma}(\alpha_t + \omega_{t+1}) + \gamma(\varphi_t + \chi_{t+1}) - \gamma \mathcal{M}_{t+1}^{a+1}\right)}{\exp\left(-\gamma \frac{1+\sigma}{\sigma+\gamma}\alpha_t + \gamma\varphi_t - \gamma \mathcal{M}_t^a\right)} \\
&= \beta \exp(-\gamma \Delta \mathcal{M}_t) \exp\left(\gamma \chi_{t+1} - \gamma \frac{1 + \sigma}{\sigma + \gamma} \omega_{t+1}\right),
\end{aligned}$$

and the marginal rate of substitution between t and $t + 1$, it is straightforward to derive from the above equation. The prices of the one-period contingent claim, and one-period non-contingent bond must then be given by equations (8) and (9), respectively.

Verifying the no bond trade guess We must now verify the initial conjecture that islands do not trade. The marginal rate of substitution between t and $t + 1$ (hence, bond price) implied by the no trade guess 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., $\int \int b_t(s^t) dF_{\kappa t}^a dF_{\theta t} =$

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 history s^t and age a in period t is given by:

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

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:

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

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

PROOF OF PROPOSITION 2

There are several alternative set of moments which identify the model: we pick a simple and intuitive combination. The proof is organized in four sequential steps.

1. The five (sets of) parameters σ , γ , $\{v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$, $\{v_{\omega t}\}_{t=2}^T$ and $\{v_{\chi t}\}_{t=2}^T$ are identified from five within-cohort changes in the macro moments, $\Delta var_t^a(\log \hat{w})$, $\Delta var_t^a(\log \hat{h})$, $\Delta cov_t^a(\log \hat{w}, \log \hat{h})$, $\Delta var_t^a(\log \hat{c})$, and $\Delta cov_t^a(\log \hat{w}, \log \hat{c})$, all available from $t = 2$ to $t = T$.³⁸ Moreover, these parameters can be identified recursively as follows:

³⁸Note that this set of five within-cohort changes constitutes the entire set of independent moments of this type, because $\Delta cov_t^a(\log \hat{h}, \log \hat{c})$ can be expressed as a linear function of other moments: $\Delta var_t^a(\log \hat{c}) = \Delta cov_t^a(\log \hat{h}, \log \hat{c}) + \Delta cov_t^a(\log \hat{w}, \log \hat{c})$. This result follows from the fact that the uninsurable component of log earnings equals log consumption.

(i) The Frisch elasticity is identified by³⁹

$$\frac{1}{\sigma} = \frac{\Delta cov_t^a(\log \hat{w}, \log \hat{h}) + \Delta var_t^a(\log \hat{h}) + \Delta cov_t^a(\log \hat{w}, \log \hat{c}) - \Delta var_t^a(\log \hat{c})}{\Delta var_t^a(\log \hat{w}) + \Delta cov_t^a(\log \hat{w}, \log \hat{h}) - \Delta cov_t^a(\log \hat{w}, \log \hat{c})}. \quad (35)$$

(ii) Given σ , $\{v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$ is identified from the denominator of eq. (35):

$$\Delta var_t^a(\log \hat{w}) + \Delta cov_t^a(\log \hat{w}, \log \hat{h}) - \Delta cov_t^a(\log \hat{w}, \log \hat{c}) = \frac{1 + \sigma}{\sigma} (v_{\eta t} + \Delta v_{\theta t}). \quad (36)$$

(iii) Given σ and $\{v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$, $\{v_{\omega t}\}_{t=2}^T$ is identified from $\Delta var_t^a(\log \hat{w})$, eq. (24).

(iv) Given σ and $\{v_{\omega t}\}_{t=2}^T$, γ is identified from $\Delta cov_t^a(\log \hat{w}, \log \hat{c})$, eq. (29).

(v) Given σ , γ and $\{v_{\omega t}\}_{t=2}^T$, $\{v_{\chi t}\}_{t=2}^T$ is identified from $\Delta var_t^a(\log \hat{c})$, eq. (27).

2. The variances of transitory shocks $\{v_{\theta t}\}_{t=1}^{T-1}$ are identified from the difference between the dispersion in growth rates (“micro moments”) and the growth rate of within-cohort dispersion (“macro moments”):

$$cov_t^a(\Delta \log \hat{w}, \Delta \log \hat{h}) + var_t^a(\Delta \log \hat{h}) - \Delta cov_t^a(\log \hat{w}, \log \hat{h}) - \Delta var_t^a(\log \hat{h}) = \frac{2(1 + \sigma)}{\sigma^2} v_{\theta, t-1}.$$

Combining the sequence $\{v_{\theta t}\}_{t=1}^{T-1}$ with $\{v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^T$ identifies $\{v_{\eta t}\}_{t=2}^{T-1}$. Substituting the value for $v_{\theta, T-1}$ into $(v_{\eta T} + \Delta v_{\theta T})$ from step 1 identifies $(v_{\eta T} + v_{\theta T})$.

3. The following macro moments, available for all t and evaluated for the youngest age group identify the sequence of cohort effects $\{v_{\varphi^0 t}, v_{\alpha^0 t}\}_{t=1}^T$:

$$\begin{aligned} cov_t^0(\log \hat{w}, \log \hat{c}) &= \frac{1 + \sigma}{\sigma + \gamma} v_{\alpha^0 t} \\ cov_t^0(\log \hat{h}, \log \hat{c}) &= v_{\varphi^0 t} + \frac{(1 - \gamma)(1 + \sigma)}{(\sigma + \gamma)^2} v_{\alpha^0 t}. \end{aligned}$$

Then $\{v_{\kappa^0 t}\}_{t=1}^{T-1}$ and $(v_{\kappa^0 T} + v_{\theta T})$ are identified from

$$cov_t^0(\log \hat{w}, \log \hat{h}) + var_t^0(\log \hat{h}) = v_{\varphi^0 t} + \frac{(1 - \gamma)(1 + \sigma)}{(\sigma + \gamma)^2} v_{\alpha^0 t} + \frac{1 + \sigma}{\sigma^2} (v_{\kappa^0 t} + v_{\theta t})$$

4. Finally, measurement error can be identified from moments in levels, for example those corresponding to the youngest age group:

$$\begin{aligned} cov_t^0(\log \hat{w}, \log \hat{h}) &= \frac{1 - \gamma}{\sigma + \gamma} v_{\alpha^0 t} + \frac{1}{\sigma} (v_{\kappa^0 t} + v_{\theta t}) - v_{\mu h} \\ var_t^0(\log \hat{w}) &= v_{\alpha^0 t} + (v_{\kappa^0 t} + v_{\theta t}) + v_{\mu y} + v_{\mu h} \\ var_t^0(\log \hat{c}) &= v_{\varphi^0 t} + \left(\frac{1 + \sigma}{\sigma + \gamma}\right)^2 v_{\alpha^0 t} + v_{\mu c}. \end{aligned}$$

³⁹If $\gamma = 1$ the relationship simplifies to

$$\frac{1}{\sigma} = \frac{\Delta cov_t^a(\log \hat{w}, \log \hat{h})}{\Delta var_t^a(\log \hat{w}) - \Delta cov_t^a(\log \hat{w}, \log \hat{c})}$$

so the Frisch elasticity is identified by the increase in the covariance between wages and hours, relative to the increase in the variance of insurable wage risk.

This concludes the proof.⁴⁰ ■

PROOF OF COROLLARY TO PROPOSITION 2

We prove the three parts of the corollary one by one.

(A) Follow first step 1 of the main proof. Then, for any cohort a identify $\{v_{\alpha t}^a\}_{t=1}^T$ from eq. (22), identify $\{v_{\varphi t}^a\}_{t=1}^T$ from eq. (21), identify $\{v_{\eta t}^a + v_{\theta t}\}_{t=1}^T$ from $cov_t^a(\log \hat{h}, \log \hat{y}) = var_t^a(\log \hat{h}) + cov_t^a(\log \hat{h}, \log \hat{w})$, then identify $v_{\mu h}$ from eq. (19), identify $v_{\mu y}$ from eq. (17), and identify $v_{\mu c}$ from eq. (20). Finally, repeat step 3 in the main proof for all periods $t = 1, \dots, T$ to identify $\{v_{\varphi^0 t}, v_{\alpha^0 t}, (v_{\kappa^0 t} + v_{\theta t})\}_{t=1}^T$.

(B) The parameters $\{\sigma, \gamma, v_{\mu y}, v_{\mu h}, v_{\mu c}\}$ and the parameters for $t > t^*$ are identified as in the main proof using data from $t = t^*, \dots, T$. The sequence $\{v_{\theta t}\}_{t=1}^{t^*}$ is identified as in step 2 of the main proof. Given values for (σ, γ) , equations (24)–(26) identify $\{v_{\chi t}, v_{\omega t}, v_{\eta t} + \Delta v_{\theta t}\}_{t=2}^{t^*}$. The moments $var_t^0(\log \hat{w})$ and $cov_t^0(\log \hat{w}, \log \hat{h})$ in equations (17)–(19) jointly identify $\{v_{\alpha^0 t}, v_{\kappa^0 t}\}_{t=1}^{t^*}$. Finally, $\{v_{\varphi^0 t}\}_{t=1}^{t^*}$ is identified from $var_t^0(\log \hat{h})$ in eq. (18).

(C) Data for different cohorts is used only to identify different cohort effects (step 3). This concludes the proof. ■

⁴⁰Note that steps 3 and 4 use all six macro moments which, in levels, are all independent.

Table 5: Parameter Estimates: Alternative Models

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
σ	2.624	2.710	2.651	3.212	4.537	1.731	2.644	2.741	2.223
γ	2.016	2.319	2.020	2.377	2.113	1.984	1.952	1.918	1.920
$v_{\mu h}$	0.036	0.036	0.036	0.025	0.026	0.039	0.036	0.034	0.037
$v_{\mu y}$	0.0000	0.000	0.000	0.000	0.010	0.000	0.000	0.029	0.000
$v_{\mu c}$	0.033	0.068	0.034	0.028	0.028	0.095	0.034	0.029	0.035
\bar{v}_{φ^0}	0.045	0.039	0.045	0.057	0.048	0.000	0.053	0.047	0.047
\bar{v}_{χ}	0.001	0.001	0.001	0.001	0.001	0.000	0.000	0.001	0.001
\bar{v}_{α^0}	0.083	0.064	0.081	0.086	0.080	0.080	0.083	0.079	0.079
\bar{v}_{ω}	0.004	0.005	0.005	0.005	0.004	0.005	0.004	0.004	0.005
\bar{v}_{κ^0}	0.081	0.099	0.083	0.084	NI	0.085	0.078	0.081	0.082
\bar{v}_{η}	0.004	0.004	0.003	0.005	NI	0.003	0.004	0.006	0.004
\bar{v}_{θ}	0.032	0.033	0.033	0.035	NI	0.028	0.032	0.000	0.031

Notes: (1) baseline model (see Table 1), (2) consumption definition excludes durables, (3) constant $v_{\alpha^0}, v_{\kappa^0}, v_{\varphi^0}$, (4) no moments in first-differences (NI = not identified), (5) includes transitory preference shocks (set $v_{\mu y} = 0.01$), (6) no initial preference dispersion, $v_{\varphi^0} = 0$, (7) no life-cycle preference shocks, $v_{\chi} = 0$, (8) no transitory wage shocks, $v_{\theta} = 0$, (9) baseline model with weights proportional to number of observations.