

Testing Full Consumption Insurance in the Frequency Domain

Paulo Santos Monteiro*
University of Warwick, United Kingdom

October 24, 2008

Abstract

Full consumption insurance implies that consumers are able to perfectly share risk by equalizing state by state their inter-temporal marginal rates of substitution in the presence of idiosyncratic endowment shocks. In this paper I test the implications of full consumption insurance using band spectrum regression methods. I argue that moving to the frequency domain provides a possible solution to many difficulties tied to tests of perfect risk sharing. In particular, it provides a unifying framework to test consumption smoothing, both over time and across states of nature. Full consumption insurance is soundly rejected at business cycle frequencies.

JEL Classifications: D1, E21.

Keywords: Consumption insurance, Idiosyncratic risk, Frequency domain

*Contact Details: Department of Economics, University of Warwick, Coventry CV4 7AL, United Kingdom. Tel: +44 24761 50587. Email Address: p.santos-monteiro@warwick.ac.uk. Financial support by the Portuguese Science and Technology Foundation is gratefully acknowledged.

1 Introduction

“Consumption insurance studies households’ ability to smooth consumption over states of nature; the permanent income hypothesis (PIH) studies their ability to smooth consumption over time. Households may be able to do each and not the other”. This quote can be found in Cochrane (1991) a seminal paper which tests the full consumption insurance hypothesis using household level data. However, the literature has seldom attempted to distinguish between the predictions of the full risk sharing model and the ones of the Life-cycle/PIH model. In particular, tests of full consumption insurance have traditionally been carried under the maintained assumption of absence of liquidity constraints and consequently perfect consumption smoothing across time.

However, if the Life-cycle/PIH hypothesis is not valid, individual consumption is likely to respond to changes in individual income, even in the presence of perfect risk sharing. For example, young consumers, whose income is expected to rise in the future, will wish to borrow. In the presence of liquidity constraints, they will not be able to do so, and as a consequence, their consumption will track their income. This is caused by a failure to smooth consumption across time and therefore will happen even if full risk sharing is possible; however an econometrician who does not distinguish between the deterministic life-cycle and the random idiosyncratic components of income would reject full consumption insurance. This paper aims at providing a unifying approach to test consumption smoothing both over time and across states of nature.

In particular, I test the implications of full consumption insurance using band spectrum regression methods. This allows me to distinguish between both consumption smoothing across time and across states of nature. This is because liquidity constraints, which prevent consumption smoothing across time, translate into a low frequency parallel between consumption growth and income growth. On the other hand lack of risk sharing, which prevents consumption smoothing across states of nature, translates into a high frequency parallel between consumption and income growth. Band spectrum regression allows to uncover features of the data at different frequencies and thus to tell apart the two models.

Spectral linear regression methods have first been applied to econometrics by Engle (1974) who discusses estimation of linear regression models in the context of measurement error and seasonality and, in particular, suggests as an application a test of the permanent income hypothesis. This technique allows models to be evaluated over particular frequencies, such as business cycles, seasonal frequencies or long horizons. Rather than simply rejecting or accepting the empirical fit of a model, it becomes possible to identify the frequencies in which the model performs well and the ones where it is rejected.

Moreover, since band spectrum regression possesses standard small sample properties, I argue that moving to the frequency domain provides a possible solution to some difficulties tied to tests of perfect risk sharing. In particular, the band spectrum regression approach may be suitable for the estimation of errors-in-variable models. This is because band spectrum regression allows for models to be estimated over frequencies where measurement error is a less pervasive problem, much in the same way as moving average filters are used to eliminate high frequency noise components from measured income in tests of the PIH, for example. Thus, provided that frequencies at which measurement error is pervasive are excluded, the full consumption insurance test proposed in this paper will be robust to the presence of measurement error.

I use data from the Panel Study of Income Dynamics (PSID) to examine the degree of consumption smoothing among North American households. Many authors have tested the complete markets model using North American household level data and findings are ambiguous. Mace (1991), using data from the Consumer Expenditure Survey (CEX), finds mixed evidence in favor of the full risk sharing assumption. Cochrane (1991), in the same spirit but using data from the PSID, measures the correlation between individual consumption growth and some indicators of households resources other than income and also finds mixed evidence. Recently, Guvenen (2007) using PSID data, rejects full consumption insurance among stockholders, but fails to reject it among non-stockholders.

Altug and Miller (1990), using PSID data, attempt to address the problem of nonseparability between food consumption and leisure, which they claim might

bias other tests, and fail to reject full risk sharing. However, Hayashi et al. (1996), also using PSID data, reject inter as well as intra-family risk sharing and argue that the results of Altug and Miller are explained by the lack of power of their test against the self-insurance hypothesis. The paper by Hayashi et al. (1996) is particularly interesting because, to my knowledge, it is the only paper where filtering the data has been suggested as a method to discriminate between risk sharing and self-insurance. These authors suggest taking long time differences of consumption and income as a method to construct a test of the complete markets assumption which has power against the alternative of self-insurance as well as to adjust for measurement error.

Moreover, the present paper is also related to the literature which tests the Life-Cycle/PIH using micro panel data sets and in particular attempt to measure the response of consumption to income shocks of different persistence. Prominent examples of this stream of work includes, Hall and Mishkin (1982), Bernanke (1984), Altonji and Siow (1987), and Attanasio and Weber (1995).

The tests of full consumption insurance described above and also the one performed in this paper fundamentally require estimating a regression equation which includes as dependent variable the changes in household consumption and on the right-hand side includes aggregate level variables and idiosyncratic endowment variables. Under the null hypothesis of full insurance, the latter should not significantly help in predicting changes in individual consumption. However, the single most important difficulty faced by studies of full insurance is to find idiosyncratic variables which are good proxies for individual endowments and are orthogonal to the error term of the estimated reduced form regression equation. Unfortunately, one variable which is unlikely to satisfy this requirement is household income because the error term possibly incorporates changes in preferences which simultaneously affect the inter-temporal allocation of consumption chosen by households as well as the inter-temporal allocation of leisure and consequently household income.

To try to solve this problem, I have adopted a two-step instrumental variables procedure. I first regress the household income on a number of variables which are likely to be exogenous such as days of work lost because of strikes or unemployment

and hourly average earnings, for example. In the second step, I use the predicted growth rate of income as a continuous scale proxy for idiosyncratic endowment changes and I implement the test of full insurance applying the band spectrum regression method. Because of the nature of the null hypothesis being tested (under full insurance the impact of idiosyncratic shocks is zero), the tests statistics should be asymptotically consistent, despite the regressors being estimated which leads to a downward bias in the estimation of the regression coefficient standard errors.¹

The findings of the paper are not supportive of the full risk sharing, complete markets hypothesis, but they are broadly consistent with the self-insurance hypothesis. In particular, full consumption insurance is soundly rejected at business cycle frequencies. Importantly, a rejection of consumption insurance at business cycle frequencies suggests that the representative agent construct may be an inappropriate paradigm for business cycle models. However, consumption is found to respond more strongly to long lasting income shocks than to rapid, high frequency, shocks in accordance with the permanent income hypothesis.

The remainder of the paper is organized as follows. In section 2 I build a framework to test the complete markets assumption and in particular I show that given the standard inter-temporal consumer choice model, consumption growth can be represented by a well defined factor structure. Section 3 describes the econometric methodology employed, in particular the band spectrum regression method. In section 4 I briefly describe the data used in the paper and I discuss some issues involving the empirical implementation. Finally, in section 5 I report my results and section 6 concludes.

2 Testing Full Consumption Insurance

Assuming complete markets and absence of private information, the solution to the social planner problem is the same as the resource allocation rule that solves the decentralized competitive equilibrium problem. Therefore, the conditions necessary to solve the planning problem provide testable implications of full consumption insurance. The derivation of the testable implications of full consumption insurance

¹see Pagan (1984)

which follows is well known and in particular can be found in Cochrane (1991). It is included here for completeness.

Consider the planning problem for an endowment economy with N households

$$\begin{aligned} \max \sum_{j=1}^N \lambda^j \sum_{t=0}^{\infty} \sum_{s_t} (\beta^j)^t \pi(s_t) u(c^j(s_t), b_t^j) \\ \text{s.t.} \quad \sum_j c^j(s_t) \leq \sum_j y^j(s_t), \end{aligned} \quad (2.1)$$

Where j indexes households, t indexes time, λ^j is household j 's Pareto weight, β^j is household j 's subjective discount factor, $\pi(s_t)$ is the probability that state s_t occurs, and $c^j(s_t)$ and $q^j(s_t)$ are, household j 's consumption and endowment in state s_t , respectively. Finally, b_t^j is a household specific parameter that captures preference heterogeneity. The first-order conditions for this problem are

$$\beta^j \lambda^j u_c(c_t^j, b_t^j) = \mu(s_t) \quad \forall \quad s_t. \quad (2.2)$$

Notice that $\mu(s_t)$, the Lagrange multiplier normalized by the probability of the state of the world s_t occurring, does not depend on j and therefore each individual's optimal consumption path is independent of her idiosyncratic endowment component. Consequently, if markets are complete and there is full consumption insurance, the discounted growth of marginal utility must be perfectly cross-sectionally correlated

$$\beta^j \frac{u_c(c_{t+1}^j, b_{t+1}^j)}{u_c(c_t^j, b_t^j)} = \frac{\mu_{t+1}}{\mu_t}. \quad (2.3)$$

Assuming that households have power utility functions with risk aversion coefficient (ρ) common across individuals, modeling preferences heterogeneity through multiplicative shocks, $u(c_t^j, b_t^j) = b_t^j \frac{(c_t^j)^{1-\rho}}{1-\rho}$, and replacing into (2.2) we obtain the following relationship

$$\log(\beta^j) + \log(b_t^j) + \log(\lambda^j) - \rho \log(c_t^j) = \log(\mu_t). \quad (2.4)$$

Aggregating over the N households, yields

$$\sum_j \frac{1}{N} \log(\mu_t) = \sum_j \rho \frac{\log(c_t^j)}{N} + \sum_j \frac{\log(\beta^j)}{N} + \sum_j \frac{\log(b_t^j)}{N} + \sum_j \frac{\log(\lambda^j)}{N}. \quad (2.5)$$

Finally, substituting (2.5) into (2.4) we obtain

$$\log(c_t^j) = \log c_t^A + \phi^j + \omega_t^j, \quad (2.6)$$

where $\log c_t^A \equiv \sum_j \frac{\log(c_t^j)}{N}$ is the logarithm of aggregate consumption and where

$$\begin{aligned} \phi^j &= \frac{1}{\rho} \left[\log(\beta^j) - \frac{1}{N} \sum_j \log(\beta^j) \right] + \frac{1}{\rho} \left[\log(\lambda_j) - \frac{1}{N} \sum_j \log(\lambda^j) \right], \\ \omega_t^j &= \frac{1}{\rho} \left[\log(b_t^j) - \frac{1}{N} \sum_j \log(b_t^j) \right]. \end{aligned}$$

Taking the first difference of equation (2.6), we obtain the following factor structure representation for the growth rate of individual consumption

$$\Delta \log c_t^j = \Delta \log c_t^A + \Delta \omega_t^j, \quad (2.7)$$

where $\Delta \log c_t^j$ and $\Delta \log c_t^A$ are the growth rates of individual j 's consumption and aggregate consumption, respectively. Notice that any household fixed effects, captured by ϕ^j , are removed when the model is taken in first differences.

Tests of perfect risk-sharing are based on the proposition that, given full consumption insurance, individual consumption growth should be perfectly cross-sectionally correlated. Hence, individual consumption growth should respond to aggregate risk but not to idiosyncratic shocks captured, for example, by variations in household income or employment status. Therefore, equation (2.7) suggests estimation of the following regression equation

$$\Delta \log c_t^j - \Delta \log c_t^A = \alpha'_1 \Delta \log c_t^A + \alpha'_2 x_t^j + \xi_t^j, \quad (2.8)$$

where $\Delta \log c_t^j - \Delta \log c_t^A$ is household j 's non-durable consumption growth net of aggregate consumption growth and x_t^j is a $q \times 1$ vector of household specific variables

meant to capture idiosyncratic endowment shocks. Finally, ξ_t^j is defined as $\delta'v_t^j + \epsilon_t^j$, where the first term captures observable changes in preferences and the second term, ϵ_t^j , captures unobservable changes in preferences as well as measurement error. Consequently, unobservable determinants of individual consumption growth will be captured in the regression by the error term ϵ_t^j .

The variables included in x_t^j are meant to capture exogenous variations in household endowments, which are not expected to affect the growth rate of consumption if full insurance is implemented. Hence, to test full consumption insurance, the right-hand side variables included in x_t^j should be independent of changes in individual preferences, captured by the error term. Equation (2.8) yields a relatively straight forward test of full consumption insurance, which has been explored before in work by Mace (1991) and Townsend (1994). Thus, equation (2.8) can be recast in the following form

$$y_t^j = \alpha' z_t^j + \delta' v_t^j + \epsilon_t^j, \quad (2.9)$$

where $y_{jt} \equiv \Delta \log c_t^j - \Delta \log c_t^A$ and

$$z_{jt} \equiv \begin{bmatrix} \Delta \log c_t^A \\ x_t^j \end{bmatrix}, \quad \alpha \equiv \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix}.$$

Finally, v_t^j is a $p \times 1$ vector of household demographic variables, meant to control for observable shifts in preferences caused, for example, by changes in household size or in the number of children in the household.

Provided that the exogeneity assumptions are satisfied, full risk-sharing requires the simple testable restriction $\alpha \equiv \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = 0$.

3 Methodology: Band Spectrum Regression

Let the data set contain T observations on each household. The complex finite Fourier transform is based on the $T \times T$ matrix W , in which each element (k, s) is

given by

$$w_{k,s} = \frac{1}{\sqrt{T}} e^{i s \theta_k} \quad s = 0, 1, \dots, T - 1, \quad (3.1)$$

where $\theta_k = \frac{2\pi k}{T}$, $k = 0, 1, \dots, T - 1$ and $i = \sqrt{-1}$. Pre-multiplying the vector of observations in the regression equation (2.9) by W , produces a finite Fourier transform of the time domain vectors, which yields the model

$$\tilde{y}_t^j = \alpha' \tilde{z}_t^j + \delta' \tilde{v}_t^j + \tilde{\epsilon}_t^j, \quad (3.2)$$

where $\tilde{y}_t^j = W y_t^j$, $\tilde{z}_t^j = W z_t^j$, $\tilde{v}_t^j = W v_t^j$ and $\tilde{\epsilon}_t^j = W \epsilon_t^j$.

Model (3.2) is a standard linear regression model made of T independent observations on \tilde{y} conditioned on \tilde{x} , each of which corresponds to a different frequency. The elements are amplitudes and phases of sine waves of different frequencies which reflect the importance of each frequency component in the original time series. If the disturbance vector in (2.9) are spherical and zero mean, that is $E[\epsilon] = 0$ and $E[\epsilon\epsilon'] = \sigma^2 I_T$, then the transformed disturbance vector, $\tilde{\epsilon}$, will have identical properties. This follows because the matrix W is unitary, that is, $WW' = I$, where W' is the transpose of the complex conjugate of W . Given the standard exogeneity assumptions and assuming spherical disturbances, application of OLS to (3.2) yields the best linear unbiased estimator (BLUE) of α . Thus, band spectrum regression possesses standard small sample properties. This estimator is, of course, identical to the OLS estimator in (2.9), a result which follows directly from the property that W is a unitary matrix.

However, when the relationship implied by (3.2) is only assumed to hold for certain frequencies, band spectrum regression allows to test a restricted version of the model in which some frequencies are ignored. This may be carried out by omitting the observations in (3.2) corresponding to the remaining frequencies. Since the variables in (3.2) are complex, Engle suggests performance of an inverse Fourier transform in order to convert the variables into real terms again, thereby making the use of standard regression routines feasible. This is done by first defining a $T \times T$ matrix A which has zeros everywhere except in the positions on the leading diagonal corresponding to the included frequencies and next taking the inverse

fourier transformed of the fourier transform times A .² Applying this method to the regression equation in (2.9) yields

$$y_t^{j*} = \alpha' z_t^{j*} + \delta' v_t^{j*} + \epsilon_t^{j*}, \quad (3.3)$$

where $y_t^{j*} = W' A \tilde{y}_t^j$, $z_t^{j*} = W' A \tilde{z}_t^j$, $v_t^{j*} = W' A \tilde{v}_t^j$ and $\epsilon_t^{j*} = W' A \tilde{\epsilon}_t^j$.

The model given by (3.3) should be treated as a standard linear regression model, except that if A is not full rank or, equivalently, if some frequencies are excluded, the model's degrees of freedom are only $T' - (1 + q + p)$ instead of $T - (1 + q + p)$, where T' is the number of included frequencies. Thus, an unbiased estimator of σ^2 is given by $\hat{\sigma}^2 = \frac{\sum_t \epsilon_t^{j*} \epsilon_t^{j*}}{T' - (1 + q + p)}$ and the sampling statistics must be adjusted accordingly. Let the transformed observation on each household be collected on the $T \times 1$ column vectors $y^{j*} = (y_1^{j*}, \dots, y_T^{j*})'$, $z^{j*} = (z_1^{j*}, \dots, z_T^{j*})'$, $v^{j*} = (v_1^{j*}, \dots, v_T^{j*})'$, $\epsilon^{j*} = (\epsilon_1^{j*}, \dots, \epsilon_T^{j*})'$, with $j = 1, \dots, N$ the household unit. The econometric methodology applied in this paper relies on pooling the individual observations, transformed to exclude the frequencies which are not of interest. Thus, defining

$$Y^* = \begin{bmatrix} y^{1*} \\ \vdots \\ y^{N*} \end{bmatrix}, \quad Z^* = \begin{bmatrix} z^{1*} \\ \vdots \\ z^{N*} \end{bmatrix}, \quad V^* = \begin{bmatrix} v^{1*} \\ \vdots \\ v^{N*} \end{bmatrix}, \quad E^* = \begin{bmatrix} \epsilon^{1*} \\ \vdots \\ \epsilon^{N*} \end{bmatrix},$$

I will perform a pooled panel regression by estimating the linear regression model given by

$$Y^* = Z^* \alpha + V^* \delta + E^*. \quad (3.4)$$

Accordingly, the null hypothesis of full consumption insurance corresponds to the linear restriction $\alpha = 0$. If the error term, capturing measurement error and changes in preferences shifts, is homoskedastic and uncorrelated across time and across households, ordinary least squares (OLS) estimates and standard t-tests can be performed. However, the model's degrees of freedom are given by $NT' - (1 + q + p)$ instead of the usual $NT - (1 + q + p)$, where T' is the number of included frequencies

²The estimator will only be real if both sines and cosines are included at each frequency. That is, if frequency component k is included, than $T - k$ must be included as well (Engle [1974]).

and an unbiased estimator of σ^2 is given by

$$\hat{\sigma}^2 = \frac{\sum_{t=1}^T \sum_{j=1}^N \hat{\epsilon}_t^{j*} \hat{\epsilon}_t^{j*}}{NT' - (1 + q + p)}. \quad (3.5)$$

4 Data and Empirical Implementation Issues

4.1 The Panel Survey of Income Dynamics

This paper uses household level data from the Panel Study of Income Dynamics (PSID). The PSID is a longitudinal study of nearly 8000 US households, following the same families and individuals since 1968. The original PSID sample consisted of two subsamples, a representative cross-section of 3000 U.S. families and a subsample of 2000 low-income families sampled from the Survey of Economic Opportunity (SEO). I drop the SEO subsample in order to work with a representative sample of the U.S. population. Thereafter, both the original households and their split offs have been interviewed each year. The survey includes a variety of socioeconomic variables, including age, education, family structure and earnings. Let t be the calendar year ($1973 + t$). I have used information on each household $j \in (1, 2, \dots, N)$ income and consumption between 1974 and 1986, corresponding to the calendar years $t = 1, 2, \dots, 13$, to build a balanced panel of observations.

The sample selection procedure, fully detailed in the Appendix, yields $N = 966$ households. Descriptive statistics about demographic characteristics of the households included in the panel are shown in table 1. The most comprehensive measure of consumption which is available from the PSID and the one used in this paper is total expenditure in food which is defined as food expenditure at home, plus food stamps, plus meals away from home. The measure of aggregate consumption used is the average of individual total food expenditure, taken over the N households. An important aspect of the PSID data is that the earnings questions are retrospective. The interviews are conducted around March and many questions and in particular those about family income refer to the previous calendar year. Thus, I date the observations according to the year corresponding to the earnings, instead of the year of the interview.

Table 1: Household demographic characteristics

Variable	Mean	Std. Dev.	Min.	Max.
1974 age of head	32.09	8.59	20	52
1974 age of wife	30.36	8.81	16	72
Household size	3.45	1.43	1	10
Children under 18	1.36	1.26	0	8
<u>Percentage of individual-year pairs</u>				
Household head is male	93.69 %			
Married households	85.56 %			
Household owns house	76.56 %			
Head finished high school	56.58 %			
Head finished college	26.96 %			

Note: The number of households included in the sample is 966. Sample period is 1974-1987.

However, the timing of the survey questions on food expenditure is much less clear (see Hall and Mishkin [1982]; Blundell et al. [2002]). Households are asked to report how much they spend on average in a given week and the answer to this question is used to make inference about household yearly expenditure on food. Since interviews are usually conducted from March onwards, it has been argued that people report their food expenditure for an average week around that period, rather than for the previous calendar year as for income.³ Consequently, and following Altug and Miller (1990), I have defined food expenditure for year t as 25% of food expenditure reported in survey year t plus 75% of food expenditure reported in survey year $t + 1$.

The strongest evidence in favor of this procedure comes from Hall and Mishkin (1982) who assume that new information about income, that the family uses to decide on consumption dated in year t , includes a fraction Φ of new information on

³According to Hayashi et al. (1996) most interviews are conducted in March, April, or May.

income that is not recorded by the survey until the following year. These authors estimate Φ , the fraction of advanced information, to be equal to 0.25.

Crucially, the use of band spectrum regression reduces the eventual bias resulting from possibly over-estimating the amount of information available to households when they make their consumption decisions. This is because, by removing high frequency noise and focusing on long lasting shocks (periodicities greater than three years), I make sure that the consumption and income fluctuations captured are “contemporaneous”. This is another important advantage of the method proposed in this paper.

4.2 Measuring Endowment Shocks

The variables used to measure changes in household endowments were, Total Household Income Net of Transfers, Average Hourly Earnings of the Head, Head’s Annual Hours of Unemployment plus the Head’s Annual Hours of Work Lost to Strikes and a dummy variable which takes value one when the household was forced to move in response to outside events (e.g.: evictions; health reasons; divorce). All dollar valued variables were deflated using the food price component of the CPI. Summary statistics on all these variables as well as on the different components of food expenditure (from now on dubbed consumption) are shown in table 2. The reason why I have chosen to work with the household income net of transfers is because the transfer component of income includes some of the payments that implement a consumption-insured allocation, such as worker’s compensation, child support and help from relatives and are therefore state contingent payments instead of strictly exogenous idiosyncratic endowment shocks.

However, as attractive as it may seem as a proxy for household endowment, the household income net of transfers is unlikely to satisfy the orthogonality conditions required for obtaining consistent estimates of the regression coefficients of interest. This is because unobservable changes in household preferences, which are part of the error term, such as an increase in the taste for leisure, simultaneously change the household income and the consumption decisions, also under the null hypothesis of full insurance.

Table 2: Household summary statistics

Variable	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986
<u>Endowments</u>													
Family income	27617	28953	30900	31276	31877	33483	35294	36343	37397	39493	42860	44311	44973
net of transfers	(17791)	(18679)	(19175)	(18196)	(17891)	(21468)	(32236)	(26329)	(26806)	(29124)	(38505)	(36076)	(37529)
Head avg	9.83	10.38	10.73	10.92	10.94	11.44	11.39	12.16	12.53	12.94	13.58	13.78	13.88
hourly earnings	(5.86)	(7.13)	(6.50)	(7.23)	(8.00)	(8.67)	(6.24)	(8.13)	(8.51)	(8.66)	(9.51)	(9.66)	(10.46)
Hours lost to strike	60	31	59	57	36	28	42	59	72	63	45	49	43
or unemployment	(191)	(251)	(215)	(211)	(158)	(142)	(172)	(211)	(254)	(233)	(177)	(197)	(176)
% Involuntary move	2.9%	1.7%	2.6%	1.8%	2.2%	1.8%	2.0%	1.9%	1.2%	1.3%	1.2%	1.9%	1.7%
<u>Consumption</u>													
Expenditure Meals	3551	3705	3676	3707	3741	3715	3743	3737	3750	3751	3807	3745	3651
at home	(1715)	(1866)	(1829)	(1794)	(1868)	(1832)	(1828)	(1797)	(1846)	(1854)	(1873)	(1871)	(1841)
Expenditure Meals	633	774	832	858	865	857	851	899	912	1027	1062	1010	1068
away from home	(743)	(910)	(962)	(927)	(970)	(975)	(927)	(978)	(925)	(1039)	(1106)	(998)	(1048)
Value of	28	32	23	18	10	13	23	21	29	21	13	14	12
food stamps	(224)	(230)	(181)	(151)	(78)	(120)	(172)	(166)	(226)	(174)	(121)	(155)	(133)

Note: Summary statistics for household consumption and endowment variables. In parenthesis are standard deviations. Dollar valued variables are in 1983 dollars.

In contrast to the total family income net of transfers, the other three variables can reasonably be expected to be independent from shifts in individual preferences and, consequently, orthogonal to the error term.⁴

Therefore, the strategy adopted in the empirical implementation of the full insurance test is a two-step instrumental variables procedure. I first estimate a regression model for household log income that includes as explanatory variables, apart from demographic control variables, the three instrumental variables mentioned above as well as two interaction terms, one between the head wage rate and the number of days lost to strike and unemployment, and another one between the wage rate and the involuntary move dummy variable. In the second step I use the changes in the predicted level of log income to proxy for idiosyncratic endowment shocks and estimate model (2.9) in order to implement the full consumption insurance test. Hence, the two equations estimated are

$$\begin{aligned}\log INC_t^j &= \lambda' INST_t^j + e_t^j, \\ \Delta \log c_t^j - \Delta \log c_t^A &= \alpha_1' \Delta \log c_t^A + \alpha_2' \Delta \log \widehat{INC}_t^j + \delta' \tilde{v}_t^j + \epsilon_t^j.\end{aligned}$$

Unfortunately, full information IV methods are not feasible because of the band spectrum regression procedure. Therefore, the two-step procedure fails to account for the fact that the generated regressors have been estimated, when the second-step coefficients and the standard errors are calculated. However, given the nature of the null hypothesis being tested (full insurance requires $\alpha = 0$), our test statistics are asymptotically consistent (Pagan [1984], Theorem 3).

The strategy just described, resembles the procedure suggested by Altonji and Siow (1987) who identify endowment shocks through other reported measures of income such as hours of involuntary inactivity and wage rates, and also follow a two-step procedure. In particular, an important maintained assumption, which is also made by Altonji and Siow (1987), is that the income determinants used are exogenous with respect to unobservable changes in the marginal utility of consumption. A similar assumption is made in Dynarski and Gruber (1997). The results of the first

⁴I will label the exogenous variables: Involuntary Inactivity which equals Head's Annual Hours of Unemployment plus the Head's Annual Hours of Work Lost to Strikes; Involuntary Move which equals one when the household was forced to move in response to outside events and zero elsewhere; and the Wage Rate which equals the Log of Average Head's Hourly Earnings.

Table 3: First stage regression

Variable	Coefficient	(Std. Err.)
Log of Head Hourly Earnings (Wage Rate)	0.6653**	(0.0069)
Involuntary Inactivity (Days Lost)	-0.0004**	(0.0000)
Involuntary Move	-0.0572	(0.0922)
Wage Rate \times Days Lost	-0.0001**	(0.0000)
Wage Rate \times Involuntary Move	-0.0201	(0.0417)
Household Size	0.0360**	(0.0024)
Female Head	-0.4537**	(0.0142)
Age of Head	0.0103**	(0.0003)
High School Education	0.1237**	(0.0095)
College Education	0.2230**	(0.0112)
Head is Black	-0.0585**	(0.0098)
Head is Hispanic	0.0055	(0.0224)
Sample Size		12558
Adjusted R^2		0.66
F-stat		2003.21
Significance levels : † : 10% * : 5% ** : 1%		

Note: The dependent variable is the logarithm of Total Household Income Net of Transfers.

The specification includes an intercept.

stage regression are shown in table 3.

4.3 Measurement Error

When studying the relationship between changes in consumption and changes in income, the presence of measurement error in income might bias the estimates in various ways. One might expect the presence of measurement error to result in a downward bias in the estimated response of consumption to income. On the other hand Altonji and Siow (1987) argue that previous work by Hall and Mishkin (1982) might have overestimated the impact of transitory income shocks on consumption because the presence of measurement error can make income to

appear more transitory than it actually is.⁵ An important application of band spectrum regression, first suggested in Engle (1974), is the estimation of linear regression models where the right-hand side variables are possibly measured with error.⁶

Thus, suppose that in the model described in (3.4)

$$Y = Z\alpha + V\delta + E,$$

some variables in X are measured with error. To be precise, suppose that X is a single variable (say exogenous changes in income) which can not be observed but instead $v = X + u$ is observed. If the signal to noise ratio of v is higher at some frequencies, then by just including these frequencies and removing the remaining ones, it is possible to increase the precision of the estimates.

When estimating model (3.4) using the generated regressors from the first stage, it is natural to assume that the explanatory variable used in the second stage regression to measure idiosyncratic changes in household endowments has a stronger signal to noise ratio at business cycle frequencies because of the nature of the idiosyncratic explanatory variables used in the first stage regression, in particular involuntary inactivity periods and the real wage rate, which are business cycle variables at the aggregate level. Hence, by removing high frequency components and very low frequency components I am likely to increase the precision of my estimates.

Moreover, by removing high frequency components I also align the observation on consumption and on income, as previously discussed. In general, removing frequency bands where the right-hand side variables have a low signal to noise is analogous to the use of moving average filters to eliminate high frequency noise components from measured income in tests of the PIH.

⁵See Altonji and Siow (1987), page 318.

⁶For a specific application example see Engle and Foley (1975).

5 Results

In this section I will present results for two groups of models: the baseline models, for which the measure of household endowments shocks used is the change in the log household total income net of transfers; and the two-step models, for which the proxy for endowment shocks is the changes in the predicted log household income. Apart from the change of the log aggregate consumption and the endowment shock proxy, the other right-hand side variables are the change in family size, the change in the number of underage children and the change in the log family food standards. These three variables are meant to capture observable changes in preferences.

All models are estimated using both standard time domain OLS and applying the band spectrum regression technique. Moreover, for the models estimated using standard OLS, three specifications are considered: one in which one year first differences are taken; next, five year first differences; and finally, ten year first differences. This was done in an effort to capture features of the data present at different frequencies and in particular to obtain power against the self-insurance alternative hypothesis. As for the models estimated using band spectrum regression methods, I consider three different frequency bands. Since the models are all taken in differences, the panel time series dimension is $T - 1 = 12$.

Consequently, the identifiable frequencies are

$$\theta_k = \frac{2\pi k}{T-1} : \theta_k \leq \pi, k = 1, 2, \dots, 11.$$

This yields six frequencies corresponding to the periodicities: 2 years; $2\frac{2}{5}$ years; 3 years; 4 years; 6 years; and 12 years. I therefore defined three different bands, the very short run (2 years, $2\frac{2}{5}$ years), the business cycle (3 years, 6 years) and the long run (12 years, ∞).

Results are reported in table 4 for the baseline model estimated using standard OLS. The three different columns give results for the three different first difference horizons considered. The t -statistics of the estimated coefficients are shown below the estimates, in parenthesis. And in the last two rows, the F -statistic as well as the p -value is shown for the null hypothesis of full consumption insurance,

Table 4: Baseline regression (Time Domain)

Variable	One Year	Five Years	Ten Years
	Time Differences	Time Differences	Time Differences
$\Delta \log$ Aggregate Consumption $\rightarrow \alpha_1$	-0.0897 (-0.75)	-0.2255 (-1.25)	-0.3335* (-2.10)
$\Delta \log$ Household Income $\rightarrow \alpha_2$	0.0665** (9.62)	0.1613** (18.99)	0.1808** (14.09)
Δ Family Size	-0.0044 (-0.49)	0.0163 [†] (1.83)	-0.0691** (-4.20)
Δ Number Children under 18	0.0251** (3.75)	0.0495** (7.54)	0.0876** (8.37)
$\Delta \log$ Family Needs	0.0483* (1.96)	0.2174** (9.76)	0.5097** (12.45)
H0: $\alpha_1 - \alpha_2 = 0; \alpha_2 = 0$			
F-stat	46.35	180.52	100.08
p-value	0.000	0.000	0.000
Significance levels : † : 10% * : 5% ** : 1%			

Note: The dependent variable is the growth rate of individual consumption minus the growth rate of aggregate consumption ($\Delta \log c - \Delta \log c^A$). In parenthesis are t-values. The specification includes an intercept.

corresponding to $\alpha \equiv \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = 0$.

Clearly the null hypothesis of full insurance is rejected, no matter the length of the first difference horizon chosen. Thus, at first glance, households appear unable to smooth consumption when exposed to idiosyncratic endowment shocks. Moreover, the idiosyncratic income coefficient estimate, $\hat{\alpha}_2$, is always significantly different from zero, and this is clearly causing the rejection of the full insurance hypothesis. However, the size of this coefficient increases with the horizon considered, suggesting that the growth rate of consumption reacts more strongly to shocks that are more long lasting. The permanent income hypothesis with no risk sharing predicts

that a purely transitory 1 percent income shock increases consumption by 1 percent times the interest rate.⁷ This is broadly consistent with the point estimate of α_2 for the one year horizon model. Insurance against more long lasting shocks seems to be much less. Mace (1991) performs the same test using data from the CEX and reports somewhat similar findings to the ones reported in the first column of table 4.⁸

Also noteworthy, the size of the aggregate consumption growth coefficient estimate, $\hat{\alpha}_1$, is not significantly different from zero, except at very long horizons (at which the 5% level test leads to a rejection). If the maintained assumption of a common coefficient of relative risk aversion, the coefficient ρ , was false, this would lead to an estimate of α_1 different from zero. Thus, this constitutes a finding which provides evidence consistent with the maintained assumption of a common coefficient of relative risk aversion. Another interesting finding is that the demographic control variables included affect the growth rate of consumption significantly only for the model corresponding to the ten years horizon. This suggests that such demographic changes, which might translate into long lasting changes in living standards, affect the low frequency component of consumption.

However, as was argued before, changes in household income are not an appropriate right-hand side variable because they will likely be correlated with the error term. Therefore, I now turn to the results in table 5 where the standard OLS estimates of the two-step model are presented.

Again the full consumption insurance hypothesis is clearly rejected at all horizons and the joint test p -values are essentially zero, although the F statistics are smaller than before. However, the coefficient estimate $\hat{\alpha}_2$ is smaller at the one year horizon than the one year baseline estimate and has a smaller t -stat as well, which suggests that the endogeneity of income biases upwards the estimated impact of endowment shocks, as expected if leisure and consumption are substitutes. At longer horizons, however, the reverse is true.⁹ Moreover, the size of the point estimate increases

⁷See Cochrane (1991), page 973.

⁸Mace (1991), table 3.

⁹This might be interpreted as evidence in favor of different elasticities of substitution in the short run and the long run.

Table 5: Two steps model (Time Domain)

Variable	One Year	Five Years	Ten Years
	Time Differences	Time Differences	Time Differences
$\Delta \log$ Aggregate Consumption $\rightarrow \alpha_1$	-0.0740 (-0.62)	-0.1848 (-1.02)	-0.2874 [†] (-1.79)
Δ Predicted log Household Income $\rightarrow \alpha_2$	0.0546** (6.37)	0.1624** (14.31)	0.1982** (11.23)
Δ Family Size	-0.0023 (-0.25)	0.0270** (2.99)	-0.0581** (-3.51)
Δ Number Children under 18	0.0216** (3.23)	0.0273** (4.20)	0.0575** (5.55)
$\Delta \log$ Family Needs	0.0247* (2.16)	0.2520** (11.24)	0.5696** (13.81)
H0: $\alpha_1 - \alpha_2 = 0$; $\alpha_2 = 0$			
F-stat	20.41	102.55	63.83
p-value	0.000	0.000	0.000
Significance levels : † : 10% * : 5% ** : 1%			

Note: The dependent variable is the growth rate of individual consumption minus the growth rate of aggregate consumption ($\Delta \log c - \Delta \log c^A$). In parenthesis are t-values. The specification includes an intercept.

with the horizon considered. The conclusions regarding $\hat{\alpha}_1$ are essentially the same as in the previous specification. On balance, the standard OLS estimates clearly lead to a rejection of the complete markets assumption. Thus, I now turn to the results obtained using the proposed pooled band spectrum regression method.

Table 6 shows estimates for the baseline specification. Each column corresponds to one of the three different frequency bands, the very short run (2 years, $2\frac{2}{5}$ years), the business cycle (3 years, 6 years) and the long run (12 years, ∞). The test statistics are the standard ones, however, as described before, an adjustment has to be made for the degrees of freedom. Thus, the estimator for the standard deviation of the error component is computed using (3.5). Turning first to the very

Table 6: The Baseline model (Band Spectrum Regression)

Variable	Very Short Run	Business Cycle	Long Run
	2 years - $2\frac{2}{5}$ years	3 years - 6 years	12 years - ∞
$\Delta \log$ Aggregate Consumption $\rightarrow \alpha_1$	0.0148 (0.05)	-0.1175 (-0.81)	-0.1732 (-0.20)
$\Delta \log$ Household Income $\rightarrow \alpha_2$	0.0007 (0.06)	0.0844** (8.12)	0.1620** (11.63)
Δ Family Size	-0.0107 (-0.64)	-0.0118 (-0.85)	0.0168 (1.13)
Δ Number Children under 18	-0.0015 (-0.12)	0.0151 (1.48)	0.0535** (4.96)
$\Delta \log$ Family Needs	-0.0915* (-1.98)	0.0824* (2.14)	0.2132** (5.66)
H0: $\alpha_1 - \alpha_2 = 0$; $\alpha_2 = 0$			
F-stat	0.00	33.08	67.69
p-value	0.997	0.000	0.000
Significance levels : † : 10% * : 5% ** : 1%			

Note: The dependent variable is the growth rate of individual consumption minus the growth rate of aggregate consumption ($\Delta \log c - \Delta \log c^A$). In parenthesis are t-values. The specification includes an intercept but no degrees of freedom adjustment was made.

short run, the full insurance hypothesis clearly cannot be rejected. Moreover, in the very short run, the growth rate of consumption is not responsive to idiosyncratic endowment changes at any significance level. As argued before, this can be the result of measurement error leading to a downward bias in the α_2 estimate; however it also highlights the fact that at high frequencies, the full insurance test has no power against the null hypothesis of self insurance. On the other hand these findings suggest that liquidity constraints do not seem to prevent households from smoothing high frequency consumption fluctuations.

Turning to the business cycle frequencies, which I identify as corresponding to the periodicities between three years and six years, full insurance is soundly rejected.

The contrast with the very short run, suggests that the findings for the higher frequencies may be either an artifact of measurement error, or alternatively the result of self-insurance, which allows households to smooth consumption fluctuations in the very short run. The coefficient on the growth rate of income, α_2 , which is again significantly different from zero and thus is causing the rejection of the joint hypothesis of full insurance, is higher than the estimates obtained through standard OLS at the one year horizon. Again, this is consistent with the PIH benchmark. Using the certainty equivalence model and thus Hall's (1978) martingale hypothesis as a benchmark, purely transitory (one year) 1 percent income shocks should increase consumption by 1 percent times the interest rate and equivalently, the response of consumption to a shock which is more persistence, lasting for example 3 years, should be greater and equal to the annuity value of the shock.

As for the changes in demographic characteristics, except for changes in family needs, the other variables are not important in explaining changes in the growth rate of consumption. This is probably explained by the small variance of demographic variables at business cycle frequencies. Thus, to the extent that unobservable preference shifts are also driven by demographic factors, it is reasonable to expect that the endogeneity problems may be a smaller source of asymptotic bias at business cycle frequencies. Finally, in the very long run, the full insurance hypothesis is again rejected. Moreover, the demographic variables are significant at explaining changes in the growth rate of consumption. Again, the findings are consistent with the PIH benchmark, but no insurance above what would be predicted by the PIH seems to be achieved, in particular at business cycle frequencies.

I finally turn my attention to the two-step model, estimated using the band spectrum method. The findings are reported in table 7. The results of the full consumption insurance test are very similar to the ones obtained using the baseline model in stead of the two-step procedure.

The full insurance hypothesis is rejected at business cycle frequencies and at lower frequencies, but it cannot be rejected in the very short run. Therefore, results again suggest that most consumption insurance is due to self-insurance. Moreover, the findings are broadly consistent with the PIH theory, and in particular liquid-

Table 7: Two steps model (Band Spectrum Regression)

Variable	Very Short Run	Business Cycle	Long Run
	2 years - $2\frac{2}{5}$ years	3 years - 6 years	12 years - ∞
$\Delta \log$ Aggregate Consumption $\rightarrow \alpha_1$	0.0154 (0.05)	-0.0931 (-0.64)	0.0124 (0.01)
Δ Predicted log Household Income $\rightarrow \alpha_2$	0.0166 (1.19)	0.0633** (4.84)	0.1753** (8.95)
Δ Family Size	-0.0114 (-0.68)	-0.0079 (-0.57)	0.0284 [†] (1.90)
Δ Number Children under 18	-0.0011 (-0.09)	0.0092 (7.54)	0.0286** (2.68)
$\Delta \log$ Family Needs	-0.0926* (-2.01)	0.0913* (2.36)	0.2488** (6.57)
H0: $\alpha_1 - \alpha_2 = 0$; $\alpha_2 = 0$			
F-stat	0.71	11.83	40.08
p-value	0.494	0.000	0.000
Significance levels : † : 10% * : 5% ** : 1%			

Note: The dependent variable is the growth rate of individual consumption minus the growth rate of aggregate consumption ($\Delta \log c - \Delta \log c^A$). In parenthesis are t-values. The specification includes an intercept but no degrees of freedom adjustment was made.

ity constraints apparently do not prevent households from achieving the level of consumption smoothing predicted by the PIH.

Comparing the estimates of α_2 obtained using the two steps instrumental variable procedure with the ones presented previously, the evidence again suggests that the endogeneity of income may lead to an upward bias in the estimation, at high frequencies and at business cycle frequencies, suggesting substitutability between consumption and leisure, however the reverse is true for the very long run. This result is consistent with the evidence provided by the standard OLS estimates.

As for the estimates of α_1 , which are never significantly different from zero, the results are again consistent with the maintained assumption of a common coeffi-

cient of relative risk aversion, at all the frequency bands. Finally, as before, the changes in demographic factors only appear to significantly affect the growth rate of consumption in the long run, consistent with the interpretation that demographic changes translate into long lasting changes in living standards, which affect the low frequency component of consumption.

6 Conclusion

This paper attempts to test full consumption insurance and at the same time it distinguishes between consumption smoothing through self-insurance and consumption insurance through risk-sharing. Thus, the central contribution of this paper is to offer a framework which simultaneously allows testing consumption smoothing across time and across states of nature.

The method employed, band spectrum regression, allows identifying features of measured consumption growth over particular frequencies such as business cycles, seasonal frequencies or long horizons, without requiring parametric assumptions to distinguish between transitory and permanent shocks. Another advantage of the econometric approach is that spectrum regression allows estimating errors-in-variables models with more precision when it is suspected that the signal-to-noise ratio of some explanatory variables might be greater at some particular frequencies. Therefore, provided that frequencies at which the signal-to-noise ratio is weaker are removed, the full consumption insurance test proposed will be free of bias from measurement error. This approach resembles the use of moving average filters used to eliminate high frequency noise components from measured income in tests of the PIH. However, it is a non-parametric approach and therefore more robust to the presence of measurement error.

To summarize the results of the paper, the full consumption insurance hypothesis is soundly rejected at business cycle frequencies and at lower frequencies, but it cannot be rejected in the short run. Therefore, the findings are broadly consistent with the PIH theory. In particular liquidity constraints do not appear to prevent households from achieving the level of high frequency consumption smoothing predicted by the PIH. However, no evidence is found in favor of insurance implementation beyond

self-insurance. The rejection of the full insurance hypothesis at business cycle frequencies raises questions about the appropriateness of business cycle models which are built on the assumption of complete markets and hence full risk-sharing.

A Appendix

A.1 Variables Definition

Consumption: is total food consumption (food at home + food stamps + meals away from home) in 1983 dollars deflated using the food price component of the CPI.

Family Income Net of Transfers: Taxable Income of Head and Wife + Total Transfers of Head and Wife + Taxable Income of Others + Transfer Income of Others – Total transfer income of Head and Wife, in 1983 dollars deflated using the food price component of the CPI.

Wage Rate: \log (Labor Income of Head in 1983 dollars deflated using the food price component of the CPI/Hours of Work For Money of Head)

Involuntary Inactivity: Head's annual hours of unemployment + Head's annual hours on strike

Involuntary Move: Dummy equaled one if head moved because of response to outside events (involuntary reasons): (household unit) coming down; being evicted; armed services; health reasons; divorce; retiring because of health; etc...

Family Size: Actual number of persons in Family Unit.

Number of Underage Children: Actual number of children aged 0-17.

Food Needs: This variable is generated by multiplying the weekly food needs by 52 and then making the following adjustments for economies of scale: add 20 percent for one-person families, 10 percent for two-person families, 5 percent for three-person families and subtract 5 percent for five-person families and 10 percent for families with six or more persons.

A.2 Sample Selection Procedure

The sample of households used in this paper were drawn from the 1974-1988 family files. The sample selection procedure leading to the balanced panel used throughout the paper was the following:

1. households that were part of the original SEO subsample were excluded: 28080 observations deleted.
2. observations for which the total family income was negative or missing were excluded: 88605 observations deleted.
3. observations for which the head average hourly earnings were negative or missing were excluded: 14705 observations deleted.
4. observations for which the real total family income increased more than 300% or decreased more than 75% were excluded: 880 observations deleted.
5. observations for which the real total family food expenditure increased by more than 300% or decreased by more than 75% were excluded: 1183 observations deleted.
6. households in which there occurred a change in the head were excluded: 29544 observations deleted.
7. households whose head is a farmer were excluded: 379 observations deleted.
8. households whose head age was misreported were excluded: 271 observations deleted.
9. households whose head was older than 65 or younger than 20 were excluded: 1891 observations deleted.
10. households whose head does not know how to read or write or for which the education is misreported were excluded: 391 observations deleted.

This selection procedure allowed to construct a balanced panel of 966 households over 13 years, corresponding to 12558 observations.

References

- [1] J. G. ALTONJI AND A. SIOW, *Testing the response of consumption to income changes with (noisy) panel data*, The Quarterly Journal of Economics, 102 (1987), pp. 293–328.
- [2] S. ALTUG AND R. A. MILLER, *Household choices in equilibrium*, Econometrica, 58 (1990), pp. 543–70.
- [3] O. ATTANASIO AND S. J. DAVIS, *Relative wage movements and the distribution of consumption*, Journal of Political Economy, 104 (1996), pp. 1227–62.
- [4] O. P. ATTANASIO AND G. WEBER, *Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey*, Journal of Political Economy, 103 (1995), pp. 1121–57.
- [5] M. BAXTER AND R. G. KING, *Measuring business cycles: Approximate band-pass filters for economic time series*, The Review of Economics and Statistics, 81 (1999), pp. 575–593.
- [6] J. BERKOWITZ, *Generalized spectral estimation of the consumption-based asset pricing model*, Journal of Econometrics, 104 (2001), pp. 269–288.
- [7] B. S. BERNANKE, *Permanent income, liquidity, and expenditure on automobiles: Evidence from panel data*, The Quarterly Journal of Economics, 99 (1984), pp. 587–614.
- [8] R. BLUNDELL AND I. PRESTON, *Consumption inequality and income uncertainty*, The Quarterly Journal of Economics, 113 (1998), pp. 603–640.
- [9] R. W. BLUNDELL, L. PISTAFERRI, AND I. PRESTON, *Consumption inequality and partial insurance*, American Economic Review, (forthcoming 2008).
- [10] L. J. CHRISTIANO AND R. J. VIGFUSSON, *Maximum likelihood in the frequency domain: the importance of time-to-plan*, Journal of Monetary Economics, 50 (2003), pp. 789–815.
- [11] J. H. COCHRANE, *A simple test of consumption insurance*, Journal of Political Economy, 99 (1991), pp. 957–76.
- [12] D. CORBAE, S. OULIARIS, AND P. C. B. PHILLIPS, *A reexamination of the consumption function using frequency domain regressions*, Empirical Economics, 19 (1994), pp. 595–609.

- [13] D. M. CUTLER AND L. F. KATZ, *Rising inequality? changes in the distribution of income and consumption in the 1980's*, American Economic Review, 82 (1992), pp. 546–51.
- [14] A. DEATON AND C. PAXSON, *Intertemporal choice and inequality*, Journal of Political Economy, 102 (1994), pp. 437–67.
- [15] S. DYNARSKI AND J. GRUBER, *Can families smooth variable earnings?*, Brookings Papers on Economic Activity, (1997), pp. 229–84.
- [16] R. F. ENGLE, *Band spectrum regression*, International Economic Review, 15 (1974), pp. 1–11.
- [17] —, *Testing price equations for stability across spectral frequency bands*, Econometrica, 46 (1978), pp. 869–81.
- [18] R. F. ENGLE AND D. K. FOLEY, *An asset price model of aggregate investment*, International Economic Review, 16 (1975), pp. 625–47.
- [19] R. F. ENGLE AND R. GARDNER, *Some finite sample properties of spectral estimators of a linear regression*, Econometrica, 44 (1976), pp. 149–65.
- [20] F. GUVENEN, *Do stockholders share risk more effectively than nonstockholders?*, The Review of Economics and Statistics, 89 (2007), pp. 275–288.
- [21] R. E. HALL, *Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence*, Journal of Political Economy, 86 (1978), pp. 971–87.
- [22] R. E. HALL AND F. S. MISHKIN, *The sensitivity of consumption to transitory income: Estimates from panel data on households*, Econometrica, 50 (1982), pp. 461–81.
- [23] A. C. HARVEY, *Linear regression in the frequency domain*, International Economic Review, 19 (1978), pp. 507–12.
- [24] F. HAYASHI, J. ALTONJI, AND L. KOTLIKOFF, *Risk-sharing between and within families*, Econometrica, 64 (1996), pp. 261–94.
- [25] D. KRUEGER AND F. PERRI, *Does income inequality lead to consumption inequality? evidence and theory*, Review of Economic Studies, 73 (2006), pp. 163–193.
- [26] B. J. MACE, *Full insurance in the presence of aggregate uncertainty*, Journal of Political Economy, 99 (1991), pp. 928–56.

- [27] J. A. NELSON, *On testing for full insurance using consumer expenditure survey data: Comment*, *Journal of Political Economy*, 102 (1994), pp. 384–94.
- [28] A. PAGAN, *Econometric issues in the analysis of regressions with generated regressors*, *International Economic Review*, 25 (1984), pp. 221–47.
- [29] K. STORESLETTEN, C. I. TELMER, AND A. YARON, *How important are idiosyncratic shocks? evidence from labor supply*, *American Economic Review*, 91 (2001), pp. 413–417.
- [30] R. M. TOWNSEND, *Risk and insurance in village india*, *Econometrica*, 62 (1994), pp. 539–91.