

How Much Does Household Collateral Constrain Regional Risk Sharing?

Hanno Lustig*

UCLA Anderson and NBER

Stijn Van Nieuwerburgh†

NYU Stern, NBER, and CEPR

August 25, 2009

Abstract

We construct a new data set of consumption and income data for the largest U.S. metropolitan areas, and we show that the extent of risk-sharing between regions varies substantially over time. In times when US housing collateral is scarce nationally, regional consumption is about twice as sensitive to income shocks. We also document higher sensitivity in regions with lower housing collateral. Household-level borrowing frictions can explain this new stylized fact. When the value of housing relative to human wealth falls, loan collateral shrinks, borrowing (risk-sharing) declines, and the sensitivity of consumption to income increases. Our model aggregates heterogeneous, borrowing-constrained households into regions characterized by a common housing market. The resulting regional consumption patterns quantitatively match those in the data.

*email: hlustig@anderson.ucla.edu, Anderson School of Management, UCLA, Box 951477 Los Angeles, CA 90095-1477

†email: svnieuwe@stern.nyu.edu, Dept. of Finance, NYU, 44 West Fourth Street, Suite 9-120, New York, NY 10012. First version May 2002. The material in this paper circulated earlier as "Housing Collateral and Risk Sharing Across US Regions." (NBER Working Paper). The authors thank Thomas Sargent, David Backus, Dirk Krueger, Patrick Bajari, Timothy Cogley, Marco Del Negro, Robert Hall, Lars Peter Hansen, Patrick Kehoe, Martin Lettau, Sydney Ludvigson, Fabrizio Perri, Luigi Pistaferri, Laura Veldkamp, Pierre-Olivier Weill, and Noah Williams. We are very grateful to Narayana Kocherlakota, the Editor, the Associate Editor, as well as the referees, for their help. We also benefited from comments from seminar participants at various institutions. **Keywords:** Regional risk sharing, housing collateral, **JEL** F41,E21

1 Introduction

On average, US metropolitan areas share only a modest fraction of region-specific income risk. But this fraction varies substantially over time. According to our estimates, the fraction of regional income risk that is traded away, more than doubles when we compare the lowest to the highest collateral scarcity period in postwar US data. A second and related stylized fact is that the dispersion of regional consumption exceeds the dispersion of regional income.¹ We will refer to this as the dispersion anomaly.

We propose an equilibrium model of household risk sharing that produces the time-variation in regional risk sharing as well as the dispersion anomaly. The model adds a regional dimension to the model of Lustig and Van Nieuwerburgh (2005), a crucial extension to generate the dispersion anomaly. Within each region, households face a stochastic income process that has a household-specific and a region-specific component. What prevents perfect consumption insurance is that households can share income risk only to the extent that borrowing is collateralized by housing wealth. Human wealth is not collateralizable. The key ingredient for replicating the dispersion anomaly is that borrowing constraints operate at the household level. Such constraints are much tighter than the constraints that would be faced by a stand-in agent at the regional level. Because there is some *intra-regional* risk-sharing, household consumption as a share of regional consumption is usually less dispersed than household income within a region. Aggregation to the regional level produces *inter-regional* consumption dispersion that exceeds regional income dispersion, at least when housing collateral is sufficiently scarce. The key ingredient for replicating the time-variation in the degree of risk sharing is variation in the value of housing collateral. Variation in the housing supply and the equilibrium house price shift the effectiveness of the household risk sharing technology over time. A reduction in the value of housing collateral tightens the household collateral constraints, causing regional consumption to respond more to regional income shocks. The ratio of income-to-consumption dispersion increases as collateral becomes scarcer.

The null hypothesis of perfect insurance is usually tested by projecting regional consumption growth on income growth. The collateral effect in our model introduces an additional interaction term of region-specific income growth with housing collateral. According to the theory, the sign on this interaction term should be negative. When collateral is scarce, a shock to region-specific income leads to a larger change in region-specific consumption. We run this linear regression on actual data and on data generated by our calibrated model. In the actual data, the sign on the interaction term is indeed negative. The housing collateral effect is economically significant.

¹A related but distinct quantity anomaly –the correlation of consumption growth is lower than that of output growth– has previously been documented in international (e.g. Backus, Kehoe and Kydland (1992), and Lewis (1996)) and in state-level data (e.g. Atkeson and Bayoumi (1993), Hess and Shin (2000) and Crucini (1999)).

Housing collateral scarcity in the 95th percentile of the empirical distribution is associated with 42% of region-specific income shocks being shared, while collateral scarcity in the 5th percentile level corresponds to regions sharing 86% of income risk. The same regression on model-generated data for consumption and income replicates these results. The advantage of this risk-sharing test, based on the interaction effect of the collateral measure and income growth, is that it is more specific than the standard regression, and the appropriate test for our collateral model. There is evidence from the cross-section as well. The income elasticity of consumption growth doubles in the quartile of regions with the least collateral, compared to those regions in the highest quartile.

The rest of the paper is organized as follows. Section 2 sets up the model, characterizes equilibrium allocations and prices, calibrates, and computes it. Section 3 describes the data and compares the results from the linear consumption growth regressions in the model and in the data. Section 4 presents additional evidence for the housing collateral mechanism. We find similar results for Canadian provinces and find that there is also a positive relationship between the degree of risk sharing and regional measures of collateral. Section 5 concludes.

2 A Theory of Time-Varying Risk Sharing

In this section we provide a model that replicates two key features of the observed regional consumption and income distribution. First, the average ratio of the cross-sectional consumption dispersion to income dispersion is larger than one, i.e. the dispersion anomaly. Second, this ratio increases as housing collateral becomes scarcer.

The model is a dynamic general equilibrium model that approximates the modest frictions inhibiting perfect risk-sharing in advanced economies like the US. The model is based on two ideas: that debts can only be enforced to the extent that they can be collateralized, and that the primary source of collateral is housing. Our emphasis on housing, rather than financial assets, reflects three features of the US economy: the participation rate in housing markets is very high (2/3 of households own their home), the value of the residential real estate makes up over seventy-five percent of total assets for the median household (Survey of Consumer Finances, 2001), and housing is a prime source of collateral.

We relax the assumption in the Lucas (1978) endowment economy that contracts are perfectly enforceable, following Alvarez and Jermann (2000), and allow households to file for bankruptcy, following Chien and Lustig (2009). Each household owns part of the housing stock. Housing provides both utility services and collateral services. When a household chooses not to honor its debt repayments, it loses all housing collateral but its labor income is protected from creditors. Defaulting households regain immediate access to credit markets. The lack of commitment gives

rise to collateral constraints whose tightness depends on the relative abundance of housing collateral. As a result, the effectiveness of the household risk sharing technology endogenously varies over time due to movements in the value of housing collateral.²

The section starts with a description of the environment in 2.1 and market structure in 2.2. We then provide a characterization of equilibrium allocations in section 2.3. The model gives rise to a simple, non-linear risk-sharing rule. The model has two levels of heterogeneity: households and regions. The key friction, collateralized borrowing, operates at the household level. We construct regional consumption and income by aggregating across households in a region. We show in 2.4 that the household collateral constraints give rise to tighter constraints at the regional level than those that would arise if there was a representative agent in each region. Section 2.5 calibrates the model and section 2.6 explains the computational procedure. Section 2.7 simulates the model. It shows that the aggregation from the household to the regional level generates the dispersion anomaly at the regional level. In the next section, we use the same simulated data to estimate linear consumption growth regressions at the regional level.

2.1 Uncertainty, Preferences and Endowments

We consider an economy with a continuum of regions. There are two types of infinitely lived households in each of these regions, and households cannot move between regions.

Uncertainty There are three layers of uncertainty: an event s consists of x , y , and z . We use s^t to denote the history of events $s^t = (x^t, y^t, z^t)$, where $x^t \in X^t$ denotes the history of household events, $y^t \in Y^t$ denotes the history of regional events and $z^t \in Z^t$ denotes the history of aggregate events. $\pi(s^t|s_0)$ denotes the probability of history s^t , conditional on observing s_0 .

The household-level event x is first-order Markov, and the x shocks are independently and identically (henceforth i.i.d.) distributed across regions. In our calibration below, x takes on one of two values, high (hi) or low (lo). When $x = hi$, the first household in that region is in the high state, and, the second household is in the low state. When $x = lo$, the first household is in the low state. The region-level event y is also first-order Markov and it is i.i.d. across regions. We will appeal to a law of large numbers (LLN) when integrating across households in different regions.³

Preferences The households j in each region i rank consumption plans consisting of (non-durable) non-housing consumption $\{c_t^{ij}(s^t)\}$ and housing services $\{h_t^{ij}(s^t)\}$ according to the objective func-

²Ortalo-Magne and Rady (2002), Ortalo-Magne and Rady (2006) and Pavan (2005) have also developed models that deliver this feature.

³The usual caveat applies when applying the LLN; we implicitly assume the technical conditions outlined by Uhlig (1996) are satisfied.

tion in equation (1).

$$U(c, h) = \sum_{s^t | s^0} \sum_{t=0}^{\infty} \beta^t \pi(s^t | s^0) u(c_t(s^t), h_t(s^t)), \quad (1)$$

where β is the time discount factor, common to all regions. The households have power utility over a CES-composite consumption good:

$$u(c_t, h_t) = \frac{1}{1-\gamma} \left[c_t^{\frac{\varepsilon-1}{\varepsilon}} + \psi h_t^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{(1-\gamma)\varepsilon}{\varepsilon-1}},$$

The preference parameter $\psi > 0$ converts the housing stock into a service flow, γ is the coefficient of relative risk aversion, and ε is the intra-temporal elasticity of substitution between non-durable and housing services consumption.⁴

Endowments Each of the households, indexed by j , in a region, indexed by i , is endowed with a claim to a labor income stream $\{\eta_t^{ij}(x_t, y_t, z^t)\}$. The aggregate non-housing endowment $\{\eta_t^a(z^t)\}$ is the sum of the household endowments in all regions:

$$\eta_t^a(z^t) = \sum_{y_t} \pi_z(y_t) \eta_t^i(y_t, z^t)$$

where $\pi_z(y_t)$ denotes the fraction of regions that draws aggregate state z . Likewise, the regional non-housing endowment $\{\eta_t^i(y_t, z^t)\}$ is the sum of the individual endowments of the households in that region:

$$\eta_t^i(y_t, z^t) = \sum_{j=1,2} \eta_t^{ij}(x_t, y_t, z^t).$$

The left hand side does not depend on x_t , because the two household endowments always sum to the regional endowment, regardless of whether the first household is in the high or the low state.

Each region i receives a share of the aggregate non-housing endowment denoted by $\hat{\eta}_t^i(y_t, z_t) \gg 0$. Thus, regional income shares are defined as in the empirical section: $\hat{\eta}_t^i(y_t, z_t) = \frac{\eta_t^i(y_t, z^t)}{\eta_t^a(z^t)}$. Household j 's labor endowment share in region i , measured as a fraction of the regional endowment share, is denoted $\hat{\eta}_t^j(x_t) \gg 0$. The shares add up to one within each region: $\hat{\eta}_t^1(x_t) + \hat{\eta}_t^2(x_t) = 1$. The level of the labor endowment of household j in region i can be written as:

$$\eta_t^{ij}(x_t, y_t, z^t) = \hat{\eta}_t^j(x_t) \hat{\eta}_t^i(y_t, z_t) \eta_t^a(z^t).$$

In addition, each region is endowed with a stochastic stream of non-negative housing services

⁴These preferences belong to the class of homothetic power utility functions of Eichenbaum and Hansen (1990). The special case of separability corresponds to $\gamma\varepsilon = 1$.

$\chi_t^i(y^t, z^t) \gg 0$. In contrast to non-housing consumption, the housing services cannot be transported across regions. We will come back to the assumptions we make on χ^i at the end of section 2.3. So far, we have made the following assumptions about the endowment processes:

Assumption 1. *The household-specific labor endowment share $\hat{\eta}^j$ only depends on x_t . The regional income share $\hat{\eta}_t^i$ only depends on (y_t, z_t) . The events (x, y, z) follow a first-order Markov process.*

2.2 Trading

We set up an Arrow-Debreu economy where all trade takes place at time zero, after observing the initial state s_0 .⁵ We denote the present discounted value of any endowment stream $\{d\}$ after a history s^t as $\Pi_{s^t} [\{d_\tau(s^\tau)\}]$, defined by $\sum_{s^\tau|s^t} \sum_{\tau=t}^{\infty} [\rho_\tau(s^\tau|s^t) d_\tau(s^\tau|s^t)]$, where $\rho_t(s^t)$ denotes the Arrow-Debreu price of a unit of non-housing consumption in history s^t .

Households in each region purchase a complete, state-contingent consumption plan

$$\{c_t^{ij}(\theta_0^{ij}, s^t), h_t^{ij}(\theta_0^{ij}, s^t)\}_{t=0}^{\infty}$$

where θ_0^{ij} denotes initial non-labor wealth.⁶ They are subject to a single time zero budget constraint which states that the present discounted value of non-housing and housing consumption must not exceed the present discounted value of the labor income stream and the initial non-labor wealth:

$$\Pi_{s_0} [\{c_t^{ij}(\theta_0^{ij}, s^t) + \rho_t^i(s^t) h_t^{ij}(\theta_0^{ij}, s^t)\}] \leq \theta_0^{ij} + \Pi_{s_0} [\{\eta_t^{ij}(s^t)\}], \quad (2)$$

where $\rho_t^i(s^t)$ denotes the rental price of housing services in region i .

Collateral Constraints In this time-zero-trading economy, collateral constraints restrict the value of a household's consumption claim net of its labor income claim to be non-negative:

$$\Pi_{s^t} [\{c_\tau^{ij}(\theta_0^{ij}, s^\tau) + \rho_\tau^i(s^\tau) h_\tau^{ij}(\theta_0^{ij}, s^\tau)\}] \geq \Pi_{s^t} [\{\eta_\tau^{ij}(x_\tau, y_\tau, z^\tau)\}]. \quad (3)$$

The left hand side denotes the value of adhering to the contract following node s^t ; the right hand side the value of default. Default implies the loss of all housing collateral wealth, and a fresh start with the present value of future labor income. The households in each region are subject to a sequence of collateral constraints, one for each future state s^τ . These constraints are *not too tight*,

⁵The same allocation can also be decentralized with sequential trade.

⁶ θ_0^{ij} denotes the value of household j 's initial claim to housing wealth, as well as any other financial wealth that is in zero net aggregate supply. We refer to this as non-labor wealth. The initial distribution of non-labor wealth is denoted Θ_0 .

in the sense of Alvarez and Jermann (2000), in an environment where agents cannot be excluded from trading.⁷

These constraints differ from the typical solvency constraints that decentralize constrained efficient allocations in environments with exclusion from trading upon default.⁸

Definition 1. A Kehoe-Levine equilibrium is a list of allocations $\{c_t^{ij}(\theta_0^{ij}, s^t)\}, \{h_t^{ij}(\theta_0^{ij}, s^t)\}$ and prices $\{\rho_t^i(s^t)\}, \{p_t(s^t)\}$ such that, for a given initial distribution Θ_0 over non-labor wealth holdings and initial states (θ_0, s_0) , (i) the allocations solve the household problem, (ii) the markets clear in all states of the world:

Consumption markets clear for all t, z^t :

$$\sum_{j=1,2} \sum_{x^t, y^t} \int c_t^{ij}(\theta_0^{ij}, x^t, y^t, z^t) \frac{\pi(x^t, y^t, z^t | x_0, y_0, z_0)}{\pi(z^t | z_0)} d\Theta_0 = \eta_t^a(z^t)$$

Housing markets in each region i clear for all t, x^t, y^t, z^t :

$$\sum_{j=1,2} h_t^{ij}(\theta_0^{ij}, x^t, y^t, z^t) = \chi_t^i(y^t, z^t).$$

2.3 Equilibrium Allocations

To characterize the equilibrium consumption dynamics we use stochastic consumption weights that describe the consumption of each household as a fraction of the aggregate endowment (see appendix A for a complete derivation). Instead of solving a social planner problem, we characterize equilibrium allocations and prices directly off the household's necessary and sufficient first order conditions. The household problem is a standard convex problem: the objective function is concave and the constraint set is convex. In equilibrium, for any two households j and j' in any two regions i and i' , the level of marginal utilities satisfies:

$$\xi_{t+1}^{ij} u_c(c_{t+1}^{ij}(\theta_0^{ij}, s^t, s'), h_{t+1}^{ij}(\theta_0^{ij}, s^t, s')) = \xi_{t+1}^{i'j'} u_c(c_{t+1}^{i'j'}(\theta_0^{i'j'}, s^t, s'), h_{t+1}^{i'j'}(\theta_0^{i'j'}, s^t, s')),$$

at any node (s^t, s') , where ξ^{ij} is the consumption weight of household j in region i . Our model provides an equilibrium theory of these consumption weights. We focus here on equilibrium allocations

⁷See Chien and Lustig (2009) for a formal proof.

⁸Most other authors in this literature take the outside option upon default to be exclusion from future participation in financial markets (e.g. Kehoe and Levine (1993), Krueger (1999), Krueger and Perri (2006), and Kehoe and Perri (2002)). If we imposed exclusion from trading instead, the solvency constraints would be looser on average, but the same mechanism would operate. The reason is that in autarchy the household would still have to buy housing services with its endowment of non-housing goods. An increase in the relative price of housing services would worsen the outside option and loosen the solvency constraints, as it does in our model.

in the model where preferences over non-durable consumption and housing services are separable ($\gamma\varepsilon = 1$), but all results carry over to the case of non-separability.

Cutoff Rule The equilibrium dynamics of the consumption weights are non-linear. They follow a simple cutoff rule, which follows from the first order conditions of the constrained optimization problem. The weights start off at $\xi_0^{ij} = \nu^{ij}$ at time zero; this initial weight is the multiplier on the initial promised utility constraints. The new weight ξ_t^{ij} of a generic household ij that enters period t with weight ξ_{t-1}^{ij} equals the old weight as long as the household does not switch to a state with a binding collateral constraint. When a household enters a state with a binding constraint, its new weight ξ_t^{ij} is re-set to a cutoff weight $\underline{\xi}_t(x_t, y_t, z^t)$.

$$\xi_t^{ij}(\nu^{ij}, s^t) = \begin{cases} \xi_{t-1}^{ij} & \text{if } \xi_{t-1}^{ij} > \underline{\xi}_t(x_t, y_t, z^t) \\ \underline{\xi}_t(x_t, y_t, z^t) & \text{if } \xi_{t-1}^{ij} \leq \underline{\xi}_t(x_t, y_t, z^t) \end{cases} \quad (4)$$

$\underline{\xi}_t(x_t, y_t, z^t)$ is the consumption weight at which the collateral constraint (3) holds with equality. It does not depend on the entire history of household-specific and region-specific shocks (x^t, y^t) , only the current shock (x_t, y_t) . This amnesia property crucially depends on assumption 1. The reason is that the right hand side of the collateral constraint in (3) only depends on the current shock (x_t, y_t) when the constraint binds. This immediately implies that household ij 's consumption share cannot depend on the region's history of shocks (see proposition 3 in appendix A for a formal proof).

The consumption in node s^t of household ij is fully pinned down by this cutoff rule:

$$c_t^{ij}(s^t) = \frac{(\xi_t^{ij}(\nu^{ij}, s^t))^{\frac{1}{\gamma}}}{\xi_t^a(z^t)} c_t^a(z^t). \quad (5)$$

Its consumption as a fraction of aggregate consumption equals the ratio of its individual stochastic consumption weight ξ_t^{ij} raised to the power $\frac{1}{\gamma}$ to the aggregate consumption weight ξ_t^a . This aggregate consumption weight is computed by integrating over the new household weights across all households, at aggregate node z^t :

$$\xi_t^a(z^t) = \sum_{j=1,2} \sum_{x^t, y^t} \int (\xi_t^{ij}(\nu^{i,j}, s^t))^{\frac{1}{\gamma}} \frac{\pi(x^t, y^t, z^t | x_0, y_0, z_0)}{\pi(z^t | z_0)} d\Phi_0^j, \quad (6)$$

where Φ_0^j is the cross-sectional joint distribution over initial household consumption weights ν and the initial shocks (x_0, y_0) for households of type $j = 1, 2$. By the law of large numbers, the aggregate weight process only depends on the *aggregate history* z^t .

The *risk sharing rule* for non-housing consumption in (5) clears the market for non-durable consumption by construction, because the re-normalization of consumption weights by the aggregate

consumption weight ξ_t^a guarantees that the consumption shares integrate to one across regions. It follows immediately from (4), (5), and (6) that in a stationary equilibrium, each household's consumption share is drifting downwards as long as it does not switch to a state with a binding constraint, while the consumption share of the constrained households jump up. The rate of decline of the consumption share for all unconstrained households is the same, and equal to the *aggregate weight shock* $g_{t+1} \equiv \xi_{t+1}^a / \xi_t^a$. When none of the households is constrained between nodes z^t and z^{t+1} , the aggregate weight shock g_{t+1} equals one. In all other nodes, the aggregate weight shock is strictly greater than one. The risk-sharing rule for housing services is linear as well:

$$h_t^{ij}(s^t) = \frac{(\xi_t^{ij}(\nu^{ij}, s^t))^{\frac{1}{\gamma}}}{\xi_t^i(x^t, y^t, z^t)} \chi_t^i(y_t, z^t), \quad (7)$$

where the denominator is now the *regional weight shock*, defined as

$$\xi_t^i(x^t, y^t, z^t) = \sum_{j=1,2} (\xi_t^{ij}(\nu^{ij}, s^t))^{\frac{1}{\gamma}}.$$

The appendix verifies that this rule clears the housing market in each region.⁹

Equilibrium State Prices In each date and state, random payoffs are priced by the unconstrained household, who have the highest intertemporal marginal rate of substitution (Alvarez and Jermann 2000). The price of a unit of a consumption in state s^{t+1} in units of s^t consumption is their intertemporal marginal rate of substitution, which can be read off directly from the risk sharing rule in (5):

$$\frac{p_{t+1}(s^{t+1})}{p_t(s^t)\pi(s_{t+1}|s_t)} = \beta \left(\frac{c_{t+1}^a}{c_t^a} \right)^{-\gamma} g_{t+1}^\gamma. \quad (8)$$

This derivation relies only on the invariance of the unconstrained household's weight between t and $t + 1$. The first part is the representative agent pricing kernel under separability. The collateral constraints contribute a second factor to the stochastic discount factor, the aggregate weight shock raised to the power γ .

Regional Rental Prices The equilibrium relative price of housing services in region i , ρ^i , equals the marginal rate of substitution between consumption and housing services of the households in

⁹In the case of non-separable preferences between non-housing and housing consumption ($\gamma \neq 1$), the equilibrium consumption allocations also follow a cutoff rule, similar to the one in equations (4), (5), and (7). In this case, the consumption weight changes when the non-housing expenditure share changes, even if the region does not enter a state with a binding constraint. The derivation is in a separate appendix, available on the authors' web sites.

that region:

$$\rho_t^i(y^t, z^t) = \frac{u_h(c_t^{ij}(\theta_0^{ij}, s^t), h_t^{ij}(\theta_0^{ij}, s^t))}{u_c(c_t^{ij}(\theta_0^{ij}, s^t), h_t^{ij}(\theta_0^{ij}, s^t))} = \psi \left(\frac{h_t^{ij}}{c_t^{ij}} \right)^{\frac{-1}{\epsilon}} = \psi \left(\frac{\xi_t^a \chi_t^i}{\xi_t^i c_t^a} \right)^{\frac{-1}{\epsilon}}. \quad (9)$$

The second equality follows from the CES utility kernel; the last equality substitutes in the equilibrium risk sharing rules (5) and (7). Because each region consumes its own housing services endowment, the rental price is region-specific and depends on the region-specific shocks y^t .

Non-Housing Expenditure Shares Using the risk sharing rule under separable utility, it is easy to show that the non-housing expenditure share is the same for all households j in region i (see appendix A):

$$\frac{c_t^{ij}}{c_t^{ij} + \rho_t^i h_t^{ij}} \equiv \alpha_t^{ij} = \alpha_t^i$$

In the remainder of the paper, we focus on the case of a perfectly elastic supply of housing services at the regional level. To do so, we impose an additional restriction on the regional housing endowments.

Assumption 2. *The regional housing endowments χ_t^i are chosen such that $\rho_t^i(y^t, z^t) = \rho_t(z^t)$ in equation (9) for all y^t, z^t .*

Under this assumption on regional housing endowments, regional rental prices only depend on the aggregate state history z^t , and are therefore equal across regions. Likewise, the equilibrium expenditure shares α^i are equal across regions and a function of the aggregate history z^t only: $\alpha_t^i = \alpha_t(z^t)$. The reason for this assumption is that, without it, the expenditure shares would also depend on the history of region-specific shocks. This would impute too much volatility to regional housing expenditures shares. The data suggest that housing expenditure shares are not very volatile over time and quite similar across regions (Davis and Ortalo-Magne 2007).

Tightness of the Collateral Constraints Because of the collateral constraints, labor income shocks cannot be fully insured in spite of the full set of consumption claims that can be traded. How much risk sharing the economy can accomplish depends on the ratio of aggregate housing collateral wealth to non-collateralizable human wealth. Integrating housing wealth and human across all households in all regions, that ratio can be written as:

$$\frac{\Pi_{z^t} \left[\left\{ c_t^a(z^t) \left(\frac{1}{\alpha_t(z^t)} - 1 \right) \right\} \right]}{\Pi_{z^t} \{ c_t^a(z^t) \}}, \quad (10)$$

where in the numerator we used the assumption that the housing expenditure shares are identical across regions. In the model, we define the collateral ratio $my_t(z^t)$ as the ratio of housing wealth

to *total* wealth:

$$my_t(z^t) = \frac{\Pi_{z^t} \left[\left\{ c_t^a(z^t) \left(\frac{1}{\alpha_t(z^t)} - 1 \right) \right\} \right]}{\Pi_{z^t} \left[\left\{ c_t^a(z^t) \frac{1}{\alpha_t(z^t)} \right\} \right]}.$$

If the aggregate non-housing expenditure share is constant, the collateral ratio is constant at $1 - \alpha$. Suppose the aggregate endowment $\eta^a = c^a$ is constant as well. Then my or α index the risk-sharing capacity of the economy. When $\alpha = 1$, $my = 0$ is zero and there is no collateral in the economy. All the collateral constraints necessarily bind at all nodes and households are in autarchy.¹⁰ On the other hand, as α becomes sufficiently small, my becomes sufficiently large, and perfect risk sharing becomes feasible, because the solvency constraints no longer bind in any of the nodes s^t .

2.4 Tighter Constraints

A region is just a unit of aggregation. We define regional consumption as the sum of consumption of the households in a region:

$$c_t^i(\theta_0^{i1}, \theta_0^{i2}, y^t, z^t) = \sum_{j=1,2} c_t^{ij}(\theta_0^{ij}, x^t, y^t, z^t).$$

The regional consumption share is defined as a fraction of total non-durable consumption, as in the empirical analysis: $\hat{c}_t^i = \frac{c_t^i}{c_t^a}$.

The constraints faced by these households are tighter than those faced by a stand-in agent, who consumes regional consumption and earns regional labor income, in each region: By the linearity of the pricing functional $\Pi(\cdot)$, the aggregated regional collateral constraint for region i is just the sum of the household collateral constraints over households j in region i :

$$\begin{aligned} \sum_{j=1,2} \Pi_{s^t} [\{c_t^{ij}(\theta_0^{ij}, s^t) + \rho_t^i(y^t, z^t) h_t^{ij}(\theta_0^{ij}, s^t)\}] &= \Pi_{s^t} [\{c_t^i(\theta_0^{i1}, \theta_0^{i2}, y^t, z^t) + \rho_t^i(y^t, z^t) \chi_t^i(y_t, z^t)\}] \\ &\geq \sum_j \Pi_{s^t} [\{\eta_t^{ij}(x_t, y_t, z^t)\}] = \Pi_{s^t} [\{\eta_t^i(y_t, z^t)\}] \text{ for all } s^t \end{aligned}$$

This condition is necessary, but not sufficient: If household net wealth is non-negative in all states of the world for both households, then regional net wealth is too, but not vice-versa. In particular, it is the household in the $x = hi$ state whose constraint is crucial, not the average household's.

Regional consumption shares depend on the history of household-specific income shocks x^t , but only in a limited sense. The changes in the regional consumption shares $\hat{c}_t^i(x^t, y^t) = \frac{\xi_t^i(x^t, y^t, z^t)}{\xi_t^a(z^t)}$ are

¹⁰ *Proof:* If a set of households with non-zero mass had a non-binding solvency constraint at some node (x^t, y^t, z^t) , there would have to be another set of households with non-zero mass at node (x^t, y^t, z^t) that violate their solvency constraint.

governed by the growth rate of the regional weight ξ_t^i relative to that of the aggregate weights ξ_t^a . This is a measure of how constrained the households in this region are relative to the rest of the economy. When one of the households switches from the low to the high state, her weight increases, causing regional consumption to increase even when the regional income share stays constant ($\hat{\eta}_t^i$ increases but $\hat{\eta}^i$ may be constant). As we show in our simulations below, this is why the cross-sectional dispersion of regional consumption shares exceeds the cross-sectional dispersion of regional income shares. In section 2.7, we explain that this effect depends on the redistributive nature of idiosyncratic shocks at the household level. But because these household shocks are i.i.d across regions, their effects disappear when we integrate over all household-specific histories by the law of large numbers:

$$\int_{x^t \in X^t} \hat{c}_t^i(x^t, y^t) d\Pi(x^t) = \int_{x^t \in X^t} \frac{\xi^i(x^t, y^t)}{\xi_t^a} d\Pi(x^t) \simeq \hat{c}_t^i(y^t). \quad (11)$$

Even though the collateral constraints pertain to households and households within a region are heterogeneous, on average, the regional consumption share $\hat{c}_t^i(y^t)$ behaves as if it is the consumption share of a representative household in the region facing a single, but tighter, collateral constraint. This insight is quantitatively important. If we simply considered constraints at the regional level and calibrated the model to regional income shocks, the constraints would hardly bind. To an econometrician with only regional data generated by the model, it looks as if the stand-in agent's consumption share is subject to preference shocks or measurement error. These preference shocks follow from switches in the identity of the constrained household within the region. This provides one structural justification for our assumption of measurement error in regional consumption shares introduced in section 3.2.

2.5 Calibration

Preference Parameters We consider the case of separable utility by setting γ at 2 and ϵ at .5, the estimate of the intratemporal elasticity of substitution by Yogo (2006).¹¹ In the benchmark calibration, the discount factor β is set equal to .95. We also explore lower values for β .

Aggregate Endowment Processes Following Mehra and Prescott (1985), the aggregate non-housing endowment growth rate follows an AR(1) with mean 0.0183, standard deviation 0.0357, and autocorrelation -.14. It is discretized as a two-state Markov chain. The aggregate housing endowment process has the same average growth rate. Following Piazzesi, Schneider and Tuzel

¹¹Yogo estimates this elasticity off the cointegration relationship between the relative price of durables to non-durables and the quantities of durable and non-durable consumption.

(2007), we assume that the log of the aggregate non-housing expenditure ratio $\ell = \log\left(\frac{\alpha}{1-\alpha}\right)$ follows an autoregressive process:

$$\ell_t = \mu_\ell + .96 \log \ell_{t-1} + \epsilon_t,$$

with $\sigma_\epsilon = .03$ and μ_ℓ was chosen to match the average US post-war non-housing expenditure ratio of 4.41. Denote by \mathcal{L} the domain of ℓ .

Average Housing Collateral Ratio To keep the model as simple as possible, we abstracted from financial assets or other kinds of capital (such as cars) that households may use to collateralize loans. According to Flow of Funds data, 75% of household borrowing in the data is collateralized by housing wealth. However, to take into account other sources of collateral, we calibrate the collateral ratio to a broader measure of collateral than housing alone.

We use two approaches to calibrate the average US ratio of housing wealth to housing plus human wealth: a factor payments and an asset values approach. First, we examine the factor payments on both sources of wealth. Between 1946 and 2002, the average ratio of total US rental income to labor income (compensation of employees) plus rental income $\frac{\rho h}{\rho h + \eta^a}$ was 0.034 (data from NIPA Table 1.12). This measure of rental income includes imputed rents for owner-occupied housing. Second, we look at asset values (Flow of Funds data). Over the same period, the average ratio of US residential wealth to labor income is 1.66. We match this ratio in a stationary equilibrium with a collateral ratio of 0.025. Both approaches suggest a ratio smaller than five percent.

The above calculation ignores non-housing sources of collateral. A broader collateral measure also includes financial wealth as a source of collateral. Its factor payments are net dividends and interest payments by domestic corporations. We treat proprietary income as a flow to non-collateralizable human wealth. The factor payment ratio is now 0.08. In terms of asset values, the average ratio of the market value of US non-farm, non-financial corporations plus residential wealth to labor income is 2.68 (see Lustig and Van Nieuwerburgh (2009) for data construction). We match this ratio in a stationary equilibrium with a collateral ratio of 0.05. Both approaches suggest a collateral ratio smaller than ten percent.

We calibrate to the broad measure of collateral and set the average collateral ratio equal to 0.10. We scale up the quantity of labor income in the model to simultaneously match an average collateral ratio of 10 percent and a non-housing expenditure ratio of 4.41.

Region-Specific and Household-Specific Income We use a 5-state first-order Markov process to approximate the regional labor income share dynamics (Tauchen and Hussey 1991): $\log \hat{\eta}_t^i =$

.94 $\log \hat{\eta}_{t-1}^i + e_t^i$ with the standard deviation of the shocks σ_e set to 1 percent. The estimation details are in appendix B. We do not model permanent income differences between regions. Finally, as is standard in this literature, we use a 2-state Markov process to match the level of household labor income share $\hat{\eta}_t^i$ (as a fraction of regional labor income) dynamics. The persistence is .9 and the standard deviation of the shock is 3 percent (Heaton and Lucas 1996).

2.6 Computation of Markov Stationary Equilibrium

When aggregate shocks move the non-housing expenditure share α and the collateral ratio around, the joint measure over consumption shares and states changes over time. Instead of keeping track of the entire measure or the entire history of aggregate shocks in the state space, we compute policy functions that depend on a truncated history of aggregate weight shocks: $\vec{g}_k = [g_{-1}, g_{-2}, \dots, g_{-k}] \in \mathcal{G}$.¹²

We assign each household a label \hat{c} , which is this household's consumption share at the end of the last period. Let \mathcal{C} denote the domain of the normalized consumption weights. Consider a household of type 1. Its new consumption *weight* at the start of the next period follows the cutoff rule $\varpi^1(\hat{c}, x, y, \ell, \vec{g}_k) : \mathcal{C} \times X \times Y \times \mathcal{L} \times \mathcal{G} \rightarrow \mathcal{C}$:

$$\begin{aligned} \varpi^1(\hat{c}, x, y, \ell, \vec{g}_k) &= \hat{c} \text{ if } \hat{c} > \underline{\varpi}^1(x, y, \ell, \vec{g}_k) \\ &= \underline{\varpi}^1(x, y, \ell, \vec{g}_k) \text{ elsewhere,} \end{aligned}$$

where $\underline{\varpi}^1(x, y, \ell, \vec{g}_k)$ is the cutoff consumption share for which the collateral constraints hold with equality, or equivalently, net wealth is zero. The cutoff consumption share satisfies

$$C^1(\underline{\varpi}^1(x, y, \ell, \vec{g}_k), x, y, \ell, \vec{g}_k) = 0,$$

where $C^1(\hat{c}, x, y, \ell, \vec{g}_k) : \mathcal{C} \times X \times Y \times \mathcal{L} \times \mathcal{G} \rightarrow R^+$ is the net wealth function. The policy functions for a household of type 2 are defined analogously. Next period's consumption shares are:

$$\hat{c}' = \frac{\varpi^1(\hat{c}, x, y, \ell, \vec{g}_k)}{g},$$

where $g = \sum_{j=1,2} \int_{\mathcal{C} \times X \times Y \times \mathcal{L} \times \mathcal{G}} \varpi^j(\hat{c}, x, y, \ell, \vec{g}_k) d\Phi^j(\hat{c}, x, y, \ell, \vec{g}_\infty)$ is the actual aggregate weight shock. Let $\Phi^j(\hat{c}, x, y, \ell, \vec{g}_\infty)$ denote the joint measure over \hat{c} and (x, y) which depends on the *infinite* history of shocks, and let $\Xi(\ell, \vec{g}_\infty)$ denote the joint measure over ℓ and g .

Definition 2. An approximate k^{th} -order Markov stationary equilibrium consists of a forecasting func-

¹²The model tells us which moment of the distribution in the last period to keep track of: if many agents were severely constrained last period and g_{-1} was large, very few are constrained this period and g is small.

tion $g(\ell, \vec{g}_k)$, a measure $\Phi^j(\hat{c}, x, y; \ell, \vec{g}_\infty)$ for each type j and a policy function $\{\varpi^j(\hat{c}, x, y; \ell, \vec{g}_k)\}_{j=1,2}$ that implements the cutoff rule $\{\underline{\varpi}^j(x, y, \ell, \vec{g}_k)\}_{j=1,2}$, where the forecasting function has zero average prediction errors:

$$g(\ell, \vec{g}_k) = \sum_{j=1,2} \int_{\vec{g}_\infty | \vec{g}_k} \int_{C \times X \times Y \times \mathcal{L} \times \mathcal{G}} \varpi^j(\hat{c}, x, y; \ell, \vec{g}_k) d\Phi^j(\hat{c}, x, y; \ell, \vec{g}_\infty) d\Xi(\ell, \vec{g}_\infty)$$

To approximate the household's net wealth function $C(\cdot)$, we use 5th-degree Tchebychev polynomials in the two continuous state variables, the consumption weights ϖ and the log expenditure ratio ℓ . We compute a first-order Markov equilibrium with $k = 5$. The prediction errors are percentage deviations of *actual* from *spent* aggregate consumption. These approximation errors are small. They never exceed 1.9% in absolute value, they are .3% on average and their standard deviation is about .4%. The computation is accurate.

2.7 Results from Model Simulation

This section shows that the model generates an equilibrium distribution of regional consumption, income and housing collateral that closely resembles that in the data, and we do so by comparing the ratio of consumption-to-income dispersion in the model and the data. However, we want to emphasize that the sensitivity of consumption growth to income growth is a better measure of risk sharing, as we explain below. This is the risk sharing measure we focus on in section 3.

Our model generates the dispersion anomaly. Not only is the ratio of consumption-to-income dispersion greater than one on average, it also increases when collateral is scarce. We simulate a panel of $T = 15,000$ periods and $N = 100$ regions. On average, the ratio of housing wealth to total wealth, my , is 10%. In order to compare model and data more easily in the rest of the paper, we define a re-normalized collateral ratio that it is always positive: $\widetilde{my}_{t+1} = \frac{my^{max} - my_{t+1}}{my^{max} - my^{min}}$. The re-normalized housing collateral ratio \widetilde{my}_{t+1} is a measure of *collateral scarcity*; when the collateral ratio is at its maximum value $\widetilde{my} = 0$, whereas a reading of 1 means that collateral is at its lowest level. We construct \widetilde{my} by setting my^{max} and my^{min} equal to the maximum and minimum value in simulation. The resulting collateral scarcity measure \widetilde{my} is 0.71 on average.

Figure 1 shows the cross-sectional dispersion of regional consumption relative to the cross-sectional dispersion of regional income in the model. Two features are important. First, the model generates the dispersion anomaly. The average ratio of consumption-to-income dispersion exceeds one. In our model, that average is 1.22, while its is 1.28 for the 23 US Metropolitan Statistical Areas that we use in our 1952-2002 sample.

[Figure 1 about here.]

Second, when housing collateral is scarce, the cross-sectional consumption-to-income dispersion is higher. The ratio of consumption dispersion to income dispersion is almost twice as high when collateral scarcity is at its highest value in the simulation. We find the same variation in the data. The dashed line in Figure 2 plots the ratio of the regional cross-sectional consumption to income dispersion in the data. This measure falls by half between 1978 and 1988, while it doubles between 1988 and 1995 before falling back to its 1988 level in 2002. This stylized fact presents a new challenge to standard models, because it reveals that the departures from complete market allocations fluctuate over time. Conditioning on a measure of housing collateral helps to understand this aspect of consumption in the data. Our empirical measure of housing collateral scarcity (solid line) broadly tracks the variation in this regional consumption-to-income dispersion ratio. It is close to its highest level in 1978, falls by half between 1978 and 1988, increases again until 1996, and falls back to its 1988 level in 2002. Finally, the turning points in the cross-sectional dispersion of consumption coincide with the turning points in the housing collateral ratio. For example, between periods 325 and 375 the dispersion ratio increases by 40 percent, from .15 to .23 as the collateral scarcity increases from .5 to .9.

[Figure 2 about here.]

Understanding the Dispersion Anomaly Regional consumption is very sensitive to regional income shocks, in spite of the fact that most of the risk faced by households has been traded away in equilibrium, even at low collateral ratios. This is apparent in Figure 3. Its two panels contrast risk-sharing at the regional and at the household level. The upper panel plots the ratio of *regional* consumption dispersion to income dispersion, while the lower panel plots the same ratio but for *household* consumption and income. The dispersion measures are conditional cross-sectional standard deviations. The collateral scarcity measure is on the horizontal axis. Since the housing collateral ratio moves around over time, we display a scatter plot. As is apparent from the bottom panel of Figure 3, two-thirds of total household income risk is insured on average. The ratio of consumption to income dispersion at the household level is around one-third on average, well below one. Yet, in the top panel, the standard deviation of consumption to income dispersion at the regional level exceeds one when housing collateral is sufficiently scarce. What explains this dispersion anomaly?

[Figure 3 about here.]

First, the cross-sectional standard deviations of consumption shares (as a fraction of the regional endowment) at the household level, denoted \hat{c}^{ij} for household $j = 1, 2$ in region i , are smaller

than the cross-sectional standard deviation of the endowment shares (as a fraction of the regional endowment) as long as some risk sharing is feasible in equilibrium:

$$std(\hat{c}_{t+1}^{ij}) < std(\hat{\eta}_{t+1}^{ij}). \quad (12)$$

Second, the following inequality holds for the household consumption shares within a region:

$$std(\hat{c}^{i1} + \hat{c}^{i2}) > std(\hat{\eta}_{t+1}^{i1} + \hat{\eta}_{t+1}^{i2}) = 0, \quad (13)$$

where the last step follows because the endowment shares $\hat{\eta}^{i1} + \hat{\eta}^{i2} = 1$ add up to one at the regional level, but the consumption shares do not: $\hat{c}^{i1} + \hat{c}^{i2} \neq 1$. Hence, the sign reversal between equations (12) and (13) comes about because (i) the household income share shocks $\Delta \log \hat{\eta}_{t+1}^{ij}$ are perfectly negatively correlated across the households within region by construction, while (ii) the individual household weight shocks that result from these shocks are not perfectly negatively correlated, because of risk sharing. As a result, in much of the parameter space we find that the cross-sectional standard deviation of regional consumption (as a share of the aggregate endowment) exceeds that of income:

$$std(\hat{c}_{t+1}^i) > std(\hat{\eta}_{t+1}^i),$$

where $\hat{c}^i = (\hat{c}^{ij} + \hat{c}^{ij}) \hat{\eta}^i$. More generally, household-level income growth is more negatively correlated within a region than consumption growth *because of risk-sharing*. Therefore, when we aggregate from the household to the regional level, household risk sharing gives rise to regional consumption growth volatility that exceeds regional income growth volatility.

Figure 4 illustrates this aggregation result. The top panel plots the consumption shares as a fraction of the aggregate endowment \hat{c}_{t+1}^{ij} (full line) and income $\hat{\eta}_{t+1}^{ij}$ (dashed line) shares for two households $j = 1, 2$ living in the same region i . Because of risk-sharing, household-level consumption dispersion is lower than household-level income dispersion; the solid lines are closer together than the dashed lines. The bottom panel plots regional consumption \hat{c}_{t+1}^i and income shares $\hat{\eta}_{t+1}^i$ for the same region i and for the same simulated sequence of shocks. The negative correlation of income shocks reduces the volatility of regional income shares relative to regional consumption; the solid and dashed lines are about equally volatile.

[Figure 4 about here.]

The redistributive nature (negative correlation) of household-level income shocks is important. If there were a continuum of households and idiosyncratic shocks were completely independent, our aggregation effect would not be operative. Our two-state, two-agent specification with income states that reverse between agents is standard in the literature (Heaton and Lucas 1996). Also,

household-level income shocks within a region are unlikely to be purely idiosyncratic in the data, for example because there are shocks that disproportionately affect one industry or one sector.

The link between risk sharing and the ratio of consumption dispersion to income dispersion is not monotone. There are two off-setting effects. On the one hand, as the supply of housing collateral decreases, the dispersion of household consumption growth increases and it approaches the cross-sectional standard deviation of household income growth from below in equation (12). In the case of autarchy (no risk-sharing), the inequality becomes an equality. On the other hand, as the supply of collateral decreases, the cross-sectional standard deviation of regional consumption growth decreases and it approaches the standard deviation of regional income growth from above in equation (13). The latter effect is because regional consumption growth becomes more negatively correlated across households within a region. To see these two effects at work, we consider an economy without aggregate uncertainty; it grows at a constant rate. Figure 5 plots the ratio of consumption to income dispersion against the housing collateral ratio. Each dot represents a different equilibrium of an economy with a different collateral ratio. The graph reveals that, for housing collateral ratios below 11%, the first effect dominates and the regional consumption-to-income dispersion ratio decreases as the collateral supply increases. However, when the housing collateral ratio is above 11%, the second effect dominates and the regional consumption-to-income dispersion ratio increases with the housing collateral ratio. Importantly, this non-monotonicity does not affect the slope coefficient in a regression of regional consumption growth on income growth, and hence does not hamper our empirical work in Section 3. Figure 6 shows this slope for the same equilibria as in Figure 5. The elasticity of consumption to income shocks decreases monotonically as we increase the housing collateral ratio. This explains why we focus on this measure of risk sharing in the empirical section.

[Figure 5 about here.]

[Figure 6 about here.]

3 Testing the Collateral Mechanism

In this section we link our model to the traditional risk-sharing tests based on linear consumption growth regressions, the workhorse of the consumption insurance literature (Cochrane (1991), Mace (1991), Nelson (1994), Attanasio and Davis (1996), Blundell, Pistaferri and Preston (2008), and ensuing work).¹³ These regressions are a useful diagnostic of the key relationship between the

¹³Our paper also makes contact with the large literature on the excess sensitivity of consumption to predictable income changes, starting with Flavin (1981), who interpreted her findings as evidence for borrowing constraints, and followed by Hall and Mishkin (1982), Zeldes (1989), Attanasio and Weber (1995) and Attanasio and Davis (1996), all of which examine at micro consumption data.

degree of risk sharing and the scarcity of housing collateral that we set out to test. Section 3.1 describes the US metropolitan data that we use. Section 3.2 then estimates the linear consumption regressions in the data. Consistent with the regional risk-sharing literature that uses state level data (Van Wincoop (1996), Hess and Shin (1998), DelNegro (1998), Asdrubali, Sorensen and Yosha (1996), Athanasoulis and Wincoop (1998), and Del Negro (2002)), we reject full consumption insurance among US metropolitan regions. More importantly, and new to this literature, we find that collateral scarcity increases the correlation between income growth shocks and consumption growth. These collateral effects are economically significant. Finally, section 3.3 runs the same regressions, but on model-generated data. The size of the coefficients, and the regression R^2 in the model are similar to the ones in the data. In sum, we replicate the variation in the income elasticity of regional consumption growth that we document in the data.

The previous section delivered a formal theory of regional consumption weights ξ_{t+1}^i that tied the distribution of these weights to the housing collateral ratio. We saw that the weights followed a cut-off rule, where the cut-off depended on the current income shock η_{t+1}^i and the housing collateral ratio, in addition to the history of aggregate shocks. Equivalently, regions i 's consumption share in deviation from the cross-sectional average, $\hat{\xi}_{t+1}^i = \xi_{t+1}^i / \xi_{t+1}^a$, is a non-linear function of the region-specific income shock $\hat{\eta}_{t+1}^i$ and the housing collateral scarcity measure $\widetilde{m}y_{t+1}$. All growth rates of hatted variables denote the growth rates in the region in deviation from the cross-regional average, and the averages are population-weighted.

To make contact with the linear consumption growth regressions in the literature, we assume here that the growth rate of the log regional consumption share is linear in the product of the housing collateral ratio and the regional income share shock: $\Delta \log \hat{\xi}_{t+1}^i = -\gamma \widetilde{m}y_{t+1} \Delta \log \hat{\eta}_{t+1}^i$. Under our assumption of separable preferences, this assumption delivers a linear consumption growth equation which simply involves regional income share growth interacted with the collateral ratio:

$$\Delta \log \hat{c}_{t+1}^i = \widetilde{m}y_{t+1} \Delta \log \hat{\eta}_{t+1}^i. \quad (14)$$

The interpretation is straightforward. If $\widetilde{m}y_{t+1}$ is zero, this region's consumption growth equals aggregate consumption growth. There is perfect insurance. On the other hand, if $\widetilde{m}y_{t+1}$ is one, this region's consumption wedge is at its largest, and the region is in autarchy: its non-housing consumption c_t^i (growth) equals its labor income η_t^i (growth). While simple, this specification captures the important features of the link between consumption, income, and housing collateral in the model. Put differently, this linear specification of the consumption weights turns out to work well inside the model.

3.1 Data

We construct a new data set of US metropolitan area level macroeconomic variables, as well as standard aggregate macroeconomic variables. All of the series are annual for the period 1951-2002.

We believe that metropolitan area data are a good choice to study the question of risk-sharing and the role of housing collateral. First, metropolitan area data have not been used before to study risk-sharing and are an interesting addition to the literature. Second, compared to state-level data, each MSA is a relatively homogenous region in terms of rental price shocks. Since we do not have good data on household-level variation in housing prices, metropolitan areas are a natural choice. If housing prices are strongly correlated within a region, there are only small efficiency gains from looking at household instead of regional consumption data if the objective is to identify the collateral effect. Second, many have argued that household level data contain substantial measurement error (e.g., Cogley (2002)). Aggregation to the regional level should alleviate this problem.

Aggregate Macroeconomic Data We use two distinct measures of the nominal housing collateral stock HV : the market value of residential real estate wealth (HV^{rw}) and the net stock current cost value of owner-occupied and tenant occupied residential fixed assets (HV^{fa}). The first series is from the Flow of Funds (Federal Board of Governors) for 1945-2002 and from the Bureau of the Census (Historical Statistics for the US) prior to 1945. The last series is from the Fixed Asset Tables (Bureau of Economic Analysis) for 1925-2001. Appendix C provides detailed sources. HV^{rw} is a measure of the value of residential housing owned by households, while HV^{fa} which is a measure of the total value of residential housing. Real per household variables are denoted by lower case letters. The real, per household housing collateral series hv^{rw} and hv^{fa} are constructed using the all items consumer price index from the Bureau of Labor Statistics, p^a , and the total number of households from the Bureau of the Census. Aggregate nondurable and housing services consumption, and labor income plus transfers data are from the National Income and Product Accounts (NIPA). Real per household labor income plus transfers is denoted by η^a and real per capita aggregate consumption is c^a .

Measuring the Housing Collateral Ratio In the model the housing collateral ratio my is defined as the ratio of collateralizable housing wealth to housing wealth plus non-collateralizable human wealth.¹⁴ In Lustig and Van Nieuwerburgh (2005), we show that the log of real per household real

¹⁴Human wealth is an unobservable. We assume that the non-stationary component of human wealth H is well approximated by the non-stationary component of labor income Y . In particular, $\log(H_t) = \log(Y_t) + \epsilon_t$, where ϵ_t is a stationary random process. This is the case if the expected return on human capital is stationary (see Jagannathan and Wang (1996) and Campbell (1996)). The housing collateral ratio then is measured as the deviation from the co-integration relationship between the value of the aggregate housing collateral measure and aggregate labor income.

estate wealth ($\log hv$) and labor income plus transfers ($\log \eta$) are non-stationary in the data. This is true for both hv^{rw} and hv^{fa} . We compute the housing collateral ratio as $myhv = \log hv - \log \eta$ and remove a constant and a trend. The resulting 1925-2002 time series $myrw$ and $myfa$ are mean zero and stationary, according to an ADF test. Formal justification for this approach comes from a likelihood-ratio test for co-integration between $\log hv$ and $\log \eta$ (Johansen and Juselius (1990)). We refer the reader to Lustig and Van Nieuwerburgh (2005) for details of the estimation. The trend removal is necessary to end up with a stationary variable that can be used in the regression analysis below. We discuss the implication of the trend in the housing wealth-to-income ratio for risk-sharing in the conclusion. The housing collateral ratios display large and persistent swings between 1925 and 2002. The correlation between $myrw$ and $myfa$ is 0.86. In the empirical work, we construct the collateral scarcity measures \widetilde{myrw} and \widetilde{myfa} by setting my^{max} and my^{min} equal to the respective 1925-2002 sample maximum and minimum of $myrw$ and $myfa$.

Regional Macroeconomic Data We construct a new panel data set for the 30 largest metropolitan areas in the US. The regions combine for 47 percent of the US population. The metropolitan data are annual for 1951-2002. Thirteen of the regions are metropolitan statistical areas (MSA). The other seventeen are consolidated metropolitan statistical areas (CMSA), comprised of adjacent and integrated MSA's. Most CMSA's did not exist at the beginning of the sample. For consistency we keep track of all constituent MSA's and construct a population weighted average for the years prior to formation of the CMSA. We use regional sales data to measure non-durable consumption. Sales data have been used by DelNegro (1998) at the state level, but never at the metropolitan level. The appendix compares our new data to other data sources that partially overlap in terms of sample period and definition, and we find that they line up. The elimination of regions with incomplete data leaves us with annual data for 23 metropolitan regions from 1951 until 2002. We denote real per capita regional income and consumption by η^i and c^i , and we define consumption and income *shares* as the ratio of regional to aggregate consumption and income: $\hat{c}_t^i = \frac{c_t^i}{c_t}$ and $\hat{\eta}_t^i = \frac{\eta_t^i}{\eta_t}$. The details concerning the consumption, income and price data we use are in the data appendix C.

3.2 Linear Consumption Growth Regressions in Data

To bring the theory to the data, we consider the consumption growth regression in equation (14). In all regressions, we include regional fixed effects to pick up unobserved heterogeneity across regions, and we take into account measurement error in non-durable consumption. We express observed consumption shares with a tilde and assume that income shares are measured without error. The

linear model collapses to the following equation for observed consumption shares \tilde{c} :

$$\Delta \log (\tilde{c}_{t+1}^i) = a_0^i + a_1 \widetilde{my}_{t+1} \Delta \log (\hat{\eta}_{t+1}^i) + \nu_{t+1}^i,$$

where the left hand side variable is *observed consumption share* growth and a_0^i are region-specific fixed effects. All measurement error terms are absorbed in ν_{t+1}^i . This equation resembles the standard consumption growth equation in the consumption literature, except for the collateral interaction term. We can rewrite this specification once more with a separate regional income growth term, using the actual housing collateral ratio instead of the collateral scarcity measure \widetilde{my}_{t+1} :

$$\Delta \log (\tilde{c}_{t+1}^i) = b_0^i + b_1 \Delta \log (\hat{\eta}_{t+1}^i) + b_2 my_{t+1} \Delta \log (\hat{\eta}_t^i) + \nu_{t+1}^i.$$

The parameter b_1 in the second specification corresponds to $a_1 \frac{my^{max}}{my^{max} - my^{min}}$ in the first specification and the coefficient b_2 corresponds to $-a_1 \frac{1}{my^{max} - my^{min}}$. We focus on the estimation results for this second specification.¹⁵

Estimation Specifics We assume that the measurement error in regional consumption share growth, ν_t^i , is orthogonal to lagged values housing collateral ratio: $E [\nu_t^i \widetilde{my}_{t-k}] = 0, \forall k \geq 0$. Since only aggregate variables affect the aggregate housing collateral ratio my and only region-specific measurement error enters in ν^i , this assumption follows naturally from the theory.

The benchmark estimation method is generalized least squares (GLS), which takes into account cross-sectional correlation in the residuals ν^i and heteroscedasticity. If the residuals and regressors are correlated, the GLS estimators of the parameters in the consumption growth regressions are inconsistent. To address this possibility, we report instrumental variables estimation results (by three-stage least squares) in addition to the GLS results. Because of the autoregressive nature of \widetilde{my} , we use two, three and four-period leads of the dependent and independent variables as instruments (Arellano and Bond (1991)).

The estimation results are in table 1. The first two lines report the results for the entire sample 1952-2002 and the two different collateral measures. Lines 3-4 report the results for the 1970-2002 sub-sample; lines 5-6 use labor income plus transfers, only available for 1970-2000, instead of disposable income. Finally, lines 7-8 report the instrumental variables (IV) estimates.

[Table 1 about here.]

First, the null hypothesis of full insurance among U.S. regions, $H_0 : b_1 = b_2 = 0$, is strongly rejected. The p -value for a Wald test is 0.00 for all rows in table 1. This is consistent with the findings of the regional risk-sharing literature for US states (see e.g. Hess and Shin (1998)).

¹⁵A previous version of the paper presented consistent results across both specifications.

Second, the correlation of region-specific consumption growth and region-specific income growth is higher when housing collateral is scarce: $b_2 < 0$ is negative in all rows. The coefficient b_2 is estimated precisely in most rows. The coefficients b_1 and b_2 , together with the average housing collateral ratio, imply that one-third of disposable income growth shocks end up in consumption growth, while two-thirds of shocks are insured away on average. Most importantly, there is substantial time variation in the degree of risk sharing depending on the level of the collateral ratio. For example, the estimates in row 2 imply that the income elasticity of consumption share growth varies between .58, when $my = my^{min} = -.124$, and .13, when $my = my^{max} = +.13$, using $myfa$ as the collateral measure. The fifth percentile value for $myrw$ and the coefficient on $[b_1, b_2]$ in row 1 imply a degree of risk-sharing of 42 percent. The 95th percentile implies a degree of risk-sharing of 86 percent. The time variation is stronger in the 1970-2000 period and estimated more precisely, regardless of which income measure we use (rows 3-6). Rows 7-8 of table 1 report IV estimates where income changes are instrumented by 2 and 3-period leads of independent and dependent variables. The instrumental variables estimates reject full insurance, and the coefficient estimates are close to the ones obtained by GLS. Again, these lend support to the collateral channel. Overall, the point estimates imply large shocks to the regional risk sharing technology in the US induced by changes in the housing collateral ratio.

3.3 Linear Consumption Growth Regression in Model

Finally, we use the same simulation to re-estimate the consumption share growth regressions that we ran on the regional consumption share data in section 3.2. The results are reported in Table 2.

The slope coefficients vary between $[\.38, -1.59]$, for $\beta = .95$, and $[\.62, -1.88]$, for $\beta = .75$. Because my is .10 on average in the simulation, the average fraction of income shocks that ends up in consumption is 22% for $\beta = .95$. That implies that 78% of income risk is insured on average. For $\beta = .75$, the average fraction of risk that is shared among regions is 57%. The 66% estimate for the average fraction of income risk shared in the data (see Table 1) corresponds to a value for β between .95 and .90. More importantly, the slope coefficients imply a lot of time-variation in the degree of risk sharing. In the model, the 5th and 95th percentile of \widetilde{my} are .55 and .95. That distribution implies a 90% confidence interval for the degree of risk-sharing of $[69, 83]$ percent for $\beta = .95$ and $[48, 66]$ percent for $\beta = .75$.

The estimates reveals that the income elasticity coefficient in the model-generated sample varies between $-.04$ when $my = my^{max}$ and $.34$ when $my = my^{min}$, in the case of $\beta = .95$. In the case of $\beta = .75$, the coefficient varies between $.09$ and $.54$. In the data, the slope coefficients varied between $.28$ and $.45$ (see Table 1). Also, the regression R^2 are close to those in the data, around 7%. They are low because regional risk is small compared to household risk.

[Table 2 about here.]

To understand the regression results, recall that in equilibrium, the growth rate of the regional consumption shares is determined by the difference between the growth rates of the regional weight and the growth rate of the aggregate weight: $\Delta \log(\hat{c}_{t+1}^i) = \Delta \log \xi_{t+1}^i - \Delta \log \xi_{t+1}^a$. As argued in section 2.4, $\Delta \log \xi_{t+1}^i$ only responds to regional income shocks on average ($\Delta \log \hat{\eta}_{t+1}^i$). The effect of household-specific shocks x is absorbed in the regression error term ν_{t+1}^i . The slope coefficients in Table 2 reflect two forces. First, in case of a positive shock to household or regional income, the cutoff shares ξ_{t+1}^i are much higher when housing collateral is scarce. Second, in case of a negative income shock, the household consumption shares drift down at a higher rate $\Delta \log \xi_{t+1}^a$ in the low collateral economy. The same logic applies to the regional consumption shares because it is the sum of the shares for the two types of households. The effects are more pronounced for lower discount rates.

4 Additional Evidence for Collateral Channel

In this section, we provide additional support for the housing collateral mechanism. First, our empirical results continue to hold for a non-separable utility function specification. Second, we find evidence that the degree of risk-sharing is also tied to regional collateral measures. Using regional measures of the housing collateral stock to sort regions into bins, we find that the income elasticity of consumption growth for regions in the lowest housing collateral quartile of US metropolitan areas is more than twice the size of the same elasticity for areas in the highest quartile, and their consumption growth is only half as correlated with aggregate consumption growth. Linear consumption growth regressions that use regional instead of aggregate collateral measures produce similar results. Third, we look at province data for Canada and find the same positive relationship between housing collateral and consumption insurance, both for aggregate and regional collateral measures.

4.1 Non-Separable Utility

Our previous results are robust to the inclusion of expenditure share growth terms which arise from the non-separability of the utility function. The point estimates for the slope coefficients on income growth interacted with the collateral ratio are very similar, but the expenditure share growth terms are not significant. The results are reported in a separate appendix, downloadable from the authors' web sites.

4.2 Estimation of the Linear Model using Regional Collateral Measures

While solving a model where the housing collateral ratio is different across regions is beyond the scope of the current paper, we find support in the data for a similar relationship between regional consumption data and *regional* measures of collateral.

For each of the US metropolitan areas we construct a measure of regional housing collateral, combining information on regional repeat sale price indices with Census estimates on the housing stock. The data construction of the regional housing wealth follows Case, Quigley and Shiller (2001) and is detailed in appendix C.4. The regional housing collateral ratios for each metropolitan area are constructed in the same way as the national measure, but from regional housing wealth and regional income measures. In the consumption growth regressions below, we also use the regional home ownership rate as a second measure of housing collateral.

To explore the cross-sectional variation in housing collateral, we conduct two exercises. First, we sort the 23 MSA's by their collateral ratio in each year and look at average population-weighted consumption growth and income growth for the 6 regions with the lowest and the 6 regions with the highest regional collateral ratio. Table 3 shows the results. Regions in the first group (highest collateral scarcity, \widetilde{my}^i is 0.84 on average, reported in column 1) experience more volatile consumption growth (column 2) that is only half as correlated with US aggregate consumption growth (column 3) than for the group with the most abundant collateral (\widetilde{my}^i is 0.23 on average). The last three columns report the result of a time-series regression of group-averaged consumption share growth on group-averaged income share growth. The income elasticity of consumption share growth is 0.66 (with t-stat 1.9) for the group with the most scarce collateral, whereas it is only 0.32 (with t-stat 1.3) for the group with the most abundant collateral. For the first group full insurance can be rejected, whereas for the last group it cannot.

[Table 3 about here.]

Second, we estimate linear consumption growth regression results for the case of separable preferences:

$$\Delta \log (\hat{c}_{t+1}^i) = b_0^i + b_1 \Delta \log (\hat{\eta}_{t+1}^i) + b_2 X_{t+1}^i \Delta \log (\hat{\eta}_{t+1}^i) + \nu_{t+1}^i.$$

Table 4 presents the results. The regional collateral measure X^i is the home-ownership rate in region i in the first row and the *regional* housing collateral ratio my^i in the second row. For both variables, we find that the correlation between consumption and income share growth is lower when the region-specific collateral measure is higher. The effects are large and the coefficients are precisely measured. For example, the region-specific collateral measures $X^i = my^i$ vary between -.25 and .25. The implied variation in the degree of risk sharing is between 45 and 74 percent.

This paper is not about a direct housing wealth effect on regional consumption: For an average unconstrained household that is not about to move, there is no reason to consume more when its housing value increases, simply because it has to live in a house and consume its services (see Sinai and Souleles (2005) for a clear discussion). In the third row of the table, we add the regional collateral measure as a separate regressor to check for a regional housing wealth effect on consumption. The coefficient, b_3 , is significant, but it has the wrong sign. After controlling for the risk-sharing role of housing, we find no separate increase in regional consumption growth when regional housing collateral becomes more abundant. In sum, regions consume more when total regional labor income increases and this effect is larger when housing wealth is *smaller* relative to human wealth in that region.

We also used bankruptcy indicators as a regional collateral measure and found that they were insignificant. US states have different levels of homestead exemptions that households can invoke upon declaring bankruptcy under Chapter 7. We used both the amount of the exemption and a dummy for MSA's in a state with an exemption level above \$20,000. In neither regression did we find a significant coefficient.

[Table 4 about here.]

Finally, measurement error may be a concern for the regional consumption data. However, as long as the the standard deviation of consumption measurement error does not systematically increase in times or regions with scarce collateral, measurement error would bias the coefficient estimates downwards, strengthening the case for the collateral mechanism in US regional data.

4.3 Canadian Data

As a robustness check, we repeat the analysis with data from Canadian provinces. While we only have data available for ten provinces from 1981-2003, the consumption data are arguably more standard. The data are on non-durable consumption (personal expenditures on goods and services less expenditures on durable goods) instead of retail sales. The income measure is personal disposable income. We construct real per capita consumption and income shares, using the provincial CPI series. The housing wealth series measure the market value of the net stock of fixed residential capital, a measure corresponding to hv^{fa} . These housing wealth series are available for Canada, as well as for the ten provinces. The housing collateral ratio is constructed in the same way as for the U.S. data. Appendix C.5 describes these data in more detail.

[Table 5 about here.]

Table 5 confirms our finding for the U.S. that the degree of risk-sharing varies substantially with the housing collateral ratio. In the first row, we use the aggregate collateral ratio. Since \widetilde{myfa} is .5 on average and $myfa$ is zero on average, they show that Canadian provinces share 85% of income risk on average. This is higher than in the U.S., presumably because there is more government redistribution. More importantly, the degree of risk sharing varies over time. When housing collateral is at its lowest point in the sample (in 1985), only 63% of income risk is shared, whereas in 2003, the degree of risk-sharing is 95%. In rows 2 and 3 we use the same collateral measure, but now measured at the regional level. Again we find a precisely estimated slope coefficient with the right sign. Lastly, we confirm our finding for the U.S. data, that these results are not driven by a wealth effect. In row 3, the coefficient on the housing collateral ratio b_3 shows up with the wrong sign.

Finally, in UK data, Campbell and Cocco (2007) also find evidence in favor of a collateral effect on regional consumption using aggregate measures of housing wealth.

5 Concluding Remarks

The availability of housing collateral significantly impacts regional risk sharing. We construct a new data set of consumption and income data for the largest US metropolitan areas. Not only do we reject perfect consumption insurance among these regions, we also find that times in which collateral is scarce are associated with significantly less risk-sharing. Canadian data show similar patterns. This time-varying degree of risk-sharing is a new stylized fact that standard models are unable to address.

A model with limited commitment and default resulting in the loss of housing collateral generates the same positive co-movement between the consumption-to-income dispersion ratio and housing collateral scarcity. Importantly, it jointly generates the dispersion anomaly: the fact that the consumption-to-income dispersion ratio is above one on average, and why this ratio co-moves positively with housing collateral scarcity. To generate this dispersion anomaly, the model has two dimensions of heterogeneity: households and regions. This structure enables us to translate a modest friction at the household level into a substantial deviations of perfect risk-sharing at the regional level.

This approach is useful because it provides a single explanation for the apparent lack of consumption insurance at different levels of aggregation. But it differs from most of the work in regional or international risk sharing which adopts the representative agent paradigm. That literature typically relies on frictions impeding the international flow of capital resulting from the government's ability to default on international debt or to tax capital flows (e.g. Kehoe and Perri (2002)), or resulting

from transportation costs (e.g. Obstfeld and Rogoff (2003)). Such frictions cannot account for the lack of risk sharing between regions within a country or between households within a region.

The collateral mechanism explored here may also help explain low-frequency patterns in household risk-sharing. In recent work, Krueger and Perri (2006) document that the dramatic increase in labor income inequality in the US between 1970 and 2002 was not accompanied by a similar increase in household consumption inequality. Our housing collateral effect seems consistent with these trends in household consumption and income inequality. In the US, the raw ratio of residential wealth to labor income increased from 1.4 in 1980 to 1.9 in 2002 and the ratio of mortgages to income increased from .45 to .80. A persistent increase in housing collateral of that magnitude would give a substantial boost to risk sharing and bring about a reduction in the cross-sectional dispersion of consumption relative to income.

References

- Alvarez, Fernando and Urban Jermann**, "Efficiency, Equilibrium, and Asset Pricing with Risk of Default.," *Econometrica*, 2000, 68 (4), 775–798.
- Arellano, Manuel and Stephen Bond**, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, 1991, 58, 277–297.
- Asdrubali, Pierfederico, Bent Sorensen, and Oved Yosha**, "Channels of Interstate Risk Sharing: United States 1963-1990," *Quarterly Journal of Economics*, 1996, 111, 1081–1110.
- Athanasoulis, Stefano and Eric Van Wincoop**, "Risk Sharing Within the United States: What Have Financial Markets and Fiscal Federalism Accomplished?," *Working Paper Federal Reserve Bank of New York*, 1998, 9808.
- Atkeson, Andy and Tamim Bayoumi**, "Do Private Capital Markets Insure Regional Risk? Evidence from the United States and Europe.," *Open Economies Review*, 1993, 4, 303–324.
- Attanasio, Orazio P. and Guglielmo Weber**, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 1995, 103 (6), 1121–57.
- and **Steven J. Davis**, "Relative Wage Movements and the Distribution of Consumption," *The Journal of Political Economy*, December 1996, 104 (6), 1127–1262.
- Backus, David, Patrick Kehoe, and Finn Kydland**, "International Real Business Cycles," *Journal of Political Economy*, 1992, 100 (4), 745–75.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008. Forthcoming.
- Campbell, John Y.**, "Understanding Risk and Return," *The Journal of Political Economy*, April 1996, 104 (2), 298–345.
- and **Joao Cocco**, "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics*, 2007, 54(3), 591–621.
- Case, Karl E. and Robert J. Shiller**, "Prices of Single-Family Homes Since 1970: New Indexes for Four Cities," *New England Economic Review*, September/October 1987, pp. 46–56.

- , **John M. Quigley**, and **Robert J. Shiller**, “Comparing Wealth Effects: The Stock Market versus The Housing Market,” October 2001. UC Berkeley Working Papers.
- Chien, Yi-Li and Hanno Lustig**, “The Market Price of Aggregate Risk and the Wealth Distribution,” *Review of Financial Studies*, 2009, *forthcoming*.
- Cochrane, John H.**, “A Simple Test of Consumption Insurance,” *The Journal of Political Economy*, October 1991, 99 (5), 957–976.
- Cogley, Timothy**, “How Fast Can the New Economy Grow? A Bayesian Analysis of the Evolution of Trend Growth,” August 2002. Mimeo Arizona State University.
- Crucini, Mario J.**, “On International and National Dimensions of Risk Sharing,” *Review of Economics and Statistics*, 1999, 81, 73–84.
- Davis, Morris and Francois Ortalo-Magne**, “Household Expenditures, Wages, Rents,” November 2007. Working Paper University of Maddison, Wisconsin.
- Del Negro, Marco**, “Asymmetric Shocks Among U.S. States,” *Journal of International Economics*, 2002, 56, 273–297.
- DelNegro, Marco**, “Aggregate Risk Sharing Across US States and Across European Countries,” January 1998. Yale University mimeo.
- Eichenbaum, Martin and Lars Peter Hansen**, “Estimating Models with Intertemporal Substitution Using Aggregate Time Series Data,” *Journal of Business and Economic Statistics*, 1990, 8 (1), 53–69.
- Flavin, Marjorie**, “The Adjustment of Consumption to Changing Expectations About Future Income,” *Journal of Political Economy*, 1981, 89 (5), 974–1009.
- Hall, Robert E. and Frederic S. Mishkin**, “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households,” *Econometrica*, 1982, 50, 461–481.
- Heaton, John and Deborah Lucas**, “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing,” *Journal of Political Economy*, 1996, 104 (3), 443–87.
- Hess, Gregory D. and Kwanho Shin**, “Intranational Business Cycles in the United States,” *Journal of International Economics*, April 1998, 44 (2), 289–313.
- and — , “Risk Sharing by Households Within and Across Regions and Industries,” *Journal of Monetary Economics*, 2000, 45, 533–560.

- Jagannathan, Ravi and Zhenyu Wang**, “The Conditional CAPM and the Cross-Section of Expected Returns,” *Journal of Finance*, 1996, 51, 3–54.
- Johansen, Soren and Katarina Juselius**, “Maximum Likelihood Estimation and Inference on Cointegration with Applications to the Demand for Money,” *Oxford Bulletin of Economics and Statistics*, May 1990, 52, 169–210.
- Kehoe, Patrick J. and Fabrizio Perri**, “International Business Cycles with Endogenous Incomplete Markets,” *Econometrica*, May 2002, 70 (3), 907–928.
- Kehoe, Timothy and David Levine**, “Debt-Constrained Asset Markets.,” *Review of Economic Studies*, 1993, 60, 865–888.
- Krueger, Dirk**, “Risk Sharing in Economies with Incomplete Markets.” PhD dissertation, University of Minnesota 1999.
- ___ **and Fabrizio Perri**, “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *The Review of Economic Studies*, 2006, 73, 163–193.
- Lewis, Karen**, “What Can Explain The Apparent Lack of International Consumption Risk-Sharing?,” *Journal of Political Economy*, 1996, 104, 267–297.
- Lucas, Robert**, “Asset Prices in an Exchange Economy,” *Econometrica*, 1978, 46 (6), 1429–54.
- Lustig, Hanno and Stijn Van Nieuwerburgh**, “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 2005, 60 (3), 1167–1219.
- ___ **and ___**, “The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street,” *Review of Financial Studies*, 2009, 21(5), 2097–2137.
- Mace, Barbara J.**, “Full Insurance in the Presence of Aggregate Uncertainty,” *The Journal of Political Economy*, October 1991, 99 (5), 928–956.
- Marcet, Albert and Ramon Marimon**, “Recursive Contracts,” 1999. working paper Universitat Pompeu Fabra.
- Mehra and Edward Prescott**, “The Equity Premium: A Puzzle.,” *Journal of Monetary Economics*, 1985, 15 (2), 145–161.
- Nelson, Julie A.**, “On Testing for Full Insurance using Consumer Expenditure Survey Data,” *The Journal of Political Economy*, April 1994, 102 (2), 384–394.

- Obstfeld, Maurice and Kenneth Rogoff**, “The Six Major Puzzles in International Economics: Is There a Common Cause?,” *Working Paper Harvard and UC Berkeley*, 2003.
- Ortalo-Magne, Francois and Sven Rady**, “Homeownership: Low Household Mobility, Volatile Housing Prices, and High Income Dispersion,” November 2002. Working Paper University of Maddison Wisconsin.
- and —, “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints,” *Review of Economic Studies*, 2006, 73, 459–485.
- Pavan, Marina**, “Consumer Durables and Risky Borrowing: The Effects of Bankruptcy Protection,” *Working Paper Boston College*, 2005.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel**, “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 2007, 83(3), 531–569.
- Pollakowski, Henry O.**, “Data Sources for Measuring House Price Changes,” *Journal of Housing Research*, 1995, 6 (3), 377–389.
- Sinai, Todd and Nicholas Souleles**, “Owner-Occupied Housing as a Hedge Against Rent Risk,” *Quarterly Journal of Economics*, 2005, 120(2), 763–789.
- Tauchen, George and Robert Hussey**, “Quadrature-based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, 1991, 59, 371–396.
- Uhlig, Harald**, “A Law of Large Numbers for Large Economies,” *Economic Theory*, 1996, 8, 41–50.
- Van Wincoop, Eric**, “Regional Risk Sharing,” *European Economic Review*, 1996, 40, 219–254.
- Yogo, Motohiro**, “A Consumption-Based Explanation of Expected Stock Returns,” *Journal of Finance*, 2006, 61 (2), 539–580.
- Zeldes, Stephen**, “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy*, 1989, 97, 305–346.

A Technical Appendix

This appendix spells out the household problem in an economy where all trade takes place at time zero.

Household Problem A household of type (θ_0^{ij}, s_0) purchases a complete contingent consumption plan $\{c_0^{ij}(\theta_0^{ij}, s_0), h^{ij}(\theta_0^{ij}, s_0)\}$ at time-zero market state prices $\{p, p\rho^i\}$. The household solves:

$$\sup_{\{c^{ij}, h^{ij}\}} U(c^{ij}(\theta_0^{ij}, s_0), h^{ij}(\theta_0^{ij}, s_0))$$

subject to the time-zero budget constraint

$$\Pi_{s_0} \left[\left\{ c_0^{ij}(\theta_0^{ij}, s_0) + \rho^i(s_0) h^{ij}(\theta_0^{ij}, s_0) \right\} \right] \leq \theta_0^{ij} + \Pi_{s_0} \left[\left\{ \eta_0^{ij} \right\} \right],$$

and an infinite sequence of collateral constraints for each t and s^t

$$\Pi_{s^t} \left[\left\{ c_t^{ij}(\theta_0^{ij}, s^t) + \rho^i(s^t) h_t^{ij}(\theta_0^{ij}, s^t) \right\} \right] \geq \Pi_{s^t} \left[\left\{ \eta_t^{ij}(s^t) \right\} \right], \forall s^t.$$

Dual Problem Given Arrow-Debreu prices $\{p, p\rho^i\}$ the household with label (θ_0^{ij}, s_0) minimizes the cost $C(\cdot)$ of delivering initial utility w_0^{ij} to itself:

$$\begin{aligned} C(w_0^{ij}, s_0) = & \min_{\{c, h\}} \left(c_0^{ij}(w_0^{ij}, s_0) + h_0^{ij}(w_0^{ij}, s_0) \rho_0^i(s_0) \right) \\ & + \sum_{s^t} p(s^t | s_0) \left(c_t^{ij}(w_0^{ij}, s^t | s_0) + h_t^{ij}(w_0^{ij}, s^t | s_0) \rho_t^i(s^t | s_0) \right) \end{aligned}$$

subject to the promise-keeping constraint

$$U_0(\{c^{ij}\}, \{h^{ij}\}; w_0^{ij}, s_0^i) \geq w_0^{ij}$$

and the collateral constraints

$$\Pi_{s^t} \left[\left\{ c_t^{ij}(w_0^{ij}, s^t) + \rho_t^i(s^t) h_t^{ij}(w_0^{ij}, s^t) \right\} \right] \geq \Pi_{s^t} \left[\left\{ \eta_t^{ij}(s^t) \right\} \right], \forall s^t.$$

The initial promised value w_0^{ij} is determined such that the household spends its entire initial wealth: $C(w_0^{ij}, s_0) = \theta_0^{ij} + \Pi_{s_0} [\{\eta^{ij}(s^t)\}]$. There is a monotone relationship between θ_0^{ij} and w_0^{ij} . The above problem is a standard, convex programming problem. We set up the saddle point problem and then make it recursive by defining cumulative multipliers (Marcet and Marimon (1999)). Let ν^{ij} be the Lagrange multiplier on the promise keeping constraint and $\gamma_t^{ij}(w_0^{ij}, s^t)$ be the Lagrange multiplier on the collateral constraint in history s^t . Define a cumulative multiplier at each node: $\zeta_t^{ij}(w_0, s^t) = 1 - \sum_{s^t} \gamma_t^{ij}(w_0^{ij}, s^t)$. Finally, we rescale the market state price $\hat{p}_t(s^t) = p_t(z^t) / \beta^t \pi_t(s^t | s_0)$. By using Abel's partial summation formula and the law of iterated expectations to the Lagrangian, we obtain an objective function that is a function of the cumulative multiplier process ζ^i :

$$D(c, h, \zeta^{ij}; w_0^{ij}, s_0) = \sum_{t \geq 0} \sum_{s^t} \left\{ \beta^t \pi(s^t | s_0) \left[\zeta_t^{ij}(w_0^{ij}, s^t | s_0) \hat{p}_t(s^t) \left(c_t^{ij}(w_0^{ij}, s^t) + \rho_t^i(s^t) h_t^{ij}(w_0^{ij}, s^t) \right) + \gamma_t^{ij}(w_0^{ij}, s^t) \Pi_{s^t} [\{\eta^i\}] \right] \right\}$$

such that

$$\zeta_t^{ij}(w_0^{ij}, s^t) = \zeta_{t-1}^{ij}(w_0^{ij}, s^{t-1}) - \gamma_t^{ij}(w_0^{ij}, s^t), \zeta_0^{ij}(w_0^{ij}, s_0) = 1$$

Then the recursive dual saddle point problem is given by:

$$\inf_{\{c^{ij}, h^{ij}\}} \sup_{\{\zeta^{ij}\}} D(c^{ij}, h^{ij}, \zeta^{ij}; w_0^{ij}, s_0)$$

such that

$$\sum_{t \geq 0} \sum_{s^t} \beta^t \pi(s^t | s_0) u(c_t^{ij}(w_0^{ij}, s^t), h_t^{ij}(w_0^{ij}, s^t)) \geq w_0^{ij}$$

To keep the mechanics of the model in line with standard practice, we re-scale the multipliers. Let

$$\xi_t^{ij}(\nu, s^t) = \frac{\nu^{ij}}{\zeta_t^{ij}(w_0^{ij}, s^t)},$$

The cumulative multiplier $\xi^{ij}(\nu, s^t)$ is a non-decreasing stochastic sequence, which is initialized at ν^{ij} at time zero. We can use ν^{ij} as the household label. If the constraint for household (ν^{ij}, s_0) binds, it goes up, else it stays put. This follows immediately from the complementary slackness condition for the solvency constraint.

Optimal Non-Housing Consumption The first order condition for $c(\nu^{ij}, s^t)$ is :

$$\hat{p}_t(s^t) = \xi_t^{ij}(\nu^{ij}, s^t) u_c(c_t^{ij}(\nu^{ij}, s^t), h_t^{ij}(\nu^{ij}, s^t)).$$

Upon division of the first order conditions for any two households ij and kl , the following restriction on the joint evolution of marginal utilities over time and across states must hold:

$$\frac{u_c(c_t^{ij}(\nu^{ij}, s^t), h_t^{ij}(\nu^{ij}, s^t))}{u_c(c_t^{kl}(\nu^{kl}, s^t), h_t^{kl}(\nu^{kl}, s^t))} = \frac{\xi_t^{kl}(\nu^{kl}, s^t)}{\xi_t^{ij}(\nu^{ij}, s^t)}. \quad (15)$$

Growth rates of marginal utility of non-durable consumption, weighted by the multipliers, are equalized across agents:

$$\frac{\xi_{t+1}^{ij}(\nu^{ij}, s^{t+1})}{\xi_t^{ij}(\nu^{ij}, s^t)} \frac{u_c(c_{t+1}^{ij}(\nu^{ij}, s^{t+1}), h_{t+1}^{ij}(\nu^{ij}, s^{t+1}))}{u_c(c_t^{ij}(\nu^{ij}, s^t), h_t^{ij}(\nu^{ij}, s^t))} = \frac{\hat{p}_{t+1}(s^{t+1})}{\hat{p}_t(s^t)} = \frac{\xi_{t+1}^{kl}(\nu^{kl}, s^{t+1})}{\xi_t^{kl}(\nu^{kl}, s^t)} \frac{u_c(c_{t+1}^{kl}(\nu^{kl}, s^{t+1}), h_{t+1}^{kl}(\nu^{kl}, s^{t+1}))}{u_c(c_t^{kl}(\nu^{kl}, s^t), h_t^{kl}(\nu^{kl}, s^t))}.$$

In the case of separable preferences between non-housing and housing consumption, there is a simple mapping from the multipliers ξ at s^t to the equilibrium allocations of both commodities. We refer to this mapping as the risk-sharing rule:

$$c_t^{ij}(\nu^{ij}, s^t) = \frac{\xi_t^{ij}(\nu^{ij}, s^t)^{\frac{1}{\gamma}}}{\xi_t^a(z^t)} c_t^a(z^t) \quad (16)$$

where

$$\xi_t^a(z^t) = \sum_{j=1,2} \sum_{x^t, y^t} \int \left(\xi_t^{ij}(\nu^{ij}, s^t) \right)^{\frac{1}{\gamma}} \frac{\pi(x^t, y^t, z^t | x_0, y_0, z_0)}{\pi(z^t | z_0)} d\Phi_0^j,$$

where Φ_0^j is the cross-sectional joint distribution over initial consumption weights and initial endowments for a household of type j . By the law of large numbers, the aggregate weight process only depends on the aggregate history z^t . It is easy to verify that this rule satisfies the optimality condition and the market clearing conditions.

The time zero ratio of marginal utilities is pinned down by the ratio of multipliers on the promise-keeping constraints. For $t > 0$, it tracks the stochastic weights ξ . From the first order condition w.r.t. $\xi_t^{ij}(\nu^{ij}, s^t)$ and the complementary

slackness conditions, we obtain a reservation weight policy:

$$\xi_t^{ij} = \xi_{t-1}^{ij} \text{ if } \xi_{t-1}^{ij} > \underline{\xi}_t(x_t, y_t, z^t), \quad (17)$$

$$\xi_t^{ij} = \underline{\xi}_t(x_t, y_t, z^t) \text{ otherwise.} \quad (18)$$

where the cutoff $\underline{\xi}_t$ is defined such that the collateral constraints hold with equality:

$$\Pi_{s^t} \left[\left\{ c_t^{ij}(\nu^{ij}, s^t; \underline{\xi}_t(\nu^{ij}, s^t)) + \rho^j(s^t) h^j(\nu^{ij}, s^t; \underline{\xi}_t(\nu^{ij}, s^t)) \right\} \right] = \Pi_{s^t} \left[\left\{ \eta_t^{ij}(s^t) \right\} \right].$$

The history-independence of the cutoff is established in proposition 3.

Optimal Housing Consumption The risk-sharing rule for housing services also follows a cutoff rule:

$$h_t^{ij}(s^t) = \frac{\left(\xi_t^{ij}(s^t) \right)^{\frac{1}{\gamma}}}{\xi_t^i(x^t, y^t, z^t)} \chi_t^i(y^t, z^t), \quad (19)$$

where the denominator is now the *regional weight shock*, defined as

$$\xi_t^i(x^t, y^t, z^t) = \sum_{j=1,2} \left(\xi_t^{ij}(s^t) \right)^{\frac{1}{\gamma}}.$$

To minimize notation, we dropped the ν in the ξ functions. Given this risk sharing rule and the form of the utility function, the regional rental price for any region i is given by:

$$\rho_t^i = \psi \left(\frac{h_t^{ij}}{c_t^{ij}} \right)^{\frac{-1}{\varepsilon}} = \psi \left(\frac{\xi_t^a \chi_t^i}{\xi_t^i c_t^a} \right)^{\frac{-1}{\varepsilon}}$$

We now verify that this risk-sharing rule clears the housing market in each region and satisfies the first order condition for housing services consumption.

Proof. First, note that these risk sharing rules clear the housing market in each region because $(\xi_t^{i1}(s^t))^{\frac{1}{\gamma}} + (\xi_t^{i2}(s^t))^{\frac{1}{\gamma}} = \xi_t^i$ by definition. Second, we check that it satisfies the first order condition for non-durable and durable consumption:

$$\begin{aligned} \xi_t^{ij} u_c(c_t^{ij}(s^t), h_t^{ij}(s^t)) &= \hat{p}_t(s^t | s_0) \\ \xi_t^{ij} u_h(c_t^{ij}(s^t), h_t^{ij}(s^t)) &= \rho_t^i(y^t, z^t) \hat{p}_t(s^t | s_0) \end{aligned}$$

Recall that the marginal utility of non-housing consumption and housing consumption are:

$$\begin{aligned} u_c(c_t^{ij}(s^t), h_t^{ij}(s^t)) &= (c_t^{ij})^{\frac{-1}{\varepsilon}} \left[(c_t^{ij})^{\frac{\varepsilon-1}{\varepsilon}} + \psi (h_t^{ij})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{1-\varepsilon\gamma}{\varepsilon-1}} \\ u_h(c_t^{ij}(s^t), h_t^{ij}(s^t)) &= \psi (h_t^{ij})^{\frac{-1}{\varepsilon}} \left[(c_t^{ij})^{\frac{\varepsilon-1}{\varepsilon}} + \psi (h_t^{ij})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{1-\varepsilon\gamma}{\varepsilon-1}} \end{aligned}$$

In the case of separability, $\varepsilon = \frac{1}{\gamma}$, and the marginal utility of housing services becomes: $u_h(c_t^{ij}(s^t), h_t^{ij}(s^t)) = \psi (h_t^{ij})^{\frac{-1}{\varepsilon}}$.

Substituting this into the optimality condition for housing produces the following expression:

$$\xi_t^{ij} \psi (h_t^{ij})^{-\frac{1}{\varepsilon}} = \xi_t^{ij} \psi \left[\frac{(\xi_t^{ij})^{\frac{1}{\gamma}}}{\xi_t^a} \chi_{i,t} \right]^{-\frac{1}{\varepsilon}} = \psi \left[\frac{\chi_{i,t}}{\xi_t^a} \right]^{-1/\varepsilon} = \rho_t^j(y^t, z^t) \hat{p}_t(s^t | s_0)$$

where the second equality follows from inserting the risk sharing rule for housing services, and the last equality follows from separability, $\gamma = \frac{1}{\varepsilon}$. Likewise, inserting the risk sharing rule for non-durable consumption into the optimality condition gives:

$$\xi_t^{ij} \psi \left[\frac{(\xi_t^{ij})^{\frac{1}{\gamma}}}{\xi_t^a} c_t^a \right]^{-\frac{1}{\varepsilon}} = \left[\frac{c_t^a}{\xi_t^a} \right]^{-\frac{1}{\varepsilon}} = \hat{p}_t(s^t | s_0)$$

Dividing through by the last line of the preceding equation, we obtain the following result: $\rho_t^j = \psi \left(\frac{\xi_t^a \chi_t^j}{\xi_t^j c_t^a} \right)^{-\frac{1}{\varepsilon}}$ for any household j in region i . This is exactly the rental price we conjectured at the start, together with the risk sharing rule, which confirms that the risk sharing rule satisfies the first order condition for optimality. The risk sharing rule also clears the housing market in every region and it clears the market for non-durable consumption. \square

The Non-Housing Expenditure Share The non-housing expenditure share is the same for all households j in region i :

$$\frac{c_t^{ij}}{c_t^{ij} + \rho_t^j h_t^{ij}} \equiv \alpha_t^{ij} \equiv \alpha_t^i.$$

Proof. To show this, we use the equilibrium risk-sharing rule for non-housing and housing consumption, as well as the expression for ρ_t^j to obtain:

$$\begin{aligned} \alpha_t^{ij} &= \frac{\frac{\xi_t^{ij} (\nu^{ij, s^t})^{\frac{1}{\gamma}}}{\xi_t^a(z^t)} c_t^a(z^t)}{\frac{\xi_t^{ij} (\nu^{ij, s^t})^{\frac{1}{\gamma}}}{\xi_t^a(z^t)} c_t^a(z^t) + \psi \left[\frac{\xi_t^a}{\xi_t^j(y^t, z^t)} \frac{\chi_t^j(y^t, z^t)}{c_t^a(z^t)} \right]^{-\frac{1}{\varepsilon}} \frac{(\xi_t^{ij}(s^t))^{\frac{1}{\gamma}}}{\xi_t^j(y^t, z^t)} \chi_t^j(y^t, z^t)} \\ &= \frac{1}{1 + \psi \left[\frac{\xi_t^a}{\xi_t^j(y^t, z^t)} \frac{\chi_t^j(y^t, z^t)}{c_t^a(z^t)} \right]^{-\frac{1}{\varepsilon}}} \end{aligned}$$

Note that this expression is the same for all households j in region i . \square

Assumption 2 imposes that the regional shares α^j only depend on the aggregate history z^t : $\alpha_t^j = \alpha_t(z^t)$. Hence, we assume that the ratio $\frac{\xi_t^j}{\xi_t^a} c_t^a = \chi_t^j$ for all regions, and all aggregate histories. Note that all regions have the same rental price as well, as a result of this assumption.

History Independence of the Cutoff Rule

Proposition 3. *In a state with a binding collateral constraint, the equilibrium consumption share, $\hat{c}_t^{ij} = \frac{c_t^{ij}}{c_t^a}$, only depends on (x_t, y_t) and z^t .*

Proof. When the collateral constraint binds for household ij ,

$$\begin{aligned} \Pi_{st} \left[\left\{ c_t^{ij}(w_0^{ij}, s^t) \left[1 + \frac{\rho_t^i(s^t) h_t^{ij}(w_0^{ij}, s^t)}{c_t^{ij}(w_0^{ij}, s^t)} \right] \right\} \right] &= \Pi_{st} \left[\left\{ \eta_t^{ij}(x_t, y_t, z^t) \right\} \right], \\ \Pi_{st} \left[\left\{ \hat{c}_t^{ij} c_t^a(z^t) \frac{1}{\alpha_t(z^t)} \right\} \right] &= \Pi_{st} \left[\left\{ \hat{\eta}_t^{ij}(x_t) \hat{\eta}_t^i(y_t, z_t) \eta_t^a(z^t) \right\} \right], \end{aligned}$$

where the second line follows from the definition of the non-housing expenditure share, and we use assumption 2.

Obviously, the right hand side does not depend on (x^{t-1}, y^{t-1}) , only on (x_t, y_t) . Fix an arbitrary aggregate history z^t . We can take two households with histories $(x^{t-1}I, x_t, y^{t-1}I, y_t)$ and $(x^{t-1}II, x_t, y^{t-1}II, y_t)$. The right hand side is the same for both, because the labor endowment share process is first order Markov in (x, y, z) (see assumption 1), and the pricing functional only depends on z^t . So, the left hand side has to be the same for both regions as well. Since the non-housing expenditure share only depends on the aggregate history z^t , this immediately implies that the household's consumption share \hat{c}_t^{ij} can only depend on (x_t, y_t, z^t) when the collateral constraint binds. \square

B Calibration of Regional Labor Income Shocks

We use the regional data set described in appendix (C) to calibrate the persistence of the regional income share process, used in section 2.5. We estimate an AR(1) process for the log disposable income share between 1952 and 2002:

$$\begin{aligned} \log \hat{\eta}_{t+1}^i &= .9434 \log \hat{\eta}_t^i + \nu_{t+1}^i \\ &\quad (0.0092) \quad (0.0286) \end{aligned}$$

If we introduce fixed effects, to correct for permanent income differences, the slope coefficient drops to .85. Based on these estimates, we set the AR(1) coefficient equal to 0.94 and the standard deviation of the innovation equal to 0.01. We use the Tauchen and Hussey (1991) method to discretize the AR(1) process into a 5-state Markov chain. The grid points are

$$[-0.0879, -0.0440, 0, 0.0440, 0.0879]$$

and the transition matrix is:

$$\begin{bmatrix} 0.9526 & 0.0474 & 0.0000 & 0 & 0 \\ 0.0069 & 0.9666 & 0.0265 & 0.0000 & 0 \\ 0.0000 & 0.0140 & 0.9721 & 0.0140 & 0.0000 \\ 0.0000 & 0.0000 & 0.0265 & 0.9666 & 0.0069 \\ 0.0000 & 0.0000 & 0.0000 & 0.0474 & 0.9526 \end{bmatrix}$$

Likewise, we calibrate the household income share process (as a fraction of the regional income), $\hat{\eta}^{i,j}$, as a two state Markov chain. The states are $[\cdot 6, \cdot 1.4]$ and the transition matrix is $[\cdot 9, \cdot 1; \cdot 1, \cdot 9]$.

Table 1: Income Growth Elasticity of Consumption Shares in Data

Coll. Measure	b_1	σ_{b_1}	b_2	σ_{b_2}	R^2
1 <i>myrw</i>	.35	(.03)	-.30	(.26)	6.5
2 <i>myfa</i>	.36	(.03)	-1.74	(.50)	6.8
3 <i>myrw</i>	.33	(.02)	-.64	(.17)	4.7
4 <i>myfa</i>	.37	(.02)	-2.12	(.31)	5.0
5 <i>myrw</i>	.48	(.02)	-1.03	(.23)	10.5
6 <i>myfa</i>	.51	(.03)	-1.13	(.30)	10.4
7 <i>myrw</i>	.31	(.04)	-.32	(.38)	
8 <i>myfa</i>	.32	(.04)	-1.75	(.64)	

Notes: We estimate: $\Delta \log(\tilde{c}_{t+1}^i) = b_0^i + b_1 \Delta \log(\hat{\eta}_{t+1}^i) + b_2 my_{t+1} \Delta \log(\hat{\eta}_{t+1}^i) + \nu_{t+1}^i$. Rows 1-2 are for the period 1952-2002 (1166 observations). Rows 3-4 are identical to rows 1-2 but are for the period 1970-2002 (759 observations). The measure of regional income is disposable personal income in rows 1-4 and 7-8. Regressions 5-6 use labor income plus transfers, available only for 1970-2000. In each block, the rows use the variables *myrw* and *myfa*, estimated for the period 1925-2002. my^{max} (my^{min}) is the sample maximum (minimum) in 1925-2002. The coefficients on the fixed effect are not reported. Estimation is by feasible Generalized Least Squares, allowing for both cross-section heteroscedasticity and contemporaneous correlation. Rows 7-8 are the results for the instrumental variable estimation by 3SLS. Instruments are a constant, $\log(\hat{\eta}_{t+2}^i)$, $\log(\hat{\eta}_{t+3}^i)$, $\log(\hat{\eta}_{t+4}^i)$, $\Delta \hat{\rho}_{t+2}^i$, $\Delta \hat{\rho}_{t+3}^i$, $\Delta \hat{\rho}_{t+4}^i$, $\log(\tilde{c}_{t+2}^i)$, $\log(\tilde{c}_{t+3}^i)$, $\log(\tilde{c}_{t+4}^i)$, and my_{t+2} , my_{t+3} , my_{t+4} . The sample is 1952-1998 (1051 observations). All results are for 23 US metropolitan areas.

Table 2: Income Growth Elasticity of Consumption Shares in Model

β	b_1	b_2	R^2	my_{min}	my_{max}	$mean(my)$
.95	0.385	-1.596	0.077	0.026	0.267	.106
.90	0.552	-1.498	0.074	0.034	0.284	.106
.85	0.553	-1.434	0.068	0.034	0.266	.106
.75	0.628	-1.883	0.071	0.042	0.277	.106

Notes: The sample is a model-simulated panel for 1000 years (annual data) and 100 regions with $\gamma = 2$, $\epsilon = .5$ and the AR(1) process for the non-housing expenditure share in equation (2.5). Each row corresponds to a different value of the time discount factor β . We estimate: $\Delta \log(\tilde{c}_{t+1}^i) = b_0^i + b_1 \Delta \log(\hat{\eta}_{t+1}^i) + b_2 my_{t+1} \Delta \log(\hat{\eta}_{t+1}^i) + \nu_{t+1}^i$. The first 3 columns report the slope coefficient and the regression's R^2 . The three last columns of the table report the min, max and mean of the collateral ratio my_t over the simulated sample. The mean of my is .10 and the mean of \tilde{my} is .71.

Table 3: Cross-Regional Variation in Collateral.

	$\widetilde{m}y^i$	$std(\Delta \log(c_t^i))$	$corr(\Delta \log(c_t^i), \Delta \log(c_t^2))$	<i>Slope</i>	[<i>t-stat</i>]	R^2
1	0.842	0.033	0.257	0.659	[1.896]	.13
2	0.577	0.032	0.233	0.354	[0.987]	.04
3	0.407	0.018	0.278	0.472	[1.757]	.11
4	0.226	0.028	0.502	0.319	[1.283]	.06

Notes: Quartiles ranked from high to low collateral scarcity. The sample is 1975-2000 (annual data). All results are for 23 US metropolitan areas sorted each year into quartiles based on that period/region's collateral scarcity measure $\widetilde{m}y_t^i$. The first column reports the average collateral scarcity over the sample for each quartile. The second column reports the standard deviation of average population-weighted non-durable consumption growth in each quartile. The third column reports the correlation with real per capita US non-durable consumption growth (NIPA). The fourth column reports the slope coefficient in a time series regression of average population-weighted consumption share growth on average population-weighted income share growth for each quartile. The regional income measure is disposable personal income. The fifth and sixth columns reports the t-stat and regression R^2 .

Table 4: Risk-Sharing Tests with Regional Collateral Measures.

Coll. Measure	b_1	σ_{b_1}	b_2	σ_{b_2}	b_3	σ_{b_3}	R^2
1 HO^i	.45	(.02)	-.11	(.03)			6.1
2 my^i	.40	(.02)	-.57	(.12)			6.2
3 my^i	.39	(.02)	-.45	(0.14)	-0.03	(0.003)	6.6

Notes: Rows 1 and 2 of the table reports estimation results for $\Delta \log(\hat{c}_{t+1}^i) = b_0^i + b_1 \Delta \log(\hat{\eta}_{t+1}^i) + b_2 X_{t+1}^i \Delta \log(\hat{\eta}_{t+1}^i) + \nu_{t+1}^i$. Rows 3 of the table reports estimation results for $\Delta \log(\hat{c}_{t+1}^i) = b_0^i + b_1 \Delta \log(\hat{\eta}_{t+1}^i) + b_2 X_{t+1}^i \Delta \log(\hat{\eta}_{t+1}^i) + b_3 X_{t+1}^i + \nu_{t+1}^i$. In row 1, X^i is the region-specific home-ownership rate (575 observations). In row 2 and row 3, $X^i = my^i$ is the region-specific housing collateral ratio (569 observations). It is measured as the residual from a regression of the log ratio of real per capita regional housing wealth to real per capita labor income, $\log(hv_t^i) - \log(\eta_t^i)$, on a constant and a time trend. A higher my^i means more abundant collateral in region i . In all regressions η is disposable income. The coefficients on the fixed effect b_0^i is not reported. Estimation is by feasible Generalized Least Squares allowing for both cross-section heteroscedasticity and contemporaneous correlation. All regressions are for the period 1975-2000 for 23 US metropolitan areas, the longest period with metropolitan housing data.

Table 5: Risk-Sharing Tests with Canadian Data.

		Panel A						
Coll. Measure		b_1	σ_{b_1}	b_2	σ_{b_2}	R^2		
1	$myfa$.15	(.02)	-2.03	(.41)	37.3		
2	my^i	.18	(.02)	-.83	(.28)	34.9		
		Panel B: Wealth Effect						
Coll. Measure		b_1	σ_{b_1}	b_2	σ_{b_2}	b_3	σ_{b_3}	R^2
3	my^i	.18	(.02)	-.78	(0.29)	-0.008	(0.002)	35.1

Notes: Row 1 (panel A) reports estimation results for $\Delta \log (\hat{c}_{t+1}^i) = b_0^i + b_1 \Delta \log (\hat{\eta}_{t+1}^i) + b_2 X_{t+1}^i \Delta \log (\hat{\eta}_{t+1}^i) + \nu_{t+1}^i$. Finally, row 3 (panel B) reports estimation results for $\Delta \log (\hat{c}_{t+1}^i) = b_0^i + b_1 \Delta \log (\hat{\eta}_{t+1}^i) + b_2 X_{t+1}^i \Delta \log (\hat{\eta}_{t+1}^i) + b_3 X_{t+1}^i + \nu_{t+1}^i$. Rows 1 uses the aggregate collateral measure for Canada $myfa$. In rows 2 and 3, X^i is the regional collateral measure my^i in Canadian province i . Both the aggregate and regional housing collateral ratios are measured as the residual from a regression of the log ratio of real per capita regional housing wealth to real per capita labor income on a constant and a time trend. The coefficients on the fixed effect, a_0^i or b_0^i are not reported. Estimation is by feasible Generalized Least Squares allowing for both cross-section heteroscedasticity and contemporaneous correlation. All regressions are for the period 1981-2003 for 10 Canadian provinces. The panel contains 220 observations.




Figure 1: Housing Collateral Scarcity and Consumption/Income Dispersion in Model.

The figure plots a simulated time path for $T = 500$ of the collateral scarcity measure $\tilde{m}y$ (solid line, measured against the right axis) against the ratio of regional consumption dispersion to regional income dispersion (dashed line, measured against the left axis). The parameters are $\gamma = 2, \epsilon = .5, \beta = .95$. The average collateral ratio is 10 percent.




Figure 2: Housing Collateral Scarcity and Consumption/Income Dispersion in Data.

This figure plots the ratio of regional consumption-to-income dispersion (dashed line, plotted against the right axis). Both consumption and income are measured in deviation from the cross-regional mean. The solid line is our collateral scarcity measure, plotted against the left axis. The sample consists of annual data from 1952 until 2002 for 23 US Metropolitan Statistical Areas. The data are discussed in section 3.1.




Figure 3: The Dispersion Anomaly.

Scatter diagram of collateral and consumption-to-income dispersion ratios. The upper panel is for regions, while the lower panel is for households. The figure plots, for a simulated time path ($T = 2,500$), the collateral scarcity measure $\bar{m}y_t$ on the horizontal axis against the ratio of consumption dispersion (cross-sectional standard deviation of regional consumption in levels) to income dispersion (cross-sectional standard deviation of regional income in levels) on the vertical axis. In the upper panel, consumption and income dispersion are measured at the regional level, in the lower panel at the household level. The parameters are $\gamma = 2, \epsilon = .5, \beta = .95$. The average collateral ratio is 10 percent.




Figure 4: Household and Regional Consumption Dynamics.

Simulation of 100 observations from equilibrium for benchmark economy. The parameters are $\gamma = 2, \epsilon = .5, \beta = .95$. The average collateral ratio is 10 percent. The top panel plots household consumption \hat{c}_{t+1}^{ij} (full line) against household income \hat{y}_{t+1}^{ij} (dotted line) as a share of the aggregate endowment for household $j = 1, 2$ in region i . The bottom panel plots regional consumption \hat{c}_{t+1}^i (full line) against regional income \hat{y}_{t+1}^i (dashed line) as a share of the aggregate endowment for region i .




Figure 5: Collateral Supply and Consumption/Income Dispersion.

Scatter Plot of the ratio of the cross-sectional standard deviation of regional consumption to the cross-sectional standard deviation of regional income $\frac{std(\hat{c}^i)}{std(\hat{\eta}^i)}$ against the collateral ratio my . Simulation from steady-state equilibria for an economy without aggregate risk. Each dot represents an equilibrium for the economy with the housing collateral ratio displayed on the horizontal axis. The parameters are $\gamma = 2, \epsilon = .5, \beta = .95$. The collateral ratio my varies from 1 percent to 18 percent.




Figure 6: Risk Sharing and Sensitivity of Consumption Growth to Income Growth.

Scatter Plot of the slope coefficient in a consumption growth regression against the collateral ratio my . We run the following cross-sectional regression: $\Delta \log \hat{c}_{t+1}^i = a_0 + a_1 \Delta \log \hat{\eta}_{t+1}^i + \epsilon_{t+1}^i, i = 1, \dots, 5000$ with a panel of 5000 households. The figure plots a_1 against my . Simulation from steady-state equilibria for an economy without aggregate risk. Each dot represents an equilibrium for the economy with the housing collateral ratio displayed on the horizontal axis. The parameters are $\gamma = 2, \epsilon = .5, \beta = .95$. The collateral ratio my varies from 1 percent to 18 percent.