

- Christiano, Lawrence J., and Martin S. Eichenbaum, "Temporal Aggregation and Structural Inference in Macroeconomics," NBER working paper no. T0060 (1986).
- Christiano, Lawrence J., and Terry J. Fitzgerald, "The Band Pass Filter," Northwestern University manuscript (1999).
- Cogley, Timothy, and James M. Nason, "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research," *Journal of Economic Dynamics and Control* 19 (1995), 253-278.
- Cooley, Thomas F., and Gary D. Hansen, "The Inflation Tax in a Real Business Cycle Model," *American Economic Review* 79:4 (1989), 733-748.
- Cooley, Thomas J., and Lee E. Ohanian, "The Cyclical Behavior of Prices," *Journal of Monetary Economics* 28:1 (1991), 25-60.
- Correia, Isabel H., Joao L. Neves, and Sergio T. Rebelo, "Business Cycles from 1850-1950: New Facts about Old Data," *European Economic Review* 36:2/3 (1992), 459-467.
- Danthine, Jean-Pierre, and John B. Donaldson, "Methodological and Empirical Issues in Real Business Cycle Theory," *European Economic Review* 37:1 (1993), 1-35.
- Danthine, Jean-Pierre, and Michel Girardin, "Business Cycles in Switzerland: A Comparative Study," *European Economic Review* 33:1 (1989), 31-50.
- Fiorito, Ricardo, and Tryphon Kollintzas, "Stylized Facts of Business Cycles in the G7 from a Real Business Cycles Perspective," *European Economic Review* 38:2 (1994), 235-269.
- Hansen, Gary D., "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics* 16:3 (1985), 309-327.
- Harvey, Andrew C., and A. Jaeger, "Detrending, Stylized Facts and the Business Cycle," *Journal of Applied Econometrics* 8 (1993), 231-247.
- Hassler, John, Petter Lundvik, Torsten Persson, and Paul Söderlind, "The Swedish Business Cycle: Stylized Facts over 130 Years," monograph no. 22, IIES, Stockholm University (1992).
- Hodrick, Robert J., and Edward C. Prescott, "Postwar U.S. Business Cycles: An Empirical Investigation," Carnegie Mellon University discussion paper no. 451 (1980).
- , "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking* 29:1 (1997), 1-16.
- King, Robert G., and Sergio T. Rebelo, "Low Frequency Filtering and Real Business Cycles," *Journal of Economic Dynamics and Control* 17:1/2 (1993), 207-231.
- Kydland, Finn E., and Edward C. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica* 50:6 (1982), 1345-1370.
- , "Business Cycles: Real Facts and a Monetary Myth," *Federal Reserve Bank of Minneapolis Quarterly Review* 14 (Spring 1990), 3-18.
- Lucas, Robert E., Jr., "Understanding Business Cycles" (pp. 7-29), in K. Brunner and Allan H. Meltzer (Eds.), *Stabilization of the Domestic and International Economy*, Carnegie-Rochester Conference Series 5 (Amsterdam: North-Holland, 1977).
- Ravn, Morten O., and Martin Sola, "Stylized Facts and Regime Changes: Are Prices Procyclical?" *Journal of Monetary Economics* 36:3 (1995), 497-526.
- Ravn, Morten O., and Harald Uhlig (1997), "On Adjusting the HP-Filter for the Frequency of Observations," CentER discussion paper no. 50.
- , "On Adjusting the HP-Filter for the Frequency of Observations," CEPR discussion paper no. 2858 (2001).
- Romer, Christina, "The Prewar Business Cycle Reconsidered: New Estimates of GNP 1869-1908," *Journal of Political Economy* 97:1 (1989), 1-37.
- Söderlind, Paul, "Cyclical Properties of a Real Business Cycle Model," *Journal of Applied Econometrics* 9 (1994), S113-S122.
- Tobin, James, *Asset Accumulation and Economic Activity* (Chicago: University of Chicago Press, 1980).

IDIOSYNCRATIC RISK AND VOLATILITY BOUNDS, OR CAN MODELS WITH IDIOSYNCRATIC RISK SOLVE THE EQUITY PREMIUM PUZZLE?

Martin Lettau*

I. Introduction

RECENTLY, there has been a lot of interest in computing asset prices in incomplete market models; see, for example, Constantinides and Duffie (1996), Heaton and Lucas (1996), den Haan (1996), Krusell and Smith (1997) and Storesletten, Telmer, and Yaron (1997). These papers have shown that market incompleteness can affect prices of financial assets qualitatively. In this paper, I propose a simple method to check whether these effects are quantitatively important enough to solve the equity premium puzzle.

Received for publication April 20, 1999. Revision accepted for publication May 10, 2001.

* Federal Reserve Bank of New York and Centre for Economic Policy Research.

This paper was written during a visit at the Department of Economics at New York University; I am grateful for its hospitality. John Campbell, Mark Gertler, Blake LeBaron, Sydney Ludvigson, Anthony Lynch, Harald Uhlig, two anonymous referees, and seminar participants at Humboldt University, New York University, and the Federal Reserve Bank of New York provided helpful comments. The views are those of the author and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

The main argument is as follows. Most incomplete market models specify endogenous endowment (labor income) shocks that are not fully insurable. Agents are allowed to trade in a small number of securities and solve for their optimal portfolio and consumption policies. It is difficult to test these types of models directly because the quality of household-level consumption data is very poor.¹ Instead of this direct approach using consumption data, I use data on individual income, which is measured more precisely than is individual consumption. In other words, I assume that agents cannot smooth idiosyncratic income shocks at all and are forced to consume their endowment. If agents were allowed to trade using some restricted set of securities, they would be able to smooth, at least partially, their individual shocks. Hence, the income process provides an upper bound on the volatility of individual consumption. If models with idiosyncratic risk are not able to generate large risk premia, they will most likely not be able to perform better with consumption data. I find even very volatile income shocks

¹ One exception is Cogley (1998).

are not able to quantitatively to generate high-enough risk premia.

The paper is organized as follows. First, I briefly demonstrate how idiosyncratic shocks can affect the Sharpe ratio.² Intuitively, idiosyncratic shocks can affect asset prices if their distribution depends on aggregate state variables (Mankiw, 1986; Constantinides & Duffie, 1996). Second, I use the estimated processes for idiosyncratic income estimated by Heaton and Lucas (1996), Krusell and Smith (1997), and Storesletten et al. (1997) and compute the Sharpe ratios as if agents had to consume their endowment. I find that none of the idiosyncratic income processes is able to generate a Sharpe ratio that is found in the data.³

II. The Sharpe and Idiosyncratic Shocks

Let C_t^j denote consumption of consumer j in period t . Lowercase letters denote log variables, so c_t^j is log consumption of agent j . M_{t+1}^j is the stochastic discount factor (SDF) of agent j , R_{t+1} is the return of an asset, and RP_{t+1} is the risk premium. The first-order condition is

$$E_t[M_{t+1}^j R_{t+1}] = 1. \quad (1)$$

The Sharpe ratio is defined as the price of risk of the market portfolio and depends on the volatility of the SDF and its correlation with the market portfolio:

$$\begin{aligned} SR(M_{t+1}^j) &= \max_{\text{all assets}} \frac{E_t[RP_{t+1}]}{\sigma_t[RP_{t+1}]} \\ &= -\rho_t(M_{t+1}^j, RP_{t+1}^m) \frac{\sigma_t[M_{t+1}^j]}{E_t[M_{t+1}^j]} \leq \frac{\sigma_t[M_{t+1}^j]}{E_t[M_{t+1}^j]}. \end{aligned} \quad (2)$$

Note that, in an economy with complete markets, the SDF of the representative agent is perfectly negatively correlated with the market portfolio, so that the Sharpe ratio is determined only by the volatility of the SDF. This is, however, no longer true in a model with idiosyncratic shocks because returns on aggregate assets cannot be correlated with idiosyncratic shocks. Hence, despite the fact that the SDF of an individual might be very volatile, the net effect on the Sharpe ratio is ambiguous.

I assume that consumption growth is of the form

$$\Delta c_{t+1}^j = E_t \Delta c_{t+1}^j + \epsilon_{t+1} + \eta_{t+1}^j,$$

with $\epsilon_{t+1} \sim N(-\sigma_\epsilon^2/2, \sigma_\epsilon^2)$ and $\eta_{t+1}^j \sim N(\sigma_{\eta,t+1}^j/2, \sigma_{\eta,t+1}^j)$. To illustrate that it is crucial that the variance of idiosyncratic shocks depends on the aggregate state, I consider the simplest example:

² I focus on the Sharpe ratio instead of the equity premium because it is a measure of the risk-return tradeoff that is independent of any specific asset. Computing prices of individual asset would require the specification of a dividend process.

³ The working-paper version of this paper, Lettau (2001), contains a more detailed analysis.

$$\sigma_{\eta,t+1}^2 = a_0 + a_1 \epsilon_{t+1}. \quad (3)$$

Lettau (2001) shows that

$$SR(M_{t+1}^j) \approx \gamma \sigma_\epsilon \left(1 - a_1 \frac{1 + \gamma}{2} \right).$$

The Sharpe ratio increases if the variance of idiosyncratic shocks depends negatively on the aggregate state. This is the same mechanism highlighted by Mankiw (1986) and Constantinides and Duffie (1996).

III. Volatility Bounds with Income Data

Next, I evaluate the empirical importance of the mechanism presented in the previous section. The problem with a direct empirical analysis of this channel is that individual consumption data (for example, in the PSID) is known to be of poor quality. In addition, consumption data in the PSID span only a short time series. Hence, estimating equations like equation (3) directly from consumption data is flawed with measurement error. However, individual income data in the PSID is of much higher quality, and there is a large empirical literature that estimates income processes for individual households (Abowd & Card, 1989; Carroll, 1992; MaCurdy, 1982). As an alternative to a direct test using consumption data, I use these income processes instead of evaluating the volatility bounds on an individual level. This strategy is appropriate if individual consumption is smoother than individual income. Of course, because there is no data set with a long time series of high-quality household consumption data, it is impossible to check directly whether individual consumption is indeed smoother than income. However, we can find theoretical conditions under which we expect this to be the case. I will argue next that these conditions are likely to be fulfilled in the data.⁴

Start with an individual who consumes her permanent income. This is a useful point of departure because much of the consumption literature is cast in this framework. (See Deaton (1992) for a survey.) Intuitively, if income growth is negatively autocorrelated, consumption is smoother than income because a shock to income has less than a one-to-one effect on permanent income. On the other hand, if income growth is positively autocorrelated, then permanent income is more volatile than current income. Because, theoretically, permanent income can be more or less volatile than income, it is an empirical question which case is more realistic. The evidence in the empirical literature points overwhelmingly in the direction of negatively autocorrelated income growth on the household level (Abowd & Card, 1989; Carroll, 1992; MaCurdy, 1982; Pischke, 1995). In this case, permanent income is smoother than current

⁴ Lettau (2001) contains a more detailed derivation of the following statements.

income, and therefore consumption of a permanent income agent will be smoother than her income.⁵

If consumption is indeed smoother than income, then risk premia and volatility bounds computed using income data provide a useful lower bound. If risk premia are low for income data, then there is surely little hope in generating higher-risk premia using smoother consumption data. Hence, this approach provides a straightforward benchmark to evaluate the potential of models with idiosyncratic risk. The punchline is that risk premia are too small even if individual income data are used to compute them. The reason is that the covariance of the volatility of idiosyncratic income risk and the aggregate state is too small to increase risk premia, even when income is used instead of consumption.

Because agents receive different income in each period and they cannot trade with each other, they will in general disagree on asset prices. One agent might require an expected return of 5% for some asset given her particular income process, and another agent might require only a 3% expected return for the same asset. If there were open asset markets, they would trade the asset and smooth their consumption paths so that in equilibrium they agree on expected returns of all assets. Because I abstract from any asset trading in this paper, I compute the expected returns implied by the consumption process (which is equal to the income process by assumption) of each individual agent. To obtain an upper bound for the ability of idiosyncratic risk to increase risk premia, I consider the highest expected return of a given asset across all agents. In terms of the Sharpe ratio, using the individual income processes as consumption will result in different Sharpe ratios for different agents. An upper bound for the Sharpe ratio that would result from trading is given by the highest Sharpe ratio computed from the individual income processes. To give the models the best possible shot, I report this maximal Sharpe ratio across agents. The Sharpe ratio, estimated from postwar data using S&P 500 excess returns, is 0.27 per quarter or 0.54 per annum.

A. Income Process from Heaton and Lucas

Heaton and Lucas (1996) (hereafter, HL) present a model with idiosyncratic labor income shocks, transaction costs, borrowing, and short-sale constraints with eight discrete states. Table 1 summarizes the results for idiosyncratic income shocks estimated by HL. They estimate $\sigma_{\eta,t+1} = \bar{a}_0 + \bar{a}_1 \log(\gamma_{t+1}^a)$, where γ^a is the growth rate of aggregate

⁵ One possible objection to these arguments is that consumption behavior in models with incomplete markets might deviate from permanent income behavior. But, even in more-complex models, it is unlikely that consumption is more volatile than current income if income growth is negatively autocorrelated.

TABLE 1.—THE SHARPE RATIO WITH HEATON-LUCAS INCOME SHOCKS

Income Process	Heaton and Lucas	
	With Dividends	No Dividends
	$\gamma = 1$	
HL1: $\bar{a}_1 = 0$	0.0110	0.0111
HL2: $\bar{a}_1 = -1.064$	0.0299	0.0301
HL3: $\bar{a}_1 = -4.450$	0.0767	0.0769
	$\gamma = 5$	
HL1: $\bar{a}_1 = 0$	0.0552	0.0553
HL2: $\bar{a}_1 = -1.064$	0.1813	0.1820
HL3: $\bar{a}_1 = -4.450$	0.5773	0.5785

Table reports the average Sharpe ratio across the eight states of the economy weighted with the probabilities of the stationary distribution. The maximal Sharpe ratio across agents is used in each state. The model is calibrated for annual data.

income.⁶ HL consider three sets of estimates: $\bar{a}_0 = 0.251$, $\bar{a}_1 = -0$; $\bar{a}_0 = 0.360$, $\bar{a}_1 = -1.064$; and $\bar{a}_0 = 0.290$, $\bar{a}_1 = -4.450$. The second set of parameters corresponds to the PSID regressions results. The aggregate shocks are chosen to match the moments of aggregate consumption. In each of the eight states, I compute the Sharpe ratio for both types of agents. Then I compute the average Sharpe ratio (weighted with the probabilities of the stationary distribution) across states using the higher of the two Sharpe ratios. The transitions probabilities are taken from HL, who consider two cases: one with aggregate dividends and one without.

For low risk aversion of unity, the picture looks very bleak. None of the three cases reach a Sharpe ratio of 0.1. Note that the effect of dividend uncertainty is very small. It is clear that higher risk aversion is needed to match the Sharpe ratio in the data. The bottom panel uses $\gamma = 5$. Here, Sharpe ratios increase substantially when \bar{a}_1 is negative. For the case corresponding to the PSID data, the Sharpe ratio is approximately 0.18. This value is still only about one-third of the Sharpe ratio in the data. Only if $\bar{a}_1 = -4.45$ (that is, more than four times the PSID value) is the Sharpe ratio high enough to be consistent with the data.

Is this a success? As mentioned before, this analysis computes only a very loose upper limit for the Sharpe ratio: consumers cannot trade against their idiosyncratic labor shocks at all; the Sharpe ratio of the poor agents is used despite the fact that those agents are likely to be borrowing constraint and hence do not determine asset prices; and, lastly, the required value for \bar{a}_1 is approximately four times the value found in the data. Only in this case does the model generate a reasonable Sharpe ratio. Thus, it is not surprising that HL found low Sharpe ratios and risk premia once they allow agents to trade. The model is not able to produce higher risk premia even if agents are forced to consume their labor income.

⁶ HL specify the standard deviation of the idiosyncratic shock as a linear function of the aggregate growth rate. This implies $\sigma_{\eta,t+1}^2 = \bar{a}_0^2 + 2\bar{a}_0\bar{a}_1 \log(\gamma_{t+1}^a) + \bar{a}_1^2 (\log(\gamma_{t+1}^a))^2$.

TABLE 2.—THE SHARPE-RATIO WITH KRUSELL-SMITH INCOME SHOCKS

Agent	Aggregate State	
	Bad	Good
Employed	0.0641	0.1088
Unemployed	0.1759	0.2840
No employment shocks	0.0037	0.0037

Table reports the Sharpe ratio with KS unemployment shocks. The model is calibrated for quarterly data.

B. Income Process from Krusell and Smith

Krusell and Smith (1997) (hereafter, KS) propose a model in which agents have a higher probability to become unemployed when aggregate times are bad. This yields again a negative covariance of idiosyncratic uncertainty and state aggregate that has the potential to increase risk premia. Let aggregate output per capita be given by Y^a . If an agent becomes unemployed, she receives unemployment insurance of gY^a . Hence, an employed agent's income is $(1 - g)Y^a$. Aggregate income can take on two values, so that each agent has four possible states.

KS calibrate the model as follows. The unemployment rate in the good aggregate state is 4% and 10% in the low state. Unemployed agents receive $g = 9\%$ unemployment insurance. The Markov transition probabilities are chosen so that unemployment shocks are fairly persistent: that is, an unemployed agent today has a smaller probability to find a job tomorrow than an agent who is employed today. Again, I refer to KS for the exact transition probabilities. All other parameters are taken from KS, who calibrate the model for quarterly data.

Table 2 reports the Sharpe ratios in this economy. (Recall that the Sharpe ratio in quarterly data is 0.27.) Unemployed agents have a much higher Sharpe ratio than do employed agents because their current and expected future income is much lower. Note that the Sharpe ratio of both types is higher in good aggregate times; that is, it is procyclical. This comes from the specific choice of the transition probabilities. However, the lowest Sharpe ratio (of the unemployed agents in bad times) just matches the Sharpe ratio in the data.

C. Income Process from Storesletten, Telmer, and Yaron

Storesletten, Telmer, and Yaron (1997) (hereafter, STY) study an OLG model with idiosyncratic income shocks and borrowing constraints. They calibrate various Markov processes for idiosyncratic shocks to match the PSID data. Labor supply of agent j , N^j , is determined by (potentially multiple) idiosyncratic shocks: $N^j_t = \exp(\eta^j_t)$. The labor input is supplied inelastically and combined with an aggregate Cobb-Douglas production function to produce the consumption good.

STY propose different processes for η^j_t . Their first three processes differ mainly in the persistence of the idiosyncratic shocks, which turns out to be important for the risk

sharing in their OLG model. In their fourth version, they allow the aggregate state to affect the variance of the idiosyncratic shock. They estimate equation (3) using PSID data. Using two different estimation techniques, they estimate the coefficients as $\hat{a}_0 = 0.054$, $\hat{a}_1 = -0.011$; and $\hat{a}_0 = 0.054$, $\hat{a}_1 = -0.046$. Both slope coefficients are not negative enough to increase the Sharpe ratio. Finally, STY propose a different technique to estimate cross-sectional dispersion: they estimate a model with a high and low idiosyncratic variance regime depending on whether the aggregate economy is above or below trend. STY find that the variance of idiosyncratic shocks is 0.032 when the aggregate economy is above trend and 0.184 when it is below trend. Table 3 reports the Sharpe ratio for all five cases considered by STY. Again, all other parameters are taken from STY.

The case without any idiosyncratic shocks generates a Sharpe ratio of 0.0594. The following three cases produce the same Sharpe ratio as without idiosyncratic risk. Here, the distribution of the idiosyncratic shocks are assumed to be independent of the aggregate state. The next two versions incorporate the dependence of the variance of the idiosyncratic shock on the aggregate state as estimated from the PSID data set. Because the estimated a_1 's are negative, the Sharpe ratio increases, but the increase is only minute because the parameters are not negative enough. The Sharpe ratio in the case with the changing variance regime is approximately 0.1, which is still only one-fifth of the required value. None of the income processes considered by STY can produce high Sharpe ratios even if agents cannot smooth them.

IV. Conclusion

To evaluate the empirical relevance of idiosyncratic uncertainty, I compute Sharpe ratios using individual income processes from the PSID data set. Using income data provides an upper bound for individual consumption data in the absence of reliable consumption data. I take the estimated income processes from recent studies of Heaton and Lucas (1996), Krusell and Smith (1997), and Storesletten, Telmer,

TABLE 3.—THE SHARPE-RATIO WITH STORESLETTEN, TELMER, AND YARON INCOME SHOCKS

Income Process	Sharpe Ratio
No idiosyncratic shocks	0.0594
Unit root	0.0594
High persistence	0.0594
Moderate persistence	0.0594
Low persistence	0.0594
PSID $a_1 = -0.011$	0.0604
PSID $a_1 = -0.046$	0.0635
High/low variance regimes	0.1013

Table reports the Sharpe ratio for income processes considered by Storesletten, Telmer, and Yaron (1997). The parameters for the different AR1 income shocks are as follows. STY approximate the processes with discrete Markov chains. Let ρ be the AR1 coefficient and σ_η the innovation standard deviation. Unit root: $\rho = 1$, $\sigma_\eta = 0.201$; high persistence: $\rho = 0.929$, $\sigma_\eta = 0.230$; moderate persistence: $\rho = 0.529$, $\sigma_\eta = 0.251$. The two PSID cases use the unit root parameters and the estimated coefficients in equation (3). The last case assumes a variance of 0.032 when the economy is in the high-aggregate state and 0.184 when the aggregate state is low. The model is calibrated for annual data.

and Yaron (1997) and compute Sharpe ratios as if agents had to consume their endowment. Because there is no market that equilibrates the SDF of different agents, they will disagree on asset prices. To obtain an upper bound on Sharpe ratios, I compare the Sharpe ratio in the data to that of the maximal Sharpe ratio across agents. I find that Sharpe ratios are generally smaller than in the data; only some extreme cases just reach the data equivalent. Hence, it seems unlikely that idiosyncratic consumption risk can generate an SDF that pass the Hansen-Jagannathan bounds test.

REFERENCES

- Abowd, John, and David Card, "On the Covariance Structure of Earnings and Hours Changes," *Econometrica* 57:2 (1989), 411-445.
- Aiyagari, Rao, and Mark Gertler, "Asset Returns and Uninsured Individual Risk," *Journal of Monetary Economics* 27:3 (1991), 311-331.
- Carroll, Christopher D., "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers of Economic Activity* 2 (1992), 61-135.
- Cogley, Timothy, "Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey," Federal Reserve Bank of San Francisco mimeograph (1998).
- Constantinides, George M., and Darrel Duffie, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104 (1996), 219-240.
- Deaton, Angus, *Understanding Consumption* (Oxford: Clarendon Press, 1992).
- Den Haan, Wouter, "Understanding Equilibrium Models with a Small and Large Number of Agents," UCSD working paper (1996).
- Heaton, John, and Deborah Lucas, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy* 104 (1996), 443-487.
- Krusell, Per, and Anthony A. Smith, "Income and Wealth Heterogeneity, Portfolio Selection, and Equilibrium Asset Returns," *Macroeconomic Dynamics* 1 (1997), 387-422.
- Lettau, Martin, "Idiosyncratic Risk and Volatility Bounds, or Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?" Federal Reserve Bank of New York working paper (2001).
- MaCurdy, Thomas E., "The Use of Time-Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis," *Journal of Econometrics* 18 (1982), 83-114.
- Mankiw, N. Gregory, "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics* 17 (1986), 211-219.
- Pischke, Jörn-Steffen, "Individual Income, Incomplete Information and Aggregate Consumption," *Econometrica* 63:4 (1995), 805-840.
- Storesletten, Kjetil, Chris Telmer, and Amir Yaron, "Persistent Idiosyncratic Shocks and Incomplete Markets," Carnegie Mellon University mimeograph (1997).